The Poisson Equation

Michael Kazhdan

(600.657)

Outline

Continuous Laplacian
- Gradients and Divergence
- Kernel of the Laplacian

Discrete Laplacian
- Finite Differences
- Matrix Representation

Normal Equation

Continuous Laplacian

Gradient, Divergence, and the Laplacian:

Given a function F:$\mathbf{R}^2 \rightarrow \mathbf{R}$, the gradient of F is the vector field ∇F:$\mathbf{R}^2 \rightarrow \mathbf{R}^2$ defined by the partial derivatives:

$$\nabla F(x,y) = \left(\frac{\partial F}{\partial x}, \frac{\partial F}{\partial y}\right)$$

Continuous Laplacian

Gradient, Divergence, and the Laplacian:

Given a function F:$\mathbf{R}^2 \rightarrow \mathbf{R}$, the gradient of F is the vector field ∇F:$\mathbf{R}^2 \rightarrow \mathbf{R}^2$ defined by the partial derivatives:

$$\nabla F(x,y) = \left(\frac{\partial F}{\partial x}, \frac{\partial F}{\partial y}\right)$$

Intuitively: At the point p_0, the vector $\nabla F(p_0)$ points in the direction of greatest change of f.

Continuous Laplacian

Examples:

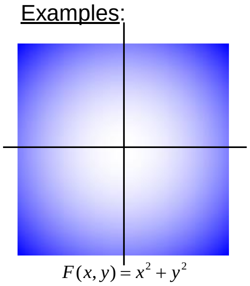

$F(x,y) = x^2 + y^2$

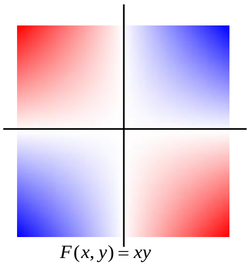

$F(x,y) = xy$

Continuous Laplacian

Examples:

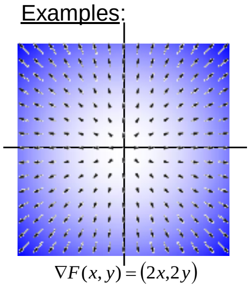

$\nabla F(x,y) = (2x, 2y)$

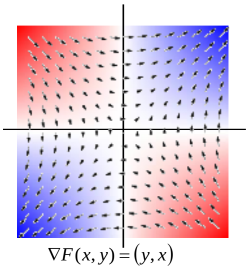

$\nabla F(x,y) = (y,x)$

Continuous Laplacian

Gradient, Divergence, and the Laplacian:

Given a function $F:\mathbf{R}^2\to\mathbf{R}$, the gradient of F is the vector field $\nabla F:\mathbf{R}^2\to\mathbf{R}^2$ defined by the partial derivatives:

$$\nabla F(x,y)=\left(\frac{\partial F}{\partial x},\frac{\partial F}{\partial y}\right)$$

Formally: Fixing the point p_0, for any direction v_0, the 1D function:

$$f(t)=F\left(p_0+tv_0\right)$$

has derivative:

$$f'(0)=\left\langle\nabla F(p_0),v_0\right\rangle$$

Continuous Laplacian

Examples:

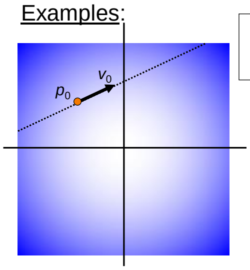

$$F(x,y)=x^2+y^2 \qquad p_0=(-1,1)$$
$$\nabla F(x,y)=(2x,2y) \qquad v_0=(1,2)$$

Continuous Laplacian

Examples:

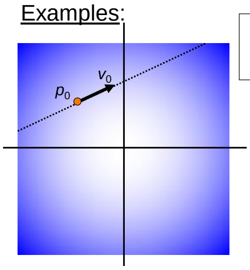

$$F(x,y)=x^2+y^2 \qquad p_0=(-1,1)$$
$$\nabla F(x,y)=(2x,2y) \qquad v_0=(1,2)$$

$$f(t)=F(-1+t,1+2t)$$

Continuous Laplacian

Examples:

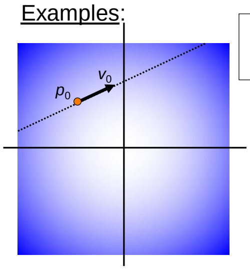

$$F(x,y)=x^2+y^2 \qquad p_0=(-1,1)$$
$$\nabla F(x,y)=(2x,2y) \qquad v_0=(1,2)$$

$$\begin{aligned}f(t)&=F(-1+t,1+2t)\\&=(-1+t)^2+(1+2t)^2\end{aligned}$$

Continuous Laplacian

Examples:

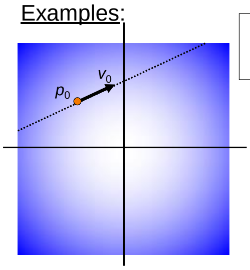

$$F(x,y)=x^2+y^2 \qquad p_0=(-1,1)$$
$$\nabla F(x,y)=(2x,2y) \qquad v_0=(1,2)$$

$$\begin{aligned}f(t)&=F(-1+t,1+2t)\\&=(-1+t)^2+(1+2t)^2\\&=2+2t+2t^2\end{aligned}$$

Continuous Laplacian

Examples:

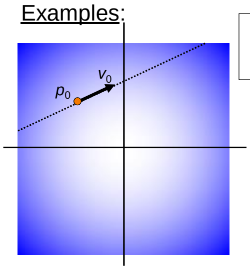

$$F(x,y)=x^2+y^2 \qquad p_0=(-1,1)$$
$$\nabla F(x,y)=(2x,2y) \qquad v_0=(1,2)$$

$$\boxed{f(t)=2+2t+2t^2}$$

$$\left\langle\nabla F(p_0),v_0\right\rangle=\left\langle\nabla F(-1,1),(1,2)\right\rangle$$

Continuous Laplacian

Examples:

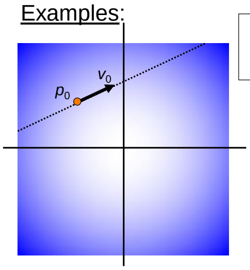

$$F(x,y) = x^2 + y^2 \qquad p_0 = (-1,1)$$
$$\nabla F(x,y) = (2x,2y) \qquad v_0 = (1,2)$$

$$f(t) = 2 + 2t + 2t^2$$

$$\langle \nabla F(p_0), v_0 \rangle = \langle \nabla F(-1,1),(1,2) \rangle = \langle (-2,2),(1,2) \rangle$$

Continuous Laplacian

Examples:

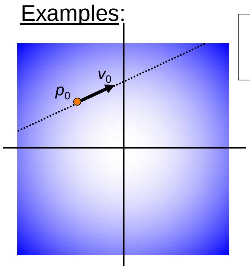

$$F(x,y) = x^2 + y^2 \qquad p_0 = (-1,1)$$
$$\nabla F(x,y) = (2x,2y) \qquad v_0 = (1,2)$$

$$f(t) = 2 + 2t + 2t^2$$

$$\langle \nabla F(p_0), v_0 \rangle = \langle \nabla F(-1,1),(1,2) \rangle = \langle (-2,2),(1,2) \rangle = 2 = f'(0)$$

Continuous Laplacian

Gradient, Divergence, and the Laplacian:

Given a vector field $\vec{F}=(F_1,F_2)$:$\mathbf{R}^2 \to \mathbf{R}^2$, the divergence of $\vec{F}$ is the function $\nabla \cdot \vec{F}$:$\mathbf{R}^2 \to \mathbf{R}$ defined by the partial derivatives:

$$\nabla \cdot \vec{F}(x,y) = \frac{\partial F_1}{\partial x} + \frac{\partial F_2}{\partial y}$$

Continuous Laplacian

Gradient, Divergence, and the Laplacian:

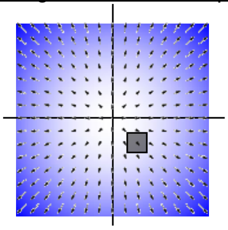

Intuitively: At the point p_0, the divergence $\nabla \cdot F(p_0)$ is a measure of the extent to which the flow (de)compresses at p_0.

Continuous Laplacian

Gradient, Divergence, and the Laplacian:

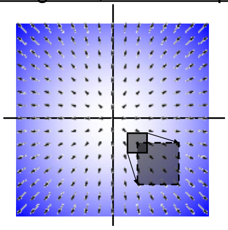

Intuitively: At the point p_0, the divergence $\nabla \cdot F(p_0)$ is a measure of the extent to which the flow (de)compresses at p_0.

Continuous Laplacian

Gradient, Divergence, and the Laplacian:

Given a function F:$\mathbf{R}^2 \to \mathbf{R}$, the Laplacian of F is the function ΔF:$\mathbf{R}^2 \to \mathbf{R}$ (or $\nabla^2 F$) defined by the partial derivatives:

$$\Delta F(x,y) = \nabla \cdot (\nabla F(x,y))$$

Continuous Laplacian

Gradient, Divergence, and the Laplacian:

Given a function $F:\mathbf{R}^2\to\mathbf{R}$, the Laplacian of F is the function $\Delta F:\mathbf{R}^2\to\mathbf{R}$ (or $\nabla^2 F$) defined by the partial derivatives:

$$\Delta F(x,y)=\nabla\cdot\left(\nabla F(x,y)\right)=\frac{\partial^2 F}{\partial x^2}+\frac{\partial^2 F}{\partial y^2}$$

Continuous Laplacian

Gradient, Divergence, and the Laplacian:

Given a function $F:\mathbf{R}^2\to\mathbf{R}$, the Laplacian of F is the function $\Delta F:\mathbf{R}^2\to\mathbf{R}$ (or $\nabla^2 F$) defined by the partial derivatives:

$$\Delta F(x,y)=\nabla\cdot\left(\nabla F(x,y)\right)=\frac{\partial^2 F}{\partial x^2}+\frac{\partial^2 F}{\partial y^2}$$

Intuitively: The Laplacian of F at the point p_0 measures the extent to which the value of F at p_0 differs from the average value of F its neighbors.

Continuous Laplacian

Gradient, Divergence, and the Laplacian:

Given a function $F:\mathbf{R}^2\to\mathbf{R}$, the Laplacian of F is the function $\Delta F:\mathbf{R}^2\to\mathbf{R}$ (or $\nabla^2 F$) defined by the partial derivatives:

$$\Delta F(x,y)=\nabla\cdot\left(\nabla F(x,y)\right)=\frac{\partial^2 F}{\partial x^2}+\frac{\partial^2 F}{\partial y^2}$$

Formally: The Laplacian of F can be defined by considering the family of smoothed functions:

$$G(t,x,y)=F(x,y)*e^{-(x^2+y^2)/t}$$

Continuous Laplacian

Gradient, Divergence, and the Laplacian:

Given a function $F:\mathbf{R}^2\to\mathbf{R}$, the Laplacian of F is the function $\Delta F:\mathbf{R}^2\to\mathbf{R}$ (or $\nabla^2 F$) defined by the partial derivatives:

$$\Delta F(x,y)=\nabla\cdot\left(\nabla F(x,y)\right)=\frac{\partial^2 F}{\partial x^2}+\frac{\partial^2 F}{\partial y^2}$$

Formally: The Laplacian of F can be defined by considering the family of smoothed functions:

$$G(t,x,y)=F(x,y)*e^{-(x^2+y^2)/t}$$

Then the Laplacian has the property that:

$$\left.\frac{\partial G}{\partial t}\right|_{t=0}=\Delta F$$

Continuous Laplacian

$$G(t,x,y)=F(x,y)*e^{-(x^2+y^2)/t}$$
$$\left.\frac{\partial G}{\partial t}\right|_{t=0}=\Delta F$$

Applications to Smoothing:

If we want to perform a small amount of function smoothing on the function F, we can update the function F by setting:

$$F(x,y)\leftarrow F(x,y)+\varepsilon\Delta F(x,y)$$

Continuous Laplacian

Why Do We Care?

Continuous Laplacian

Why Do We Care?

A key component of human perception is the detection of local changes.

Continuous Laplacian

Why Do We Care?

A key component of human perception is the detection of local changes.

↓

It may be easier to work in the space of “local changes” than in the space of “absolute” values.

Continuous Laplacian

Why Do We Care?

A key component of human perception is the detection of local changes.

↓

It may be easier to work in the space of “local changes” than in the space of “absolute” values.

↓

Instead of trying to model a sound/image/shape *F*, it may be easier to model the way the values of *F* change.

Continuous Laplacian

Example

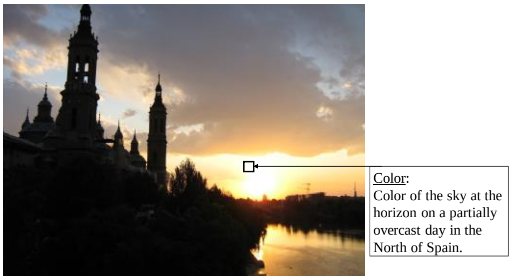

Modeling in image space

Continuous Laplacian

Example

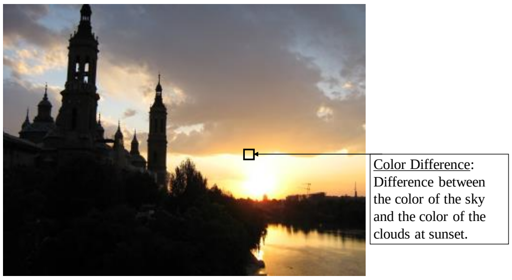

Modeling in Laplacian space

Continuous Laplacian

Challenge:

Given a “difference based” representation, convert it back to a “value based” representation.

Continuous Laplacian

Challenge:

Given a "difference based" representation, convert it back to a "value based" representation.

↓

Poisson Equation:

Given some know "difference function" *G*, solve for the function *F* with the property that:

$$\Delta F(x,y) = G(x,y)$$

Continuous Laplacian

Question: For a given function *G* Is there a unique function *F* with the property that:

$$\Delta F(x,y) = G(x,y)$$

Continuous Laplacian

Question: For a given function *G* Is there a unique function *F* with the property that:

$$\Delta F(x,y) = G(x,y)$$

Equivalently (by linearity of the Laplacian), is there a unique function *F* with the property that:

$$\Delta F(x,y) = 0$$

Continuous Laplacian

Question: For a given function *G* Is there a unique function *F* with the property that:

$$\Delta F(x,y) = G(x,y)$$

Equivalently (by linearity of the Laplacian), is there a unique function *F* with the property that:

$$\Delta F(x,y) = 0$$

Answer: No! (In general)

Continuous Laplacian

Examples $\Delta F(x,y)=0$:

1. $F(x,y) = a$

Continuous Laplacian

Examples $\Delta F(x,y)=0$:

1. $F(x,y) = a$
2. $F(x,y) = ax + by + c$

Continuous Laplacian

Examples $\Delta F(x,y)=0$:

1. $F(x,y)=a$
2. $F(x,y)=ax+by+c$
3. $F(x,y)=axy+bx+cy+d$

Continuous Laplacian

Examples $\Delta F(x,y)=0$:

1. $F(x,y)=a$
2. $F(x,y)=ax+by+c$
3. $F(x,y)=axy+bx+cy+d$
4. $F(x,y)=\left(a\cos(kx)+b\sin(kx)\right)e^{ky}$

Continuous Laplacian

Examples $\Delta F(x,y)=0$:

1. $F(x,y)=a$
2. $F(x,y)=ax+by+c$
3. $F(x,y)=axy+bx+cy+d$
4. $F(x,y)=\left(a\cos(kx)+b\sin(kx)\right)e^{ky}$

A function *F* with the property that $\Delta F(x,y)=0$ is called Harmonic.

Continuous Laplacian

Given a bounded region Ω, and a function, *G* defined inside of Ω, there are many ways to define a function *F* inside of Ω so that $\Delta F(x,y)=G$.

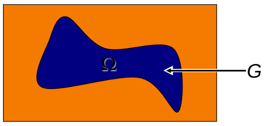

Continuous Laplacian

Given a bounded region Ω, and a function, *G* defined inside of Ω, there are many ways to define a function *F* inside of Ω so that $\Delta F(x,y)=G$.

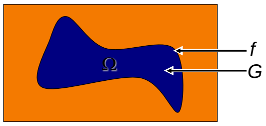

Given a function *f* defined on the boundary ∂Ω, then there is a unique function *F* such that:

$$\Delta F(x,y)=G(x,y) \quad \text{for} \quad (x,y)\in\Omega$$
$$F(x,y)=f(x,y) \quad \text{for} \quad (x,y)\in\partial\Omega$$

Outline

Continuous Laplacian

Discrete Laplacian
- Finite Differences
- Matrix Representation

Normal Equation

Discrete Laplacian

In general, solving the continuous formulation of the Poisson equation:

$$\Delta F(x,y) = G(x,y)$$

is difficult as it would require explicit integration.

However, we can approach the problem of solving the Poisson equation by discretizing.

Discrete Laplacian

Rather than thinking of functions as defined over the continuous domain, we will think of functions as a discrete set of samples over a regular grid:

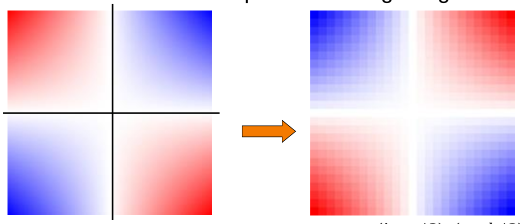

$$F(x,y) = xy \qquad F[i][j] = \frac{(i-w/2)\cdot(y-h/2)}{(w/2)\cdot(h/2)}$$

Discrete Laplacian

In order to formulate the Poisson equation, we use finite differences to define the gradient and the divergence of an array.

Discrete Laplacian (1D)

Set-Up:

In the 1D case, a scalar function and a vector field are equivalent:

1. The gradient is a map from an *N*-dimensional space to an *N*-dimensional space, and
2. The divergence is a map from an *N*-dimensional space to an *N*-dimensional space.

Discrete Laplacian (1D)

We define the discrete gradient $\nabla f[\,]$ as:

$$\nabla f[i] = f[i+1] - f[i]$$

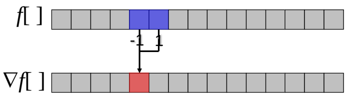

Discrete Laplacian (1D)

We define the discrete gradient $\nabla f[\,]$ as:

$$\nabla f[i] = f[i+1] - f[i]$$

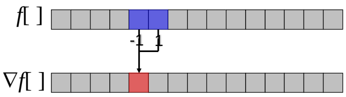

And we define the discrete divergence $\nabla\cdot\vec{f}[\,]$ as:

$$\nabla\cdot\vec{f}[i] = \vec{f}[i] - \vec{f}[i-1]$$

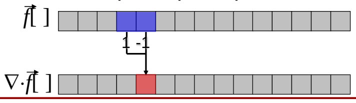

Discrete Laplacian (1D)

Expressed in matrix notation, we get:

$$\nabla = \begin{pmatrix} -1 & 1 & \cdots & 0 & 0 \\ 0 & -1 & & 0 & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & & -1 & 1 \\ 0 & 0 & \cdots & 0 & -1 \end{pmatrix} \qquad \nabla\cdot = \begin{pmatrix} 1 & 0 & \cdots & 0 & 0 \\ -1 & 1 & & 0 & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & & 1 & 0 \\ 0 & 0 & \cdots & -1 & 1 \end{pmatrix}$$

so that the gradient and divergence operators are negative transposes of each other.

Discrete Laplacian (1D)

Thus the discrete Laplacian $\Delta f[\,]$ becomes:

$$\Delta f[i] = f[i+1] - 2f[i] + f[i-1]$$

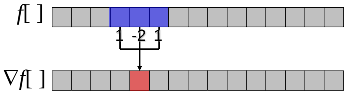

Discrete Laplacian (1D)

And in matrix notation, we get:

$$(\nabla\cdot)\nabla = \Delta = \begin{pmatrix} -2 & 1 & \cdots & 0 & 0 \\ 1 & -2 & & 0 & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & & -2 & 1 \\ 0 & 0 & \cdots & 1 & -2 \end{pmatrix}$$

Note:
Since the matrices corresponding to $\nabla\cdot$ and ∇ are negative transposes of each other, the product of the two matrices (the Laplacian) is symmetric.

Discrete Laplacian (2D)

Set-Up:

In the 2D case, a scalar function and a vector field are different:

1. The gradient is a map from an *N*x*N*-dimensional space to an 2*N*x*N*-dimensional space, and
2. The divergence is a map from an 2*N*x*N*-dimensional space to an *N*x*N*-dimensional space.

Discrete Laplacian (2D)

We define the "gradient" $\nabla f[\,][\,]$ as:

$$\nabla f[i][j] = \big(f[i+1][j] - f[i][j],\ f[i][j+1] - f[i][j]\big)$$

Discrete Laplacian (2D)

We define the "gradient" $\nabla f[\,][\,]$ as:

$$\nabla f[i][j] = \big(f[i+1][j] - f[i][j],\ f[i][j+1] - f[i][j]\big)$$

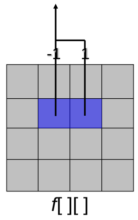

Discrete Laplacian (2D)

We define the “gradient” $\nabla f[\,][\,]$ as:

$$\nabla f[i][j] = \left(f[i+1][j] - f[i][j],\; f[i][j+1] - f[i][j]\right)$$

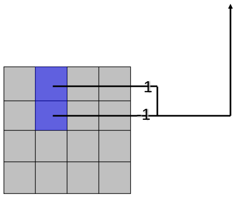

$f[\,][\,]$

Discrete Laplacian (2D)

We define the “gradient” $\nabla f[\,][\,]$ as:

$$\nabla f[i][j] = \left(f[i+1][j] - f[i][j],\; f[i][j+1] - f[i][j]\right)$$

And we define the “divergence” $\nabla \cdot \vec{f}[\,][\,]$ as:

$$\nabla \cdot \vec{f}[i][j] = \vec{f}[i][j].x - \vec{f}[i-1][j].x + \vec{f}[i][j].y - \vec{f}[i][j-1].y$$

Discrete Laplacian (2D)

We define the “gradient” $\nabla f[\,][\,]$ as:

$$\nabla f[i][j] = \left(f[i+1][j] - f[i][j],\; f[i][j+1] - f[i][j]\right)$$

And we define the “divergence” $\nabla \cdot \vec{f}[\,][\,]$ as:

$$\nabla \cdot \vec{f}[i][j] = \underbrace{\vec{f}[i][j].x - \vec{f}[i-1][j].x}_{} + \vec{f}[i][j].y - \vec{f}[i][j-1].y$$

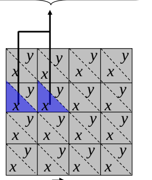

$\vec{f}[\,][\,]$

Discrete Laplacian (2D)

We define the “gradient” $\nabla f[\,][\,]$ as:

$$\nabla f[i][j] = \left(f[i+1][j] - f[i][j],\; f[i][j+1] - f[i][j]\right)$$

And we define the “divergence” $\nabla \cdot \vec{f}[\,][\,]$ as:

$$\nabla \cdot \vec{f}[i][j] = \vec{f}[i][j].x - \vec{f}[i-1][j].x + \underbrace{\vec{f}[i][j].y - \vec{f}[i][j-1].y}_{}$$

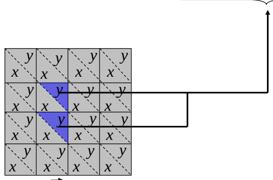

$\vec{f}[\,][\,]$

Discrete Laplacian (1D)

Thus the discrete Laplacian $\Delta f[\,][\,]$ becomes:

$$\Delta f[i][j] = f[i+1][j] + f[i][j+1] - 4f[i][j] + f[i-1][j] + f[i][j-1]$$

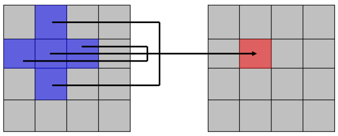

$f[\,][\,]$ $\Delta f[\,][\,]$

Discrete Laplacian (1D)

Thus the discrete Laplacian $\Delta f[\,][\,]$ becomes:

$$\Delta f[i][j] = f[i+1][j] + f[i][j+1] - 4f[i][j] + f[i-1][j] + f[i][j-1]$$

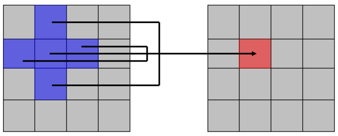

As in the 1D case, the gradient and divergence operators are negative transposes of each other so the Laplacian matrix is symmetric.

Outline

Continuous Laplacian

Discrete Laplacian

Normal Equation

Normal Equation

Question: Since we want to model using "local changes", why not stop with the gradient?

Why not represent a model by its gradient field $\vec{G}$ and then solve for the function *F* such that:

$$\nabla F(x,y) = \vec{G}(x,y)$$

Normal Equation

Question: Since we want to model using "local changes", why not stop with the gradient?

Why not represent a model by its gradient field $\vec{G}$ and then solve for the function *F* such that:

$$\nabla F(x,y) = \vec{G}(x,y)$$

Answer: If we discretize 2D space using a regular *N*x*N* grid, the function *F* becomes an N^2 vector, while the gradient field $\vec{G}$ becomes a $2N^2$ vector.

⇒ The linear system is over-constrained and there may not be any solutions.

Normal Equation

In general, for *m*>*n*, given an *m*-dimensional vector *b* and an *nxm* matrix *A*, we do not expect there to be an *n*-dimensional vector *x* such that:

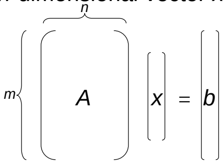

Normal Equation

In general, for *m*>*n*, given an *m*-dimensional vector *b* and an *nxm* matrix *A*, we do not expect there to be an *n*-dimensional vector *x* such that:

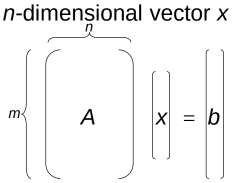

However, we can still try to solve for the vector *x* that minimizes the norm of the residual:

$$x = \arg\min_{x_0} \|Ax_0 - b\|^2$$

Normal Equation

$$\boxed{x = \arg\min_{x_0} \|Ax_0 - b\|^2}$$

Writing out the square norm as:

$$\|Ax_0 - b\|^2 = \langle Ax_0 - b, Ax_0 - b\rangle$$

Normal Equation

$$x = \arg\min_{x_0} \|Ax_0 - b\|^2$$

Writing out the square norm as:

$$\|Ax_0 - b\|^2 = \langle Ax_0 - b, Ax_0 - b\rangle$$

we are looking for the value of *x* minimizing:

$$\Phi(x) = \langle Ax - b, Ax - b\rangle$$
$$= \langle Ax, Ax\rangle - 2\langle Ax, b\rangle + \langle b, b\rangle$$

Normal Equation

$$x = \arg\min_{x_0} \|Ax_0 - b\|^2$$

$$\Phi(x) = \langle Ax, Ax\rangle - 2\langle Ax, b\rangle + \langle b, b\rangle$$

To minimize $\Phi(x)$, we need to find the value of *x* at which the gradient is equal to zero:

$$0 = \nabla\Phi(x)$$

Normal Equation

$$x = \arg\min_{x_0} \|Ax_0 - b\|^2$$

$$\Phi(x) = \langle Ax, Ax\rangle - 2\langle Ax, b\rangle + \langle b, b\rangle$$

To minimize $\Phi(x)$, we need to find the value of *x* at which the gradient is equal to zero:

$$0 = \nabla\Phi(x)$$
$$= 2A^tAx - 2A^tb$$

Normal Equation

$$x = \arg\min_{x_0} \|Ax_0 - b\|^2$$

$$\Phi(x) = \langle Ax, Ax\rangle - 2\langle Ax, b\rangle + \langle b, b\rangle$$

$$A^tb = A^tAx$$

Since *A* is an *nxm* matrix, and since *b* is an *m* dimensional matrix:

- A^tA is an *nxn* matrix
- A^tb is an *n*-dimensional vector

Normal Equation

$$x = \arg\min_{x_0} \|Ax_0 - b\|^2$$

$$\Phi(x) = \langle Ax, Ax\rangle - 2\langle Ax, b\rangle + \langle b, b\rangle$$

$$A^tb = A^tAx$$

Since *A* is an *nxm* matrix, and since *b* is an *m* dimensional matrix:

- A^tA is an *nxn* matrix
- A^tb is an *n*-dimensional vector

$\Rightarrow$ Minimizing the residual amounts to solving a set of *n* equations in *n* unknowns.

Normal Equation

Implications:

If we want to represent a model by its gradient field *G* then we can still solve for the function *F* minimizing:

$$F = \arg\min_{F_0} \left\|\nabla F_0(x,y) - \vec{G}(x,y)\right\|^2$$

Normal Equation

Implications:

If we want to represent a model by its gradient field *G* then we can still solve for the function *F* minimizing:

$$F = \arg\min_{F_0} \left\| \nabla F_0(x,y) - \vec{G}(x,y) \right\|^2$$

To do this, we need to apply the transpose of the gradient and solve:

$$\nabla^t \vec{G}(x,y) = \nabla^t \nabla F(x,y)$$

Normal Equation

Implications:

But we know that the transpose of the gradient is the negative divergence so this gives:

$$\begin{aligned} -\nabla \cdot \vec{G}(x,y) &= -\nabla \cdot \nabla F(x,y) \\ &= -\Delta F(x,y) \end{aligned}$$

Normal Equation

Implications:

But we know that the transpose of the gradient is the divergence so the equation becomes:

$$\begin{aligned} -\nabla \cdot \vec{G}(x,y) &= -\nabla \cdot \nabla F(x,y) \\ &= -\Delta F(x,y) \end{aligned}$$

So even when the difference constraints are given as the gradients of the model, we are still required to solve the Poisson equation.