Lecture 40 - UM 102 (Spring 2021)

Vamsi Pritham Pingali

IISc

Recap

Recap

- Interpretation of curl, conservative vector fields.

Recap

- Interpretation of curl, conservative vector fields.
- Divergence theorem, examples.

Recap

- Interpretation of curl, conservative vector fields.
- Divergence theorem, examples.
- Interpretation of divergence.

Constrained optimisation

Constrained optimisation

- Most real world problems

Constrained optimisation

- Most real world problems require the "best answer"

Constrained optimisation

- Most real world problems require the “best answer” subject to some constraints.

Constrained optimisation

- Most real world problems require the “best answer” subject to some constraints.
- A pure maths example:

Constrained optimisation

- Most real world problems require the “best answer” subject to some constraints.
- A pure maths example: Find the maximum possible value of

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$.

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$.
- Firstly, the constraining set

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$.
- Firstly, the constraining set is closed (it is a sphere)

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded.

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded. $x + y + z$ being a continuous function

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded. $x + y + z$ being a continuous function attains its maximum and minimum somewhere.

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x+y+z$ subject to $x^2+y^2+z^2=1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded. $x+y+z$ being a continuous function attains its maximum and minimum somewhere. So the question makes sense.

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded. $x + y + z$ being a continuous function attains its maximum and minimum somewhere. So the question makes sense.
- One way to solve it:

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x+y+z$ subject to $x^2+y^2+z^2=1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded. $x+y+z$ being a continuous function attains its maximum and minimum somewhere. So the question makes sense.
- One way to solve it: Solve for $z=\pm\sqrt{1-x^2-y^2}$ and $x^2+y^2\leq 1$.

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded. $x + y + z$ being a continuous function attains its maximum and minimum somewhere. So the question makes sense.
- One way to solve it: Solve for $z = \pm\sqrt{1 - x^2 - y^2}$ and $x^2 + y^2 \leq 1$. Since we want the maximum of $x + y + z$,

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded. $x + y + z$ being a continuous function attains its maximum and minimum somewhere. So the question makes sense.
- One way to solve it: Solve for $z = \pm\sqrt{1 - x^2 - y^2}$ and $x^2 + y^2 \leq 1$. Since we want the maximum of $x + y + z$, we choose $f(x, y) = x + y + \sqrt{1 - x^2 - y^2}$ over $x^2 + y^2 \leq 1$.

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded. $x + y + z$ being a continuous function attains its maximum and minimum somewhere. So the question makes sense.
- One way to solve it: Solve for $z = \pm\sqrt{1 - x^2 - y^2}$ and $x^2 + y^2 \leq 1$. Since we want the maximum of $x + y + z$, we choose $f(x, y) = x + y + \sqrt{1 - x^2 - y^2}$ over $x^2 + y^2 \leq 1$. This is an unconstrained problem.

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x+y+z$ subject to $x^2+y^2+z^2=1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded. $x+y+z$ being a continuous function attains its maximum and minimum somewhere. So the question makes sense.
- One way to solve it: Solve for $z=\pm\sqrt{1-x^2-y^2}$ and $x^2+y^2\leq 1$. Since we want the maximum of $x+y+z$, we choose $f(x,y)=x+y+\sqrt{1-x^2-y^2}$ over $x^2+y^2\leq 1$. This is an unconstrained problem. We look at critical points on the interior:

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x+y+z$ subject to $x^2+y^2+z^2=1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded. $x+y+z$ being a continuous function attains its maximum and minimum somewhere. So the question makes sense.
- One way to solve it: Solve for $z=\pm\sqrt{1-x^2-y^2}$ and $x^2+y^2\leq 1$. Since we want the maximum of $x+y+z$, we choose $f(x,y)=x+y+\sqrt{1-x^2-y^2}$ over $x^2+y^2\leq 1$. This is an unconstrained problem. We look at critical points on the interior: $f_x=1-\frac{x}{\sqrt{1-x^2-y^2}}=0=f_y=1-\frac{y}{\sqrt{1-x^2-y^2}}$.

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x+y+z$ subject to $x^2+y^2+z^2=1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded. $x+y+z$ being a continuous function attains its maximum and minimum somewhere. So the question makes sense.
- One way to solve it: Solve for $z=\pm\sqrt{1-x^2-y^2}$ and $x^2+y^2\leq 1$. Since we want the maximum of $x+y+z$, we choose $f(x,y)=x+y+\sqrt{1-x^2-y^2}$ over $x^2+y^2\leq 1$. This is an unconstrained problem. We look at critical points on the interior: $f_x=1-\frac{x}{\sqrt{1-x^2-y^2}}=0=f_y=1-\frac{y}{\sqrt{1-x^2-y^2}}$. Thus $x=y=\frac{1}{\sqrt{3}}$. $f(\frac{1}{\sqrt{3}})=\sqrt{3}$.

Constrained optimisation

- Most real world problems require the "best answer" subject to some constraints.
- A pure maths example: Find the maximum possible value of $x+y+z$ subject to $x^2+y^2+z^2=1$.
- Firstly, the constraining set is closed (it is a sphere) and bounded. $x+y+z$ being a continuous function attains its maximum and minimum somewhere. So the question makes sense.
- One way to solve it: Solve for $z=\pm\sqrt{1-x^2-y^2}$ and $x^2+y^2\leq 1$. Since we want the maximum of $x+y+z$, we choose $f(x,y)=x+y+\sqrt{1-x^2-y^2}$ over $x^2+y^2\leq 1$. This is an unconstrained problem. We look at critical points on the interior: $f_x=1-\frac{x}{\sqrt{1-x^2-y^2}}=0=f_y=1-\frac{y}{\sqrt{1-x^2-y^2}}$. Thus $x=y=\frac{1}{\sqrt{3}}$. $f(\frac{1}{\sqrt{3}})=\sqrt{3}$. We now look at the maximum of the function $f(x,y)$ on the boundary $x^2+y^2=1$.

Constrained optimisation

Constrained optimisation

Constrained optimisation

- This is another constrained problem.

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$.

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$.

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$.

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$.

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$. Thus the maximum possible value of $x + y + z$

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$. Thus the maximum possible value of $x+y+z$ subject to $x^2+y^2+z^2 = 1$ occurs at $x = y = z = \frac{1}{\sqrt{3}}$ and equals $\sqrt{3}$.

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$. Thus the maximum possible value of $x+y+z$ subject to $x^2+y^2+z^2 = 1$ occurs at $x = y = z = \frac{1}{\sqrt{3}}$ and equals $\sqrt{3}$.
- This strategy (besides being

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$. Thus the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$ occurs at $x = y = z = \frac{1}{\sqrt{3}}$ and equals $\sqrt{3}$.
- This strategy (besides being the equivalent of self-flagellation)

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$. Thus the maximum possible value of $x+y+z$ subject to $x^2+y^2+z^2 = 1$ occurs at $x = y = z = \frac{1}{\sqrt{3}}$ and equals $\sqrt{3}$.
- This strategy (besides being the equivalent of self-flagellation) cannot work if

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$. Thus the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$ occurs at $x = y = z = \frac{1}{\sqrt{3}}$ and equals $\sqrt{3}$.
- This strategy (besides being the equivalent of self-flagellation) cannot work if we cannot solve for z explicitly.

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$. Thus the maximum possible value of $x+y+z$ subject to $x^2+y^2+z^2 = 1$ occurs at $x = y = z = \frac{1}{\sqrt{3}}$ and equals $\sqrt{3}$.
- This strategy (besides being the equivalent of self-flagellation) cannot work if we cannot solve for z explicitly.
- Here is a geometric picture:

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$. Thus the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$ occurs at $x = y = z = \frac{1}{\sqrt{3}}$ and equals $\sqrt{3}$.
- This strategy (besides being the equivalent of self-flagellation) cannot work if we cannot solve for z explicitly.
- Here is a geometric picture: $x + y + z = c$ is a plane at a distance of $\frac{c}{\sqrt{3}}$ from the origin.

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1 - x^2}$ and $g(x) = x + \sqrt{1 - x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$. Thus the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$ occurs at $x = y = z = \frac{1}{\sqrt{3}}$ and equals $\sqrt{3}$.
- This strategy (besides being the equivalent of self-flagellation) cannot work if we cannot solve for z explicitly.
- Here is a geometric picture: $x + y + z = c$ is a plane at a distance of $\frac{c}{\sqrt{3}}$ from the origin. So we want to find a plane with a given normal

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$. Thus the maximum possible value of $x+y+z$ subject to $x^2+y^2+z^2 = 1$ occurs at $x = y = z = \frac{1}{\sqrt{3}}$ and equals $\sqrt{3}$.
- This strategy (besides being the equivalent of self-flagellation) cannot work if we cannot solve for z explicitly.
- Here is a geometric picture: $x+y+z = c$ is a plane at a distance of $\frac{c}{\sqrt{3}}$ from the origin. So we want to find a plane with a given normal that is farthest from the origin

Constrained optimisation

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$. Thus the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$ occurs at $x = y = z = \frac{1}{\sqrt{3}}$ and equals $\sqrt{3}$.
- This strategy (besides being the equivalent of self-flagellation) cannot work if we cannot solve for z explicitly.
- Here is a geometric picture: $x + y + z = c$ is a plane at a distance of $\frac{c}{\sqrt{3}}$ from the origin. So we want to find a plane with a given normal that is farthest from the origin but intersects the sphere somewhere.

- This is another constrained problem. Again, $y = \sqrt{1-x^2}$ and $g(x) = x + \sqrt{1-x^2}$ on $-1 \leq x \leq 1$. Its interior critical points correspond to $g'(x) = 1 - \frac{x}{\sqrt{1-x^2}} = 0$ and hence $x = \frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}}) = \sqrt{2}$. The boundary values are $g(1) = 1, g(-1) = -1$. Thus the maximum possible value of $x + y + z$ subject to $x^2 + y^2 + z^2 = 1$ occurs at $x = y = z = \frac{1}{\sqrt{3}}$ and equals $\sqrt{3}$.
- This strategy (besides being the equivalent of self-flagellation) cannot work if we cannot solve for z explicitly.
- Here is a geometric picture: $x + y + z = c$ is a plane at a distance of $\frac{c}{\sqrt{3}}$ from the origin. So we want to find a plane with a given normal that is farthest from the origin but intersects the sphere somewhere. Clearly the tangent plane at some point is

Constrained optimisation

- This is another constrained problem. Again, $y=\sqrt{1-x^2}$ and $g(x)=x+\sqrt{1-x^2}$ on $-1\leq x\leq 1$. Its interior critical points correspond to $g'(x)=1-\frac{x}{\sqrt{1-x^2}}=0$ and hence $x=\frac{1}{\sqrt{2}}$. $g(\frac{1}{\sqrt{2}})=\sqrt{2}$. The boundary values are $g(1)=1, g(-1)=-1$. Thus the maximum possible value of $x+y+z$ subject to $x^2+y^2+z^2=1$ occurs at $x=y=z=\frac{1}{\sqrt{3}}$ and equals $\sqrt{3}$.
- This strategy (besides being the equivalent of self-flagellation) cannot work if we cannot solve for z explicitly.
- Here is a geometric picture: $x+y+z=c$ is a plane at a distance of $\frac{c}{\sqrt{3}}$ from the origin. So we want to find a plane with a given normal that is farthest from the origin but intersects the sphere somewhere. Clearly the tangent plane at some point is answer.

The set up of the problem

The set up of the problem

- Def:

The set up of the problem

- Def: Let f be a function from a

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x, y, z) = 0$ of a C^1 function to $\mathbb{R}$.

- Def: Let f be a function from a neighbourhood of the level-set $g(x, y, z) = 0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x, y, z) = 0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0, y_0, z_0) lying on the level set, subject to the constraint $g = 0$

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x,y,z)=0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0,y_0,z_0) lying on the level set, subject to the constraint $g=0$ if g is defined on a neighbourhood of (x_0,y_0,z_0),

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x,y,z)=0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0,y_0,z_0) lying on the level set, subject to the constraint $g=0$ if g is defined on a neighbourhood of (x_0,y_0,z_0), $\nabla g(x_0,y_0,z_0)\neq 0$,

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x, y, z) = 0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0, y_0, z_0) lying on the level set, subject to the constraint $g = 0$ if g is defined on a neighbourhood of (x_0, y_0, z_0), $\nabla g(x_0, y_0, z_0) \neq 0$, and $f(x, y, z) \leq f(x_0, y_0, z_0)$ for all

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x, y, z) = 0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0, y_0, z_0) lying on the level set, subject to the constraint $g = 0$ if g is defined on a neighbourhood of (x_0, y_0, z_0), $\nabla g(x_0, y_0, z_0) \neq 0$, and $f(x, y, z) \leq f(x_0, y_0, z_0)$ for all (x, y, z) in a neighbourhood of (x_0, y_0, z_0) lying on the level set.

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x, y, z) = 0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0, y_0, z_0) lying on the level set, subject to the constraint $g = 0$ if g is defined on a neighbourhood of (x_0, y_0, z_0), $\nabla g(x_0, y_0, z_0) \neq 0$, and $f(x, y, z) \leq f(x_0, y_0, z_0)$ for all (x, y, z) in a neighbourhood of (x_0, y_0, z_0) lying on the level set. Likewise for a local minimum.

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x, y, z) = 0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0, y_0, z_0) lying on the level set, subject to the constraint $g = 0$ if g is defined on a neighbourhood of (x_0, y_0, z_0), $\nabla g(x_0, y_0, z_0) \neq 0$, and $f(x, y, z) \leq f(x_0, y_0, z_0)$ for all (x, y, z) in a neighbourhood of (x_0, y_0, z_0) lying on the level set. Likewise for a local minimum.
- The point of

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x,y,z)=0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0,y_0,z_0) lying on the level set, subject to the constraint $g=0$ if g is defined on a neighbourhood of (x_0,y_0,z_0), $\nabla g(x_0,y_0,z_0)\neq 0$, and $f(x,y,z)\leq f(x_0,y_0,z_0)$ for all (x,y,z) in a neighbourhood of (x_0,y_0,z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0,y_0,z_0)\neq 0$ is

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x,y,z)=0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0,y_0,z_0) lying on the level set, subject to the constraint $g=0$ if g is defined on a neighbourhood of (x_0,y_0,z_0), $\nabla g(x_0,y_0,z_0)\neq 0$, and $f(x,y,z)\leq f(x_0,y_0,z_0)$ for all (x,y,z) in a neighbourhood of (x_0,y_0,z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0,y_0,z_0)\neq 0$ is the implicit function theorem:

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x,y,z)=0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0,y_0,z_0) lying on the level set, subject to the constraint $g=0$ if g is defined on a neighbourhood of (x_0,y_0,z_0), $\nabla g(x_0,y_0,z_0)\neq 0$, and $f(x,y,z)\leq f(x_0,y_0,z_0)$ for all (x,y,z) in a neighbourhood of (x_0,y_0,z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0,y_0,z_0)\neq 0$ is the implicit function theorem: If $\frac{\partial g}{\partial z}(x_0,y_0,z_0)$, then near this point,

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x,y,z)=0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0,y_0,z_0) lying on the level set, subject to the constraint $g=0$ if g is defined on a neighbourhood of (x_0,y_0,z_0), $\nabla g(x_0,y_0,z_0)\neq 0$, and $f(x,y,z)\leq f(x_0,y_0,z_0)$ for all (x,y,z) in a neighbourhood of (x_0,y_0,z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0,y_0,z_0)\neq 0$ is the implicit function theorem: If $\frac{\partial g}{\partial z}(x_0,y_0,z_0)$, then near this point, $z=f(x,y)$ for some C^1 function f, i.e.,

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x,y,z)=0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0,y_0,z_0) lying on the level set, subject to the constraint $g=0$ if g is defined on a neighbourhood of (x_0,y_0,z_0), $\nabla g(x_0,y_0,z_0)\neq 0$, and $f(x,y,z)\leq f(x_0,y_0,z_0)$ for all (x,y,z) in a neighbourhood of (x_0,y_0,z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0,y_0,z_0)\neq 0$ is the implicit function theorem: If $\frac{\partial g}{\partial z}(x_0,y_0,z_0)$, then near this point, $z=f(x,y)$ for some C^1 function f, i.e., the level set is a regular parametrised surface near this point.

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x,y,z)=0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0,y_0,z_0) lying on the level set, subject to the constraint $g=0$ if g is defined on a neighbourhood of (x_0,y_0,z_0), $\nabla g(x_0,y_0,z_0)\neq 0$, and $f(x,y,z)\leq f(x_0,y_0,z_0)$ for all (x,y,z) in a neighbourhood of (x_0,y_0,z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0,y_0,z_0)\neq 0$ is the implicit function theorem: If $\frac{\partial g}{\partial z}(x_0,y_0,z_0)$, then near this point, $z=f(x,y)$ for some C^1 function f, i.e., the level set is a regular parametrised surface near this point.
- As a consequence,

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x,y,z)=0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0,y_0,z_0) lying on the level set, subject to the constraint $g=0$ if g is defined on a neighbourhood of (x_0,y_0,z_0), $\nabla g(x_0,y_0,z_0)\neq 0$, and $f(x,y,z)\leq f(x_0,y_0,z_0)$ for all (x,y,z) in a neighbourhood of (x_0,y_0,z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0,y_0,z_0)\neq 0$ is the implicit function theorem: If $\frac{\partial g}{\partial z}(x_0,y_0,z_0)$, then near this point, $z=f(x,y)$ for some C^1 function f, i.e., the level set is a regular parametrised surface near this point.
- As a consequence, if $g(x,y,z)=0$ is a regular closed surface that is a bounded set,

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x, y, z) = 0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0, y_0, z_0) lying on the level set, subject to the constraint $g = 0$ if g is defined on a neighbourhood of (x_0, y_0, z_0), $\nabla g(x_0, y_0, z_0) \neq 0$, and $f(x, y, z) \leq f(x_0, y_0, z_0)$ for all (x, y, z) in a neighbourhood of (x_0, y_0, z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0, y_0, z_0) \neq 0$ is the implicit function theorem: If $\frac{\partial g}{\partial z}(x_0, y_0, z_0)$, then near this point, $z = f(x, y)$ for some C^1 function f, i.e., the level set is a regular parametrised surface near this point.
- As a consequence, if $g(x, y, z) = 0$ is a regular closed surface that is a bounded set, then a C^1 function f attains a global max/min

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x, y, z) = 0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0, y_0, z_0) lying on the level set, subject to the constraint $g = 0$ if g is defined on a neighbourhood of (x_0, y_0, z_0), $\nabla g(x_0, y_0, z_0) \neq 0$, and $f(x, y, z) \leq f(x_0, y_0, z_0)$ for all (x, y, z) in a neighbourhood of (x_0, y_0, z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0, y_0, z_0) \neq 0$ is the implicit function theorem: If $\frac{\partial g}{\partial z}(x_0, y_0, z_0)$, then near this point, $z = f(x, y)$ for some C^1 function f, i.e., the level set is a regular parametrised surface near this point.
- As a consequence, if $g(x, y, z) = 0$ is a regular closed surface that is a bounded set, then a C^1 function f attains a global max/min and does so at

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x, y, z) = 0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0, y_0, z_0) lying on the level set, subject to the constraint $g = 0$ if g is defined on a neighbourhood of (x_0, y_0, z_0), $\nabla g(x_0, y_0, z_0) \neq 0$, and $f(x, y, z) \leq f(x_0, y_0, z_0)$ for all (x, y, z) in a neighbourhood of (x_0, y_0, z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0, y_0, z_0) \neq 0$ is the implicit function theorem: If $\frac{\partial g}{\partial z}(x_0, y_0, z_0)$, then near this point, $z = f(x, y)$ for some C^1 function f, i.e., the level set is a regular parametrised surface near this point.
- As a consequence, if $g(x, y, z) = 0$ is a regular closed surface that is a bounded set, then a C^1 function f attains a global max/min and does so at local extrema.

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x,y,z)=0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0,y_0,z_0) lying on the level set, subject to the constraint $g=0$ if g is defined on a neighbourhood of (x_0,y_0,z_0), $\nabla g(x_0,y_0,z_0)\neq 0$, and $f(x,y,z)\leq f(x_0,y_0,z_0)$ for all (x,y,z) in a neighbourhood of (x_0,y_0,z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0,y_0,z_0)\neq 0$ is the implicit function theorem: If $\frac{\partial g}{\partial z}(x_0,y_0,z_0)$, then near this point, $z=f(x,y)$ for some C^1 function f, i.e., the level set is a regular parametrised surface near this point.
- As a consequence, if $g(x,y,z)=0$ is a regular closed surface that is a bounded set, then a C^1 function f attains a global max/min and does so at local extrema. On the other hand,

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x,y,z)=0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0,y_0,z_0) lying on the level set, subject to the constraint $g=0$ if g is defined on a neighbourhood of (x_0,y_0,z_0), $\nabla g(x_0,y_0,z_0)\neq 0$, and $f(x,y,z)\leq f(x_0,y_0,z_0)$ for all (x,y,z) in a neighbourhood of (x_0,y_0,z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0,y_0,z_0)\neq 0$ is the implicit function theorem: If $\frac{\partial g}{\partial z}(x_0,y_0,z_0)$, then near this point, $z=f(x,y)$ for some C^1 function f, i.e., the level set is a regular parametrised surface near this point.
- As a consequence, if $g(x,y,z)=0$ is a regular closed surface that is a bounded set, then a C^1 function f attains a global max/min and does so at local extrema. On the other hand, if the level set has a boundary,

The set up of the problem

- Def: Let f be a function from a neighbourhood of the level-set $g(x, y, z) = 0$ of a C^1 function to $\mathbb{R}$. It is said to achieve a local maximum at a point (x_0, y_0, z_0) lying on the level set, subject to the constraint $g = 0$ if g is defined on a neighbourhood of (x_0, y_0, z_0), $\nabla g(x_0, y_0, z_0) \neq 0$, and $f(x, y, z) \leq f(x_0, y_0, z_0)$ for all (x, y, z) in a neighbourhood of (x_0, y_0, z_0) lying on the level set. Likewise for a local minimum.
- The point of $\nabla g(x_0, y_0, z_0) \neq 0$ is the implicit function theorem: If $\frac{\partial g}{\partial z}(x_0, y_0, z_0)$, then near this point, $z = f(x, y)$ for some C^1 function f, i.e., the level set is a regular parametrised surface near this point.
- As a consequence, if $g(x, y, z) = 0$ is a regular closed surface that is a bounded set, then a C^1 function f attains a global max/min and does so at local extrema. On the other hand, if the level set has a boundary, then the global extrema can occur on the boundary too.

Lagrange's multipliers

- Theorem:

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$.

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$.

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$ be another C^1 function.

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$,

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.
- Geometrically,

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.
- Geometrically, at the place where a local extremum is achieved,

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.
- Geometrically, at the place where a local extremum is achieved, the direction of fastest change of f is

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.
- Geometrically, at the place where a local extremum is achieved, the direction of fastest change of f is along a normal to $g = 0$.

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.
- Geometrically, at the place where a local extremum is achieved, the direction of fastest change of f is along a normal to $g = 0$. Indeed, roughly speaking,

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.
- Geometrically, at the place where a local extremum is achieved, the direction of fastest change of f is along a normal to $g = 0$. Indeed, roughly speaking, if f could increase/decrease along $g = 0$,

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.
- Geometrically, at the place where a local extremum is achieved, the direction of fastest change of f is along a normal to $g = 0$. Indeed, roughly speaking, if f could increase/decrease along $g = 0$, then we are not at a local extremum!

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.
- Geometrically, at the place where a local extremum is achieved, the direction of fastest change of f is along a normal to $g = 0$. Indeed, roughly speaking, if f could increase/decrease along $g = 0$, then we are not at a local extremum!
- One can generalise

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.
- Geometrically, at the place where a local extremum is achieved, the direction of fastest change of f is along a normal to $g = 0$. Indeed, roughly speaking, if f could increase/decrease along $g = 0$, then we are not at a local extremum!
- One can generalise the theorem to

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.
- Geometrically, at the place where a local extremum is achieved, the direction of fastest change of f is along a normal to $g = 0$. Indeed, roughly speaking, if f could increase/decrease along $g = 0$, then we are not at a local extremum!
- One can generalise the theorem to more than one constraint

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \rightarrow \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.
- Geometrically, at the place where a local extremum is achieved, the direction of fastest change of f is along a normal to $g = 0$. Indeed, roughly speaking, if f could increase/decrease along $g = 0$, then we are not at a local extremum!
- One can generalise the theorem to more than one constraint by demanding $\nabla f = \lambda_1 \nabla g_1 + \lambda_2 \nabla g_2 + \ldots$ provided

Lagrange's multipliers

- Theorem: Suppose $g(x_1, x_2, \ldots) = 0$ is a level set of a C^1 function $g : U \subset \mathbb{R}^n \to \mathbb{R}$. Let $\vec{r}_0$ be a point on $g^{-1}(0)$ such that $\nabla g(\vec{r}_0) \neq \vec{0}$. Let $f : U \subset \mathbb{R}^n \to \mathbb{R}$ be another C^1 function. If f attains a local extremum subject to the constraint $g = 0$ at the point $\vec{r}_0$, then $\nabla f(\vec{r}_0) = \lambda \nabla g(\vec{r}_0)$ for some constant $\lambda \in \mathbb{R}$.
- Geometrically, at the place where a local extremum is achieved, the direction of fastest change of f is along a normal to $g = 0$. Indeed, roughly speaking, if f could increase/decrease along $g = 0$, then we are not at a local extremum!
- One can generalise the theorem to more than one constraint by demanding $\nabla f = \lambda_1 \nabla g_1 + \lambda_2 \nabla g_2 + \ldots$ provided $\nabla g_1(\vec{r}_0), \nabla g_2(\vec{r}_0), \ldots$ are linearly independent.

Proof

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$.

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem,

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.
- Now $\tilde{f} = f(x_1, \ldots, x_{n-1}, h(x_1, \ldots, x_{n-1}))$ attains an unconstrained local extremum

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.
- Now $\tilde{f} = f(x_1, \ldots, x_{n-1}, h(x_1, \ldots, x_{n-1}))$ attains an unconstrained local extremum at $\vec{r}_0$.

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.
- Now $\tilde{f} = f(x_1, \ldots, x_{n-1}, h(x_1, \ldots, x_{n-1}))$ attains an unconstrained local extremum at $\vec{r}_0$. Thus $\frac{\partial \tilde{f}}{\partial x_1}(\vec{r}_0) = \ldots = \frac{\partial \tilde{f}}{\partial x_{n-1}}(\vec{r}_0) = 0$.

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.
- Now $\tilde{f} = f(x_1, \ldots, x_{n-1}, h(x_1, \ldots, x_{n-1}))$ attains an unconstrained local extremum at $\vec{r}_0$. Thus $\frac{\partial \tilde{f}}{\partial x_1}(\vec{r}_0) = \ldots = \frac{\partial \tilde{f}}{\partial x_{n-1}}(\vec{r}_0) = 0$. By the chain rule,

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.
- Now $\tilde{f} = f(x_1, \ldots, x_{n-1}, h(x_1, \ldots, x_{n-1}))$ attains an unconstrained local extremum at $\vec{r}_0$. Thus $\frac{\partial \tilde{f}}{\partial x_1}(\vec{r}_0) = \ldots = \frac{\partial \tilde{f}}{\partial x_{n-1}}(\vec{r}_0) = 0$. By the chain rule, $\frac{\partial f}{\partial x_1} + \frac{\partial f}{\partial x_n}\frac{\partial h}{\partial x_1} = 0 = \ldots$.

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.
- Now $\tilde{f} = f(x_1, \ldots, x_{n-1}, h(x_1, \ldots, x_{n-1}))$ attains an unconstrained local extremum at $\vec{r}_0$. Thus $\frac{\partial \tilde{f}}{\partial x_1}(\vec{r}_0) = \ldots = \frac{\partial \tilde{f}}{\partial x_{n-1}}(\vec{r}_0) = 0$. By the chain rule, $\frac{\partial f}{\partial x_1} + \frac{\partial f}{\partial x_n}\frac{\partial h}{\partial x_1} = 0 = \ldots$.
- However, since $g(x_1, \ldots, x_{n-1}, h(x_1, \ldots)) = 0$,

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.
- Now $\tilde{f} = f(x_1, \ldots, x_{n-1}, h(x_1, \ldots, x_{n-1}))$ attains an unconstrained local extremum at $\vec{r}_0$. Thus $\frac{\partial \tilde{f}}{\partial x_1}(\vec{r}_0) = \ldots = \frac{\partial \tilde{f}}{\partial x_{n-1}}(\vec{r}_0) = 0$. By the chain rule, $\frac{\partial f}{\partial x_1} + \frac{\partial f}{\partial x_n}\frac{\partial h}{\partial x_1} = 0 = \ldots$.
- However, since $g(x_1, \ldots, x_{n-1}, h(x_1, \ldots)) = 0$, differentiating both sides using the chain rule,

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.
- Now $\tilde{f} = f(x_1, \ldots, x_{n-1}, h(x_1, \ldots, x_{n-1}))$ attains an unconstrained local extremum at $\vec{r}_0$. Thus $\frac{\partial \tilde{f}}{\partial x_1}(\vec{r}_0) = \ldots = \frac{\partial \tilde{f}}{\partial x_{n-1}}(\vec{r}_0) = 0$. By the chain rule, $\frac{\partial f}{\partial x_1} + \frac{\partial f}{\partial x_n}\frac{\partial h}{\partial x_1} = 0 = \ldots$.
- However, since $g(x_1, \ldots, x_{n-1}, h(x_1, \ldots)) = 0$, differentiating both sides using the chain rule, $g_{x_1} + g_{x_n} h_{x_1} = 0 = \ldots$.

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.
- Now $\tilde{f} = f(x_1, \ldots, x_{n-1}, h(x_1, \ldots, x_{n-1}))$ attains an unconstrained local extremum at $\vec{r}_0$. Thus $\frac{\partial \tilde{f}}{\partial x_1}(\vec{r}_0) = \ldots = \frac{\partial \tilde{f}}{\partial x_{n-1}}(\vec{r}_0) = 0$. By the chain rule, $\frac{\partial f}{\partial x_1} + \frac{\partial f}{\partial x_n}\frac{\partial h}{\partial x_1} = 0 = \ldots$.
- However, since $g(x_1, \ldots, x_{n-1}, h(x_1, \ldots)) = 0$, differentiating both sides using the chain rule, $g_{x_1} + g_{x_n} h_{x_1} = 0 = \ldots$.
- Thus, $\frac{g_{x_i}}{g_{x_n}} = \frac{\tilde{f}_{x_i}}{\tilde{f}_{x_n}}$ for all $1 \leq i \leq n-1$.

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.
- Now $\tilde{f} = f(x_1, \ldots, x_{n-1}, h(x_1, \ldots, x_{n-1}))$ attains an unconstrained local extremum at $\vec{r}_0$. Thus $\frac{\partial \tilde{f}}{\partial x_1}(\vec{r}_0) = \ldots = \frac{\partial \tilde{f}}{\partial x_{n-1}}(\vec{r}_0) = 0$. By the chain rule, $\frac{\partial f}{\partial x_1} + \frac{\partial f}{\partial x_n}\frac{\partial h}{\partial x_1} = 0 = \ldots$.
- However, since $g(x_1, \ldots, x_{n-1}, h(x_1, \ldots)) = 0$, differentiating both sides using the chain rule, $g_{x_1} + g_{x_n} h_{x_1} = 0 = \ldots$.
- Thus, $\frac{g_{x_i}}{g_{x_n}} = \frac{\tilde{f}_{x_i}}{\tilde{f}_{x_n}}$ for all $1 \leq i \leq n-1$. Therefore, $\nabla f = \lambda \nabla g$ at $\vec{r}_0$.

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.
- Now $\tilde{f} = f(x_1, \ldots, x_{n-1}, h(x_1, \ldots, x_{n-1}))$ attains an unconstrained local extremum at $\vec{r}_0$. Thus $\frac{\partial \tilde{f}}{\partial x_1}(\vec{r}_0) = \ldots = \frac{\partial \tilde{f}}{\partial x_{n-1}}(\vec{r}_0) = 0$. By the chain rule, $\frac{\partial f}{\partial x_1} + \frac{\partial f}{\partial x_n}\frac{\partial h}{\partial x_1} = 0 = \ldots$.
- However, since $g(x_1, \ldots, x_{n-1}, h(x_1, \ldots)) = 0$, differentiating both sides using the chain rule, $g_{x_1} + g_{x_n} h_{x_1} = 0 = \ldots$.
- Thus, $\frac{g_{x_i}}{g_{x_n}} = \frac{\tilde{f}_{x_i}}{\tilde{f}_{x_n}}$ for all $1 \leq i \leq n-1$. Therefore, $\nabla f = \lambda \nabla g$ at $\vec{r}_0$.
- By the way,

Proof

- Assume (WLog) that $\frac{\partial g}{\partial x_n}(\vec{r}_0) \neq 0$. By the implicit function theorem, $x_n = h(x_1, \ldots, x_{n-1})$ for a C^1 function h near $\vec{r}_0$.
- Now $\tilde{f} = f(x_1, \ldots, x_{n-1}, h(x_1, \ldots, x_{n-1}))$ attains an unconstrained local extremum at $\vec{r}_0$. Thus $\frac{\partial \tilde{f}}{\partial x_1}(\vec{r}_0) = \ldots = \frac{\partial \tilde{f}}{\partial x_{n-1}}(\vec{r}_0) = 0$. By the chain rule, $\frac{\partial f}{\partial x_1} + \frac{\partial f}{\partial x_n}\frac{\partial h}{\partial x_1} = 0 = \ldots$.
- However, since $g(x_1, \ldots, x_{n-1}, h(x_1, \ldots)) = 0$, differentiating both sides using the chain rule, $g_{x_1} + g_{x_n} h_{x_1} = 0 = \ldots$.
- Thus, $\frac{g_{x_i}}{g_{x_n}} = \frac{\tilde{f}_{x_i}}{\tilde{f}_{x_n}}$ for all $1 \leq i \leq n-1$. Therefore, $\nabla f = \lambda \nabla g$ at $\vec{r}_0$.
- By the way, the second derivative test for constrained local extrema is too complicated for us.

Example 1

Example 1

- Let's go back to

Example 1

- Let's go back to maximising $f = x + y + z$ subject to $g = x^2 + y^2 + z^2 - 1 = 0$.

Example 1

- Let's go back to maximising $f = x + y + z$ subject to $g = x^2 + y^2 + z^2 - 1 = 0$. Note that

Example 1

- Let's go back to maximising $f = x + y + z$ subject to $g = x^2 + y^2 + z^2 - 1 = 0$. Note that the level set is a regular closed surface that is a bounded set.

Example 1

- Let's go back to maximising $f = x + y + z$ subject to $g = x^2 + y^2 + z^2 - 1 = 0$. Note that the level set is a regular closed surface that is a bounded set. f is continuous

Example 1

- Let's go back to maximising $f = x + y + z$ subject to $g = x^2 + y^2 + z^2 - 1 = 0$. Note that the level set is a regular closed surface that is a bounded set. f is continuous and hence the global max is a local max

Example 1

- Let's go back to maximising $f = x + y + z$ subject to $g = x^2 + y^2 + z^2 - 1 = 0$. Note that the level set is a regular closed surface that is a bounded set. f is continuous and hence the global max is a local max and we can use Lagrange's multipliers:

Example 1

- Let's go back to maximising $f = x + y + z$ subject to $g = x^2 + y^2 + z^2 - 1 = 0$. Note that the level set is a regular closed surface that is a bounded set. f is continuous and hence the global max is a local max and we can use Lagrange's multipliers: $\nabla f = (1, 1, 1)$ and $\nabla g = (2x, 2y, 2z)$.

Example 1

- Let's go back to maximising $f = x + y + z$ subject to $g = x^2 + y^2 + z^2 - 1 = 0$. Note that the level set is a regular closed surface that is a bounded set. f is continuous and hence the global max is a local max and we can use Lagrange's multipliers: $\nabla f = (1, 1, 1)$ and $\nabla g = (2x, 2y, 2z)$. Thus, $\nabla f = \lambda \nabla g$ implies that $x = y = z$.

Example 1

- Let's go back to maximising $f = x + y + z$ subject to $g = x^2 + y^2 + z^2 - 1 = 0$. Note that the level set is a regular closed surface that is a bounded set. f is continuous and hence the global max is a local max and we can use Lagrange's multipliers: $\nabla f = (1, 1, 1)$ and $\nabla g = (2x, 2y, 2z)$. Thus, $\nabla f = \lambda \nabla g$ implies that $x = y = z$. Since $x^2 + y^2 + z^2 = 1$, $x = y = z = \pm\frac{1}{\sqrt{3}}$.

Example 1

- Let's go back to maximising $f = x + y + z$ subject to $g = x^2 + y^2 + z^2 - 1 = 0$. Note that the level set is a regular closed surface that is a bounded set. f is continuous and hence the global max is a local max and we can use Lagrange's multipliers: $\nabla f = (1, 1, 1)$ and $\nabla g = (2x, 2y, 2z)$. Thus, $\nabla f = \lambda \nabla g$ implies that $x = y = z$. Since $x^2 + y^2 + z^2 = 1$, $x = y = z = \pm\frac{1}{\sqrt{3}}$. One of these corresponds to the maximum and the other to the minimum.

Example 2

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$.

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*.

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary corresponding to one of the x_i being 0.

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary corresponding to one of the x_i being 0. At these boundary points,

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary corresponding to one of the x_i being 0. At these boundary points, $f = 0$.

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary corresponding to one of the x_i being 0. At these boundary points, $f = 0$.
- Let us look at the interior local extrema:

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary corresponding to one of the x_i being 0. At these boundary points, $f = 0$.
- Let us look at the interior local extrema: $\nabla g = (1, 1, \ldots, 1) \neq 0$. Hence

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary corresponding to one of the x_i being 0. At these boundary points, $f = 0$.
- Let us look at the interior local extrema: $\nabla g = (1, 1, \ldots, 1) \neq 0$. Hence we can use Lagrange's multipliers.

Example 2

- Maximise $f(x_1, \dots, x_n) = x_1 x_2 \dots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \dots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary corresponding to one of the x_i being 0. At these boundary points, $f = 0$.
- Let us look at the interior local extrema: $\nabla g = (1, 1, \dots, 1) \neq 0$. Hence we can use Lagrange's multipliers.
- $\nabla f = (x_2 x_3 \dots x_n, x_1 x_3 \dots, \dots) = \lambda(1, 1, \dots)$.

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary corresponding to one of the x_i being 0. At these boundary points, $f = 0$.
- Let us look at the interior local extrema: $\nabla g = (1, 1, \ldots, 1) \neq 0$. Hence we can use Lagrange's multipliers.
- $\nabla f = (x_2 x_3 \ldots x_n, x_1 x_3 \ldots, \ldots) = \lambda(1, 1, \ldots)$. Thus $x_1 = x_2 = \ldots = \frac{a}{n}$.

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary corresponding to one of the x_i being 0. At these boundary points, $f = 0$.
- Let us look at the interior local extrema: $\nabla g = (1, 1, \ldots, 1) \neq 0$. Hence we can use Lagrange's multipliers.
- $\nabla f = (x_2 x_3 \ldots x_n, x_1 x_3 \ldots, \ldots) = \lambda(1, 1, \ldots)$. Thus $x_1 = x_2 = \ldots = \frac{a}{n}$. Thus the maximum value of f is

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary corresponding to one of the x_i being 0. At these boundary points, $f = 0$.
- Let us look at the interior local extrema: $\nabla g = (1, 1, \ldots, 1) \neq 0$. Hence we can use Lagrange's multipliers.
- $\nabla f = (x_2 x_3 \ldots x_n, x_1 x_3 \ldots, \ldots) = \lambda(1, 1, \ldots)$. Thus $x_1 = x_2 = \ldots = \frac{a}{n}$. Thus the maximum value of f is $\frac{a^n}{n^n}$.

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary corresponding to one of the x_i being 0. At these boundary points, $f = 0$.
- Let us look at the interior local extrema: $\nabla g = (1, 1, \ldots, 1) \neq 0$. Hence we can use Lagrange's multipliers.
- $\nabla f = (x_2 x_3 \ldots x_n, x_1 x_3 \ldots, \ldots) = \lambda(1, 1, \ldots)$. Thus $x_1 = x_2 = \ldots = \frac{a}{n}$. Thus the maximum value of f is $\frac{a^n}{n^n}$. This means that $f^{1/n} \leq \frac{g}{n}$ for all $x_i \geq 0$.

Example 2

- Maximise $f(x_1, \ldots, x_n) = x_1 x_2 \ldots x_n$ on $x_i \geq 0$ subject to $x_1 + x_2 + \ldots + x_n = a > 0$. The constraining surface is a closed *set* and a bounded *set*. However it is NOT a closed *surface* and has a boundary corresponding to one of the x_i being 0. At these boundary points, $f = 0$.
- Let us look at the interior local extrema: $\nabla g = (1, 1, \ldots, 1) \neq 0$. Hence we can use Lagrange's multipliers.
- $\nabla f = (x_2 x_3 \ldots x_n, x_1 x_3 \ldots, \ldots) = \lambda(1, 1, \ldots)$. Thus $x_1 = x_2 = \ldots = \frac{a}{n}$. Thus the maximum value of f is $\frac{a^n}{n^n}$. This means that $f^{1/n} \leq \frac{g}{n}$ for all $x_i \geq 0$. This inequality is the famous AM-GM inequality.