FFTs in Graphics and Vision

The Laplace Operator

Outline

Math Stuff

- Symmetric/Hermitian Matrices
- Lagrange Multipliers
- Diagonalizing Symmetric Matrices

The Laplacian Operator

Linear Operators

Definition:

Given a real inner product space $(V, \langle\cdot,\cdot\rangle)$ and a linear operator $L: V \to V$, the adjoint of the L is the linear operator L^*, with the property that:

$$\langle v, Lw\rangle = \langle L^* v, w\rangle \qquad \forall v, w \in V$$

Linear Operators

<u>Note</u>:

If V is the space of n-dimensional, real-valued, arrays with the standard inner product:

$$\langle v[\cdot], w[\cdot] \rangle = \sum_{i=1}^{n} v[i] \cdot w[i] = v^t w$$

then the adjoint of a matrix M is its transpose:

$$\begin{aligned} \langle v, Mw \rangle &= v^t M w \\ &= (M^t v)^t w \\ &= \langle M^t v, w \rangle \end{aligned}$$

Linear Operators

Definition:

A real linear operator L is self-adjoint if it is its own adjoint, i.e.:

$$\langle v, Lw \rangle = \langle Lv, w \rangle \qquad \forall v, w \in V$$

Linear Operators

Note:

If V is the space of n-dimensional, real-valued, arrays with the standard inner product:

$$\langle v[\cdot], w[\cdot]\rangle = \sum_{i=1}^{n} v[i]\cdot w[i] = v^t w$$

then a matrix M is self-adjoint if it is symmetric:

$$M = M^t$$

Linear Operators

Definition:

Given a complex inner product space $(V, \langle\cdot,\cdot\rangle)$ and given a linear operator $L: V \to V$, the adjoint of the L is the linear operator L^*, with the property that:

$$\langle v, Lw \rangle = \langle L^* v, w \rangle \qquad \forall v, w \in V$$

Linear Operators

Note:

If V is the space of n-dimensional, complex-valued, arrays with the standard inner product:

$$\langle v[\cdot], w[\cdot]\rangle = \sum_{i=1}^{n} v[i]\cdot\bar{w}[i] = v^t\bar{w}$$

then the adjoint of a matrix M is just the complex conjugate of the transpose:

$$\begin{aligned}\langle v, Mw\rangle &= v^t\overline{Mw}\\ &= (\bar{M}^t v)^t\bar{w}\\ &= \langle\bar{M}^t v, w\rangle\end{aligned}$$

Linear Operators

Definition:

A complex linear operator L is self-adjoint if it is its own adjoint, i.e.:

$$\langle v, Lw \rangle = \langle Lv, w \rangle \qquad \forall v, w \in V$$

Linear Operators

Note:

If V is the space of n-dimensional, complex-valued, arrays with the standard inner product:

$$\langle v[\cdot], w[\cdot]\rangle = \sum_{i=1}^{n} v[i]\cdot\bar{w}[i] = v^t\bar{w}$$

then a matrix M is self-adjoint if it is Hermitian:

$$M = \bar{M}^t$$

Outline

Math

- Symmetric/Hermitian Matrices
- Lagrange Multipliers
- Diagonalizing Symmetric Matrices

The Laplacian Operator

Implicit Surface

Definition:

Given a function $g(x,y,z)$, the implicit surface or iso-surface defined by $g(x,y,z)$ is the set of points in 3D satisfying the condition:

$$g(x,y,z)=0$$

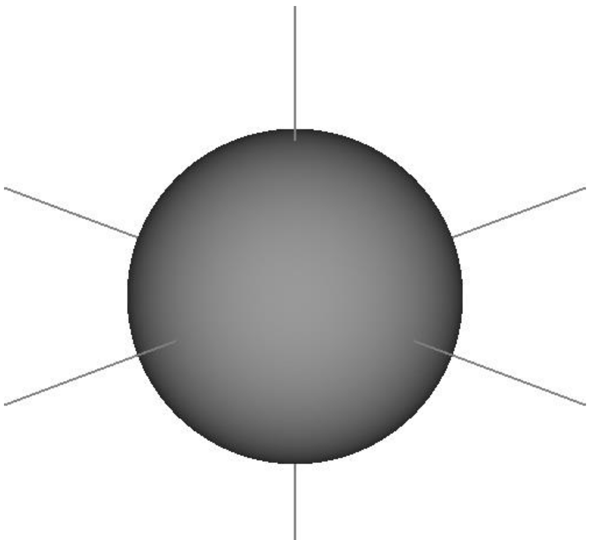

$g(x,y,z)=x^2+y^2+z^2-1$

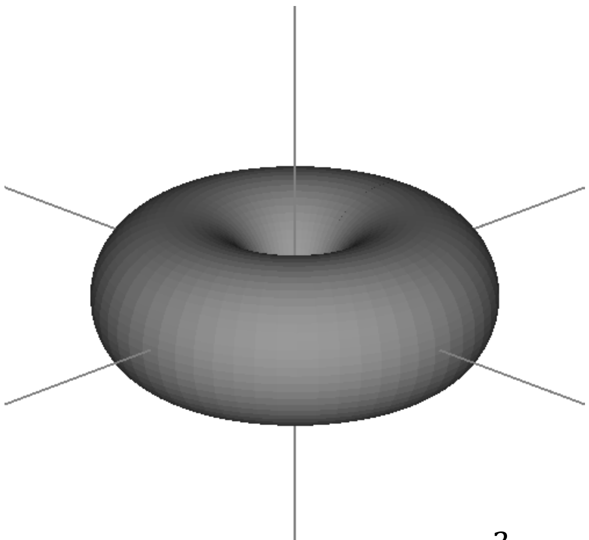

$g(x,y,z)=\left(x^2+y^2+z^2-(R^2+r^2)\right)^2-4R^2(r^2-z^2)$

Lagrange Multipliers

Goal:

Given an implicit surface defined by a function $g(x,y,z)$ and given a function $f(x,y,z)$, we want to find the extrema of f on the surface.

This can be done by finding the points on the surface where the gradient of f is parallel to the surface normal.

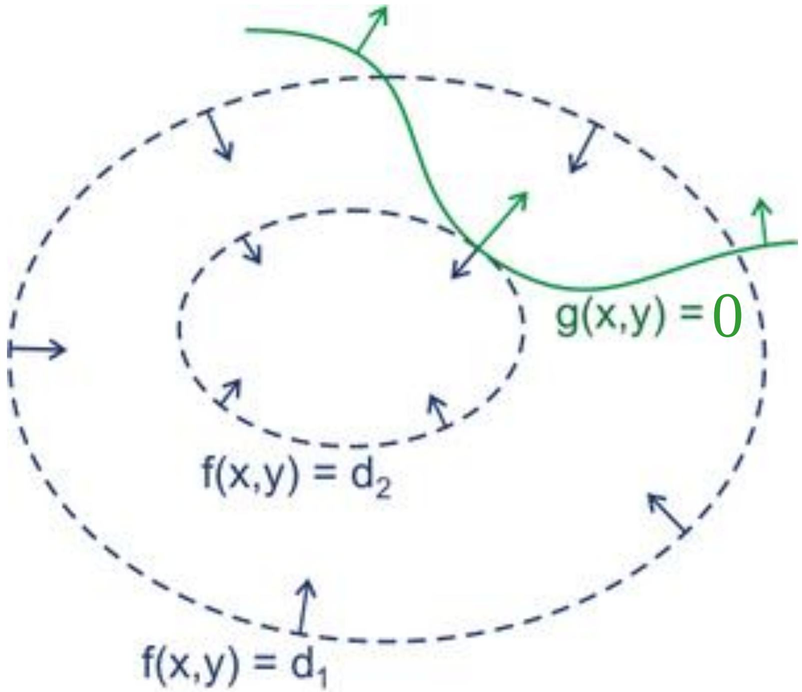

http://www.answers.com

Lagrange Multipliers

Since the implicit surface is the set of points with:

$$g(x,y,z)=0$$

the normal at a point on the surface is parallel to the gradient of g.

Finding the extrema amounts to finding the points (x,y,z) such that:

- $g(x,y,z)=0$ (the point is on the surface)
- $\lambda\nabla f=\nabla g$ (the point is a local extrema)

Outline

Math

- Symmetric/Hermitian Matrices
- Lagrange Multipliers
- Diagonalizing Symmetric Matrices

The Laplacian Operator

Diagonalizing Symmetric Matrices

Claim:

Given the space of n-dimensional, real-valued, arrays and given a symmetric matrix M:

The eigenvectors of M form an orthogonal basis

Diagonalizing Symmetric Matrices

<u>The Eigenvectors Form an Orthogonal Basis</u>:

To show this we will show two things:

1. If v is an eigenvector, then the space of vectors orthogonal to v is fixed by M.
2. At least one eigenvector exists.

Diagonalizing Symmetric Matrices

1. If v is an eigenvector, then the space of vectors orthogonal to v is fixed by M.

Suppose that v is an eigenvector and w is some other vector that is orthogonal to v:

$$\langle v, w \rangle = 0$$

Since v is an eigenvector, this implies that:

$$\langle Mv, w \rangle = \langle \lambda v, w \rangle = 0$$

Since M is symmetric, we have:

$$\langle v, Mw \rangle = \langle Mv, w \rangle = 0$$

Diagonalizing Symmetric Matrices

1. If v is an eigenvector, then the space of vectors orthogonal to v is fixed by M.

If W is the subspace of vectors orthogonal to v:

$$W = \{w \in V | \langle w, v \rangle = 0\}$$

then we have:

$$\langle v, Mw \rangle = 0 \qquad \forall w \in W$$

$$\Updownarrow$$

$$M(w) \in W \qquad \forall w \in W$$

Diagonalizing Symmetric Matrices

1. If v is an eigenvector, then the space of vectors perpendicular to v is fixed by M.

<u>Implications</u>:

If we know that we can find one eigenvector v, we can consider the restriction of M to W.

We know that:

- $M(W) \subset W$
- $\langle Mu, w \rangle = \langle u, Mw \rangle \qquad \forall u, w \in W$

So we can repeat to find the next eigenvector.

Diagonalizing Symmetric Matrices

2. At least one eigenvector must exist

We will show this using Lagrange multipliers:

- The implicit surface will be the sphere in $\mathbb{R}^n$:

$$g(x_1,\cdots,x_n) = x_1^2 + \cdots x_n^2 - 1$$
$$\Updownarrow$$
$$g(v) = \|v\|^2 - 1$$

- The function we optimize will be:

$$f(v) = \langle v, Mv \rangle$$

Because the sphere is compact, an extrema must exist.

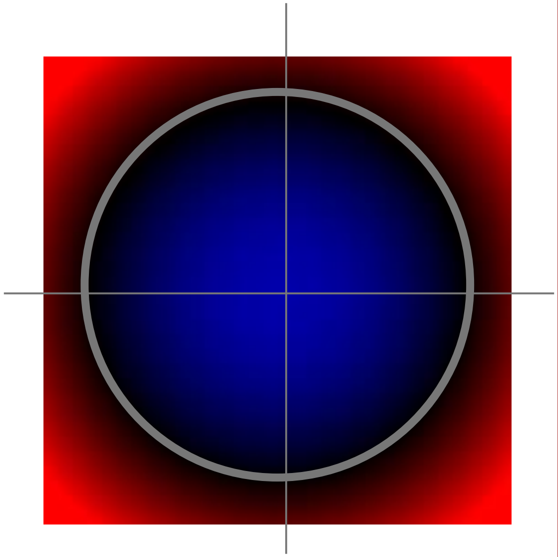

$g(x,y) = x^2 + y^2 - 1$

Diagonalizing Symmetric Matrices

2. At least one eigenvector must exist

The normal of a point on the sphere is parallel to the gradient of g:

$$\nabla g(x_1, \cdots, x_n) = 2(x_1, \cdots, x_n)$$

$$\Updownarrow$$

$$\nabla g(v) = 2v$$

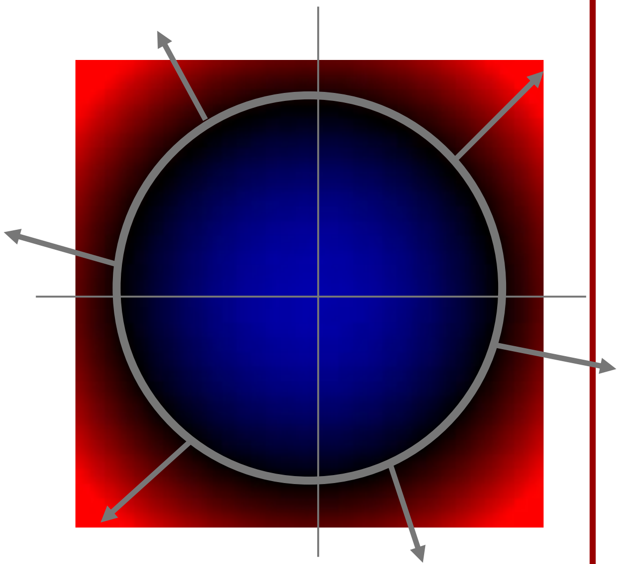

$g(x, y) = x^2 + y^2 - 1$

Diagonalizing Symmetric Matrices

2. At least one eigenvector must exist

Claim:

The gradient of f is:

$$\nabla f(v) = 2Mv$$

Diagonalizing Symmetric Matrices

2. At least one eigenvector must exist

<u>Proof</u>:

Let e_i be the vector with zeros everywhere but in the i-th entry:

$$e_i = (0, \cdots, 0, 1, 0, \cdots, 0)$$

i-th entry

Diagonalizing Symmetric Matrices

2. At least one eigenvector must exist

Proof:

Let e_i be the vector with zeros everywhere but in the i-th entry. Then the i-th coefficient of the gradient is:

$$\frac{\partial}{\partial x_i} f(v) = \frac{d}{ds}\bigg|_{s=0} f(v + se_i)$$
$$= \frac{d}{ds}\bigg|_{s=0} (v + se_i)^t M (v + se_i)$$
$$= \frac{d}{ds}\bigg|_{s=0} v^t M v + s e_i^t M v + s v^t M e_i + s^2 e_i^t M e_i$$

Diagonalizing Symmetric Matrices

2. At least one eigenvector must exist

Proof:

Let e_i be the vector with zeros everywhere but in the i-th entry. Then the i-th coefficient of the gradient is:

$$\frac{\partial}{\partial x_i} f(v) = e_i^t M v + v^t M e_i$$
$$= \langle e_i, Mv \rangle + \langle M^t v, e_i \rangle$$
$$= 2\langle e_i, Mv \rangle$$
$$= 2(Mv)_i$$

Diagonalizing Symmetric Matrices

2. At least one eigenvector must exist

<u>Proof</u>:

Since, the i-th coefficient of the gradient of f at v is twice the i-th coefficient of Mv, we have:

$$\nabla f(v) = 2Mv$$

Diagonalizing Symmetric Matrices

2. At least one eigenvector must exist

We know that the normal of the point v on the unit sphere is parallel to the gradient, which is:

$$\nabla g(v) = 2v$$

And we know that the gradient of f is:

$$\nabla f(v) = 2Mv$$

Diagonalizing Symmetric Matrices

2. At least one eigenvector must exist

Since the function g must have a maximum on the sphere, we know that there exists v s.t.:

$$\lambda \nabla g(v) = \nabla f(v)$$
$$\Updownarrow$$
$$\lambda v = Mv$$

So at the maximum, we have our eigenvalue.

Outline

Math
- Symmetric/Hermitian Matrices
- Lagrange Multipliers
- Diagonalizing Symmetric Matrices

The Laplacian Operator

The Laplacian Operator

<u>Recall</u>:

The Laplacian of a function f at a point measures how similar the value of f at the point is to the average values of its neighbors.

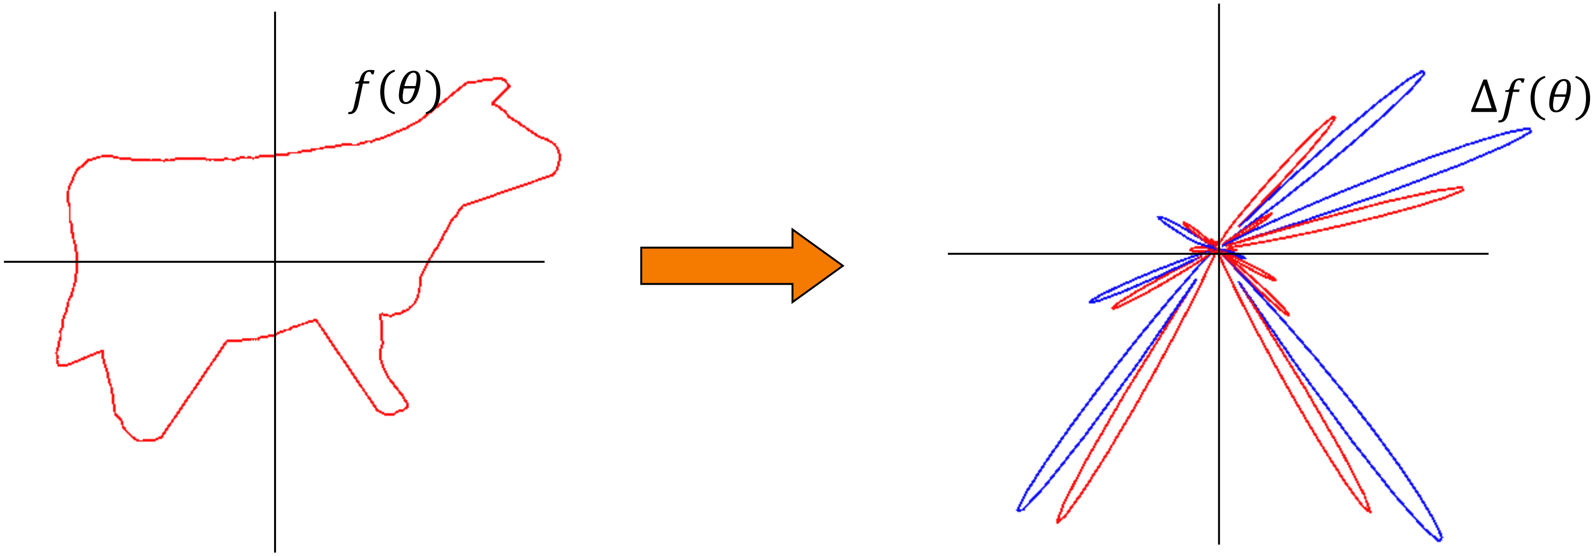

The Laplacian Operator

Recall:

Formally, for a function in 2D, the Laplacian is the sum of unmixed partial second derivatives:

$$\Delta f(x,y) = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$$

The Laplacian Operator

Observation 1:

The Laplacian is a self-adjoint operator.

To show this, we need to show that for any two functions f and g, we have:

$$\langle f, \Delta g \rangle = \langle \Delta f, g \rangle$$

The Laplacian Operator

Observation 1:

First, we show this in the 1D case, for functions $f(\theta)$ and $g(\theta)$:

$$\langle f, g'' \rangle = \langle f'', g \rangle$$

Writing the dot-product as an integral gives:

$$\langle f, g'' \rangle = \int_0^{2\pi} f(\theta) \cdot g''(\theta)\, d\theta$$

The Laplacian Operator

<u>Observation 1</u>:

Using the product rule for derivatives:

$$(f \cdot g)' = f' \cdot g + f \cdot g'$$

$$\Downarrow$$

$$\int_0^{2\pi} (f \cdot g)'(\theta)\, d\theta = \int_0^{2\pi} f'(\theta) \cdot g(\theta)\, d\theta + \int_0^{2\pi} f(\theta) \cdot g'(\theta)\, d\theta$$

Since f and g are functions on a circle, their values at 0 and 2π are the same:

$$\int_0^{2\pi} (f \cdot g)'(\theta)\, d\theta = (f \cdot g)(2\pi) - (f \cdot g)(0) = 0$$

The Laplacian Operator

Observation 1:

Thus, we have:

$$\int_0^{2\pi} f(\theta)\cdot g'(\theta)\,d\theta = -\int_0^{2\pi} f'(\theta)\cdot g(\theta)\,d\theta$$

"Moving" the derivative twice gives:

$$\int_0^{2\pi} f''(\theta)\cdot g(\theta)\,d\theta = -\int_0^{2\pi} f'(\theta)\cdot g'(\theta)\,d\theta$$

$$= (-1)^2\int_0^{2\pi} f(\theta)\cdot g''(\theta)\,d\theta$$

$$\Updownarrow$$

$$\langle f'', g\rangle = \langle f, g''\rangle$$

The Laplacian Operator

Observation 1:

To generalize this to higher dimensions, we write out the dot-product as:

$$\langle \Delta f, g \rangle = \int_0^{2\pi}\int_0^{2\pi} \frac{\partial^2 f}{\partial \theta^2} \cdot g \, d\theta \;\, d\phi + \int_0^{2\pi}\int_0^{2\pi} \frac{\partial^2 f}{\partial \phi^2} \cdot g \, d\phi \;\, d\theta$$

$$= \int_0^{2\pi}\int_0^{2\pi} f \cdot \frac{\partial^2 g}{\partial \theta^2} \, d\theta \;\, d\phi + \int_0^{2\pi}\int_0^{2\pi} f \cdot \frac{\partial^2 g}{\partial \phi^2} \, d\phi \;\, d\theta$$

$$= \langle f, \Delta g \rangle$$

The Laplacian Operator

Observation 2:

The Laplacian operator commutes with rotation – i.e. computing the Laplacian and rotating gives the same function as first rotating and then computing the Laplacian:

$$\Delta\big(\rho_R(f)\big) = \rho_R\big(\Delta(f)\big)$$

The Laplacian Operator

Implications:

- **Observation 1**: Since the Laplacian operator is self-adjoint, it must be diagonalizable.

  $\Rightarrow$ There is an orthogonal basis of eigenvectors.

  $\Rightarrow$ If we group the eigenvectors with the same eigenvalues together, we get a set of vector spaces F_λ such that for any function $f \in F_\lambda$:

$$\Delta f = \lambda f$$

The Laplacian Operator

Implications:

- **Observation 2**: Since the Laplacian operator commutes with rotation, rotations map vectors in F_λ back into F_λ.

$$\begin{aligned}\Delta\big(\rho_R(f)\big) &= \rho_R\big(\Delta(f)\big)\\ &= \rho_R(\lambda f)\\ &= \lambda\big(\rho_R(f)\big)\end{aligned}$$

$\Rightarrow$ The space F_λ fixed under the action of rotation.

$\Rightarrow$ The space F_λ is a sub-representations for the group of rotation.

The Laplacian Operator

Going back to the problem of finding the irreducible representations, this means we can begin by looking for the eigenspaces of the Laplacian operator.

Computing the Laplacian

We know how to compute the Laplacian of a circular function represented by parameter:

$$\Delta f(\theta) = f''(\theta)$$

How do we compute the Laplacian for a function represented by restriction?

Computing the Laplacian

If we define a function on a circle as the restriction of a 2D function, the 2D Laplacian is not the same as the circular Laplacian!

<u>Example</u>:

Consider the function $f(x,y) = x$.

- In the plane, the Laplacian is:
$$\Delta f(x,y) = 0$$
- On the circle this is the function $f(\theta) = \cos(\theta)$:
$$\Delta f(\theta) = -\cos(\theta)$$

Computing the Laplacian

If we define a function on a circle as the restriction of a 2D function, the 2D Laplacian is not the same as the circular Laplacian!

Intuitively:

The Laplacian measures the difference between the value of a point and the average value of the "neighbors".

Who the "neighbors" are changes depending on whether we are considering the plane or the circle.

Computing the Laplacian

Recall:

For a vector field:

$$\vec{F}(x,y) = \big(F_1(x,y), F_2(x,y)\big)$$

the divergence is defined:

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

We can also express the Laplacian as the divergence of the gradient:

$$\Delta f = \nabla \cdot (\nabla f)$$

Computing the Gradient

In general, the gradient of the function $f(x,y)$ need not lie along the unit-circle.

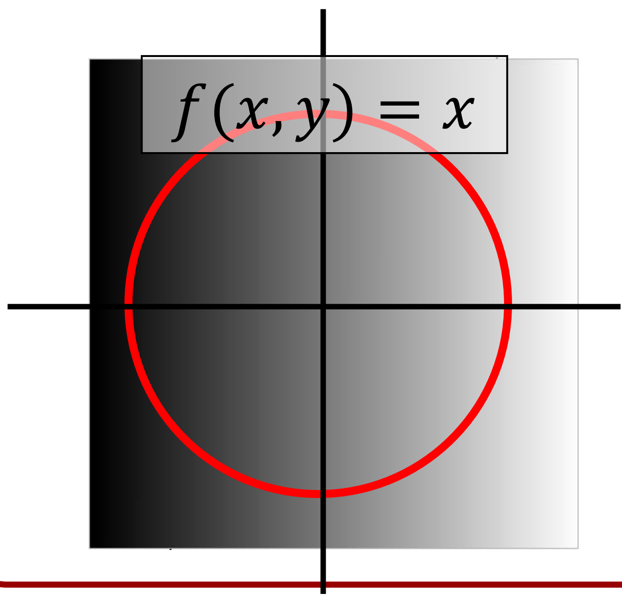

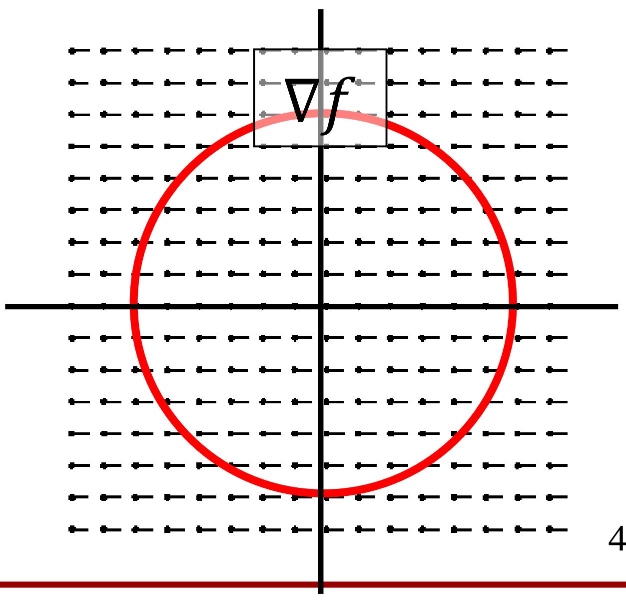

Computing the Gradient

In general, the gradient of the function $f(x, y)$ need not lie along the unit-circle.

We can fix this by projecting the gradient on to the unit circle:

$$\pi(\nabla f) = \nabla f - \langle \nabla f, (x, y) \rangle (x, y)$$

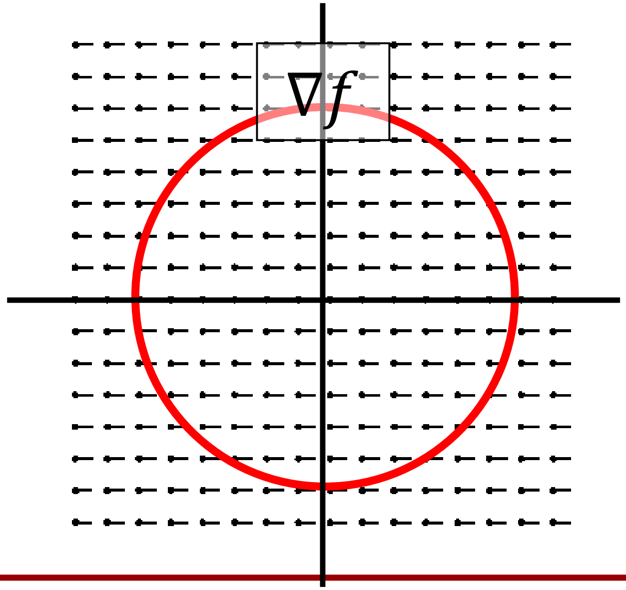

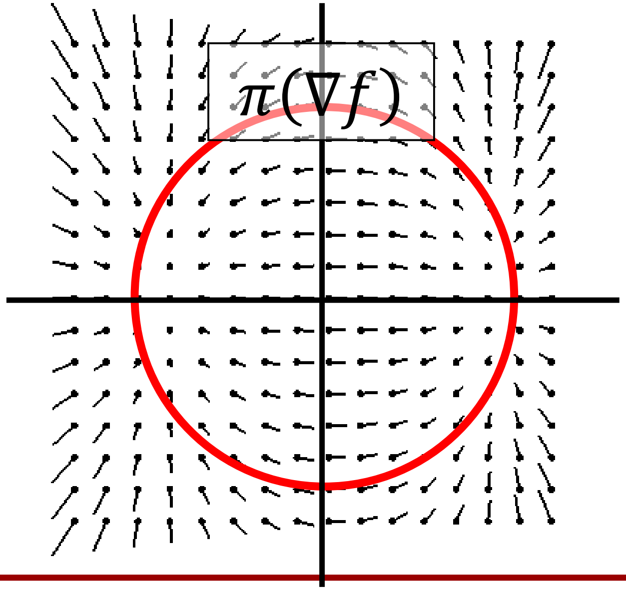

Computing the Laplacian

The divergence of a vector field $\vec{F}$ can be expressed as the sum of partials:

$$\operatorname{div}\left(\vec{F}\right) = \nabla \cdot \vec{F} = \frac{\partial F_1}{\partial x} + \frac{\partial F_2}{\partial y}$$

Computing the Laplacian

Given any orthogonal basis $\{v, w\}$, the divergence is the derivative of the v-component of the vector field in the v-direction, plus the derivative of the w-component of the vector field in the w-direction:

$$\operatorname{div}\left(\vec{F}\right) = \nabla \cdot \vec{F} = \frac{\partial \langle \vec{F}, v \rangle}{\partial v} + \frac{\partial \langle \vec{F}, w \rangle}{\partial w}$$

Computing the Laplacian

Thus, to compute the divergence of the vector field along the circle, we can compute the 2D divergence, and subtract off the contribution from the normal direction:

$$\mathrm{div}_{\mathrm{circle}}\left(\vec{F}\right) = \mathrm{div}_{\mathrm{2D}}\left(\vec{F}\right) - \frac{\partial\langle\vec{F},n\rangle}{\partial n}$$

Since the component of $\vec{F}$ in the normal direction is a scalar function, its derivative in the normal direction can be expressed as a gradient:

$$\frac{\partial\langle\vec{F},n\rangle}{\partial n} = \langle\nabla\langle\vec{F},n\rangle,n\rangle$$

Computing the Laplacian ($f(x,y)=x$)

Example:

Consider the function:

$$f(x,y)=x$$

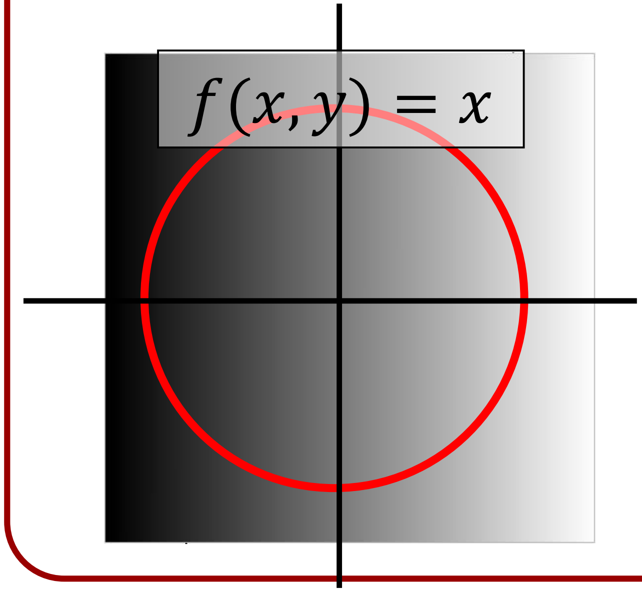

Computing the Laplacian ($f(x, y) = x$)

Example:

Its gradient is:

$$\nabla f = (1,0)$$

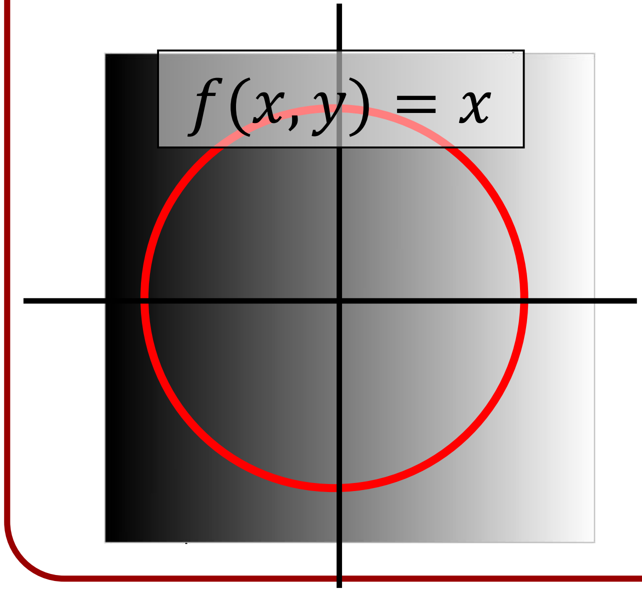

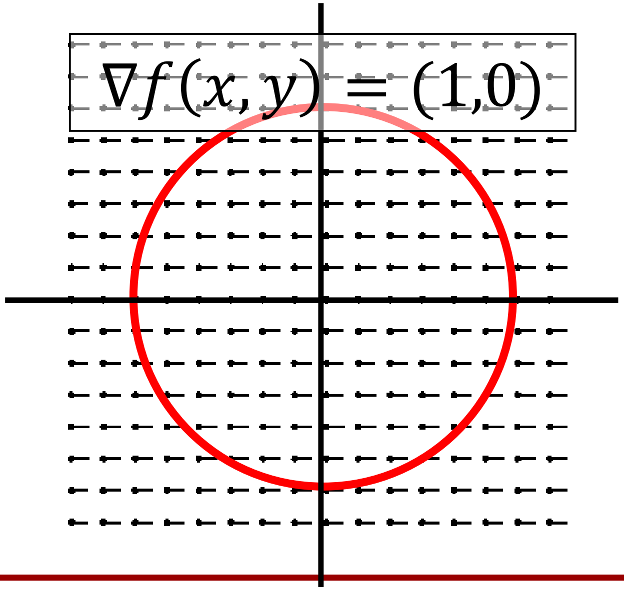

Computing the Laplacian ($f(x,y) = x$)

Example: $\nabla f = (1,0)$

Projecting the gradient onto the unit-circle we get:

$$\begin{aligned}\pi(\nabla f) &= \nabla f - \langle \nabla f, n \rangle n \\ &= \nabla f - \langle \nabla f, (x,y) \rangle (x,y) \\ &= (1,0) - x(x,y)\end{aligned}$$

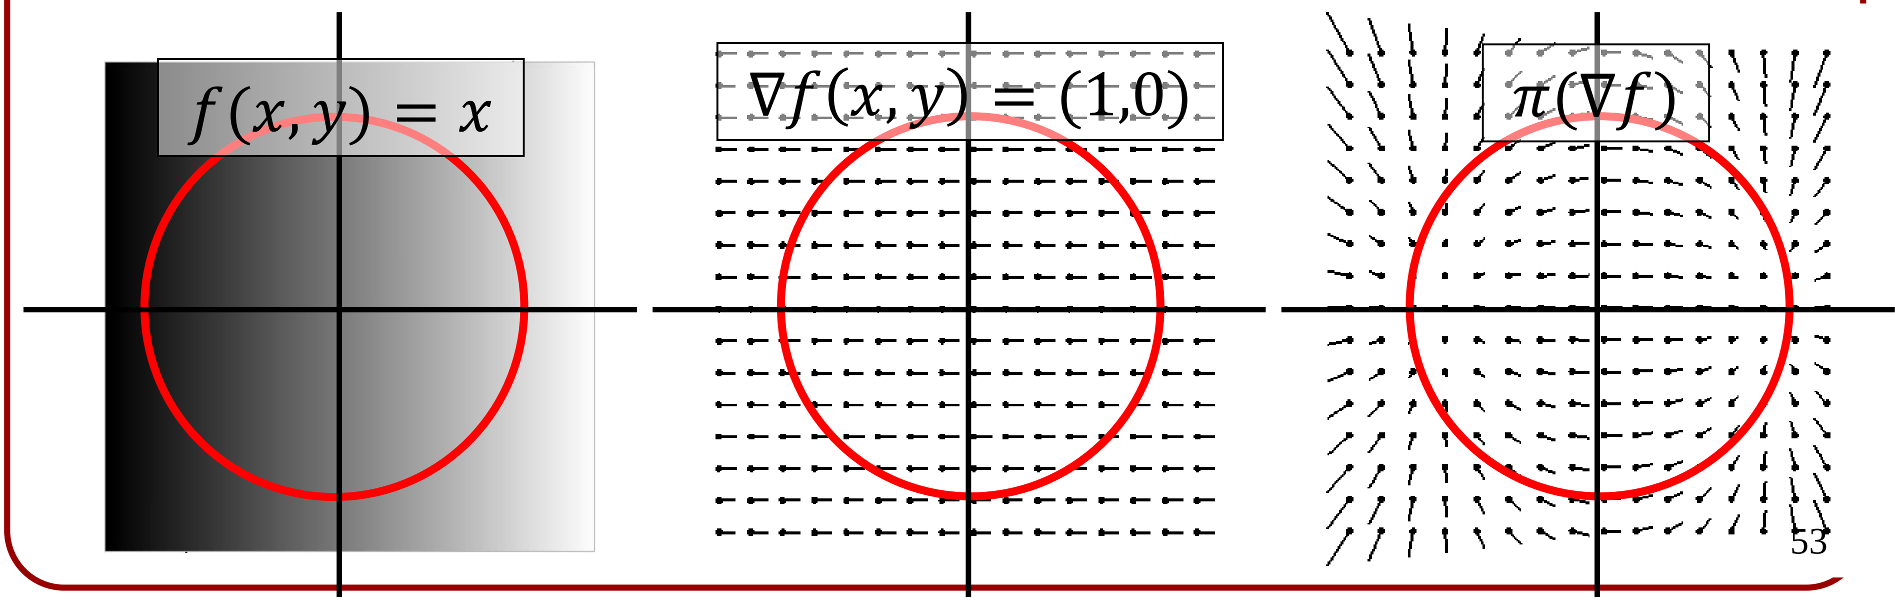

Computing the Laplacian ($f(x,y) = x$)

Example: $\pi(\nabla f) = (1,0) - x(x,y)$

The divergence of the vector field $\pi(\nabla f)$ is:

$$\begin{aligned}\mathrm{div}_{2D}\big(\pi(\nabla f)\big) &= -2x - x \\ &= -3x\end{aligned}$$

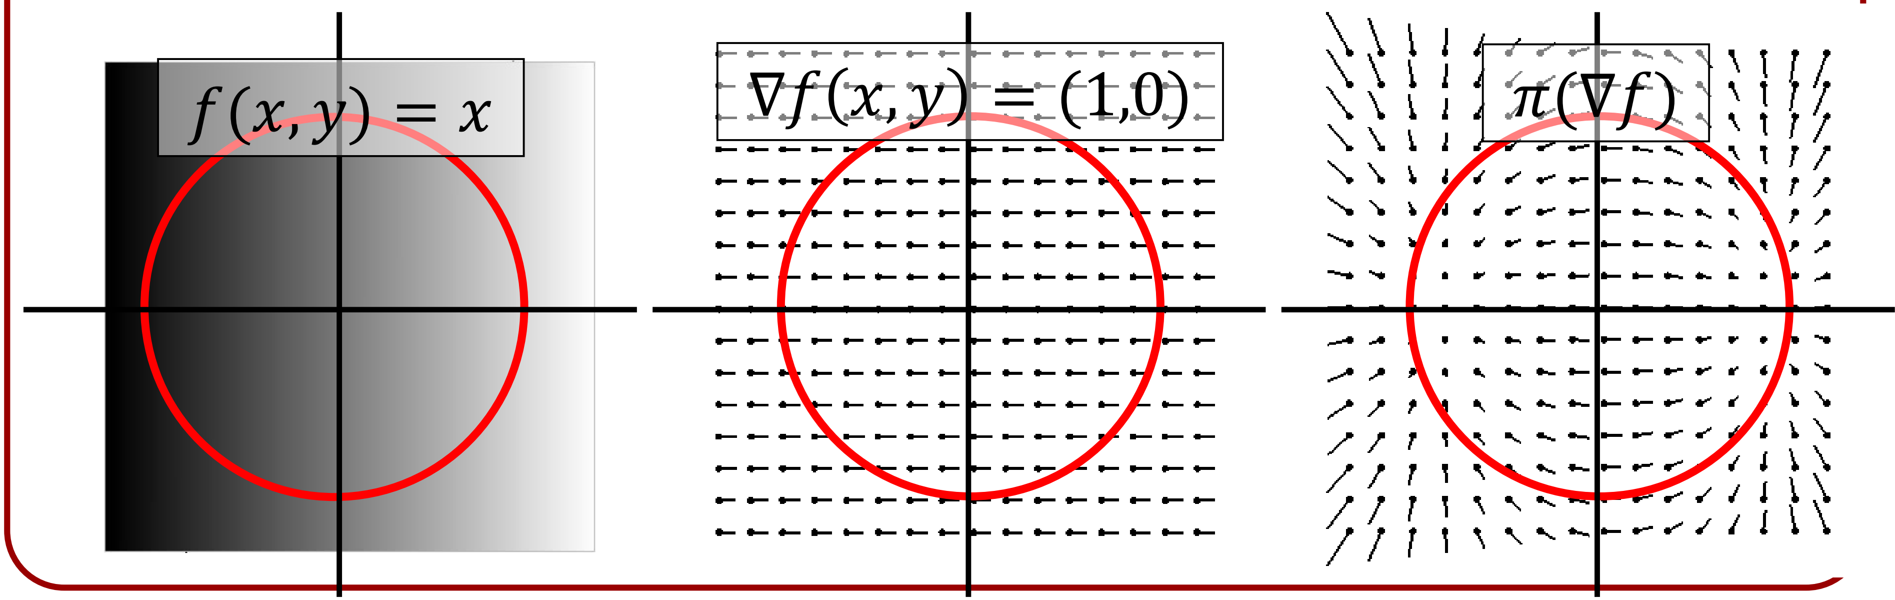

Computing the Laplacian ($f(x, y) = x$)

Example: $\pi(\nabla f) = (1,0) - x(x, y)$

Projecting $\pi(\nabla f)$ on the normal direction gives:

$$\begin{aligned}\langle \pi(\nabla f), n \rangle &= \langle (1,0) - x(x, y), (x, y) \rangle \\ &= x - x(x^2 + y^2) \\ &= x - x^3 + xy^2\end{aligned}$$

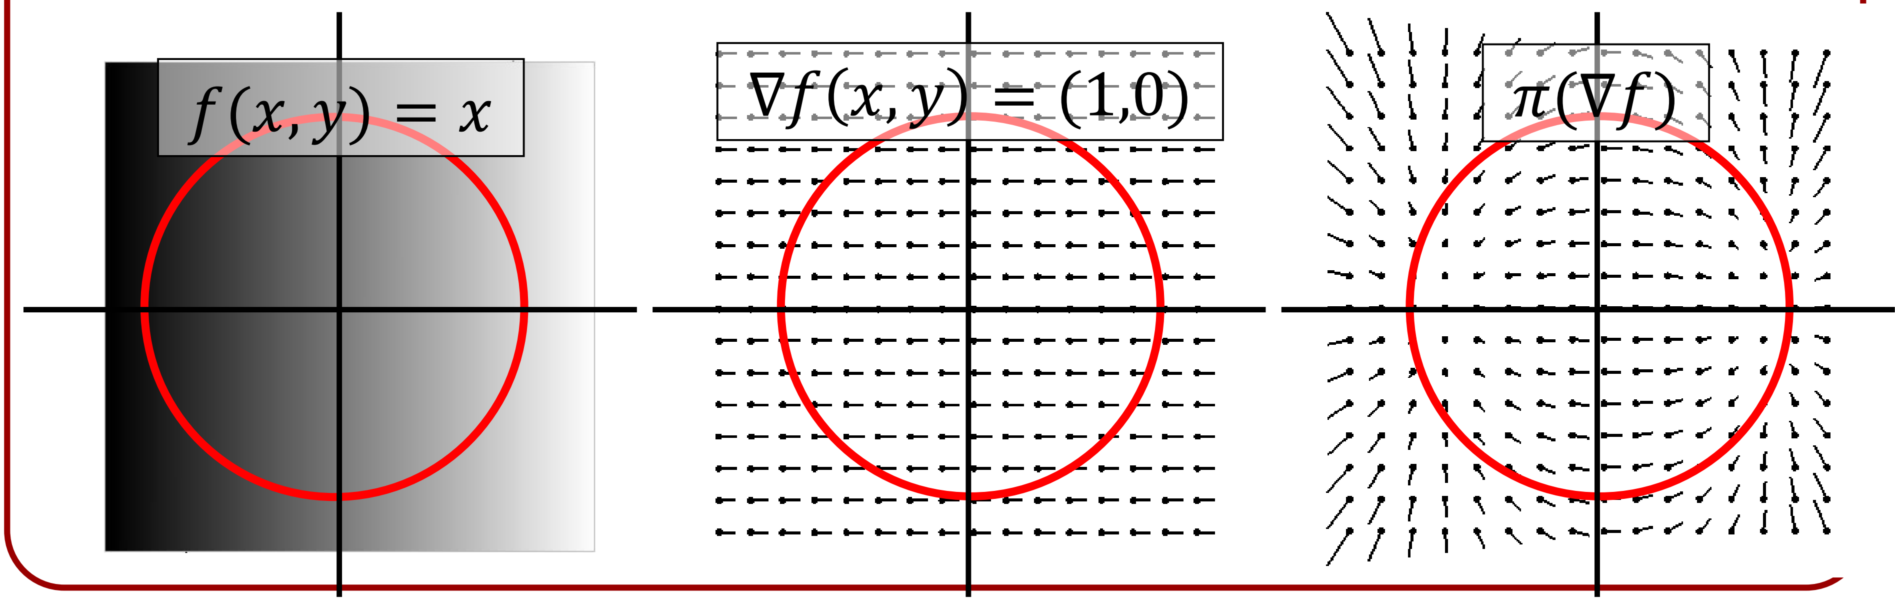

Computing the Laplacian ($f(x,y)=x$)

Example: $\langle\pi(\nabla f),n\rangle = x - x^3 + xy^2$

The gradient of the projection is:

$$\nabla\langle\pi(\nabla f),n\rangle = (1,0) - (3x^2+y^2,2xy)$$

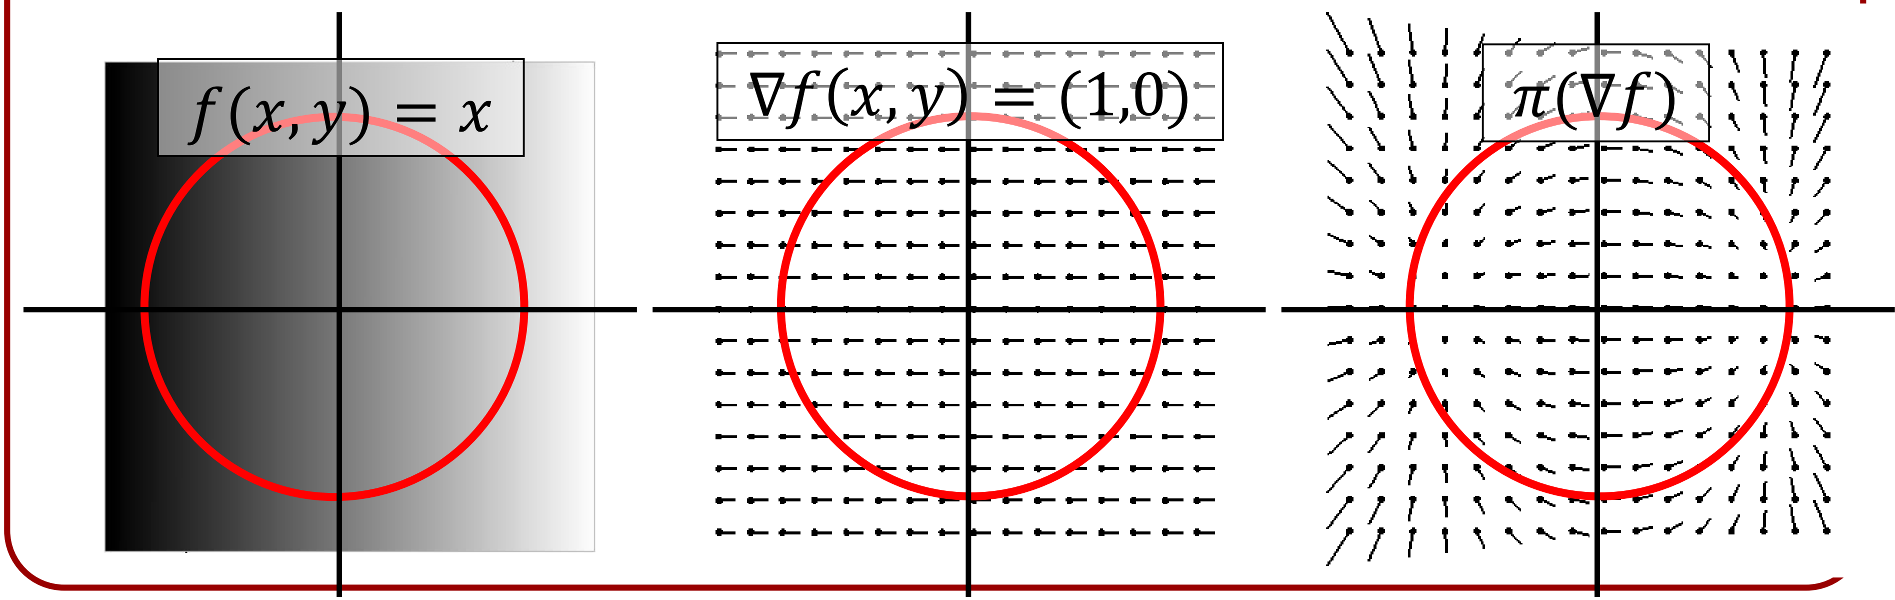

Computing the Laplacian ($f(x,y) = x$)

<u>Example</u>: $\nabla\langle\pi(\nabla f), n\rangle = (1,0) - (3x^2 + y^2, 2xy)$

So the divergence in the normal direction is:

$$\begin{aligned}\text{div}_n\big(\pi(\nabla f)\big) &= \langle(1,0) - (3x^2 + y^2, 2xy), (x,y)\rangle \\ &= x - 3x - xy^2 - 2xy^2 \\ &= x - 3x^3 - 3xy^2\end{aligned}$$

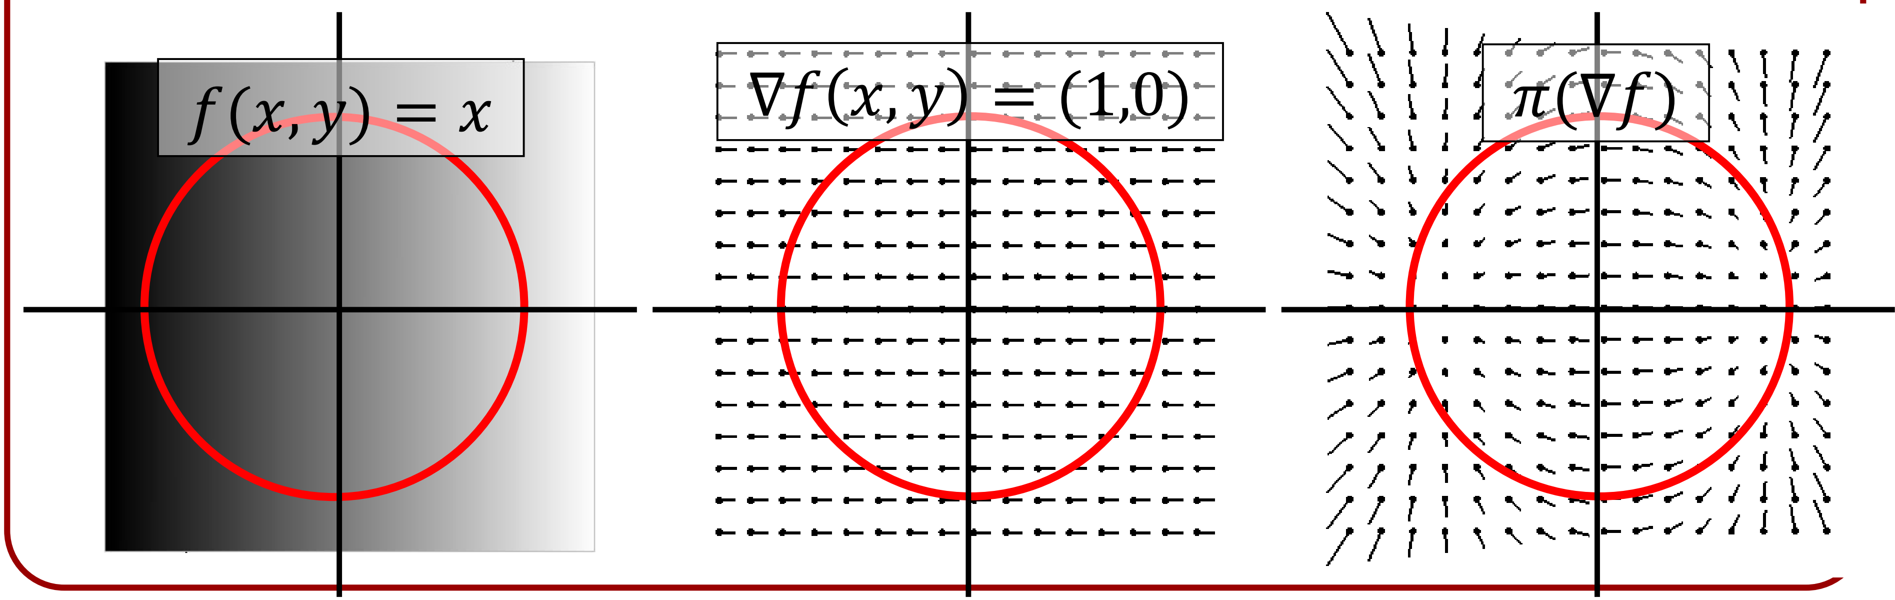

Computing the Laplacian ($f(x,y) = x$)

Example:

$$\text{div}_{\text{2D}}\big(\pi(\nabla f)\big) = -3x \qquad \text{div}_n\big(\pi(\nabla f)\big) = x - 3x^3 - 3xy^2$$

Thus the circular Laplacian can be expressed as the difference between the 2D divergence and the divergence in the normal direction:

$$\begin{aligned}\Delta_{\text{circle}} f(x,y) &= \text{div}_{2D}\big(\pi(\nabla f)\big) - \text{div}_n\big(\pi(\nabla f)\big)\\ &= -3x - (x - 3x^3 - 3xy^2)\\ &= -4x + 3x(x^2 + y^2)\end{aligned}$$

Since points on the circle satisfy $x^2 + y^2 = 1$, this implies that for (x,y) on the circle:

$$\Delta_{\text{circle}} f(x,y) = -x$$

Computing the Laplacian ($f(x,y) = x$)

Example:

$$\mathrm{div}_{2\mathrm{D}}\big(\pi(\nabla f)\big) = -3x \qquad \mathrm{div}_n\big(\pi(\nabla f)\big) = x - 3x^3 - 3xy^2$$

Thus the circular Laplacian can be expressed as the difference between the 2D divergence and the divergence in the normal direction:

$$\begin{aligned}\Delta_{\mathrm{circle}} f(x,y) &= \mathrm{div}_{2D}\big(\pi(\nabla f)\big) - \mathrm{div}_n\big(\pi(\nabla f)\big)\\ &= -3x - (x - 3x^3 - 3xy^2)\\ &= -4x + 3x(x^2+y^2)\end{aligned}$$

Since points on the circle satisfy $x^2 + y^2 = 1$, this implies that for (x,y) on the circle:

$$\Delta_{\mathrm{circle}} f(x,y) = -f(x,y)$$

Computing the Laplacian ($\boldsymbol{f(x,y)=x}$)

Example:

Thus just as in the parameter case the function f, is an eigenvector of the circular Laplacian operator, with eigenvalue -1.