Inverse Modelling of Incompressible Gas Flow in Subspace

Xiao Zhai[1], Fei Hou[1,2,*], Hong Qin[3] and Aimin Hao[1]

[1]State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, China
{zhaixiao, houfei, ham}@buaa.edu.cn
[2]School of Computer Engineering, Nanyang Technological University, Singapore
[3]Department of Computer Science, Stony Brook University, USA
qin@cs.sunysb.edu

Abstract
This paper advocates a novel method for modelling physically realistic flow from captured incompressible gas sequence via modal analysis in frequency-constrained subspace. Our analytical tool is uniquely founded upon empirical mode decomposition (EMD) and modal reduction for fluids, which are seamlessly integrated towards a powerful, style-controllable flow modelling approach. We first extend EMD, which is capable of processing 1D time series but has shown inadequacies for 3D graphics earlier, to fit gas flows in 3D. Next, frequency components from EMD are adopted as candidate vectors for bases of modal reduction. The prerequisite parameters of the Navier–Stokes equations are then optimized to inversely model the physically realistic flow in the frequency-constrained subspace. The estimated parameters can be utilized for re-simulation, or be altered toward fluid editing. Our novel inverse-modelling technique produces real-time gas sequences after precomputation, and is convenient to couple with other methods for visual enhancement and/or special visual effects. We integrate our new modelling tool with a state-of-the-art fluid capturing approach, forming a complete pipeline from real-world fluid to flow re-simulation and editing for various graphics applications.

Keywords: model reduction, EMD, parameter estimation

ACM CCS: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism Animation

1. Introduction

Since early 1990s, various gas/fluid simulation methods for graphics and animation have been developed, and the majority of existing techniques are founded upon solving the incompressible Navier–Stokes equations or their possible variants. Along with the rapid advancement in graphics, computational fluid dynamics and relevant engineering disciplines, tremendous fluid-capture methods have been proposed in recent years, such as particle image velocimetry (e.g. Tomo-PIV [ESWVO06]), time-resolved Schlieren system [AIH*08], surface modelling [DLJY11, YJLY12], optical-flow-based methods [LS08, GKHH12], etc. Strongly inspired by the most recent works [WLZ*09, LPS*13, GIT*14] attempting to bridge the gap between simulation and capturing, in this paper, we propose to seek novel strategies to further integrate captured data and generated flow sequences.

Our inverse-modelling technique aims at generating gas sequences that best match the existing/captured flows by simulation. Different from previous works on fluid capturing [GIT*14] and guiding control [NCZ*09], we fulfil it by infering the prerequisite parameters of the flows according to the incompressible Navier–Stokes equations, so the flows can be naturally reproduced and edited through simulation without any prior knowledge on those parameters.

A pipeline of our technique can be found in Figure 2. This pipeline and our work are based on previous works [GKHH12, GIT*14], and the parameters estimation step is the core of our work, providing the ability of physical re-simulation and editing. Given that the effect of viscosity and that of external forces are hard to differentiate, we model the flows by an optimization process to best

---

*Corresponding author: Fei Hou (houfei@buaa.edu.cn)

| (a) | (b) | (c) | (d) | (e) |
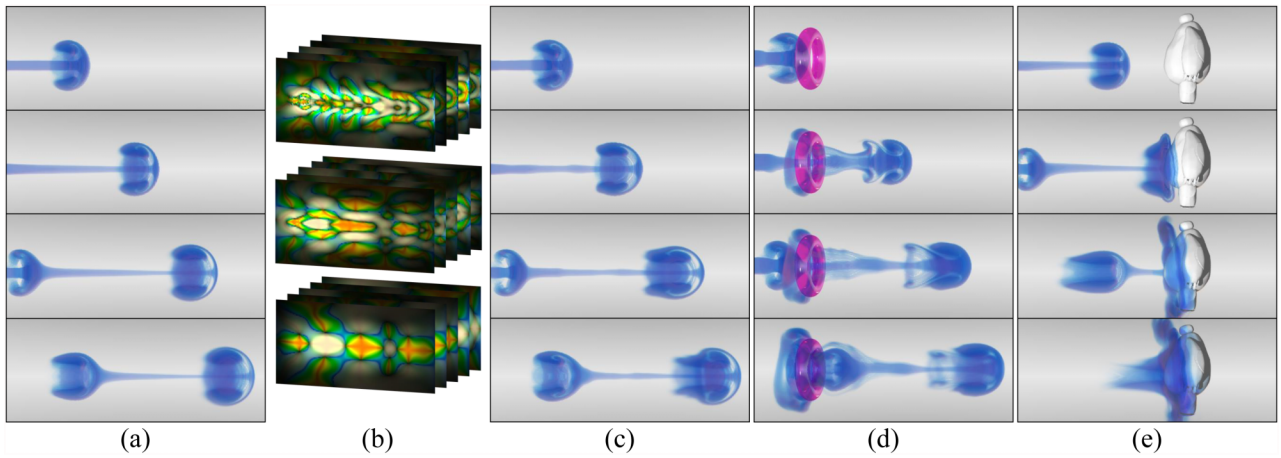
Figure 1: *Results on a simulated 'ejection' dataset. (a) The existing flow. (b) Sample empirical mode decomposition (EMD) bases in different frequency scales for modal reduction (mapped to HSV colour space for visualization). (c) Resimulation of the flow from physical inverse modelling. (d) Flow edited in subspace by adding a torus boundary. (e) Flow edited in the original space by applying a complex boundary.*
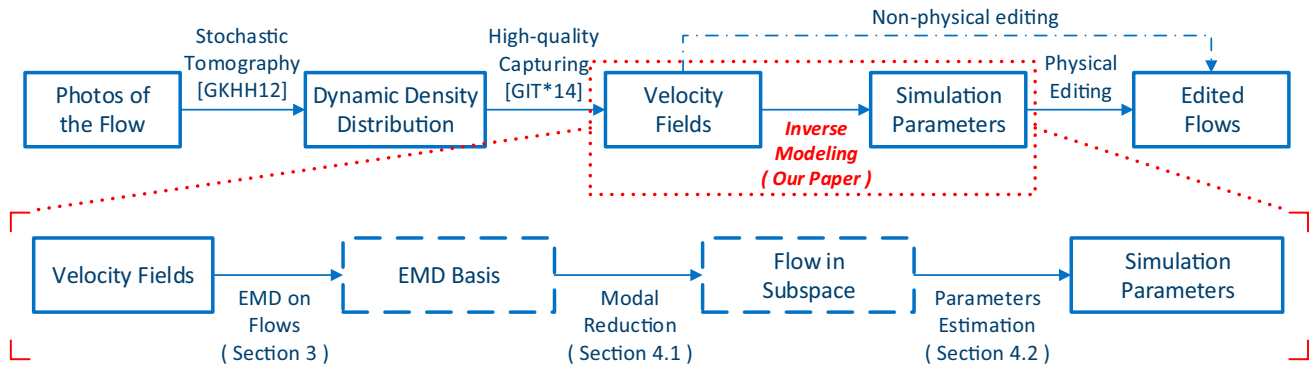
Figure 2: *The pipeline from real-world flows to fluid simulation/editing.*

fit the formula. The parameters estimation is conducted in subspace for acceleration. Flows after this step can be physically modified towards desired behaviours through multiple approaches, for instance, various driving forces, thickened air, varied boundaries, etc. More complex modifications, e.g. non-linearities like vorticity confinement, can be implemented in the original space to combat the limited degree of freedom in subspace.

Traditionally, the subspace of flow fields is constructed by principal component analysis (PCA) bases [TLP06], which are far less meaningful in frequency domain and may have inter-frequency energy mixing in some circumstances. Instead, we seek new bases via EMD [HSL*98], which separates flows into frequency bands to better capture the intrinsic features. The original EMD algorithm is proposed for 1D time series and 2D signal processing [HS05], and proves ideal for handling non-stationary and non-linear signals. However, the usage of EMD in fluid area is far from being adequate, since previous techniques only serialize the 3D flow field with space-filling curves [GLRH13, RLL*13]. We leverage thin-plate spline (TPS) in 3D interpolation and apply TPS-based 3D EMD to better expose the intrinsic features of fluid flows.

Most fluid capturing methods that provide velocity fields of the flows, such as optical-flow-based methods, are compatible with our technique. We integrate our tool with a state-of-the-art fluid capturing method from Gregson *et al.* [GKHH12, GIT*14]. The real-world fluid is captured with stochastic tomography [GKHH12], then the velocity fields are reconstructed by the fluid tracking algorithm [GIT*14]. We analyse the velocity fields with our proposed technique to acquire re-simulated sequences, forming a complete pipeline from real-world fluid to flow re-simulation and flow editing for various graphics applications, see Figure 2. These lead to the main contributions of this paper:

(1) We apply TPS-based 3D EMD algorithm for fluid analysis. The new algorithm leverages TPS for 3D interpolation, and takes into account the fluid-dynamical constraints, i.e. zero divergence and free-slip boundary conditions.

(2) We improve previous PCA-based model reduction on fluid flows by applying EMD and deliberately choosing modes from multiple frequency domains, which better exposes the intrinsic features through all frequencies.

(3) We devise an analytical tool based on EMD and modal reduction, capable of inversely modelling gas sequences and applying physical editing on existing flows. We further integrate it with a state-of-the-art fluid capturing method to bond the fluid simulation in graphics with real-world flows.

2. Related Work

Our work is closely relevant to fluid simulation, flow editing, fluid capturing, modal reduction and empirical mode decomposition (EMD). We briefly review them in the following categories.

Fluid simulation aims at calculating the flow motion by solving the incompressible Navier–Stokes equations. The grid-based fluid solver was first brought into graphics by Foster and Metaxas [FM96], and continued to retain its popularity with the unconditionally stable solver [Sta99]. Since then many techniques have been proposed to eliminate the numerical dissipation and to enrich the details [FSJ01, ZB05, KLLR05, SFK*08, YKH*09]. So far, even though simulators can produce astonishing visual results, it is still challenging to mimic real-world fluid motion merely by simulation due to the limitations on resolution, boundary condition and numerical issues. In this paper, we inversely model existing flows and try to provide sequences by simulation that best match them.

Flow editing on Navier–Stokes-based simulation was first proposed by Foster and Metaxas [FM97]. Keyframe control was then introduced on smoke simulation [TMPS03], and boosted by McNamara *et al.* [MTPS04]. Fattal and Lischinski [FL04] and Shi and Yu [SY05] used target shapes to guide smoke, while Rasmussen *et al.* [REN*04] applied soft or hard control over liquids by introducing particles. Thurey *et al.* [TKPR09] achieved fluid guiding in smoothed particle hydrodynamics (SPH) and lattice boltzmann method (LBM). Other techniques apply fluid control to produce finer results out of a coarse guidance. Nielsen *et al.* [NCZ*09] proposed an optimization scheme to ensure the consistency between a high-resolution smoke simulation and a low-resolution velocity field. Yuan *et al.* [YCZ11] yielded high-resolution smoke animation following the low-resolution version by Lagrangian coherent structure, and Nielsen *et al.* [NB11] guided high-resolution liquid simulation by its low-resolution version. We present a distinct physics-based editing method in this paper by inversely modelling the flows and changing their physical parameters.

Fluid capturing has gained abundant attention in graphics recently. Capture measurements roughly include surface modelling [DLJY11, YJLY12], density estimation [AIH*08, GKHH12] and velocimetry [ESWVO06, LS08]. Recent researches exhibit a trend of applying physical constraints or guidance rather than purely relying on capturing methods. Wang *et al.* [WLZ*09] utilized physically based surface optimization to reconstruct water surface from videos, while Li *et al.* [LPS*13] acquired water surface from a single viewpoint through shallow water model. Most lately, Gregson *et al.* [GIT*14] presented a proximal optimization framework that related the captured density data to corresponding velocity fields. Inspired by their work, we propose a novel method aiming at inverse modelling from existing velocity fields to the simulation progress that produces them, which can also be regarded

as an analytical measurement capable of reusing the intrinsic properties carried by the flow.

Modal reduction has widely been used in graphics, including non-linear deformation [JF03] and finite element method [BJ05]. Treuille *et al.* [TLP06] first introduced model reduction with PCA bases to fluid simulation and generated real-time flows. Subsequently, Wicke *et al.* [WST09] generalized this technique to modular tiles to fit more complicated scenes. DeWitt *et al.* [DWLF12] replaced PCA eigenvectors with Laplacian eigenfunctions which are suitable for arbitrary flows. Kim and Delaney [KD13] proposed cubature approach to support semi-Lagrangian advection and other non-linearities. The inverse modelling step in this paper requires huge computation; hence, we adopt model reduction technique for acceleration.

EMD was first introduced by Huang *et al.* [HSL*98] for temporal signal processing. Its data-dependent nature is ideal for analysing non-stationary data. In graphics and image processing researchers have extended EMD to handle 2D or 3D signals. Linderhed [Lin11] presented a 2D EMD method to compress images. Wang *et al.* [WHZQ15] applied EMD for 3D geometry processing. The difficulty of applying EMD on multi-dimensional data lies in constructing the envelopes. Wu *et al.* [WHC11] proposed multi-dimensional ensemble EMD that partitioned data into 1D slices for decomposition. To interpolate the exact envelopes, Damerval *et al.* [DMP05] used Delaunay triangulation and piecewise cubic polynomial interpolation, Xu *et al.* [XLLR06] presented a mesh fitting method and Linderhed [Lin11, Lin09] made comparisons on existing solutions and suggested using TPS for smooth result with continuous second derivative. Rehman and Mandic [RM10] even extended EMD to process vector values by projections on hyperspheres. Recently, EMD has been adapted into fluid simulation [GLRH13] and modulation [RLL*13]. Their sifting operations are achieved by serializing 3D data to 1D with a space-filling curve, and then using 1D EMD algorithm. However, this method loses connectivity information in 3D, resulting in spurious noises. In this paper, we apply a full 3D EMD variant that can eliminate such deterioration, and hence is much more appropriate for flow field. We further use the resulting components as candidate bases in model reduction.

3. Empirical Mode Decomposition on Flows

EMD, originally designed for analysing time-varying signals in 1D [HSL*98], needs to be extended to 3D and requires other modifications for processing flow fields. We briefly review the standard EMD method in Section 3.1, Section 3.2 explains the strategy to apply EMD to 3D signals on Eulerian grids and Section 3.3 discusses the EMD method further extended to process 3D flow fields. The modified EMD algorithm for 3D flows is listed in Algorithm 1.

3.1. Conventional EMD

The functionality of EMD is to decompose a signal $x(t)$ into a finite number of intrinsic mode functions (IMFs) $c_i, i = 1, 2, \ldots, q$ and a residue r_q, which characterize the features of the signal on different scales:

$$x(t) = \sum_{i=1}^{q} c_i + r_q. \tag{1}$$

Conceptually, the IMFs and the residue can be considered as the counter-part of 'frequency components' in other decomposition methods. However, in comparison with traditional Fourier and wavelet analysis, their frequencies are not pre-defined but data-dependent, and they can offer spatially-varying 'instantaneous frequencies'. This self-adaptive nature provides the suitability for handling highly non-stationary and non-linear signals [HSL*98, WHC11]. The EMD algorithm that extracts IMFs through 'sifting' is documented in [HSL*98]. In each iteration, an IMF representing features with higher frequency ranges is extracted, leaving the lower frequency part as residue or input of the subsequent iterations. Thus, IMFs are arranged spontaneously in the order of their frequency ranges.

3.2. EMD on 3D data

The original EMD method is powerful to deal with scalar data defined in 1D space, but needs modifications to fit the velocity fields of flows, three-component vector signals in 3D space. Although vector signals can be processed by multi-variate EMD [RM10], the overhead makes it infeasible since Eulerian grids for gas simulation usually contain millions of cells. Recall that in traditional fluid solvers, three components of the velocity fields are often treated as individual channels and handled respectively. Likewise, we also carry out the sifting procedure on each scalar component, and similar solution of applying EMD on vector fields can be found in [GLRH13] and [RLL*13]. Experiments demonstrate that the results are still adequate with this simplification.

Now we explain EMD for 3D signals. Since we are only concerned with discrete 3D signals, the extrema points are simply extracted by comparing the candidate data point with its nearest 26-connected neighbours. The challenge lies in fitting a space to the 3D scattered data points representing the extrema points (line 8 in Algorithm 1). In image processing, Linderhed [Lin11, Lin09] compared existing solutions, including directional EMD, triangle-based cubic spline interpolation and TPS interpolation, and suggested using TPS interpolation for smooth result with continuous second derivative. Conceptually, TPS is the generalization of the natural cubic splines in 1D. The construction is based on choosing a function $f(\mathbf{x})$ that exactly interpolates the data points and minimizes the bending energy,

$$E[f] = \int_{\mathbb{R}^n} \|D^2 f\|^2 d\mathbf{x}, \qquad (2)$$

where $D^2 f$ is the matrix of second-order partial derivatives of f and $\|\cdot\|^2$ is the sum of squares of the matrix entries. The implementation details of TPS can be found in [Ebe02]. Given that the huge computational costs prevent previous works from using TPS interpolation, we resort to CUDA for acceleration and obtain acceptable efficiency.

The boundary constraints are more important in 3D than in 1D/2D cases. The extrema points are very sparse and since the interpolation methods only interpolate between points, the boundary needs special care. Similar to Linderhed [Lin11], we add extra data points to the set of extrema points (line 7 in Algorithm 1). The extra points are placed at the corners and some additional points at the boundary equally spaced between the corners. We
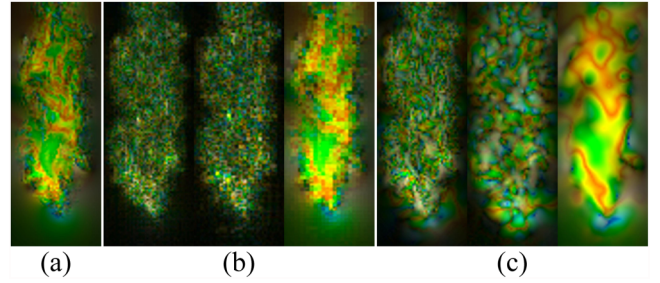
Figure 3: *EMD methods comparison at the 29th frame of the captured 'bloom' dataset. (a) is the original flow field, (b) is the result from 1D EMD with Hilbert space-filling curve and (c) is the result of our 3D EMD. For visualization, the 3D vectors are mapped to HSV colour space, and the central slice of the 3D domain is displayed. Our 3D EMD succeeds in extracting high-consistency, low-noisy frequency features, while the original method lacks global topology and induces spurious noises.*

place extra points at large blank spaces which lack control points as well, since regions without extremum implies monotonicity, which should be excluded from any IMFs and characterized as residue. Without these extra points, the spaces not covered by the interpolation yield bad results in the sifting process.

Algorithm 1: 3D EMD algorithm for flow fields.

1 Init: $u(t)$ is the x-,y- or z-component of the velocity;
2 $r_0 \leftarrow u(t), i \leftarrow 1$;
3 **for** *the i^{th} IMF c_i* **do**
4 $\quad h_0 \leftarrow r_{i-1}, k \leftarrow 1$;
5 $\quad$ **while** *true* **do**
6 $\quad\quad$ Find local extrema points of h_{k-1};
7 $\quad\quad$ Insert extra local extrema points;
8 $\quad\quad$ Get upper and lower envelopes of h_{k-1} $E_{max,k-1}, E_{min,k-1}$ by 3D TPS interpolation;
9 $\quad\quad$ Enforce $E_{max,k-1}$ and $E_{min,k-1}$ to satisfy the free-slip boundary condition;
10 $\quad\quad E_{mean,k-1} \leftarrow \frac{1}{2}(E_{max,k-1} + E_{min,k-1})$;
11 $\quad\quad h_k \leftarrow h_{k-1} - E_{mean,k-1}$;
12 $\quad\quad$ **if** *the IMF stopping criterion is satisfied* **then**
13 $\quad\quad\quad c_i \leftarrow h_k$;
14 $\quad\quad\quad$ break;
15 $\quad\quad$ **else**
16 $\quad\quad\quad k \leftarrow k+1$;
17 $\quad$ Project c_i to its divergence-free component $\bar{c}_i$;
18 $\quad r_i \leftarrow r_{i-1} - \bar{c}_i$;
19 $\quad i \leftarrow i+1$;

Our method has the advantage of being fully 3D, which retains the global topology in the original signal. A comparison between our method and previous EMD technique on flows [GLRH13, RLL*13] can be seen in Figure 3. The method we compare against is implemented by first serializing the 3D signals with Hilbert space-filling curve and then applying 1D EMD algorithm. We confine each method to only compute two IMFs and the rest

velocity is directly treated as residue. As shown in the comparison, our method manages to arrange properties in different frequency bands to their relevant frequency components, and generate high-consistency, low-noisy results, which is superior to previous technique mainly because of the global topology it retains.

3.3. Flow field mode decomposition

Making use of the 3D EMD method, the velocity field **u** can be decomposed as

$$\mathbf{u} = \sum_{i=1}^{q} \mathbf{u}^i + \mathbf{u}^r, \qquad (3)$$

where $\mathbf{u}^i, i = 1, 2, \ldots, q$ are the IMFs carrying flow features from local to global scales, and $\mathbf{u}^r$ is the residue after those IMFs have been extracted.

Nevertheless, to better fit fluid applications, especially the downstream modal reduction, we have to make a little bit more modification on the EMD pipeline. Treuille *et al.* [TLP06] specified that the velocity fields used to generate bases should meet two important properties: being free of divergence and satisfying the free-slip boundary condition. We fulfil these requirements by restricting the envelopes to the free-slip boundary condition (line 9 in Algorithm 1) and the pre-sifted IMFs divergence-free (line 17 in Algorithm 1). The free-slip boundary condition is fulfilled by setting the velocity of cells outside the boundary so that the normal velocity at the exact boundary equals zero, while the tangential velocity is left unchanged, and the divergence-free IMFs are obtained by pressure projection as the traditional Eulerian fluid solvers do. Since the input velocity fields satisfy these properties, the final residues should also do.

The IMF stopping criterion we adopt is the criterion suggested in [HSL*98], which sets an upper threshold $\sigma_{\max}$ of the standard deviation between two consecutive sifting procedures and a maximum number of iterations $k_{\max}$. During the EMD step, we do not repeatedly carry out the algorithm to the very end, because after several sifting processes, the flow field is already smooth enough as a total low frequency signal and further decomposition becomes dispensable. Instead, the number of IMFs q is fixed to a constant, and the residue signal after q sifting processes is directly saved in $\mathbf{u}^r$. In all experiments of this paper, we set $q = 2$ unless otherwise specified, so each velocity field $\mathbf{u}_i$ will have three frequency components ($\mathbf{u}_i^{low}$, $\mathbf{u}_i^{mid}$, and $\mathbf{u}_i^{high}$) after EMD step.

4. Physical Inverse Modelling on Flows

The EMD method introduced in Section 3 is a powerful tool that separates flow fields into various frequency components, which hereinafter refer to the outcome of EMD including IMFs and residues. In this section, we propose an inverse-modelling method based on subspace analysis to re-simulate or generate new sequences with varying conditions from a known gas flow.

The basic idea is to apply EMD on each frame of the original flow, reuse the features carried with frequency components by modal reduction and estimate the parameters implied within the original gas flow in subspace, see Figure 2. After estimation, the

flow can be re-simulated or modified towards desired behaviors. We concisely introduce modal reduction for fluid flows in Section 4.1, followed by the algorithm to estimate flow parameters in the incompressible Navier–Stokes equations in Section 4.2, and flow editing schemes are finally discussed in Section 4.3.

4.1. Modal reduction

Modal reduction means representing a high-dimensional vector and its time evolution in a lower dimensional space. Specifically, a flow field $\mathbf{u} \in \mathbb{R}^n$ can be represented as $\mathbf{r} \in \mathbb{R}^m$, where m is independent of the spatial resolution n and $m \ll n$. Usually, the m-dimensional subspace is *linear*, so an *orthonormal* projection matrix B can be found to connect the two spaces $\mathbf{u} = B\mathbf{r}$ and $\mathbf{r} \approx B^T \mathbf{u}$, with equality if $\mathbf{u}$ lies exactly in the subspace spanned by B. Linear differential equation $\dot{\mathbf{u}} = M\mathbf{u}$ can also be projected into subspace using Galerkin projection $\dot{\mathbf{r}} = B^T M B \mathbf{r}$, by which the incompressible Navier–Stokes equations can be solved within subspace. For more information about fluid simulation in subspace, please refer to [TLP06].

The process to find bases B is somewhat tricky. Previous methods [TLP06, KD13] apply PCA to compress a set of example velocity fields, and the first m eigenvectors are set as columns in B. However, we have found that this method fails to capture the scale information conveyed in the flows and produces modes neglecting their frequencies. Consequently, specific parts in the spectral domain cannot be represented independently within the subspace. In our experiments, the kinetic energy induced by abundant high-frequency behaviors is transferred from fine-to-coarse scales through the bases in the re-simulation. This inter-frequency energy mixing can result in an overall falsely enlarged energy in all frequency bands, see Figure 6.

We solve this problem by taking advantages of EMD, fulfilling its potential in frequency domain. First, we create a set of example fields for each frequency component, $U^{freq} = \{\mathbf{u}_i^{freq}\}$, where *freq* can be *low*, *mid* or *high*. Subsequently, PCA is performed on each set for compression, and the first several eigenvectors of each set are selected into the bases B. Finally, the combined bases B are orthonormalized by Gram–Schmidt process. Note that after orthonormalization, the basis functions will no longer be purely *low*, *mid* or *high* frequency, but this is totally fine since the subspace remains the same. Essentially, the PCA-based methods select basis functions regardless of their frequency, while our method ensures that the basis functions are chosen from all frequency bands; hence, the EMD-based bases are free from the inter-frequency energy mixing. A comparison between PCA bases and EMD bases can be seen in Figure 6, where the PCA bases fail to reproduce the exact dynamics as the EMD bases do, eliminating the feasibility of accurate physical re-simulation/editing by PCA bases. The corresponding frequency spectra are presented in Figure 7, which indicates that our EMD-based method is able to recover more accurate energy through the frequencies compared with the PCA-based methods. The composition of bases B from each frequency band can be adjusted to accomplish style control, and typically $\frac{m}{3}$ vectors from each set are used for general purpose.

4.2. Parameters estimation

Having the reduction bases B, now we present the algorithm to estimate the parameters in fluid equations. The incompressible
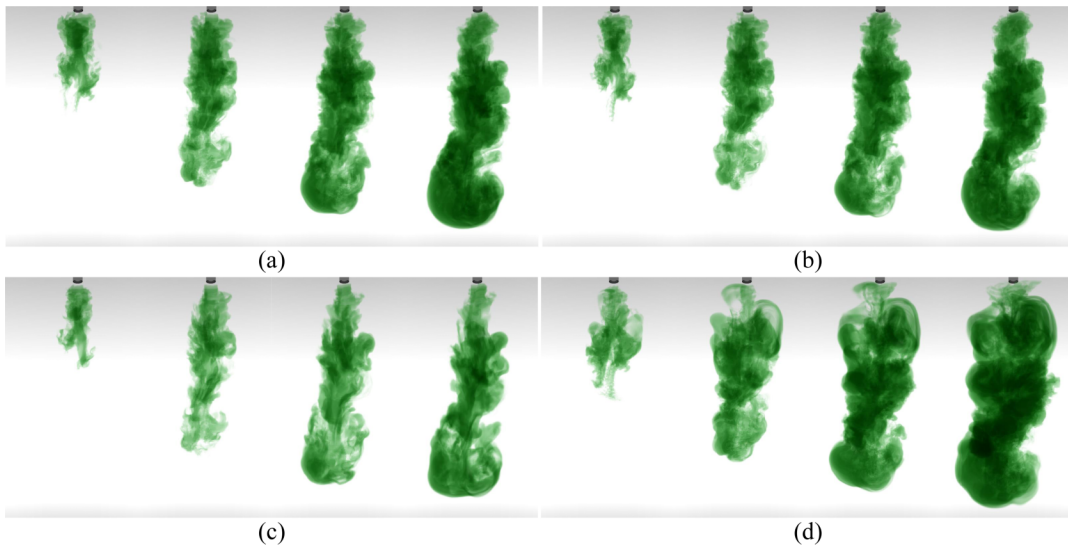
Figure 4: *Experiments on the captured 'bloom' dataset. (a) Existing flow. (b) Re-simulation of the inversely modelled flow with our technique. (c) Flow edited by doubling the viscosity. (d) Flow edited by adding extra impulse.*
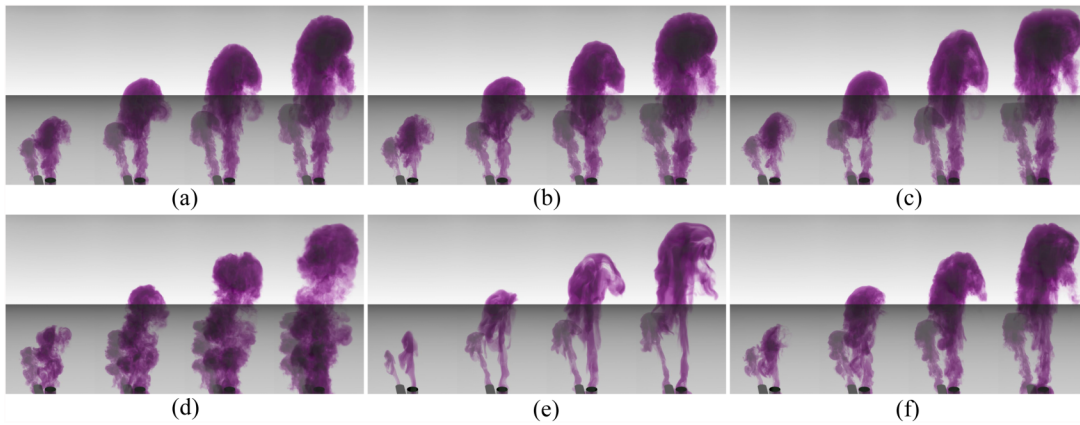
Figure 5: *Experiments on the captured 'smoke' dataset. (a) Existing flow. (b) Re-simulation of the inversely modelled flow. (c) Re-simulation with PCA bases that contain identical number of columns as our EMD bases. (d) Effect of vorticity confinement, $\epsilon = 0.2$. (e) Style control by removing high-frequency bases. (f) Style control by removing middle-frequency bases.*

Navier–Stokes equations are as follows:

$$\dot{\mathbf{u}} = -(\mathbf{u} \cdot \nabla)\mathbf{u} - \nu \nabla^2 \mathbf{u} + \nabla p + \mathbf{f}, \qquad (4)$$

$$\nabla \cdot \mathbf{u} = 0, \qquad (5)$$

where p denotes the pressure, ν the viscosity and $\mathbf{f}$ the external forces. When integrating, the pressure is solved by the Poisson equation $\nabla^2 p = \nabla \cdot \mathbf{u}$; therefore, only the viscosity and external forces are left unknown. The purpose of inverse modelling is to estimate these parameters and recreate a simulated flow sequence according to Navier–Stokes equations. Given the large size of the velocity field $\mathbf{u}$ and the time-consuming pressure projection step, the estimation is hard to implement within acceptable time without modal reduction.

Similar with Treuille *et al.* [TLP06], we discretize the advection term using finite difference as $\dot{\mathbf{u}} = A_u \mathbf{u}$, and the diffusion term as $\dot{\mathbf{u}} = \nu D \mathbf{u}$, where A_u is dependent on $\mathbf{u}$ but D is not. In subspace, each vector in the bases has its own advection matrix $A_{\mathbf{u}_i}$, and the total advection term can be denoted by $A = \sum_i r_i A_{\mathbf{u}_i}$. The corresponding linear differential equation of Navier–Stokes equations in subspace is

$$\dot{\mathbf{r}} = \hat{A}\mathbf{r} + \nu \hat{D}\mathbf{r} + \hat{\mathbf{f}}, \qquad (6)$$

where $\hat{A}$ is the contracted advection tensor $\hat{A} = \sum_i r_i B^T A_{\mathbf{u}_i} B$, $\hat{D}$ is the diffusion matrix $\hat{D} = B^T D B$ and $\hat{\mathbf{f}}$ is the reduced external force $\hat{\mathbf{f}} = B^T \mathbf{f}$. The pressure term and the divergence constraint vanish in subspace, because all vectors in the bases are inherently free of divergence. So it is with the boundary condition. The reduced matrices $\hat{D}$ and $B^T A_{\mathbf{u}_i} B$ are supposed to be pre-computed for acceleration.

(a) (b) (c)
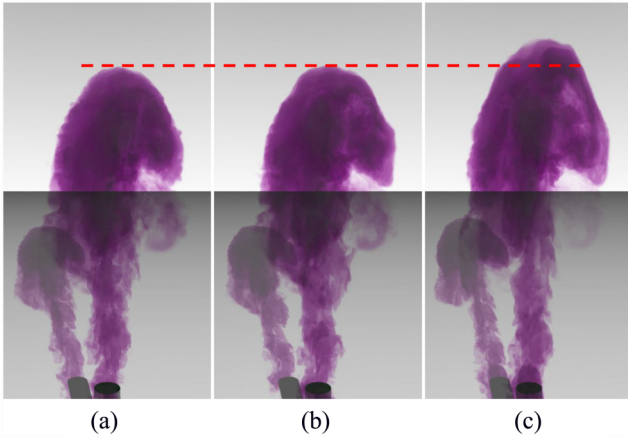
Figure 6: *Comparison about the inter-frequency energy mixing effect on PCA bases and EMD bases at the 63rd frame of the captured 'smoke' dataset. (a) Existing flow. (b) Inversely modelled with EMD bases. (c) Inversely modelled with PCA bases which have equal number of columns as our EMD bases. The horizontal red line indicates the flaw caused by inter-frequency energy mixing making that modelling with PCA bases has falsely enlarged energy.*

Table 1: *Default parameters used in our experiments.*

	Parameter	Value
EMD	q in Section 3.1	2
	σ_{max} in Section 3.3	0.2
	k_{max} in Section 3.3	10
Modal reduction	m for the captured datasets	75
	m for the 'ejection' dataset	90

Suppose an inviscid flow ($\nu = 0$), the external forces can be easily calculated from Equation (6), and vice versa. However, if the viscosity and external forces exist at the same time, it is hard to distinguish. Our solution is to assume that the external forces are as small as possible throughout all frames, that is, to minimize the objective function

$$\underset{\nu}{\text{argmin}} \sum_{i \in F} \| B(e^{-\Delta t(\hat{A}_i + \nu \hat{D})} \mathbf{r}_{i+1} - \mathbf{r}_i) \|_2^2, \text{ s.t. } \nu \geq 0, \quad (7)$$

for appropriate viscosity, where F denotes the set of frames. This method essentially maximizes the viscous effect, which is true in most cases, and yields nice results. Afterwards, the external force of each frame can be calculated through Equation (6).

For the flows on which we have prior knowledge about their inflow source, denoted by S, an ameliorant is to exclude the external forces near S (manually picked with no need to be exact) and minimize the magnitude of external forces in other regions. In practice, we remove the rows corresponding to the vicinity of S from bases B. For performance consideration, down-sampling is adopted both spatially and temporally. The bases B is down-sampled by removing rows randomly, and we select a random subset from F to apply the optimization. Despite the large size of B, this optimization can be very fast since only $B^T B$ is used, which is rather small and can be precomputed. In case of deterioration by down-sampling, this procedure can be repeated several times to average the estimated viscosity.

4.3. Flow editing

Flows after inverse modelling can be reproduced by re-simulation, and modified towards desired behaviors through multiple approaches, for instance, various driving forces, thickened air,

varied boundaries, etc. Such subspace editing approaches are detailed in [TLP06, KD13]. More complex modifications, e.g. non-linearities like vorticity confinement, can be implemented in the original space to combat the limited degree of freedom in subspace.

It should be noted that the flow editing of our technique differs from that of [GIT*14]. Although [GIT*14] and our technique can both produce flows from captured data, flow editing approaches proposed by [GIT*14], including resolution enhancement, domain change and guided simulation, could be non-physical since it only has the velocity fields but lacks the simulation parameters, which are indispensable for physics-based re-simulation and editing. In comparison, our inverse modelling technique provides the ability to apply physical editing following both the Navier–Stokes equations and the input flows, see Figure 2.

With the frequency-aware property of EMD, style control on the modelled flows is an interesting application that cannot be achieved by PCA bases. When inverse modelling, the composition of bases B can be adjusted by selecting unequal eigenvectors from different frequency bands. More vectors help capture the intrinsic features more precisely within a frequency band, and fewer vectors smooth out the specific effect. Inter-frequency energy mixing is less conspicuous with EMD-based bases, hence changing the number of vectors from a certain band will not affect the behaviors of other frequency bands. As a matter of fact, this method yields different results by altering the subspace. However, it should be noted that removing columns from the low-frequency band is usually not recommended, since the overall flow may well be altered completely if the subspace is overly shrunk that way.

5. Experimental Results

In this section, we provide numerical evaluations as well as several applications of our method. We implement our EMD algorithm with CUDA for acceleration, while the optimization in Equation (7) is solved with interior point method by calling the MATLAB Engine API. The experiments run on a PC with Geforce GTX970 GPU and Core i7 CPU, and our results are rendered by POV-Ray.

Several applications of our technique are documented below, including re-simulating and flow editing within or without subspace with convincing effects. The captured flow datasets are from Gregson *et al.* [GKHH12], and recovered by fluid tracking method documented in [GIT*14]. A simulated gas ejection is also included as ground truth in our experiments. Table 1 shows the default parameters used in our experiments and Table 2 lists the

Table 2: *Time performance (in seconds) of our experiments.*

Scene	#(Resolution, Frames)	EMD	Modal reduction	Inverse Modelling	Re-simulation
Ejection (Figure 1)	$(160 \times 80 \times 80, 120)$	816.197	935.108	1.803	0.710 ($\sim$169 fps)
Bloom (Figure 4)	$(75 \times 165 \times 75, 60)$	1250.441	492.536	0.602	0.213 ($\sim$281 fps)
Smoke (Figure 5)	$(75 \times 165 \times 75, 94)$	3665.044	490.953	1.000	0.334 ($\sim$281 fps)

experimental performance, indicating the high efficiency of the proposed inverse-modelling and editing technique.

Table 3 verifies the viscosity estimation in Equation (7). We empirically set the inflow source S on the simulated 'ejection' dataset, while on captured datasets, no inflow source is specified. B is randomly down-sampled by a factor of $\frac{1}{8}$, and F by $\frac{1}{4}$. Five results on each dataset are provided to prove the consistency, and the first one is adopted to perform the downstream re-simulation and flow editing.

Figure 1 displays the results on a simulated 'ejection' flow. The flow in Figure 1(a) is inversely modelled using our technique, and re-simulated as seen in Figure 1(c). The flow is generally well recovered, except some slight disagreement lying in the high-frequency area between the existing flow and the re-simulated one, caused by the discrepancy of their numerical integration methods. The existing flow is simulated by semi-Lagrangian method, while in subspace, the differential equation has analytical solution. Kim and Delaney [KD13] reported that the distinction of the integration methods may cause clashes on the high-frequency modes. Nonetheless, this flaw is weakened if we apply our technique on captured datasets, where there is no discrepancy of integration methods; hence, no disagreement is found in those experiments. An additional torus barrier is placed in the domain as shown in Figure 1(d) to test the ability of our subspace technique to handle editing on boundary, and the flow interacts correctly with the barrier. To handle complex boundary editing, the subspace lacks the degree of freedom, but it is still convenient to recover the flow in its original space and apply the changes, as shown in Figure 1(e).

Figure 4 presents the results on a captured 'bloom' flow, which involves dye being poured into still water. The flow is rich of turbulence in the top half of the domain, but becomes stable in the bottom half since the buoyancy and the diffusion effect of the surrounding fluid dampen those details. The re-simulation result is exhibited in Figure 4(b), which is basically the same as the original flow without noticeable distinction. We edit this captured flow by doubling its viscosity and adding an extra downward

Table 3: *Viscosity estimation.*

|	Ejection	Bloom	Smoke
Real visocity	10.000	–	–
Estimated visocity #1	10.200	3.674	1.927
Estimated visocity #2	10.210	3.665	1.913
Estimated visocity #3	10.415	3.678	1.901
Estimated visocity #4	10.308	3.641	1.927
Estimated visocity #5	10.231	3.674	1.912

impulse in subspace, and the results are shown in the bottom row. In Figure 4(c), the surrounding fluid becomes more viscous; consequently, the flow loses its turbulence quickly as soon as it enters the domain. In contrast, the flow in Figure 4(d) turns more energetic and presents more details throughout the entire domain.

Figure 5 shows the results on a captured 'smoke' flow, which has abundant high-frequency energy caused by buoyancy. The subspace re-simulation result demonstrated in Figure 5(b) is in good correspondence with the existing flow. To incorporate more complex methods of flow editing, we re-simulate the flow in its original space, and further apply editing by adding vorticity confinement [FSJ01] with $\epsilon = 0.2$ to produce more turbulent results, as seen in Figure 5(d). To stress the ability of EMD, we make comparisons of our proposed EMD-based technique against the traditional PCA-based model reduction using this dataset, so the entire pipeline is reproduced with the EMD bases substituted by PCA bases, which have the same number of columns as our EMD bases, and the result is presented in Figure 5(c). In this case, PCA bases fail to correctly capture the flow due to the effect of inter-frequency energy mixing. Large amount of kinetic energy is introduced in high-frequency modes, leading the energy transferred to low frequencies through PCA bases, which are rather meaningless in frequency domain. On the contrary, EMD bases are constructed by the extracted intrinsic features within each frequency bands, hence are powerful handling this situation and avoiding inter-frequency energy mixing, since information conveyed in different frequency bands are extracted independently in EMD bases. A more obvious view is in Figure 6, where a single frame of the original flow, the modelling using EMD bases and the modelling using PCA bases are put together. It is obvious that the result from PCA bases has larger velocity than the original flow and the result from EMD bases. The frequency spectra are provided in Figure 7, where the energy of the modelling result with PCA bases is higher than normal in all frequency bands, while the result from EMD bases is closer to the existing flow through the frequencies. Nevertheless, the PCA method generally has less distinction between the subspace simulation and the original flow since it guarantees minimizing the error, especially when fewer bases are used [TLP06, KD13]. Hence, our EMD-based method can be treated as a common technique regardless of the energy distribution, and becomes more adequate for inverse modelling if the energy is unevenly distributed through the frequency domain and the number of bases is large enough to neglect the error by model reduction, which is not a critical requirement considering the efficiency of subspace calculation. To exhibit the style controllable capability of the EMD bases, we conduct the re-simulation experiments on EMD bases with the high-frequency vectors and middle-frequency vectors removed, respectively, and the results are shown in Figures 5(e) and (f). The bases without high-frequency vectors produce much smoother results, and the bases without
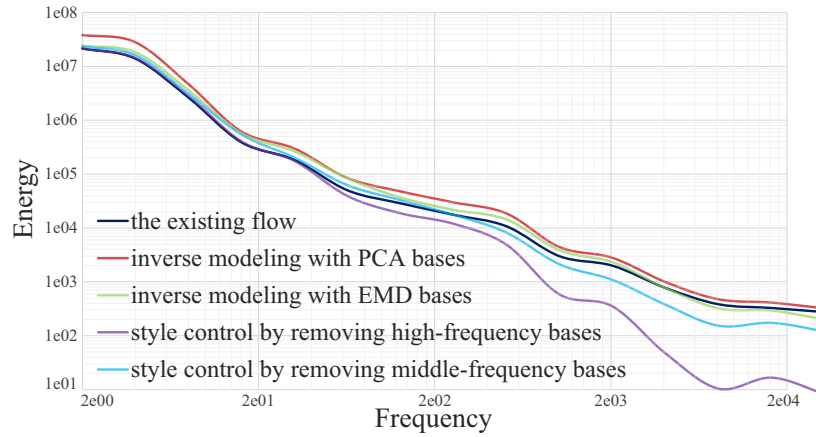
Figure 7: *Frequency spectra of the resulting sequences on the 'smoke' dataset. The inter-frequency energy mixing effect makes the energy of the modelling with PCA bases higher than the existing flow in all frequency bands, while our EMD bases can better capture the intrinsic features in each frequency. The style control method also effectively influences the spectra as expected.*

(a)   (b)
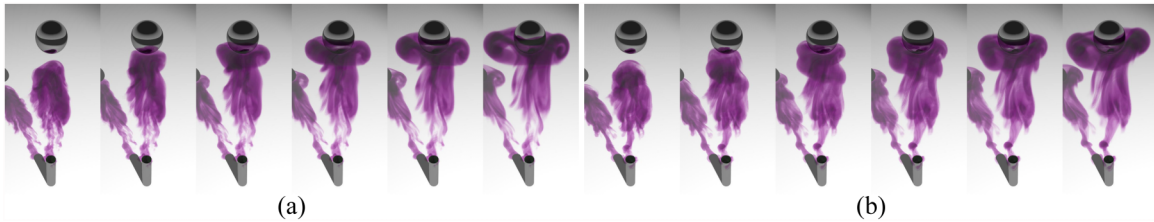
Figure 8: *Comparison of the boundary modification on the captured 'smoke' dataset. (a) Result of [GIT*14]. (b) Our result.*

middle-frequency vectors generate less energetic results, where the dominant trend of the flow is preserved. Their frequency spectra are available in Figure 7. Note that the frequency in EMD is not pre-defined but data-dependent and some frames' high-frequency bases may lie in the low-frequency domain in other frames, so the energy spectrum for 'style control by removing middle-frequency bases' does not represent any lower in the middle frequencies than the other plots. We also provide the results by changing the parameter q of Equation (3) in Figure 9. The numbers of bases m in these experiments are set as $25 \times q$, and 25 vectors are selected from each frequency band to form the bases. When increasing q (essentially this is introducing more high-frequency bases), the shape of the flow holds while the high-frequency details become a little bit more prominent, and all these results faithfully reproduce the captured flow.

Our method can be seen as a follow-up of [GIT*14], but we still provide a comparison with their method in the 'domain change' application using the captured 'smoke' dataset, as demonstrated in Figure 8, where a metal ball is placed in the domain blocking the smoke from rising. Noticeable difference can be found in the results although the same settings are used to implement both methods. The main cause is that our method achieves physical simulation and editing with the estimated flow parameters, while theirs can only convey the density passively. Their method advects the flow up to a desired time frame, after which it switches to a dynamic setting considering the new boundary, which is non-physical to be rigorous.

(a)   (b)

(c)   (d)
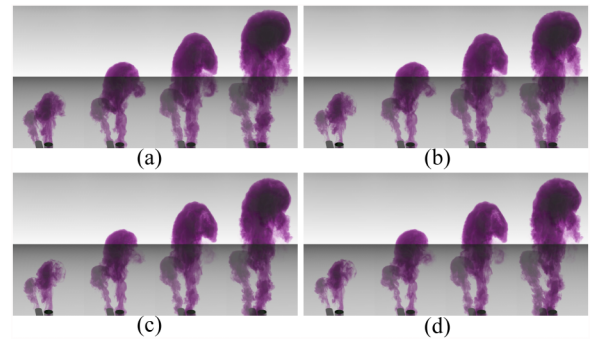
Figure 9: *Re-simulation results on the captured 'smoke' dataset by changing q in Equation 3. (a)–(d) Results using $q = 1, 2, 3, 4$, respectively.*

All other stylistic modifications that [GIT*14] brings about, such as resolution enhancement and guided simulation, even though beyond the topic of this paper, can be accomplished with our method in a more rigorous manner that strictly follows the governing equations.

6. Conclusion and Discussion

Our main goal in this paper is to make new attempts to remove the perceived barrier between fluid capturing and pure fluid simulation

in graphics applications via a powerful analytical tool. This paper has detailed a data-driven approach for fluid analysis based on EMD and modal reduction, capable of inversely modelling existing gas flows to discover the simulation parameters and editing them towards desired behaviors, hence uniting fluid capturing and fluid simulation through tight and seamless coupling. Our extensive experiments have illustrated that our extended EMD for fluid flows, accelerated by CUDA, is much more adequate for fluid decomposition compared with the previous space-filling curve approach. Making use of this EMD algorithm, we have decomposed the 3D gas flows into frequency components representing the intrinsic features, which are also candidate basis vectors for modal reduction. In subspace, we have devised a functional optimization method to isolate the effect of viscosity and external forces, aiming at performing parameters estimation and modifying the gas flow by reusing the existing characteristics extracted from the original flow. We have demonstrated that our EMD bases are superior to the traditional PCA bases in avoiding inter-frequency energy mixing, as well as bringing the flexibility of style control on the modelled flow. With the inverse-modelling technique, physical editing on the existing flows is achievable. Finally, we have integrated our technique with a state-of-the-art fluid capturing method, building a pipeline to bond the fluid simulation in graphics with real-world flows.

Our present work still has several limitations. The re-simulation scheme and flow editing approach, which depend on modal reduction, is unable to handle a new domain with an enlarged boundary yet, since the data we can reuse are confined within the domain of the existing flow. To enlarge the simulated domain, extra efforts must be made to modify the subspace and provide extra information. The proposed EMD algorithm, although proved suitable for gas flows, guarantees no temporal coherence either. A possible solution for remedy is to use some additional temporal constraints during the EMD process simultaneously considering the results from multiple adjacent frames. The fitting of TPS needs to solve a dense matrix. To accelerate the solution, we could use biharmonic B-spline [FW12] instead of TPS, whose basis functions are localized.

The approach we have proposed indeed offers the possibilities to bridge the gap between fluid capturing and simulation, and we believe more applications beyond what we have discussed in this paper shall be explored in order to do a full justice for this method. Future topics may include investigating more advanced fluid-capture systems to construct larger database from various categories of fluids, which are readily available towards flow synthesizing. In addition, we are considering possible investigation of guided simulation in an expanded domain to handle more complex scenarios with proper boundary conditions.

Acknowledgements

References

[AIH*08] ATCHESON B., IHRKE I., HEIDRICH W., TEVS A., BRADLEY D., MAGNOR M., SEIDEL H.-P.: Time-resolved 3d capture of non-stationary gas flows. *ACM Transactions on Graphics 27*, 5 (2008), 132.

[BJ05] BARBIČ J., JAMES D. L.: Real-time subspace integration for st.venant-kirchhoff deformable models. *ACM Transactions on Graphics 24*, 3 (2005), 982–990.

[DLJY11] DING Y., LI F., JI Y., YU J.: Dynamic fluid surface acquisition using a camera array. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)* (2011), IEEE, pp. 2478–2485.

[DMP05] DAMERVAL C., MEIGNEN S., PERRIER V.: A fast algorithm for bidimensional emd. *IEEE Signal Processing Letters 12*, 10 (2005), 701–704.

[DWLF12] DE WITT T., LESSIG C., FIUME E.: Fluid simulation using laplacian eigenfunctions. *ACM Transactions on Graphics 31*, 1 (2012), 571–583.

[Ebe02] EBERLY D.: *Thin Plate Splines*. Geometric Tools Inc., 2002.

[ESWVO06] ELSINGA G. E., SCARANO F., WIENEKE B., VAN OUDHEUSDEN B.: Tomographic particle image velocimetry. *Experiments in Fluids 41*, 6 (2006), 933–947.

[FL04] FATTAL R., LISCHINSKI D.: Target-driven smoke animation. *ACM Transactions on Graphics 23*, 3 (2004), 441–448.

[FM96] FOSTER N., METAXAS D.: Realistic animation of liquids. *Graphical Models and Image Processing 58*, 5 (1996), 471–483.

[FM97] FOSTER N., METAXAS D.: Controlling fluid animation. In *Proceedings of the 1997 Conference on Computer Graphics International* (1997), IEEE Computer Society, pp. 178–188.

[FSJ01] FEDKIW R., STAM J., JENSEN H. W.: Visual simulation of smoke. In *SIGGRAPH '01: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (2001), ACM, pp. 15–22.

[FW12] FENG P., WARREN J.: Discrete bi-Laplacians and biharmonic B-splines. *ACM Transactions on Graphics 31*, 4 (2012), 115:1–115:11.

[GIT*14] GREGSON J., IHRKE I., THUEREY N., HEIDRICH W.: From capture to simulation-connecting forward and inverse problems in fluids. *ACM Transactions on Graphics 33* (2014), 70–79.

[GKHH12] GREGSON J., KRIMERMAN M., HULLIN M. B., HEIDRICH W.: Stochastic tomography and its applications in 3d imaging of mixing fluids. *ACM Transactions on Graphics 31*, 4 (2012), 52:1–52:10.

[GLRH13] GAO Y., LI C.-F., REN B., HU S.-M.: View-dependent multiscale fluid simulation. *IEEE Transactions on Visualization and Computer Graphics 19*, 2 (2013), 178–188.

[HS05] HUANG N. E., SHEN S. S.: *Hilbert-Huang Transform and Its Applications*, vol. 5. World Scientific, Singapore, 2005.

[HSL*98] HUANG N. E., SHEN Z., LONG S. R., WU M. C., SHIH H. H., ZHENG Q., YEN N.-C., TUNG C. C., LIU H. H.: The empirical mode decomposition and the hilbert spectrum for nonlinear and nonstationary time series analysis. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences 454*, 1971 (1998), 903–995.

[JF03] JAMES D. L., FATAHALIAN K.: Precomputing interactive dynamic deformable scenes. *ACM Transactions on Graphics 22*, 3 (2003), 879–887.

[KD13] KIM T., DELANEY J.: Subspace fluid re-simulation. *ACM Transactions on Graphics 32*, 4 (2013), 96.

[KLLR05] KIM B., LIU Y., LLAMAS I., ROSSIGNAC J.: Flowfixer: Using bfecc for fluid simulation. In *NPH'05L: Proceedings of the 1st Eurographics Conference on Natural Phenomena* (2005), Eurographics Association, pp. 51–56.

[Lin09] LINDERHED A.: Image empirical mode decomposition: A new tool for image processing. *Advances in Adaptive Data Analysis 1*, 2 (2009), 265–294.

[Lin11] LINDERHED A.: Variable sampling of the empirical mode decomposition of two-dimensional signals. *International Journal of Wavelets, Multiresolution and Information Processing 3*, 3 (2011), 435–452.

[LPS*13] LI C., PICKUP D., SAUNDERS T., COSKER D., MARSHALL D., HALL P., WILLIS P.: Water surface modeling from a single viewpoint video. *IEEE Transactions on Visualization and Computer Graphics 19*, 7 (2013), 1242–1251.

[LS08] LIU T., SHEN L.: Fluid flow and optical flow. *Journal of Fluid Mechanics 614* (2008), 253–291.

[MTPS04] MCNAMARA A., TREUILLE A., POPOVIĆ Z., STAM J.: Fluid control using the adjoint method. *ACM Transactions on Graphics 23*, 3 (2004), 449–456.

[NB11] NIELSEN M. B., BRIDSON R.: Guide shapes for high resolution naturalistic liquid simulation. *ACM Transactions on Graphics 30*, 4 (2011), 83.

[NCZ*09] NIELSEN M. B., CHRISTENSEN B. B., ZAFAR N. B., ROBLE D., MUSETH K.: Guiding of smoke animations through variational coupling of simulations at different resolutions. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2009), ACM, pp. 217–226.

[REN*04] RASMUSSEN N., ENRIGHT D., NGUYEN D., MARINO S., SUMNER N., GEIGER W., HOON S., FEDKIW R.: Directable photorealistic liquids. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2004), Eurographics Association, pp. 193–202.

[RLL*13] REN B., LI C.-F., LIN M. C., KIM T., HU S.-M.: Flow field modulation. *IEEE Transactions on Visualization and Computer Graphics 19*, 10 (2013), 1708–1719.

[RM10] REHMAN N., MANDIC D.: Multivariate empirical mode decomposition. *Proceedings of the Royal Society A Mathematical Physical and Engineering Sciences 466*, 2117 (2010), 1291–1302.

[SFK*08] SELLE A., FEDKIW R., KIM B., LIU Y., ROSSIGNAC J.: An unconditionally stable maccormack method. *Journal of Scientific Computing 35* (2008), 350–371.

[Sta99] STAM J.: Stable fluids. In *SIGGRAPH '99: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques* (1999), ACM Press/Addison-Wesley Publishing Co., pp. 121–128.

[SY05] SHI L., YU Y.: Taming liquids for rapidly changing targets. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2005), ACM, pp. 229–236.

[TKPR09] THÜREY N., KEISER R., PAULY M., RÜDE U.: Detail-preserving fluid control. *Graphical Models 71*, 6 (2009), 221–228.

[TLP06] TREUILLE A., LEWIS A., POPOVIĆ Z.: Model reduction for real-time fluids. *ACM Transactions on Graphics 25*, 3 (2006), 826–834.

[TMPS03] TREUILLE A., MCNAMARA A., POPOVIĆ Z., STAM J.: Keyframe control of smoke simulations. *ACM Transactions on Graphics 22*, 3 (2003), 716–723.

[WHC11] WU Z., HUANG N. E., CHEN X.: The multi-dimensional ensemble empirical mode decomposition method. *Advances in Adaptive Data Analysis 1*, 3 (2011), 339–372.

[WHZQ15] WANG X., HU J., ZHANG D., QIN H.: Efficient emd and Hilbert spectra computation for 3d geometry processing and analysis via space-filling curve. *The Visual Computer 31*, 6 (2015), 1135–1141.

[WLZ*09] WANG H., LIAO M., ZHANG Q., YANG R., TURK G.: Physically guided liquid surface modeling from videos. *ACM Transactions on Graphics (TOG) 28*, 3 (2009), 90.

[WST09] WICKE M., STANTON M., TREUILLE A.: Modular bases for fluid dynamics. *ACM Transactions on Graphics 28*, 3 (2009), 341–352.

[XLLR06] XU Y., LIU B., LIU J., RIEMENSCHNEIDER S.: Two-dimensional empirical mode decomposition by finite elements. *Proceedings of the Royal Society A Mathematical Physical and Engineering Sciences 462*, 2074 (2006), 3081–3096.

[YCZ11] YUAN Z., CHEN F., ZHAO Y.: Pattern-guided smoke animation with lagrangian coherent structure. *ACM Transactions on Graphics 30*, 6 (2011), 136.

[YJLY12] Ye J., Ji Y., Li F., Yu J.: Angular domain reconstruction of dynamic 3d fluid surfaces. In *Proceedings Of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2012), IEEE, pp. 310–317.

[YKH*09] Yoon J.-C., Kam H. R., Hong J.-M., Kang S. J., Kim C.-H.: Procedural synthesis using vortex particle method for fluid simulation. *Computer Graphics Forum 28*, ffi7 (2009), 1853–1859.

[ZB05] Zhu Y., Bridson R.: Animating sand as a fluid. *ACM Transactions on Graphics 24* (2005), 965–972.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's web site:

Video SI