

Sparse Ferguson-Hermite Signed Distance Fields

Róbert Bán¹ and Gábor Valasek¹

Eötvös Loránd University, Hungary
{rob.ban,valasek}@inf.elte.hu

Abstract

We investigate Hermite interpolation in the context of discrete signed distance field filtering. Our method uses tricubic Hermite interpolation to generate a C^1 continuous approximation to the signed distance function of the input scene. Our representation is kept purely first order by setting the mixed partial derivatives to zero, similarly to how Ferguson constructed bicubic Hermite patches. Our scheme stores four scalars at each sample, the value of the signed distance function and its first three partial derivatives. We optimize storage by only storing voxels that enclose a volume boundary. We show that this provides both a significant reduction in storage and render times compared to a dense grid of Ferguson-Hermite samples. Moreover, our construct requires smaller storage than traditional zero order trilinearly filtered fields of the same visual quality, at the expense of performance.

CCS Concepts

• **Computing methodologies** → Rendering; Shape modeling; • **Mathematics of computing** → Continuous functions;

1. Introduction

Implicit functions are flexible representations of objects. Surfaces are defined by the zero level sets of $f: \mathbb{E}^3 \rightarrow \mathbb{R}$ scalar-valued functions in 3D space. In most high-performance real-time graphics applications, these implicit surfaces are discretized. Traditionally, these are represented as simple function value samples stored in 3D textures. The texture is then sampled using trilinear interpolation, resulting in a C^0 -continuous approximation. However, if a visually smooth output is desired, higher resolution textures are required.

It was shown that one can achieve a visually higher quality reconstruction or smaller storage by using higher order data, even on regular grids [VB23]. By only storing the cells that contain parts of the surface, storage can be further reduced. This work




Figure 1: Dense, regular grid around a model, and the cells used for rendering.

presents a method for detecting empty cells that are guaranteed not to contain volume boundaries, and a simple yet accurate and high-performance ray intersection algorithm for rendering a sparse structure of tricubic Ferguson-Hermite cells enclosed in boxes. These boxes correspond to $2 \times 2 \times 2$ samples of a dense 3D texture.

2. Dense Hermite fields

Our aim is to improve the smoothness of the result without significantly increasing storage, such as in Figure 3. To achieve this, we extend the traditional representation by incorporating the partial derivatives of the function along with the function values. By using tensor product Hermite interpolation in reconstruction, the gradient information is taken into account, resulting in a C^1 -continuous approximation. Although more data needs to be stored per sample, the resolution of the texture can be reduced as Hermite interpolation provides a higher order accurate approximation to the sampled function [VB23]. As a rule of thumb, a Ferguson-Hermite field of resolution K^3 offers at least the same visual quality as a trilinearly filtered $(2K)^3$ field. Practically, this halves the required number of scalars, however, we found that the Ferguson-Hermite field resolution may be often lowered even more, as illustrated in Figure 2.

3. Hermite cell rendering

Naive Hermite interpolation is computationally more expensive than trilinear filtering, especially since the latter may be hardware accelerated on GPUs while the former necessitates manual sample interpolation. As such, we employ empty space skipping by separately rendering the voxel boxes of the dense grid but only those




Figure 2: Top row: trilinear (top half: normal with central differences, bottom half: exact normal), bottom row: Hermite interpolation. Resolution left to right: 16^3 , 32^3 , 64^3 .

cells that contain volume boundaries. Moreover, this allows us to only read the field sample values at the start of the search or even obtaining them as vertex attributes.

The search process is straightforward. We subdivide the Hermite approximation of the cell at several equidistant points along the ray and look for a sign change, indicating an intersection point between the evaluated positions. We assume a linear change between the points and approximate the solution accordingly, without the need for additional evaluations. Root refinement may be also added.

4. Empty cell detection

In our rendering algorithm, we only need to process the non-empty cells, making the detection of such an important task. The simplest case is when a cell has two vertices with function values of different signs, as this guarantees that the surface crosses the cell.

To detect empty cells, we transform the Hermite polynomial into the Bernstein basis and leverage its properties [Far01]. The Bernstein coefficients can be calculated directly from the function values and gradients at the sample positions. If all coefficients have the same sign, the cell is guaranteed to be empty due to the convex hull property of the Bernstein basis. This is illustrated in Figure 1.

5. Results

Table 1 summarizes performance measurements tested on the bunny model. The tests were run on a desktop AMD RX 5700 at full HD resolution. The *Trilinear* and *Hermite* cases used sphere tracing, with a maximal error of $\epsilon = 10^{-4}$. The *Hermite Cell* column describes the proposed method; the number of subdivisions were set so the rendered result is visually indistinguishable from the former methods (2 or 3 subdivisions).




Figure 3: Comparison of low resolution approximations of an implicit curve in red: $f(x, y) = (x^2 + y^2)^2 - 2(x^2 - y^2) - 0.1255$, bilinear interpolation in blue, and Hermite interpolation in green.

Res.	Trilinear		Hermite		Hermite Cell	
16^3	0.72 ms	(4k)	2.90 ms	(16k)	1.48 ms	(20k)
32^3	0.76 ms	(32k)	2.80 ms	(130k)	1.47 ms	(80k)
64^3	0.90 ms	(260k)	2.95 ms	(1M)	1.69 ms	(310k)
128^3	1.06 ms	(2M)	2.96 ms	(8.4M)	3.45 ms	(1.3M)

Table 1: Performance measurements on the bunny model.

The parentheses contain the approximate number of stored scalars for each test case (k = thousand, M = million). The performance drop in the last test case is caused by the small size of the triangles and arithmetically heavy manual Hermite computations.

6. Conclusions

Our discrete implicit representation is well-suited for smooth surfaces, but the evaluation of the approximation requires manual filtering of the data. We reduce the number of evaluations by rendering each interpolation cell individually, thus limiting the intersection search space. Furthermore, the representation used for direct rendering requires less storage space, as the amount of stored data is proportional to the square of the resolution, rather than its cube.

Acknowledgement

We would like to thank Visual Concepts for providing the AMD GPU used in the tests.

Supported by the ÚNKP-22-3 New National Excellence Program of the Ministry for Culture and Innovation from the source of the National Research, Development and Innovation Fund.

References

- [Far01] FARIN G.: *Curves and surfaces for CAGD: A practical guide*, 5 ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2001. 2
- [VB23] VALASEK G., BAN R.: Higher Order Algebraic Signed Distance Fields. *Computer-Aided Design and Applications* (Jan. 2023), 1005–1028. doi:10.14733/cadaps.2023.1005-1028. 1