

RESEARCH ARTICLE

Novel fluid detail enhancement based on multi-layer depth regression analysis and FLIP fluid simulation

Yuxing Qiu¹ | Lipeng Yang¹ | Shuai Li¹ | Qing Xia¹ | Hong Qin² | Aimin Hao¹

¹State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, Beijing, China

²Department of Computer Science, Stony Brook University, New York, USA

Correspondence

Shuai Li, Aimin Hao, State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, Beijing, China.

Email: lishuai@buaa.edu.cn, ham_buaa@163.com

Funding information

National Natural Science Foundation of China, Grant/Award Number: 61190120, 61190125, 61190124, 61300067, 61672077, 6167214, 61602341 and 61532002

Abstract

In this paper, we propose a novel integrated method for effective modeling and realistic enhancement of scale-sensitive fluid simulation details. The core of our method is the organic of multi-layer depth image regression analysis and fluid implicit particle fluid simulation of which the regression analysis induces the criterion where the fluid details should be produced. First, we capture the depth buffer of the fluid surface dynamically from the top of scene. Second, we employ depth peeling technique to decompose the target fluid volume into multiple depth layers and conduct time-space analysis over surface layers. Third, we propose a logistic regression-based model to rigorously pinpoint the complex interacting regions, wherein multiple detail-relevant factors are taken into account based on the captured multiple depth layers. Finally, details are enhanced by animating extra diffuse materials and augmenting the air-fluid mixing phenomenon. It is evident that, with depth peeling technology, we can afford rigorous analysis not only across surface layers at different fluid depth but along the depth direction as well. After integrating the analysis results from these two sources, we are capable of performing detail enhancement both on the fluid surface and inside the fluid to obtain a great visual effect, even when large occlusion exists. Directly benefiting from the flexibility of image-space-dominant processing, our unified framework can be entirely implemented on graphics processing units and thus achieves interactive performance. For various fluid phenomena with different diffuse materials (e.g., spray, foam, and bubble), comprehensive experiments and evaluations have demonstrated its superiority in high-fidelity fluid detail enhancement and its interaction with surrounding environment.

KEYWORDS

depth peeling, FLIP, fluid detail enhancement, GPU, image space method, time-space analysis model

1 | INTRODUCTION AND MOTIVATION

During the past two decades, physically based fluid simulation has become a rapidly evolving research topic in many virtual reality (VR) applications, wherein various methods have been developed in response to different motivations, including single-phase fluid simulation, multiple fluid interaction, two-way fluid-solid coupling, and so on. With the growing demand for physical realism and interactive cross-scale detail enhancement, such as dealing with small-scale splash, spray, and other diffuse materials^{1,2} for large-scale fluid

simulations, the systematic and comprehensive studies of fluid detail simulation are beginning to regain popularity with great and revived momentum.

Although accuracy is the key concern in computational fluid dynamics, visual fidelity and efficiency are more important in graphics, animation, and many other VR applications. For the visual fidelity improvement, the state-of-the-art methods oftentimes resort to various space-grid refinement or denser particle sampling schemes in physical simulation. In addition, from the perspective of pure simulation, other key ideas to enhance details include handling different materials

separately (such as air and liquid) and simulating the crucial interaction^{3–5} in a physically plausible fashion to increase visual fidelity. Considering the unavoidable improvement of computational expenses, adaptive sampling methods are more favorable, such as octree-grid,⁶ adaptive SPH,⁷ and so on. Meanwhile, parallel methods on multi-core CPUs or graphics processing unit (GPU) are also developed in order to reduce the ever-increasing computational cost^{8,9} while still improving visual quality.

Despite the recent success of diffuse material simulation, certain difficulties still prevail and need to be addressed for high-fidelity region detection for fluid details in a more efficient way. First, most of the detail detection criteria depend on the correct identification of accurate fluid surface, which is too time-consuming to be precisely detected in any physics-driven real-time applications. Meanwhile, spatially local analysis criteria, such as curvature and velocity variance, are usually combined manually with many parameters, which are hard to be fine-tuned for different phenomena. Second, most of the criteria are hard to pinpoint scale-aware details, wherein the details are mostly determined by probability-biased parameters and their ad hoc combination. Third, the state-of-the-art analysis methods tend to ignore the informative time-domain knowledge, which is crucial for fluid detail detection and prediction.

Our rationale is that pure physics-based simulation may fall short in fluid detail enhancement during interaction, unless significantly increased computational expenses could be afforded. Hybrid strategies, which integrate simulation and analysis, must be explored towards high-fidelity fluid interaction while not sacrificing computational efficiency drastically. This paper attempts to offer a viable solution to tackle the aforementioned technical challenges. We articulate a novel integrated analysis-centric approach by rigorously analyzing the acquired data. In particular, we propose to utilize multi-level depth buffers captured during simulation as an indicator of the dynamic fluid status to analyze the region-variant necessity of detail enhancement. By analyzing the time–space features in image space, we expect to correctly identify the fluid-detail regions and then simulate the diffuse materials to augment the air–fluid interacting phenomena. Through the integration of simulation and analysis, we are hoping to achieve the competing goals of high-fidelity fluid detail enhancement and computational efficiency for VR applications. The salient contributions of this paper can be summarized as follows:

- We propose a time–space analysis model based on logistic regression to integrate the geometric, physical, and temporal factors for fluid details detection while avoiding complex parameter tuning.
- We propose to bridge simulation and analysis via the utility of dynamically captured depth buffer data, which are both low cost and versatile for quality-efficiency trade-off.

- We design a multi-layer surface information extraction method based on depth peeling to enhance the fluid simulation details over (multi-layered) surface regions.
- We develop a seeding method for diffuse particle generation, which only relies on the depth buffers and analysis results, making it applicable for all particle-based simulation methods.
- Without applying depth peeling technique, the involved data-driven analysis method and diffuse material simulation can be integrated into any grid-based method with very little extra workload with some specific surface extraction algorithm for grid-based fluid.
- We design specific GPU-based algorithms to implement our entire simulation framework in parallel.

2 | RELATED WORK

In recent years, there have been many methods developed to simulate fluids, including smoothed particle hydrodynamics (SPH) method,^{10–13} fluid implicit particle (FLIP) method,¹⁴ level set methods,^{15,16} and so on. Our approach is based on FLIP method, which combines particle-based representation with grid-based solver. FLIP method was introduced in computer graphics by Zhu and Bridson,¹⁴ and then was extended to many fields, such as splashing water simulation,¹⁷ preservation of fluid sheet,¹⁸ fluid–solid coupling applications,¹⁹ combining with particle method,²⁰ multi-scale droplet/spray simulation,²¹ and so on. Lately, Ando et al.,¹⁸ Um et al.,²² and Cornelis et al.²⁰ respectively proposes methods to improve the particle distribution of FLIP method.

Closely relevant to the central theme of this paper, we now briefly review previous works in diffuse material simulation, data-specific analysis, and depth peeling technology separately.

2.1 | Diffuse material simulation

Diffuse material simulation is crucial in fluid detail simulation and enhancement, especially for the large-scale fluid phenomenon having complex interaction with air.

Detecting diffuse materials is straightforward for grid-based methods, because the surface of a liquid can be easily tracked. Existing works essentially measure geometric information to generate splash, foam, and bubble using curvatures,⁵ markers escaped from surface,^{2,3,23} and so on. To exploit the ignored velocity information, Mihalef et al.⁴ present a Weber number threshold-based filtering method to avoid the detail loss of the marker level set method by introducing a physical factor into the criteria. Meanwhile, some other works regard air as a separate phase fluid and simulate the interaction between liquid and air directly, including volume-of-fluid,²⁴ regional level set,²⁵ multiFLIP,²⁶ two-continua approach,²⁷ and so on. Patkar et al.²⁸ propose a hybrid method for bubbles, using level set to simulate large bubbles while employing particle method to represent small bubbles.

As for particle-based methods, analyzing the diffuse materials is not trivial, because the precise fluid surface is not easy to be detected and extracted. To simulate the water–air interaction, Müller et al., Solenthaler et al., and Threy et al.^{29–31} achieve detail enhancement based on SPH model. Although these methods can obtain improved visual results, however, they still have problems when certain qualifications required in their papers cannot be guaranteed. To ameliorate, Ihmsen et al.³² simulate the bubbles and liquids separately by differentiating the regions where air is likely to be trapped with a velocity-based heuristic criterion. Losasso et al.³³ propose a two-way coupled simulation framework that uses the particle level set method to model dense liquid volumes and an SPH model to simulate diffuse regions such as sprays. Bagar et al.³⁴ detect foam particles from regular liquid particles based on the velocities of particles, which involve no foam simulation. To unify spray, foam, and air bubble, Ihmsen et al.¹ propose a post-processing model to generate and advect the diffuse materials for SPH fluids, wherein the diffuse regions are detected by analyzing the curvatures and velocities of SPH particles. Changbo et al.³⁵ propose a similar method for grid-based simulator. Although having avoided expensive inter-particle computation, such schemes heavily rely on the proper detection of surface particles and the curvature calculation over surface particles.

2.2 | Data-specific analysis methods

Image-specific analysis methods are commonly used to enhance image details, which in our method will be employed to detect the complex interacting regions in depth images. Most recently, some researchers produce plausible results in image smoothing and detail enhancement. For instance, in order to achieve data-specific wavelet, Hammond et al.³⁶ propose a method for the construction of wavelet transforms, according to the functions defined on the vertices of an arbitrary weighted graph. Narang et al.³⁷ construct the data-specific wavelet by specifically designing wavelet filters based on the spectral decomposition of the graph, and then necessary and sufficient conditions for a two-channel graph filter bank on bipartite graphs are stated, which can achieve aliasing-cancellation and perfect reconstruction. In addition, Xu et al.³⁸ design an optimization framework via L_0 gradient minimization, wherein they can sharpen major edges by increasing the steepness of transition while eliminating a manageable degree of low-amplitude structures at the same time. However, the computational costs of these methods are expensive, and they are not suitable for interactive VR applications. Guided image filtering is derived from a local linear model according to the texture pattern of guidance image. Compared with the widely used bilateral filter,³⁹ guided filter naturally is a generic linear-time algorithm, which has a better trade-off between the accuracy and efficiency. Thus, considering both the structure-preserving effect and the low computational cost demands, we introduce the guided image filter⁴⁰ into our method.

2.3 | Depth peeling techniques

According to Everitt,⁴¹ depth peeling is a fragment-level depth-sorting technique described by Mammen using Virtual Pixel Maps⁴² and by Diefenbach using a dual depth buffer.⁴³ This technique is widely used in transparent surfaces rendering domain, and each unique depth in the scene is extracted into layers and composite in depth-sorted order, to produce a correctly blended final image. With depth peeling, the rendering order-dependent problem of non-refractive transparent surfaces can be solved.⁴¹ There exists some applications and improvements of depth peeling techniques. For instance, Bernardon et al.⁴⁴ propose a new hardware-based ray casting algorithm for the depth peeling of unstructured meshes consisting of tetrahedral cells. In this way, they can capture the rays when they re-enter the mesh, without introducing imaginary cells to fill the intractable space caused by mesh non-convexities. Depth peeling techniques have also been extended to texture-based volumetric iso-surfaces rendering,⁴⁵ in which the authors originate the term *volumetric depth peeling*. In this way, the full range of volumetric effects can be achieved. Inspired by this, David et al.⁴⁶ utilize volumetric depth peeling for medical image visualization, which is flexible enough to handle multiple region occlusion and object's self-occlusion and requires no pre-segmentation over the dataset. Based on depth peeling method, the joint surface occlusion problem has been solved, which has been successfully used in urology and visual arthroscopic studies.⁴⁶ However, the classical depth peeling algorithm has performance bottleneck for large and complex scenes; to improve it, Fang et al.⁴⁷ exploit multiple render targets as bucket array in pixel-wise way. With the use of bucket sort and the efficient schemes to reduce collisions in the same bucket, Fang et al.⁴⁷ achieve 32 times speedup in depth peeling while guaranteeing faithful visual results, especially for complex scenes.

To integrate time-domain information into our method, we refer to time–space analysis, which is commonly utilized in video object tracking. The relevant techniques include Bayesian model,^{48,49} logistic regression,⁵⁰ Gaussian mixture model,⁵¹ and so on. Our detail region detection is similar with video tracking in the aspects that both time and space information are crucial, but our aim is to guarantee the continuous evolution of enhanced regions. It is appropriate to use the logistic regression to combine multiple time–space factors in order to obtain the probability of each pixel in depth buffer, indicating whether it should be enhanced or not.

3 | ALGORITHMIC OVERVIEW

Figure 1 illustrates the entire algorithmic flow of our method. Each simulation cycle starts with regular FLIP; by way of analyzing the captured result, we enhance fluid details via animating diffuse materials. We outline the algorithm as follows:

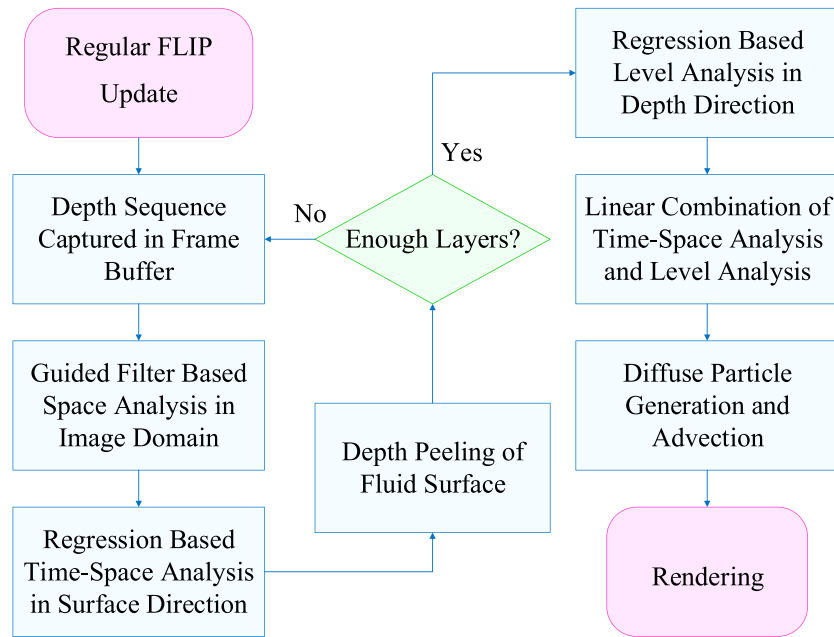


FIGURE 1 The pipeline of each simulation cycle. FLIP, fluid implicit particle

FLIP update. Compute the density, pressure, and force for each fluid particle, and then update its velocity and position.

Depth capture. Capture and store the depth buffer dynamically from the top view point with an *orthogonal virtual camera*.

Depth peeling. Peel the multi-layer fluid surfaces according to the obtained depth buffer and capture the depth of each layer for the subsequent analysis.

Space analysis. Analyze the 2D depth buffer spatially with data-specific guided filter to obtain the geometric features within each layer of the near-surface.

Time–space analysis. Integrate temporal information and spatial analysis result into a logistic regression model to obtain the criteria for detail enhancement along surface direction.

Level analysis. Conduct similar regression analysis over multi-layered depth maps to induce the criteria along depth direction.

Final criteria. Integrate the regression results from the aforementioned two sources in a weighted way in order to obtain the final scale-sensitive criteria for detail enhancement over multi-layered surface regions.

Diffuse particle generation and simulation. Generate diffuse particles based on the analysis results, then advect them according to the material-specific rules.

Rendering. Render the scene with POV-Ray software.

essentially based on Navier–Stokes equations (N-S equations) that conserve both mass and momentum:

$$\frac{\partial \rho}{\partial t} + \rho \nabla \cdot \mathbf{u} = 0, \quad (1)$$

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) = -\nabla p + \mu \nabla^2 \mathbf{u} + \mathbf{f}. \quad (2)$$

Here, ρ is the density, \mathbf{u} is the velocity, p is the pressure, μ is the kinematic viscosity, and \mathbf{f} is the external force. ∇ means the differential operator of vector.

In FLIP model, fluid is discretized as particles, and traditional Eulerian method is employed to solve the N-S equations instead. Unlike particle-in-cell (PIC) method, the velocity changes on grid; rather than computing the velocity on grid directly from the aforementioned equations, the grid velocity is interpolated from the surrounding particles. As a result, the numerical dissipation problem is avoided, making FLIP more suitable for violent fluid simulation.

To alleviate the noisy-like behaviors of FLIP, similar to Zhu et al.,¹⁴ we linearly blend the PIC and FLIP velocities via

$$\mathbf{v} = a\mathbf{v}_{FLIP} + (1 - a)\mathbf{v}_{PIC}, \quad (3)$$

where the blending factor a is set to be 0.95 in all of our experiments.

4 | TIME–SPACE ANALYSIS

4.1 | Depth image sequence capture

As shown in Figure 2, the depth buffer is captured at the end of each regular FLIP update cycle. Similar to Van der Laan et al.⁵² (but with a downward *orthogonal camera* at the top of the scene), we render the particles as spheres, and then

3.1 | Brief FLIP review

Because our model is based on the adaptive FLIP model,¹⁸ we now briefly review the basic idea of fluid simulation, FLIP model, and its possible improvement. Fluid dynamics are

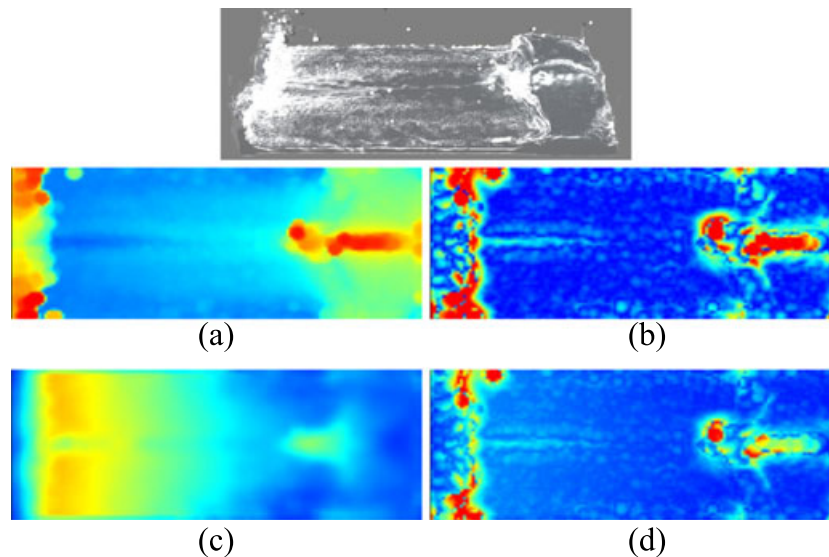


FIGURE 2 Time-space analysis results. Warm color indicates high value and vice versa. (a) The captured depth buffer, (b) the space analysis result, (c) velocity field, and (d) time-space analysis result

capture the depth buffer from frame buffer. The captured depth buffer is stored as texture first and then is mapped into GPU memory as 2D array, which is the basis of conducting analysis. The depth buffer actually forms a 2D *projected grid* in simulation domain, which enables us to extract velocity field and seed diffuse particles. Meanwhile, we also keep a record of some sequential frame buffers to analyze the temporal information involved in the dynamics of the underlying fluid surface.

The resolution of captured depth buffer is crucial to the trade-off of analysis quality and computational cost. An appropriate resolution could accommodate sharp/fine features on the fluid surface while guaranteeing efficiency. We will discuss our resolution selection strategy and demonstrate the results of different resolutions in Section 7.

4.2 | Guided filter-based space analysis

Geometric features of fluid surface is critical to determining where the details should be enhanced. By converting the fluid surface feature detection to a problem of depth image analysis, it cannot only reduce the scale of computation but also take advantages of the abundance of techniques developed in image processing. We employ the content-specific guided filter⁴⁰ to perform detail detection. Although some simple linear translation-invariant filters (such as the mean, Gaussian, Laplacian, and Sobel filters) have been widely used in image processing, they are weak in making the smoothed image preserve original structures and edges. Utilizing guidance image instead of a fixed filter, we can take into account more data-specific information to make the smoothed image preserve structures and edges. In our application, we tend to extract the potential detail features from the differences of smoothed multi-level depth buffers and original ones. According to this, the edges of depth buffer need to be

consistently persisted. Some sharp edges, which might represent the overlaps between different fluid flows and fluid pools, should not be considered as regions needing significant enhancement. Unlike bilateral filter, guided filter has good edge-preserving properties without suffering from gradient reversal artifacts. Besides, efficiency is also an important factor for us to choose guided filter, because the filtering output of guided filter is locally a linear transform of the guidance image. In sharp contrast, other content-specific filters usually utilize some optimizations to pursue better quality, and this oftentimes comes with the price of expensive computational time. Considering our eager expectation of simplicity and efficiency, we finally choose the guided filter in our method.

The key idea of guided filter is to develop a filter kernel that can be applied onto the input image to produce a new image. In our framework, we apply its smoothing function to obtain smoothed depth buffers, whose differences could pinpoint scale-aware detail region with the high-frequency characteristics.

We assume there is a local linear transformation between guidance image I and output q in a window w_k centered at the pixel k ,

$$q_i = a_k I_i + b_k, \forall i \in w_k, \quad (4)$$

where a_k and b_k are linear coefficients. To determine a_k and b_k , we minimize the difference between q and the input image p ,

$$E(a_k, b_k) = \sum_i ((a_k I_i + b_k - p_i)^2 + \epsilon a_k^2), \quad (5)$$

where ϵ is a regularization parameter, and we define the same depth buffer as the guidance image I and the input image p . According to He et al.,⁴⁰ we set ϵ with an empirical value 0.01 in all our experiments, which can produce satisfactory results.

Then a_k, b_k can be directly solved by linear regression as

$$a_k = \frac{\frac{1}{|w|} \sum_{i \in w_k} I_i p_i - \mu_k \bar{p}_k}{\sigma_k^2 + \epsilon}, \quad (6)$$

$$b_k = \bar{p}_k - a_k \mu_k, \quad (7)$$

where μ_k and σ_k^2 are the mean and variance of I in w_k , $|w|$ is the number of pixels in window w_k , and \bar{p}_k is the mean of p in w_k .

After that, q can be computed using Equation (4). Because each pixel i belongs to all the windows centered at the surrounding pixels k , we average a_k, b_k to guarantee consistency.

$$q_i = \frac{1}{|w|} \sum_{k: i \in w_k} (a_k I_i + b_k) = \bar{a}_k I_i + \bar{b}_k, \quad (8)$$

where \bar{a}_k and \bar{b}_k are the mean of a_k and b_k in window w_i . We compute the differences of smoothed multi-level depth buffers to obtain potential detail features, which reside in the high-frequency regions of the fluid surface (refer to Figure 2(b)). The window size w_k allows us to identify features of different scales (refer to Figure 6 and Video S1).

4.3 | Regression-based time–space analysis

Besides the aforementioned geometric features, we further exploit velocity field as a physical criterion to determine the detail region. Specifically, we only need the velocities at the projected grid nodes, which can be extracted by finding the fluid particles around each node and computing the weighted sum of the fluid velocities. Figure 2(c) shows the examples of the extracted velocity field.

Once we obtain the geometric and velocity information of the depth image, we need to combine them together to form a criterion, which indicates the regions where diffuse materials are prone to generate. To complete this task, we explore the logistic regression model, which is a classic machine-learning model to estimate the probability of binary response based on one or more predictor variables. In our method, logistic regression model takes the geometric and physical information as inputs, and generates a probability for each pixel in depth image, indicating its probability of generating diffuse material. We define a logistic function to compute the probability that a depth buffer pixel should fall in a detail region: $h_\theta(x) = 1/(1 + \exp(-\sum_{j=1}^n \theta_j x_j))$, where θ_j is the corresponding coefficient to be learned and x_j represents all the factors being taken into consideration. In regression-based time–space analysis of certain surface layer, x_1, x_2 , and x_3 are separately set to be the difference of the depth images, the curvature of the depth images, and the difference of the velocity fields. As for the regression-based level analysis, which will be introduced in the following sections, we set x_1 and x_2 as the differences and curvatures of the cross-layer depth images, and set x_3 as the velocity gradient in depth direction.

The same as the classic logistic regression model, we use maximum likelihood estimation to learn the coefficients θ_j , so the cost function in our model is as follows:

$$E_s = -\frac{1}{m} \sum_{i=1}^m [y^{(i)} \log h_\theta(x^{(i)}) + (1 - y^{(i)}) \log(1 - h_\theta(x^{(i)}))] + r \|\theta\|^2. \quad (9)$$

Here, m denotes the number of samples, which is the number of pixels in the image, $y^{(i)}$ denotes the probability of i -pixel position that should be enhanced, and r is a user-defined coefficient to prevent θ from being too large. In our experiments, we found that 0.01 is a good choice for r , because the decrease of r does not obviously accelerate the convergence speed anymore; however, increasing r can slightly contribute to the visual fidelity improvement. As a result, we finally choose to set r as 0.01 for all phenomena in this paper. By minimizing the E_s in Equation (9) using gradient descent method, we can obtain the coefficients θ_j .

However, in the temporal aspect, the features of fluid are changing continuously; we expect the adjacent detected features would not differ too much. So we introduce a regularization penalty item to measure the inconsistency between adjacent frames, $E_t = \sum_{i=1}^m (y^{(i)} - y')^2$, where y_t is the analysis result of the previous frame.

By integrating the above aspects, the proposed cost function is formulated as follows:

$$E = E_s + \mu E_t, \quad (10)$$

where μ is a non-negative trade-off parameter to adjust the weight of the logistic regression cost and inconsistency cost between adjacent frames. Considering that consistency is important for the visual fidelity of fluid simulation, we set μ with a relatively large weight as 5. In fact, it is difficult to quantitatively evaluate the visual fidelity for the absence of a specific standard. Thus, we just use the same value for all scenes to achieve the simplification, and the results are considered to be satisfactory enough. Considering y and θ are the independent variables to be learned from Equation (10), we employ the gradient descent approach to minimize the cost function, which can be formulated as follows:

$$\begin{aligned} \theta_j^{n+1} &= \theta_j^n - \alpha_1 \frac{\partial}{\partial \theta_j} E \\ &= \theta_j^n - \alpha_1 \left[\frac{1}{m} \sum_{i=1}^m (h_\theta(x^{(i)}) - y^{(i)}) x_j^{(i)} + 2r\theta_j \right], \end{aligned} \quad (11)$$

$$y^{n+1} = y^n - \alpha_2 \frac{\partial}{\partial y} E = y^n - \alpha_2 \left(-\sum_{j=1}^n \theta_j x_j + 2\mu(y^n - y') \right). \quad (12)$$

In each simulation step, y and θ are initialized with the result from the previous step, while at the first step, y is set as 0, and θ is 1. As for α_1 and α_2 , they represent the step length of the gradient descent method; we set them according to the

empirical values, and the resulted convergence performance and the time consumption are satisfied. It should be noted that these values are also kept same for all scenes (α_1 is 0.1 and α_2 is 0.01). By iterating θ and y until they converge, we obtain y to finally indicate whether a pixel of depth buffer (i.e., a node on the projected grid) is in the detail regions or not. The algorithmic details are shown in Algorithm 1 and Section 6. Figure 2(d) shows our time–space analysis results as an example.

4.4 | Depth peeling method

We have discussed how to induce the criteria for detail enhancement from depth buffer captured from a downward orthogonal camera. However, this depth buffer only captures the information of the fluid surface, which is apparently far from being useful, because the bubbles may appear underwater. In order to calculate the full-scale generation criteria for diffuse particles and equip our model with much stronger ability to handle more complex cases (e.g., plunging breaker), we employ depth peeling technology⁴¹ to acquire multi-layer depth buffers inside fluid. That is, we can get n depth layers deeper from the fluid particles beneath the surface via n -pass rendering of this scene.

In practice, after capturing the depth image of current surface, we peel the top layer particles according to this depth image (it is used to locate surface particles), so that we can expose the subsurface of the fluid. After peeling, the method detailed in Section 4.1 will be applied again on the new surface to capture the depth buffer. Then the guided filter-based space analysis (Section 4.2) and regression-based time–space analysis (Section 4.3) are employed to conduct analysis across the surface layers, so that we can also obtain the velocity field, as well as the regression analysis result of this layer. Figure 3(c–f) illustrate the time–space analysis results of the different layers obtained from depth peeling over surface regions. In this scene, two violent water streams on the left side are injected into a tank from underwater, while another two are spouted above the water on the right. We choose to use five layers to conduct analysis as an example; however, this method does not have a limitation in the number

of depth layers, which is decided by the number of passes of depth peeling process. If we only take the surface information into account, the streams inside the fluid will be ignored, and no diffuse particles will be generated for them. As shown in Figure 3(b), it looks unreal and should have been avoided. It shows that different layer holds unique physical and geometric information at different fluid depth, which makes depth peeling step both meaningful and necessary. As we can see in this figure, the differences among the velocity fields of different layers are especially remarkable, which decide the imparities of the time–space analysis results to a large extent. In Section 7, we will conduct comparisons between the detail enhancement results with and without depth peeling.

The trade-off between the computational cost and analysis quality of depth peeling is controlled by two parameters: the number of peeling passes and the range of height for one-time peeling, which is represented by a distance threshold λ from underwater particles to fluid surface. The more the peeling passes, the deeper we can reach underwater, and the more computational cost will be needed. On the other hand, the height range for one-time peeling can influence the quality of detail enhancement, because the depth and velocity field extraction results of the current surface are accurate and make the diffuse particle generation analysis more reliable only at certain positions. However, the probability of detail enhancement around the two adjacent layers is calculated via linearly interpolating the two-layer regression results, which may be error-prone. Our parameter selection strategy will be discussed in Section 7.

4.5 | Regression-based level analysis

During the several rounds of depth peeling, we can gradually extract the depth information as well as the velocity field of deeper layers under the fluid surface, which comprise a depth buffer sequence and a velocity sequence along the depth direction. For more thorough analysis, we propose regression-based level analysis, using the two sequences documented above to induce the criteria for detail enhancement along the depth direction. It should be noted that the




FIGURE 3 Time-space analysis results and level analysis results aided by depth peeling. Warm color indicates high value and vice versa. Figure 3(c–f) are analyzed over surface regions while (h–j) are processed along the depth direction. (a) The scene analyzed by the depth peeling technique, (b) scene analyzed without depth peeling, (c) captured multi-layer depth buffers, (d) space analysis results, (e) velocity fields, (f) time–space analysis results, (g) the final analysis result used for seeding diffuse particles, (h) space analysis result between neighboring layers, (i) velocity analysis result, and (j) level analysis result

analysis introduced in Section 4.3 is executed over (multi-layered) surfaces only.

By analyzing the curvatures of these sequential depth layers, we can leverage the variation rate of depth as geometrical criterion to focus on the locations with high-frequency depth changes along the depth direction. And we consider that, the faster the depth changes in vertical direction, the more likely the fluid is to interact with air. As the diffuse particles have higher probability to appear at fluid–air mixed positions, the second-order derivatives of the depth layers should be taken as one of the indicators to enhance details as Equation (13), and the computed results are shown in Figure 3(h).

$$g_i^{2d} = d_{i+1} - 2d_i + d_{i-1}, \quad (13)$$

where g_i^{2d} means the pixel-wise second-order derivative of depth layers and d_i is the depth value of the i -th layer. Considering the depth gradient ($d_i - d_{i-1}$), we can use Equation (13) to approximate the second-order derivative.

In addition, once we obtain the velocity fields of different fluid layers, we can calculate the differences of fluid velocity among adjacent layers, which serve as physical features as

$$g_i^v = \frac{v_{i+1} - v_i}{2} + \frac{v_i - v_{i-1}}{2} = \frac{v_{i+1} - v_{i-1}}{2}, \quad (14)$$

where g_i^v is the pixel-wise velocity gradient in depth field and v_i is the velocity at this position in the i -th depth layer. In common sense, the bubbles are more likely to appear around the locations where the fluid velocities in the neighboring layers are quite different. The results of velocity analysis in depth direction can be found in Figure 3(i).

Both the geometric and the physical characteristics calculated above will be used to determine the potential detail regions. Similar to the method used in Section 4.3, we employ the logistic regression model to estimate the probability of diffuse particles generation along the depth direction, with the aforementioned geometric and the physical analysis results as inputs. And the time–space analysis result along the depth direction is shown in Figure 3(j), while Figure 3(f) illustrates the time–space analysis results over surface regions only.

After level analysis, we can obtain the regression results both over surface regions and along the depth direction, and thus the final criterion for detail enhancement can be calculated by $y[i] = \omega_1 * y_1[i] + \omega_2 * y_2[i]$ for certain depth layer, wherein $y_1[i]$ is the regression result of time–space analysis on surface, $y_2[i]$ represents the result of level analysis in the depth direction, and i is the layer index. The weighted combination of $y_1[i]$ and $y_2[i]$ gives rise to the final regression results, and the weight coefficients ω_1 and ω_2 satisfy $\omega_1 + \omega_2 = 1.0$. In all experiments, ω_1 is set to be 0.7 and ω_2 is set to be 0.3. The probability of the diffuse particles generation between two adjacent depth layers is calculated via linearly interpolating the criteria over those two layers. Figure 3(g) shows the final analysis results for seeding diffuse particles at certain depth layer.

5 | DIFFUSE MATERIAL SIMULATION

After determining the detail regions by time–space analysis, we seed diffuse particles correspondingly to augment the complex air–fluid interaction, and then advect them with the flow of fluid. This simulation strategy is similar to Ihmsen et al.,¹ but we have modified the method to accommodate FLIP framework and GPU acceleration.

5.1 | Seeding of diffuse particles

To be consistent with the time–space analysis result, we seed diffuse particles directly around the nodes of grid projected by depth buffer. For each projected grid node, we calculate the number of diffuse particles to be seeded and their 3D positions in the simulation space, and then seed the diffuse particles randomly in a specific spatial range around the projected grid node. We utilize the depth image to generate diffuse particles, rather than relying on the particle distribution,¹ which reduces the computational time.

The newly generated diffuse particles are given an initial velocity v_{dp} to flow with the surrounding fluid particles. Meanwhile, we add a small random disturbance to avoid uniform movement of the diffuse particles via $v_{dp} = v_f + v_{random}$, where v_f denotes velocity of nearby fluid and v_{random} is the disturbing velocity.

Meanwhile, to efficiently model the dissolution of the diffuse materials, we initially set a life time $t_{life} = \|v_f\|$ for each diffuse particle, which is related to the velocity of surrounding fluid particles v_f . In each simulation step, we decrease t_{life} when the diffuse particle is classified as a foam particle, and if $t_{life} \leq 0$, we remove the particle. The diffuse particle classification will be described in Section 5.2. Although this lifetime model is simple to implement and can reflect the crack of bubbles and foams, it ignores the complex physical factors such as temperature, size, and material, which still may bring some artifacts in the simulation.

5.2 | Advection of diffuse particles

To advect diffuse particles in a more specific way, we classify them into spray, foam, and bubble according to their relative position to the fluid surface. Spray, foam, and bubble particles corresponding have different advection formulations, which will be introduced in the following paragraphs. In our experiments, when a diffuse particle is apart from the fluid surface and flying in the air, it is empirically classified as spray particle. And the diffuse particles floating around the surface are classified as foam particle. What is more, the diffuse particles under the fluid surface are regarded as bubble particles. The classification task can be accomplished by computing the liquid particle number in the neighboring volume. Figure 4 illustrates the criterion and the result of our classification method.

Air bubbles mean that air is trapped inside the liquid, which are mainly affected by buoyancy force and drag force from

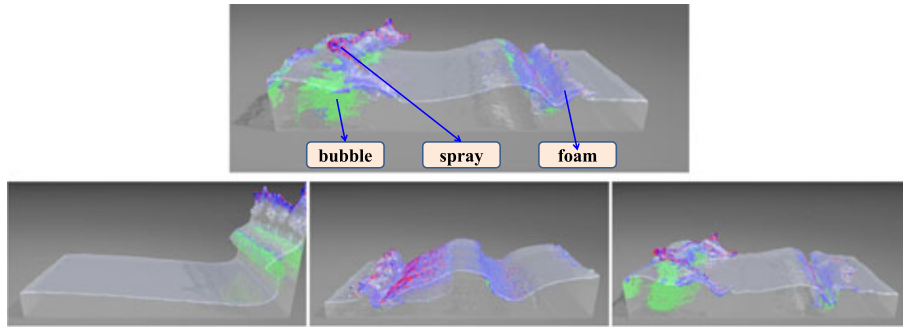


FIGURE 4 The classification of diffuse particles (red for spray particles, green for bubble particles, and blue for foam particles) using dam-breaking simulation as an example

liquid particles. The drag force is determined by the relative velocities between the bubble particles and the liquid particles, and the buoyancy force is enforced in the opposite direction of gravity acceleration \mathbf{g} .

$$\mathbf{f}_{dp} = k_d(\mathbf{v}_{dp} - \mathbf{v}_f) - k_b\mathbf{g}, \quad (15)$$

where k_d , k_b separately correspond to the coefficients of drag force and buoyancy force used to advect bubble particles. The buoyancy force on a submerged body points to the opposite direction to gravity and is equal to $\rho V\mathbf{g}$, where ρ is the density of the fluid, V measures the volume of the underwater bubbles, and \mathbf{g} means the gravity acceleration. Compared with our proposed formulation, k_b is exactly calculated by $\rho * V$. According to this, k_b should be a constant for all scenes mentioned in this paper, as we use bubble particles with the same size, and all diffuse particles are generated for water scenes. If the fluid property is changed (especially the density of fluid), the value of k_b also needs to be adjusted. And as for k_d , we take it as a constant to control drag effects. And k_d is chosen as 1 in order to make air bubbles immediately dragged into the fluid flow direction to enhance the visual fidelity according to Ihmsen et al.,¹ which are also kept the same for all scenes. The subscript dp is always associated with a variable of diffuse particle. Here, \mathbf{v}_{dp} and \mathbf{v}_f separately represent the velocity of diffuse particles and fluid particles.

Foam particles represent the foam at the surface of fluid; they are mainly affected by the drag force $k_d(\mathbf{v}_{dp} - \mathbf{v}_f)$ due to the flowing of liquid. Spray particles indicate the liquid particles that depart from the liquid volume, and they are only affected by the gravity force $m_{dp}\mathbf{g}$. Meanwhile, all the diffuse particles will be affected by a coupling force if they collide with solids in the scene. At last, the velocity and position can be updated based on Euler integration.

6 | CUDA-BASED IMPLEMENTATION

Since FLIP-based simulations have been implemented on GPU efficiently,²¹ in the interest of space, we only detail GPU implementation for the depth buffer-based analysis and diffuse particle seeding/advection, which collectively guarantee

the interactive performance of our framework. Algorithm 1 documents the pseudocode, and we detail the implementation challenges and solutions as follows.

6.1 | Guided filter-based space analysis

Guided filter method is suitable for GPU parallelization, mainly because each pixel is only related to its surrounding pixels. We invoke a CUDA kernel for each pixel, compute ω_i, μ_i according to the surrounding pixels, and then compute a_k, b_k and store them in global memory cache. We invoke another kernel to compute the output image q and compute the difference simultaneously.

6.2 | Regression-based time-space analysis

The key steps of time-space analysis are shown on lines 8–12 of Algorithm 1. We need two kernel functions for lines 10 and 12; both will invoke a kernel thread for each pixel's parallel computation. In the first kernel function, we compute $\sum_{i=1}^m (h_\theta(x^{(i)} - y^{(i)}))$ with the parallel sum reduction, and then update θ_1 and verify whether the convergence condition is satisfied. In the second kernel function, we update y_1 for each pixel in a parallel way. When the computation is completed, the analysis result y_1 is stored in the GPU array to be used in the following steps.

6.3 | Depth peeling and level analysis

We employ a CUDA kernel to compute the coordinates of fluid surface according to the depth map and invoke another CUDA kernel to compute the distances from each particle to the fluid surface, and then decide whether the particle needs to be peeled or not based on a distance threshold λ .

To conduct level analysis, we first need to calculate the second-order derivatives of depth maps and the differences of velocity field in the depth direction; therefore, we invoke two CUDA kernels at each pixel to compute these two items separately. As shown on the lines 18–22 of Algorithm 1, we employ the two kernels mentioned in Section 6.2 to compute θ_2 and y_2 . And a CUDA thread for weighted combination of y_1 and y_2 is executed at last.

6.4 | Diffuse particle seeding

Because diffuse particles are dynamically generated from projected grid according to the analysis results, we must design an appropriate data structure management on GPU to improve the algorithmic efficiency. We not only, first, estimate a maximum number of diffuse particles n_{seed} that each projected grid node may produce but also estimate a maximum number of the diffuse particles n_{max} with positive t_{life} to be handled simultaneously in the animation. Then we allocate memory for $n_{seed} * n_{node} + n_{max}$ diffuse particles, where n_{node} denotes the number of projected grid nodes.

In each simulation cycle, we invoke a CUDA thread for each node of the projected grid, in which we produce diffuse particles around the node, and further set the velocities and life time, respectively. The generated diffuse particles are dynamically inserted into the data structure proposed above, forming a sparse array. We then employ the thrust library to partition the array, moving the particles with positive duration time to the head of the array. At last, we invoke a thread for each particle to count the valid diffuse particles n_{now} , which is required in advection and rendering steps.

6.5 | Diffuse particle advection

For diffuse particle advection, we invoke one thread for each particle to update the force, velocity, and position. Specifically, we utilize the fast neighborhood searching algorithm to find the surrounding fluid particles of diffuse particles, and then compute the average velocity of the neighboring fluid particles (by using Equation 15). Meanwhile, we update the life time of diffuse particle in this thread. If the life time of

a diffuse particle is negative, it will be overwritten after the partitioning operation, without explicitly deleting it.

7 | EXPERIMENTS, EVALUATIONS, AND DISCUSSIONS

We have implemented our method on a PC with a GeForce GTX 780 GPU, Intel Core I7 CPU based on C++, CUDA, and GLSL APIs. In addition, we have re-rendered our experimental result using POV-Ray. Considering that the rendering algorithm does not truly dominate our subject, we finally chose the same particle-based rendering method to reduce the complexity as much as possible, and the resulted visual results are satisfactory. However, more complex rendering algorithm is also significant for better visual fidelity; for example, methods introduced by Van der Laan et al.⁵³ and Bagar⁵⁴ can be used to obtain the more faithful visual fidelity. For the issue of determining the resolution of depth buffer, we find that the best analysis result is achieved when each pixel of depth buffer represents almost one particle, which means that the width of depth buffer can be computed by $w_{db} = w_{scene} / (d_p)$, wherein w_{scene} denotes the width of scene and d_p is the diameter of fluid particle. From top to bottom, Figure 5 shows the results of the captured depth buffer, guided filter-based space analysis, time-space analysis, and simulation. Column (a) shows that higher-resolution results fail to represent the required local geometric information, while only reflecting the sphere shapes of particles in rendering. Meanwhile, the computational cost is more expensive compared with (b), which is the best in our experiments. Column (c) shows lower-resolution results, where the surrounding pixels involve analysis across a long distance in the simulation space. To handle details of

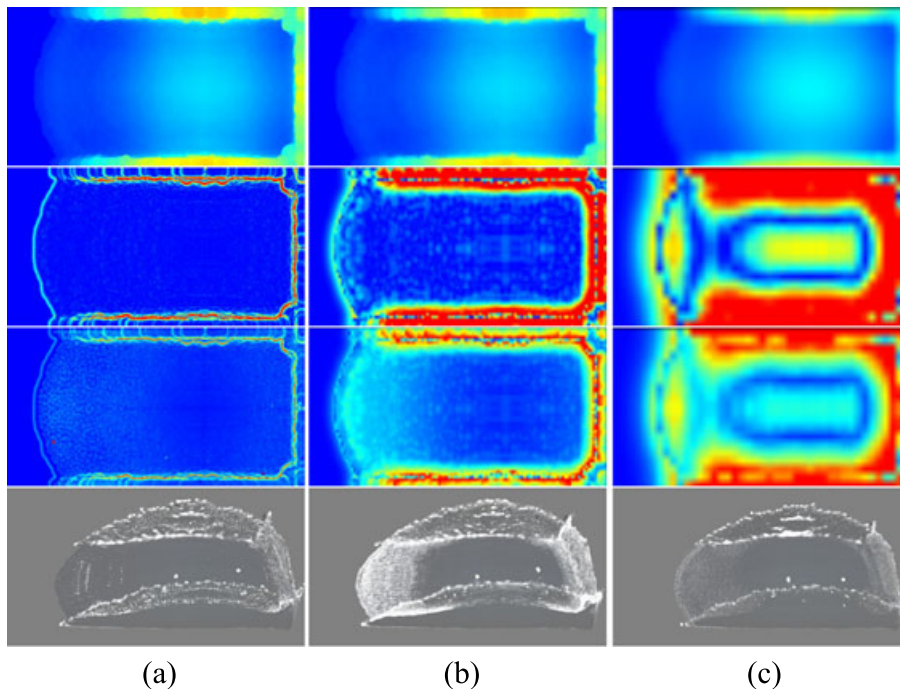


FIGURE 5 Results corresponding to different resolutions. Warm color indicates high value and vice versa. The resolution in simulation space is 140×60 , and the depth buffer resolutions are (a) 700×300 , (b) 140×60 , and (c) 70×30

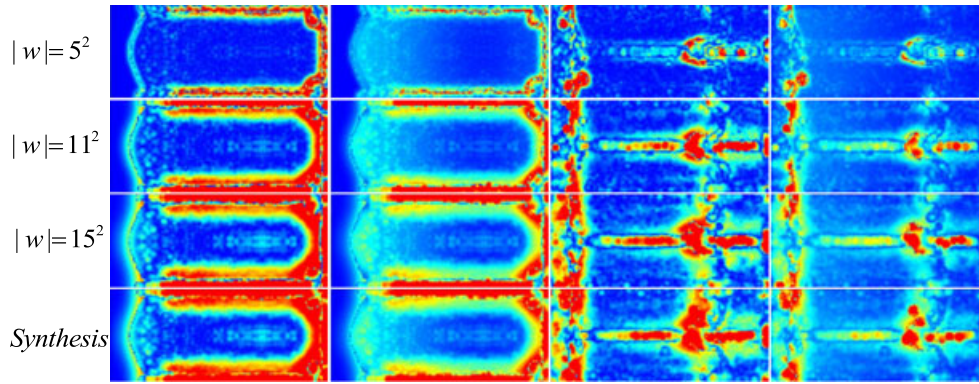


FIGURE 6 Our scale-aware analysis results. Warm color indicates high value and vice versa

different scales simultaneously, we detect guided filter features with different window sizes ($|w_k| = 5^2, 11^2, \text{ or } 15^2$) then synthesize the corresponding results together as the final detected features. As shown in Figure 6, our result reveals scale-aware features effectively.

Algorithm 1: Analysis-driven Detail Enhancement.

```

1   $n$  in the pass number of depth peeling
2  for  $i = 1$  to  $n$  do
3      capture the depth buffer  $d$  directly from scene.
4      smooth  $d$  with guided filter method.
5      compute the difference  $diff[i]$ .
6      compute the velocity field  $\mathbf{v}[i]$ .
7      initialize  $y_1[i], \theta_1[i]$ ;
8      while not converged do
9          while not converged do
10             update  $\theta_1[i]$  with Equation (11)
11             while not converged do
12                 update  $y_1[i]$  with Equation (12)
13             peel surface particles according to  $d[i]$  and distance
                threshold  $\lambda$ .
14 for  $i = 2$  to  $n$  do
15     compute the second-order derivative  $g_i^{2d}[i]$  in the
                depth direction.
16     compute the difference of velocity field  $\mathbf{g}_i^v[i]$  in the
                depth direction.
17     initialize  $y_2[i], \theta_2[i]$ ;
18     while not converged do
19         while not converged do
20             update  $\theta_2[i]$  with Equation (11)
21             while not converged do
22                 update  $y_2[i]$  with Equation (12)
23         compute a weighted combination of  $y_1[i]$  and  $y_2[i]$  to
                obtain  $y[i]$ .
24 base on linear interpolation, compute  $y_{All}$  according to  $y$ .
25 seed diffuse particles at the projected grid node
                according to  $d$  and  $y_{All}$ .
26 partition the array by thrust library.
27 advect all the diffuse particles with positive  $t_{life}$ .

```

Table 1 documents the statistics for the average testing time (in milliseconds) of each simulation cycle. We record the number of FLIP particles (#FP) and the maximum number of diffuse particles (#DPM) for each case. To be clear, (AN) and (DP) are the time costs for time–space analysis and diffuse particle generation and advection separately, (GR) means the grid resolution while (DBR) refers to *depth buffer resolution*, and (*Other*) is the time cost for deformable solid simulation and fluid–solid coupling. In comparison with the computational time of FLIP simulation (refer to (FLIP) column), the overhead of our time–space analysis and diffuse material simulation is small, while improving the visual effects significantly. The depth peeling technique is introduced into the simulation of dam-break scene and pouring scene. Some parameters and time statistics are shown in the subsequent columns, wherein (#D) means the passes of depth peeling (i.e., the number of depth layers) and (λ) represents the distance threshold for one-time peeling. The time cost of time–space analysis for all layers (AND) is approximately equal to (#D) the times of the analysis cost for one layer (AN), that is, the time cost of time–space analysis will increase in nearly linear way with the number increasing of depth layers. And level analysis (ANL) costs similar computational time as (AN), no matter how many depth layers there are. DPP refers to the time used for diffuse particle generation after depth peeling. Considering that the diffuse particles generated for each grid in one frame is limited, there is no definite multiple relationship between (DP) and (DPP), as shown in the last column of Table 1. In order to show how significantly the computational cost will increase when the layer number increases, we separately document the time cost of time–space analysis(AND), level analysis(ANL), and diffuse particle generation after depth peeling (DPP) in Table 2. And the time cost change caused by varying the distance threshold for one-time peeling(λ) is also considered in this table. As we can see from the last line of Table 2, λ does not significantly influence the computation cost. And Figure 7 documents the parameter values used in our experiments.

To further explore how significantly the parameter settings can influence the final visual results, we conduct some experiments under different parameter settings, and the results are

TABLE 1 Time performance (in milliseconds) of our experiments

Scene	#FP	GR	#DPM	DBR	FLIP	AN	DP	Other	#D	λ	AND	ANL	DPP
Dam-break (Figure 4)	248.4k	$96 \times 48 \times 64$	288×144	108.2k	47.7	20.2	5.3	-	4	$3d_p$	82.8	19.2	9.6
Interaction (Figure 10)	355.3k	$96 \times 48 \times 64$	288×144	49.6k	54.4	23.2	5.5	0.6	-	-	-	-	-
Pouring (Figure 12)	242.1k	$96 \times 64 \times 64$	288×192	38.7k	41.5	18.4	5.3	-	4	$2d_p$	74.5	18.9	9.7
Waterfall (Figure 11)	212.9k	$64 \times 64 \times 64$	192×192	40.8k	40.8	16.6	5.2	36.9	-	-	-	-	-

TABLE 2 Time performance (in milliseconds) with different parameters in depth peeling method in pouring scene

#D	λ	AN	DP	AND	ANL	DPP
1	$3d_p$	21.5	5.6	21.5	0	5.6
2	$3d_p$	20.1	5.8	41.3	17.6	6.2
3	$3d_p$	20.7	5.2	60.8	18.7	8.1
4	$3d_p$	20.2	5.3	82.8	19.2	9.6
5	$3d_p$	21.1	5.5	97.9	21.4	8.9
4	d_p	19.9	4.9	79.5	18.9	6.5
4	$2d_p$	20.4	5.1	81.8	19.8	7.6
4	$4d_p$	20.8	5.5	80.7	20.4	9.2

shown in Figure 7. Considering that the visual effects of detail enhancement are determined by the regression results, we only record and compare the output of logistic regression model when it adopts different parameter values, taking the logistic regression results of the first depth layer of frame 250 in pouring scene as example. Using the parameter listed in the top right corner of Figure 7 as basic setting, we can obtain the regression results as Figure 7(a). When setting r as 0.001 and 0.1 and keeping other parameter unchanged, we can obtain Figure 7(b1) and Figure 7(c1) separately. The difference between (b1) and (a) is shown in Figure 7(b2); the same numbering rule is used for other sub-figures. It is remarkable that only little difference can be observed when we change the value of parameter r , and the similar thing also happens when we change other parameter values. This phenomenon can reflect that our method are insensitive to parameter changes. Accordingly, to keep simplicity, we use the same parameter setting in all our experiments, as the basic parameter settings shown in Figure 7.

Figure 4 shows the simulation results of a dam-breaking scenario. In this scene, about 10k diffuse particles are generated nearby the wave crack, indicating the spray, foam, and small bubbles. As the fluid moves fiercely, the diffuse particles move correctly along with the wave. We can observe that the scale of the generated diffuse particles closely relates to the height (determining the value of depth buffer at each grid) and velocity of the fluid.

Our method is also effective for a fluid–solid coupling scene, as shown in Figure 8. A sphere is moving periodically on the fluid surface, pushing water aside and forming a small wave. Diffuse particles are generated according to the analysis result; these foams represent the trapped air in water, showing the trail of the sphere’s movement; such details cannot be captured without diffuse materials (as shown in the second row).

The third row shows the involved different types of diffuse particles, which mainly represent foams (colored in blue).

Figure 9 shows a waterfall scene. As the water stream falls down rapidly, the analysis result indicates complex interaction happens, wherein the diffuse materials are generated due to the complex interaction with the underlying terrain. Comparing with the result without diffuse particle simulation (refer to the second row), our method enhances the simulation plausibility by adding spray, bubble, and foam. This scene also demonstrates the effectiveness of our analysis for large-scale simulation. In comparison with Lipeng et al.,²¹ our method ignores the smoke-like sprays; nonetheless, we are able to simulate spray, bubble, and foam in a unified framework, which indicates that our method is more flexible and has a stronger capability for more complex scenarios.

Figure 10 demonstrates a pouring scene, where two violent above-water-level streams are spouted into the water tank on the right and another two streams on the left are injected from underwater. As the streams pouring down over the water from the right, numerous spray particles are generated to enhance the visual details. Furthermore, when streams flowing into the water volume, it is expected that plentiful bubbles should appear. Comparing with the depth peeling involved result (illustrated in the first row), we can clearly see that the simulation results without depth peeling cannot faithfully exhibit impetuous torrent of underwater (refer to the second row), which definitely looks less realistic. The same things happen in Figure 11, where the sea waves roll over surfaces again and again. Without depth peeling, we cannot deal with the details under the overlapping sea waves, and the bubbles would not appear under the sea surface even if the undercurrents have great velocities (refer to the first column of Figure 11).

In order to evaluate how the parameters of depth peeling impact the simulation results, we conduct some experiments with sharp contrast in Figure 12. Considering the water-pouring scene as an example, the first line of Figure 12 has a varying distance threshold ($\lambda = d_p, 2d_p, 3d_p, 4d_p$, and $5d_p$) for each pass of depth peeling. As we can see, when the number of peeling layer is fixed, particles (after each depth peeling) with small distance threshold may cause incomplete diffuse materials below the water level (refer to Figure 12(a)), while the results would not improve gradually as the distance threshold value grows constantly, and even over-fitting phenomena may appear (see Figure 12(e)). If we make the distance threshold λ be constant and set the number of depth layers $\#D$ to range from 1 to 5, the second line of Figure 12

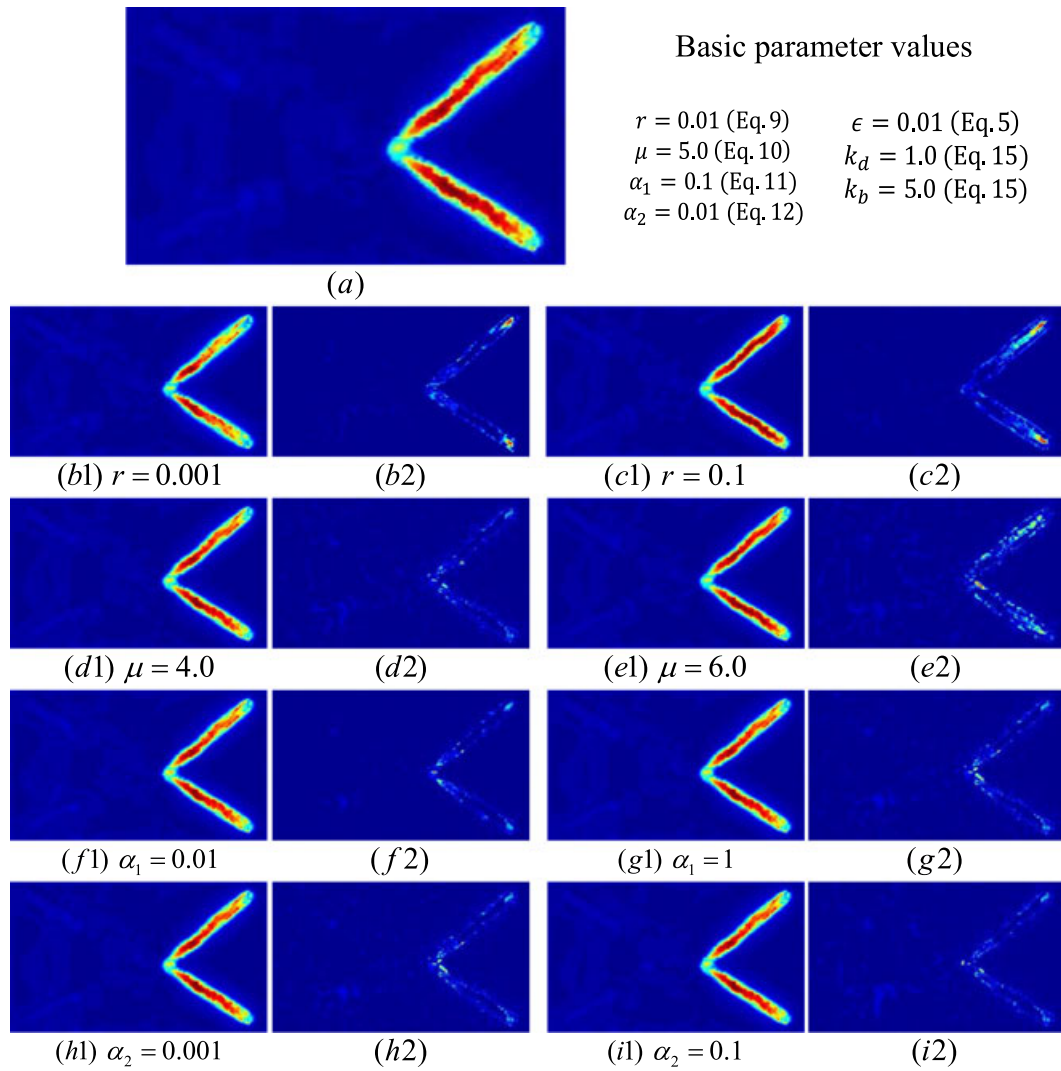


FIGURE 7 The logistic regression results of the first depth layer of frame 250 in pouring scene with different parameter settings. (a) The regression result with basic parameter settings; the parameter values are shown in the top right corner, the first and the third columns show the regression results with specific parameter settings (documented below the corresponding sub-figure), the second and the fourth columns show the difference between sub-figure (in the first and the third columns) and sub-figure (a). Warm color indicates high value and vice versa. The basic and suggested parameter values used in the experiment are given in the top right corner




FIGURE 8 The interaction effects between solid and water. The first row shows simulation result with diffuse materials, the second row shows the original FLIP simulation result, and the third row shows the classification of diffuse materials

can be obtained. As shown in Figure 12(f–g), too few layers may influence the experimental results; however, too many layers are meaningless for a certain scene, because

no obvious performance improvement could be recognized from Figure 12(h–j), while the computational cost increases significantly.




FIGURE 9 Waterfall. Liquid stream is pouring down into a tank of water, with generated spray, bubble, and foam (refer to the first row), and the second row shows the original simulation result of FLIP method




FIGURE 10 Pouring water into a water tank from the right (above the water level) and injecting water from the left (below the water level). Depth peeling is used in the first row, while the second row shows the result using the surface information only. The original FLIP simulation result is illustrated in the last row

Figure 13 shows the comparison between our method and that in Ihmsen et al.¹ Comparing with Ihmsen et al.,¹ our method defines completely different criteria to facilitate diffuse material generation and can achieve competitive visual effects. It may be noted that Ihmsen et al.,¹ requires extracting surface particles and computing the local curvature from the neighboring particles, which is more suitable for densely sampled fluids. In sharp contrast, our method does not rely on the particles' information when computing the fluid surface curvature and other features. Directly benefiting from our time-space analysis, the results evolve continuously during the simulation (please refer to Video S1). Moreover, our method has a significant improvement in terms of efficiency, primarily benefiting from our image-space-based GPU algorithmic architecture. However, when the method in Ihmsen et al.¹ is implemented on GPU, its diffuse particle generation and simulation costs around 30 ms for a scenario involving 100 k fluid particles.

According to Figures 10 and 11, it is easy to observe that, when fluid occlusion happens, *without depth peeling* the depth buffer can only capture the surface information, and thus cannot generate diffuse materials around the occluded fluid surface and inside water volume as well. However, after introducing *depth peeling technology*, we can obtain the depth and velocity information as deep as we desire inside the fluid volume through flexible parameter control, which affords us a new way for underwater analysis and enables us to handle more scenes even with fluid occlusion and overlapping phenomena. On the other hand, even though using depth peeling could improve the visual effect for detail enhancement to a great extent, it necessarily increases the computational cost for peeling and analysis.

In order to analyze how the adopted grid size can influence the visual effect, as shown in Figure 14, we take frame 350 of the pouring scene as example to show the experiment results under different grid sizes. Suppose that




FIGURE 11 Sea waves are rolling over surfaces again and again in this scene. The first row illustrates the simulation results with depth peeling, but in the second row no layer has been peeled. In the third row, FLIP method is employed only

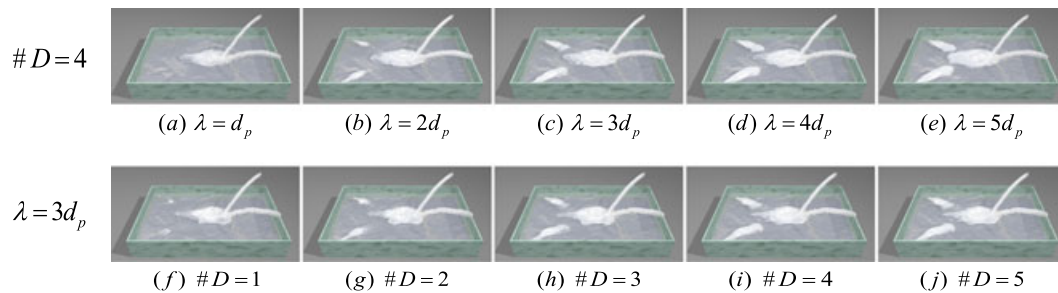


FIGURE 12 The first row shows the simulation results corresponding to different distance thresholds (λ) for one-time depth peeling, while the second row shows the simulation results corresponding to different passes of depth peeling ($\#D$) for one-time depth peeling




FIGURE 13 The comparison of the results from our paper (top two rows) and that from Ihmsen et al.¹ (bottom two rows)

the total size of this scene is represented as $1.5 \times 1 \times 1$, and the FLIP grid resolution is calculated using the total size divide by grid size. As shown in Figure 14, the grid size can significantly influence the FLIP simulation results. Considering that our diffuse particle generalization

method is actually based on the FLIP fluid simulation results, the final visual effects will be affected by the grid size. However, as shown in Figure 14, no matter what the grid size is, the fluid details can be enhanced properly with our model.

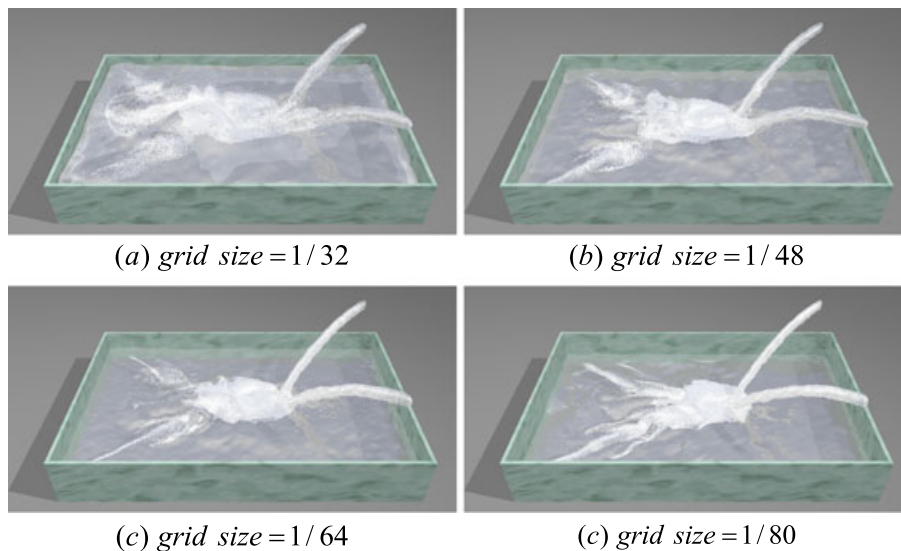


FIGURE 14 Comparison of the results based on different fluid implicit particle grid sizes. The total size of this scene is supposed as $1.5 \times 1 \times 1$. (a) The grid size is set as $1/32$. (b) The grid size is set as $1/48$. (c) The grid size is set as $1/64$. (d) The grid size is set as $1/80$

8 | CONCLUSION AND FUTURE WORK

In this paper, we have detailed a novel integrated framework for diffuse material animation and its visual detail enhancement by introducing a new analysis-and-simulation approach. The technical essence of our novel approach is the unification of time–space analysis in image domain and 3D physical simulation, built upon a CUDA-centric computational framework. The key innovation is that the geometry- and physics-based criteria, together with time–space integrated strategy, can be coupled into a logistic regression model. Moreover, the depth peeling method could improve the applicability and reality of visual detail enhancement beneath fluid surface and help us overcome difficulties when handling underwater scenes caused by fluid occlusion. Our method showcases the detail enhancements of complex fluid interaction phenomena and affords detail-preserving interaction while guaranteeing high efficiency even for scenes with 212.9k liquid particles and 40.8k diffuse particles. The involved data-driven analysis method and diffuse material simulation can be integrated into other fluid simulation grid-based method with very little extra workload, because our approach only depends on the captured depth buffer and certain local information during simulation. However, it is still difficult to apply depth peeling to grid-based method unless particles are automatically combined with the grid like FLIP, while generalizing depth peeling scheme to particle-based fluid algorithm is simple.

Our ongoing efforts include directly extending our time–space analysis model to process video data recording real fluid, which should offer more realistic diffuse material and more detail-informative free surfaces with better visual effects. Our novel approach is also readily available to be integrated with other available VR techniques to handle other types of visual data. In addition, the two-way coupling simula-

tion between diffuse materials and liquid also deserves further investigation.

ACKNOWLEDGMENT

This research is supported by the National Natural Science Foundation of China (61190120, 61190125, 61190124, 61300067, 61672077, 6167214, 61602341, and 61532002). We also thank the anonymous reviewers for their constructive critiques.

REFERENCES

1. Ihmsen M, Akinci N, Akinci G, Teschner M. Unified spray, foam and air bubbles for particle-based fluids. *Vis Comput*. 2012;28(6–8): 669–677.
2. Losasso F, Talton J, Kwatra N, Fedkiw R. Two-way coupled sph and particle level set fluid simulation. *IEEE TVCG*. 2008;14(4): 797–804.
3. Kim J, Cha D, Chang B, Koo B, Ihm I. Practical animation of turbulent splashing water. *Proc. of SCA'06, Vienna, Austria*; 2006. p. 335–344.
4. Mihalef V, Metaxas D, Sussman M. Simulation of two-phase flow with sub-scale droplet and bubble effects. *Comput Graph Forum*. 2009;28(2): 229–238.
5. Takahashi T, Fujii H, Kunimatsu A, Hiwada K, Saito T, Tanaka K, Ueki H. Realistic animation of fluid with splash and foam. *Comput Graph Forum*. 2003;22(3): 391–400.
6. Losasso F, Gibou F, Fedkiw R. Simulating water and smoke with an octree data structure. *ACM SIGGRAPH, Los Angeles, California, USA*; 2004. p. 457–462.
7. Adams B, Pauly M, Keiser R, Guibas LJ. Adaptively sampled particle fluids. *ACM Trans Graph*. 2007;26(3): 481–487.
8. Goswami P, Schlegel P, Solenthaler B, Pajarola R. Interactive SPH simulation and rendering on the GPU. *Proc. of SCA'10, Madrid, Spain*; 2010. p. 55–64.
9. Krog ØE, Elster AC. Fast GPU-based fluid simulations using SPH. In Manninen P, Öster P, eds. *Applied Parallel and Scientific Computing*, vol. 7134, Springer, Helsinki, Finland; 2012. p. 98–109.
10. Ihmsen M, Orthmann J, Solenthaler B, Kolb A, Teschner M. SPH fluids in computer graphics. *Eurographics 2014-State of the Art Reports*; 2014. p. 21–42.

11. Lipeng Y, Shuai L, Aimin H, Hong Q. Realtime two-way coupling of meshless fluids and nonlinear FEM. *Comput Graph Forum*. 2012;31(7): 2037–2046.
12. Monaghan JJ. Smoothed particle hydrodynamics. *Annu Rev Astron Astrophys*. 1992;30(1): 543–574.
13. Müller M, Charypar D, Gross M. Particle-based fluid simulation for interactive applications. *Proc. ACM SIGGRAPH/Eurographics Symp. Computer Animation, San Diego, California, USA; 2003*. p. 154–159.
14. Zhu Y, Bridson R. Animating sand as a fluid. *ACM Trans Graph*. 2005;24(3): 965–972.
15. Enright D, Marschner S, Fedkiw R. Animation and rendering of complex water surfaces. *ACM Trans Graph*. 2002;21(3): 736–744.
16. Foster N, Fedkiw R. Practical animation of liquids. *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '01; New York, NY, USA; 2001*. p. 23–30. ACM.
17. Gerszewski D, Bargteil AW. Physics-based animation of large-scale splashing liquids. *ACM Trans Graph*. 2013;32(6): 185:1–185:6.
18. Ando R, Thurey N, Tsuruno R. Preserving fluid sheets with adaptively sampled anisotropic particles. *IEEE Trans Vis Comput Graph*. 2012;18(8): 1202–1214.
19. Selino A, Jones MD. Large and small eddies matter: Animating trees in wind using coarse fluid simulation and synthetic turbulence. *Comput Graph Forum*. 2013;32(1): 75–84.
20. Cornelis J, Ihmsen M, Peer A, Teschner M. Iisph-flip for incompressible fluids. *Comput Graph Forum*. 2014;33(2): 255–262.
21. Lipeng Y, Shuai L, Aimin H, Hong Q. Hybrid particle-grid modeling for multi-scale droplet/spray simulation. *Comput Graph Forum*. 2014;33(7): 199–208.
22. Um K, Baek S, Han J. Advanced hybrid particle grid method with subgrid particle correction. *Comput Graph Forum*. 2014;33(7): 209–218.
23. Hong J-M, Lee H-Y, Yoon J-C, Kim C-H. Bubbles alive. *ACM SIGGRAPH, Los Angeles, California, USA; 2008*. p. 48:1–48:4.
24. Hong J-M, Kim C-H. Animation of bubbles in liquid. *Comput Graph Forum*. 2003;22(3): 253–262.
25. Zheng W, Yong J-H, Paul J-C. Simulation of bubbles. *Proc. of SCA'06, Vienna, Austria; 2006*. p. 325–333.
26. Boyd L, Bridson R. Multiflip for energetic two-phase fluid simulation. *ACM Trans Graph*. 2012;31(2): 16:1–16:12.
27. Nielsen MB, Osterby O. A two-continua approach to eulerian simulation of water spray. *ACM Trans Graph*. 2013;32(4): 67:1–67:10.
28. Patkar S, Aanjaneya M, Karpman D, Fedkiw R. A hybrid Lagrangian-Eulerian formulation for bubble generation and dynamics. *Proc. ACM SIGGRAPH/Eurographics Symp. Computer Animation, Anaheim, California, USA; 2013*. p. 105–114.
29. Müller M, Solenthaler B, Keiser R, Gross M. Particle-based fluid-fluid interaction. *Proc. of SCA'05, Los Angeles, California, USA; 2005*. p. 237–244.
30. Solenthaler B, Pajarola R. Density contrast SPH interfaces. *Proc. of SCA'06, Vienna, Austria; 2008*. p. 211–218.
31. Threy N, Sadlo F, Schirm S, Miller-Fischer M, Gross M. Real-time simulations of bubbles and foam within a shallow water framework. *ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA 2007; San Diego, California, USA, August; 2007*. p. 191–198.
32. Ihmsen M, Bader J, Akinci G, Teschner M. Animation of air bubbles with SPH. *GRAPP, Vancouver, Canada; 2011*. p. 225–234.
33. Losasso F, Talton J, Kwatra N, Fedkiw R. Two-way coupled SPH and particle level set fluid simulation. *IEEE Trans Vis Comput Graph*. 2008;14(4): 797–804.
34. Bagar F, Scherzer D, Wimmer M. A layered particle-based fluid model for real-time rendering of water. *Comput Graph Forum*. 2010;29(4): 1383–1389.
35. Changbo W, Qiang Z, Fanlong K, Hong Q. Hybrid particle-grid fluid animation with enhanced details. *Vis Comput*. 2013;29(9): 937–947.
36. Hammond DK, Vandergheynst P, Gribonval R. Wavelets on graphs via spectral graph theory. *Appl Comput Harmon A*. 2011;30: 129–150.
37. Narang SK, Ortega A. Compact support biorthogonal wavelet filterbanks for arbitrary undirected graphs. *IEEE TSP*. 2013;61(19): 4673–4685.
38. Xu L, Lu C, Xu Y, Jia J. Image smoothing via 10 gradient minimization. *ACM Trans Graph*. 2011;30(6): 174:1–174:12.
39. Tomasi C, Manduchi R. Bilateral filtering for gray and color images. *ICCV, Bombay, India; 1998*. p. 839–846.
40. He K, Sun J, Tang X. Guided image filtering. *IEEE Trans PAMI*. 2013;35(6): 1397–1409.
41. Everitt C. Interactive order-independent transparency. *White Paper*. 2001;7(7-8): 491–503.
42. Abraham M. Transparency and antialiasing algorithms implemented with the virtual pixel maps technique. *IEEE Comput Graph Appl*. 1989;9(4): 43–55.
43. Diefenbach PJ. Pipeline rendering: Interaction and realism through hardware-based multi-pass rendering; 1996. IRCS Technical Reports.
44. Bernardon FF, Pagot CA, Comba JLD, Silva CT. Gpu-based tiled ray casting using depth peeling. *J Graph Tools*. 2006;11(4): 1–16.
45. Nagy Z, Klein R. Depth-peeling for texture-based volume rendering. *Proceedings of the 11th Pacific Conference on Computer Graphics and Applications, PG '03; IEEE Computer Society, Washington, DC, USA; 2003*. p. 429–433.
46. Borland D. Volumetric depth peeling for medical image display. *Proceedings of SPIE - The International Society for Optical Engineering*. 2006;6060: 606004–606004–11.
47. Fang L, Mengcheng H, Xuehui L, Enhua W. Bucket depth peeling. *International Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 2009, August 3–7, 2009, Talks Proceedings; New Orleans, Louisiana, USA; 2009*. p. 1–1.
48. Cagniard C, Boyer E, Ilie S. Probabilistic deformable surface tracking from multiple videos. *ECCV, Hersonissos, Heraklion, Crete, Greece; 2010*. p. 326–339.
49. Khan ZH, Gu I. Y.-H.. Bayesian online learning on Riemannian manifolds using a dual model with applications to video object tracking. *ICCV Workshops, Barcelona, Spain; 2011*. p. 1402–1409.
50. Ciocca G, Cusano C, Schettini R. Image orientation detection using LBP-based features and logistic regression[J]. *Multimedia Tools and Applications*. 2015;74(9): 1–22.
51. Carmi A, Septier F, Godsill SJ. The Gaussian mixture MCMC particle algorithm for dynamic cluster tracking. *Automatica*. 2012;48(10): 2454–2467.
52. Van der Laan WJ, Green S, Sainz M. Screen space fluid rendering with curvature flow. *Proc. of I3D, Boston, USA; 2009*. p. 91–98.
53. Van der Laan WJ, Green S, Sainz MA. Screen space fluid rendering with curvature flow. *Proc. of I3D, Boston, USA; 2009*. p. 91–98.
54. Bagar F, Scherzer D, Wimmer M. A layered particle-based fluid model for real-time rendering of water. *Comput Graph Forum*. 2010;29(4): 1383–1389.

SUPPORTING INFORMATION

Additional Supporting Information may be found online in the supporting information tab for this article.

How to cite this article: Qiu, Y., Yang, L., Li, S., Xia, Q., Qin, H., and Hao, A. (2016), Novel Fluid Detail Enhancement based on Multi-Layer Depth Regression Analysis and FLIP Fluid Simulation, *Comp Anim Virtual Worlds*, doi:10.1002/cav.1741


AUTHORS BIOGRAPHIES



Yuxing Qiu is now a graduate student at the State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, China. Her research interests include computer graphics, physics-based modeling and simulation, particularly in fluid.



Lipeng Yang received his PhD degree in Computer Science from at the State Key Laboratory of Virtual Reality Technology and Systems at Beihang University, China. His research interests include computer graphics, physics-based modeling and simulation, fluid simulation, and virtual surgery simulation.



Shuai Li received his PhD degree in cComputer sScience from Beihang University. He is currently an assistant professor at the State Key Laboratory of Virtual Reality Technology and Systems, Beihang University. His research interests include computer graphics, physics-based modeling and simulation, virtual surgery simulation, computer vision, and medical image processing.



Qing Xia is now a PhD candidate at the State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, China. His research interests include computer graphics, physics-based modeling, geometry modeling, and processing and shape analysis.



Hong Qin is a full professor of Computer Science in the Department of Computer Science at Stony Brook University (SUNY). He received his BS and his MS degrees in Computer Science from Peking University, China. He received his PhD in Computer Science from the University of Toronto. Currently, he serves as an associate editor for *The Visual Computer*, *Graphical Models*, and *Journal of Computer Science and Technology*. His research interests include geometric and solid modeling, graphics, physicsbased modeling and simulation, computer aided geometric design, human-computer interaction, visualization, and scientific computing.



Aimin Hao is a professor in Computer Science School and the associate director of State Key Laboratory of Virtual Reality Technology and Systems at Beihang University. He received his BS, MS, and PhD degrees in Computer Science at Beihang University. His research interests are on virtual reality, computer simulation, computer graphics, geometric modeling, image processing, and computer vision.