

**LA-UR-24-30258**

Accepted Manuscript

## Multi-plane moment-of-fluid interface reconstruction in 3D

Spainhour, Jacob

Shashkov, Mikhail Jurievich

Provided by the author(s) and the Los Alamos National Laboratory (1930-01-01).

**To be published in:** Journal of Computational Physics

**DOI to publisher's version:** 10.1016/j.jcp.2025.114239

**Permalink to record:**

<https://permalink.lanl.gov/object/view?what=info:lanl-repo/lareport/LA-UR-24-30258>




Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA000001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

## Journal Pre-proof

Multi-plane moment-of-fluid interface reconstruction in 3D

Jacob Spainhour, Mikhail Shashkov

PII: S0021-9991(25)00522-4  
DOI: <https://doi.org/10.1016/j.jcp.2025.114239>  
Reference: YJCPH 114239



To appear in: *Journal of Computational Physics*

Received date: 24 October 2024  
Revised date: 25 April 2025  
Accepted date: 10 July 2025


Please cite this article as: Jacob Spainhour, Mikhail Shashkov, Multi-plane moment-of-fluid interface reconstruction in 3D, *Journal of Computational Physics* (2025), doi: <https://doi.org/10.1016/j.jcp.2025.114239>

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.


© 2025 Published by Elsevier Inc.

### Multi-plane moment-of-fluid interface reconstruction in 3D

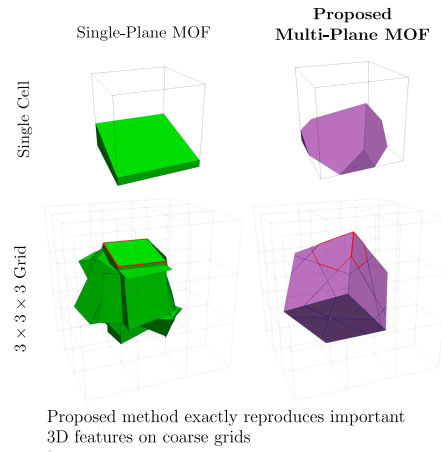
Select plane parameters to minimize error in up to 3<sup>rd</sup> order geometric moments



Ellipsoid constructed from 2<sup>nd</sup> order moments defines normalization scheme



Faces of a single inscribed polyhedron are initial guess for non-linear optimization



Journal Pre-proof

# Multi-plane moment-of-fluid interface reconstruction in 3D

Jacob Spainhour<sup>a,\*</sup>, Mikhail Shashkov<sup>b</sup>

<sup>a</sup> *Department of Applied Mathematics, University of Colorado Boulder  
11 Engineering Dr, Boulder, CO 80309, USA*

<sup>b</sup> *X-Computational Physics Division, XCP-4  
Los Alamos National Laboratory, Los Alamos, NM 87545, USA*

---

## Abstract

Moment-of-fluid (MOF) methods for interface reconstruction approximate the region occupied by material in each mesh element only through reference to its geometric moments. We present a 3D MOF method that represents the material (POM) in each cell as the convex intersection of the cell and multiple half-spaces, each selected to minimize the least-squares error between computed moments of the approximated material and provided reference moments.

This optimization problem is highly non-linear and non-convex, making the numerical result very sensitive to the initial guess. To create an effective initial guess in each cell, we construct an ellipsoid from 0<sup>th</sup>-2<sup>nd</sup> order reference moments such that its shape corresponds with that of the POM. Within this ellipsoid we inscribe a polyhedron, and initialize the minimization problem with the half-spaces defined by each of its faces. The inscribed polyhedron has minimally 4 faces, and using up to 3<sup>rd</sup> order moments permits optimization over up to 20 unknown values. We therefore define MOF methods that utilize 4, 5, or 6 half-spaces, correspondingly initialized with the faces of a single inscribed tetrahedron, triangular prism, or hexahedron. Stability of the non-linear optimization is further improved with a preprocessing step that normalizes the reference moments according to the axes of the reference ellipsoid.

Using this approach, the non-linear least-squares solver reliably converges to a near-global minimum from a single initial guess. We demonstrate accuracy and robustness using single-cell and multi-cell examples over a wide spectrum of geometry. In particular, we demonstrate our ability to exactly reproduce several important and complex features defined by up to four half-spaces, such as corners, filaments, filament tips, and embedded material in the cell.

*Keywords:* 3D interface reconstruction, Moment-of-fluid, Convex polyhedra, Non-linear optimization

---

## 1. Introduction

### 1.1. Background and Rationale

In simulating the flow of multiple materials, their shape must be accurately captured over a discrete computational mesh. The goal of material interface reconstruction is then to use data provided by the simulation of this flow to approximate the intersection of the region occupied by material with each mesh element—the per-cell pieces of material (POM)—while conserving the volume of each POM.

Among volume tracking techniques, methods for piecewise linear interface calculation (PLIC) are widely used for their simplicity and utility in adjacent applications. Such methods define the interface in each cell with a single linear element whose position is informed by the tracked volume, thereby enforcing a physically conservative solution. At the same time, the orientation of this segment is chosen to resemble that of the underlying material fragment. In the volume-of-fluid family of methods, the orientation is informed by the volume fraction of neighboring cells, such as via an estimate of the gradient of the volume fraction field [1, 2], a coupled level-set method [3, 4], or more recently methods of machine learning [5].

There is also longstanding interest in VOF methods which represent the per-cell interface with more complex geometry, which serves to increase their accuracy or permit more direct analysis of derived

---

\*Corresponding author [Jacob.Spainhour@colorado.edu](mailto:Jacob.Spainhour@colorado.edu)

This work was performed at LANL while first author was a summer intern. LA-UR-24-30258

quantities like curvature. For example, there are 2D PLIC methods which represent the interface as a single “bent” line [6, 7, 8], or even multiple disjoint segments [9, 10]. Other 2D methods utilize high-order curves to approximate the interface, such as parabolas [11, 12, 13], splines [14, 15], or circular arcs [16, 17, 18]. In three spatial dimensions, such methods are extended to utilize paraboloids [12] and spheres [19].

However, in all of these examples, the placement of the interface is necessarily informed by the data of neighboring cells. In contrast, a key feature of moment-of-fluid (MOF) methods is that each POM is approximated using only data from the cell it occupies. In particular, we consider the geometric moments of the piece of material, the most general geometric characteristics of any shape [20].

In this work, we assume that the moments of each POM up to some order (referred to hereafter as reference moments) are given to us according to the specific setting. As a motivating example, we consider hydrodynamic simulation. At the initial time, moments can be obtained from the known initial material shape by its exact or approximate intersection with cells of the mesh [21, 22]. At later times, the material moments have been advanced in time according to the governing physics of the problem, the simplest process being advection [23, 24, 25, 26]. Here, however, we consider the origin of these moments to be irrelevant to their use in the described MOF reconstruction methods.

In general, the MOF approach for interface reconstruction defines an approximation of each POM by intersecting the computational cell with some shape. These shapes are parameterized in such a way that the difference between the moments of their intersection with the cell and that cell’s reference moments can be minimized via numerical optimization. This process is repeated across all cells in the mesh to produce an approximation of the boundary of the full material [27]. In general, this approximation will be discontinuous, as constructing a continuous boundary across an entire mesh requires considering non-local information, either within a local stencil [28, 8] or globally across the mesh [7].

As an illustrative example, the prototypical MOF-PLIC method in 2D uses the material volume (zeroth-order moment) and centroid (ratios of first-order and zeroth-order moments) in each computational cell to determine parameters of a single half-plane. These parameters are selected to minimize the error between the reference centroid and the computed centroid of the material polygon which approximates the POM in the cell [27]. This methodology extends readily to 3D, using the intersection of a single half-space and the cell to approximate the material polyhedron for each POM [23, 29, 30, 31].

While one-plane methods are particularly numerically efficient (and can currently be made *more* efficient using an advanced implementation [32] or analytic formulae [33, 29, 31]), there is also considerable interest in extensions of local MOF methods which utilize a more complex interface to approximate each POM, particularly in two spatial dimensions. This includes methods that utilize circular arcs [34], more general quadratic curves [35, 36, 37], or multiple linear segments [38, 24, 39, 40]. Importantly, many of these methods incorporate higher-order moments in the error function to ensure that the ensuing minimization problem is not underdetermined. In this work, we extend this pursuit in 3D using a new MOF method for interface reconstruction that uses locally defined, high-order moment data to approximate the material boundary in each cell with a shape that is far more geometrically expressive than than a single plane.

### 1.2. Motivation

There are many incentives to develop computationally efficient interface reconstruction methods in 3D, particularly those that utilize local information to reproduce geometrically complex material. For example, the exclusive use of per-cell data makes such a method more robust to unstructured computational meshes.

Furthermore, the computational benefits of a parallelizable method like MOF become even more pronounced as the number of cells increases with spatial dimension.

However, there is considerably more variety in the kinds of sub-cell geometric features that cannot be adequately captured by the single plane used in a typical MOF-PLIC method, several of which we categorize in Figure 1. Many of these features are described by extruding their 2D analogues along the third dimension, such as “wall”-type shapes and edge-corners that require two planes to be reproduced exactly. Others are more unique to 3D space, such as 3D filaments and tips, as well as vertex-corners, each of which require three planes to accurately reproduce. Finally, four planes are necessary to adequately approximate regions that are entirely embedded within a computational cell. With these shapes in mind, we target a MOF method that represents the interface by the convex intersection of, at minimum, four half-spaces. To our knowledge, there exist no contemporary methods of performing a 3D MOF interface reconstruction with interface geometry more complex than a single plane.




Figure 1: Shape types that can be described by up to four planes. We target a method that can reproduce these types of geometric features in a single computational cell.

We prioritize a reconstruction method with planar boundary components because doing so ensures it is compatible with adjacent physics applications. For example, performing dynamic reconstructions for advection through Lagrangian remapping requires computing intersections between a set of likely disjoint POMs and arbitrary backtraced computational cells, and so there is little fundamental difference applying such a technique to a single plane (as in [23]) and in the proposed multi-plane technique. Similarly, such a method generalizes far more readily to multi-material scenarios using the strategy of nested dissection (again described in [23]), as exact moments can be computed from the intersection of arbitrarily many material interfaces. This is often not the case for methods which consider curved interfaces: even if one can compute exact moments for the intersection of the computational cell and a sphere [41] or paraboloid [13], the problem becomes intractable when even two separate curved interfaces are considered. Because using the proposed MOF method within these advection and multi-material frameworks is, by design, largely a direct application of existing techniques, we consider their specific implementation to be beyond the scope of this work. Instead, we restrict our attention solely to the reconstruction problem.

As in the typical MOF approach, we formulate interface reconstruction as a non-linear optimization problem that finds a representation of the approximate POM to best match its computed moments to some provided reference moment data.

As an aside, we follow the notation of [34] and distinguish between MOF methods with subscripts and superscripts. For example, consider the  $\text{MOF}_{2hp}^2$  method for 2D interface reconstruction, where the superscript represents the maximum order of the moments considered, and the subscript represents the type of bounding geometry used, in this case 2 half-planes. Despite the slight abuse of notation, it will be clear from context whether the method is designed for use 2D or 3D space. For example,  $\text{MOF}_{2hp}^2$  is

assumed to be a 2D method using two half-planes. The analogous 3D method is  $\text{MOF}_{2hs}^2$ , which uses the same order of moments, but instead bounds the POM by two 3D half-spaces.

In  $\text{MOF}_{2hp}^2$  and many derived 2D techniques, an ellipse is used to create the initial guess for the relevant non-linear optimization problem [24]. A unique reference ellipse can be constructed from moment data up to second-order so that its volume and centroid coincide with the material, and so that the length and orientation of its main axes correspond with the principle directions and magnitude of elongations of the ground-truth material. Although the construction of this reference ellipse generalizes readily to three spatial dimensions in the form of a reference ellipsoid, the analogous *usage* is necessarily different in 3D, in large part due to same complexities of 3D space that motivate high-order MOF methods in the first place. The least-squares error between computed moments and reference moments used in the objective function of many MOF methods is often both highly non-linear and non-convex as a function of the chosen parameterization, and so typical descent methods tend towards local minima if the initial guess is far from the desired global minimum. This makes the numerical result of the optimization highly sensitive to the choice of initial guess. To compensate for this, the final interface in  $\text{MOF}_{2hp}^2$  is generated through repeated solutions of the optimization problem, each of which is initialized using a pair of edges from a polygon inscribed in the reference ellipse. Ideally, the intersection of one pair with the cell will sufficiently approximate the material so that the subsequent descent algorithm converges to the true global minimum [24].

While this is already an expensive process in 2D, the analogous procedure in 3D is largely infeasible. This is due to the additional combinatorial burden of selecting faces from an inscribed polyhedron combined with the increased cost of solving the non-linear optimization. A possible alternative to the combinatorial approach in 2D is presented in the  $\text{MOF}_{PIE}^2$  method of [40], where the interface is instead represented by the intersection of the computational cell with *all* half-planes defined by the edges of the polygon inscribed in an ellipse. The position of the polygon is fixed relative to the ellipse, and then the ellipse's 5 parameters are optimized in a typical MOF fashion, minimizing the difference between computed and reference moments. This permits the use of a single initial guess for the non-linear optimization, the reference ellipse itself. This approach potentially generalizes directly to 3D, with an interface represented by a polyhedron inscribed in a 9 parameter ellipsoid, initialized using reference moments. However, such a method is limited by the large number of configurations possible for the connectivity of vertices and edges for general polyhedron, making it difficult for the same initial guess to converge to all possible geometric features of interest.

It is through comparison to these methods that we motivate our proposed strategy for interface reconstruction, guided by a desire to represent complex 3D geometry with as few solutions of the non-linear optimization problem as possible. As in  $\text{MOF}_{PIE}^2$ , we initialize the optimization procedure with an interface represented by the faces of a polyhedron inscribed in the reference ellipsoid, each of which define a half-space in the intersection with the cell. However, rather than parameterizing the optimization through the parameters of the circumscribing ellipsoid, we instead optimize over the full set of parameters of *all* planes defined by the polyhedron. In doing so, we essentially encode geometric information from a larger set of potential initial guesses, which in the framework of  $\text{MOF}_{2hp}^2$  would have been described by subsets of faces from an inscribed polyhedron, into a single initial guess for numerical optimization.

Properly defining a closed, inscribed polyhedron requires at minimum 4 faces to form a tetrahedron, which is also conveniently sufficient to appropriately recreate our target geometries depicted in Figure 1.

However, this results in a non-linear least-squares problem with minimally 12 parameters, which is underdetermined when only the 10 zeroth- to second-order moments are considered. To compensate for this, we instead minimize the least-squares error over the 20 moments up to third order. Incorporating higher-order moments is also useful from a purely geometric perspective, as third-order moments encode a greater amount of information about the material's orientation, a fact that has motivated [40] to consider a

2D  $\text{MOF}_{PIE}^3$  method alongside  $\text{MOF}_{PIE}^2$ .

This fully describes the  $\text{MOF}_{4hs}^3$  method, which utilizes 3<sup>rd</sup> order moments to construct an interface out of 4 half-spaces, initialized from a single tetrahedron inscribed in the reference ellipsoid. Furthermore, minimizing the error between 20 third-order moments also permits definition of a 15 parameter  $\text{MOF}_{5hs}^3$  method and a 18 parameter  $\text{MOF}_{6hs}^3$  before the optimization problem becomes underdetermined. These methods are initialized with the faces of a triangular prism and a hexahedron respectively. Although the minimization problem becomes proportionally more expensive with each additional plane parameterized,

each of these methods is capable of recreating shapes with complex sub-cell features from a single initial guess, thereby avoiding the even greater cost of repeating the minimization procedure multiple times to achieve the desired interface.

### 1.3. Objectives and Article Structure

Our primary objective in this work is for our MOF strategy to consistently reproduce the geometry of the shapes depicted in Figure 1 from a single guess. With respect to this objective, the three proposed  $\text{MOF}_{Nhs}^3$  methods are somewhat interchangeable: Although the number of faces in the initial polyhedron varies, each method is similarly “over-parameterized” for this set of target features. Regardless of the method, we *expect* that some planes will be redundant or inactive at the desired minimum, particularly in cases where an exact reconstruction is possible. Furthermore, we will show in such contexts that the global minimum is consistently achieved from a single guess with even the four-plane  $\text{MOF}_{4hs}^3$  method.

However, there are clearly cases where the ground-truth shape is described by more planes than are available during optimization, and many more where the shape cannot be exactly reproduced with *any* number of planes. In such cases, we will show that the robustness of our approach ensures that any of our three methods will generate an interface that adequately reconstructs the shape’s position and orientation within the cell, as well as many important sub-cell geometric features, such as the location of its intersection with the cell’s faces.

Furthermore, we will also show that the additional planes considered by  $\text{MOF}_{4hs}^3$ ,  $\text{MOF}_{5hs}^3$ , and  $\text{MOF}_{6hs}^3$  are often utilized to represent shapes that are *more* complex than those depicted in Figure 1. In scenarios where only an inexact reconstruction is possible, there is an important trade-off among these three methods between computational cost and the maximum geometric complexity of the optimized interface.

While we do not expect these inexact reconstructions to always represent the *global* minimum of their corresponding optimization problem, the computed *local* minimum is consistently higher quality than any conventional alternative. As reproduction of the shapes in Figure 1 represent the primary objective of this work, it is in some sense an unexpected benefit of our multi-plane framework for interface reconstruction that it also effectively depicts these more complex features.

This paper proceeds as follows. In Section 2, we present an overview of the relevant details on 3D moments including formulas related to the translation, scaling, and rotation transformations. These are used to normalize given moment data and in the construction of the reference ellipsoid. Next, in Section 3 we discuss details of the new MOF interface reconstruction method. This includes definition of the objective function for non-linear optimization, our choice of initial guess, and the customized Levenberg-Marquardt algorithm used, among other details. We summarize the proposed MOF algorithms in Section 4. We present numerical results in Section 5, including the results of reconstruction on a single computational cell as well as across the full mesh. Discussion related to practical implementation is presented in Section 6.

This includes the reconstruction of non-convex material, the appearance of geometric artifacts, 3D polyhedra with identical moments up to some order, and sensitivity to noise in reference moment data. Conclusions and a discussion on future work to be done is in Section 7. Finally, acknowledgments are given in Section 8.

## 2. Moments Primer

### 2.1. Translation, Scaling, and Rotation of Volume Moments

To introduce the necessary notation, we restate the objective of the proposed method as finding an approximation of the material  $\Omega$  within a convex computational cell  $C$ . We denote this approximation as  $\Omega_r$ .

In this work, we consider standard 3D geometric moments of  $\Omega$  given by

$$M_{ijk} = \int_{\Omega} x^i y^j z^k dV. \quad (1)$$

These “raw” moments describe important geometric features of the material, such as the volume  $M_{000}$  and centroid  $\vec{c} = (M_{100}/M_{000}, M_{010}/M_{000}, M_{001}/M_{000})$ . However, it is useful and often necessary to consider alternate formulations and transformations of moment data. For example, consider a shift transform

$S(s_x, s_y, s_z) : \Omega \rightarrow \Omega'$ . If the material is shifted by  $\vec{s} = (s_x, s_y, s_z)$  according to  $\widehat{\Omega} = \Omega - \vec{s}$ , then its moments can be computed as

$$M'_{ijk} = \int_{\Omega - \vec{s}} x^i y^j z^k dV = \int_{\Omega} (x - s_x)^i (y - s_y)^j (z - s_z)^k dV. \quad (2)$$

Importantly, this latter formulation allows one to compute these moments directly from their raw counterparts without reference to the underlying material  $\Omega$ . For example, moments up to second order are given by

$$\begin{aligned} M'_{000} &= M_{000} \\ M'_{100} &= M_{100} - s_x M_{000} \\ M'_{010} &= M_{010} - s_y M_{000} \\ M'_{001} &= M_{001} - s_z M_{000} \\ M'_{200} &= M_{200} - 2s_x M_{100} + s_x^2 M_{000} \\ M'_{020} &= M_{020} - 2s_y M_{010} + s_y^2 M_{000} \\ M'_{002} &= M_{002} - 2s_z M_{001} + s_z^2 M_{000} \\ M'_{110} &= M_{110} - s_x M_{010} - s_y M_{100} + s_x s_y M_{000} \\ M'_{101} &= M_{101} - s_x M_{001} - s_z M_{100} + s_x s_z M_{000} \\ M'_{011} &= M_{011} - s_z M_{010} - s_y M_{001} + s_y s_z M_{000}. \end{aligned} \quad (3)$$

Similarly, one can apply a scaling  $L(l_x, l_y, l_z) : \Omega \rightarrow \Omega'$  to moment data along each coordinate axis by  $(l_x, l_y, l_z)$  using the transform

$$M'_{ijk} = l_x^{1+i} l_y^{1+j} l_z^{1+k} \cdot M_{ijk} \quad (4)$$

Finally, we also consider *rotations* around the origin,  $R(\alpha, \beta, \gamma) : \Omega \rightarrow \Omega'$  defined by the following rotation matrix:

$$R(\alpha, \beta, \gamma) = \begin{bmatrix} c_\beta c_\gamma & s_\alpha s_\beta c_\gamma - c_\alpha s_\gamma & c_\alpha s_\beta c_\gamma + s_\alpha s_\gamma \\ c_\beta s_\gamma & s_\alpha s_\beta s_\gamma + c_\alpha c_\gamma & c_\alpha s_\beta s_\gamma - s_\alpha c_\gamma \\ -s_\beta & s_\alpha c_\beta & c_\alpha c_\beta \end{bmatrix}, \quad (5)$$

where  $\alpha, \beta$ , and  $\gamma$  are rotations around the  $x, y$ , and  $z$  axes respectively.

Written explicitly, this gives the transformation

$$\begin{aligned} M'_{ijk} &= \int_{R(\Omega; \alpha, \beta, \gamma)} x^i y^j z^k dV \\ &= \int_{\Omega} (c_\beta c_\gamma x + (s_\alpha s_\beta c_\gamma - c_\alpha s_\gamma) y + (c_\alpha s_\beta c_\gamma + s_\alpha s_\gamma) z)^i \cdot \\ &\quad (c_\beta s_\gamma x + (s_\alpha s_\beta s_\gamma + c_\alpha c_\gamma) y + (c_\alpha s_\beta s_\gamma - s_\alpha c_\gamma) z)^j \cdot \\ &\quad (-s_\beta x + s_\alpha c_\beta y + c_\alpha c_\beta z)^k dV \end{aligned} \quad (6)$$

Under this transformation, the rotated first-order moment data are given by

$$\begin{aligned} M'_{000} &= M_{000} \\ M'_{100} &= c_\beta c_\gamma M_{100} + (s_\alpha s_\beta c_\gamma - c_\alpha s_\gamma) M_{010} + (c_\alpha s_\beta c_\gamma + s_\alpha s_\gamma) M_{001} \\ M'_{010} &= c_\beta s_\gamma M_{100} + (s_\alpha s_\beta s_\gamma + c_\alpha c_\gamma) M_{010} + (c_\alpha s_\beta s_\gamma - s_\alpha c_\gamma) M_{001} \\ M'_{001} &= -s_\beta M_{100} + s_\alpha c_\beta M_{010} + c_\alpha c_\beta M_{001}. \end{aligned} \quad (7)$$

Higher order moments can be computed similarly. We also note for future use that the inverse transformation  $R^{-1}(\alpha, \beta, \gamma)$  has the same form as Equation 6, but with the inverse of the rotation matrix in Equation 5.

## 2.2. Reference Ellipsoid and Normalization

To ensure stability during numerical optimization, we normalize the reference moments  $M_{ijk}^r$  with the transformations described in Section 2.1 such that the underlying material  $\Omega$  is translation-invariant and scale-invariant. This is a critical step for the proposed MOF methods, as it makes subsequent optimization steps far more robust to unimportant features of the material, such as spatial position and size within the cell.

As our first transformation, we create ‘‘central moments’’  $\overline{M}_{ijk}^r$  which correspond to the reference moments  $M_{ijk}^r$  shifted by the centroid  $\vec{c} = (M_{100}^r/M_{000}^r, M_{010}^r/M_{000}^r, M_{001}^r/M_{000}^r)$  according to Equation 2. Under this transformation, we can then consider other useful properties of the second-order central moments, which encode information about the orientation of the material. Specifically, we define a covariance matrix for the material in terms of its central moments, given by

$$A = \frac{1}{\overline{M}_{000}^r} \begin{bmatrix} \overline{M}_{200}^r & \overline{M}_{110}^r & \overline{M}_{101}^r \\ \overline{M}_{110}^r & \overline{M}_{020}^r & \overline{M}_{011}^r \\ \overline{M}_{101}^r & \overline{M}_{011}^r & \overline{M}_{002}^r \end{bmatrix}. \quad (8)$$

This matrix is necessarily symmetric and positive semi-definite, and so it defines a quadratic form  $x^T A x$  whose zero level set  $x^T A x = 0$  corresponds to a *reference* ellipsoid in  $\mathbb{R}^3$ . Historically, quadratic forms of this kind have been used in 2D to define effective initial guesses for the non-linear optimization problem, as the position and orientation of the derived ellipse roughly corresponds to the position and orientation of the underlying material [24, 40]. In this work, we normalize the material according to the corresponding reference ellipsoid through both a rotation transformation and a scaling transformation. Done properly, the centered material  $\overline{\Omega}$  and its corresponding moments  $\overline{M}_{ijk}^r$  are normalized so that the transformed moments  $\widehat{M}_{ijk}^r$  of the transformed  $\widehat{\Omega}$  generate a diagonal covariance matrix, and therefore a spherical reference ellipsoid.

To this end, let  $\lambda_0, \lambda_1, \lambda_2$  be the three eigenvalues of  $A$  with associated eigenvectors  $v^{(0)}, v^{(1)}, v^{(2)}$ . These eigenvectors can be used to define rotation angles

$$\begin{aligned} \alpha &= \text{atan2}(v_z^{(1)}, v_z^{(2)}), \\ \beta &= \text{asin}(-v_z^{(0)}), \\ \gamma &= \text{atan2}(v_y^{(0)}, v_x^{(0)}). \end{aligned} \quad (9)$$

These angles are selected so that, when used in the *inverse* rotation transformation  $R^{-1}(\alpha, \beta, \gamma)$ , the orientation of the centered material  $\overline{\Omega}$  is ‘‘undone’’, and the resulting intermediate material is aligned with the Cartesian axes. Furthermore, each axis-length  $a, b, c$  of this ellipsoid is proportional to the square root of an eigenvalue of  $A$ , which permits a scaling of each such that the respective lengths are equal. At the same time, we resolve the ambiguity in the constant of proportionality by selecting  $l_x, l_y$ , and  $l_z$  so that the result is also normalized to have unit material volume  $\widehat{M}_{000}^r$ . Together, this results in the following scaling coefficients  $l_x, l_y$ , and  $l_z$ :

$$l_x = \frac{1}{\sqrt{\lambda_0}} \left( \frac{\sqrt{\lambda_0 \lambda_1 \lambda_2}}{\overline{M}_{000}^r} \right)^{1/3}, \quad l_y = \frac{1}{\sqrt{\lambda_1}} \left( \frac{\sqrt{\lambda_0 \lambda_1 \lambda_2}}{\overline{M}_{000}^r} \right)^{1/3}, \quad l_z = \frac{1}{\sqrt{\lambda_2}} \left( \frac{\sqrt{\lambda_0 \lambda_1 \lambda_2}}{\overline{M}_{000}^r} \right)^{1/3}. \quad (10)$$

By construction, applying these to the reference central volume moment  $\overline{M}_{000}^r$  through Equation 4 results in

$$\widehat{M}_{000}^r = l_x l_y l_z \overline{M}_{000}^r = \frac{1}{\sqrt{\lambda_0 \lambda_1 \lambda_2}} \left( \frac{\sqrt{\lambda_0 \lambda_1 \lambda_2}}{\overline{M}_{000}^r} \right) \overline{M}_{000}^r = 1 \quad (11)$$

as desired.

Taken together, this sequence of transformations turn the material  $\Omega$  into  $\widehat{\Omega}$  that is centered at the origin, has unit volume, and is not stretched along any particular axis. We can see an example of each stage of

this normalization in Figure 2. From another perspective, we can consider this mapping to apply not to  $\Omega$ , but to the entire space in which it is embedded. This means that during optimization, the *raw* moments of the reconstruction are matched to the *raw* moments of  $\widehat{\Omega}$ , since these are equal to the centered, rotated, and scaled moments of  $\Omega$  by construction. Importantly, this makes the process of “normalization” wholly disjoint from any manipulation of the objective function (see Section 3.1). After the optimal reconstruction of  $\widehat{\Omega}$  is found, we reverse this sequence of transforms to place it back in physical space.

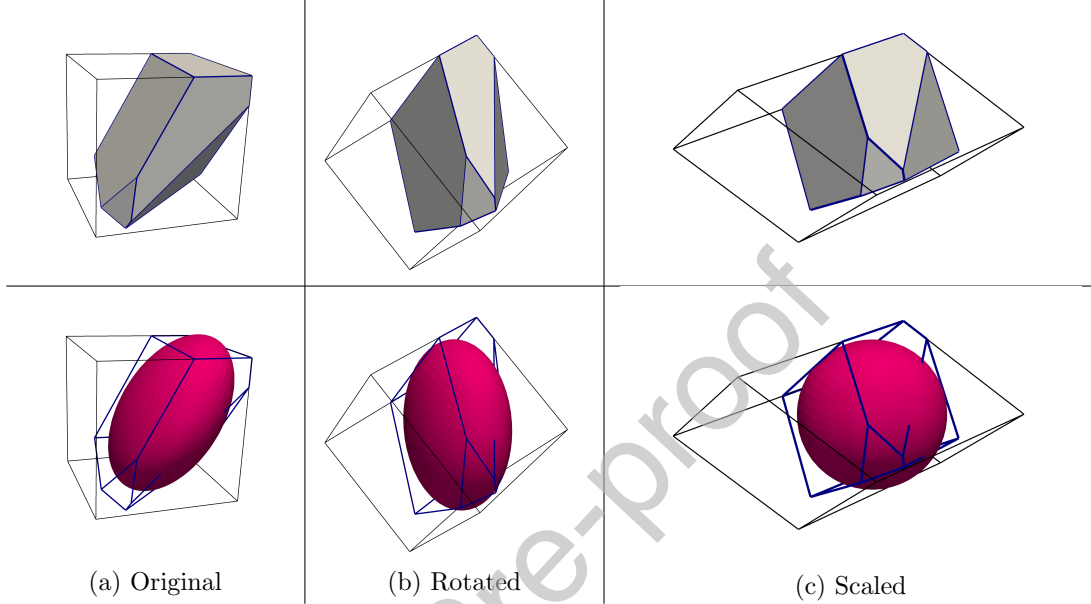


Figure 2: Example of orientation-aligned material normalization. Given the original material (a), we apply a rotation (b) and scaling (c) to the space so that optimization can be performed in a reference space for which the reference ellipsoid (red) for the true material (gray) is spherical.

### 3. MOF Algorithms Details

#### 3.1. Non-linear Optimization

We construct our objective function from the non-linear least-squares error in raw moments up to third order between the given reference moments and the actual moments of the reconstructed approximation to the material, but do so in reference space. In 3D, there are a total of 20 moments up to third order, which permits optimization over the 12 parameters of  $\text{MOF}_{4hs}^3$ , 15 parameters of  $\text{MOF}_{5hs}^3$ , or 18 parameters of

$\text{MOF}_{6hs}^3$  before the system becomes underdetermined. Each individual oriented plane  $\widehat{\mathbf{n}} \cdot \mathbf{x} = d$  is parameterized by angles  $\theta \in [0, 2\pi)$ ,  $\phi \in [0, \pi)$  that define normal  $\widehat{\mathbf{n}} = (\cos(\theta) \sin(\phi), \sin(\theta) \sin(\phi), \cos(\phi))$ , and the signed distance from the origin  $d$ . This provides us with the following objective function:

$$E(\Omega_r; \Omega) = \sum_{0 \leq i+j+k \leq 3} \left( \widehat{M}_{ijk}^r - \widehat{M}_{ijk} \right)^2 \quad (12)$$

with  $\widehat{M}_{ijk}^r$  representing the raw moments calculated for the reconstruction and  $\widehat{M}_{ijk}$  representing the reference moments scaled to the reference space. While we assume that reference moments will be provided via an outside application (i.e., hydrodynamic simulation), in this work we calculate reference moments for static tests from the intersection of high-polygon surface models and the cells of a mesh for demonstration purposes in this work. The computational cell  $C$  is also transformed into the same reference space, which we denote  $\widehat{C}$ .

We minimize this objective function using the Levenberg-Marquardt algorithm for solving unconstrained non-linear least-squares problems [42]. This algorithm is implemented in the C++ linear algebra library

Eigen [43], which computes gradients of the objective function using centered finite differences in lieu of an analytic representation. In general, the dampening term used in the Levenberg-Marquardt algorithm makes it quite robust to an ill-conditioned Jacobian relative to standard Gauss-Newton solvers, making it particularly well-suited to this problem where only approximate gradients are available, and spurious local minima are common and undesirable. The use of finite differences also means that the computational cost of optimization largely scales with the number of input parameters, as material moments up to third-order must be re-computed for changes in each. Doing so at each iteration, combined with the number of iterations necessary for convergence, requires the calculation of moments up to third-order be performed efficiently. For this reason we use the Interface Reconstruction Library (IRL), which provides routines for performing intersections between planes and general polyhedra using efficient half-edge data structures [44]. Fortunately, once the material polyhedron is computed by clipping the computational cell with each plane in the reconstruction, the cost of computing additional high-order moments is marginal.

The primary feature of the proposed method is to initialize this optimization procedure using the planes generated by inscribing a polyhedron in the reference ellipsoid, which is a unit sphere in reference space by construction. For  $\text{MOF}_{4hs}^3$  and  $\text{MOF}_{6hs}^3$  we use a single regular tetrahedron and cube respectively for this purpose, and for  $\text{MOF}_{5hs}^3$  we use a triangular prism. We show an example of each kind of inscribed polyhedron in Figure 3. The exact orientation of the initial polyhedron in the sphere is somewhat arbitrary, as the nature of the orientation-normalization means that no specific configuration is more likely to converge to the global minimum than any other. In the provided numerical experiments, we initialize the relevant polyhedron with the vertices in Table 1, defined in spherical coordinates.

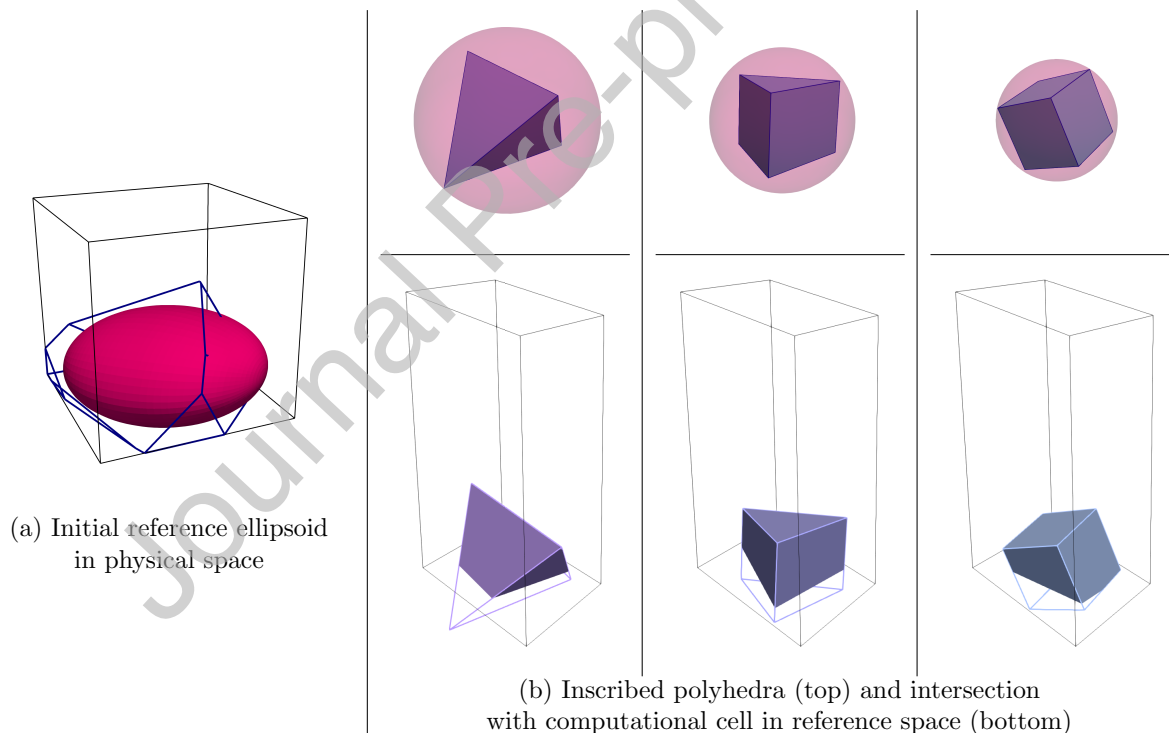


Figure 3: Example of initial inscribed polyhedra. We normalize reference moments so that the reference ellipsoid in physical space (a) is spherical in reference space. We then inscribe a single polyhedron with the corresponding number of faces (b) as the initial guess for optimization. Note that each sphere has been scaled in reference space so that the volume of the intersection of the inscribed shape with the cell is equal to the reference value.

Although the normalization scheme is also designed so that the reference ellipsoid (a sphere) has volume equal to the volume of the reference material, the inscribed polyhedron will naturally have a significantly smaller volume. As a result, we modify this shape prior to the initiation of the Levenberg-Marquardt algorithm so that the intersection of the inscribed polyhedron and  $\hat{C}$  has volume equal to that of the

Tetrahedron	Triangular Prism	Cube
$(\pi/2, \arctan(\sqrt{2}/2))$	$(0, \arctan(\sqrt{2}))$	$(0, \arctan(\sqrt{2}))$
$(\pi/2, \pi - \arctan(\sqrt{2}/2))$	$(2\pi/3, \arctan(\sqrt{2}))$	$(\pi/2, \arctan(\sqrt{2}))$
$(\pi - \arctan(\sqrt{2}/2), 3\pi/2)$	$(4\pi/3, \arctan(\sqrt{2}))$	$(\pi, \arctan(\sqrt{2}))$
$(\arctan(\sqrt{2}/2), 3\pi/2)$	$(0, \pi - \arctan(\sqrt{2}))$	$(3\pi/2, \arctan(\sqrt{2}))$
	$(2\pi/3, \pi - \arctan(\sqrt{2}))$	$(0, \pi - \arctan(\sqrt{2}))$
	$(4\pi/3, \pi - \arctan(\sqrt{2}))$	$(\pi/2, \pi - \arctan(\sqrt{2}))$
		$(\pi, \pi - \arctan(\sqrt{2}))$
		$(3\pi/2, \pi - \arctan(\sqrt{2}))$

Table 1: Spherical coordinates  $(u, v)$  on a unit sphere parameterized by  $S(u, v) = (\cos(u) \sin(v), \sin(u) \sin(v), \cos(v))$  for each initial inscribed polyhedron.

reference material  $\widehat{M}_{000}^r$ . This is an important preprocessing step for this specific implementation of Levenberg-Marquardt, which uses the initial error in each individual moment to define internal scaling factors. When this volume correction is not performed, the large difference in volume between the reference and approximate interface relative to the errors in first- and second- order moments (which are initially close to zero by construction) causes the algorithm to become unstable.

### 3.1.1. Customized Levenberg-Marquardt Algorithm

While the standard unconstrained Levenberg-Marquardt solver using the default parameters supplied by Eigen is usually sufficient, we empirically observe improved convergence to a global solution in numerical tests when we impose geometric constraints on the involved planes. For example, we restrict the angle parameters of each plane,  $\theta, \phi$ , to the intervals  $[0, 2\pi)$  and  $[0, \pi)$  respectively using the modulo operator after each step of optimization to avoid numerical blowup that can occur when computed gradients become unstable, as well as restricting the maximum step size used during the calculation of finite differences for such periodic parameters.

More impactful is our restriction of the distance parameter  $d$  for each plane. Ordinarily, it is possible for a plane to become entirely disconnected from the computational cell. When this happens, perturbations in the planes' parameters do not result in any changes to the computed moments of the interface, and so all local gradients are zero. This means that planes become "stuck" when the corresponding half-space has no overlap with the cell, and do not influence the material interface for all subsequent optimization iterations. As a result, we restrict the parameter  $d$  for each plane to an interval that ensures the plane intersects the computational cell.

To construct this interval for a given plane with unit normal  $\hat{\mathbf{n}}$ , we note the plane defined by  $\hat{\mathbf{n}} \cdot \mathbf{x} = \hat{\mathbf{n}} \cdot \mathbf{v} = d$  necessarily contains the point  $\mathbf{v}$ , and moves continuously with  $d$  along the normal direction  $\hat{\mathbf{n}}$ . Because the computational cell is polyhedral, we need only check its vertices  $\mathbf{v}$  to find the maximum and minimum values of  $d = \hat{\mathbf{n}} \cdot \mathbf{v}$  which still correspond to planes  $\hat{\mathbf{n}} \cdot \mathbf{x} = d$  that intersect the cell at least one point of the cell. Altogether, this means that for each plane, we must clamp the value of the distance  $d$  to the interval

$$d \in [\min\{\hat{\mathbf{n}} \cdot \mathbf{v}\}, \max\{\hat{\mathbf{n}} \cdot \mathbf{v}\}] \quad (13)$$

after each optimization step to ensure that at least one point in the computational cell still intersects the resulting plane. While this step does not completely remove the possibility of the converged-to interface involving fewer planes than are parameterized by the method, it does improve the overall stability and robustness of the optimization algorithm.

### 3.2. Enforcing the Volume Constraint

It is necessary that the volume of the reconstructed interface  $\Omega_r$  exactly matches that of the reference material  $\Omega$  in order for other applications of interest utilizing the reconstruction to be properly volume conserving. While the specific implementation of Levenberg-Marquardt in Eigen solves the unconstrained non-linear least squares problem without accounting for this restriction on the volume, we observe empirically that the volume of the reconstruction immediately produced is nevertheless very close to the reference volume. To correct this further, we follow the procedure in [40], which contracts or dilates all planes of the reconstruction simultaneously along their normal direction to directly enforce the volume constraint. Specifically, we begin with an initial set of distances  $\{d_i\}_{i=1}^N$  for each of  $N$  planes in the reconstruction, and find the scalar displacement  $\delta$  for which the convex intersection of the cell and the same planes with distances  $\{d_i + \delta\}_{i=1}^N$  has volume equal to the reference. Effectively, this becomes a non-linear root-finding problem that is solved with the Illinois method [45], a modified version of the Regula-Falsi method that is well suited for this application, as the non-linear function is monotonic, and an initial bracketing interval is readily available by considering the vertices of the cell (See Figure 4). We note that this necessarily introduces very small deviations in the remaining moments, but this effect is ultimately negligible to the overall quality of the reconstruction. It is for this reason that we do not apply the same constraint at every step in the optimization, as to do so would be unnecessarily computationally expensive.

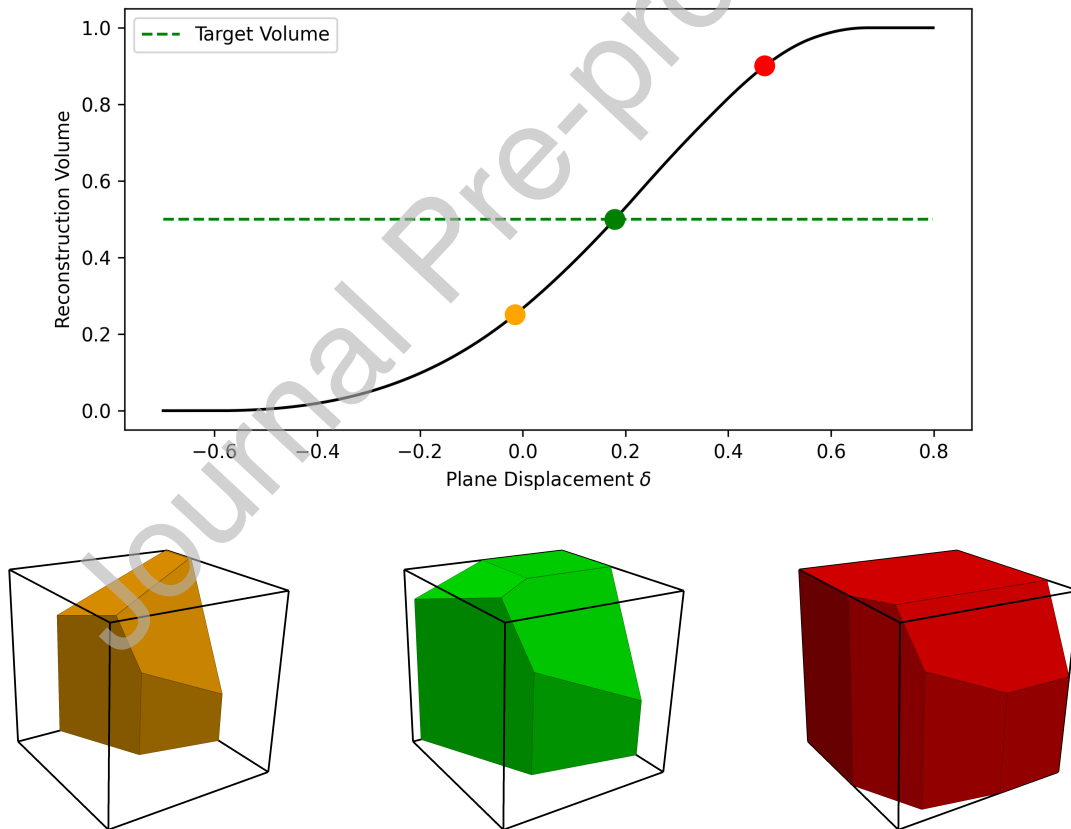


Figure 4: Example of root-finding problem for volume conservation. We plot the volume of captured material as a function of plane displacement parameter  $\delta$ , and show the corresponding interface for different values of  $\delta$ . We use the Illinois method to solve for the value of  $\delta$  which gives the target volume, as the volume varies monotonically with the  $\delta$ , and an initial bracketing interval is simple to compute.

### 3.3. Nonconvex Reconstruction

In this work, we have enforced that the material be equal to the convex intersection of the reconstruction planes, primarily for reasons of simplicity and efficiency in implementation. However, this can clearly only describe a subset of possible shapes. We extend this somewhat by allowing the half-spaces to represent the boundary of the *complement* of the convex intersection as well. While this too is only a second subset of possible shapes, their union is often sufficient for practical applications.

To the best of authors knowledge, there is no way of knowing *a priori* whether the underlying material is convex or not. To address this, we calculate two solutions for the non-linear optimization problem in each cell, one that assumes a convex reconstruction and one that assumes a reconstruction whose complement is convex. The first is performed using the original reference data  $M_{ijk}$ , while the second uses their complement in the unnormalized cell  $C$ , given by

$$M_{ijk}^c = M_{ijk}^C - M_{ijk}, \quad (14)$$

where  $M_{ijk}^C$  are the raw moments of the unnormalized cell  $C$ . We note that these two reconstructions could be performed in parallel.

The results of each optimization are stored and then compared. Importantly, the proposed method uses a different normalization for each set of moment data, and so it is necessary for the compared errors to each be considered in the original, unnormalized, physical space. If it is found that the convex-complement method more accurately captures the reference moment data, then an internal flag is set to denote for downstream applications that utilize the reconstruction. Often, such applications will require decomposition of the non-convex material into a union of convex shapes, which we demonstrate in Figure 11.

## 4. Summary of the MOF Algorithm

We summarize the main steps of the proposed family of MOF interface reconstruction algorithms below. Given a computational cell, its reference moments, and the number of half-spaces  $N$  used in the MOF $_{Nhs}^3$  reconstruction, perform the following steps:

1. Normalize the reference moments so that the reference ellipsoid is spherical (Section 2.2)
2. Construct an initial polyhedron inscribed in the sphere (Table 1)
3. Enforce that the polyhedron has the same volume in the cell as the normalized reference data (Section 3.2)
4. Apply the Levenberg-Marquardt algorithm on the planes defined by the polyhedron (Section 3.1)
5. Again, enforce that the reconstruction has the same volume in the cell as the normalized reference data (Section 3.2)
6. Repeat the above steps for the complement of the reference moment data (Section 3.3)
7. Evaluate the error for each reconstruction against their respective reference moment data:
  - If the first error is smaller, return the intersection of the half-spaces in the first reconstruction.
  - If the second error is smaller, return the *complement* of the intersection of the half-spaces in the second reconstruction as the union of a set of convex polyhedra.

## 5. Results

### 5.1. Single Cell Tests

We begin by demonstrating this method on simple single-cell examples for known shapes where an exact reconstruction is possible. In such cases where the target shape is known, we can use an error metric derived from the *symmetric difference*, which is defined for arbitrary sets as the total region contained in either set, but not their intersection. In the context of material interface reconstruction, this symmetric difference error metric measures the sum of the volume of the true material that our approximation has

failed to capture, and the volume of the approximation that does not capture the true material. Written mathematically, we have

$$S(\Omega_r; \Omega) := \int_{\Omega_r \setminus \Omega} dV + \int_{\Omega \setminus \Omega_r} dV. \quad (15)$$

In practice, however, we only have access to the moment-derived least-squares error as an error metric for the reconstruction  $E(\Omega_r; \Omega)$  (see Section 3.1).

### 5.1.1. Exact Reconstructions

We show in Figure 5 the “vertex-corner” shape introduced in Section 1. The computational cell is taken to be the cube  $[-0.5, 0.5]^3$ , with the three planes with normals and distances defined by

$$\begin{aligned} \hat{n}_1 &= (1.0, -0.1, -0.1), & d_1 &= 0.0, \\ \hat{n}_2 &= (-0.1, 1.0, -0.1), & d_2 &= 0.0, \\ \hat{n}_3 &= (-0.1, -0.1, 1.0), & d_3 &= 0.0. \end{aligned} \quad (16)$$

As expected,  $\text{MOF}_{4hs}^3$  is able to perfectly recreate this shape, both exactly matching each of the 20 reference moments and resulting in zero symmetric difference error, shown in Table 2. We note that although  $\text{MOF}_{4hs}^3$  is in some sense over-parameterized for this geometry, as we allow the placement of four planes when three is sufficient to achieve an exact reconstruction, we are able to exactly reconstruct the shape. In this example, the non-active plane is tangent to the resulting material polyhedron at its optimum placement, but does not influence its shape. Similar results are observed with  $\text{MOF}_{5hs}^3$  and  $\text{MOF}_{6hs}^3$ , but with more redundant planes in the final reconstruction. For an additional point of comparison, we make reference to a simple one-plane MOF method utilizing third-order moments. We initialize this  $\text{MOF}_{1hs}^3$  method with a plane whose normal points from the material centroid to the cell centroid, a reasonably effective initial guess for such a method. As shown in Figure 5, this provides a poor approximation of the geometry, and will perform even worse in following examples, where even more portions of the shape are detached from the cell boundary.

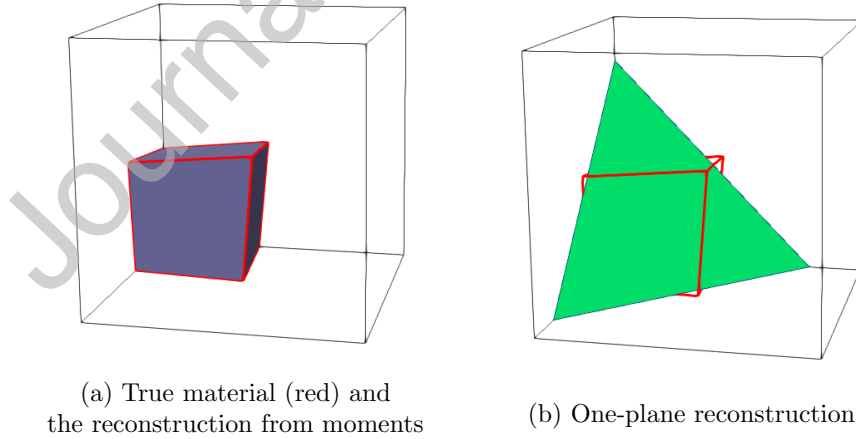


Figure 5: A comparison between the proposed multi-plane MOF method (a) and a single-plane MOF method (b) on a corner feature. We exactly reconstruct the “vertex-corner” using three of the four planes used in  $\text{MOF}_{4hs}^3$ .

We perform a similar test on the same computational cell for the “tip” shape in Figure 6, defined by the

Moment	MOF <sub>1hs</sub> <sup>3</sup>	MOF <sub>4hs</sub> <sup>3</sup>
$M_{000}$	$-9.57776 \times 10^{-13}$	$-1.38778 \times 10^{-17}$
$M_{100}$	$-0.00264148$	$6.93889 \times 10^{-18}$
$M_{010}$	$-0.00293966$	$1.73472 \times 10^{-17}$
$M_{001}$	$-0.00232491$	$1.38778 \times 10^{-17}$
$M_{200}$	$0.00218141$	$-1.73472 \times 10^{-18}$
$M_{110}$	$0.000683377$	$-5.20417 \times 10^{-18}$
$M_{101}$	$0.000475685$	$-3.46945 \times 10^{-18}$
$M_{020}$	$0.00228779$	$-8.67362 \times 10^{-18}$
$M_{011}$	$0.000575576$	$-5.20417 \times 10^{-18}$
$M_{002}$	$0.00207197$	$-6.93889 \times 10^{-18}$
$M_{300}$	$-0.00107547$	$-2.1684 \times 10^{-18}$
$M_{210}$	$-0.000503406$	$8.67362 \times 10^{-19}$
$M_{201}$	$-0.000422462$	0
$M_{102}$	$-0.0004949$	$1.30104 \times 10^{-18}$
$M_{111}$	$0.000259171$	$1.30104 \times 10^{-18}$
$M_{102}$	$-0.00043277$	$8.67362 \times 10^{-19}$
$M_{030}$	$-0.00112707$	$-4.33681 \times 10^{-19}$
$M_{021}$	$-0.000452824$	$1.30104 \times 10^{-18}$
$M_{012}$	$-0.000471645$	$8.67362 \times 10^{-19}$
$M_{003}$	$-0.00102168$	0
$E(\Omega_r; \Omega)$	0.0574	$1.923 \times 10^{-17}$
Symmetric Difference	0.645554	$2.498 \times 10^{-15}$

Table 2: Per-moment errors between reference data and the reconstructed interface for the “vertex-corner” example in Figure 5.

$$\begin{aligned}
& \text{planes} \\
& \hat{n}_1 = (-1.0, 0.0, 0.2), \quad d_1 = 0.1, \\
& \hat{n}_2 = (\sqrt{2}, \sqrt{2}, 0.2), \quad d_2 = 0.0, \\
& \hat{n}_3 = (\sqrt{2}, \sqrt{2}, 0.2), \quad d_3 = 0.0.
\end{aligned} \tag{17}$$

Again, MOF<sub>4hs</sub><sup>3</sup> achieves an exact reconstruction while a one-plane method necessarily makes the shape unrecognizable to maintain the volume constraint.

To further test the robustness of our method, we repeat this test in Figure 7 using the same “tip” shape, but occupying a much smaller volume in the cell. Specifically, we translate the cell along the positive  $z$  direction to occupy the  $[-0.5, 0.5] \times [-0.5, 0.5] \times [0.2986, 1.2986]$  so that the section of filament that intersects with the cell has the same geometry, but a calculated volume of  $3.19117 \times 10^{-14}$ . The resulting tetrahedron has coordinates at

$$\begin{aligned}
p_1 &= (-0.04226039, 3.1973229 \times 10^{-5}, 0.298600) \\
p_2 &= (-0.04226039, -3.1973229 \times 10^{-5}, 0.298600) \\
p_3 &= (-0.042228417, 0, 0.298600) \\
p_4 &= (-0.042241661, 0, 0.29869365)
\end{aligned} \tag{18}$$

Despite needing to handle a volume that is close to machine precision, our family of methods is nevertheless able to match the shape exactly. This is largely due to our unique normalization procedure, which causes the method to be extremely robust to the size of the material. As a result of this normalization, the optimization performed previous example and this one are near-identical in reference space, with the only difference being a larger normalized cell for the smaller material.

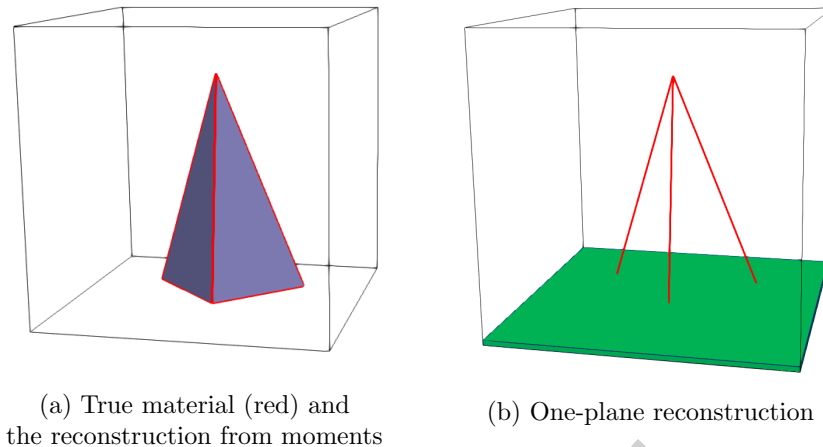


Figure 6: A comparison between the proposed multi-plane MOF method (a) and a single-plane MOF method (b) on a “tip” feature. We exactly reconstruct the shape using three of the four planes in  $\text{MOF}_{4hs}^3$ .

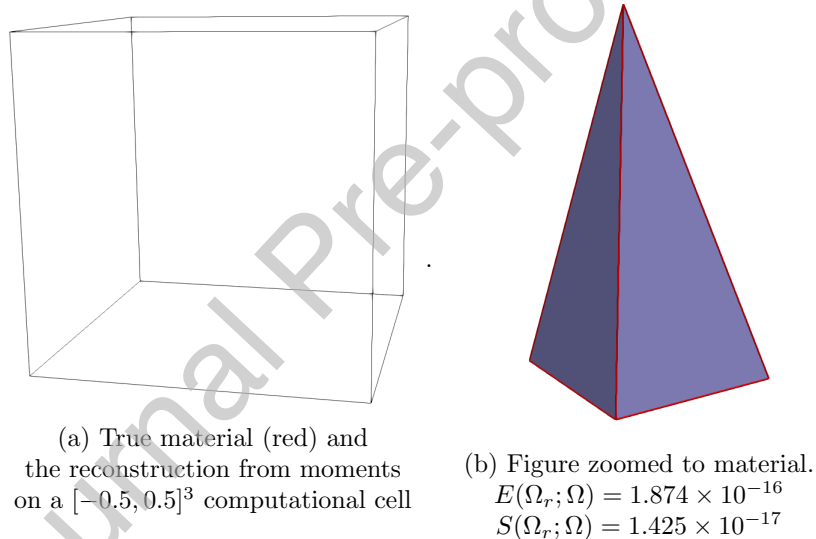


Figure 7: Robustness test for small material volume. Note that relative to the cell, the true material (a) is so small as to occupy only a single pixel in this rendering. Nevertheless, the proposed method achieves an exact reconstruction (b).

Finally, we consider in Figure 8 an example that is even more adversarial for  $\text{MOF}_{1hs}^3$ , in which the material is entirely disconnected from the cell boundary. As before, the proposed  $\text{MOF}_{4hs}^3$  method is able to accurately recover the shape, while having, in some sense, the maximum error possible with respect to the symmetric difference.

### 5.1.2. Inexact Reconstructions

In addition to the sharp target features depicted in Figure 1, our method is also capable of approximating shapes with smooth features, which no planar reconstruction can recover exactly. We perform these tests using the  $\text{MOF}_{6hs}^3$  method as a representative of the proposed family.

In the following examples, we consider the intersection of our computational cell with ellipsoids of varying size, position, and orientation. To obtain the reference moments, we construct a polyhedron that approximates the geometry of the curved object, from which we can compute the moments of its intersection with the cell using the clipping methods from the IRL software package described in

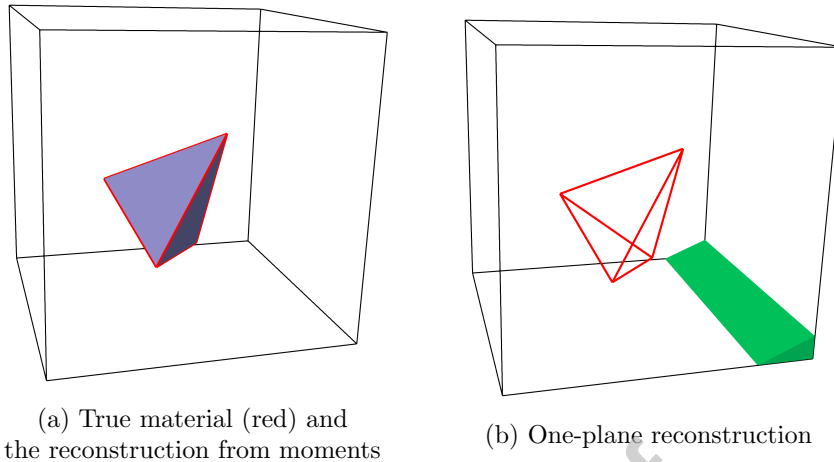


Figure 8: A comparison between the proposed multi-plane MOF method (a) and a single-plane MOF method (b) on material that is fully embedded within the computational cell. We exactly reconstruct the shape using all four planes in  $\text{MOF}_{4hs}^3$ .

	$a$	$b$	$c$	$\alpha$	$\beta$	$\gamma$	$x_0$	$y_0$	$z_0$
Ellipsoid 1	0.32	0.40	0.98	4.9	-1.4	3.3	-0.10	0.34	0.20
Ellipsoid 2	0.35	0.23	0.58	5.1	2.3	4.0	0.37	0.41	0.59
Ellipsoid 3	0.48	0.60	0.57	4.7	4.0	1.7	0.10	-0.18	0.48

Table 3: Parameters for single-cell ellipsoid tests in Figure 9.

Section 3.1. While these do not exactly match the moments of a true truncated ellipsoid, the symmetric difference errors are computed with respect to this approximating polyhedron, and not the ground truth geometry.

In Figure 9, we consider three such ellipsoids, and define each for reproducibility with parameters defined by the values in Table 3. Numerically, we parameterize these ellipsoids through the lengths of their three axes  $a, b, c$  (parallel to the  $x$ -,  $y$ -, and  $z$ -axis respectively) and the rotation matrix  $R(\alpha, \beta, \gamma)$  defined in Equation 5, followed by a shift to have the center  $(x_0, y_0, z_0)$ .

As seen in Figure 9, we see qualitative agreement with the ground truth material. Although an exact reconstruction cannot possibly be achieved for such smooth shapes, our method has correctly recovered the geometric features of each shape, an observation which is supported by the computed error metrics on each. In particular, each reconstruction intersects the same cell faces as the ground-truth ellipsoid, which is not possible with only a single plane.

We can also use these smooth shapes to illustrate differences between the  $\text{MOF}_{4hs}^3$ ,  $\text{MOF}_{5hs}^3$ , and  $\text{MOF}_{6hs}^3$  methods. Because an exact recovery is not possible, there is an important trade-off between computational cost and the potential accuracy of the optimized interface. As a representative example, we see in Figure 10 that increasing the number of planes in the reconstruction naturally improves the error, but also the total cost of the method. It is also worth noting that while the cost of each method scales directly with the number of planes in the reconstruction, the complexity of the polyhedral cell also has a considerable impact on the performance of the method, as additional cell faces each have the same performance impact on the moment calculation as an additional reconstruction plane. While any of these three methods naturally incur a greater computational cost than the analogous one-plane strategy, this cost is largely compensated for by the ability to accurately represent complex shapes on coarse meshes. As a particularly illustrative example, consider the shape in Figure 5, for which our multi-plane methods are able to achieve a level of detail that is not simply expensive for a one-plane method to reproduce, but entirely impossible within the same computational cell.

We note here that with any of the proposed methods, there are scenarios for which it is more difficult for the optimized interface to represent the true, global minimum of the corresponding minimization problem.

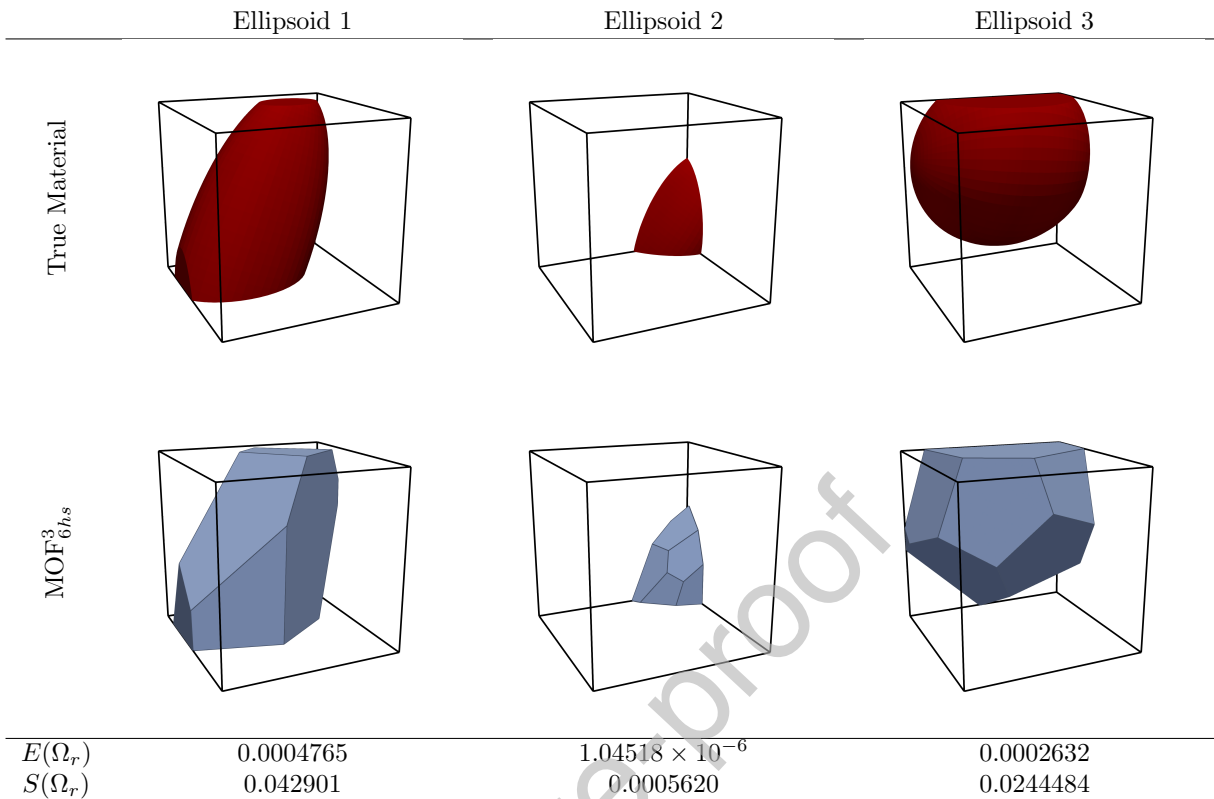


Figure 9: Example reconstructions on curved shapes. We test MOF<sub>6hs</sub><sup>3</sup> on shapes for which an exact reconstruction is not possible. Despite this, the reconstructed interface accurately captures important geometric features.

Empirically, we find this most often to be the case when the ground-truth shape has *only* flat features. For example, we observe in Ellipsoid 2 of Figure 9 that the optimized interface found by our method is defined by five planes. However, it is not inconceivable that the global minimum of the six-plane optimization problem solved by the MOF<sub>6hs</sub><sup>3</sup> method fully utilizes all six planes. Because this and other predominantly flat shapes are reasonably well-approximated by even a single plane, the relevant optimization problem is subject to additional local minima, many of which utilize fewer planes in the reconstruction than are available. Importantly, the already simple nature of the ground truth shape means that the resulting error is comparable to that of the unreachable global minimum, even if the interface itself has less complexity, and certainly outperforms any one-plane alternative.

Finally, we can also test our method on the *complement* of the third ellipsoid to demonstrate the ability of the proposed methods to reproduce non-convex shapes. We plot in Figure 11 both the trial convex reconstruction and convex-complement reconstruction, observing that because the convex-complement reconstruction has a lower moment error, it is accepted as the final reconstruction. Importantly, because the moment data over which the optimization is performed is identical between this case and that demonstrated for Ellipsoid 3 in Figure 9, the material approximations returned by the two methods are themselves complements. It is for this reason that we primarily restrict our numerical tests to convex material, as the problem of reconstructing materials whose complement is convex in the cell is *equivalent* to the problem of reconstruction convex material.

## 5.2. Mesh Tests

As would be the case in a practical application, it is necessary that our method is performant when considered across *all* computational cells in a mesh. We note that the local nature of any moment-of-fluid method means that accuracy across an entire mesh be equivalent to accuracy across each individual cell. However, in addition to providing a wide breadth of single-cell examples, we evaluate the proposed method

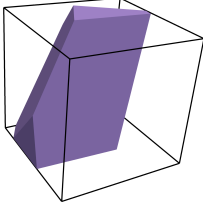
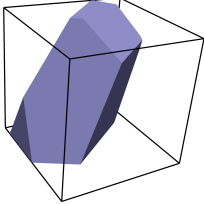
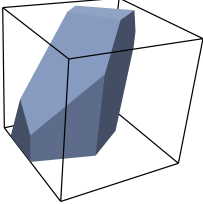
	MOF <sub>4hs</sub> <sup>3</sup>	MOF <sub>5hs</sub> <sup>3</sup>	MOF <sub>6hs</sub> <sup>3</sup>
Reconstructed Ellipsoid			
$E(\Omega_r)$	0.00188321	0.00102675	0.000476561
$S(\Omega_r)$	0.0675061	0.0525182	0.0429109
Total Time (s)	0.0805	0.174	0.203
Iteration Count	20	33	21
Time per Iteration (s)	0.004025	0.00527	0.00967

Figure 10: Comparison of accuracy vs. wall-clock run-time among multi-plane MOF methods. In cases where an exact reconstruction is not possible, the use of additional planes results in an approximated interface that is more computationally expensive, but can result in considerably lower error.

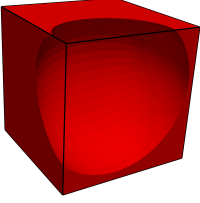
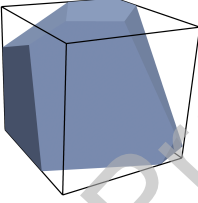
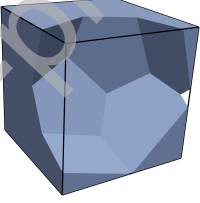
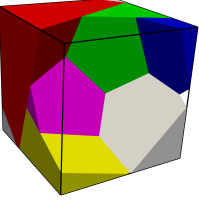
True Material	Convex MOF <sub>6hs</sub> <sup>3</sup>	Non-convex MOF <sub>6hs</sub> <sup>3</sup>	Decomposition into Convex Polyhedra
			
$E(\Omega_r)$	0.0197032	0.000237875	
$S(\Omega_r)$	0.190205	0.0233722	

Figure 11: Example of reconstruction of non-convex material. Because the interface generated by the complement of a convex shape produces lower moment error, it is returned as the final interface. We also demonstrate the decomposition of a reconstructed non-convex material into convex polyhedra for related applications utilizing the reconstruction.

across an entire mesh to demonstrate the quantitative advantages of a multi-plane reconstruction scheme.

### 5.2.1. Cube Test

In our first example in Figure 12, we consider MOF<sub>1hs</sub><sup>3</sup> and MOF<sub>4hs</sub><sup>3</sup> on the cube  $[-0.5, 0.5]^3$  that we have rotated by  $R(1.2, -0.82, 1.0)$  around the origin (See Equation 5) and then translated by  $(0.05, 0.10, -0.05)$  to avoid unfairly exploiting the shape's symmetry. This shape is then reconstructed on an  $N \times N \times N$  uniform grid of cells across the larger cube  $[-1, 1]^3$ . Although the MOF<sub>1hs</sub><sup>3</sup> method used for comparison does approach the ground truth with successive refinement, the nature of the shape's corner features makes it impossible to represent the shape exactly. Furthermore, the orientation of the shape makes it impossible for a PLIC method recover its edges exactly at *any* refinement level on an axis-aligned grid, even in the presence of a more sophisticated, adaptive refinement scheme.

In these examples, we consider in Table 4 the maximum of both the moment error  $E(\Omega_r; \Omega)$  and the symmetric difference error  $S(\Omega_r; \Omega)$  across all cells in the mesh. We also consider their average across all mixed cells in the mesh, that is to say, cells which are neither empty nor completely filled by the material. In Figure 12 we see that even at the coarsest level of refinement, our MOF<sub>4hs</sub><sup>3</sup> method visually recovers the shape near-identically, in contrast to the unrecognizable geometry produced MOF<sub>1hs</sub><sup>3</sup> method. These results are supported quantitatively by Table 4, which shows that we indeed produce an interface with considerably less error along each of the described metrics.

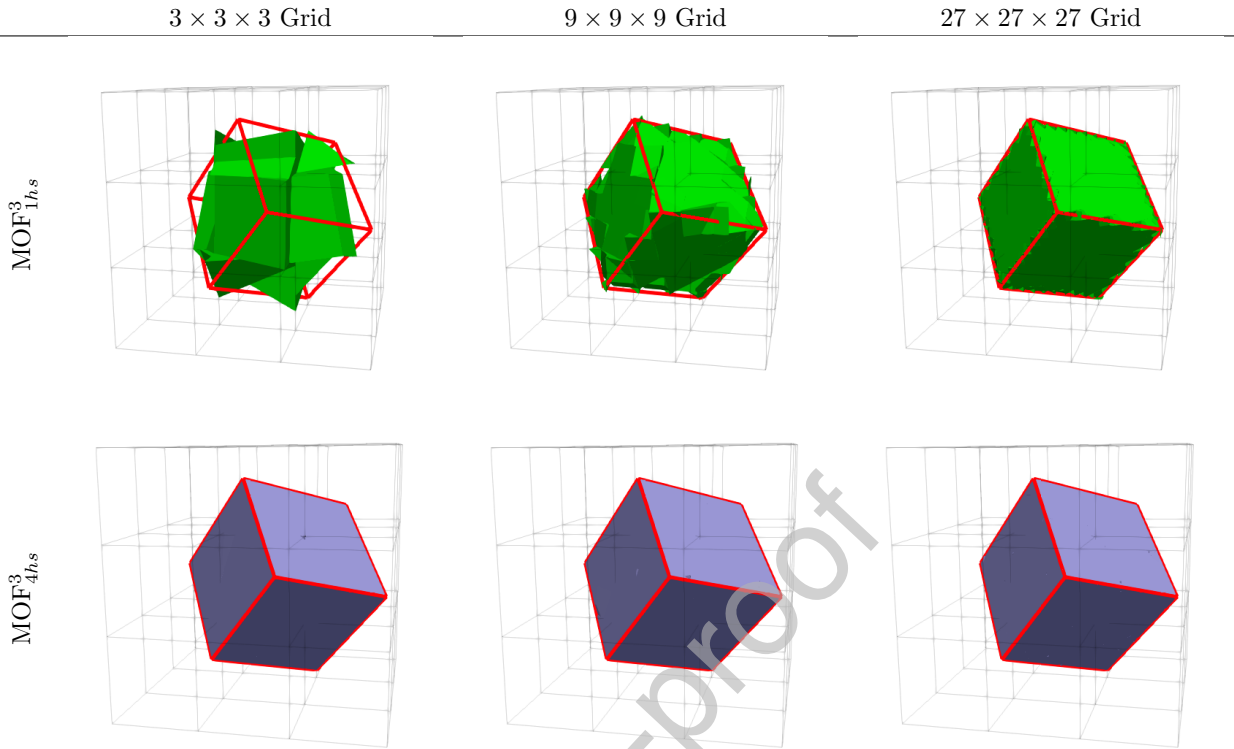


Figure 12: Comparison between single-plane and multi-plane reconstruction schemes for a cube across a Cartesian mesh. Using a multi-plane reconstruction scheme allows for more accurate reproductions on coarser grids. The true material is outlined in red. The relevant error metrics are presented in Table 4.

### 5.2.2. Ellipsoid Test

To consider an example that better resembles practical application, we apply  $\text{MOF}_{4hs}^3$ ,  $\text{MOF}_{5hs}^3$ , and  $\text{MOF}_{6hs}^3$  on a convex shape that cannot be recovered exactly in Figure 13, which depicts an ellipsoid defined by the parameters in Table 5. As in the previous example, we consider an orientation of the shape that is not aligned with the grid of cells, which are placed in the cube  $[-1, 1]^3$ .

We see in this example that the multi-plane reconstruction offered by each of  $\text{MOF}_{4hs}^3$ ,  $\text{MOF}_{5hs}^3$ , and  $\text{MOF}_{6hs}^3$  each offer a dramatic improvement over the single-plane technique, adequately recovering the primary geometric features even at remarkably low resolution grids. Even in the case of a single-cell grid, we can see that our method is able to recover a crude approximation of the completely embedded material, which is impossible when only a single plane is used in the reconstruction.

Additionally, considering reconstruction accuracy across an entire mesh allows for a more comprehensive analysis of how grid refinement impacts each of the proposed methods. As one would expect, Figure 13 demonstrates that the reconstruction is more accurate both with increasing levels of grid refinement and increasing numbers of planes used in the representation. However, we also observe that it is for the coarser grids that increasing the number of planes brings about the most improvement. For example,  $\text{MOF}_{6hs}^3$  offers a dramatic improvement over  $\text{MOF}_{1hs}^3$ , and even over  $\text{MOF}_{4hs}^3$ , on the  $2 \times 2 \times 2$  grid, being clearly recognizable as depicting the original ellipsoid shape. On the other hand, the different methods visually perform very similarly on a  $32 \times 32 \times 32$  grid. This is because as the cell-size decreases, the local geometry of the shape becomes increasingly flat, and the POM in each cell is already well approximated by a single plane. We note, though, that this is not always the case, as shapes with sharp features (as in Figure 12) cannot be well approximated by a PLIC method at any level of refinement for an arbitrary mesh.

We further explore this point on the same ellipsoid shape by considering reconstruction accuracy as function of cell index through the heatmaps in Figure 14, where the two error metrics are evaluated over a  $16 \times 16 \times 16$  grid. We note that although there is a considerable amount of cell-to-cell variability for each




Figure 13: Comparison between single-plane and several multi-plane MOF reconstruction schemes for an ellipsoid over a Cartesian mesh. Using a multi-plane reconstruction scheme allows for more accurate reproductions on coarser grids. The target ellipsoid is suggested in red wireframe on the first row, although the shape from which ground truth moments are computed is of considerably higher resolution.

		$E(\Omega_r; \Omega)$		
Grid Resolution (# Mixed Cells)		$3 \times 3 \times 3$ (23)	$9 \times 9 \times 9$ (189)	$27 \times 27 \times 27$ (1752)
MOF <sub>1hs</sub> <sup>1</sup>	Cell Average	1.3284e-03	6.0892e-06	2.9583e-08
	Maximum	5.4000e-03	8.3435e-05	1.2974e-06
MOF <sub>4hs</sub> <sup>1</sup>	Cell Average	2.0053e-07	2.4080e-07	1.1596e-10
	Maximum	4.3996e-06	1.2727e-07	5.0307e-08

		$S(\Omega_r; \Omega)$		
Grid Resolution (# Mixed Cells)		$3 \times 3 \times 3$ (23)	$9 \times 9 \times 9$ (189)	$27 \times 27 \times 27$ (1752)
MOF <sub>1hs</sub> <sup>1</sup>	Cell Average	4.1915e-02	1.8975e-04	2.5371e-06
	Maximum	1.2922e-02	1.8596e-03	6.7378e-05
MOF <sub>4hs</sub> <sup>1</sup>	Cell Average	1.0440e-05	7.2676e-06	1.9463e-08
	Maximum	2.1353e-04	1.2985e-05	9.7976e-06

Table 4: Evaluation of our moment error metric  $E(\Omega_r; \Omega)$  and the symmetric difference error metric  $S(\Omega_r; \Omega)$  for the cube example in Figure 12. We see that the MOF<sub>4hs</sub><sup>3</sup> method results in a significantly error across both metrics.

	$a$	$b$	$c$	$\alpha$	$\beta$	$\gamma$	$x_0$	$y_0$	$z_0$
Ellipsoid	0.1	-0.1	0.1	0.9	-1.07	0.7	3.30719	0.18485	4.93365

Table 5: Parameters for multi-cell ellipsoid test in Figures 13, 14, and 15

error metric within a single shape, this can be attributed to varying material volume in each cell. Nevertheless, cell-to-cell comparisons between the different methods are appropriate, as the interface in each cell's reconstruction is attempting to match the same data.

We observe in Figure 14(a) that the accuracy of the MOF<sub>1hs</sub><sup>3</sup> method, both in terms of moment error  $E(\Omega_r; \Omega)$  and the symmetric difference error  $S(\Omega_r; \Omega)$ , generally correlates with the curvature of the shape, where areas of high curvature tend to have higher reconstruction error. At the same time, the proposed multi-plane methods are far more robust to such qualities and, as seen in Figure 13, can be observed to outperform MOF<sub>1hs</sub><sup>3</sup> most in regions of high curvature. We consider the overall distribution of these errors more directly through histograms over the same data in Figure 14(b), where we observe that we see a broad improvement in error across the domain of each of our methods against MOF<sub>1hs</sub><sup>3</sup>. At the same time, we observe through in this example that the difference between the MOF<sub>4hs</sub><sup>3</sup>, MOF<sub>5hs</sub><sup>3</sup>, and MOF<sub>6hs</sub><sup>3</sup> methods is somewhat minimal. This is largely due to the resolution of the grid, as the curvature across a single cell within  $16 \times 16 \times 16$  grid is (on average) not high enough to see strong visual improvement between the three methods. On the other hand, we can repeat this analysis for a coarser  $4 \times 4 \times 4$  grid in Figure 15 and see more clearly the advantage of using MOF<sub>6hs</sub><sup>3</sup> over MOF<sub>5hs</sub><sup>3</sup> and MOF<sub>4hs</sub><sup>3</sup>.

### 5.2.3. Nonconvex Test

Next, we consider the proposed family of methods on a shape with many non-convex features in the form of five "creases" positioned around an ellipsoid shape. We generate this shape according to the following parametric surface:

$$\begin{aligned}
 x(u, v) &= 0.04 (1.5 \sin(u)(6 \cos(v) - \cos(6v)), \\
 y(u, v) &= 0.04 (1.5 \sin(u)(6 \sin(v) - \sin(6v)), \\
 z(u, v) &= 0.24 \cos(u).
 \end{aligned} \tag{19}$$

The shape generated by this surface is then rotated around the origin by the rotation matrix  $R(-0.1, -1.67, 0.0)$  and translated by the point  $(0.2, 0.2, -0.2)$ . We see as before in Figure 16 that proposed family of methods can far more accurately capture subgrid and curved features of the shape, and that successive refinement of the mesh improves the quality of this reconstruction relative to the ground truth.

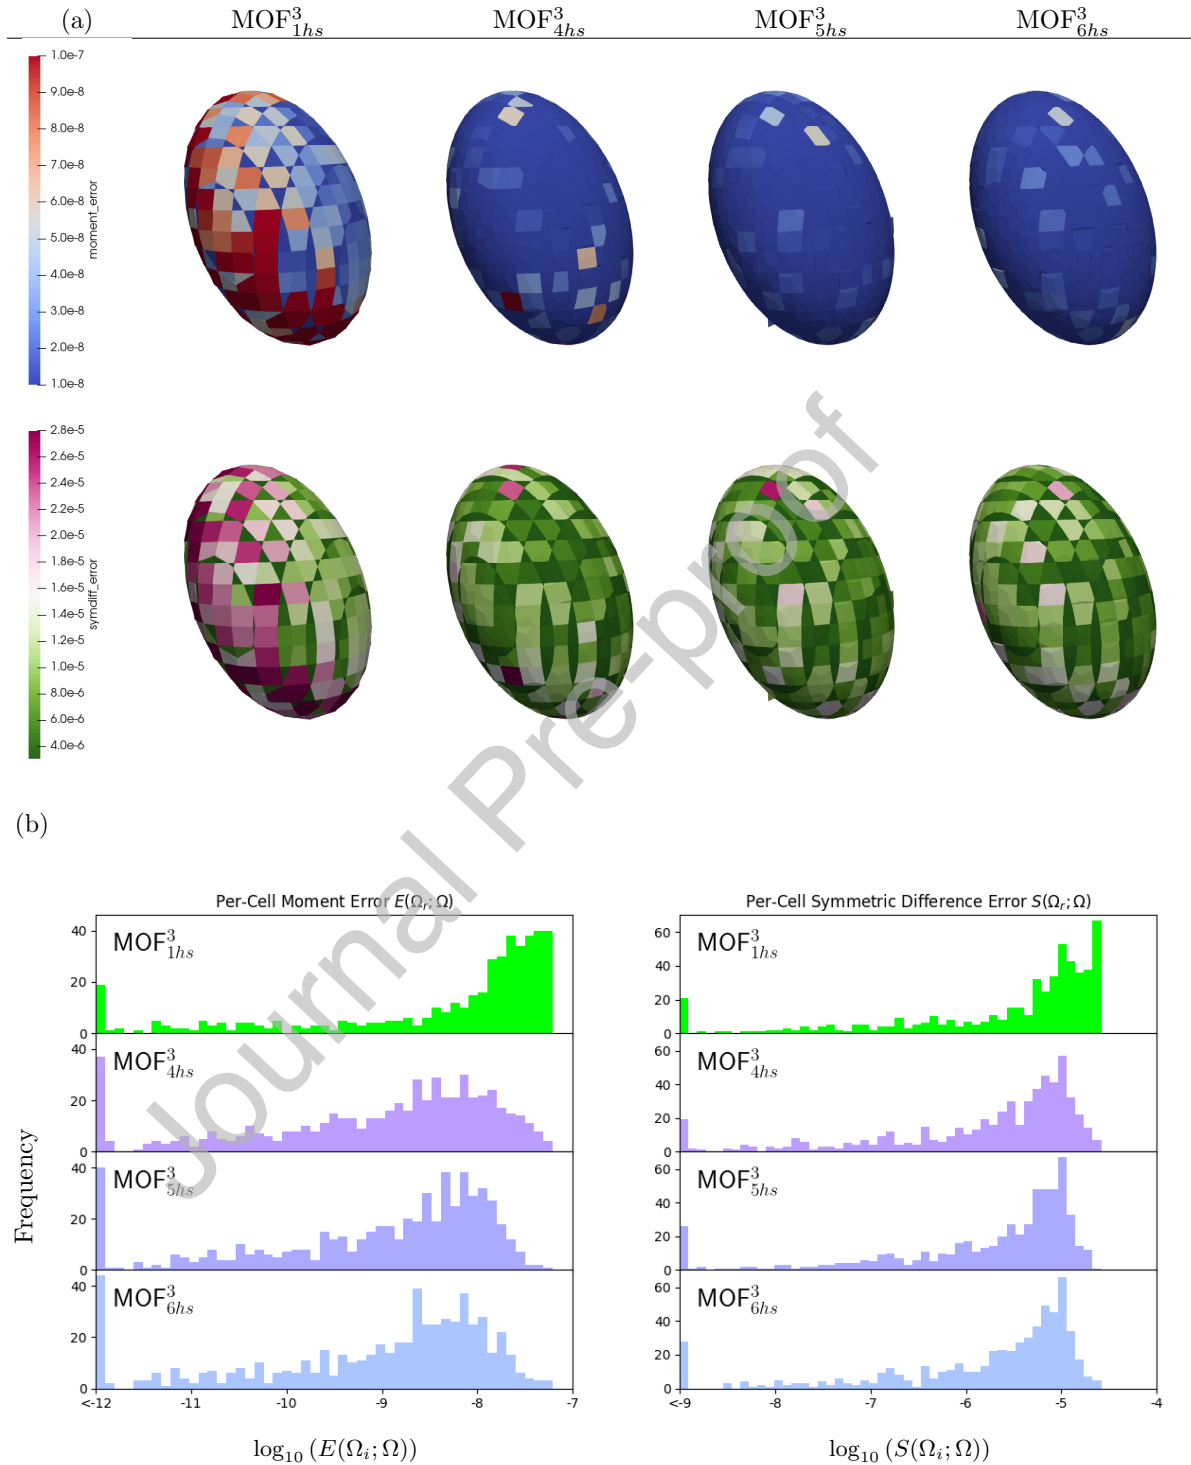
Per-Cell Errors on  $16 \times 16 \times 16$  Mesh Grid

Figure 14: Error comparison between methods evaluated for an ellipsoid over a Cartesian mesh. Although  $\text{MOF}_{1hs}^3$  can provide a sufficient representation for regions of low curvature, the use of a multi-plane scheme makes the accuracy of the reconstruction more tolerant to regions of higher curvature. Even still, the proposed family of multi-plane schemes still improves accuracy over one-plane alternatives when considered across the entire shape.

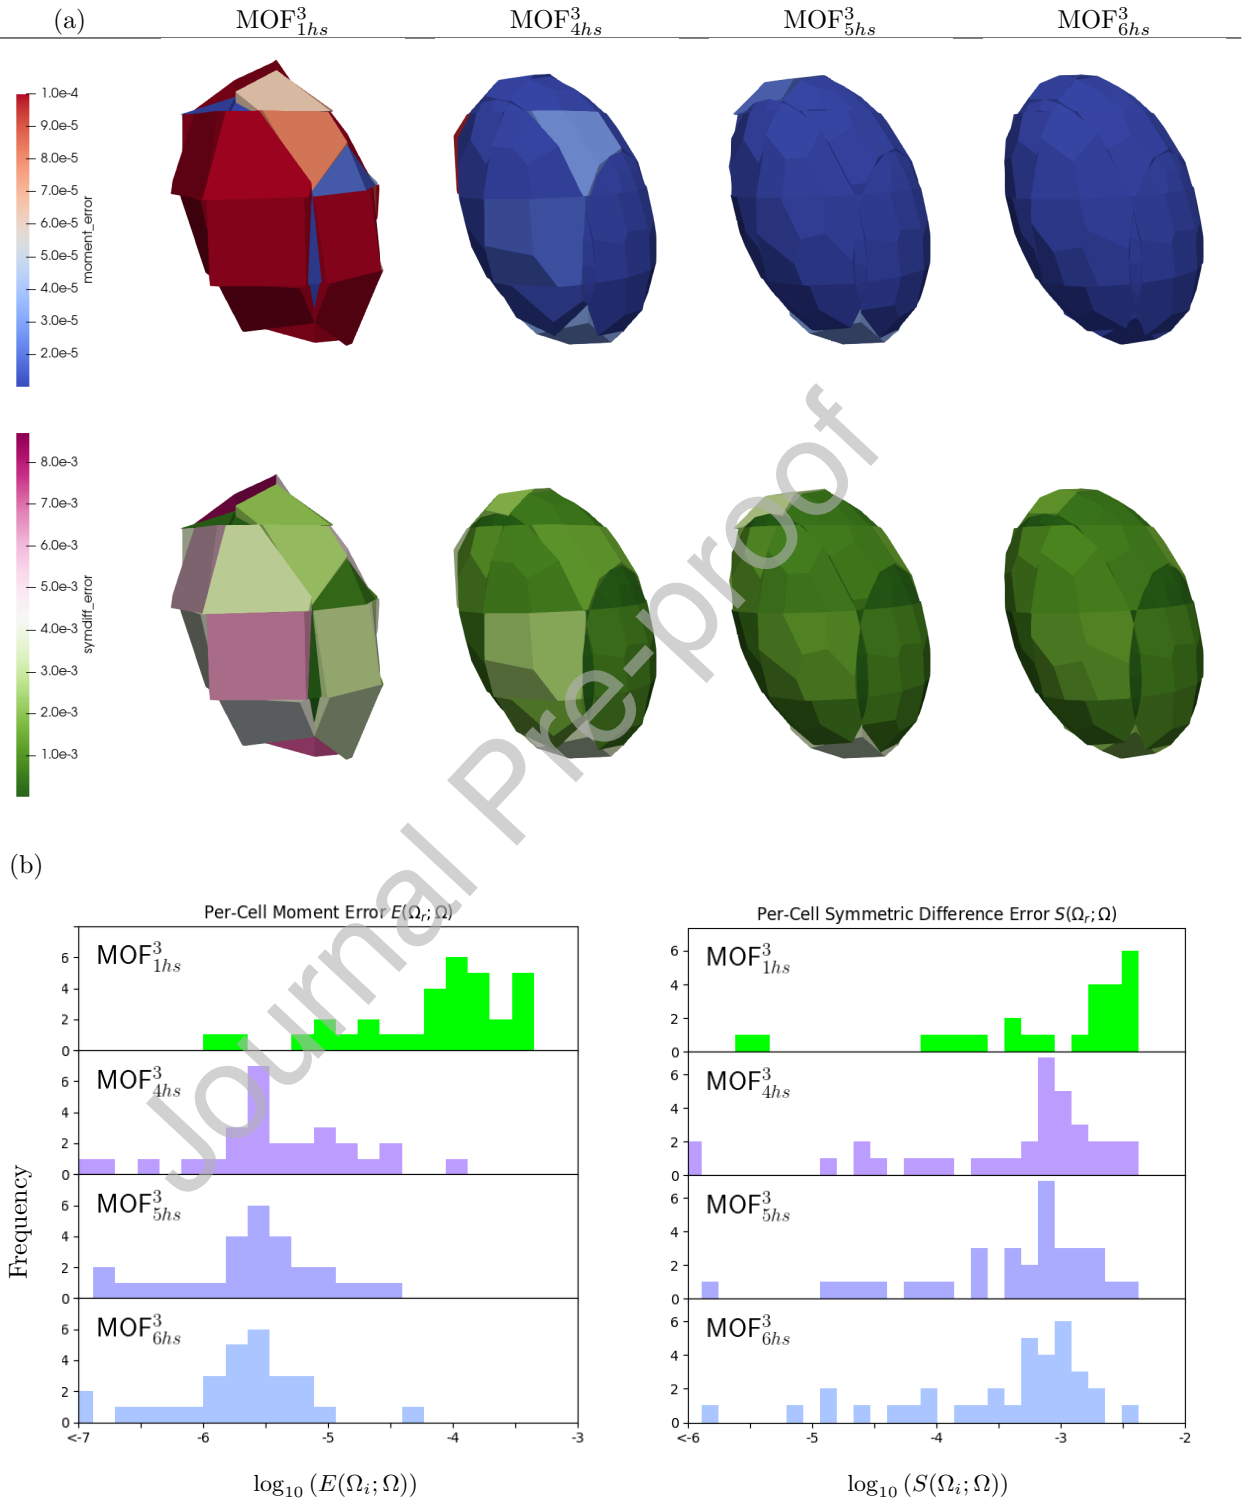
Per-Cell Errors on  $4 \times 4 \times 4$  Mesh Grid

Figure 15: The results of Figure 14 reproduced on a coarser grid. We can see more clearly the advantages of using more planes per cell in the reconstruction, as  $\text{MOF}_{5hs}^3$  and  $\text{MOF}_{6hs}^3$  are shown to be significantly more performant than  $\text{MOF}_{4hs}^3$ .




Figure 16: Comparison between reconstruction methods on a non-convex shape over a Cartesian mesh. Using a multi-plane reconstruction scheme permits reconstruction of non-convex regions within a single cell. The target shape is suggested in red wireframe on the first row, although the shape from which ground truth moments are computed is of considerably higher resolution. For the  $4 \times 4 \times 4$  grid, we also emphasize in red the  $\text{MOF}_{1hs}^3$  and  $\text{MOF}_{5hs}^3$  reconstruction in one cell for further discussion in Section 6.1, and in cyan the  $\text{MOF}_{4hs}^3$  reconstruction for further discussion in Section 6.2.

It is through this example that we demonstrate the principle advantage of  $\text{MOF}_{4hs}^3$ ,  $\text{MOF}_{5hs}^3$ , and  $\text{MOF}_{6hs}^3$ : the ability to achieve the same level of accuracy on a coarser grid. To do this, we explore the minimum level of mesh refinement needed for the  $\text{MOF}_{1hs}^3$  PLIC method to achieve the same overall level of accuracy as each of our multi-plane methods.

We first consider the shape in its entirety over a  $5 \times 5 \times 5$  grid, chosen heuristically as the minimum level over which the  $\text{MOF}_{Nhs}^3$  methods reasonably approximate the curved features of the shape. For each grid, we sum the symmetric difference error across cells in the mesh. In contrast to the moment error, the sum of the symmetric difference error metric directly represents the total error in the shape. In Figure 17, we see that when the error is aggregated across all cells in the mesh, there is similar performance among the proposed family of methods, particularly relative to the single-plane alternative. However, we can see that an  $7 \times 7 \times 7$  grid of cells is necessary to reconstruct the shape with accuracy roughly equal to the  $\text{MOF}_{4hs}^3$ ,

and an  $8 \times 8 \times 8$  grid is needed to surpass the accuracy of  $\text{MOF}_{5hs}^3$  and  $\text{MOF}_{6hs}^3$ . This correlates to a nearly three-fold increase in the number of mixed-material cells (58 vs. 163), serving in part to justify the increased cost of a multi-plane MOF method compared to the analogous PLIC procedure.

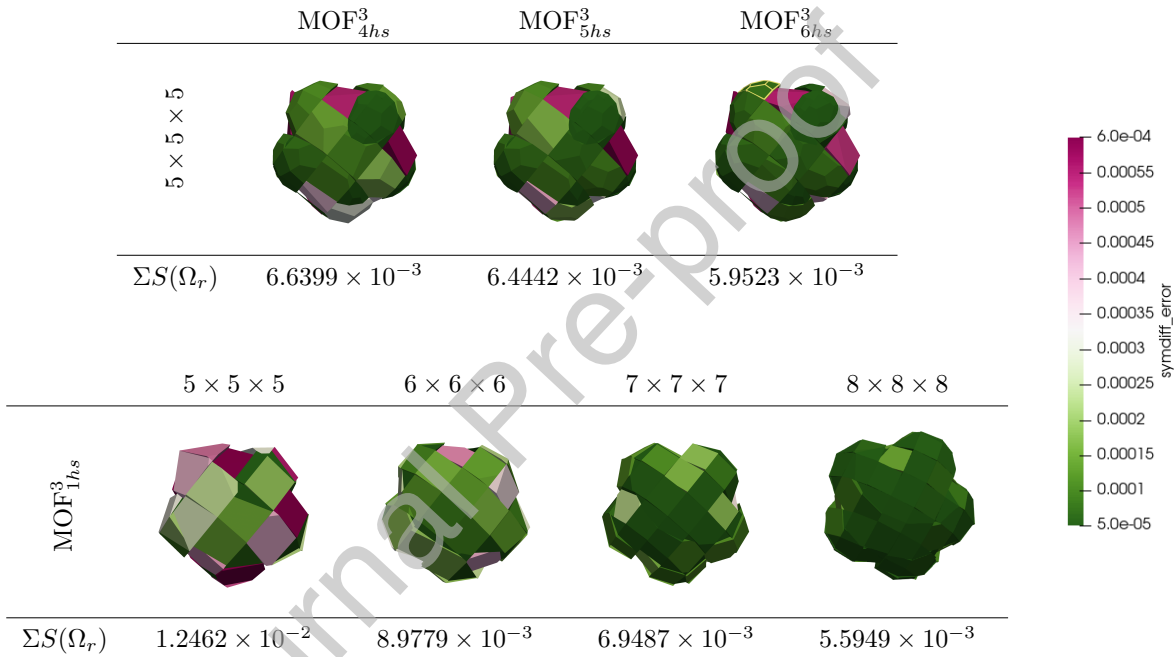


Figure 17: Error comparison across varying levels of refinement evaluated for a non-convex material. Additional levels of refinement are needed for a PLIC method to achieve the same aggregate levels of accuracy. In this example, an  $8 \times 8 \times 8$  grid is required for  $\text{MOF}_{1hs}^3$  to achieve the same total symmetric difference error as  $\text{MOF}_{6hs}^3$ . At the top of the  $\text{MOF}_{6hs}^3$  reconstruction we have emphasized one POM in yellow for further discussion in Figure 18.

In many ways, we consider such a demonstration to be more effective at illustrating the merits of our reconstruction technique than formal convergence analysis, as the particulars of the ground truth geometry often have a far more significant influence on reconstruction accuracy than the resolution of the mesh. For example, we observe in Figure 17 that the error of each  $\text{MOF}_{Nhs}^3$  method (including the PLIC method) on the  $5 \times 5 \times 5$  grid varies considerably with the complexity of the geometric features present in a given cell.

However, we also observe in this example that the error for the PLIC method becomes more uniformly distributed among the cells as the mesh resolution is increased. More generally, it is almost always the case that at sufficiently high mesh resolutions, the portion of material covered by a single cell is simple enough that it can be accurately recovered by even a single plane. On the other hand, this means that there are comparatively fewer advantages to using our multi-plane methods when the geometry in a cell already has no distinct geometric features.

To see this more clearly, we consider in Figure 18 example of mesh refinement the POM in the  $\text{MOF}_{6hs}^3$  example of Figure 17 which we have highlighted in yellow. At the coarsest level of resolution, a single

computational cell, the ground truth shape has relatively high curvature. As such there is over an order of magnitude of difference between the errors of the PLIC method and the six-plane reconstruction on the original cell. However, as the mesh resolution increases, the comparative improvement becomes markedly smaller. As observed previously in the “Ellipsoid 2” example of Figure 9, considering the surface of this shape on increasingly small cells results in an optimization problem that is subject to additional spurious

local minima, many of which use fewer planes than are available during optimization. Indeed, the five-plane method in this example actually outperforms the six-plane method on the  $2 \times 2 \times 2$  and  $3 \times 3 \times 3$  grid, despite having fewer parameters which can be optimized. Ultimately, however, the key result of this figure is that the  $\text{MOF}_{6hs}^3$  method is able to reconstruct features on a single cell with an accuracy that the PLIC method requires a  $3 \times 3 \times 3$  grid to achieve.

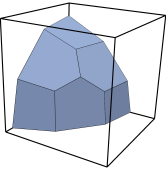
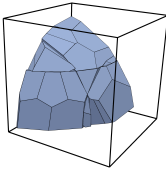
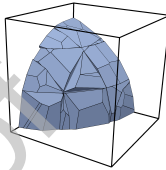
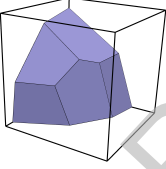
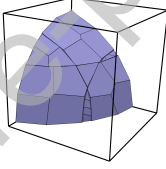
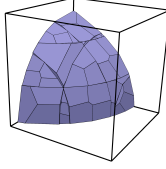
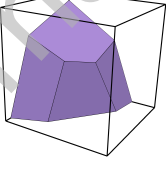
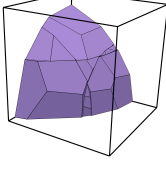
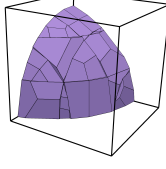
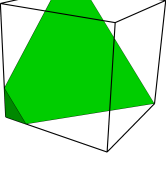
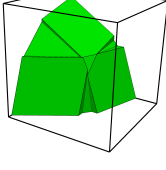
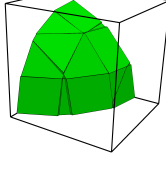
	$1 \times 1 \times 1$	$2 \times 2 \times 2$	$3 \times 3 \times 3$
$\text{MOF}_{6hs}^3$			
$\Sigma S(\Omega_r)$	$6.96227 \times 10^{-5}$	$3.6633 \times 10^{-5}$	$2.7864 \times 10^{-5}$
$\text{MOF}_{5hs}^3$			
$\Sigma S(\Omega_r)$	$9.5327 \times 10^{-5}$	$2.8328 \times 10^{-5}$	$1.8954 \times 10^{-5}$
$\text{MOF}_{4hs}^3$			
$\Sigma S(\Omega_r)$	$1.2266 \times 10^{-4}$	$4.3528 \times 10^{-5}$	$2.2858 \times 10^{-5}$
$\text{MOF}_{1hs}^3$			
$\Sigma S(\Omega_r)$	$4.8481 \times 10^{-4}$	$1.0026 \times 10^{-4}$	$4.8324 \times 10^{-5}$

Figure 18: Error comparison for a single cell across varying levels of refinement. We can see that additional levels of refinement are needed to capture important geometric features accurately, particularly in areas of high or variable curvature.

## 6. Discussion

### 6.1. Reconstruction of Nonconvex Material

The proposed family of methods are intended to accurately represent both convex material *or* material whose complement is convex, but many shapes have features for which neither is the case. This is shown somewhat adversarially in the example of Section 5.2.3, which has the property where all non-convex material localized to a cell is *also* not the complement of a convex shape. We highlight all such cells in Figure 19, and can observe that the error in the method along either error metric is concentrated to cells which do not contain convex material. Indeed, in all cells that *do* contain convex material, the  $\text{MOF}_{6hs}^3$  method achieves significant agreement with the ground-truth geometry.




Figure 19: Reconstruction of material that is neither convex nor has a convex complement. Although the proposed technique can effectively approximate convex material, it is difficult to use such a reconstruction strategy to represent material which is neither convex nor the complement of a convex shape, as shown by the increased error for such cells.

We also observe through Figure 19 that the majority of cells use a convex material in the approximation, despite the material itself being non-convex. For a particularly difficult example, we refer back to the material highlighted in red in Figure 16, for which  $\text{MOF}_{5hs}^3$  uses a convex approximation that, visually at least, appears to be a poor approximation of the natural crease in the ground-truth shape. We investigate this particular cell in Figure 20, where we see that the complex geometric features of the material fragment make both the convex and convex-complement reconstructions largely undesirable. Nevertheless, we see that the convex reconstruction (which has functionally combined the two components of ground-truth material) provides a better approximation along each of the two measured error metrics. We also see in this single cell example that both trial materials generated by  $\text{MOF}_{5hs}^3$  outperform the PLIC reconstruction, as such a technique can (at best) only accurately represent convex material in a convex cell. In either case, we can see throughout the provided examples that the multi-plane method dramatically outperforms the PLIC method numerically, despite the result appearing visually unintuitive.




Figure 20: Reconstruction of a single POM that is neither convex nor has a convex complement. Although the selection of a nonconvex material appears unintuitive for this type of shape, it nevertheless provides a considerably lower error both in terms of the moment error and the symmetric difference error.

### 6.2. Geometric Artifacts

In each of the proposed MOF methods, we occasionally observe that one or more planes used in the reconstruction only exclude a very small volume from the convex intersections of half-spaces that represents the approximated material. When the material itself is convex, these artifacts manifests through “gaps” or “holes” between adjacent cells that are otherwise filled in the ground-truth shape. The same type of artifact is more prominent visually, however, in the case where the material is non-convex, such as that emphasized in cyan in Figure 16, where thin “slices” of material appear catastrophic for the quality of the interface.

The primary source of these artifacts is mathematical. The geometric moments used in the objective function vary continuously with our parameterization of the reconstructed interface with no direct influence from cell boundaries. These “gaps” or “slices” in the reconstruction have near-zero volume, and so this means that during optimization, there is no little to no distinction between returned interfaces that have or lack such artifacts, as they functionally do not contribute to the error.

From the perspective of the reconstruction problem, we consider these artifacts to be a simple visual error whose visually striking appearance can be attributed almost entirely to the particular viewing angle at which the shape is observed.

From a simulation perspective, they are similarly inert with respect to Lagrangian remapping. Such methods naturally handle discontinuous material components, and the the fact that artifacts of this kind have near-zero volume means their individual moments do not meaningfully contribute to the intersection with a backtraced cell. However, we acknowledge that there are circumstances in which material interactions between the *surface* of such spurious components can be impactful for physics applications.

We consider the more rigorous treatment of these surface interactions to be separate from the reconstruction problem, but we nevertheless suggest the approach of straight-forwardly removing from the reconstruction any plane that only clips a negligible volume from the cell. By weighting this cutoff volume by the reference material volume, we also ensure that cells with *only* small volume components are not completely removed. While by definition this procedure only removes material with otherwise insignificant volume, the total volume within the cell must still be tightly controlled to preserve physically in numerical simulation. In such cases where the volume has diverged from the reference, we can re-enforce the volume constraint on the remaining clipping planes through the same root-finding problem employed during optimization at a cost that is less than that of a single optimization step (See Section 3.2). We demonstrate the results of this strategy on each of the proposed family of methods on an  $8 \times 8 \times 8$  grid of cells for the shape in Figure 16. In Figure 21 we see that even an extremely conservative tolerance for the cut volume works well to remove these artifacts.

### 6.3. Shapes with Identical Moments

Although the entire set of geometric moments is sufficient to uniquely identify any shape, numerical restrictions require that a method such as ours only consider a finite subset of available moment data during optimization. This leads to cases in which multiple geometrically distinct shapes have identical moments up to a certain order, typically when the shape itself posses symmetry of some kind. A possible




Figure 21: Example of visual artifact removal. Because the optimization procedure is indifferent to the number of components that compose each material in a cell, there are often superficial artifacts introduced by bounding planes that define regions of near-zero volume. This is especially noticeable in the case of convex-complement reconstructions, as such artifacts have the appearance of thin sheets or lines. Nevertheless, we consider these extraneous features to be largely irrelevant to downstream physics applications, and suggest a simple strategy for removing them.

solution would be to simply use additional high-order moments in the construction of the interface to distinguish such shapes. Indeed, it is for this reason that we only propose MOF methods that use, at minimum, third-order moments, as we observe empirically that there exist too many possible symmetries for geometric objects in 3D for second-order moments to be sufficient.

However, we note that even using third-order moments, there are many shapes that remain indistinguishable. For example, this can be observed in any regular cube inscribed in the same sphere (See Figure 22), all of which have identical moments up to and including third-order. Indeed, any MOF\_{6hs}^3 method attempted on these shapes will terminate immediately, as the initial inscribed cube matches the *available* moments exactly.




Figure 22: Example of distinct shapes with some identical moments. These shapes have identical moments up to third order, and can therefore not be distinguished by any of the proposed methods.

At the same time, we have observed that considering *more* than third-order moments also leads to undesirable results, likely because fourth-order moments individually encode less important information than quantities derived from lower-order moments like the centroid. This is particularly problematic when the proposed optimization procedure weights all moments equally. For example, a reconstruction that more closely matches the 15 fourth-order moments than the 10 third-order moments will likely be less accurate

to the underlying geometry.

Nevertheless, we observe throughout this work that cases where moments up to third-order are incapable of distinguishing shapes are largely adversarial, and the reconstructions generated by the provided  $\text{MOF}_{Nhs}^3$  methods are suitably capable of expressing realistic geometric features.

#### 6.4. Sensitivity to Noise in Reference Moment Data

In a practical application, the provided reference moments are subject to numerical error. For example, one possible source of error is the physics simulation scheme under which the material evolves over time, from which moment data is calculated. This causes the material itself to erroneously drift, which impacts the accuracy of the interface reconstruction relative to the unknown ground truth.

Another potentially more manageable source is the calculation of moment data itself from the polyhedral approximation after each step of the simulation. Here too the accuracy of the interface reconstruction method is affected, but now instead relative to data from a previous state of the simulation. While this error cannot be accounted for by interface reconstruction techniques, it is important that the non-linear optimization scheme used by the proposed family of methods is not unnecessarily sensitive to slight numerical issues with reference moment data.

To demonstrate this, we consider a sphere placed on a  $3 \times 3 \times 3$  grid, a shape for which the exact shape cannot be recovered exactly. The sphere itself is centered at  $(0.1, -0.1, 0.1)$  and has a radius of 0.7 to avoid unfairly exploiting the symmetry of the shape in the reconstruction.

After computing ground-truth moment data for this shape, we add an increasing amount of normally distributed noise to each raw moment to simulate this second kind of numerical error. We then perform a reconstruction with the proposed family of methods, as well as the simple  $\text{MOF}_{1hs}^3$  PLIC method for further comparison. Specifically, we add a sample  $n_i$  from the normal distribution  $\mathcal{N}(0, 10^{-L})$  to each moment for varying levels of noise  $L$ , in effect perturbing the moment data in the  $L^{\text{th}}$  decimal place. For context, the ground-truth moments vary amongst themselves on the order of  $10^{-3} \sim 10^{-5}$ . We compile these results for various levels of noise in Figure 23. The first row of noisy results shows a fairly realistic case, where a noise level  $L = 8$  corresponds roughly to the rounding that would be observed if double precision values were converted to single precision. This is essentially a perturbation in the 6<sup>th</sup> significant figure. As shown in the figure, this has functionally no impact on the accuracy of the method.

By  $L = 5$ , the inaccuracy borders on unrealistic, as we have now perturbed most pieces of data in at least their 3<sup>rd</sup> significant figure. By this point the quality of the reconstruction has suffered significantly, particularly with the increased presence of visible geometric artifacts stemming from the use of non-convex reconstructions, yet we nevertheless can recognize the original sphere shape. By the 3<sup>rd</sup> level of noise, all values have zero digits of accuracy after perturbation, and it is only at this point does the shape become completely unrecognizable and unusable. Altogether, these results demonstrate that the non-linear optimization scheme is fairly robust to imprecision in moment data, and more generally, that moment information is a useful method of encoding geometric features of polyhedra.

## 7. Conclusions and Future Work

In this work we have presented a novel approach to 3D interface reconstruction, in which we approximate material as the convex intersection (or the complement thereof) of 4, 5, and 6 half-spaces. The parameters of each half-space are selected through the minimization of a non-linear optimization problem over zeroth-order through third-order reference moments. In many ways, this work represents the natural progression of the existing body of literature in 2D multi-plane interface reconstruction, where we have combined the geometric expressiveness derived from using multiple independently parameterized planes (as in  $\text{MOF}_{2hp}^2$ ) with the computational efficiency derived from using a single initial condition for the non-linear optimization (as in  $\text{MOF}_{PIE}^3$ ). At the same time, our approach is aided by a custom normalization procedure that uses data from a reference ellipsoid to transform the space over which optimization is performed, thereby improving robustness and computational efficiency of the method. Through these techniques, we are able to create material approximations that far outpace PLIC methods on the same level of resolution, in particular being able to more accurately resolve areas of high curvature or sharp features.

The development of this family of techniques leaves much room for further improvements, both computationally and algorithmically. For example, the largest barrier to reaching a true global minimum as




Figure 23: Example of the effect of noise in moment data on interface reconstruction. For noise level  $L$ , we add to each piece of moment data a random sample from the normal distribution  $\mathcal{N}(0, 10^{-L})$ . While the presence of noise in reference moment data makes it impossible to reconstruct the shape accurately with respect to the ground truth, our non-linear optimization scheme is nevertheless stable enough to withstand small perturbations in moment data before becoming unusable.

opposed to a somewhat acceptable local minimum during optimization is the presence of “inactive planes,” in which one of the planes reaches a plateau of the objective function, for which the local gradient is zero.

In such cases, the plane becomes fixed in this position for the remainder of the optimization procedure. While the current approach of “clamping” planes to be tangent to the cell alleviates this issue, it is an *ad hoc* solution that could be improved in future work. Furthermore, while the generic Levenberg-Marquardt algorithm used for numerical optimization is effective at solving this problem, it is possible that more specified procedures could achieve an improved optimal solution at a faster rate, particularly those that utilize automatic differentiation. Additionally, one could improve computational performance across a mesh with a method that adaptively selects the number of planes to use in the reconstruction (akin to [40] or [34]), as areas of low curvature are often sufficiently captured by a single plane, while the advantages between  $\text{MOF}_{4hs}^3$ ,  $\text{MOF}_{5hs}^3$ , and  $\text{MOF}_{6hs}^3$  increase with increasing curvature.

Finally, there exist many techniques in the MOF literature which could be more-or-less immediately applied to the proposed multi-plane method. As described in Section 1.2, defining the interface with half-spaces makes the proposed method highly compatible with contemporary advection strategies, and in principle generalizes well to multi-material scenarios.

## 8. Acknowledgment

This work was done under the auspices of the National Nuclear Security Administration of the US Department of Energy at Los Alamos National Laboratory under Contract No. 89233218CNA000001. The authors gratefully acknowledges the support of the US Department of Energy National Nuclear Security Administration Advanced Simulation and Computing Program.

The Authors would like to thank Roxana Bujack and Lucas Blakeslee for useful discussions. A special thanks to Robert Chiodi for advice and help with using his IRL library and useful discussions.

## Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## References

- [1] James Edward Pilliod and Elbridge Gerry Puckett. Second-order accurate volume-of-fluid algorithms for tracking material interfaces. *Journal of Computational Physics*, 199(2):465–502, 2004. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2003.12.023>.
- [2] Tomislav Marić, Douglas B. Kothe, and Dieter Bothe. Unstructured un-split geometrical volume-of-fluid methods – a review. *Journal of Computational Physics*, 420:109695, 2020. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2020.109695>.
- [3] D.L. Sun and W.Q. Tao. A coupled volume-of-fluid and level set (voset) method for computing incompressible two-phase flows. *International Journal of Heat and Mass Transfer*, 53(4):645–655, 2010. ISSN 0017-9310. doi: <https://doi.org/10.1016/j.ijheatmasstransfer.2009.10.030>. URL <https://www.sciencedirect.com/science/article/pii/S0017931009005717>.
- [4] Matthew Jemison, Eva Loch, Mark Sussman, Mikhail Shashkov, Marco Arienti, Mitsuhiro Ohta, and Yaohong Wang. A coupled level set-moment of fluid method for incompressible two-phase flows. *Journal of Scientific Computing*, 54:454–491, 2013.
- [5] Andrew Cahaly, Fabien Eyraud, and Olivier Desjardins. Plic-net: A machine learning approach for 3d interface reconstruction in volume of fluid methods. *International Journal of Multiphase Flow*, 178:104888, Aug 2024. ISSN 0301-9322. doi: 10.1016/j.ijmultiphaseflow.2024.104888. URL <http://dx.doi.org/10.1016/j.ijmultiphaseflow.2024.104888>.
- [6] Ruben Scardovelli and Stephane Zaleski. Interface reconstruction with least-square fit and split eulerian–lagrangian advection. *International Journal for Numerical Methods in Fluids*, 41(3):251–274, 2003. doi: <https://doi.org/10.1002/flid.431>.
- [7] Shengping Liu, Heng Yong, Shaodong Guo, Yiqing Shen, and Guoxi Ni. An improved continuity-preserving interface reconstruction method for multi-material flow. *Computers & Fluids*, 224:104960, 2021. ISSN 0045-7930. doi: <https://doi.org/10.1016/j.compfluid.2021.104960>.
- [8] Martin van der Eijk and Peter Wellens. An efficient bilinear interface reconstruction algorithm and consistent multidimensional unsplit advection scheme for accurate capturing of highly-curved interfacial shapes on structured grids. *Journal of Computational Physics*, 498:112656, 2024. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2023.112656>. URL <https://www.sciencedirect.com/science/article/pii/S0021999123007519>.
- [9] J. López, J. Hernández, P. Gómez, and F. Faura. An improved plic-vof method for tracking thin fluid structures in incompressible two-phase flows. *Journal of Computational Physics*, 208(1):51–74, 2005. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2005.01.031>.

- [10] Austin Han, Robert Chiodi, and Olivier Desjardins. Capturing thin structures in vof simulations with two-plane reconstruction. *Journal of Computational Physics*, 519:113453, 2024. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2024.113453>. URL <https://www.sciencedirect.com/science/article/pii/S0021999124007010>.
- [11] Glenn Robert Price. *A piecewise parabolic volume tracking method for the numerical simulation of interfacial flows*. University of Calgary Calgary, AB, 2000.
- [12] Fabien Evrard, Robert Chiodi, Berend van Wachem, and Olivier Desjardins. Simulating interfacial flows: a farewell to planes. *arXiv preprint arXiv:2401.15012*, 2024.
- [13] Fabien Evrard, Robert Chiodi, Austin Han, Berend van Wachem, and Olivier Desjardins. First moments of a polyhedron clipped by a paraboloid. *SIAM Journal on Scientific Computing*, 45(5):A2250–A2274, 2023. doi: [10.1137/22M1524308](https://doi.org/10.1137/22M1524308). URL <https://doi.org/10.1137/22M1524308>.
- [14] I. Ginzburg and G. Wittum. Two-phase flows on interface refined grids modeled with vof, staggered finite volumes, and spline interpolants. *Journal of Computational Physics*, 166(2):302–335, 2001. ISSN 0021-9991. doi: <https://doi.org/10.1006/jcph.2000.6655>.
- [15] S.V. Diwakar, Sarit K. Das, and T. Sundararajan. A quadratic spline based interface (quasi) reconstruction algorithm for accurate tracking of two-phase flows. *Journal of Computational Physics*, 228(24):9107–9130, 2009. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2009.09.014>. URL <https://www.sciencedirect.com/science/article/pii/S0021999109004999>.
- [16] Stewart Mosso, Christopher Garasi, and Richard Drake. A smoothed two- and three-dimensional interface reconstruction method. *Comput. Vis. Sci.*, 12(7):365–381, Sep 2009. ISSN 1432-9360.
- [17] Ram Kumar Maity, T. Sundararajan, and K. Velusamy. An accurate interface reconstruction method using piecewise circular arcs. *International Journal for Numerical Methods in Fluids*, 93(1):93–126, 2021. doi: <https://doi.org/10.1002/fld.4876>.
- [18] Yujie Chen, Junhua Gong, Wei Lu, Bohong Wang, Dongliang Sun, Bo Yu, Wei Zhang, and Wenquan Tao. A circle-based interface reconstruction algorithm based on the coupled volume-of-fluid and level set method. *Physics of Fluids*, 36(3), 2024. doi: [10.1063/5.0200627](https://doi.org/10.1063/5.0200627). URL <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85189088052&doi=10.1063%2f5.0200627&partnerID=40&md5=0692950c553328bc83c4f0c80fbb0b5c>. Cited by: 4; All Open Access, Bronze Open Access.
- [19] Yujie Chen, Junhua Gong, Dongliang Sun, Dongxu Han, Peng Wang, Bo Yu, and Wen-Quan Tao. A three-dimensional curve interface reconstruction algorithm for two-phase fluid flow. *Journal of Computational Physics*, 520:113489, 2025. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2024.113489>. URL <https://www.sciencedirect.com/science/article/pii/S002199912400737X>.
- [20] Jan Flusser, Barbara Zitova, and Tomas Suk. *Moments and Moment Invariants in Pattern Recognition*. Wiley Publishing, 2009. ISBN 0470699876.
- [21] K Weiss, G Zagaris, R Rieben, and A Cook. Spatially accelerated shape embedding in multimaterial simulations. In *Proceedings of the 25th International Meshing Roundtable*. Lawrence Livermore National Lab. (LLNL), Livermore, CA (United States), 08 2016. URL <https://www.osti.gov/biblio/1357384>.
- [22] Brian N. Granzow, Stephen D. Bond, Michael J. Powell, and Daniel A. Ibanez. An efficient second-order adaptive procedure for inserting cad geometries into hexahedral meshes using volume fractions, 2025. URL <https://arxiv.org/abs/2504.03525>.
- [23] Hyung Taek Ahn and Mikhail Shashkov. Multi-material interface reconstruction on generalized polyhedral meshes. *Journal of Computational Physics*, 226(2):2096–2132, 2007. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2007.06.033>.

- [24] Mikhail Shashkov and Eugene Kikinzon. Moments-based interface reconstruction, remap and advection. *Journal of Computational Physics*, 479:111998, 2023. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2023.111998>.
- [25] Philippe Hergibo, Qiuhua Liang, Timothy N. Phillips, and Zihua Xie. A quadtree-based adaptive moment-of-fluid method for interface reconstruction with filaments. *Journal of Computational Physics*, 499:112719, 2024. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2023.112719>. URL <https://www.sciencedirect.com/science/article/pii/S0021999123008148>.
- [26] S.M. Thamil Kumaran, Rohit Kumar, and B. Premachandran. A multi-directional advection based moment of fluid method for phase change problems. *Journal of Computational Physics*, 508:113011, 2024. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2024.113011>. URL <https://www.sciencedirect.com/science/article/pii/S0021999124002602>.
- [27] Vadim Dyadechko and Mikhail Shashkov. Moment-of-fluid interface reconstruction. Technical report, Los Alamos National Laboratory, Los Alamos, NM, Oct 2005. Technical Report LA-UR-05-7571.
- [28] Zhouteng Ye, Cody Estebe, Yang Liu, Mehdi Vahab, Zeyu Huang, Mark Sussman, Alireza Moradikazrouni, Kouros Sholee, Yongsheng Lian, Mitsuhiro Ohta, and Mohammed Hussaini. An improved coupled level set and continuous moment-of-fluid method for simulating multiphase flows with phase change. *Communications on Applied Mathematics and Computation*, 6, 08 2023. doi: 10.1007/s42967-023-00286-6.
- [29] Thomas Milcent and Antoine Lemoine. Moment-of-fluid analytic reconstruction on 3d rectangular hexahedrons. *Journal of Computational Physics*, 409:109346, 2020. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2020.109346>. URL <https://www.sciencedirect.com/science/article/pii/S0021999120301200>.
- [30] Anirudh Asuri Mukundan, Thibaut Ménard, Jorge César Brändle de Motta, and Alain Berlemont. A 3d moment of fluid method for simulating complex turbulent multiphase flows. *Computers & Fluids*, 198:104364, 2020. ISSN 0045-7930. doi: <https://doi.org/10.1016/j.compfluid.2019.104364>. URL <https://www.sciencedirect.com/science/article/pii/S0045793019303238>.
- [31] Thomas Milcent and Antoine Lemoine. An analytic approach for the moment-of-fluid interface reconstruction method on tetrahedral meshes. *Journal of Computational Physics*, 500:112758, 2024. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2024.112758>. URL <https://www.sciencedirect.com/science/article/pii/S002199912400007X>.
- [32] Murray Cutforth, Philip T. Barton, and Nikos Nikiforakis. An efficient moment-of-fluid interface tracking method. *Computers & Fluids*, 224:104964, 2021. ISSN 0045-7930. doi: <https://doi.org/10.1016/j.compfluid.2021.104964>. URL <https://www.sciencedirect.com/science/article/pii/S0045793021001316>.
- [33] Antoine Lemoine, Stéphane Glockner, and Jérôme Breil. Moment-of-fluid analytic reconstruction on 2d cartesian grids. *Journal of Computational Physics*, 328:131–139, 2017. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2016.10.013>. URL <https://www.sciencedirect.com/science/article/pii/S0021999116305113>.
- [34] Mikhail Shashkov. An adaptive moments-based interface reconstruction using intersection of the cell with one half-plane, two half-planes and a circle. *Journal of Computational Physics*, 494:112504, 2023. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2023.112504>.
- [35] Igor Chollet, Giulia Lissoni, Théo Corot, Philippe Hoch, Thomas Leroy, and Laurent Dumas. Curved interface reconstruction for 2d compressible multi-material flows. *ESAIM: Proceedings and Surveys*, 67: 178–190, 2020.
- [36] Jacob Spainhour. Quadratic moment-of-fluid interface reconstruction. Florida State University Undergraduate Honors Thesis, Apr 2020. URL [http://purl.flvc.org/fsu/fd/FSU\\_libsubv1\\_scholarship\\_submission\\_1587147272\\_cff86e96](http://purl.flvc.org/fsu/fd/FSU_libsubv1_scholarship_submission_1587147272_cff86e96).

- [37] Ronald A. Remmerswaal and Arthur E.P. Veldman. Parabolic interface reconstruction for 2d volume of fluid methods. *Journal of Computational Physics*, 469:111473, 2022. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2022.111473>. URL <https://www.sciencedirect.com/science/article/pii/S0021999122005356>.
- [38] Matthew Jemison, Mark Sussman, and Mikhail Shashkov. Filament capturing with the multimaterial moment-of-fluid method. *Journal of Computational Physics*, 285:149–172, 2015.
- [39] Philippe Hergibo, Timothy N. Phillips, and Zihua Xie. A moment-of-fluid method for resolving filamentary structures using a symmetric multi-material approach. *Journal of Computational Physics*, 491:112401, 2023. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2023.112401>. URL <https://www.sciencedirect.com/science/article/pii/S0021999123004965>.
- [40] Robert Chiodi and Mikhail Shashkov. A moment-of-fluid interface reconstruction using polygon inscribed in ellipse in 2d. *Journal of Computational Physics*, 528:113814, 2025. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2025.113814>. URL <https://www.sciencedirect.com/science/article/pii/S002199912500097X>.
- [41] Severin Strobl, Arno Formella, and Thorsten Pöschel. Exact calculation of the overlap volume of spheres and mesh elements. *Journal of Computational Physics*, 311:158–172, 2016. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2016.02.003>. URL <https://www.sciencedirect.com/science/article/pii/S0021999116000577>.
- [42] Donald Marquardt. An algorithm for least square estimation of non-linear parameters. *SIAM Journal on Applied Mathematics*, 11:431–441, Jun 1963. doi: 10.1137/0111030.
- [43] Gaël Guennebaud, Benoît Jacob, et al. Eigen v3. <http://eigen.tuxfamily.org>, 2010.
- [44] Robert Chiodi and Olivier Desjardins. General, robust, and efficient polyhedron intersection in the interface reconstruction library. *Journal of Computational Physics*, 449:110787, Oct 2021. doi: 10.1016/j.jcp.2021.110787.
- [45] M. Dowell and P. Jarratt. A modified regula falsi method for computing the root of an equation. *BIT*, 11(2):168–174, June 1971. ISSN 0006-3835. doi: 10.1007/BF01934364. URL <https://doi.org/10.1007/BF01934364>.