



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-850159

Spatio-Temporal Surrogates for Interaction of a Jet with High Explosives: Part II - Clustering Extremely High-Dimensional Grid-Based Data

C. Kamath, J. S. Franzman

June 12, 2023

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Spatio-Temporal Surrogates for Interaction of a Jet with High Explosives: Part II - Clustering Extremely High-Dimensional Grid-Based Data

Chandrika Kamath and Juliette S. Franzman

Lawrence Livermore National Laboratory
7000 East Avenue, Livermore, CA 94551, USA
`kamath2,franzman1@llnl.gov`

9 June 2023

Abstract

Building an accurate surrogate model for the spatio-temporal outputs of a computer simulation is a challenging task. A simple approach to improve the accuracy of the surrogate is to cluster the outputs based on similarity and build a separate surrogate model for each cluster. This clustering is relatively straightforward when the output at each time step is of moderate size. However, when the spatial domain is represented by a large number of grid points, numbering in the millions, the clustering of the data becomes more challenging. In this report, we consider output data from simulations of a jet interacting with high explosives. These data are available on spatial domains of different sizes, at grid points that vary in their spatial coordinates, and in a format that distributes the output across multiple files at each time step of the simulation. We first describe how we bring these data into a consistent format prior to clustering. Borrowing the idea of random projections from data mining, we reduce the dimension of our data by a factor of thousand, making it possible to use the iterative k-means method for clustering. We show how we can use the randomness of both the random projections, and the choice of initial centroids in k-means clustering, to determine the number of clusters in our data set. Our approach makes clustering of extremely high dimensional data tractable, generating meaningful cluster assignments for our problem, despite the approximation introduced in the random projections.

Contents

1	Introduction	3
2	Description of the data	3
3	Challenges to the analysis	6
3.1	The unsuitability of the raw data for clustering	7
3.2	The iterative nature of k-means clustering algorithm	9
4	Related work	9
5	Solution approach	10
5.1	Pre-processing the data	10
5.1.1	Creating a data file for each time step in a simulation	11
5.1.2	Aligning and cropping coordinates to a common domain	11
5.1.3	Remapping simulations to a common grid	11
5.1.4	Converting remapped data to snapshot matrices	12
5.2	Clustering the snapshots	12
5.2.1	Using an iterative method with dimension reduction	13
5.2.2	Using an iterative method with a reduced representation	14
5.2.3	Using a non-iterative method	15
6	Results and discussion	15
6.1	Results with random projections and k-means clustering	16
6.1.1	Understanding the effect of the randomness of random projections	16
6.1.2	Exploiting the randomness of k-means to determine number of clusters	17
6.2	Results with k-means clustering using weights from SVD	21
6.3	Results with hierarchical clustering	21
6.4	Discussion	23
7	Conclusions	25
8	Acknowledgment	25
A	Appendix: Data and results for x-momentum	28
B	Appendix: Data and results for y-momentum	32

1 Introduction

A common task in building surrogates for the spatio-temporal outputs from computer simulations is to transform these outputs into a lower-dimensional space using a linear method such as the principal component analysis (PCA) [14]. However, when the data do not lie on a linear manifold, the dimension of this lower-dimensional space can be large, prompting the consideration of non-linear dimension-reduction methods. A simple approach to introducing non-linearity is to cluster the simulation outputs by similarity and then use a linear transform on each cluster separately [17], creating a locally-linear surrogate. However, clustering the outputs becomes challenging when the spatial domains vary across simulations or the spatial data generated at each time step of the simulation is extremely large, composed of values at over a million grid points in the spatial domain. This report describes our experiences with addressing these challenges.

We start this report by describing the problem for which we want to build a spatio-temporal surrogate, namely, the interaction of a jet with high explosives, and the two-dimensional outputs that are generated during the simulations of this problem (Section 2). These outputs form the data set that we want to cluster. We then discuss the issues that make it challenging to cluster these outputs (Section 3) and place our contributions in the context of related work (Section 4). Our solution approach (Section 5) describes how ideas from other domains make it possible to cluster the high-dimensional data. We show how we can select the parameters used in our algorithms to generate meaningful clustering results for our data set (Section 6). We conclude this report with a summary of our experiences and the lessons learned (Section 7).

This report is the second of two reports summarizing our work on building spatio-temporal surrogates for the problem of a jet interacting with high explosives. In the first report [16], we discuss the applications aspect of our work and describe how we can build an accurate surrogate despite the small number of simulations in our data set. One of the approaches to improving the accuracy of the surrogate is by building locally-linear surrogates. This requires clustering of the data, which is the focus of this report.

2 Description of the data

We illustrate our ideas on clustering extremely high-dimensional data using simulation output from a problem describing the interaction of a jet with high explosives (HE). The domain of the problem is a right cylinder with its axis oriented horizontally as shown on the left in Figure 1. There is a steel plate, 1cm thick near the right end of the cylinder, with the LX14 high explosives to the left of the plate. Both the plate and the HE have a fixed radius of 10cm. A copper jet, aligned along the center line of the cylinder, enters the HE from the left. The simulation models what happens as the jet moves through the HE and the plate. The jet is modeled initially as uniform cylinder. It is 10 cm in length with a varying radius. The jet tip velocity is specified as an input parameter; a linearly varying velocity profile is applied to the remainder of the cylinder that represents the jet to approximate a stretching metal jet. As the problem is radially

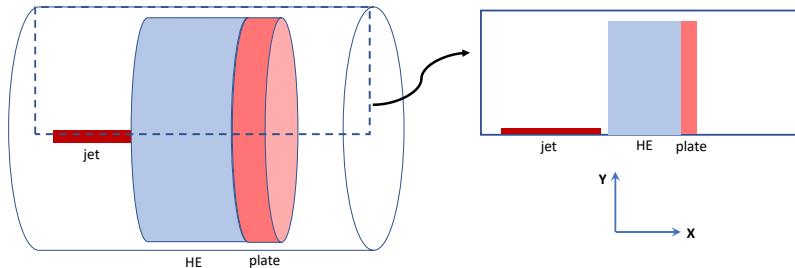


Figure 1: A schematic of the problem being simulated. On the left is the horizontal cylinder in three-dimensions, showing the plate in pink, to the left of which is the HE in blue. The jet, in red, enters the HE from the left. As the problem is radially symmetric about the axis of the cylinder, we need to simulate only the two-dimension region shown by a dotted rectangle on the left and schematically on the right.

symmetric about the axis of the cylinder, only the two-dimensional region shown by a dotted rectangle on the left, and schematically on the right, is simulated. At each time step, the simulation outputs variables of interest, such as mass and momentum, at different points on a grid in the two-dimensional region.

There are three input parameters for the simulation: the *radius* of the jet, the *length* of the high explosives to be traversed by the jet, and the *tip velocity* of the jet in the positive x direction. By running the simulations at select values of these input parameters, and collecting the output at different time steps, we can create a data set that could be used to build a surrogate model for predicting the output at a new set of input parameters and a given time step. We are interested in determining, for example, whether the plate breaks; what is the final position of the plate; and, if the plate breaks, what is the velocity of the tip of the jet as it comes out on the other side of plate. To answer these questions, we need to build an accurate spatio-temporal surrogate, a topic we discuss in the companion report [16], which focuses on the accuracy of the surrogates created using only a small number of simulations. This report focuses on the tasks of processing and clustering of extremely high-dimensional data that are crucial to building this accurate surrogate.

To illustrate the instances in our data set, we use four simulations whose parameters are listed in Table 1. Figure 2 shows the output variable, *mass*, at the first and last time steps for these four example simulations. As explained earlier, we have simplified the three-dimensional problem by assuming radial symmetry around the axis of the cylinder, so the output from the simulation is shown as two dimensional images, with the axis of the cylinder shown at the bottom, that is, at $y = 0$. The domain extent in x (along the length of the cylinder) varies as the length of the HE varies across simulations; however, the domain in y ranges from 0 to 11cm for all simulations.

Simulation key	jet radius (cm)	HE-length (cm)	jet tip velocity (cm/ μ sec)	# time steps	#grid points	outcome
r01_i004	0.15	13.67	0.894	38	2,859,387	almost break
r01_i017	0.17	12.24	0.648	41	2,759,181	no break
r02_i021	0.14	6.77	0.914	30	2,374,179	break
r02_i028	0.23	10.54	0.843	35	2,639,637	break

Table 1: Input parameters for the four example simulations shown in Figures 2. Note the very large number of grid points (over two million) at which variables of interest are output at each time step in a simulation.

In Figure 2 the vertical plate, shown in red, is stationary at time $t = 0$. To the left of the plate is the HE shown in light blue. The jet is shown in red at the bottom of the domain to the left of the HE; it is quite thin relative to the radius of the cylinder, and is barely visible in the images. As the simulation evolves, the jet moves to the right, through the HE, which expands, pushing the plate to the right. At late time, depending on the simulation input parameters, the plate could

- **break**, with the jet going through the plate and coming out clearly on the other side;
- **almost break**, with the jet either going completely through the plate but barely coming out the other side or the jet going almost all the way through the plate, leaving it barely connected at the bottom;
- **not break**, with the plate remaining attached, either partially or completely, at the bottom. The plate could have moved from its original position at time $t = 0$.

We used the last two time steps in each simulation to assign one of these three class labels to the simulation. This label was not used in building the surrogate; it was used only to ensure we had a good coverage of the design space. We selected the four example simulations in Figure 2 to illustrate these three cases.

The output at each time step of a simulation consists of the values of variables of interest that are generated at grid points in the two-dimensional rectangular domain. These grid points are on a regular grid, with $\Delta x = \Delta y = 0.0125\text{cm}$. There are three output variables: *mass*, *x-momentum*, and *y-momentum*; the latter two are shown in later in Appendix A and B, respectively. The values of these variables are defined at the

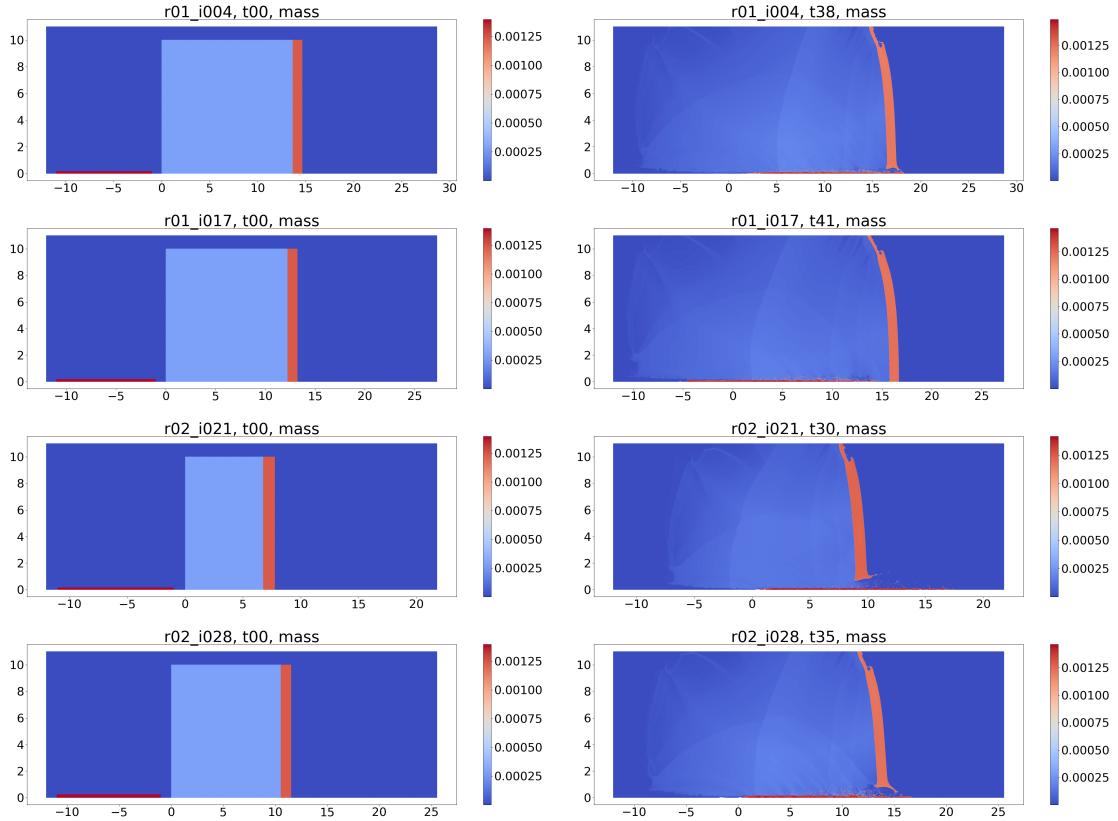


Figure 2: The variable *mass* at the first time step, (left column) and last time step (right column) for the four example simulations in Table 1. From top to bottom, simulations with keys r01_i004, r01_i017, r02_i021, r02_i028, illustrate an *almost break*, *no break*, *break*, and *break* case, respectively. The data shown are before the pre-processing steps described in Section 5.1. The vertical red region represents the plate, the light blue represents the high explosive (HE), the horizontal red region at the bottom to the left of the HE at the first time step is the jet, and the dark blue region is air. The right column shows the last time step of each simulation and the effect on the plate as the jet moves to the right, through the HE, and potentially through the plate. The range of x values is different for different simulations, while the range of y values is the same. The HE starts at $x = 0$ in each plot. Note that even when the plate does not break (second row), the plate has moved to the right from its original position.

center of the square cell formed by four nearby grid points. Thus the data appear as an image, with regularly spaced pixels. However, in general, in a simulation, the grid points need not be on a regular grid; they could form an unstructured grid, as in a finite element mesh, or a locally structured grid, as in an Adaptive Mesh Refinement (AMR) mesh, where the mesh evolves with the simulation. As a result, unlike an image, most output from simulations also include the (x, y) coordinates of the grid points. In our work, we retain this association of the coordinates with the grid points as they enable us to extract sub-domains of the larger domain for processing.

Each simulation is run for a fixed number of time steps which is determined as $(\lfloor (\text{HE-length}/\text{jet-tip-velocity}) \rfloor + 23)$, with the output generated at each time step. As both HE-length and jet tip velocity vary with the simulation, the number of time steps also varies across simulations. At early time, as the jet starts to move through the HE, there is little of interest in the simulation output. Once the jet is partway through the HE, as indicated by the first term in the equation above, it starts to influence the location of the plate, until 23 μsec later, it is expected that we should know the final status of the plate. In our work, we consider all the time steps in the analysis; an alternative would be to consider only the later 23 time steps.

The data processed in this report was obtained by running 45 simulations at select values of the input parameters. These values, or sample points in the three input dimensions, were generated incrementally using a modified version of the best candidate algorithm [21, 15], which selects samples randomly, but far apart from each other. We started with a small number of sample points and an initial guess at the range of each of the three inputs. We then restricted the ranges to focus on the break cases, and added new sample points until we had a total 45 sample points [16]. Our data set, shown in Figure 3, indicates that at high jet tip velocity, but low HE-length, the plate breaks, while at low jet tip velocity and high HE-length, the jet is unable to penetrate through the plate.

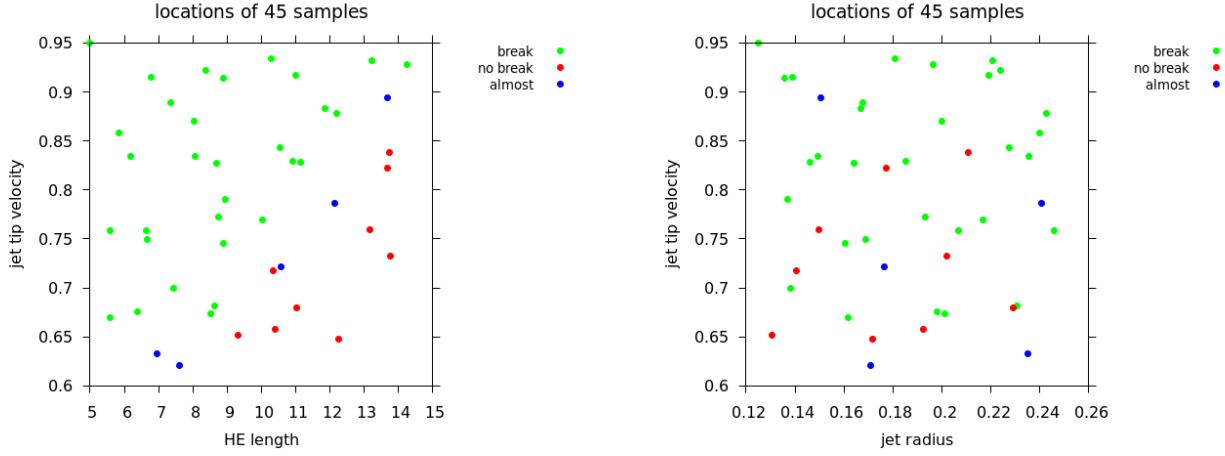


Figure 3: The 45 samples in the space of the three input parameters for the simulation, labeled by the state of the plate at the last time step of the simulation. The simulation at the extreme corner of the input space, with length, jet tip velocity, and jet radius equal to 5.0cm, 0.950cm/ μ s, and 0.125cm, respectively, is referred to as the *baseline* simulation. It has the smallest number of time steps. Our goal is to cluster all the time steps across these 45 simulations.

It is clear that our data set is unbalanced as we have 9 samples where the plate does not break, 31 where the plate breaks, and 5 that are almost break. Generating an appropriate data set for a problem like ours is challenging as time constraints limit the number of simulations we can run. In addition, we do not know *a priori* the range of inputs that will give us sample points with the desired outcome, and the boundary between the classes is poorly defined, making it difficult to generate a balanced data set. We erred on the side of having more break cases as these were of greater interest; the no-break cases also tended to have output that appeared very similar, and we expected that a small number of such cases would suffice. Admittedly, our choice of sample points will affect the clustering results.

The output data for a variable at a time step in a simulation is referred to as a *snapshot*, so named as it is a snapshot of the evolution of the simulation at a particular point in time. These 45 simulations generate a total of 1604 snapshots. The simulation at the extreme corner of the input space, with length, jet tip velocity, and jet radius equal to 5.0cm, 0.950cm/ μ s, and 0.125cm, respectively, is referred to as the *baseline* simulation. It has the smallest number of time steps, with 29 snapshots.

3 Challenges to the analysis

There are two main challenges to building a high-quality spatio-temporal surrogate for our problem:

- The first is how do we build a surrogate that is *accurate* when we can only run a small number of simulations? We discuss several options in our companion report [16], one of which is to cluster the snapshots and build a separate surrogate for each cluster.

- This leads to our second challenge - how do we cluster the snapshots in our data set? Table 1 and Figure 2 indicate that the snapshots are high-dimensional, each with more than two million grid points; the sizes of the snapshots vary across simulations; and the snapshots are not aligned in any way; and the data at each time step are available in multiple files. All these factors would make clustering the snapshots challenging. In addition, popular clustering algorithms, such as k-means, are iterative, which can lead to computational inefficiencies in processing extremely high-dimensional data.

The first of these two challenges is the focus of the companion report [16]; this report focuses on the second challenge, namely, generating a meaningful clustering of the high-dimensional snapshots from our problem of jet-HE interaction. To accomplish this goal, we need to address the two issues discussed next.

3.1 The unsuitability of the raw data for clustering

The output data generated for each of the 45 simulations, regardless of the size of the domain, are available in 360 files in HDF5 format [24] for each time step. Each HDF5 file includes five variables — x and y coordinates, mass, x-momentum, and y-momentum. Within each file, the variables are in natural ordering, that is, ordered by increasing values of the y -coordinate, and for a fixed y -coordinate, ordered by increasing values of the x -coordinate. All simulations are on a regular grid with $\Delta x = \Delta y = 0.0125\text{cm}$.

To apply our analysis algorithms to these data, we first need to re-arrange the data so that we have three snapshot matrices, one for each of the three variables of interest. Each snapshot matrix should have the grid points as the rows, listed in natural order of the (x, y) coordinates, and, as columns, the variable values at each time step of each simulation. A snapshot matrix, $\mathbf{X} \in \mathbb{R}^{D \times N}$, for any one of the three variables, can then be written as

$$\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N], \quad (1)$$

where $\mathbf{x}_i \in \mathbb{R}^D$, N is the number of snapshots, and D is the number of grid points in a snapshot. Since we have a total of 1604 time steps in the 45 simulations, there are 1604 columns in the snapshot matrix. However, identifying the rows of the snapshot matrix is more challenging. To cluster the snapshots, we first need to bring the data in each simulation to a common grid, which is not straightforward for several reasons:

- The domain over which the data are generated is different for each simulation. The y values are in the range $[0:11.0]$ cm for all simulations, but the range of x values varies as the HE-length varies, resulting in snapshots that vary in length across simulations.
- Figure 2 shows that the origin of the domain in all simulations is at the left edge of the HE, with 12 cm of air to the left of the origin. So, if we align the snapshots at the left edge at $x = -12.0$, the plate locations will not be aligned, even at time step 0, as the HE length varies across simulations. Since the plate forms an important structure in the output data, the lack of alignment of the plate across simulations would not give us the clustering results we expect.
- For a simulation, the distribution of the data in the 360 files is the same across all time steps. However, the files within a simulation have four different sizes, as shown in Figure 4. But, as the domain size varies across simulations, the sizes of these HDF5 files also vary across simulations. As each row in the snapshot matrix corresponds to the same (x, y) coordinates, it becomes non-trivial to map a point in one of the 360 files at a time step to its location in the snapshot matrix efficiently.
- All the simulations have the coordinate $(-12, 0)$ as the lower left corner. But, due to the way in which the (x, y) coordinates are generated for domains with different ranges in x , the values of Δx and Δy are not exactly 0.0125cm across simulations. While the coordinates in y , which has a fixed range of values, are identical across simulations, this is not the case for the x coordinates.
- Finally, the total number of grid points across the 360 HDF5 files for the data at any time step is nearly three million (see Table 1). A snapshot matrix with such a large number of rows and 1604 columns is too large to be read into memory for processing and we need to consider alternative ways to store the snapshot matrix; this would also affect the analysis algorithms that read in the data.

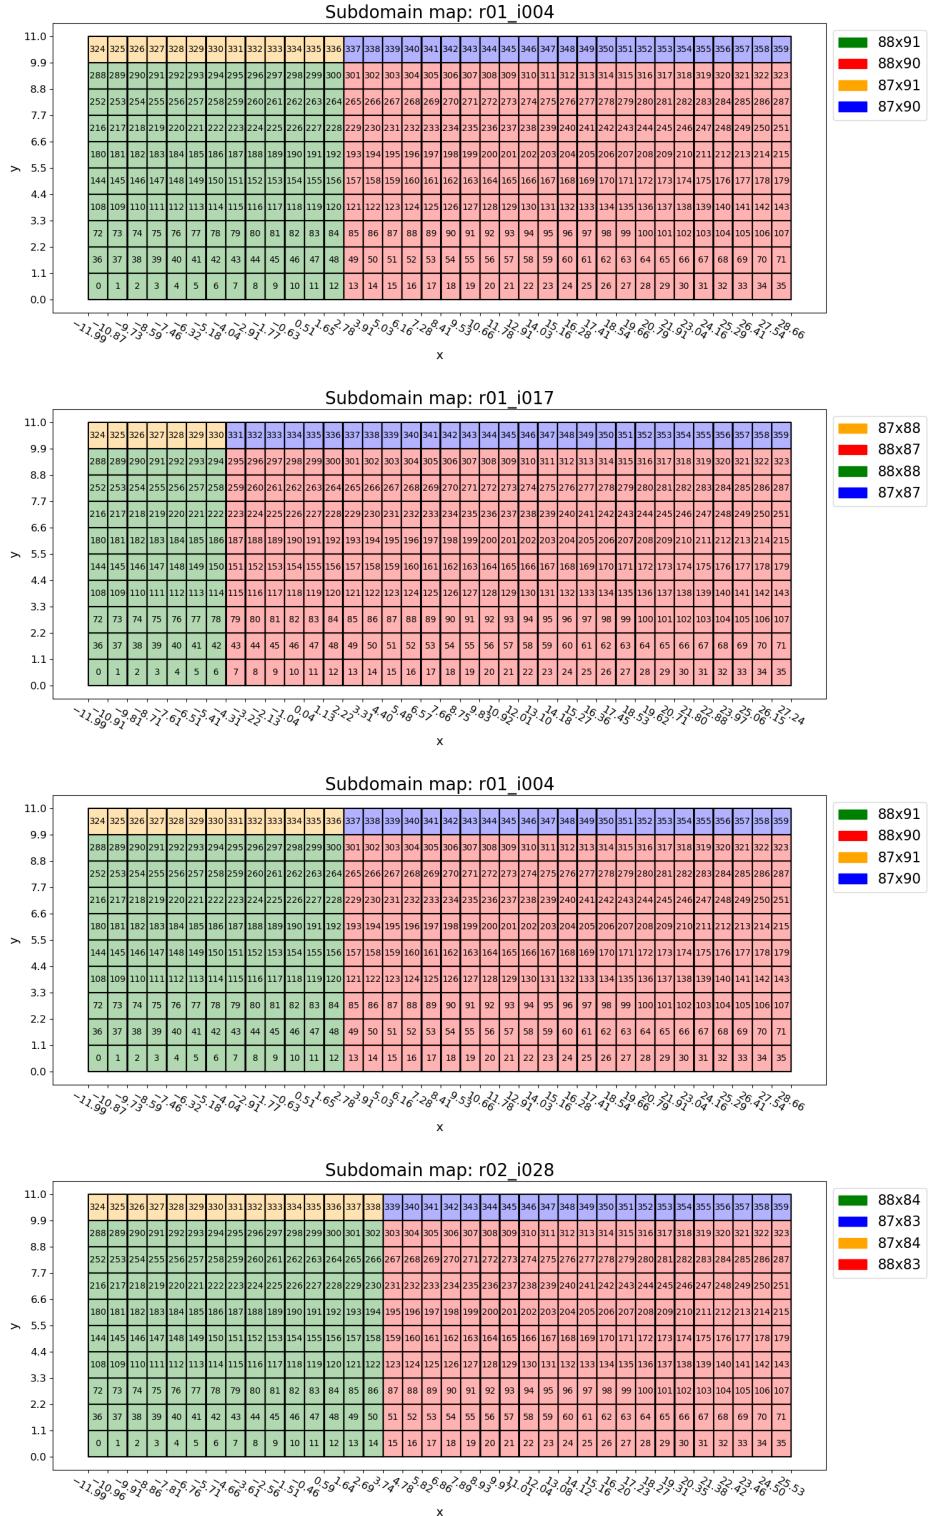


Figure 4: The distribution of data in 360 sub-domains for the four example simulations in Table 1. At each time step, the data in each sub-domain is available in a separate HDF5 file. From top to bottom, simulations with keys r01_i004, r01_i017, r02_i021, r02_i028, respectively. Note that the sizes of the sub-domains within a simulation vary. All time steps in a simulation have the same distribution of the data.

3.2 The iterative nature of k-means clustering algorithm

In addition to the challenges posed by the size of the raw data and the storage format, we also need to consider the clustering algorithm used to group the snapshots. We intuitively expect that the snapshots can be clustered because the early time steps, when the plate has barely moved, are distinctly different from the later time steps, where, depending on the input parameters of the simulation, the plate moves but does not break, or the plate moves and breaks. However, as neighboring time steps are often quite similar, any clustering will place some consecutive time steps from a simulation into different clusters, suggesting that any clusters identified are not very well separated.

In the absence of any information on the shape of possible clusters or the density of points in the very high-dimensional space, it is not obvious which clustering algorithm is the most suitable one for our data. We chose to start with the simplest clustering method, k-means clustering [13, 11], described in Algorithm 1, with the expectation that it might provide insight to guide the choice of a more appropriate algorithm.

Algorithm 1 *k*-means clustering

- 1: Goal: given the N snapshots in a D -dimensional space in the form of the matrix X in Equation 1, identify a clustering with nc clusters, such that each snapshot is assigned to a cluster, and the snapshot is closer to the centroid of its cluster than the centroid of any other cluster.
- 2: Set $niter$, the maximum number of iterations and $thresh$, the threshold that determines when a cluster center has moved too small an amount between iterations.
- 3: Choose an initial set of nc snapshots randomly as the initial cluster centroids.
- 4: **for** $iter = 1$ to $niter$ **do**
- 5: Assign each snapshot to its nearest cluster centroid using the Euclidean distance.
- 6: Update each cluster centroid to be the mean of the snapshots assigned to it.
- 7: If the maximum Euclidean distance moved by any centroid is less than $thresh$, terminate.
- 8: **end for**

We observe that at each iteration, this algorithm requires the calculation of the distance of each snapshot to the centroid of each cluster. As the snapshot matrix is too large to fit into memory, reading in the matrix, piece-meal, on each iteration, will be time consuming. We therefore need alternate ways to make the clustering of the snapshots computationally tractable.

4 Related work

We address two tasks in this report - converting the data from the HDF5 files into a snapshot matrix and clustering the columns of the high-dimensional snapshot matrix. The first task is specific to each data set and is based on the characteristics of the data. This section therefore focuses on the second task.

In the field of spatio-temporal modeling, the problems solved typically have snapshots with grid points numbering in the thousands or tens of thousands, and occasionally hundreds of thousands [16]. Clustering these relatively low-dimensional snapshots then becomes a straightforward application of a clustering algorithm, such as k-means or k-mediods, with a suitably defined distance metric, such as Euclidean distance, reconstruction distance, or Grassmann distance [17, 3, 8, 6].

To cluster extremely high-dimensional data, where each snapshot has over two million grid points, we borrowed ideas from the field of data mining, where such data sets are common place. A typical solution is to first reduce the dimension and then cluster the data. Ritter and Kohonen in 1989 suggested using random projections [2] for dimension reduction, followed by clustering using self-organizing maps, which is another iterative technique like k-means. Using the same combination of methods, Kaski in 1998 [18] showed that random projections was faster than PCA for dimension reduction; it required a slightly larger number of reduced dimensions, but the results were equally good, and similar to the results with the original data. In later work, Fern and Brodley [10] found that the randomness of random projections could result in different cluster assignments and proposed using ensemble clustering [23] to generate a stable cluster assignment.

They also found that random projections performed better than PCA, and that there was no universal single best way to combine the results of the ensemble [9]. More recently, Anderlucci et al. [4] have further investigated different ways to combine the different cluster assignments resulting from random projections.

Based on this prior work, we selected random projections as the approach for reducing the dimension of the three snapshot matrices in our data set. Our contributions in this work are as follows:

- We show how we can generate the snapshots matrices when the data available are on spatial domains with different sizes, at grid points that vary in their (x, y) locations across simulations, and in a format that distributes the output across multiple files at each time step of a simulation.
- We discuss how we can apply random projections to extremely high-dimensional data, determine the reduced dimension, and evaluate the results.
- We indicate how we can exploit the randomness resulting both from the random projections, and the initial choice of centroids in k-means clustering, to obtain the number of clusters in the data.

5 Solution approach

Our approach to clustering the very-high dimensional data set resulting from the simulations of a jet interacting with high explosives (HE) requires us first to convert the raw data into a format that we can cluster and then to identify a computationally feasible approach to clustering these high-dimensional snapshots. We next describe how we address these tasks.

5.1 Pre-processing the data

As described in Section 3.1, we have data for 1604 time steps across the 45 simulations. For each simulation, at each time step, the three variables, along with the corresponding (x, y) coordinates of the grid points, are spread across 360 files, each covering a small part of the domain. We need to convert this to a snapshot matrix for each variable, where the matrix has all the grid points, in natural ordering, along the rows, and each column contains the values of the variable at a different snapshot.

At first glance, this appears to be just a re-arrangement of data from one set of files to another. But as we described in Section 3.1, we also need to account for the different ranges of x values, the lack of alignment of the plate even at early time, and the inconsistent values for the (x, y) coordinates across the simulations. This will require computation on the data values, in addition to the re-arrangement of values in files.

We first decided that instead of storing the final data for each variable in a single snapshot matrix file as in Equation 1, we would split the matrix into blocks, with each block stored in a separate file, as follows:

$$\mathbf{X} = \begin{bmatrix} \mathbf{X}_{b1} \\ \mathbf{X}_{b2} \\ \vdots \\ \mathbf{X}_{bk} \end{bmatrix} \quad (2)$$

This would make the processing of the very large number of grid points tractable as we could read the matrix a block at a time. However, all processing would have to be modified to work with the matrix spread across several files. We chose each block to contain all grid points in a specific non-overlapping range of y values. As the grid points are stored in natural order within a block, concatenating the blocks in the order of their y values creates a single snapshot-matrix file with all the grid points stored in natural order.

There are many ways in which we can transform the HDF5 data from 360 files per time step, per simulation, into the three snapshot matrices, one for each variable that represents all the data for the variable across the 45 simulations. The solution we describe next is one we found expedient to implement. We used computationally efficient implementations, exploiting parallelism where possible. As we wanted our codes

to remain flexible for processing other data formats, we kept the steps in this multi-step process distinct, though for higher computational efficiency, we could have merged some steps, or executed them in a different order.

5.1.1 Creating a data file for each time step in a simulation

We started by first processing the raw data for each simulation. For each time step in the simulation, the data are available in 360 files, one for each sub-domain. Each file contains the (x, y) coordinates and all variable values for that sub-domain. In this first step, the 360 files were combined into a single file for each time step in a simulation. This essentially involved reading each file in turn, extracting the coordinates and variable values, and writing out this information, a row at a time, to the single data file. The rows in this file are the grid points in the same order as they appear in the HDF5 file, starting at sub-domain 0 and ending at sub-domain 359. The columns are the (x, y) coordinates and the three variable values. For each time step, we also generated a summary file with statistics on each sub-domain, including the range of x and y values, and the starting location of each sub-domain in the single data file. This file is used to improve the computational efficiency of subsequent steps in the processing of the data.

This step is just a re-arrangement of the data and involves reading a set of files, extracting the relevant information, and writing it out. The time steps across the simulations can all be processed in parallel.

5.1.2 Aligning and cropping coordinates to a common domain

After the first step, we have a large file for all variables at each time step in a simulation, along with an associated summary file. When we consider these files across the 45 simulations, the number of grid points is different as the range of x values is different for each simulation. In addition, as the origin in each simulation is at the start of the HE, the plate is not aligned across simulations at time $t=0$.

To fix this, we first changed the origin of the coordinates in each file to be at the lower right corner of each domain. This shift in x values is the same for all time steps in a simulation, but differs across simulations. This step automatically aligns the plate at time $t = 0$ across all simulations because the left edge of the 1cm wide plate is 15 cm from the right edge of the domain. Such alignment of the data is common in tasks such as face recognition using the eigenface approach [25], where each image is pre-processed so the face covers the full image, is upright, and centered in the image. In our problem, unlike the face, which is stationary, the plate and the HE move as the simulation progresses. By aligning the plate at initial time we reduce the amount of misalignment of the plate across the snapshots.

Next, we cropped the data in each file so that grid points with the new shifted x coordinate outside the range $[-32, 0.0]$ are excluded. This range was selected as it corresponds to the smallest HE length of 5.0cm. After this step, the data files for all simulations have the same range of x values and the same plate location at initial time. However, the values of the variables are at different (x, y) coordinate locations and the grid points are still in the same order as they were within each sub-domain, with the sub-domains stored in order.

These changes make it possible to generate a meaningful clustering of the snapshots; they can be applied in parallel for the time steps across the simulations.

5.1.3 Remapping simulations to a common grid

Next, we used a simple interpolation scheme to remap the data values for each simulation to a common grid. This remapping consists of two steps. We first defined the common grid by choosing an x -range of $[-31.5, -0.5]$, which is slightly smaller than the range for each simulation to ensure that all remapped values were being interpolated, not extrapolated. The y -range was selected as $[0.0063, 10.99]$. We kept the grid resolution the same at $\Delta x = \Delta y = 0.0125$, resulting in a total of $D = 2,180,799$ grid points in the common grid. Then, given the fine resolution of the common grid, we used a simple 1-nearest-neighbor algorithm for interpolation, though more complex algorithms are also an option.

We created the common grid in a block form, as in Equation 2, with each block written to a separate file. A block spans the full range of x values, but a smaller range of y values. We chose $bk = 22$ blocks, with each block representing approximately 0.5cm of the y -range. This gave approximately 100K grid points in the first 21 blocks, and a smaller number in the last block. Within a block, the grid points were listed in natural ordering so that any data remapped to this common grid would have the rows in the form required by the final snapshot matrix.

We generated the remapped data a block at a time. For each time step in a simulation, we extracted a block of data that had a range of (x, y) coordinates slightly larger than the block of the common grid to which we were remapping this data. This ensured that we were interpolating, not extrapolating, to the common grid. For computational efficiency, we used the summary file associated with the time step to identify and process only those sub-domains that had coordinates in the required range.

We also observed that for our data, all time steps of a simulation could be remapped together because the data extracted for a block for the different time steps had values at the same (x, y) coordinates in the same order. So, for each block, we first concatenated the columns at different time steps to create a file for each simulation with $(2 + 3 * \#timesteps)$ columns to store the (x, y) coordinates and the three outputs across all the time steps. Merging these columns is possible as we had maintained the ordering of the data in the sub-domains and a fixed grid was used for all time steps.

This re-mapping for the 45 simulations can be performed in parallel. Once we have remapped the data from each simulation to each block of the common grid, we have all data for the 45 simulations in the same order as in the common grid, which simplifies the next steps discussed in Section 5.1.4

This step of remapping has several benefits. It introduces some flexibility in the analysis as we can reduce the data size by mapping to a coarser grid or remapping only the data in a smaller sub-region of the full domain when the rest of the domain is of lesser interest. By remapping to the common grid that is in the block form of the final snapshot matrix, we can generate the remapped data for each block in parallel, and with the smaller file sizes, process the data in memory as well. However, we observe that the computational efficiencies in the remapping are the result of the same grid being used across time steps in a simulation. Generating the common grid will be more challenging for AMR grids; the remapping will also be more time consuming as it has to be done separately for each time step.

5.1.4 Converting remapped data to snapshot matrices

After the remapping step, we have 22 files for each simulation. Each file corresponds to a block, with columns containing all time steps for the three variables, along with the (x, y) coordinates. The remapping ensures that these coordinates are the same across the simulations.

To create the final snapshot matrices, we split each block for a simulation into three files, one for each variable, and then merge the columns of these files across simulations. For each variable, this gives the values at 2,180,799 grid points, split across 22 files, with each file having 1604 columns, which is the total number of time steps across the simulations. These snapshot matrices do not include the (x, y) coordinates.

Figure 5 shows the mass variable for the first and last snapshot, for each of our four example simulations, after the raw output data have been aligned, cropped, and remapped. The corresponding images for x-momentum and y-momentum are shown in Figures 13 and 18 in Appendix A and B, respectively.

5.2 Clustering the snapshots

Once the three snapshot matrices, corresponding to the three output variables, are created in block form, we can cluster the columns in the matrices. In Section 3.2, we described how the high dimensionality of our data can be an issue for iterative clustering algorithms as the full snapshot matrix is too large to fit into memory. There are two obvious solutions to this problem of high dimensionality that we discuss next.

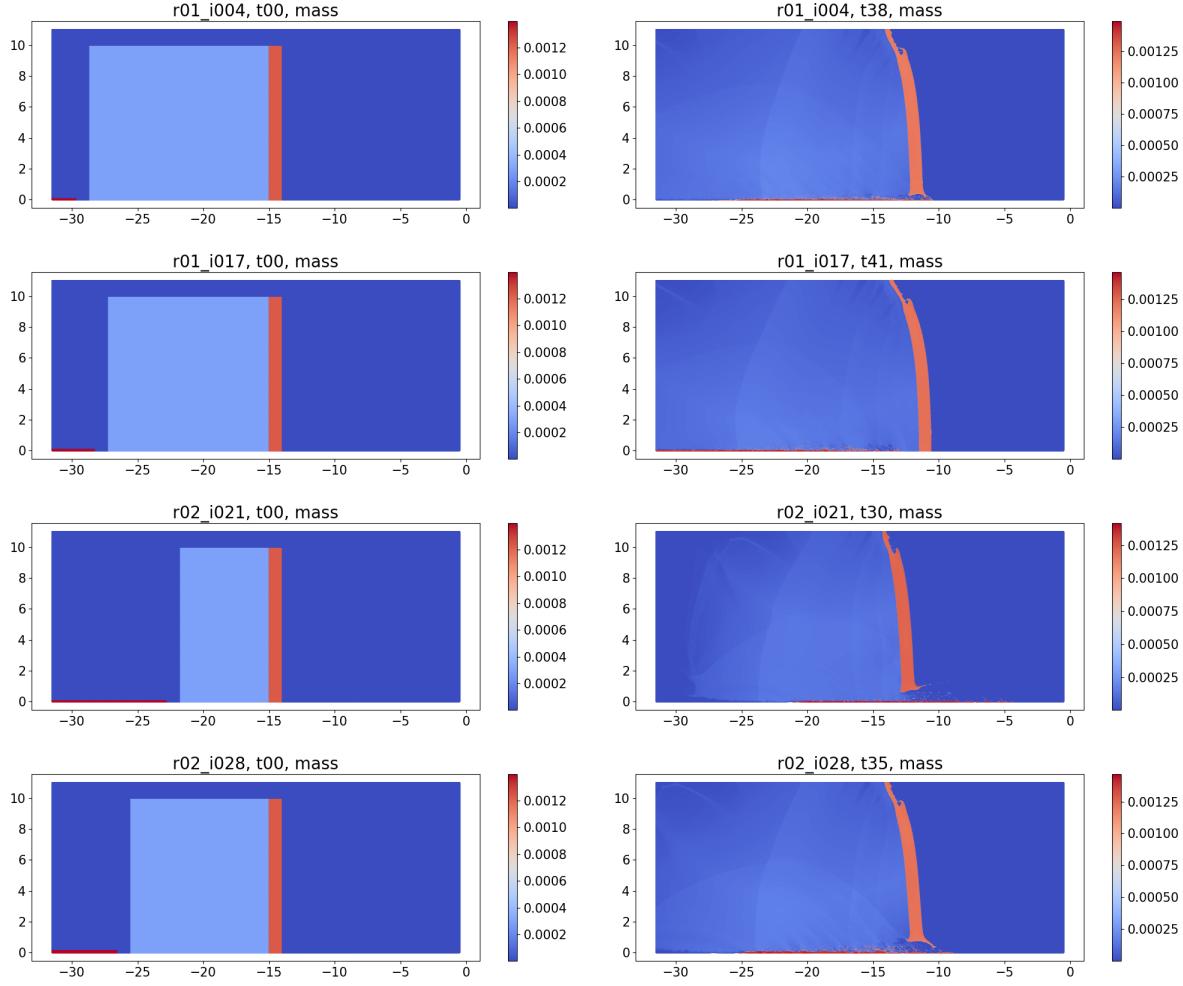


Figure 5: The variable *mass* for the four example simulations in Table 1, after the snapshots have been aligned, cropped, and remapped. Left column shows the first time step and right column shows the last time step. From top to bottom, simulations with keys r01_i004, r01_i017, r02_i021, r02_i028, respectively. The color bars are different between simulations and across time steps.

5.2.1 Using an iterative method with dimension reduction

Our first solution is to use the iterative k-means clustering algorithm, but with modifications to account for the high dimensionality of the data. One way to achieve this is by using a distributed-memory version of the algorithm, which operates on data divided into blocks (as in Equation 2). We can assign each data block, and the corresponding block of the k-means centroids, to a processor and calculate, in parallel, the partial distance-squared of each snapshot to each centroid. Then a communication step would add the partial distances and identify the closest centroid to each snapshot, followed by an update of each block of the centroids. This approach has limited parallelism equal to the number of blocks and would require implementing a distributed-memory version of the k-means algorithm.

Alternately, we could reduce the dimension of the data so that the transformed matrix fits into the memory of a single processor. To accomplish this, we used random projections [1, 2] which projects the D -dimensional matrix X onto a d -dimensional subspace, with $d \ll D$ using a random matrix $R \in \mathbb{R}^{d \times D}$

$$X_{d \times N}^{RP} = R_{d \times D} X_{D \times N} \quad (3)$$

We can justify the use of random projections as the Johnson-Lindenstrauss lemma [7] states that for any distortion ϵ , with $0 < \epsilon < 1$, and any integer N , if d is a positive integer such that

$$d \geq 4 \left(\frac{\epsilon^2}{2} - \frac{\epsilon^3}{3} \right)^{-1} \ln N, \quad (4)$$

then, for any set X of N points in \mathbb{R}^D , there is a map, $f : \mathbb{R}^D \rightarrow \mathbb{R}^d$ such that for all $u, v \in X$,

$$(1 - \epsilon) \|u - v\|^2 \leq \|f(u) - f(v)\|^2 \leq (1 + \epsilon) \|u - v\|^2. \quad (5)$$

In other words, points in a sufficiently high-dimensional space can be projected onto a suitable lower dimensional space, while approximately preserving the distances between the points. As the k-means algorithm is based on Euclidean distances, this means we can cluster the randomly-projected snapshot matrix X^{RP} instead of the original snapshot matrix X and expect to get approximately similar clustering results.

To apply the map f in Equation 5, in the form of a random projection of our snapshot matrix, X , we need to define the random matrix, R , and determine the new reduced dimension, d . As our snapshot matrix is extremely high dimensional, we prefer a sparse matrix, such as the one proposed by Achlioptas [1, 2] with i.i.d. entries as:

$$r_{i,j} = \sqrt{s} \begin{cases} 1 & \text{with probability } \frac{1}{2s} \\ 0 & \text{with probability } 1 - \frac{1}{s} \\ -1 & \text{with probability } \frac{1}{2s} \end{cases} \quad (6)$$

where $s = 1$ or $s = 3$, with the former resulting in a dense projection matrix. Li, Hastie, and Church [19], proposed using an even higher value of $s = \sqrt{D}$. For our problem, with $D = 2,180,799$, this choice results in a very sparse matrix, whose sparsity can be exploited for computational efficiency. As the random projection changes the distances between two snapshots by $(1 \pm \epsilon)$, we can obtain the new dimension d by first choosing a value of ϵ , that is, the distortion we can tolerate, and use $N = 1604$ in Equation 4. For $\epsilon = 0.1$, we get $d > 6325$, while $\epsilon = 0.05$ gives $d \geq 24431$ and $\epsilon = 0.01$ gives $d \geq 594383$.

It has been observed that values of d less than the ones suggested by Equation 4 work well in practice [5]. We chose d based on both the distortion in distances we can tolerate and the space available to store the randomly-projected snapshot matrix X^{RP} in memory. We generated X^{RP} incrementally by reading in each block of the snapshot matrix X a row at a time, generating a column of the sparse random matrix R , and calculating the outer product of the column and row and adding it to X^{RP} . By not storing either R or a block of X , we can choose a larger value for d as we require storage only for X^{RP} .

Once we have obtained the randomly projected snapshot matrix X^{RP} of size $d \times 1604$, we can apply the k-means algorithm to cluster the randomly-projected snapshots. The clustering results so obtained could vary based on the randomness of the projection and the random initial choice of the cluster centroids. We discuss these issues, and our choice of parameters for random projections and k-means, further in Section 6.

5.2.2 Using an iterative method with a reduced representation

An alternative that is similar in spirit to the use of random projections, but specific to our approach to building spatio-temporal surrogates, is to exploit an intermediate step in the building of these surrogates. As explained in the companion report [16], we perform a singular value decomposition (SVD) on the snapshot matrix \mathbf{X} from Equation 1 to obtain

$$\mathbf{X} = \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T \quad (7)$$

where

$$\mathbf{X} \in \mathbb{R}^{D \times N}, \quad \mathbf{U} \in \mathbb{R}^{D \times D}, \quad \boldsymbol{\Sigma} \in \mathbb{R}^{D \times N}, \quad \text{and} \quad \mathbf{V} \in \mathbb{R}^{N \times N}.$$

Here, $\boldsymbol{\Sigma}$ is a diagonal matrix whose non-zero diagonal elements, σ_{ii} , are the singular values of the matrix \mathbf{X} and the columns of the orthonormal matrices \mathbf{U} and \mathbf{V} are the left and right singular vectors of \mathbf{X} . The rows and the columns of the \mathbf{U} and \mathbf{V} matrices are ordered such that the singular values are in descending order in $\boldsymbol{\Sigma}$. Since $D \gg N$, only the top N rows of $\boldsymbol{\Sigma}$ will have non-zero diagonal elements, assuming $\text{rank}(\mathbf{X}) = N$.

N . If the rank, r , is less than N , then only the first r diagonal elements are non-zero. The columns of the matrix

$$\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_N] \quad \text{where } \mathbf{u}_i \in \mathbb{R}^D \quad (8)$$

form a basis in \mathbb{R}^D for the data, so each snapshot can be written as a linear combination of the \mathbf{u}_i :

$$\mathbf{x}_i = \sum_{k=1}^N w_{ki} \mathbf{u}_k \quad (9)$$

where the coefficient w_{ki} of the k -th basis for the i -th snapshot is

$$w_{ki} = \mathbf{u}_k^T \mathbf{x}_i \quad \text{for } k = 1, \dots, N. \quad (10)$$

In building the spatio-temporal surrogate, we consider only the more important \mathbf{u}_i corresponding to the larger singular values, and, by truncating the summation in Equation 9, we create a reduced representation. However, for the purpose of clustering the snapshots, we can view the reduced representation of the i -th snapshot as the vector of the N weights

$$\mathbf{w}_i = [w_{1i}, w_{2i}, \dots, w_{Ni}], \quad (11)$$

and then obtain a clustering of the snapshots by clustering the vectors of weights that define each snapshot. If only the few initial weights are calculated, they can be used as approximate representations of a snapshot.

In efforts where we want to compare the quality of a global spatio-temporal surrogate, built using all snapshots, with a set of local surrogates created after clustering the snapshots, we can generate the cluster assignment directly using the weights obtained from the SVD of the full data set (Equations 7 and 10), avoiding any additional computation. However, if we are creating only the local surrogates, this approach would be computationally more expensive than using random projections with sparse matrices (Section 5.2.1).

In this report, we considered the option of clustering the snapshot using the weights obtained from the SVD as it allows us to compare the cluster assignment obtained from an exact representation of the data with that obtained from an approximate representation generated by random projections.

5.2.3 Using a non-iterative method

An alternative to using an iterative clustering algorithm, which is made tractable using distributed memory processing or a dimension reduction technique, is a clustering algorithm that does not require iterating over the snapshot matrix. One such method is hierarchical clustering [13, 11], a greedy algorithm which, in its agglomerative form, begins with every data point in its own cluster. At each step a linkage criterion is used to measure the similarity between all the clusters, and the two most similar clusters are merged together. This process of merging clusters continues until the desired number of clusters are obtained or until the clusters no longer satisfy a desired level of similarity [27]. Hierarchical clustering is usually accompanied by a dendrogram, so this stopping condition can be visualized as making a cut across the dendrogram at the specified similarity threshold to create a set of clusters. Hierarchical clustering requires the following parameters: a distance metric, often Euclidean; a linkage criterion; and the number of clusters.

One of the benefits of hierarchical clustering in processing extremely high-dimensional data is that it is not an iterative algorithm, requiring only the pairwise distance matrix between the data points. This allowed us to generate results using the original data set, without the need to create an approximation using random projections. While calculating the pairwise distance matrix is computationally expensive, it can be computed once and then used to test the algorithm with several different parameters.

6 Results and discussion

We next present the results of clustering the snapshots for the three variables using three methods: i) random projections followed by k-means, ii) k-means using the weight vectors from SVD, and iii) hierarchical

Algorithm 2 Hierarchical clustering

- 1: Goal: given the N snapshots in a D -dimensional space in the form of the matrix X , identify a clustering with nc clusters, such that each snapshot is assigned to a cluster, by merging clusters with the highest similarity.
- 2: Set the linkage criterion and the number of clusters, nc .
- 3: Initialize the clusters by putting every snapshot in its own cluster.
- 4: Compute the similarity between all pairs of clusters using the selected linkage criterion and distance metric.
- 5: Merge the two clusters that are the most similar into one cluster.
- 6: Repeat steps 4 and 5 until the desired number of nc clusters has been obtained or the similarity between clusters no longer satisfies a specified level.

clustering on the original snapshot matrix. We also discuss how we set various parameters in the algorithms to generate a cluster assignment for each snapshot.

All codes, unless otherwise stated, were written in C++ and Python. We exploited parallelism where possible through the use of the sub-process capability in Python, but did not explicitly optimize the conversion of the HDF5 files to the snapshot matrices, which was the more time consuming part of our solution.

We begin by first taking a closer look at the data to be clustered, now focusing on multiple time steps in two example simulations, key r01_i017 and key r02_i028, after the original output has been pre-processed as discussed in Section 5.1. Figure 6 shows how the mass variable evolves with time in these two simulations.

We again observe that the first and last time steps in the example simulations are very different, suggesting there are at least two clusters in the data. In addition, there is similarity in intermediate time steps, even though one simulation is a break case and the other a no-break case. This suggests that there is inherent clustering in the data, though we expect the clusters will not be well separated as any clustering would assign some neighboring time steps, which are very similar, to different clusters. The corresponding images for x-momentum and y-momentum are shown in Figures 14 and 19 in Appendix A and B, respectively. Unlike the mass variable, the intermediate time steps are less similar between the break and no-break case.

6.1 Results with random projections and k-means clustering

Our approach to clustering using random projections followed by k-means clustering, introduces randomness in the projections and in the choice of initial cluster centers, both of which can influence the assignment of snapshots to clusters. If this influence is large, it can result in uncertainty in cluster assignment.

6.1.1 Understanding the effect of the randomness of random projections

We first empirically evaluated the effect of the randomness in the projection by repeating the projection of the snapshot matrix twice for a reduced dimension $d = 2200$. This new dimension gives nearly a 1000-fold reduction from the original dimension of $D = 2,180,799$. We calculated the original and reduced-dimensional distances between all snapshots and the 29 snapshots of the *baseline* simulation with HE length, jet tip velocity, and jet radius equal to 5.0cm, 0.950cm/ μ s, and 0.125cm, respectively. This simulation is at an extreme corner of the input space. The results for the mass variable are shown in Figure 7; the corresponding figures for the x- and y-momentum are shown in Figures 15 and 20 in Appendix A and B, respectively.

These figures indicate that for our data set, the randomness of the random projections has little effect on the distances between snapshots, unlike the results of Fern and Brodley [10], who found random projections to be unstable for clustering. Though we present the results only for two repetitions, we observed similar behavior in multiple repetitions. In addition, as has been observed by others [5], we can use a smaller reduced dimension than the one suggested by the Johnson-Lindenstrauss lemma (Equation 4) for a specific value of the distortion ϵ . Our choice of $d = 2200$, which is much smaller than the dimension of 6325 suggested by the

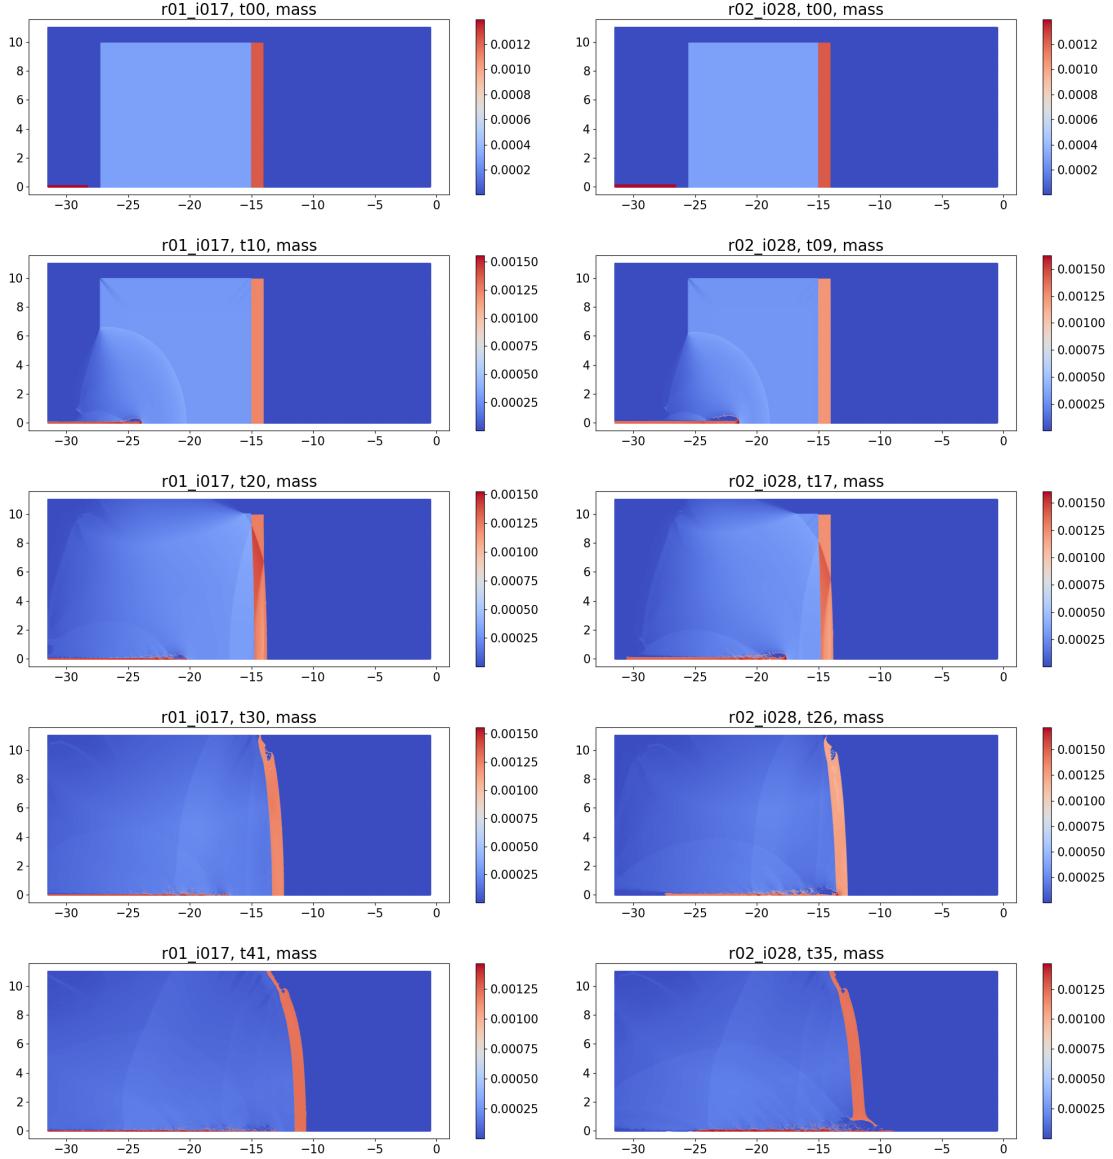


Figure 6: The variable *mass*, after the snapshots have been aligned, cropped, and remapped, at different time steps in two simulations showing the evolution of the data over time. Left: key r01_i017 (no break case) at time steps t0, t10, t20, t30, and t41. Right: key r02_i028 (break case) at time steps t0, t09, t17, t26, and t35. The color bars are different between simulations and across time steps.

lemma for $\epsilon = 0.1$, results in a distortion less than 0.05 for most snapshots.

6.1.2 Exploiting the randomness of k-means to determine number of clusters

Next we applied the k-means algorithm to the randomly-projected snapshot matrix. As the results depend on the random choice of the initial centroids, we used ensemble clustering, where we repeat the algorithm to generate results for 10 different initial cluster centroids. In our implementation of k-means (Algorithm 1), we set *niter*, the maximum number of iterations to 100 and *thresh*, the maximum distance moved by any centroid between iterations to be 0.0. For our data, the latter condition is satisfied before 100 iterations.

Another parameter we need to select is the number of clusters, *nc*. A typical approach used in building

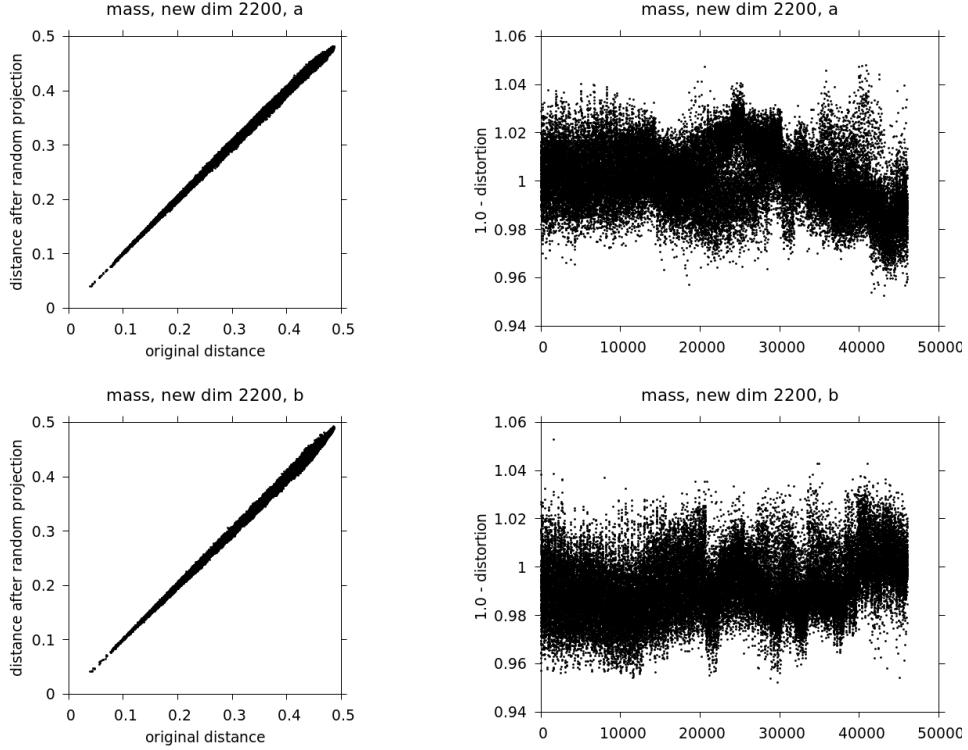


Figure 7: For the mass variable, evaluating the choice of reduced dimension and randomization on the distances between the 29 snapshots from the baseline simulation to all 1604 snapshots. Left - actual vs. random-projected distances for two repetitions of random projections indicating that the randomness of the random projections has little effect on the new distances. Right - the corresponding values of $(1.0 - \epsilon)$ indicating that the distortion ϵ is much less than 10%, even though our choice of reduced dimension of 2200 is smaller than the 6325 suggested by the Johnson-Lindenstrauss lemma for $\epsilon = 0.1$.

spatio-temporal surrogates is to generate results for different values of nc , and select the best based on some metric. This metric can be the quality of reconstruction of test snapshots [3], which is computationally expensive for a large data set like ours, or silhouette analysis [6], which is useful when the clusters are compact and clearly separated [22], which is not the case for our data. In practice, the number of clusters is a trade-off between larger values that give a better approximation of the non-linear manifold, and smaller values that make the class assignment of the snapshots more stable.

To determine the number of clusters, we first tried the iterative consensus clustering (ICC) method [20], which had its own parameters that were difficult to set. It required generating results with a range of values for nc and combining them into a consensus matrix, which identifies the number of times two snapshots are clustered together. We found the results to be sensitive to the range of values of nc ; they also did not give a clear indication of the number of clusters inherent in the data. In addition, once the number of clusters had been identified, the method did not provide a way to obtain a cluster assignment for the snapshots.

However, we realized that a consensus matrix could be exploited to understand the sensitivity of the clustering results to both the number of clusters and the randomness of the random projections, and, in addition, to identify the number of clusters. For our data set with 1604 snapshots, we were interested in moderate-to large-sized clusters where possible. We ran the k-means ensemble clustering for small values of $nc = 3, 4, 5$, and 6, with 10 repetitions each, changing the initial centroids each time. Unlike the ICC method that generated a single consensus matrix combining all these results, we generated a separate consensus matrix for each nc and for each of the two repetitions of the random projections (referred to as a and b). Our consensus matrices represented the fraction of times two snapshots were in the same cluster; with matrix entries between 0.0 and 1.0, it became easier to set parameters for subsequent processing.

Using the mass variable as an example, we next describe how we analyze the consensus matrices to identify the number of clusters and the cluster assignment for the snapshots. Table 2 shows the statistics on the eight consensus matrices. For $nc = 3$ and 4, most of the values are either 0.0 (snapshots never in same cluster) or 1.0 (snapshots always in same cluster). However, as the number of clusters increases to $nc = 5$ and 6, there is a range of values, indicating that some snapshots are clustered together only occasionally. This suggests that for larger nc , the results are sensitive to the initial cluster centroids. Further, for a given value of nc , there is less difference between the results of the two random projections when $nc = 3$ and 4, than when $nc = 5$ and 6, that is, the clusters are more stable for lower nc . This last observation does not necessarily hold for the x- and y-momentum variables as shown in Tables 3 and 4 in Appendix A and B, respectively.

Values	nc3, a	nc3, b	nc4, a	nc4, b	nc5, a	nc5, b	nc6, a	nc6, b
0.0	59.18	59.34	63.72	63.75	64.55	63.43	65.01	65.73
0.1	0.0	0.0	0.91	0.01	0.01	0.0	12.86	1.15
0.2	0.0	0.0	0.65	0.89	0.0	0.91	1.67	1.65
0.3	0.0	0.0	0.0	0.63	13.81	0.59	1.19	0.40
0.4	0.0	0.0	0.0	0.0	0.09	0.003	0.52	11.62
0.5	0.41	0.0	0.0	0.0	0.0	16.95	0.04	1.50
0.6	0.0	0.0	0.0	0.0	0.01	0.03	0.63	1.49
0.7	0.0	0.0	0.0	0.39	5.27	1.31	0.80	0.41
0.8	0.0	0.0	0.51	0.63	0.0	1.01	1.14	2.12
0.9	0.0	0.0	0.57	0.02	0.09	0.03	2.57	0.69
1.0	40.41	40.66	33.65	33.68	16.17	15.74	13.58	13.24

Table 2: Distribution of values in the consensus matrices for the mass variable, with 10 repetitions of the clustering. We used four values for nc , the number of clusters, and two repetitions of the random projections, indicated by a and b. The values in the table indicate how often (in percentage) the value in the first column occurs in the matrix. Thus, for $(nc = 3, a)$ in the second column, nearly 60% of the matrix entries are 0.0, nearly 40% are 1.0, with a small percentage of entries at 0.5. In contrast, for $nc = 6$, both a and b, all values between 0.0 and 1.0 appear in the consensus matrix, and only 79% of the entries are either 0 or 1, which indicates how often snapshots are never or always in the same cluster, respectively.

Figure 8 shows a small 108×108 subset of the consensus matrices for $(nc = 3, a)$ and $(nc = 5, a)$ that further highlight the differences between the small and large values of nc . The subsets show the first three simulations with 29, 38, and 41 snapshots, respectively, in order of the time steps, starting with the first time step of the first simulation, and ending at the 41-st time step of the third simulation. As described in the caption, we can visually identify which snapshots in which simulations are in the same cluster.

This observation gives us a simple approach to cluster assignment for the snapshots. We start with the first row of a consensus matrix, and put all snapshots with a value 1.0 in this row into one cluster. We repeat with the rows corresponding to each of these snapshots, and so on, until we have identified all snapshots in the first cluster. Then, we identify a snapshot not yet assigned to a cluster, and repeat to create the second cluster. And so on. This approach gives us the class assignment for the snapshots. However, we found that if we only combine snapshots that are always grouped together (that is, matrix entries with value 1.0), we will identify more clusters than nc because a snapshot that occurs together with another only nine times out of ten (with a value of 0.9 in the matrix) would be in a different cluster.

We observed this behavior for both $(nc = 3, a)$ and $(nc = 5, a)$ as shown in Figure 9 using the full consensus matrix with the snapshots reordered by the cluster number. For $(nc = 3, a)$, we have five clusters, the three clearly seen in the figure and two small ones with 4 and 2 snapshots that appear at the end. In contrast, $(nc = 5, a)$ has nine clusters instead of 5, as indicated by the diagonal blocks with values equal to 1.0. If we identified the clusters using matrix entries greater than a threshold, say 0.7, the results were sensitive to the threshold. We also had to address the issue of entries with values equal to 0.5, which occurs when

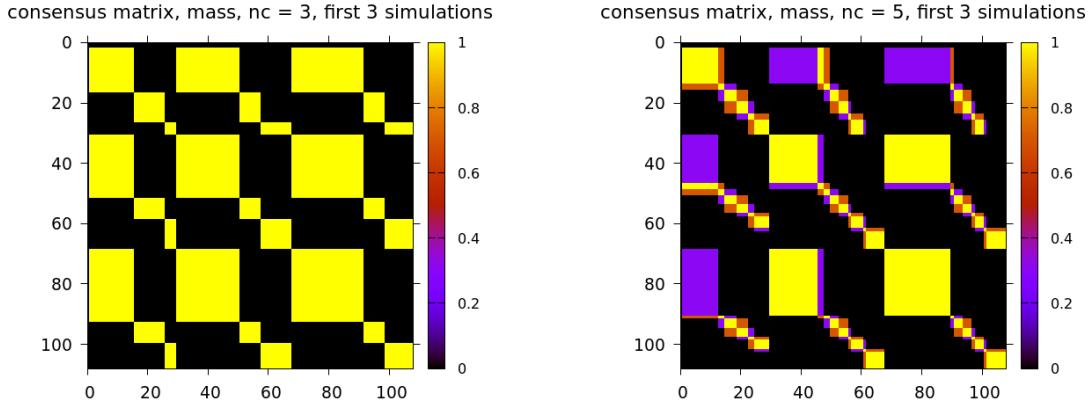


Figure 8: A 108×108 subset of the consensus matrix for mass for $(nc = 3, a)$ (left) and $(nc = 5, a)$ (right). The results are for the first three simulations with 29, 38, and 41 snapshots, respectively, shown in order, with the first snapshot of the first simulation at top left corner, and the 41-st snapshot of the third simulation at the bottom right corner. The left figure indicates that the early time steps of the three simulations are in the same cluster, and that a small number of snapshots at late time in the first simulation are in the same cluster as the late time steps in the other two simulations. The plot on the right has a larger range of values, showing that the early time steps of the first simulation are less often in the same cluster as the early time steps of the second and third simulations, while the latter two are usually in the same cluster. Also, the cluster assignment of later time steps in a simulation is less clear with smaller clusters.

neighboring time steps are at a cluster boundary and assigned to one cluster or the other half the time.

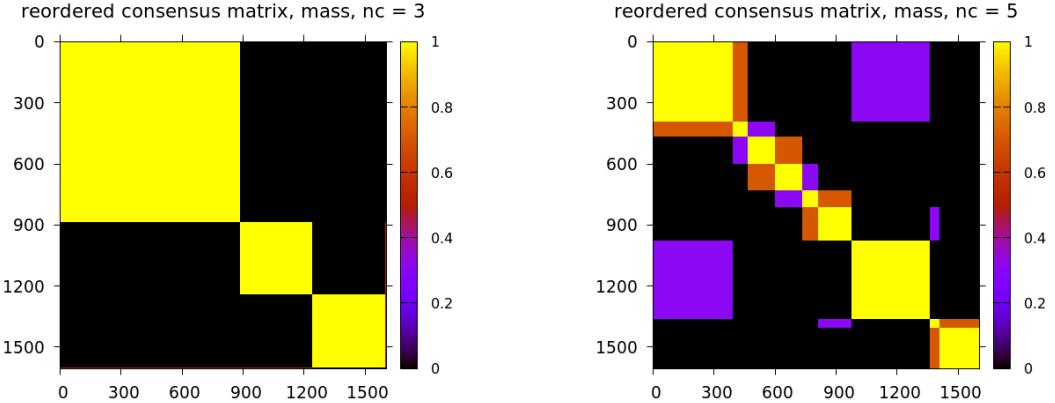


Figure 9: The reordered consensus matrix for mass showing all snapshots reordered by class. Left: matrix for $(nc3, a)$; right: matrix for $(nc5, a)$.

To understand how the time steps for the 45 simulations are distributed across these 5 and 9 clusters, Figure 10 shows the cluster assignment for each snapshot as a function of HE length. For $(nc = 3, a)$, there are two small clusters, with 4 and 2 snapshots, that are at the time step boundary between two clusters and could be assigned to either. For $(nc = 5, a)$, there are several small clusters that are connected to other clusters either weakly or strongly, with matrix values 0.3 and 0.7, respectively, based on Table 2. An option in this specific case is to use the reordered consensus matrix to merge a small cluster with another if they are strongly connected. However, when the consensus matrix entries take values closer to 0.5, it becomes less clear which clusters should be merged and why. In addition, when a small group of consecutive (in time) snapshots in a simulation spans 4 or 5 clusters, as shown in Figure 10, for $(nc = 5, a)$ at large values of HE length, it is not clear whether these snapshots should remain with their original large cluster, or be merged into the cluster of neighboring (in time) snapshots. We suspect that this poor cluster assignment is the result

of either too large a value of nc and/or the clustering of snapshots being more sensitive to the initial choice of centroids for larger nc , resulting in values in the consensus matrix other than 0.0, 0.5, or 1.0.

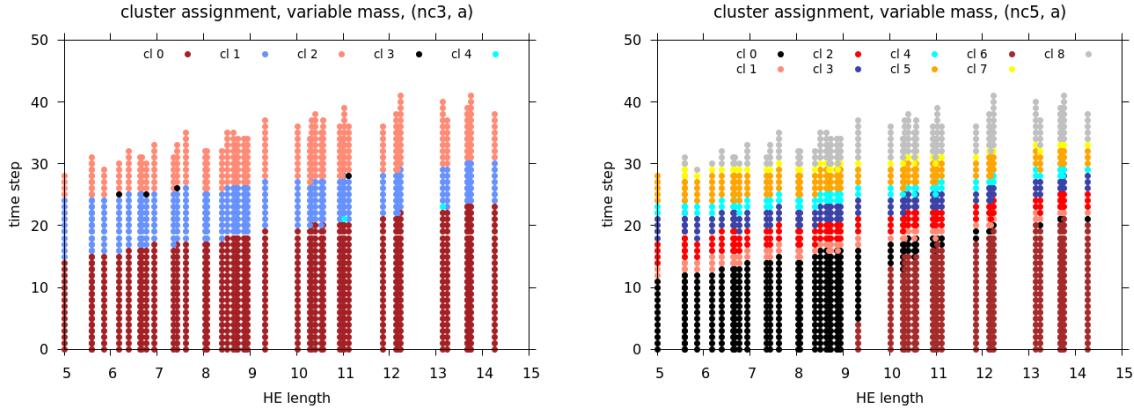


Figure 10: The cluster assignment for mass for the time steps in the simulations plot along with the HE length. The points are colored by their class. Left: cluster assignment for ($nc = 3$, a) (left) and ($nc = 5$, a) (right). In the left plot, note the two very small clusters, one in black with 4 snapshots and the other in cyan with 2 snapshots, indicating entries in the consensus matrix with a value of 0.5. On the right, the multiple small clusters are the result of a larger $nc = 5$ and the relatively large number of matrix entries at 0.7 and 0.3 that form their own clusters. Note that the simulation at the smallest HE length, which is the baseline simulation, does not have any time steps in cluster 8, shown in grey, which is of greatest interest in our problem as it has the late time data for the rest of the simulations.

This analysis suggests that we should select the number of clusters that gives more repeatable results in the ensemble as it reflects a more stable clustering. Using this criterion, we chose the cluster assignment identified by ($nc = 3$, b) for the mass variable as it has only 0.0 and 1.0 in the consensus matrix; the result is shown in Figure 11. The results for the x- and y-momentum variables, using random projections and k-means clustering, are shown in Appendix A and B, respectively, and include the distribution of values in the consensus matrix (Tables 3 and 4), along with the cluster assignment (Figures 16 and 21). These results indicate 4 clusters for the x- and y-momentum variables. We discuss these results further in Section 6.4.

6.2 Results with k-means clustering using weights from SVD

Next, we applied k-means to the weights that formed the reduced representation of each snapshot obtained using the SVD of the full snapshot matrix as described in Section 5.2.2. We used the number of clusters that were identified using random projections with k-means (Section 6.1.2) for the three variables. As a sanity check, we repeated the clustering ten times for each variable and confirmed that the clustering results presented in Figures 12, 17, and 22 for the mass, x- and y- momentum, respectively, were stable. These results are discussed further in Section 6.4.

6.3 Results with hierarchical clustering

We used the implementation of hierarchical clustering available in the SciPy Python library [26], version 1.10.1. This method requires three parameters: the number of clusters, a linkage criterion, and a distance metric between two snapshots. We considered both three and four clusters for each of the three variables as these were the most likely values for nc identified using the consensus matrix analysis (Section 6.1.2). We wanted to evaluate whether hierarchical clustering on the original matrix \mathbf{X} would give different results from the k-means method. To allow experimentation with different parameters, we precomputed the pairwise distance matrix using Euclidean distances. We used a serial implementation for the distance calculation, though the block form of the snapshot matrix (Equation 2) could be exploited for parallel computation. As there is no randomness in hierarchical clustering, we ran the method once.

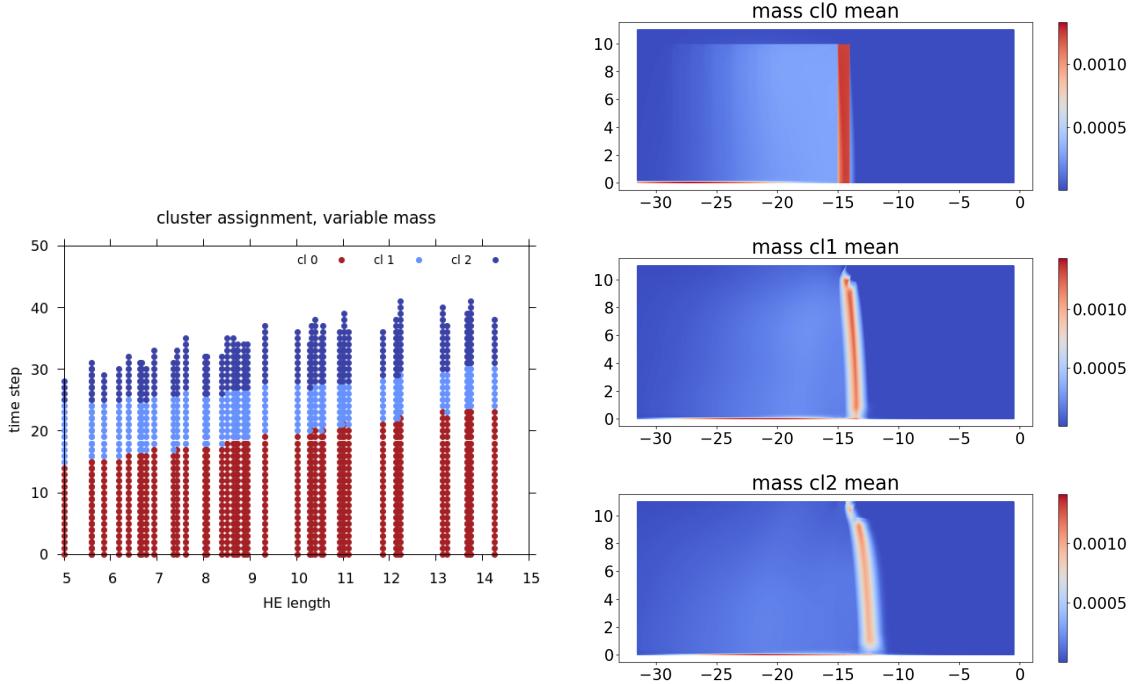


Figure 11: Left: the cluster assignment, corresponding to $(nc = 3, b)$, selected for the mass variable shown for the times steps in the simulations plot along with the HE length. The cluster sizes for clusters 0, 1, and 2, are 889, 351, and 364, respectively. Right: the mean snapshot for the three clusters for the mass variable, from top to bottom, cluster 0, 1, and 2, that show the movement of the plate over time.

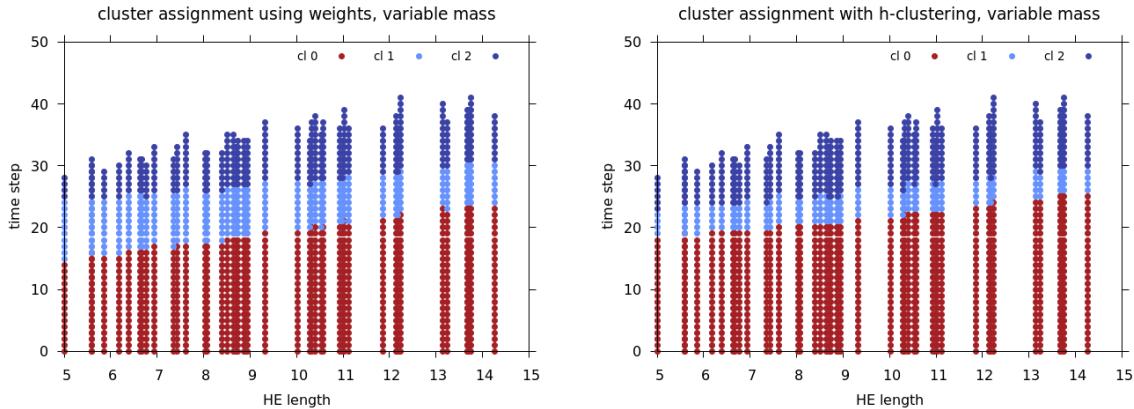


Figure 12: Cluster assignment for the mass variable shown for the times steps in the simulations plot along with the HE length. Left: clustering SVD weights, with $nc = 3$; the cluster sizes for clusters 0, 1, and 2, are 890, 352, and 362, respectively. Right: using hierarchical clustering with the Ward linkage criterion, with $nc = 3$; the cluster sizes for clusters 0, 1, and 2, are 995, 199, and 410, respectively.

SciPy provides several linkage criteria that define the similarity (or dissimilarity) between clusters; at each step, the two clusters that are the most similar, or least dissimilar, are merged. For the single, complete, and average linkage criteria, the dissimilarity between two clusters is defined as the minimum, maximum, and average Euclidean distance, respectively, between two snapshots, one in each of the two clusters. The Ward linkage defines the dissimilarity between two clusters as the increase in the error sum of squares (*ESS*) when the two clusters are merged. If \mathbf{x}_i is a snapshot in a cluster with n snapshots, whose mean is $\bar{\mathbf{x}}$, then

ESS of a cluster is defined as:

$$ESS = \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})^2. \quad (12)$$

Note the similarity between the definition of Ward linkage and the k-means algorithm, where each snapshot is assigned to the closest centroid.

We evaluated the different linkage criteria using two quality metrics. First, we wanted the clusters to be of moderate size, especially the cluster at late time that was of interest in our problem. Second, as neighboring time steps in a simulation are similar, we wanted all the snapshots in each simulation that are assigned to a cluster to be at contiguous time steps. We observed the following:

- As neighboring time steps in a simulation are similar, we expected the single linkage criterion would place all or most snapshots within a single cluster. We found this to be true for all three variables.
- The complete linkage criterion on the mass variable created a small cluster of late time snapshots, but only for the simulations with moderate to large HE length, excluding similar snapshots at late time steps for small HE length. For the x-momentum variable, complete linkage split some clusters across non-contiguous time steps in a simulation, while results for the y-momentum variable were acceptable.
- Average linkage for the mass variable gave results similar to those with complete linkage, but performed poorly for the other two variables. For the x-momentum variable, with $nc = 3$, it created a single cluster composed of the late time snapshots in two simulations at small HE length; for $nc = 4$, this cluster was split into two, one for each simulation. For the y-momentum variable, with $nc = 3$, some simulations had all snapshots assigned to one cluster, while for $nc = 4$, the early and late-time snapshots were merged into a single cluster, while the mid-time snapshots were in different clusters.
- Ward linkage generated the most meaningful cluster assignments. These are shown in Figures 12, 17, and 22 for the mass, x- and y- momentum variables, respectively, and discussed further in Section 6.4.

6.4 Discussion

We next discuss the cluster assignment results from the three clustering methods shown in Figures 11 and 12 for the mass variable, Figures 16 and 17 in Appendix A for the x-momentum variable, and Figures 21 and 22 in Appendix B for the y-momentum variable. These figures also include the mean snapshots for each cluster generated using the cluster assignment from the k-means algorithm with random projections. These results, and our experiences, indicate the following:

- We found that the theory behind random projections and the Johnson-Lindenstrauss lemma works in practice. As observed by others [5], it is possible to obtain low distortion in the projected data even though the reduced dimension is smaller than the value recommended for a specific distortion. For our data set, we obtained a thousand-fold reduction in dimension with $\approx 5\%$ distortion of the data.
- We can optimize the calculation of the randomly-projected data by using a very sparse random projection matrix [19] and reordering the computation such that the projected matrix is generated incrementally by reading in the data one row at a time and generating the random matrix on the fly, a column at a time. This initial step required for the k-means method takes roughly the same time as the calculation of the distance matrix for the hierarchical clustering, making the two methods competitive.
- The results obtained by clustering using k-means and the weights from the SVD of the snapshot matrix are close to those obtained using random projections and k-means. This indicates that the approximation resulting from random projections has little effect on the clustering results.
- Selecting a value for nc , the number of clusters, and generating a class assignment for the snapshots, remains challenging. For k-means, we used the randomness in choice of initial centroids and random projections, as well as the structure in the consensus matrix, to identify a stable clustering. For our data, we found smaller values of nc gave more consistent results, but the similarity between snapshots

at consecutive time steps resulted in uncertainty in cluster assignments for some snapshots. We could exploit domain information to assign such snapshots in very small clusters to other larger clusters, but this may not be an option if there are too many small clusters and it is not clear which clusters should be merged. For hierarchical clustering, an analysis of the dendograms did not shed any light on the appropriate number of clusters; we therefore used the number of clusters obtained from random projections and k-means clustering.

- Hierarchical clustering is sensitive to the choice of the linkage criterion, which is expected. As has long been observed [13], the Ward linkage produces the best results, generating a meaningful cluster assignment for the time steps in a simulation. This might be expected, given the similarity between the Ward linkage and the k-means algorithm.
- The plots of the cluster assignment for the snapshots and the HE length helped to evaluate whether the results are meaningful. As expected, each cluster captures the behavior over contiguous (in time) snapshots, with the early time snapshots in one cluster, the late time in another, and the mid-time snapshots in one or more clusters. This structure in the plots, where each cluster, assigned a different color, appears as a band, is clearer for the mass and y-momentum variables, but less so for the x-momentum, especially with hierarchical clustering, where the banded structure is less well defined (Figure 17).
- Comparing the clustering results for the three variables, we see that the assignment of snapshots to clusters, as well as the sizes of the clusters, are very different.

For the mass variable, the early time cluster is the largest. For the two methods based on k-means, the mid- and late-time clusters are of the same size, while hierarchical clustering gives a smaller mid-time cluster and larger early- and late-time clusters. Each simulation has snapshots in all three clusters.

This is not the case for the x-momentum variable, where the simulations with small HE length have time steps assigned to just three of the four clusters. This is because the cluster to which the very early time steps are assigned depends on the value of HE length. At early time, the x-momentum variable is dominated by the jet that enters from the left of the domain, moving to the right. Since we cropped the simulations with larger HE length on the left during pre-processing, the jet is not seen in these simulations in the very early time steps, which are therefore assigned to a different cluster.

For the y-momentum variable, we observe two unusual aspects of the cluster assignment. First, all clustering methods create a relatively small cluster just after early time; this cluster is wedge shaped, thin for small HE length and growing wider as the HE length increases (Figures 21 and 22). This cluster captures what happens as the wave that emanates from the tip of the jet (Figure 19) reaches the top of the domain. Second, when we use the k-means method, a small number of snapshots at low values of HE length, which appear to be part of the wedge-shaped cluster, are in fact assigned to the late time cluster, which is now split across time steps. This result was consistent across multiple runs of the k-means ensemble, for clustering using both the weights from the SVD and the randomly-projected snapshot matrix with runs a and b, though the apparently mis-assigned snapshots varied slightly. We suspect that these snapshots may indeed be closer to the centroid of the late-time cluster, if only by a tiny amount. Changing the cluster assignment of these 13 snapshots to avoid splitting the late-time cluster across time steps resulted in minor changes in the mean snapshots of the corresponding clusters (Figure 23); we selected this as the cluster assignment for the y-momentum variable. Note that the hierarchical clustering does not split the late time cluster.

- Comparing the different clustering methods, we find that clustering using k-means, with either the random-projected snapshots or the weights from SVD, works quite well. It is a simple method, and though it was proposed several decades ago, it is still an admissible method [12]. As it is iterative, issues such as a poor initial choice of centroids, or a snapshot assigned to the wrong cluster, are fixed in later iterations, unlike hierarchical clustering. For our data set, with the high-dimensional snapshot matrix in block form, k-means requires the reduction in dimension through random projection, while hierarchical clustering requires calculation of the distance matrix; both take roughly the same compute time. We were able to use the randomness of the initial centroids in k-means to select the number of clusters, nc , but there was no such option we could use for hierarchical clustering with Ward linkage.

7 Conclusions

In this report, we used output from simulations of a jet interacting with high explosives to address two challenges in building spatio-temporal surrogates for high-dimensional data. First, the data were available on spatial domains of different sizes, at grid points that varied in their spatial coordinates, and in a format that distributed the output across multiple files at each time step of the simulation. We described how we reorganized these large data sets into a consistent format efficiently, exploiting parallelism where possible. Second, to improve the accuracy of the surrogates, we wanted to cluster the data by similarity and build a separate surrogate for each cluster. However, as the outputs are high-dimensional, with the spatial domain represented by more than two million grid points, traditional iterative clustering methods, such as k-means, could not be applied directly. We showed how we could use random projections to make the clustering of these outputs tractable. Our experiences indicated that the approximation introduced by the random projections had little effect on the clustering results. Moreover, we could use the randomness of both the random projections and the initial choice of cluster centroids in k-means to identify the number of clusters. The effectiveness of our approach is discussed further in the companion report [16], where we show how clustering the data by similarity improves the accuracy of the spatio-temporal surrogates created from a small set of simulations.

8 Acknowledgment

We would like to thank the Defense Threat Reduction Agency (DTRA) for funding this work. The simulations of the interaction of the jet with high explosives were performed using the ARES code developed at Lawrence Livermore National Laboratory.

LLNL-TR-850159 This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

References

- [1] D. Achlioptas. Database-friendly random projections. In *PODS '01: Proceedings of the twentieth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 274–281. Association for Computing Machinery, 2001. doi: <https://doi.org/10.1145/375551.375608>.
- [2] D. Achlioptas. Database-friendly random projections: Johnson-Lindenstrauss with binary coins. *Journal of Computer and System Sciences*, 66(4):671–687, 2003. doi: [https://doi.org/10.1016/S0022-0000\(03\)00025-4](https://doi.org/10.1016/S0022-0000(03)00025-4). Special Issue on PODS 2001.
- [3] D. Amsallem, M. J. Zahr, and C. Farhat. Nonlinear model order reduction based on local reduced-order bases. *International Journal for Numerical Methods in Engineering*, 92(10):891–916, 2012. doi: <https://doi.org/10.1002/nme.4371>.
- [4] L. Anderlucci, F. Fortunato, and A. Montanari. High-Dimensional Clustering via Random Projections. *Journal of Classification*, 39(1):191–216, March 2022. doi: <https://doi.org/10.1007/s00357-021-09403->

[5] E. Bingham and H. Mannila. Random projection in dimensionality reduction: Applications to image and text data. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '01, page 245–250, New York, NY, USA, 2001. Association for Computing Machinery. doi: <https://doi.org/10.1145/502512.502546>.

[6] T. Daniel, F. Casenave, N. Akkari, and D. Ryckelynck. Model order reduction assisted by deep neural networks (ROM-net). *Advanced Modeling and Simulation in Engineering Sciences*, 7, 2020. doi: <https://doi.org/10.1186/s40323-020-00153-6>.

[7] S. Dasgupta and A. Gupta. An elementary proof of a theorem of Johnson and Lindenstrauss. *Random Structures & Algorithms*, 22(1):60–65, 2002. doi: <https://doi.org/10.1002/rsa.10073>.

[8] Q. Du and M. Gunzburger. Model reduction by proper orthogonal decomposition coupled with centroidal Voronoi tessellations (Keynote). In *Proceedings, Fluids Engineering Division Summer Meeting*, volume Volume 1: Fora, Parts A and B, pages 1401–1406, 07 2002. doi: <https://doi.org/10.1115/FEDSM2002-31051>.

[9] X. Fern and C. Brodley. Cluster ensembles for high dimensional clustering: an empirical study. *J Mach Learn Res*, 22, 01 2004.

[10] X. Z. Fern and C. E. Brodley. Random projection for high dimensional data clustering: A cluster ensemble approach. In *International Conference on Machine Learning*, pages 186–193, 2003.

[11] G. Gan, C. Ma, and J. Wu. *Data Clustering: Theory, Algorithms, and Applications*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2007. doi: <https://doi.org/10.1137/1.9780898718348>.

[12] A. K. Jain. Data clustering: 50 years beyond k-means. *Pattern Recognition Letters*, 31(8):651–666, 2010. doi: <https://doi.org/10.1016/j.patrec.2009.09.011>. Award winning papers from the 19th International Conference on Pattern Recognition (ICPR).

[13] A. K. Jain and R. C. Dubes. *Algorithms for clustering data*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1988.

[14] I. T. Jolliffe and J. Cadima. Principal component analysis: A review and recent developments. *Phil. Trans. R. Soc. A.*, 374, 2016. doi: <http://doi.org/10.1098/rsta.2015.0202>.

[15] C. Kamath. Intelligent sampling for surrogate modeling, hyperparameter optimization, and data analysis. *Machine Learning with Applications*, 9:100373, 2022. doi: <https://doi.org/10.1016/j.mlwa.2022.100373>.

[16] C. Kamath, J. S. Franzman, and B. H. Daub. Spatio-temporal surrogates for interaction of a jet with high explosives: Part I - Analysis with a small sample size. Technical Report LLNL-TR-850152, Lawrence Livermore National Laboratory CA., June 2023.

[17] N. Kambhatla and T. K. Leen. Dimension reduction by local principal component analysis. *Neural Computation*, 9:1493–1516, 1997.

[18] S. Kaski. Dimensionality reduction by random mapping: Fast similarity computation for clustering. In *Proceedings IEEE International Conference on Neural Networks*, volume 1, pages 413 – 418 vol.1, 06 1998. ISBN 0-7803-4859-1. doi: <https://doi.org/10.1109/IJCNN.1998.682302>.

[19] P. Li, T. J. Hastie, and K. W. Church. Very sparse random projections. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '06, pages 287–296, New York, NY, USA, 2006. Association for Computing Machinery. doi: <https://doi.org/10.1145/1150402.1150436>.

[20] C. Meyer, S. Race, and K. Valakuzhy. Determining the number of clusters via iterative consensus clustering. In *Proceedings of the 2013 SIAM International Conference on Data Mining (SDM)*, pages 94–102, 2013. doi: <https://doi.org/10.1137/1.9781611972832.11>.

- [21] D. P. Mitchell. Spectrally optimal sampling for distribution ray tracing. *Computer Graphics*, 25(4):157–164, 1991.
- [22] P. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 11 1987. doi: [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7).
- [23] A. Strehl and J. Ghosh. Cluster ensembles - a knowledge reuse framework for combining multiple partitions. *Journal of Machine Learning Research*, 3:583–617, 01 2002. doi: <https://doi.org/10.1162/153244303321897735>.
- [24] The HDF Group. Hierarchical Data Format, version 5, 1997-2023. <https://www.hdfgroup.org/HDF5/>.
- [25] M. Turk and A. Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1):71–86, 01 1991. doi: <https://doi.org/10.1162/jocn.1991.3.1.71>.
- [26] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: <https://doi.org/10.1038/s41592-019-0686-2>.
- [27] Joe H. Ward. Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301):236–244, 1963. doi: <https://doi.org/10.1080/01621459.1963.10500845>.

A Appendix: Data and results for x-momentum

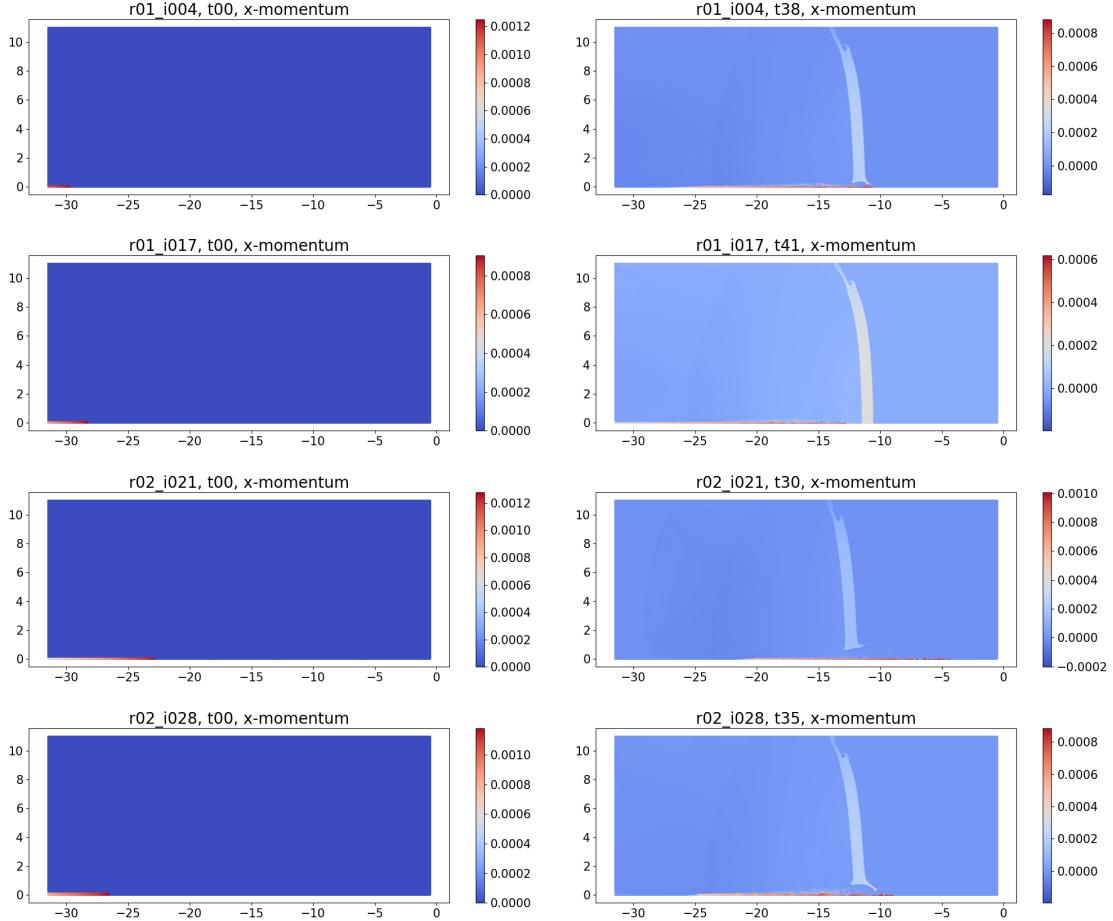


Figure 13: The variable *x-momentum*, after the snapshots have been aligned, cropped, and remapped, at the first time step (left column) and last time step (right column) for the four example simulations in Table 1. From top to bottom, simulations with keys r01_i004, r01_i017, r02_i021, r02_i028, respectively. The first time step has value zero for most of the domain, except the jet along the bottom on the left side of the region. The color bars are different between simulations and across time steps.

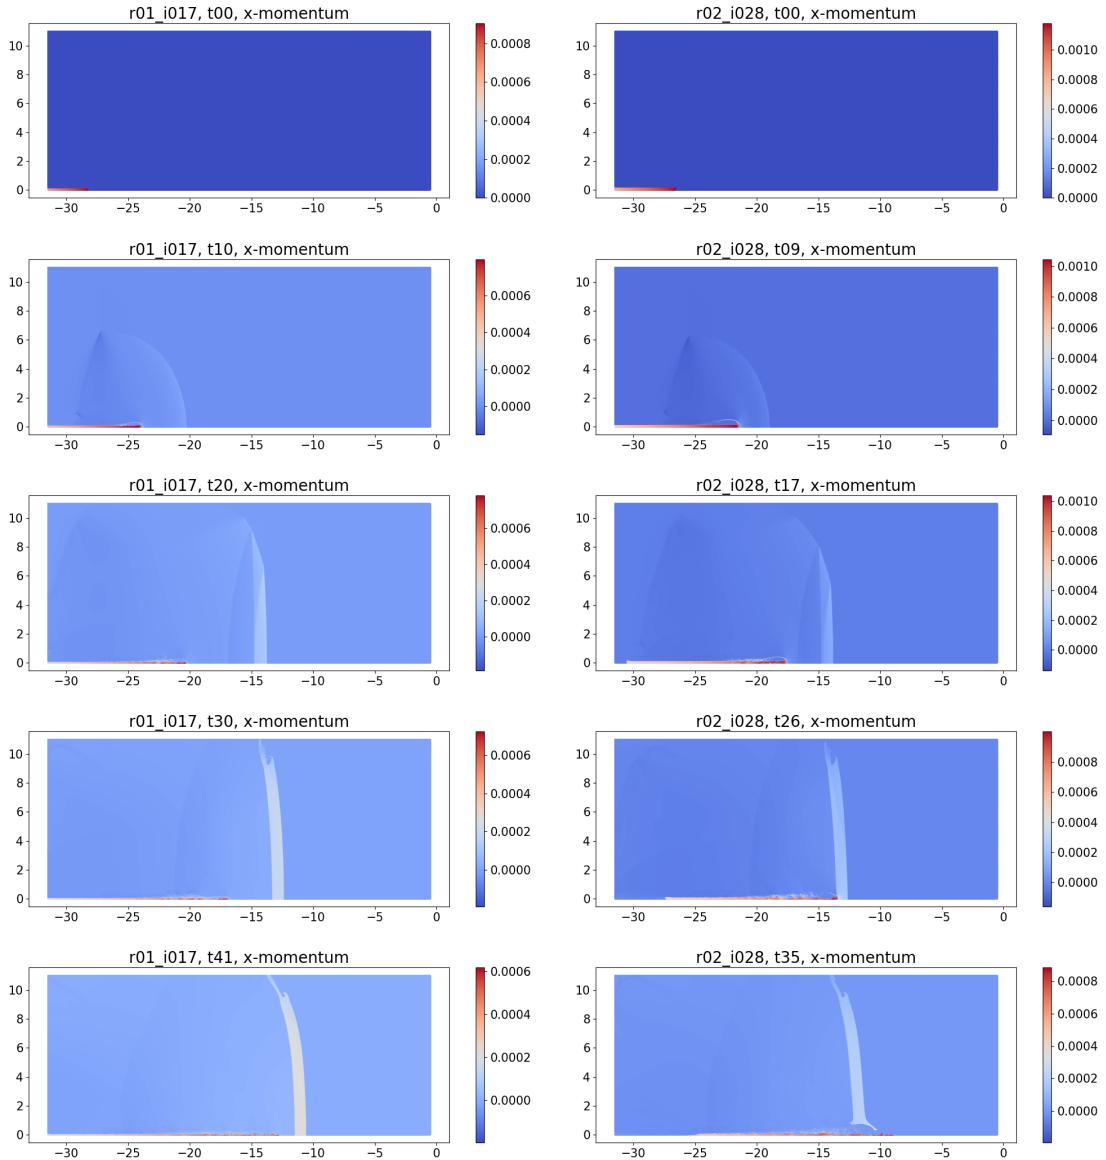


Figure 14: The variable *x-momentum*, after the snapshots have been aligned, cropped, and remapped, at different time steps in two simulations showing the evolution of the data over time. Left: key r01_i017 (no break case) at time steps t0, t10, t20, t30, and t41. Right: key r02_i028 (break case) at time steps t0, t09, t17, t26, and t35. The color bars are different between simulations and across time steps.

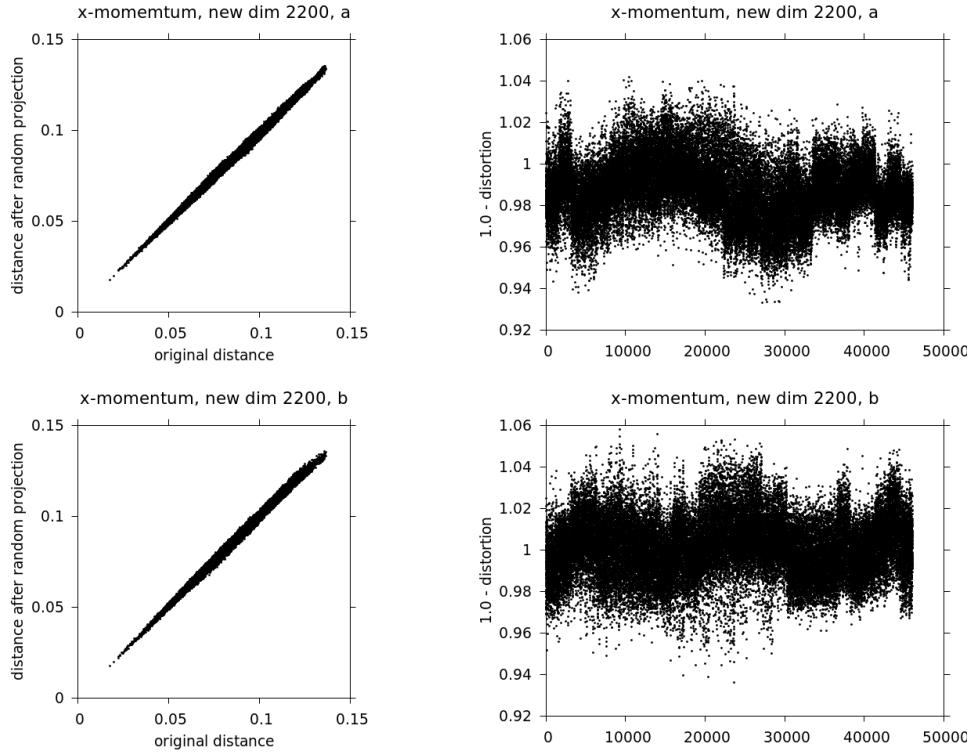


Figure 15: For the x-momentum variable, evaluating the choice of reduced dimension and randomization on the distances between the 29 snapshots from the baseline simulation to all 1604 snapshots. Left - actual vs. random-projected distances for two repetitions of random projections indicating that the randomness of the random projections has little effect on the new distances. Right - the corresponding values of $(1.0 - \epsilon)$ indicating that the distortion ϵ is less than 10%, even though our choice of reduced dimension of 2200 is smaller than the 6325 suggested by the Johnson-Lindenstrauss lemma for $\epsilon = 0.1$.

Values	nc3, a	nc3, b	nc4, a	nc4, b	nc5, a	nc5, b	nc6, a	nc6, b
0.0	58.39	63.12	71.34	71.52	70.62	71.45	70.43	68.87
0.1	0.12	0.008	0.003	0.0	0.08	0.34	3.00	3.54
0.2	0.07	0.0	2.79	0.0	1.39	0.34	2.27	3.61
0.3	2.37	0.075	0.12	0.0	2.85	2.21	2.78	2.87
0.4	3.25	2.76	0.19	0.0	0.0	1.32	2.78	2.66
0.5	4.46	0.47	0.0	6.27	8.76	7.98	2.64	2.42
0.6	3.13	3.10	0.20	0.0	0.006	1.31	2.64	2.14
0.7	2.56	0.089	0.12	0.0	2.81	2.13	2.71	3.12
0.8	0.16	0.0005	2.52	0.0	1.56	0.51	2.63	3.37
0.9	0.23	0.024	0.008	0.0	0.16	0.33	2.52	3.36
1.0	25.26	30.35	22.71	22.22	11.76	12.07	5.61	4.04

Table 3: Distribution of values in the consensus matrix generated for the x-momentum variable, with 10 repetitions of the clustering. We used four values for nc , the number of clusters, along with two repetitions of the random projections, indicated by a and b. The values in the table indicate how often the value in the first column occurs in the matrix. Thus, for $(nc = 4, b)$ in the fifth column, nearly 72% of the matrix entries are 0.0, nearly 22% are 1.0, with a small percentage of entries at 0.5.

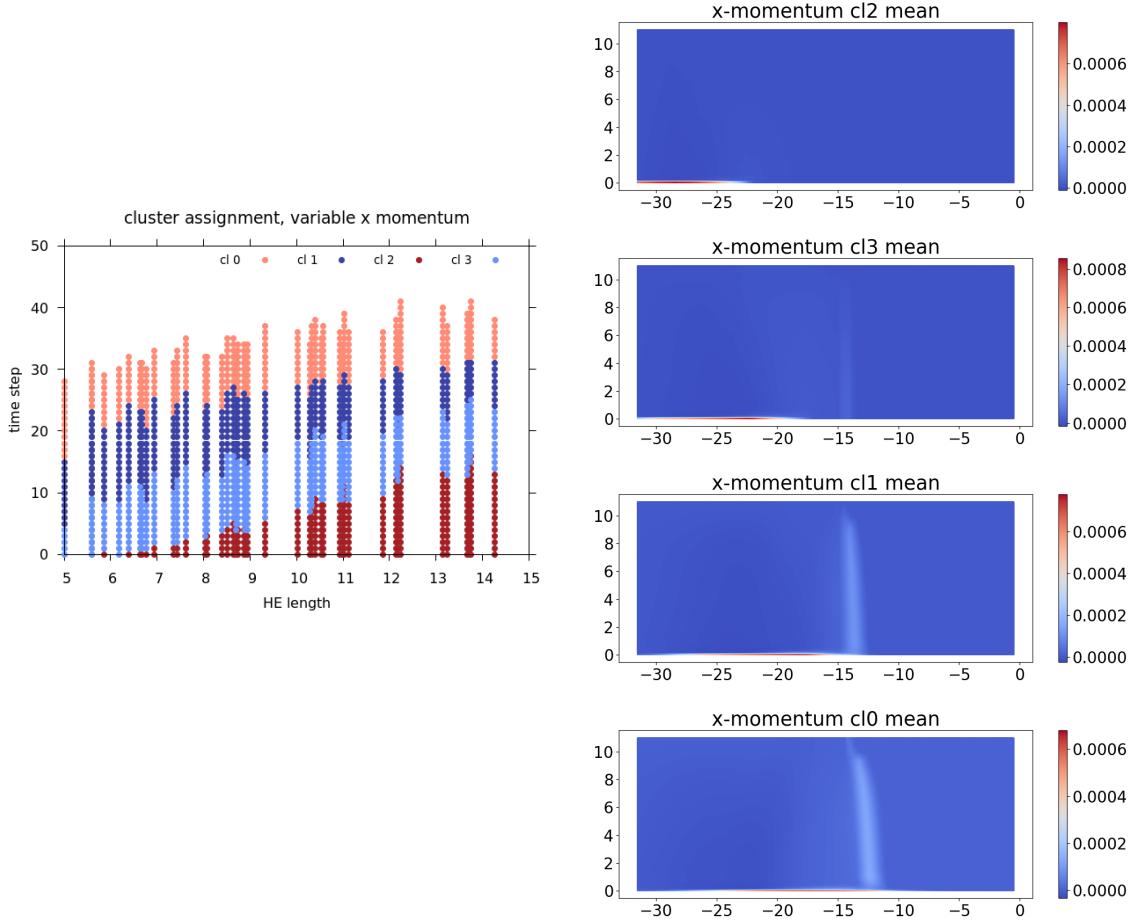


Figure 16: Left: the cluster assignment, corresponding to $(nc = 4, a)$, selected for the x-momentum variable shown for the times steps in the simulations plot along with the HE length. The cluster sizes for clusters 0, 1, 2, and 3, are 397, 454, 305, and 448, respectively. Right: the mean snapshot for the four clusters for the x-momentum variable, from top to bottom, clusters 2, 3, 1, and 0, showing the evolution with time. Note that the clusters form four bands that go across the HE length.

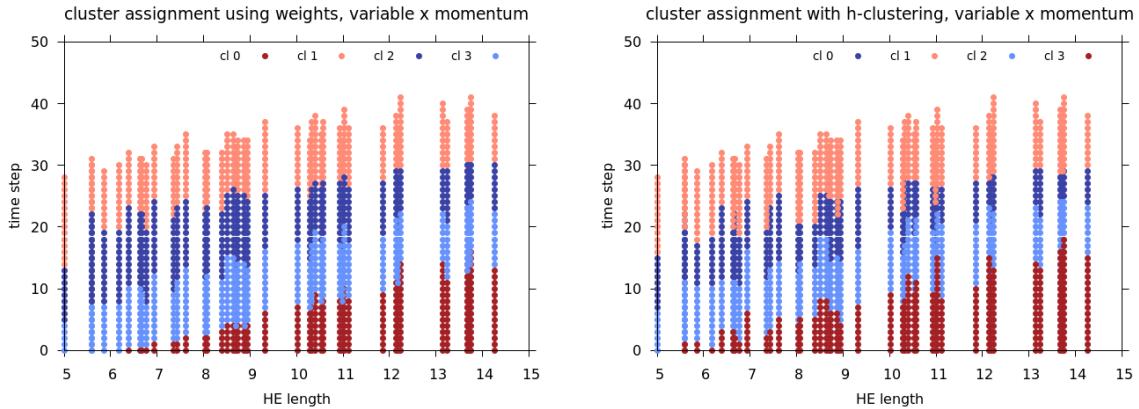


Figure 17: Cluster assignment for the x-momentum variable shown for the times steps in the simulations plot along with the HE length. Left: clustering SVD weights, with $nc = 4$; the cluster sizes for clusters 0, 1, 2, and 3, are 298, 441, 453, and 412, respectively. Right: using hierarchical clustering with the Ward linkage criterion, with $nc = 4$; the cluster sizes for clusters 0, 1, 2, and 3 are 342, 475, 404, and 383, respectively. The cluster number associated with a cluster of a specific color is different in the two plots. Note that the bands formed by the four clusters are less clear for hierarchical clustering.

B Appendix: Data and results for y-momentum

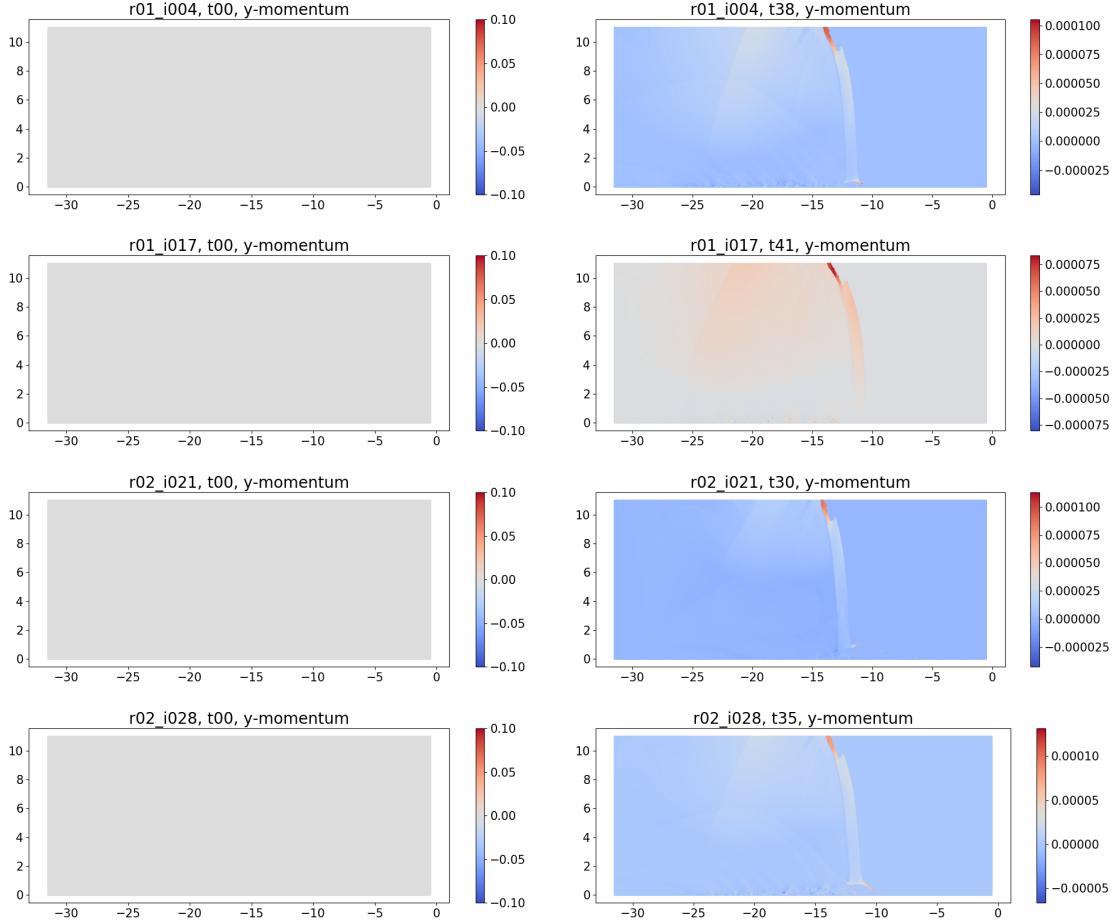


Figure 18: The variable *y-momentum*, after the snapshots have been aligned, cropped, and remapped, at the first (left column) and last (right column) time steps for the four example simulations in Table 1. From top to bottom, simulations with keys r01_i004, r01_i017, r02_i021, r02_i028, respectively. The first time step has value zero for most of the domain. Unlike the x-momentum in Figure 13, there is no y-momentum for the jet at initial time. The color bars are different between simulations and across time steps.

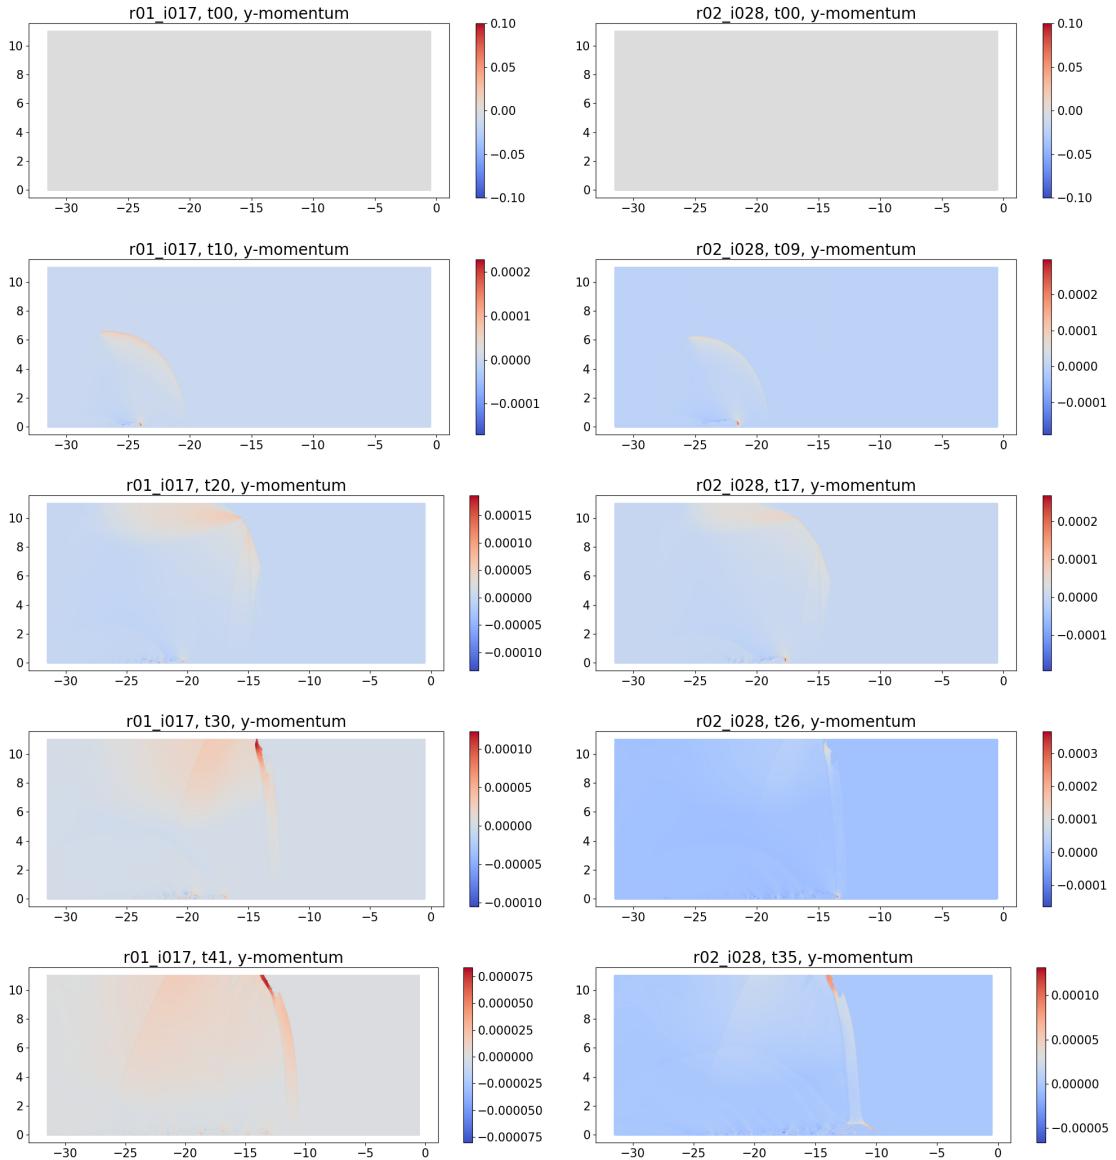


Figure 19: The variable *y-momentum*, after the snapshots have been aligned, cropped, and remapped, at different time steps in two simulations showing the evolution of the data over time. Left: key r01_i017 (no break case) at time steps t0, t10, t20, t30, and t41. Right: key r02_i028 (break case) at time steps t0, t09, t17, t26, and t35. The color bars are different between simulations and across time steps.

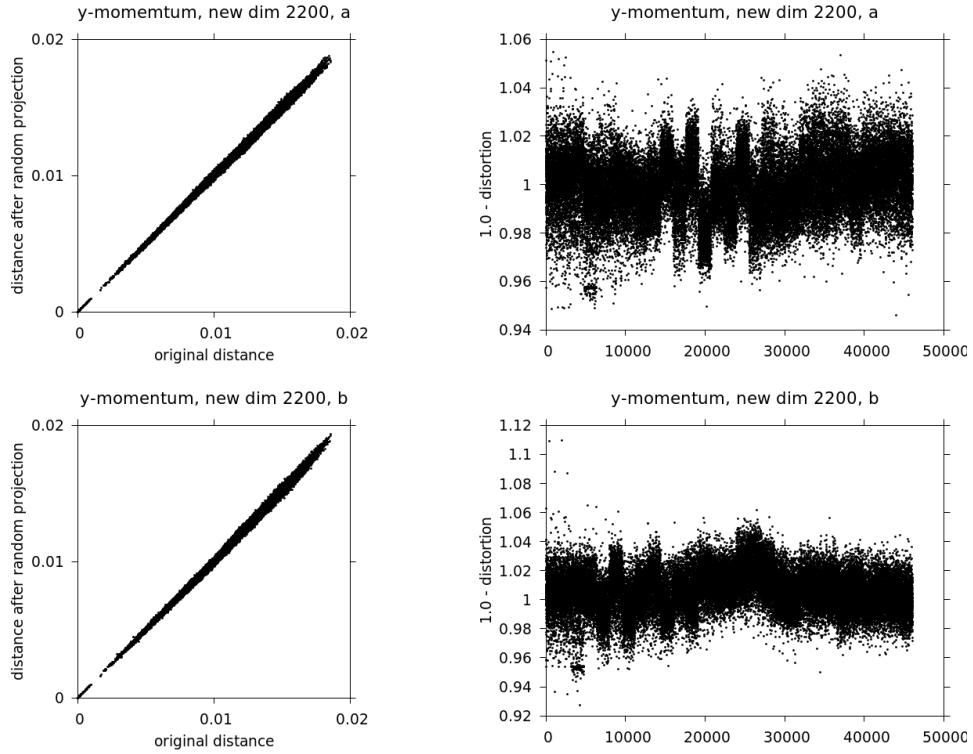


Figure 20: For the *y-momentum* variable, evaluating the choice of reduced dimension and randomization on the distances between the 29 snapshots from the baseline simulation to all 1604 snapshots. Left - actual vs. random-projected distances for two repetitions of random projections indicating that the randomness of the random projections has little effect on the new distances. Right - the corresponding values of $(1.0 - \epsilon)$ indicating that the distortion ϵ is less than 10%, even though our choice of reduced dimension of 2200 is smaller than the 6325 suggested by the Johnson-Lindenstrauss lemma for $\epsilon = 0.1$.

Values	nc3, a	nc3, b	nc4, a	nc4, b	nc5, a	nc5, b	nc6, a	nc6, b
0.0	50.30	49.53	69.88	70.65	66.64	72.12	71.85	69.97
0.1	0.29	0.62	0.001	0.11	1.44	1.32	0.33	1.27
0.2	0.50	0.008	0.57	0.0	1.78	0.99	1.95	2.19
0.3	0.0	0.10	0.66	0.0	1.74	0.82	1.03	1.03
0.4	14.32	0.92	0.0	0.0	2.77	0.009	0.52	1.54
0.5	1.47	22.31	0.0	0.0	0.68	0.06	4.65	3.62
0.6	7.37	0.75	0.0	0.0	1.88	0.03	2.09	2.18
0.7	0.0	0.16	0.25	0.0	3.88	0.68	1.59	2.12
0.8	0.21	0.06	0.47	0.0	2.16	1.38	1.72	1.31
0.9	0.58	0.30	0.003	0.18	1.52	0.94	0.48	1.82
1.0	24.95	25.23	28.17	29.06	15.52	21.65	13.79	12.93

Table 4: Distribution of values in the consensus matrix generated for the *y-momentum* variable, with 10 repetitions of the clustering. We used four values for nc , the number of clusters, along with two repetitions of the random projections, indicated by a and b. The values in the table indicate how often the value in the first column occurs in the matrix. Thus, for $(nc = 4, b)$ in the fifth column, nearly 71% of the matrix entries are 0.0, nearly 29% are 1.0, with a small percentage of entries at 0.1 and 0.9

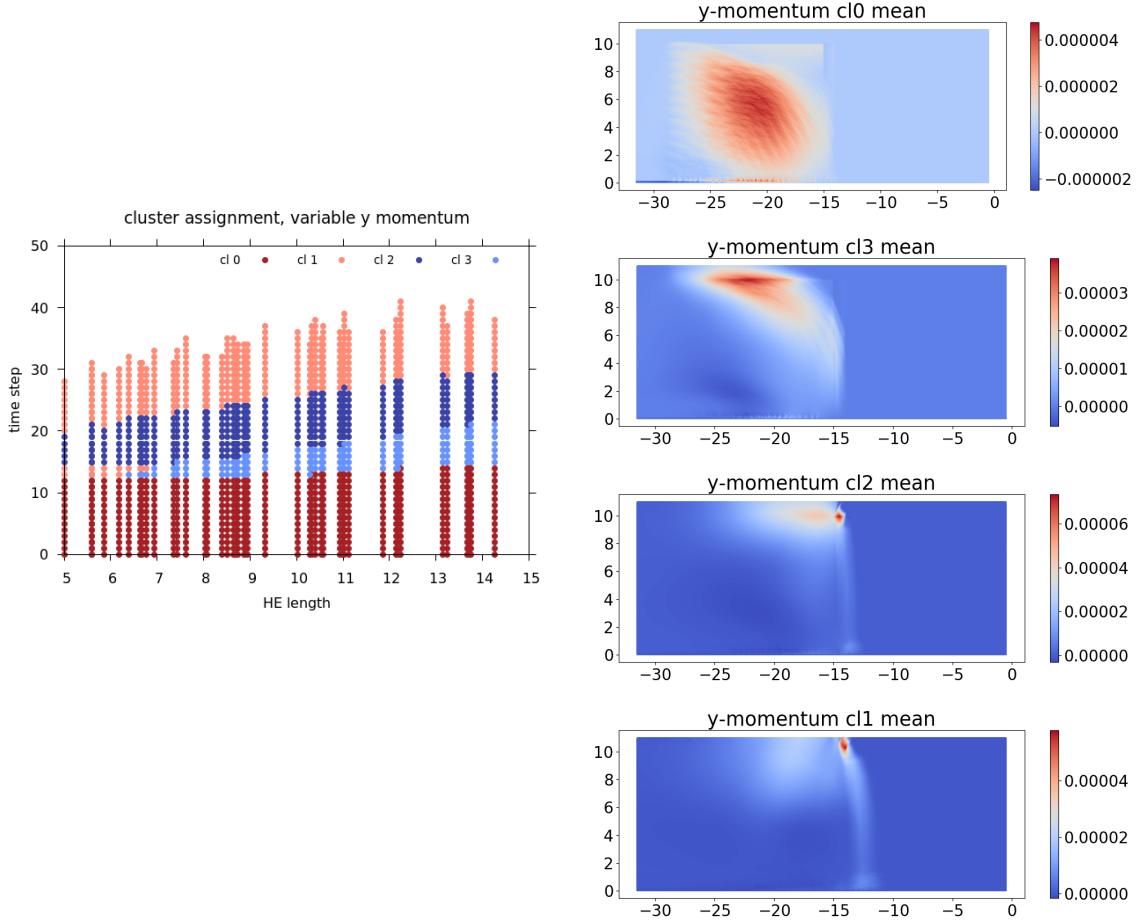


Figure 21: Left: the cluster assignment, corresponding to $(nc = 4, b)$, selected for the *y-momentum* variable shown for the time steps in the simulations plot along with the HE length. The cluster sizes for clusters 0, 1, 2, and 3, are 615, 460, 371, and 158, respectively. Right: the mean snapshot for the four clusters for the *y-momentum* variable, from top to bottom, clusters 0, 3, 2, and 1, showing the evolution with time. Note that the clusters form four bands across the HE length, with the much thinner band having some time steps at low HE length that belong to a different cluster; this is fixed in Figure 23.

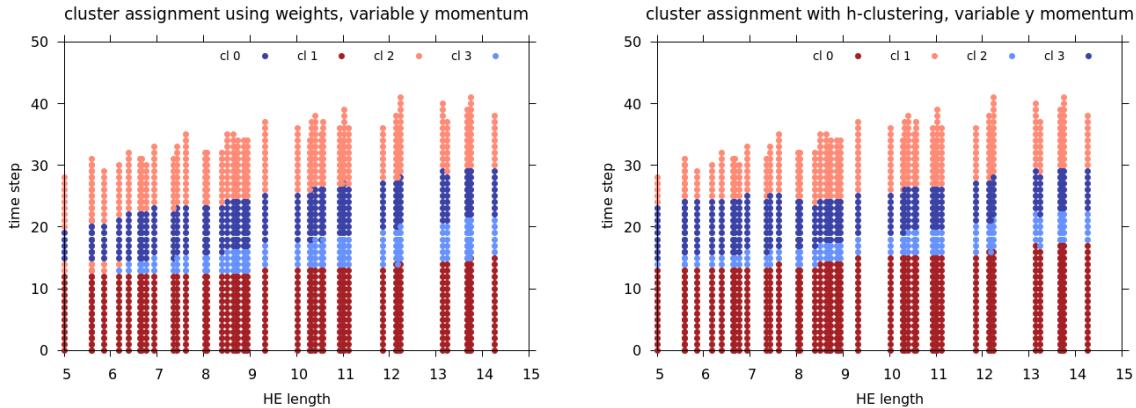


Figure 22: Cluster assignment for the *x-momentum* variable shown for the times steps in the simulations plot along with the HE length. Left: clustering SVD weights, with $nc = 4$; the cluster sizes for clusters 0, 1, 2, and 3, are 355, 619, 455, 175, respectively. Right: using hierarchical clustering with the Ward linkage criterion, with $nc = 4$; the cluster sizes for clusters 0, 1, 2, and 3 are 697, 412, 152, and 343, respectively. The cluster number associated with a cluster of a specific color is different in the two plots. Note that, unlike k-means, hierarchical clustering does not include some mid-time snapshots in the late time cluster at low values of HE length.

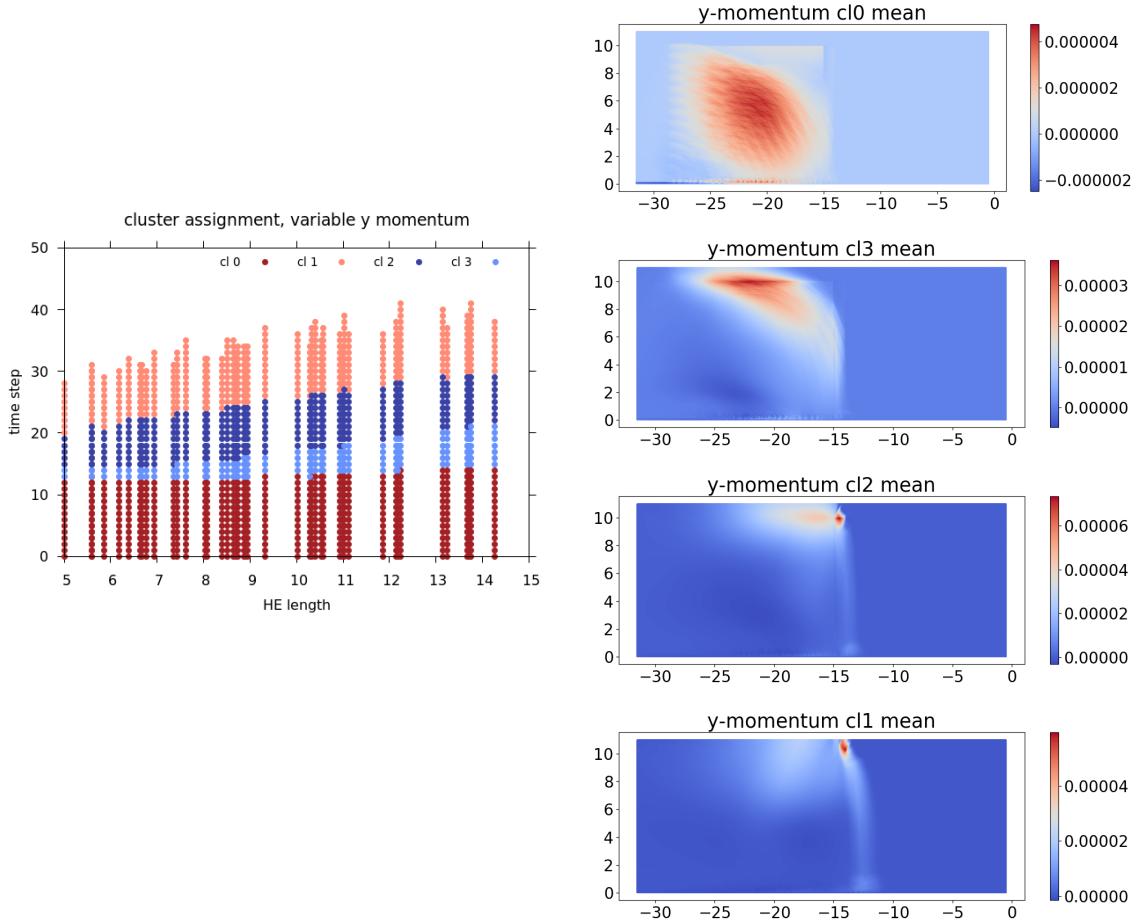


Figure 23: Left: the **fixed** cluster assignment, corresponding to $(nc = 4, b)$, selected for the *y-momentum* variable shown for the time steps in the simulations plot along with the HE length. The cluster sizes for clusters 0, 1, 2, and 3, are 615, 447, 371, and 171, respectively. 13 snapshots were moved from cluster 1 to cluster 3, resulting in a slight difference in values to the left of the plate near the top compared to the results in Figure 21. Right: the mean snapshot for the four clusters for the *y-momentum* variable, from top to bottom, clusters cl0, cl3, cl2, cl1, showing the evolution with time.