FFT-Based Surface Reconstruction

Outline

- Approach
  - What is a shape?
  - Reduction to Volume Integration
  - Implementation
  - Non-Uniform Samples
- Results
- Conclusion

What is a shape?

Plato:

Socrates: *To what then do we give the name of figure?*

Meno: *I would rather that you should answer, Socrates.*

…

Socrates: *I define figure to be that in which the solid ends; or, more concisely, the limit of solid.*

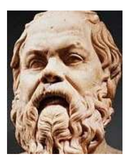

Plato, *Meno*. 380 BC

What is a shape?

Jordan Curve Theorem:

Any continuous simple closed curve in the plane, separates the plane into two disjoint regions, the inside and the outside.

Camille Jordan, *Cours d'Analyze de l'École Polytechnique*. 1887

What is a shape?

Jordan Curve Theorem:

Any continuous simple closed curve in the plane, separates the plane into two disjoint regions, the inside and the outside.

Camille Jordan, *Cours d'Analyze de l'École Polytechnique*. 1887

What is a shape?

Jordan Curve Theorem:

Any continuous simple closed curve in the plane, separates the plane into two disjoint regions, the inside and the outside.

1887: Camille Jordan gives an incorrect proof
1905: Oswald Veblen give correct proof
1906: Arthur Schönflies gives proof of the strong form of the Jordan curve theorem
1909: Errors in Schönflies's proof fixed by Luitzen Brouwer
1912: Jordan-Brouwer Separation Theorem: Any imbedding of the n-1 dimensional sphere into n-dimensional Euclidean space, separates the Euclidean space into two disjoint regions.

Camille Jordan, *Cours d'Analyze de l'École Polytechnique*. 1887

What is a shape?

Divergence (Gauss's) Theorem:

Let V be a region in space with boundary ∂V. The volume integral of the divergence ∇·F of F over V and the surface integral of F over the boundary ∂V of V are related by:

$$\int_V (\nabla \cdot F)\, dV = \int_{\partial V} F \cdot dA$$

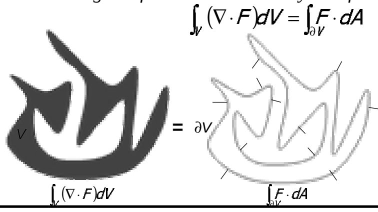

$$F(p) = (F_x(p), F_y(p), F_z(p))$$

$$\nabla \cdot F = \frac{\partial F_x}{\partial x} + \frac{\partial F_y}{\partial y} + \frac{\partial F_z}{\partial z}$$

$$F \cdot dA = \langle F, \vec{n} \rangle |dA|$$

What is a shape?

Divergence (Gauss's) Theorem:

Let V be a region in space with boundary ∂V. The volume integral of the divergence ∇·F of F over V and the surface integral of F over the boundary ∂V of V are related by:

$$\int_V (\nabla \cdot F)\, dV = \int_{\partial V} F \cdot dA$$

1762: Joseph Lagrange discovers the theorem
1813: Carl Gauss discovers the theorem (independently)
1825: George Green discovers the theorem (independently)
1831: Mikhail Ostrogradsky discovers the theorem (independently) and proves it

Outline

- Approach
  - What is a shape?
  - Reduction to Volume Integration
  - Implementation
  - Non-Uniform Samples
- Results
- Conclusion

Reduction to Volume Integration I

Characteristic Function:

The *characteristic function* χ_V of a solid V is the function:

$$\chi_V(x,y,z) = \begin{cases} 1 & \text{if } (x,y,z) \in V \\ 0 & \text{otherwise} \end{cases}$$

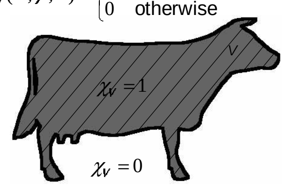

Reduction to Volume Integration II

Fourier Coefficients:

The *Fourier coefficients* of the characteristic function give an expression of χ_V as a sum of complex exponentials:

$$\chi_V(x,y,z) = \sum_{l,m,n} \hat{\chi}_V(l,m,n) e^{i(lx+my+nz)}$$

Reduction to Volume Integration III

Volume Integration:

The Fourier coefficients of the characteristic function χ_V can be obtained by integrating:

$$\hat{\chi}_V(l,m,n) = \int_{[0,1]^3} \chi_V(x,y,z) e^{-2\pi i(lx+my+nz)} dxdydz$$

Reduction to Volume Integration III

Volume Integration:
The Fourier coefficients of the characteristic function χ_V can be obtained by integrating:

$$\begin{aligned}\hat{\chi}_V(l,m,n) &= \int_{[0,1]^3} \chi_V(x,y,z)e^{-2\pi i(lx+my+nz)}dxdydz \\ &= \int_V e^{-2\pi i(lx+my+nz)}dxdydz\end{aligned}$$

since the characteristic function is one inside of V and zero everywhere else.

Outline

- Approach
  - What is a shape?
  - Reduction to Volume Integration
  - Implementation
  - Non-Uniform Samples
- Results
- Conclusion

Applying Gauss's Theorem

Surface Integration
If $\bar{F}_{lmn}(x,y,z)$ is any function whose divergence is equal to the (l,m,n)-th complex exponential:

$$\left(\nabla\cdot\vec{F}_{lmn}\right)(x,y,z) = e^{-2\pi i(lx+my+nz)}$$

applying the Divergence Theorem, the volume integral can be expressed as a surface integral:

$$\int_V e^{-2\pi i(lx+my+nz)}dxdydz = \int_{\partial V}\left\langle \vec{F}_{lmn}(p), n(p)\right\rangle dp$$

Reconstruction Algorithm

Problem 1: A direct implementation of this algorithm is on the order of $O(n^5)$.
(Assuming we sample the 3D function at $O(n^3)$ samples and we are given $O(n^2)$ input sample points.)

$$\hat{\chi}_V(l,m,n) = \int_{\partial V}\left\langle \vec{F}_{l,m,n}(p), n(p)\right\rangle dp \approx \sum_{i=1}^{k}\left\langle \vec{F}_{l,m,n}(p_i), n_i\right\rangle$$

- Apply the inverse Fourier Transform to

Problem 2: To do this, we need to find a vector valued function $F_{l,m,n}(x,y,z)$ such that:

$$\left(\nabla\cdot F_{l,m,n}\right)(x,y,z) = e^{-2\pi i(lx+my+nz)}$$

Choosing the Functions $\bar{F}_{l,m,n}$

There are many solutions to the equation:

$$\left(\nabla\cdot\vec{F}_{l,m,n}\right)(x,y,z) = e^{-2\pi i(lx+my+nz)}$$

Choosing the Functions $\bar{F}_{l,m,n}$

There are many solutions to the equation:

$$\left(\nabla\cdot\vec{F}_{l,m,n}\right)(x,y,z) = e^{-2\pi i(lx+my+nz)}$$

Examples:

$$\vec{F}_{l,m,n}(x,y,z) = \begin{pmatrix}\frac{i}{2\pi l}e^{-2\pi i(lx+my+nz)} \\ 0 \\ 0\end{pmatrix} \quad \vec{F}_{l,m,n}(x,y,z) = \begin{pmatrix}0 \\ \frac{i}{2\pi m}e^{-2\pi i(lx+my+nz)} \\ 0\end{pmatrix} \quad \vec{F}_{l,m,n}(x,y,z) = \begin{pmatrix}0 \\ 0 \\ \frac{i}{2\pi n}e^{-2\pi i(lx+my+nz)}\end{pmatrix}$$

Choosing the Functions $\bar{F}_{l,m,n}$

There are many solutions to the equation:

$$\left(\nabla \cdot \vec{F}_{l,m,n}\right)(x,y,z) = e^{-2\pi i(lx+my+nz)}$$

<u>Examples</u>:

$$\vec{F}_{l,m,n}(x,y,z) = \begin{pmatrix} \frac{i}{2\pi l} e^{-2\pi i(lx+my+nz)} \\ 0 \\ 0 \end{pmatrix} \quad \vec{F}_{l,m,n}(x,y,z) = \begin{pmatrix} 0 \\ \frac{i}{2\pi m} e^{-2\pi i(lx+my+nz)} \\ 0 \end{pmatrix} \quad \vec{F}_{l,m,n}(x,y,z) = \begin{pmatrix} 0 \\ 0 \\ \frac{i}{2\pi n} e^{-2\pi i(lx+my+nz)} \end{pmatrix}$$

$$\vec{F}_{l,m,n}(x,y,z) = \frac{1}{6\pi}\begin{pmatrix} \frac{i}{l} e^{-2\pi i(lx+my+nz)} \\ \frac{i}{m} e^{-2\pi i(lx+my+nz)} \\ \frac{i}{n} e^{-2\pi i(lx+my+nz)} \end{pmatrix} \quad \vec{F}_{l,m,n}(x,y,z) = \frac{1}{2\pi}\begin{pmatrix} \frac{i}{l+m+n} e^{-2\pi i(lx+my+nz)} \\ \frac{i}{l+m+n} e^{-2\pi i(lx+my+nz)} \\ \frac{i}{l+m+n} e^{-2\pi i(lx+my+nz)} \end{pmatrix}$$

Choosing the Functions $\bar{F}_{l,m,n}$

There are many solutions to the equation:

$$\left(\nabla \cdot \vec{F}_{l,m,n}\right)(x,y,z) = e^{-2\pi i(lx+my+nz)}$$

We choose the functions $\bar{F}_{l,m,n}$ so that they commute with translation and rotation:

$$\vec{F}_{l,m,n}(x,y,z) = \begin{pmatrix} \frac{i}{2\pi}\frac{l}{l^2+m^2+n^2} e^{-2\pi i(lx+my+nz)} \\ \frac{i}{2\pi}\frac{m}{l^2+m^2+n^2} e^{-2\pi i(lx+my+nz)} \\ \frac{i}{2\pi}\frac{n}{l^2+m^2+n^2} e^{-2\pi i(lx+my+nz)} \end{pmatrix}$$

Choosing the Functions $\bar{F}_{l,m,n}$

Does not Commute:

$$F_{l,m,n}(x,y,z) = \begin{pmatrix} \frac{i}{l+m+n} e^{-i(lx+my+nz)} \\ \frac{i}{l+m+n} e^{-i(lx+my+nz)} \\ \frac{i}{l+m+n} e^{-i(lx+my+nz)} \end{pmatrix}$$

Commutes:

$$F_{l,m,n}(x,y,z) = \begin{pmatrix} \frac{il}{l^2+m^2+n^2} e^{-i(lx+my+nz)} \\ \frac{im}{l^2+m^2+n^2} e^{-i(lx+my+nz)} \\ \frac{in}{l^2+m^2+n^2} e^{-i(lx+my+nz)} \end{pmatrix}$$

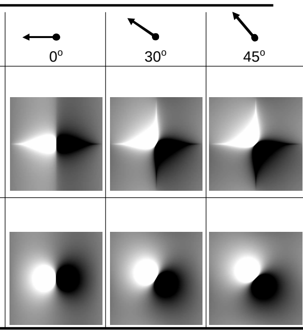

Choosing the Functions $\bar{F}_{l,m,n}$

Does not Commute:

$$F_{l,m,n}(x,y,z) = \begin{pmatrix} \frac{i}{l+m+n} e^{-i(lx+my+nz)} \\ \frac{i}{l+m+n} e^{-i(lx+my+nz)} \\ \frac{i}{l+m+n} e^{-i(lx+my+nz)} \end{pmatrix}$$

Commutes:

$$F_{l,m,n}(x,y,z) = \begin{pmatrix} \frac{il}{l^2+m^2+n^2} e^{-i(lx+my+nz)} \\ \frac{im}{l^2+m^2+n^2} e^{-i(lx+my+nz)} \\ \frac{in}{l^2+m^2+n^2} e^{-i(lx+my+nz)} \end{pmatrix}$$

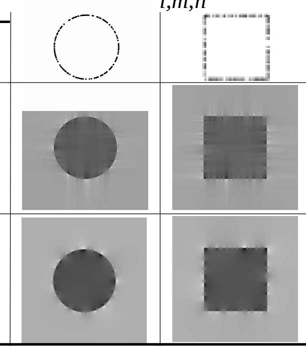

Choosing the Functions $\bar{F}_{l,m,n}$

<u>Properties</u>:

- Commutes with translation and rotation.
- Unique $\bar{F}_{l,m,n}$ with these properties.
- Provides an extended notion of convolution, turning the reconstruction problem into an $O(n^3\log(n))$ problem.

Outline

- Approach
  - What is a shape?
  - Reduction to Volume Integration
  - Implementation
  - Non-Uniform Samples
- Results
- Conclusion

Non-Uniform Samples

Challenge:

In a direct implementation of Monte-Carlo integration, we assume uniform sampling:

$$\int f(x)dx \approx \frac{1}{n}\sum_{i=1}^{n} f(x_i)$$

Non-Uniform Samples

Challenge:

In a direct implementation of Monte-Carlo integration, we assume uniform sampling:

$$\int f(x)dx \approx \frac{1}{n}\sum_{i=1}^{n} f(x_i)$$

However, often the oriented point samples may be non-uniformly distributed:

- Parts of scans may overlap
- Patches parallel to the view plane may be more densely sampled
- Representations may store fewer points in regions of low curvature

Non-Uniform Samples

Challenge:

If we have a sampling density ρ_i associated to each sample x_i, we can modify the equation:

$$\int f(x)dx \approx \frac{1}{\left(\sum_{i=1}^{n}\rho_i\right)}\sum_{i=1}^{n} f(x_i)\rho_i$$

Non-Uniform Samples

Challenge:

If we have a sampling density ρ_i associated to each sample x_i, we can modify the equation:

$$\int f(x)dx \approx \frac{1}{\left(\sum_{i=1}^{n}\rho_i\right)}\sum_{i=1}^{n} f(x_i)\rho_i$$

However, when we get an oriented point sample, we usually aren't given the sampling density at each sample.

Non-Uniform Samples

Challenge:

If we have a sampling density ρ_i associated to each sample x_i, we can modify the equation:

Ho
san g
der

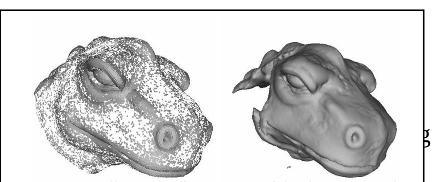

Non-Uniform Samples   Unweighted Reconstruction

Non-Uniform Samples

Approach:

Compute the sampling density at each sample by counting the number of samples around it:

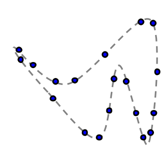

Non-Uniform Samples

Approach:

Compute the sampling density at each sample by counting the number of samples around it:

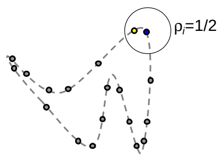

Non-Uniform Samples

Approach:

Compute the sampling density at each sample by counting the number of samples around it:

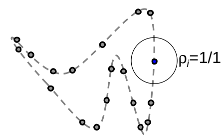

Non-Uniform Samples

Approach:

Compute the sampling density at each sample by counting the number of samples around it:

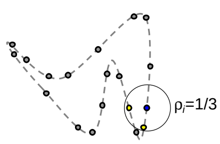

Non-Uniform Samples

Approach:

Compute the sampling density at each sample by counting the number of samples around it:

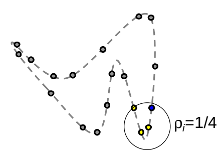

Non-Uniform Samples

Approach:

Compute the sampling density at each sample by counting the number of samples around it.

To do this quickly, we "splat" the points into a voxel grid and convolve with a low-pass filter.

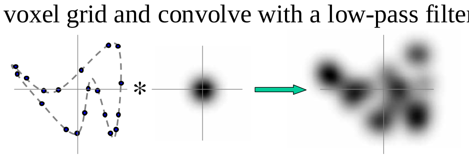

Non-Uniform Samples

Approach:

Compute the sampling density at each sample by counting the number of samples around it.

To do this quickly, we "splat" the points into a voxel grid and convolve with a low-pass filter.

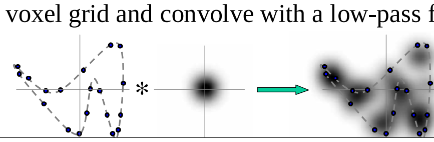

Set ρ_i equal to the reciprocal of the value of the convolution at point p_i.

Non-Uniform Samples

Approach:

Compute the sampling density at each sample by counting the number of samples around it.

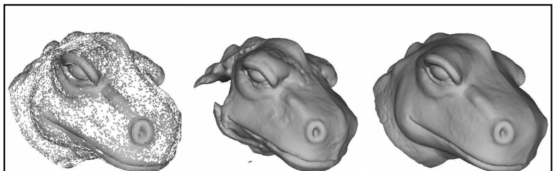

Non-Uniform Samples | Unweighted Reconstruction | Weighted Reconstruction

Outline

- Approach
  - What is a shape?
  - Reduction to Volume Integration
  - Implementation
  - Non-Uniform Samples
- Results
- Conclusion

Results (Resolution)

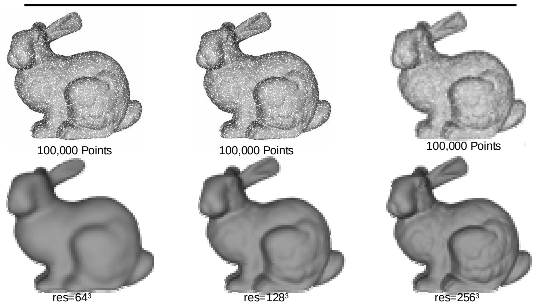

100,000 Points | 100,000 Points | 100,000 Points

res=64^3, tris=11,672, time<0:01 | res=128^3, tris=49,008, time=0:01 | res=256^3, tris=199,796, time=0:07

Results (Sample Count)

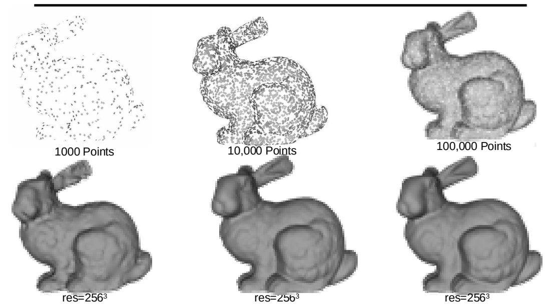

1000 Points | 10,000 Points | 100,000 Points

res=256^3, tris=206,216, time=0:07 | res=256^3, tris=200,704, time=0:07 | res=256^3, tris=199,796, time=0:07

Results (Non-Uniform Sampling)

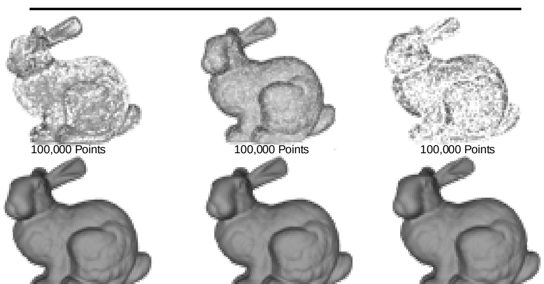

100,000 Points | 100,000 Points | 100,000 Points

res=256^3, tris=220,324, time=0:09 | res=256^3, tris=199,712, time=0:09 | res=256^3, tris=111,680, time=0:09

Results (Non-Uniform / Related)

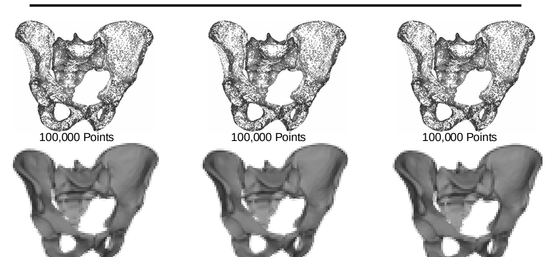

100,000 Points | 100,000 Points | 100,000 Points

RBF Reconstruction† res=256^3, tris=302,000, time=5:23 | MPU Reconstruction‡ res=256^3, tris=288,000, time=0:39 | FFT-Based Reconstruction res=256^3, tris=286,916, time=0:09

†Carr *et al.* SIGGRAPH '01

‡Ohtake *et al.* ACM TOG '03

Results (Noise / Related)

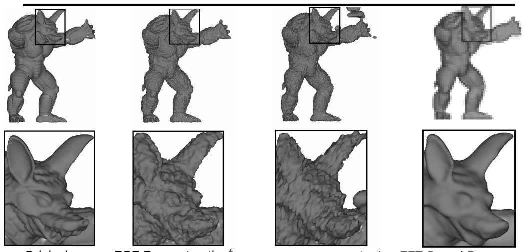

Original | RBF Reconstruction† | MPU Reconstruction‡ | FFT-Based Reconstruction

	RBF Reconstruction†	MPU Reconstruction‡	FFT-Based Reconstruction
res	256^3	256^3	256^3
tris	200,000	205,000	174,824
time	24:10	2:14	0:07
points	100,000	100,000	100,000

†Carr *et al.* SIGGRAPH '01 ‡Ohtake *et al.* ACM TOG '03

Outline

- Approach
  - What is a shape?
  - Reduction to Volume Integration
  - Implementation
  - Non-Uniform Samples
- Results
- Conclusion

Conclusion

Properties:

✓ Fast and simple to compute

✓ Independent of topology

✓ Robust to non-uniform sampling

✓ Robust to noise

✗ $O(n^3)$ space, $O(n^3\log(n))$ time footprint for $O(n^2)$ reconstruction

Conclusion

Theoretical Contribution:

- Transformed the surface reconstruction problem into a Volume integral
- Used the Divergence Theorem to express the integral as a surface integral
- Used Monte-Carlo integration to approximate the surface integral as a summation over an oriented point sample

Conclusion

An algorithm for reconstruction that proceeds in three simple steps:

Conclusion

Three simple steps:

1. Splat the oriented points into a voxel grid

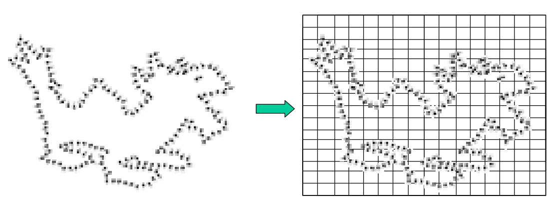

Conclusion

Three simple steps:

1. Splat the oriented points into a voxel grid
2. Convolve with a fixed filter

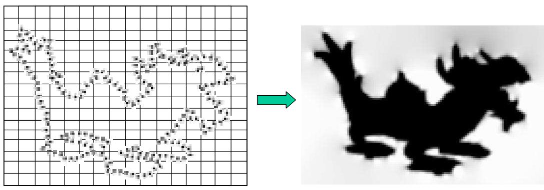

Conclusion

Three simple steps:

1. Splat the oriented points into a voxel grid
2. Convolve with a fixed filter

3. Extract the iso-surface

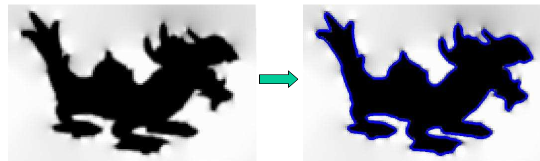

Surface Reconstruction = Integration

One way to look at this surface reconstruction problem is as an integration problem:

> The gradient field of a characteristic function is zero everywhere but at the boundary points. At these points, the gradient field is just equal to the surface normal.

Surface Reconstruction = Integration

One way to look at this surface reconstruction problem is as an integration problem:

> The gradient field of a characteristic function is zero everywhere but at the boundary points. At these points, the gradient field is just equal to the surface normal.

> Given a (sample of the) gradient field we want to compute the scalar function which has the same gradient field.

Relationship to the Poisson Equation

If f is a 3D scalar, we can express f in terms of its Fourier decomposition:

$$f(p) = \sum_{\vec{l} \in Z^3} \hat{f}(\vec{l}) e^{-2\pi i \langle \vec{l}, p \rangle}$$

Relationship to the Poisson Equation

If f is a 3D scalar, we can express f in terms of its Fourier decomposition:

$$f(p) = \sum_{\vec{l} \in Z^3} \hat{f}(\vec{l}) e^{-2\pi i \langle \vec{l}, p \rangle}$$

Then, the Laplacian of f is given by:

$$(\nabla^2 f)(p) = (-2\pi i)^2 \|\vec{l}\|^2 \sum_{\vec{l} \in Z^3} \hat{f}(\vec{l}) e^{-2\pi i \langle \vec{l}, p \rangle}$$

Relationship to the Poisson Equation

Thus, if g is a scalar function we can solve for the function f satisfying $\nabla^2 f=g$ by computing the Fourier decomposition of g:

$$g(p)=\sum_{\vec{l}\in Z^3}\hat{g}(\vec{l})e^{-2\pi i\langle \vec{l},p\rangle}$$

and then dividing to get the Fourier coefficients of f:

$$\hat{f}(\vec{l})=\frac{\hat{g}(\vec{l})}{(-2\pi i)^2\|\vec{l}\|^2}$$

Relationship to the Poisson Equation

If $\bar{F}$ is a 3D vector field, we can express $\bar{F}$ in terms of its Fourier decomposition:

$$\vec{F}(p)=\sum_{\vec{l}\in Z^3}\vec{a}(\vec{l})e^{-2\pi i\langle \vec{l},p\rangle}$$

Then, the divergence of $\bar{F}$ is given by:

$$\left(\nabla\cdot\vec{F}\right)(p)=-2\pi i\sum_{\vec{l}\in Z^3}\left\langle \vec{l},\vec{a}(\vec{l})\right\rangle e^{-2\pi i\langle \vec{l},p\rangle}$$

Relationship to the Poisson Equation

Thus, given a gradient field, we can reconstruct the characteristic function by computing the divergence and then solving the Laplacian:

$$\sum_{\vec{l}\in Z^3}\vec{a}(\vec{l})e^{-2\pi i\langle \vec{l},p\rangle}\xrightarrow{\text{Divergence}}-2\pi i\sum_{\vec{l}\in Z^3}\left\langle \vec{l},\vec{a}(\vec{l})\right\rangle e^{-2\pi i\langle \vec{l},p\rangle}\xrightarrow{\text{Laplacian}^{-1}}\frac{i}{2\pi}\sum_{\vec{l}\in Z^3}\frac{\left\langle \vec{l},\vec{a}(\vec{l})\right\rangle}{\|\vec{l}\|^2}e^{-2\pi i\langle \vec{l},p\rangle}$$

Relationship to the Poisson Equation

In our discussion, we turn the samples into a vector field by taking the weighted sum of delta functions:

$$\vec{F}(p)=\sum_{k=1}^{n}\delta_{p_i}(p)n_i$$

Relationship to the Poisson Equation

In our discussion, we turn the samples into a vector field by taking the weighted sum of delta functions:

$$\vec{F}(p)=\sum_{k=1}^{n}\delta_{p_i}(p)n_i$$

Since the Fourier decomposition of the delta function is:

$$\vec{\delta}_{p_i}(\vec{l})=e^{-2\pi i\langle p_i,\vec{l}\rangle}$$

This implies that:

$$\vec{a}(\vec{l})=\sum_{k=1}^{n}e^{-2\pi i\langle p_i,\vec{l}\rangle}n_i$$

Relationship to the Poisson Equation

This is precisely the filter $F_{l,m,n}$ that we used for reconstruction!

$$\sum_{\vec{l}\in Z^3}\vec{a}(\vec{l})e^{-2\pi i\langle \vec{l},p\rangle}\xrightarrow{\text{Divergence}}-2\pi i\sum_{\vec{l}\in Z^3}\left\langle \vec{l},\vec{a}(\vec{l})\right\rangle e^{-2\pi i\langle \vec{l},p\rangle}\xrightarrow{\text{Laplacian}^{-1}}\frac{i}{2\pi}\sum_{\vec{l}\in Z^3}\frac{\left\langle \vec{l},\vec{a}(\vec{l})\right\rangle}{\|\vec{l}\|^2}e^{-2\pi i\langle \vec{l},p\rangle}$$