Multi-Resolution Dynamic Mode Decomposition

J. Nathan Kutz[‡][*]   Xing Fu[*]   Steven L. Brunton[†]

June 2, 2015

Abstract

We demonstrate that the integration of the recently developed dynamic mode decomposition (DMD) with a multi-resolution analysis allows for a decomposition method capable of robustly separating complex systems into a hierarchy of multi-resolution time-scale components. A one-level separation allows for background (low-rank) and foreground (sparse) separation of dynamical data, or robust principal component analysis. The multi-resolution dynamic mode decomposition is capable of characterizing nonlinear dynamical systems in an equation-free manner by recursively decomposing the state of the system into low-rank terms whose temporal coefficients in time are known. DMD modes with temporal frequencies near the origin (zero-modes) are interpreted as background (low-rank) portions of the given dynamics, and the terms with temporal frequencies bounded away from the origin are their sparse counterparts. The multi-resolution dynamic mode decomposition (mrDMD) method is demonstrated on several examples involving multi-scale dynamical data, showing excellent decomposition results, including sifting the El Niño mode from ocean temperature data. It is further applied to decompose a video data set into separate objects moving at different rates against a slowly varying background. These examples show that the decomposition is an effective dynamical systems tool for data-driven discovery.

1 Introduction

Modeling of multi-scale systems, both in time and space, pervade modern theoretical and computational efforts across the engineering, biological and physical sciences. Driving innovations are methods and algorithms that circumvent the significant challenges in efficiently connecting microscale to macro-scale effects that are separated potentially by orders of magnitude spatially and/or temporally. Wavelet-based methods and/or windowed Fourier Transforms are ideally structured to perform such multi-resolution analyses (MRA) as they systematically remove temporal or spatial features by a process of recursive refinement of sampling from the data of interest [1, 2, 3]. Typically, MRA is performed in either space or time, but not both simultaneously. We propose integrating the concept of MRA with the recently developed Dynamic Mode Decomposition (DMD) [4, 5, 6, 7, 8, 9], a technique that produces low-dimensional spatio-temporal modes. The proposed multi-resolution DMD (mrDMD) is shown to naturally integrate space and time so that the multi-scale spatio-temporal features are readily separated and approximate dynamical models constructed.

[*]Department of Applied Mathematics, University of Washington, Seattle, WA. 98195-2420. [‡] (kutz@uw.edu). Questions, comments, or corrections to this document may be directed to that email address. A video summary/abstract of this work may be found at: http://youtu.be/E1dNE02LaCE

[†]Department of Mechanical Engineering, University of Washington, Seattle, WA. 98195-2420.

The origins of the DMD method, which arose from pioneering work connecting the Koopman operator to dynamical systems theory [10, 11], are associated with the fluid dynamics community and the modeling of complex flows [5, 6]. Its growing success stems from the fact that it is an *equation-free*, data-driven method [3] capable of providing accurate assessments of the spatio-temporal coherent structures in a complex system, or short-time future estimates, thus potentially allowing for control protocols to be enacted simply from data sampling. The mathematical architecture advocated here is an alternative to the equation-free, multi-scale modeling method proposed by Kevrekidis and co-workers [12, 13]. More broadly, DMD has quickly gained popularity since it provides information about nonlinear dynamical systems. DMD analysis can be considered to be a numerical approximation to Koopman spectral analysis [6, 14], and it is in this sense that DMD is applicable to nonlinear systems. In fact, the terms *DMD mode* and *Koopman mode* are often used interchangeably in the fluids literature.

At its core, the DMD method can be thought of as an ideal combination of the Proper Orthogonal Decomposition (POD), a spatial dimensionality-reduction technique, with Fourier Transforms in time. More precisely, the DMD method produces a least-square regression to a linear dynamical system over the range of data collection. In this work, our goal is to integrate the DMD decomposition with key concepts from wavelet theory and MRA. Specifically, the DMD method takes snapshots of an underlying dynamical system to construct its decomposition. However, the frequency and duration (sampling window) of the data collection process can be adapted, much as in wavelet theory, to sift out information at different scales. Indeed, an iterative refinement of progressively shorter snapshot sampling windows and recursive extraction of DMD modes from slow to increasingly fast time scales allows for the mrDMD. Moreover, it also allows for improved analytic predictions of the short-time future state of the system which is of critical importance for feedback control, for instance. Critical innovations demonstrated here are the ability of the mrDMD to handle transient phenomenon and moving (translating/rotating) structures in data, both weaknesses of SVD-based techniques.

The paper is outlined as follows: In Sec. 2 the basic DMD theory is outlined with an emphasis on its low-rank approximation of data. This is followed in Sec. 3 by the development of the mrDMD structure and algorithm used in the subsequent applications, described in Sec. 4. The applications exhibit the application to dynamical systems as well as more broadly to other applications such as video analysis. The paper is concluded in Sec. 5 with an overview and outlook of the method.

2 Dynamic Mode Decomposition

The DMD method provides a spatio-temporal decomposition of data into a set of dynamic modes that are derived from snapshots or measurements of a given system in time. The mathematics underlying the extraction of dynamic information from time-resolved snapshots is closely related to the idea of the Arnoldi algorithm [5], one of the workhorses of fast computational solvers. The data collection process involves two (integer) parameters:

$$N = \text{number of spatial measurements per time snapshot}, \tag{1a}$$
$$M = \text{number of snapshots taken in time}. \tag{1b}$$

Originally the algorithm was designed to collect data at regularly spaced intervals of time. However, new innovations, and a more general definition of the DMD, allow for both sparse spatial [15] and temporal [16] collection of data as well as irregularly spaced collection times [9]. To illustrate the

algorithm, we consider regularly spaced sampling in time. The data collection times are given by:

$$(2) \qquad t_{m+1} = t_m + \Delta t$$

where the collection time starts at t_1 and ends at t_M, and the interval between data collection times is Δt. In the mrDMD method, the total number of snapshots will vary as the algorithm extracts multi-timescale spatio-temporal structures. This will be the central focus of the next section.

The data snapshots are arranged into an $N \times M$ matrix

$$(3) \qquad \mathbf{X} = [\mathbf{x}(t_1) \ \mathbf{x}(t_2) \ \mathbf{x}(t_3) \ \cdots \ \mathbf{x}(t_M)]$$

where the vector $\mathbf{x}$ are the N measurements of the state variable of the system of interest at the data collection points. The objective is to mine the data matrix $\mathbf{X}$ for important dynamical information. For the purposes of the DMD method, the following matrix is also defined:

$$(4) \qquad \mathbf{X}_j^k = [\mathbf{x}(t_j) \ \mathbf{x}(t_{j+1}) \ \cdots \ \mathbf{x}(t_k)].$$

Thus this matrix includes columns j through k of the original data matrix.

The DMD method approximates the modes of the so-called *Koopman operator*. The Koopman operator is a linear, infinite-dimensional operator that represents nonlinear, possibly infinite-dimensional, dynamics without linearization [6, 14], and it is the adjoint of the Perron-Frobenius operator. The method can be viewed as computing, from the experimental data, the eigenvalues and eigenvectors (low-dimensional modes) of a linear model that approximates the underlying dynamics, even if the dynamics are nonlinear. Since the model is assumed to be linear, the decomposition gives the growth rates and frequencies associated with each mode. If the underlying model is linear, then the DMD method recovers the leading eigenvalues and eigenvectors normally computed using standard solution methods for linear differential equations.

The DMD involves approximating the eigendecomposition of the best-fit linear operator $\mathbf{A}$ that relates a state $\mathbf{x}_j$ at time t_j to the state $\mathbf{x}_{j+1}$ at the next timestep:

$$(5) \qquad \mathbf{x}_{j+1} \approx \mathbf{A}\mathbf{x}_j.$$

If Eq. (5) holds exactly and the data is generated by a linear system, then each column of $\mathbf{X}_1^{M-1}$ is an element of a Krylov subspace:

$$(6) \qquad \mathbf{X}_1^{M-1} = \left[\mathbf{x}_1 \ \mathbf{A}\mathbf{x}_1 \ \mathbf{A}^2\mathbf{x}_1 \ \cdots \ \mathbf{A}^{M-2}\mathbf{x}_1\right].$$

Eq. (5) may be written in matrix form as

$$(7) \qquad \mathbf{X}_2^M \approx \mathbf{A}\mathbf{X}_1^{M-1},$$

where the operator $\mathbf{A}$ is chosen to minimize the Frobenius norm of $\|\mathbf{X}_2^M - \mathbf{A}\mathbf{X}_1^{M-1}\|_F$. In other words, the operator $\mathbf{A}$ advances each snapshot column in $\mathbf{X}_1^{M-1}$ a single timestep, Δt, resulting in the future snapshot columns in $\mathbf{X}_2^M$.

In practice, when the state dimension N is large, the matrix $\mathbf{A}$ may be intractable to analyze directly. Instead, DMD circumvents the eigendecomposition of $\mathbf{A}$ by considering a rank-reduced representation in terms of a POD-projected matrix $\tilde{\mathbf{A}}$. The DMD algorithm proceeds as follows [9]:

1. First, take the SVD of $\mathbf{X}_1^{M-1}$[17]:

   (8)
   $$\mathbf{X}_1^{M-1} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^*,$$

   where $*$ denotes the conjugate transpose, $\mathbf{U} \in \mathbb{C}^{N \times K}$, $\boldsymbol{\Sigma} \in \mathbb{C}^{K \times K}$ and $\mathbf{V} \in \mathbb{C}^{M-1 \times K}$. Here K is the rank of the reduced SVD approximation to $\mathbf{X}_1^{M-1}$. The left singular vectors $\mathbf{U}$ are POD modes.

   The SVD reduction in (8) could also be exploited at this stage in the algorithm to perform a low-rank truncation of the data. Specifically, if low-dimensional structure is present in the data, the singular values of $\boldsymbol{\Sigma}$ will decrease sharply to zero with perhaps only a limited number of dominant modes. A principled way to truncate noisy data would be to use the recent hard-thresholding algorithm of Gavish and Donoho [18].

2. Next, compute $\tilde{\mathbf{A}}$, the $K \times K$ projection of the full matrix $\mathbf{A}$ onto POD modes:

   (9)
   $$\begin{aligned} \mathbf{A} &= \mathbf{X}_2^M \mathbf{V}\boldsymbol{\Sigma}^{-1}\mathbf{U}^* \\ \implies \tilde{\mathbf{A}} &= \mathbf{U}^*\mathbf{A}\mathbf{U} = \mathbf{U}^*\mathbf{X}_2^M\mathbf{V}\boldsymbol{\Sigma}^{-1}. \end{aligned}$$

3. Compute the eigendecomposition of $\tilde{\mathbf{A}}$:

   (10)
   $$\tilde{\mathbf{A}}\mathbf{W} = \mathbf{W}\boldsymbol{\Lambda},$$

   where columns of $\mathbf{W}$ are eigenvectors and $\boldsymbol{\Lambda}$ is a diagonal matrix containing the corresponding eigenvalues λ_k.

4. Finally, we may reconstruct eigendecomposition of $\mathbf{A}$ from $\mathbf{W}$ and $\boldsymbol{\Lambda}$. In particular, the eigenvalues of $\mathbf{A}$ are given by $\boldsymbol{\Lambda}$ and the eigenvectors of $\mathbf{A}$ (DMD modes) are given by columns of $\boldsymbol{\Psi}$:

   (11)
   $$\boldsymbol{\Psi} = \mathbf{X}_2^M\mathbf{V}\boldsymbol{\Sigma}^{-1}\mathbf{W}.$$

Note that Eq. (11) from [9] differs from the formula $\boldsymbol{\Psi} = \mathbf{U}\mathbf{W}$ from [5], although these will tend to converge if $\mathbf{X}_1^{M-1}$ and $\mathbf{X}_2^M$ have the same column spaces.

With the low-rank approximations of both the eigenvalues and eigenvectors in hand, the projected future solution can be constructed for all time in the future. By first rewriting for convenience $\omega_k = \ln(\lambda_k)/\Delta t$, then the approximate solution at all future times, $\mathbf{x}_{\mathrm{DMD}}(t)$, is given by

(12)
$$\mathbf{x}_{\mathrm{DMD}}(t) = \sum_{k=1}^{K} b_k(0)\psi_k(\boldsymbol{\xi})\exp(\omega_k t) = \boldsymbol{\Psi}\mathrm{diag}(\exp(\omega t))\mathbf{b}$$

where $\boldsymbol{\xi}$ are the spatial coordinates, $b_k(0)$ is the initial amplitude of each mode, $\boldsymbol{\Psi}$ is the matrix whose columns are the eigenvectors ψ_k, $\mathrm{diag}(\omega t)$ is a diagonal matrix whose entries are the eigenvalues $\exp(\omega_k t)$, and $\mathbf{b}$ is a vector of the coefficients b_k.

An alternative interpretation of (12) is that it represents the least-square fit, or regression, of a linear dynamical system $d\mathbf{x}_{\mathrm{DMD}}/dt = \mathbf{A}\mathbf{x}_{\mathrm{DMD}}$ to the data sampled. In particular, the matrix $\mathbf{A}$ constructed is such that $\|\mathbf{x}(t) - \mathbf{x}_{\mathrm{DMD}}(t)\|$ is minimized. In the context of the multi-resolution analysis that follows, each level of the multi-scale decomposition produces a linear dynamical system, or matrix $\mathbf{A}$, for the time-scale under consideration.

It only remains to compute the initial coefficient values $b_k(0)$. If we consider the initial snapshot ($\mathbf{x}_1$) at time $t_1 = 0$, let's say, then (12) gives $\mathbf{x}_1 = \mathbf{\Psi b}$. This generically is not a square matrix so that its solution

$$(13) \qquad \qquad \mathbf{b} = \mathbf{\Psi}^+ \mathbf{x}_1$$

can be found using a pseudo-inverse. Indeed, $\mathbf{\Psi}^+$ denotes the Moore-Penrose pseudo-inverse that can be accessed in MATLAB via the `pinv` command. The pseudo-inverse is equivalent to finding the best solution $\mathbf{b}$ the in the least-squares (best fit) sense. This is equivalent to how DMD modes were derived originally.

Overall then, the DMD algorithm presented here takes advantage of low dimensionality in the data in order to make a low-rank approximation of the linear mapping that best approximates the nonlinear dynamics of the data collected for the system. Once this is done, a prediction of the future state of the system is achieved for all time. Unlike the POD-Galerkin method, which requires solving a low-rank set of dynamical quantities to predict the future state, no additional work is required for the future state prediction outside of plugging in the desired future time into (12). Thus the advantages of DMD revolve around the fact that (i) it is an equation-free architecture, and (ii) a future state prediction is known for any time t (of course, provided the DMD approximation holds).

More broadly, the DMD method was shown to be a highly successful method for foreground/background subtraction in video feeds [19]. Indeed, the DMD method is a novel, dynamical systems base method for performing a *robust Principal Components Analysis* (RPCA) of data streams [19]. Importantly, the DMD-based RPCA performs the low-rank/sparse matrix separation 3-4 orders of magnitude faster than standard ℓ_1 optimization methods [20]. RPCA is extremely important for handling data outliers and/or corrupt data matrices. In the multi-resolution version of DMD, the RPCA can be used to effectively remove outliers at each level of decomposition.

3   Multi-Resolution Dynamic Mode Decomposition

The mrDMD is inspired by the observation that the slow- and fast-modes can be separated for such applications as foreground/background subtraction in video feeds [19]. The mrDMD recursively removes low-frequency, or slowly-varying, content from a given collection of snapshots. Typically, the number of snapshots M are chosen so that the DMD modes provide an approximately full rank approximation of the dynamics observed. Thus M is chosen so that all high- and low-frequency content is present. In the mrDMD, M is originally chosen in the same way so that an approximate full rank approximation can be accomplished. However, from this initial pass through the data, the slowest m_1 modes are removed, and the domain is divided into two segments with $M/2$ snapshots each. DMD is once again performed on each $M/2$ snapshot sequences. Again the slowest m_2 modes are removed and the algorithm is continued until a desired termination.

Mathematically, the mrDMD separates the DMD approximate solution (12) in the first pass as follows:

$$(14) \quad \mathbf{x}_{\mathrm{mrDMD}}(t) = \sum_{k=1}^{M} b_k(0)\psi_k^{(1)}(\boldsymbol{\xi})\exp(\omega_k t) \quad = \underbrace{\sum_{k=1}^{m_1} b_k(0)\psi_k^{(1)}(\boldsymbol{\xi})\exp(\omega_k t)}_{\text{(slow modes)}} + \underbrace{\sum_{k=m_1+1}^{M} b_k(0)\psi_k^{(1)}(\boldsymbol{\xi})\exp(\omega_k t)}_{\text{(fast modes)}}$$

where the $\psi_k^{(1)}(\mathbf{x})$ represent the DMD modes computed from the full M snapshots.

The first sum in this expression (14) represents the slow-mode dynamics whereas the second sum is everything else. Thus the second sum can be computed to yield the matrix:

$$(15) \qquad \mathbf{X}_{M/2} = \sum_{k=m_1+1}^{M} b_k(0) \psi_k^{(1)}(\boldsymbol{\xi}) \exp(\omega_k t) \,.$$

The DMD analysis outlined in the previous section can now be performed once again on the data matrix $\mathbf{X}_{M/2}$. However, the matrix $\mathbf{X}_{M/2}$ is now separated into two matrices

$$(16) \qquad \mathbf{X}_{M/2} = \mathbf{X}_{M/2}^{(1)} + \mathbf{X}_{M/2}^{(2)}$$

where the first matrix contains the first $M/2$ snapshots and the second matrix contains the remaining $M/2$ snapshots. The m_2 slow-DMD modes at this level are given by $\psi_k^{(2)}$, where they are computed separately in the first of second interval of snapshots.

The iteration process works by recursively removing slow frequency components and building the new matrices $\mathbf{X}_{M/2}, \mathbf{X}_{M/4}, \mathbf{X}_{M/8}, \cdots$ until a desired/prescribed multi-resolution decomposition has been achieved. The approximate DMD solution can then be constructed as follows:

$$
\begin{aligned}
\mathbf{x}_{\mathrm{mrDMD}}(t) \;=\; & \sum_{k=1}^{m_1} b_k^{(1)} \psi_k^{(1)}(\boldsymbol{\xi}) \exp(\omega_k^{(1)} t) + \sum_{k=1}^{m_2} b_k^{(2)} \psi_k^{(2)}(\boldsymbol{\xi}) \exp(\omega_k^{(2)} t) \\
& + \sum_{k=1}^{m_3} b_k^{(3)} \psi_k^{(3)}(\boldsymbol{\xi}) \exp(\omega_k^{(3)} t) + \cdots
\end{aligned}
$$

(17)

where at the evaluation time t, the correct modes from the sampling window are selected at each level of the decomposition. Specifically, the $\psi_k^{(k)}$ and $\omega_k^{(k)}$ are the DMD modes and DMD eigenvalues at the kth level of decomposition, the $b_k^{(k)}$ are the initial projections of the data onto the time interval of interest, and the m_k are the number of slow-modes retained at each level. The advantage of this method is readily apparent: different spatio-temporal DMD modes are used to represent key multi-resolution features. Thus there is not a single set of modes that dominates the SVD and potentially marginalizes features at other time scales.

Figure 1 illustrates the multi-resolution DMD process pictorially. In the figure, a three-level decomposition is performed with the slowest scale represented in blue (eigenvalues and snapshots), the mid-scale in red and the fast scale in green. The connection to multi-resolution wavelet analysis is also evident from the bottom panels as one can see that the mrDMD method successively pulls out time-frequency information in a principled way.

As a final remark, the sampling strategy and algorithm discussed here (See Fig. 1) can be easily modified since only the slow modes at each decomposition level need to be accurately computed. Thus one can modify the algorithm so as to sample a fixed number, for instance M, data snapshots in each sampling window. The value of M need not be large as only the slow modes need to be resolved. Thus the sampling rate would increase as the decomposition proceeds from one level to the next. This assures that the lowest levels of the mrDMD are not highly-sampled since the cost of the SVD would be greatly increased by such a fine sampling rate.

3.1 Formal mrDMD Expansion

The solution (17) can be made more precise. Specifically, one must account for the number of levels (L) of the decomposition, the number of time bins (J) for each level, and the number of modes
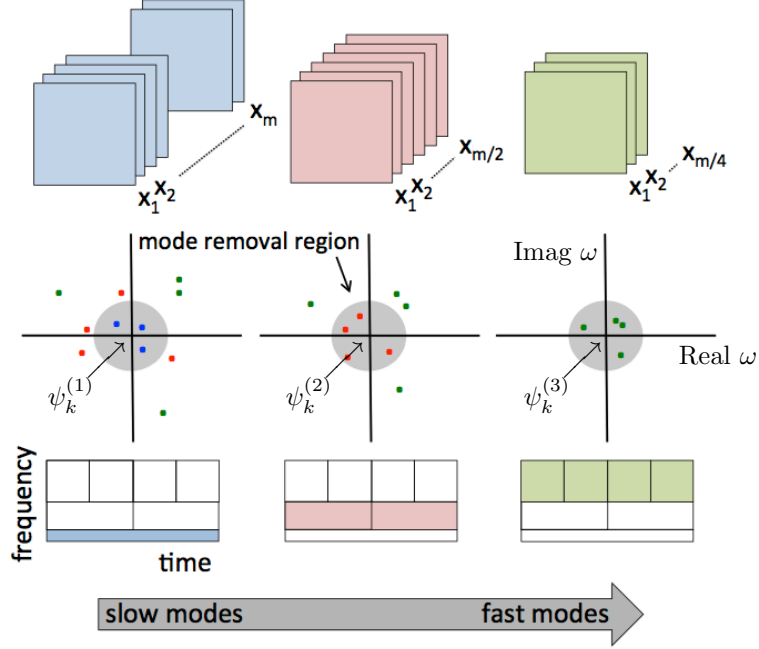
Figure 1: Representation of the multi-resolution dynamic mode decomposition where successive sampling of the data, initially with M snapshots and decreasing by a factor of two at each resolution level, is shown (top figures). The DMD spectrum is shown in the middle panel where there are m_1 (blue dots) slow-dynamic modes at the slowest level, m_2 (red) modes at the next level and m_3 (green) modes at the fastest time-scale shown. The shaded region represents the modes that are removed at that level. The bottom panels shows the wavelet-like time-frequency decomposition of the data color coded with the snapshots and DMD spectral representations.

retained at each level (m_L). This can be easily seen in Fig. 1. Thus the solution is parametrized by the following three indices:

$$\ell = 1, 2, \cdots, L \;\; \text{number of decomposition levels} \tag{18a}$$

$$j = 1, 2, \cdots, J \;\; \text{number time bins per level } (J = 2^{(\ell-1)}) \tag{18b}$$

$$k = 1, 2, \cdots, m_L \;\; \text{number of modes extracted at level } L. \tag{18c}$$

To formally define the series solution for $\mathbf{x}_{\text{mrDMD}}(t)$, we propose the following indicator function

$$\text{(19)} \qquad f_{\ell,j}(t) = \begin{cases} 1 & t \in [t_j, t_{j+1}] \\ 0 & \text{elsewhere} \end{cases} \quad \text{with} \quad j = 1, 2, \cdots, J \quad \text{and} \quad J = 2^{(\ell-1)}$$

which is only non-zero in the interval, or time bin, associated with the value of j. The parameter ℓ denotes the level of the decomposition.
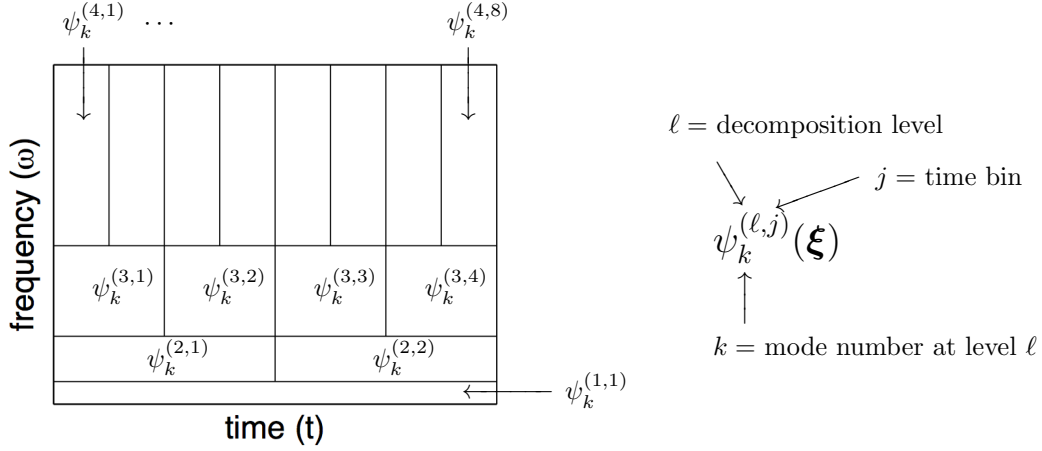
Figure 2: Illustration of the mrDMD mode decomposition and hierarchy. Represented are the modes $\psi_k^{\ell,j}(\boldsymbol{\xi})$ and their position in the decomposition structure. The triplet of integer values, ℓ, j and k, uniquely express the time level, bin and mode of the decomposition.

The three indices and indicator function (19) give the mrDMD solution expansion

$$(20) \qquad \mathbf{x}_{\mathrm{mrDMD}}(t) = \sum_{\ell=1}^{L} \sum_{j=1}^{J} \sum_{k=1}^{m_L} f_{\ell,j}(t) b_k^{(\ell,j)} \psi_k^{(\ell,j)}(\boldsymbol{\xi}) \exp(\omega_k^{(\ell,j)} t) .$$

This is a concise definition of the mrDMD solution that includes the information on the level, time bin location and number of modes extracted. Figure 2 demonstrates the mrDMD decomposition in terms of the solution (20). In particular, each mode is represented in its respective time bin and level. An alternative interpretation of this solution is that it yields the least-square fit, at each level ℓ of the decomposition, to the linear dynamical system

$$(21) \qquad \frac{d\mathbf{x}^{(\ell,j)}}{dt} = \mathbf{A}^{(\ell,j)} \mathbf{x}^{(\ell,j)}$$

where the matrix $\mathbf{A}^{(\ell,j)}$ captures the dynamics in a given time bin j at level ℓ.

The indicator function $f_{\ell,j}(t)$ acts as sifting function for each time bin. Interestingly, this function acts as the Gabór window of a windowed Fourier transform [3]. Since our sampling bin has a hard cut-off of the time series, it may introduce some artificial high-frequency oscillations. Time-series analysis, and wavelets in particular, introduce various functional forms that can be used in an advantageous way. Thus thinking more broadly, one can imagine using wavelet functions for the sifting operation, thus allowing the time function $f_{\ell,j}(t)$ to take the form of one of the many potential wavelet basis, i.e. Haar, Daubechies, Mexican Hat, etc. This will be considered in future work. For the present, we simply use the sifting function introduced in (19)
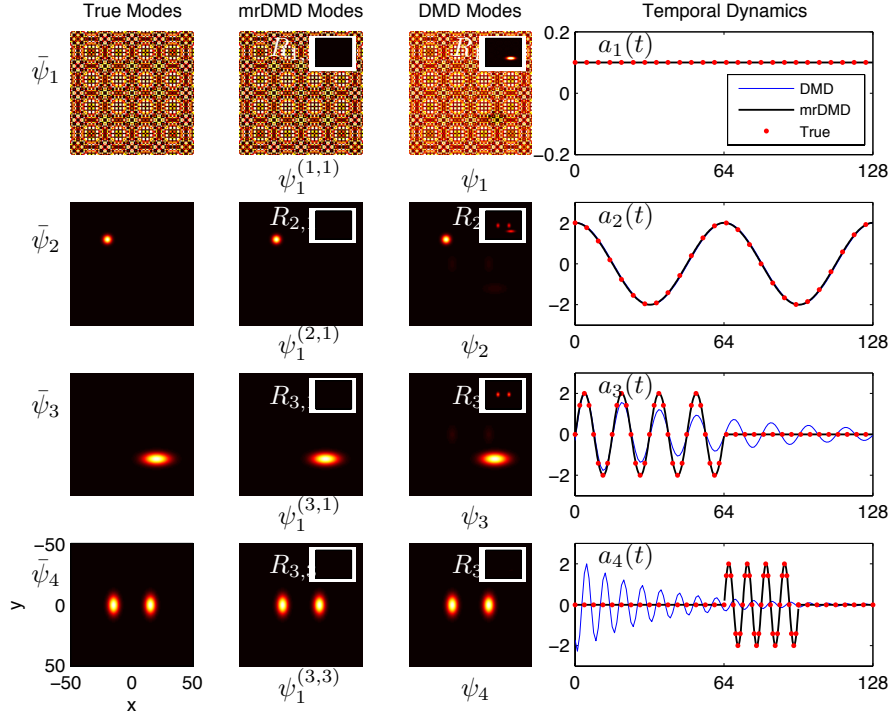
8

Figure 3: Comparison of the true modes ($\bar{\psi}_k$) to DMD (ψ_k) and mrDMD (ψ_k^j) modes. The true, mrDMD and DMD modes are showed in columns from left to right. The inset shows the relative error of the mrDMD ($R_{j,\ell} = \|\bar{\psi}_j - \psi_j^{(\ell,1)}\|$) and DMD ($R_j = \|\bar{\psi}_j - \psi_j\|$) in comparison to the true mode. Temporal dynamics are also compared, with true dynamics shown in red dots, DMD reconstructed dynamics shown in blue solid line and mrDMD reconstructed dynamics shown in black solid line.

4   Application of Method

The mrDMD developed in the last section is implemented here on three example data sets: the first being an exemplar of a video stream, the second coming from atmospheric-ocean data, and the third an exemplar of data with moving objects. This final example is especially important as it renders many SVD-based decompositions useless. For the first case, the implementation is compared against the standard DMD algorithm, highlighting the ability of the mrDMD to correctly capture multi-scale phenomena. It should be noted that in previous work with DMD, which was not framed in the mrDMD architecture advocated here, a level-1 decomposition was effectively used to separate the foreground from the background in a video [19]. Indeed, this work motivates the generalization to the multi-resolution analysis.
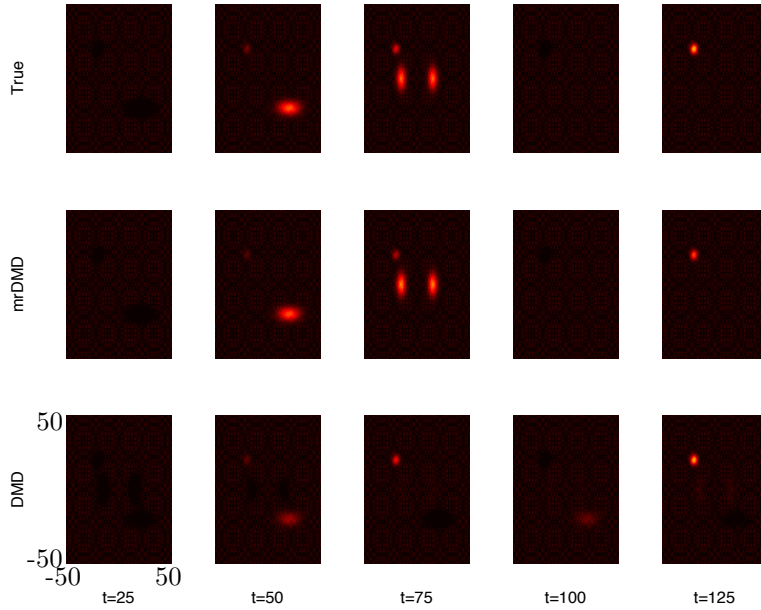
9

Figure 4: Comparison of snapshots from the original video of Fig. 3 and reconstruction from DMD and mrDMD. The true dynamics are showed in the first row, while the mrDMD reconstructed dynamics are showed in the second row and the DMD reconstructed dynamics are showed in the third row. The time labels on the bottom correspond to when each snapshot was taken. The mrDMD reconstruction matches with the original video very well while the DMD reconstruction shows inconsistencies. Indeed, the mrDMD method can easily handle signals that turn on and off in time while the DMD method cannot.

4.1 Spatio-temporal filtering of video

The first example we consider for application of the mrDMD is illustrated in Fig. 3. For this example, four spatio-temporal modes are combined into a single data set. Specifically, we combine the four modes shown in the left panels with the time dynamics given in the right panels. The four modes used to construct the true solution are represented by $\bar{\psi}_j(x, y)$ for $j = 1, 2, 3$ and 4. Their corresponding time dynamics are given by $a_j(t)$. Thus the true solution is expressed by

$$(22) \qquad \bar{\mathbf{x}} = \sum_{j=1}^{4} a_j(t)\bar{\psi}_j(x, y).$$

Both the DMD (represented by $\mathbf{x}_{\mathrm{DMD}}$ and the modes ψ_j of (12) with $j = 1, 2, 3$ and 4) and mrDMD (represented by $\mathbf{x}_{\mathrm{mrDMD}}$ and the $\psi_k^{(\ell,j)}$ of (17) where $k = 1$ and $\ell = 1, 2, 3$) attempt to reconstruct $\bar{\mathbf{x}}$.

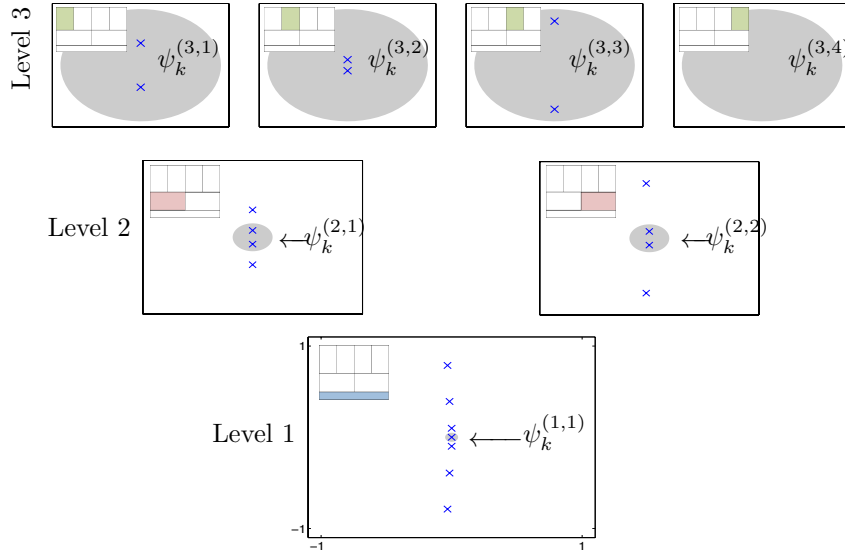Figure 3 shows the true modes along with the approximating mrDMD and DMD modes. The

Figure 5: Eigenvalues from level-1 to level-3 of the mrDMD decomposition applied to the data generated in Fig. 3. This is a specific case of the mrDMD abstraction shown in Fig. 1. The shaded circle shows the threshold radius used for background mode subtraction. The inset illustrates the time-frequency position of each sampling window. Also illustrated are the mode selection and labeling process.

inset in the mrDMD and DMD modes shows the ℓ_2 difference between the true modes and the approximating modes of the respective decomposition, i.e. $R_{j,\ell} = \|\bar{\psi}_j - \psi_j^{(\ell,1)}\|$ and $R_j = \|\bar{\psi}_j - \psi_j\|$ respectively. The mrDMD modes are almost identical to the true modes while the DMD modes show an error in mixing the modes. The error in the DMD modes is clearly illustrated in the time dynamics. Specifically, the DMD does not correctly capture the on-off dynamics of modes three and four so that $\|\bar{x} - x_{\text{DMD}}\| \sim O(1)$. In contrast, the mrDMD is able to easily capture the correct time dynamics with its windowing decomposition. Indeed, the agreement between the mrDMD and exact solution is remarkable, i.e. $\|\bar{x} - x_{\text{mrDMD}}\| \ll 1$.

To further illustrate the accuracy of the mrDMD decomposition in comparison to the DMD decomposition, consider Fig. 4. This shows the true solution $\bar{x}$ at the time snap shots of $t = 25, 50, 75, 100$ and 125. The mixing of the four modes of (22) is nicely captured by the mrDMD while the DMD fails to capture key features at these time points. Indeed, at time $t = 75$ in particular, the agreement is quite poor between the DMD and true solution.

As a final part of the analysis of the video sequence represented by $\bar{x}(t)$ in (22), we can explicitly perform the multi-resolution decomposition depicted in Fig. 1. In this case, the dominant eigenvalue are kept at each level of the decomposition. The bottom centered box is the level 1 decomposition showing that there exists a single mode in the mode removal region (shaded blue inset). This is

11

mean temperature mode

$\psi_1^{(1,1)}$

(a)

$\psi_2^{(1,1)}$

(b)

$\psi_1^{(4,7)}$
$t \approx 1997$

(c)

$\psi_1^{(4,9)}$
$t \approx 1999$
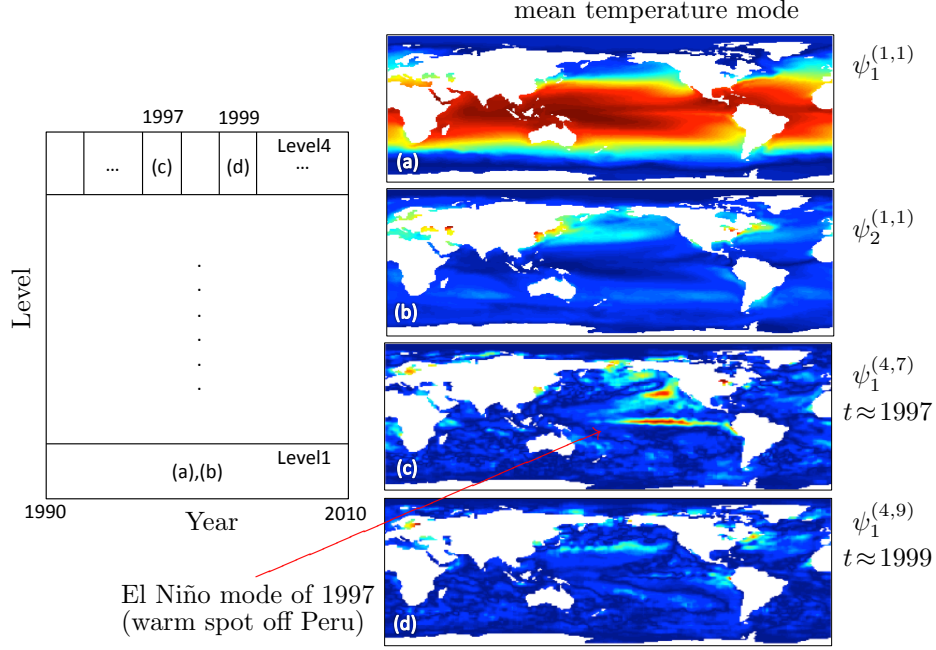
(d)

El Niño mode of 1997
(warm spot off Peru)

Figure 6: Application of mrDMD on sea surface temperature data from 1990 to 2010. The left panel illustrates the process for a 4-level decomposition. At each level, the *slowest* modes are extracted. At a given level, the zero-mode component has a period $T = \infty$. Illustrated are the (a) Level-1 mrDMD mode with period $T = \infty$, and (b) Level-1 mrDMD mode with period $T = 52$ weeks. These two modes are in the gray mode removal region of Fig. 1. Further in the decomposition we can extract (c) Level-4 mrDMD mode of 1997 with period $T = \infty$ and (d) Level-4 mrDMD mode of 1999 with period $T = \infty$. Mode (c) clearly shows the El Niño mode of interest that develops in the central and east-central equatorial Pacific. The El Niño mode was not present in 1999 as is clear from mode (d). Data source: NOAA_OI_SST_V2 data provided by the NOAA/OAR/ESRL PSD, Boulder, Colorado, USA, from their Web site at http://www.esrl.noaa.gov/psd/

mode $\psi_1^{(1,1)}$. At the next level, there are two modes (complex conjugate pairs) in the mode removal region (shaded pink inset), $\psi_1^{(2,1)}$ and $\psi_2^{(2,1)}$ (left panel) and $\psi_1^{(2,2)}$ and $\psi_2^{(2,2)}$ (right panel). These modes represent the slow oscillatory dynamics driven by $a_2(t)$. And in the final level, an additional pair of eigenvalues are present in the first three frames of the decomposition that correspond to the dynamics of modes $a_3(t)$ (frames 1 and 2) and $a_4(t)$ (frame 3) with modes $\psi_1^{(3,j)}$ and $\psi_2^{(3,j)}$. Thus the multi-resolution strategy outlined in Fig. 1 holds. Moreover, the intuition derived from this wavelet-like strategy is remarkably effective in decomposing the multi-resolution, spatio-temporal dynamics of $\bar{\mathbf{x}}$.

12

4.2 Multi-scale time separation of complex system: El Niño, Southern Oscillation

The example of the previous subsection was contrived to demonstrate the ability of the mrDMD to separate the spatio-temporal modes of Fig. 3. In this example, we use the mrDMD on a more realistic data set. Specifically, we consider global surface temperature data over the ocean. The data is open source from the NOAA/OAR/ESRL PSD, Boulder, Colorado, USA. The NOAA_OI_SST_V2 data set considered can be downloaded at http://www.esrl.noaa.gov/psd/ The data spans a 20 year period from 1990 to 2010.

Figure 6 shows the results of the mrDMD algorithm. Specifically, a 4-level decomposition is performed with the slow spatio-temporal modes pulled at each level as suggested in Fig. 1. The zero mode in each window is the DC component, or period infinity $T = \infty$ mode. At the first level of the decomposition, two modes are extracted: the zero mode ($T = \infty$) depicted in Fig. 6(a) and a yearly cycle ($T = 52$ weeks) shown in Fig. 6(b). The yearly cycle is the *slowest* mode extracted at level 1. Note that the zero mode component of level 1, $\psi_1^{(1,1)}$, is just the average ocean temperature over the entire 20-year data set.

We can continue the mrDMD analysis through to the fourth-level. At the fourth level, the approximate zero mode (period $T = \infty$) of the sampling window extracts physically interesting results. In particular, the data-driven method of the mrDMD discovers the 1997 El Niño mode generated from the well-known El Niño, Souther Oscillation (ENSO). Indeed, 1997 is known to have been a strong El Niño year, as verified by its strong modal signature in the 4th level decomposition of the mrDMD. In contrast, the same sampling window shifted down to 1999 produces no El Niño mode, which is in keeping with known ocean patterns that year. El Niño is the warm phase of the ENSO cycle and is associated with a band of warm ocean water that develops in the central and east-central equatorial Pacific (between approximately the International Date Line and 120°W), including off the Pacific coast of South America. The mrDMD mode clearly shows this band of warm ocean water as a spatio-temporal mode in 1997 in Fig. 6(c). ENSO refers to the cycle of warm and cold temperatures, as measured by sea surface temperature, SST, of the tropical central and eastern Pacific Ocean. El Niño is accompanied by high air pressure in the western Pacific and low air pressure in the eastern Pacific. The cool phase of ENSO is called La Niña with SST in the eastern Pacific below average and air pressures high in the eastern and low in western Pacific. The ENSO cycle, both El Niño and La Niña, causes global changes of both temperatures and rainfall.

These results could not have been produced with DMD unless the correct sampling windows were chosen ahead of time, thus requiring a supervised learning step not required by the mrDMD. Further, even in such a case, the previous slow modes, such as those of level-1 demonstrated in Fig. 6(a) and (b) would pollute the data at the level of investigation. Thus the mrDMD provides a principled, algorithmic approach that is capable of data-driven discovery in complex systems such as ocean/atmospheric data.

4.3 Translating and/or Rotating Structures

The final application of the mrDMD is on an example that is notoriously difficult for SVD-based methods to characterize, namely, when translational and/or rotational structures are present in the data snapshots, i.e. continuous or discrete invariances. Indeed, such invariances completely undermine our ability to compute low-rank embeddings of the data as driven by correlated structures, or POD/PCA modes. This has been the Achilles heel of many SVD based methods, thus requiring in
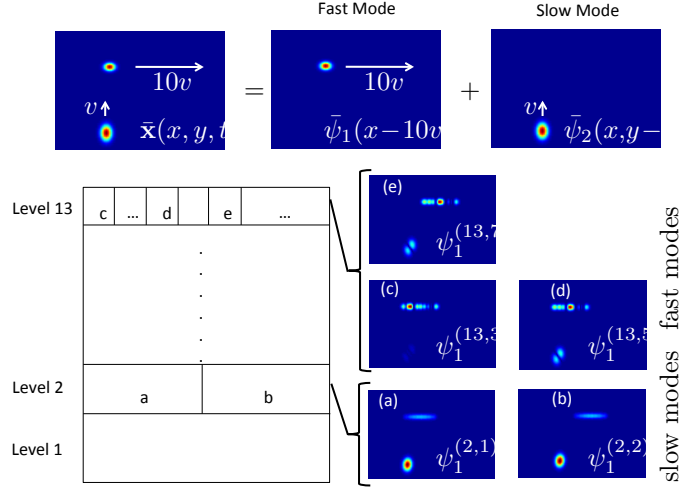
Figure 7: Application of mrDMD on moving objects separation. Two modes (top panel labeled "Fast Mode" and "Slow Mode") are combined in the data snapshots (top right panel). The "Fast Mode" moves from left to right as indicated by the arrow with a speed of $10v$ while the "Slow Mode" moves from bottom to top at the speed v. We take $v = 1/40$ without loss of generality. Thus the fast and slow mode speeds are approximately an order of magnitude different. In the mrDMD decomposition (left bottom panel), the "Slow Mode" is extracted at level 2 as represented by the panels (a) and (b). The "Fast Mode" is extracted at level 13 as represented at three representative panels (c)-(e). Although there is some shadow (residual) of the slow mode on the fast mode and vice-versa, the mrDMD is remarkably effective in extracting the correct modes. Moreover, the level at which they are extracted can allow for a reconstruction of the velocity and direction. To our knowledge, this is the best performance achieved to date with an SVD-based method with multiple time-scale objects.

applications such as PCA-based face recognition (eigenfaces) [21], well-cropped and centered faces for reasonable performance, i.e. translation and rotation are removed in an expensive pre-processing procedure.

In a dynamical systems setting, a simple traveling wave will appear to be a high-dimensional object in POD space despite the fact that it is only constructed from two modes, one associated with translational invariance. For dynamical cases exhibiting translation and/or rotation, Rowley and Marsden [22] formulated one of the only mathematical strategies to date to extract the low-dimensional embeddings. In particular, they developed a template-matching technique to first remove the invariance before applying the SVD decomposition. Although effective, it is not suited for cases where a myriad of objects and timescales are present in the data.

The mrDMD is well suited to handle invariances such as translation and rotation. Consider once again the case of a simple traveling wave. The standard DMD decomposition applied to the traveling wave would result in a solution approximation requiring many DMD modes due to the slow fall off of the singular values (8) in step 1 of the DMD algorithm. Further, the eigenvalues

ω_k in (12) would also be bounded away from the origin as there is no *background* mode for such translating data. In the mrDMD architecture, the fact that the eigenvalues are bounded away from the origin (and typically $O(1)$) in the initial snapshot window $\mathbf{X}_M$ would simply allow the mrDMD method to ignore the traveling wave at the first level of the mrDMD. Once the domain is now divided into two for the next level of analysis, the traveling wave is now effectively moving at half the speed in this new domain, i.e. the eigenvalues have migrated towards the origin and the traveling wave is now re-evaluated. The recursive procedure would eventually produce a sampling window where the traveling wave looks sufficiently stationary and low-rank so as to be extracted in the multi-resolution analysis. The level at which it is extracted also characterizes the speed of the traveling wave. Specifically, the higher the level in the decomposition where the traveling wave is extracted, the faster its speed.

Figure 7 demonstrates the application of the mrDMD method on a simple example in which there are two moving objects, one moving at a slow velocity and another moving at a high velocity. Specifically, the example constructed results from the dynamical sequence

$$(23) \qquad \bar{\mathbf{x}} = \bar{\psi}_1(x - 10vt, y) + \bar{\psi}_2(x, y - vt)$$

where the two modes (Gaussians of the form $\bar{\psi}_j = \exp(-\sigma(x - x_0)^2 - \sigma(y - y_0)^2)$ with $\sigma = 0.1$ and $(x_0, y_0) = (-18, 20)$ and $(x_0, y_0) = (-20, -9)$ for the fast and slow modes respectively) used to construct the true solution are represented by $\bar{\psi}_j(x, y)$ for $j = 1, 2$. Note that the first mode is translating from left to right at speed $10v$ whereas the second mode is translating from bottom to top at speed v. Without loss of generality, v can be set to any value. It is chosen to be $v = 1/40$ for the domain and Gaussians considered. In this case, the straightforward template matching procedure would fail due to the two distinct time-scales of the objects. As shown in the figure, the mrDMD is capable of pulling out the two modes at level 2 (slow) and 13 (fast) of the analysis. It is at these levels that the translating objects look close to stationary ($|\omega_k| \ll 1$) in the mrDMD analysis. The level 2 modes correspond to the slow-moving object whereas the level 13 modes are associated with the fast object. Although there is a residual in both the extracted slow and fast modes, it does a reasonable job in extracting the fast and slow modes. To our knowledge, this is the only SVD-based method capable of such unsupervised performance.

The motivation for this example is quite clear when considering video processing and surveillance. In particular, for many applications in this arena, background subtraction is only one step in the overall video assessment. Identification of the foreground objects is the next, and ultimately more important, task. Consider the example of a video of a street scene in which there is a pedestrian walking (slow translation) along with a vehicle driving down the road (fast translation). Not only would we want to separate the foreground from background, we would also want to separate the pedestrian from the vehicle. More precisely, we would want to make a separate video and identification of the slow and fast objects in the video. The results of Fig. 7 show the mrDMD to be well-suited and effective in this task. Moreover, one can also envision augmenting the algorithm with a recursive application of mrDMD to improve performance much like in the foreground/background separation work on DMD [19].

5 Discussion and Outlook

Data-driven strategies for analyzing complex systems are of growing interest in the mathematical sciences. Indeed, methods capable of providing principled decompositions of data arising in multi-scale, spatio-temporal systems are key enabling strategies for many applications in the engineering,

physical and biological sciences. The multi-resolution dynamic mode decomposition advocated in this work capitalizes on recent innovations in equation-free modeling via the dynamic mode decomposition. It leverages these ideas by integrating them with concepts from wavelet theory and multi-resolution analysis. By construction, the method provides a principled reconstruction of multi-resolution, spatio-temporal data sets. The methods effectiveness is demonstrated on several example data sets, highlighting its ability to extract critical information and enact data-driven discovery protocols.

The DMD method provides a spatio-temporal decomposition of data into a set of dynamic modes that are derived from snap shots or measurements of a given system in time. The DMD method approximates the modes of the so-called Koopman operator. The Koopman operator is a linear, infinite-dimensional operator that represents nonlinear, infinite-dimensional dynamics without linearization, and is the adjoint of the Perron-Frobenius operator. The method can be viewed as computing, from the experimental data, the eigenvalues and eigenvectors (low-dimensional modes) of a linear model that approximates the underlying dynamics, even if the dynamics is nonlinear. By interpreting the DMD eigenvalues as corresponding to prescribed time scale dynamics, one can extract spatio-temporal structures recursively for shorter and shorter sampling windows. Thus the slow-modes are removed first and the data is filtered for analysis of its higher frequency content. This recursive sampling structure is demonstrated to be effective in allowing for a reconstruction of several example data sets.

One can envision a number of innovations to augment the proposed mrDMD strategy. Many are particularly attractive for applications across the engineering, physical and biological sciences. Indeed, the impact that DMD is having on complex fluid flows is already known and has already been mentioned in the introduction. Fields like neuroscience, which are rich in multi-scale, complex dynamics are also ideal candidates for exploration using the mrDMD infrastructure as DMD has already had recent demonstrated success in this arena [23]. The outlook of these techniques is highlighted here:

Compressive sampling: First is the ability to leverage tools from compressive sampling [15] to facilitate the collection of considerably fewer measurements, resulting in the same multi-resolution dynamic mode decomposition, but with considerably fewer measurements. This reduction in the number of measurements may have a broad impact in situations where data acquisition is expensive and/or prohibitive. In particular, we envision these tools being used in particle image velocimetry (PIV) to reduce the data transfer requirements for each snapshot in time, increasing the maximum temporal sampling rate. Other applications include ocean and atmospheric monitoring, where individual sensors are expensive. Even if full-state measurements are available, the proposed method of compressed DMD will be computationally advantageous in situations where there is low-rank structure in the high-dimensional data.

DMD Control: A second important direction revolves around recent innovations of the DMD with control [24], which is capable of disambiguating between the underlying dynamics and the effects of actuation, resulting in accurate input-output models. The method is data-driven in that it does not require knowledge of the underlying governing equations, only snapshots of state and actuation data from historical, experimental, or black-box simulations. One can envision developing such input-output models at the various levels of spatio-temporal encoding in the data, i.e. multi-resolution input-output models can be assessed and constructed from a given system. Both these research directions, and many others, highlight the potential strength of the model in pushing forward principled, equation-free strategies for the analysis of complex multi-scale systems.

Machine Learning: Modern tools of statistical analysis and dimensionality-reduction have become the workhorses for the burgeoning field of machine learning (ML). ML techniques aim to capitalize on underlying low-dimensional patterns and clustering in data. In the dynamical applications considered here, one might exploit these patterns, or DMD modes, by building libraries of low-rank dynamical modes, much like is done with POD modes [25, 26, 27]. Such DMD libraries for different dynamical regimes partner nicely with compressive sensing strategies. Additionally, Kernel based techniques, which are at the core of support vector machines, for instance, have already found successful application in the DMD architecture when considering more accurate, nonlinear dynamical reconstructions [28]. Maximum advantage should be taken of such techniques when integrating the DMD architecture in applications.

Acknowledgment

References

[1] I. Daubechies. *Ten Lectures on Wavelets*. SIAM, 1992.

[2] L. Debnath. *Wavelet transforms and their applications*. Birkhäuser, 2002.

[3] J. N. Kutz. *Data-driven modeling and scientific computing: Methods for Integrating Dynamics of Complex Systems and Big Data*. Oxford Press, 2013.

[4] P. J. Schmid and J. Sesterhenn. Dynamic mode decomposition of numerical and experimental data. In *61st Annual Meeting of the APS Division of Fluid Dynamics*. American Physical Society, November 2008.

[5] P. Schmid. *Dynamic mode decomposition of numerical and experimental data*. *Journal of Fluid Mechanics*, 656:5–28, 2010.

[6] C. Rowley, I. Mezić, S. Bagheri, P. Schlatter, and D. Henningson. *Spectral analysis of nonlinear flows*. *Journal of Fluid Mechanics*, 641:115–127, 2009.

[7] P. Schmid, L. Li, M. Juniper, and O. Pust. *Applications of the dynamic mode decomposition*. *Theoretical and Computational Fluid Dynamics*, 25(1-4):249–259, 2011.

[8] K. Chen, J. Tu, and C. Rowley. *Variants of Dynamic Mode Decomposition: Boundary Condition, Koopman, and Fourier Analyses*. *Journal of Nonlinear Science*, 22(6):887–915, 2012.

[9] J. Tu, C. Rowley, D. Luchtenberg, S. Brunton, and J. N. Kutz. *On Dynamic Mode Decomposition: Theory and Applications*. *Journal of Computational Dynamics*, 1:391–421, 2014.

[10] Igor Mezić and Andrzej Banaszuk. Comparison of systems with complex behavior. *Physica D: Nonlinear Phenomena*, 197(12):101 – 133, 2004.

[11] Igor Mezić. Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1-3):309–325, 2005.

[12] I. G. Kevrekidis, C. W. Gear, J. M. Hyman, P. G. Keverekidis, O. Runborg, and C. Theodoropoulos. *Equation-free, course-grained multi scale computation: enabling microscopic simulators to perform system-level analysis. Communications in Mathematical Sciences*, 1:715–762, 2003.

[13] I. G. Kevrekidis, C. W. Gear, and G. Hummer. *Equation-free: The computer-aided analysis of complex multiscale systems. Am. Instit. Chem. Eng.*, 50:1346–1355, 2004.

[14] I. Mezić. *Analysis of Fluid Flows via Spectral Properties of the Koopman Operator. Annual Review of Fluid Mechanics*, 45:357–378, 2013.

[15] S. Brunton, J. Proctor, and J. N. Kutz. *Compressive sampling and dynamic mode decomposition. Journal of Computational Dynamics*, to appear (2015).

[16] Jonathan H Tu, Clarence W Rowley, J Nathan Kutz, and Jessica K Shang. Spectral analysis of fluid flows using sub-nyquist-rate piv data. *Experiments in Fluids*, 55(9):1–13, 2014.

[17] L. N. Trefethen and D. Bau. *Numerical Linear Algebra*. SIAM, Philadelphia, 1997.

[18] M. Gavish and D.L. Donoho. The optimal hard threshold for singular values is $4/\sqrt{3}$. *Information Theory, IEEE Transactions on*, 60(8):5040–5053, Aug 2014.

[19] J. Grosek and J. N. Kutz. *Dynamic Mode Decomposition for Real-Time Background/Foreground Separation in Video. arXiv*, page arXiv:1404.7592, 2014.

[20] E. J. Candés, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *J. ACM*, 58:Article No. 11, 2011.

[21] M.A. Turk and A.P. Pentland. Face recognition using eigenfaces. In *Computer Vision and Pattern Recognition, 1991. Proceedings CVPR '91., IEEE Computer Society Conference on*, pages 586–591, Jun 1991.

[22] Clarence W. Rowley and Jerrold E. Marsden. Reconstruction equations and the karhunen-loéve expansion for systems with symmetry. *Physica D: Nonlinear Phenomena*, 142(12):1 – 19, 2000.

[23] B. Brunton, L. Johnson, J. Ojemann, and J. N. Kutz. Extracting spatial-temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition. *arXiv:1409.5496*, 2014.

[24] J. Proctor, S. Brunton, and J. N. Kutz. *Dynamic mode decomposition with control. arXiv:1409.6358*, 2014.

[25] I. Bright, G. Lin, and J. N. Kutz. Compressive sensing based machine learning strategy for characterizing the flow around a cylinder with limited pressure measurements. *Physics of Fluids*, 25:127102, 2013.

[26] S. Brunton, J. Tu, I. Bright, and J. N. Kutz. Compressive sensing and low-rank libraries for classification of bifurcation regimes in nonlinear dynamical systems. *SIAM App. Dyn. Sys.*, 13:1716–1732, 2014.

[27] J. L. Proctor, B. Brunton, S. Brunton, and J. N. Kutz. Exploiting sparsity and equation-free architectures in complex systems. *European Physics Journal Special Topics*, 223:2665–2684, 2015.

[28] M. Williams, C. W. Rowley, and I. G. Kevrekidis. A kernel-based approach to data-driven koopman spectral analysis. *arXiv:1411.2260*, 2015.