# Differentiable Optimisation in Deep Learning 

Stephen Gould<br>stephen.gould@anu.edu.au<br>Australian National University

15 December, 2022

## Discovery of Ceres



## Discovery of Ceres



## Optimisation is Everywhere

- financial mathematics: maximise profits or minimise costs subject to constraints on resources and budgets


## Optimisation is Everywhere

- financial mathematics: maximise profits or minimise costs subject to constraints on resources and budgets
- mechanical engineering: maximise the span of a bridge subject to load constraints


## Optimisation is Everywhere

- financial mathematics: maximise profits or minimise costs subject to constraints on resources and budgets
- mechanical engineering: maximise the span of a bridge subject to load constraints
- electrical engineering: minimise the size of a transistor in a circuit subject to power and timing constraints


## Optimisation is Everywhere

- financial mathematics: maximise profits or minimise costs subject to constraints on resources and budgets
- mechanical engineering: maximise the span of a bridge subject to load constraints
- electrical engineering: minimise the size of a transistor in a circuit subject to power and timing constraints
- logistics and planning: find the cheapest way to distribute goods from suppliers to consumers across a network


## Optimisation is Everywhere

- financial mathematics: maximise profits or minimise costs subject to constraints on resources and budgets
- mechanical engineering: maximise the span of a bridge subject to load constraints
- electrical engineering: minimise the size of a transistor in a circuit subject to power and timing constraints
- logistics and planning: find the cheapest way to distribute goods from suppliers to consumers across a network
- statistics/data science: curve fitting and data visualisation


## Optimisation is Everywhere

- financial mathematics: maximise profits or minimise costs subject to constraints on resources and budgets
- mechanical engineering: maximise the span of a bridge subject to load constraints
- electrical engineering: minimise the size of a transistor in a circuit subject to power and timing constraints
- logistics and planning: find the cheapest way to distribute goods from suppliers to consumers across a network
- statistics/data science: curve fitting and data visualisation
- machine learning and deep learning: minimise loss functions with respect to the parameters of our model


## Overview

- Introduction to Optimisation
- Formal definition
- Least squares
- Convex sets and functions
- Convex optimisation problems
- Lagrangian
- Optimality conditions
- Algorithms
- Differentiable Optimisation and Deep Learning
- Machine learning from 10,000ft
- Automatic differentiation
- Forward and backward passes
- Imperative and declarative nodes
- Bi-level optimisation
- Implicit function theorem
- Differentiable optimisation results
- Examples and Applications
- Least squares
- Optimal transport
- Blind perspective-n-point

> accompanying lecture notes available at https://users.cecs.anu.edu.au/~sgould

## lecture 1

## Lecture 1: Introduction to Optimisation



## Assumed Background

LINEAR ALGEBRA AND ITS APPLICATIONS GILBERT STRANG


## Optimisation Problems

find the assignment to variables that minimises a measure of cost subject to some constraints ${ }^{1}$

[^0]
## Optimisation Problems

$$
\begin{array}{ll}
\text { minimize (over } x) & \text { objective }(x) \\
\text { subject to } & \text { constraints }(x)
\end{array}
$$

## Optimisation Problems

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x) \\
\text { subject to } & f_{i}(x) \leq 0, \quad i=1, \ldots, p \\
& h_{i}(x)=0, \quad i=1, \ldots, q
\end{array}
$$

- $x=\left(x_{1}, \ldots, x_{n}\right) \in \mathbb{R}^{n}$ - optimisation variables
- $f_{0}: \mathbb{R}^{n} \rightarrow \mathbb{R}$ - objective (or cost or loss) function
- $f_{i}: \mathbb{R}^{n} \rightarrow \mathbb{R}, i=1, \ldots, p$ - inequality constraint functions
- $h_{i}: \mathbb{R}^{n} \rightarrow \mathbb{R}, i=1, \ldots, q$ - equality constraint functions


## Solution and Optimal Value

A point $x$ is feasible if $x \in \operatorname{dom}\left(f_{0}\right)$ and it satisfies the constraints.
A solution, or optimal point, $x^{\star}$ has the smallest value of $f_{0}$ among all feasible $x$.

[^1]
## Solution and Optimal Value

A point $x$ is feasible if $x \in \operatorname{dom}\left(f_{0}\right)$ and it satisfies the constraints.
A solution, or optimal point, $x^{\star}$ has the smallest value of $f_{0}$ among all feasible $x$.
The optimal value is ${ }^{1}$

$$
p^{\star}=\inf _{x \in \mathcal{D}}\left\{\begin{array}{l|l}
f_{0}(x) & \begin{array}{l}
f_{i}(x) \leq 0, \quad i=1, \ldots, p \\
h_{i}(x)=0, \quad i=1, \ldots, q
\end{array}
\end{array}\right\}
$$

- $p^{\star}$ and is equal to $f_{0}\left(x^{\star}\right)$ when $x^{\star}$ exists
- $p^{\star}=\infty$ if the problem is infeasible (no $x$ satisfies the constraints)
- $p^{\star}=-\infty$ if the problem is unbounded below

[^2]
## Locally Optimal Points

A point $x$ is locally optimal if there is an $R>0$ such that $z=x$ is optimal for

$$
\begin{array}{lll}
\operatorname{minimize}(\text { over } z) & f_{0}(z) & \\
\text { subject to } & f_{i}(z) \leq 0 & i=1, \ldots, p \\
& h_{i}(z)=0 & i=1, \ldots, q \\
& \|z-x\|_{2} \leq R . &
\end{array}
$$

## Examples (1D)

$$
\begin{aligned}
& 1 / x \\
& -\log x \\
& x \log x \\
& x^{3}-3 x \\
& \operatorname{dom}\left(f_{0}\right) \text { : } \\
& p^{\star} \text { : } \\
& x^{\star} \text { : } \\
& \mathbb{R}_{++} \\
& -1 / e \\
& 1 / e \\
& \begin{array}{c}
\mathbb{R} \\
-\infty
\end{array} \\
& \text { none } \\
& \begin{array}{l}
\mathbb{R}_{++} \\
-\infty
\end{array} \\
& \begin{array}{c}
\mathbb{R}_{++} \\
-1 / e \\
1 / e
\end{array} \\
& \mathbb{R} \\
& x=1 \text { locally }
\end{aligned}
$$

## Examples (2D)





## Least Squares

$$
\operatorname{minimize}\|A x-b\|_{2}^{2}
$$

## Least Squares

$$
\operatorname{minimize}\|A x-b\|_{2}^{2}
$$

- unique solution if $A^{T} A$ is invertible, $x^{\star}=\left(A^{T} A\right)^{-1} A^{T} b$
- solution via SVD, $A=U \Sigma V^{T}$, if $A^{T} A$ not invertible, $x^{\star}=V \Sigma^{-1} U^{T} b$
- in fact, $x^{\star}+w$ for any $w \in \mathcal{N}(A)$ also a solution
- solution via QR factorisation, $x^{\star}=R^{-1} Q^{T} b$
- solved in $O\left(n^{2} m\right)$ time, less if structured
- typically use iterative solver


## Example: Polynomial Curve Fitting

fit $n$-th order polynomial $f_{a}(x)=\sum_{k=0}^{n} a_{k} x^{k}$ to set of noisy points $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{m}$ minimize (over $a) \quad \sum_{i=1}^{m}\left(f_{a}\left(x_{i}\right)-y_{i}\right)^{2}$
$\operatorname{minimize}\left\|\left[\begin{array}{ccccc}1 & x_{1} & x_{1}^{2} & \ldots & x_{1}^{n} \\ 1 & x_{2} & x_{2}^{2} & \ldots & x_{2}^{n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_{m} & x_{m}^{2} & \ldots & x_{m}^{n}\end{array}\right]\left[\begin{array}{c}a_{0} \\ a_{1} \\ \vdots \\ a_{n}\end{array}\right]-\left[\begin{array}{c}y_{1} \\ y_{2} \\ \vdots \\ y_{m}\end{array}\right]\right\|_{2}^{2}$

- special case of convex optimisation



## Lines and Line Segments

- a line through two points $x_{1}$ and $x_{2}$

$$
x=\theta x_{1}+(1-\theta) x_{2}, \quad(\theta \in \mathbb{R})
$$



- an affine set contains the line through any two distinct points in the set
- an affine hull the set formed by taking all lines through points in a set


## Lines and Line Segments

- a line through two points $x_{1}$ and $x_{2}$

$$
x=\theta x_{1}+(1-\theta) x_{2}, \quad(\theta \in \mathbb{R})
$$



- an affine set contains the line through any two distinct points in the set
- an affine hull the set formed by taking all lines through points in a set
- a line segment between $x_{1}$ and $x_{2}$

$$
x=\theta x_{1}+(1-\theta) x_{2}, \quad(0 \leq \theta \leq 1)
$$

- a convex set contains the line segment between any two distinct points in the set
- an convex hull the set formed by taking all line segments between points in a set


## Convex Sets

$$
x_{1}, x_{2} \in \text { convex set } C \quad \Longrightarrow \quad \theta x_{1}+(1-\theta) x_{2} \in C \text { for all } 0 \leq \theta \leq 1
$$


common examples in machine learning:

- nonnegative orthant, $\mathbb{R}_{+}^{n}=\left\{x \mid x_{i} \geq 0, i=1, \ldots, n\right\}$
- positive semindefinite matrices, $\mathbb{S}_{+}^{n}=\left\{X \mid z^{T} X z \geq 0, z \in \mathbb{R}^{n}\right\}$


## More Examples


hyperplane,
$\left\{x \mid a^{T} x=b\right\}$

norm ball,

$$
\left\{x \mid\left\|x-x_{c}\right\|_{p} \leq r\right\}
$$


halfspace, $\left\{x \mid a^{T} x \leq b\right\}$

ellipsoid,
$\left\{A x+b \mid\|x\|_{2} \leq 1\right\}$

$\{x \mid A x \preceq b, C x=d\}$


Lorentz cone, $\{(x, t) \mid\|x\| \leq t\}$

## Convex Functions

A function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is convex if $\operatorname{dom}(f)$ is a convex set and

$$
f(\theta x+(1-\theta) y) \leq \theta f(x)+(1-\theta) f(y)
$$

for all $x, y \in \operatorname{dom}(f), 0 \leq \theta \leq 1$.


- $f$ is concave if $-f$ is convex


## Examples








Weighted Sum and Pointwise Maximum Preserve Convexity



## Convex, Strictly Convex, and Strongly Convex






- $f_{1}$ is smooth and convex: $f(\theta x+(1-\theta) y) \leq \theta f(x)+(1-\theta) y$
- $f_{2}$ is non-differentiable and convex: $f(\theta x+(1-\theta) y) \leq \theta f(x)+(1-\theta) y$
- $f_{3}$ is strictly convex: $f(\theta x+(1-\theta) y)<\theta f(x)+(1-\theta) y$
- $f_{4}$ is strongly convex: $\exists m$ s.t. $m(y-x)^{2} \leq f(y)-f(x)$


## Epigraph

The epigraph of function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is the set

$$
\mathbf{e p i}(f)=\left\{(x, t) \in \mathbb{R}^{n+1} \mid x \in \operatorname{dom}(f), f(x) \leq t\right\}
$$



- $f$ is a convex function if and only if $\mathbf{e p i}(f)$ is a convex set


## First-order Condition

differentiable $f$ with convex domain is convex iff

$$
f(y) \geq f(x)+\nabla f(x)^{T}(y-x) \quad \text { for all } x, y \in \operatorname{dom}(f)
$$



- first-order approximation of (convex) $f$ is a global under estimator


## Second-order Condition





twice differentiable $f$ with convex domain is convex iff

$$
\nabla^{2} f(x) \succeq 0 \quad \text { for all } x \in \operatorname{dom}(f)
$$

- if $\nabla^{2} f(x) \succ 0$ for all $x \in \operatorname{dom}(f)$, then $f$ is strictly convex
- if $\nabla^{2} f(x) \succeq m I$ for some $m>0$ and all $x \in \operatorname{dom}(f)$, then $f$ is strongly convex
- strongly convex functions have a unique minimum

Worked Example: log-sum-exp is Convex

$$
f(x)=\log \sum_{k=1}^{n} \exp x_{k}
$$

## Worked Example: log-sum-exp is Convex

$$
f(x)=\log \sum_{k=1}^{n} \exp x_{k}
$$

Proof. Start by computing the gradient and Hessian,

$$
\begin{array}{rlr}
\frac{\partial f(x)}{\partial x_{i}} & =\frac{\exp x_{i}}{\sum_{k=1}^{n} \exp x_{k}} & \left.\quad \text { (derivative of } \log (z), z^{\prime} / z\right) \\
\frac{\partial^{2} f(x)}{\partial x_{i} \partial x_{j}} & =\frac{\left(\sum_{k=1}^{n} \exp x_{k}\right) \llbracket i=j \rrbracket \exp x_{i}-\exp x_{i} \exp x_{j}}{\left(\sum_{k=1}^{n} \exp x_{k}\right)^{2}} \quad \text { (quotient rule, } \frac{v \cdot \mathrm{~d} u-u \cdot \mathrm{~d} v}{v^{2}} \text { ) }
\end{array}
$$

## Worked Example: log-sum-exp is Convex

$$
f(x)=\log \sum_{k=1}^{n} \exp x_{k}
$$

Proof. Start by computing the gradient and Hessian,

$$
\begin{array}{rlr}
\frac{\partial f(x)}{\partial x_{i}} & =\frac{z_{i}}{\mathbf{1}^{T} z} \\
\frac{\partial^{2} f(x)}{\partial x_{i} \partial x_{j}} & =\frac{\left(\mathbf{1}^{T} z\right) \llbracket i=j \rrbracket z_{i}-z_{i} z_{j}}{\left(\mathbf{1}^{T} z\right)^{2}} & \left(z_{k}=\exp x_{k}\right)
\end{array}
$$

## Worked Example: log-sum-exp is Convex

$$
f(x)=\log \sum_{k=1}^{n} \exp x_{k}
$$

Proof. Start by computing the gradient and Hessian,

$$
\begin{aligned}
\nabla f(x) & =\frac{1}{\mathbf{1}^{T} z} z \\
\nabla^{2} f(x) & =\frac{1}{\left(\mathbf{1}^{T} z\right)^{2}}\left(\left(\mathbf{1}^{T} z\right) \operatorname{diag}(z)-z z^{T}\right)
\end{aligned}
$$

## Worked Example: log-sum-exp is Convex

$$
f(x)=\log \sum_{k=1}^{n} \exp x_{k}
$$

Proof. Start by computing the gradient and Hessian,

$$
\nabla^{2} f(x)=\frac{1}{\left(\mathbf{1}^{T} z\right)^{2}}\left(\left(\mathbf{1}^{T} z\right) \operatorname{diag}(z)-z z^{T}\right)
$$

$$
\left(z_{k}=\exp x_{k}\right)
$$

To show that $\nabla^{2} f(x) \succeq 0$, we must verify that $v^{T} \nabla^{2} f(x) v \geq 0$ for all $v$.

## Worked Example: log-sum-exp is Convex

$$
f(x)=\log \sum_{k=1}^{n} \exp x_{k}
$$

Proof. Start by computing the gradient and Hessian,

$$
\nabla^{2} f(x)=\frac{1}{\left(\mathbf{1}^{T} z\right)^{2}}\left(\left(\mathbf{1}^{T} z\right) \operatorname{diag}(z)-z z^{T}\right) \quad\left(z_{k}=\exp x_{k}\right)
$$

To show that $\nabla^{2} f(x) \succeq 0$, we must verify that $v^{T} \nabla^{2} f(x) v \geq 0$ for all $v$.

$$
\begin{aligned}
v^{T} \nabla^{2} f(x) v & =\frac{1}{\left(\mathbf{1}^{T} z\right)^{2}} v^{T}\left(\left(\mathbf{1}^{T} z\right) \operatorname{diag}(z)-z z^{T}\right) v \\
& =\frac{1}{\left(\mathbf{1}^{T} z\right)^{2}}\left(\left(\mathbf{1}^{T} z\right) v^{T} \mathbf{\operatorname { d i a g }}(z) v-v^{T} z z^{T} v\right)
\end{aligned}
$$

## Worked Example: log-sum-exp is Convex

$$
f(x)=\log \sum_{k=1}^{n} \exp x_{k}
$$

Proof. Start by computing the gradient and Hessian,

$$
\begin{aligned}
& \nabla^{2} f(x)=\frac{1}{\left(\mathbf{1}^{T} z\right)^{2}}\left(\left(\mathbf{1}^{T} z\right) \operatorname{diag}(z)-z z^{T}\right) \quad\left(z_{k}=\exp x_{k}\right) \\
& \begin{aligned}
v^{T} \nabla^{2} f(x) v & =\frac{1}{\left(\mathbf{1}^{T} z\right)^{2}} v^{T}\left(\left(\mathbf{1}^{T} z\right) \boldsymbol{\operatorname { d i a g }}(z)-z z^{T}\right) v \\
& =\frac{1}{\left(\mathbf{1}^{T} z\right)^{2}}\left(\left(\mathbf{1}^{T} z\right) v^{T} \boldsymbol{\operatorname { d i a g }}(z) v-v^{T} z z^{T} v\right)
\end{aligned}
\end{aligned}
$$

Therefore we need to show that $\left(\mathbf{1}^{T} z\right) v^{T} \boldsymbol{\operatorname { d i a g }}(z) v \geq\left(v^{T} z\right)^{2}$ for all $v$.

## Worked Example: log-sum-exp is Convex

$$
f(x)=\log \sum_{k=1}^{n} \exp x_{k}
$$

Proof. Start by computing the gradient and Hessian,

$$
\nabla^{2} f(x)=\frac{1}{\left(\mathbf{1}^{T} z\right)^{2}}\left(\left(\mathbf{1}^{T} z\right) \operatorname{diag}(z)-z z^{T}\right)
$$

$$
\left(z_{k}=\exp x_{k}\right)
$$

Therefore we need to show that $\left(\mathbf{1}^{T} z\right) v^{T} \boldsymbol{\operatorname { d i a g }}(z) v \geq\left(v^{T} z\right)^{2}$ for all $v$.

## Worked Example: log-sum-exp is Convex

$$
f(x)=\log \sum_{k=1}^{n} \exp x_{k}
$$

Proof. Start by computing the gradient and Hessian,

$$
\nabla^{2} f(x)=\frac{1}{\left(\mathbf{1}^{T} z\right)^{2}}\left(\left(\mathbf{1}^{T} z\right) \operatorname{diag}(z)-z z^{T}\right) \quad\left(z_{k}=\exp x_{k}\right)
$$

Therefore we need to show that $\left(\mathbf{1}^{T} z\right) v^{T} \boldsymbol{\operatorname { d i a g }}(z) v \geq\left(v^{T} z\right)^{2}$ for all $v$. That is, we need to show

$$
\left(\sum_{k=1}^{n} z_{k}\right)\left(\sum_{k=1}^{n} z_{k} v_{k}^{2}\right) \geq\left(\sum_{k=1}^{n} v_{k} z_{k}\right)^{2}
$$

## Worked Example: log-sum-exp is Convex

$$
f(x)=\log \sum_{k=1}^{n} \exp x_{k}
$$

Proof. Start by computing the gradient and Hessian,

$$
\nabla^{2} f(x)=\frac{1}{\left(\mathbf{1}^{T} z\right)^{2}}\left(\left(\mathbf{1}^{T} z\right) \operatorname{diag}(z)-z z^{T}\right) \quad\left(z_{k}=\exp x_{k}\right)
$$

Therefore we need to show that $\left(\mathbf{1}^{T} z\right) v^{T} \boldsymbol{\operatorname { d i a g }}(z) v \geq\left(v^{T} z\right)^{2}$ for all $v$. That is, we need to show

$$
\left(\sum_{k=1}^{n} z_{k}\right)\left(\sum_{k=1}^{n} z_{k} v_{k}^{2}\right) \geq\left(\sum_{k=1}^{n} v_{k} z_{k}\right)^{2}
$$

which is true by the Cauchy-Schwarz inequality, $\|a\|_{2}^{2}\|b\|_{2}^{2} \geq\left(a^{T} b\right)^{2}$, with $a=\left(\sqrt{z_{1}}, \ldots, \sqrt{z_{n}}\right)$ and $b=\left(\sqrt{z_{1}} v_{1}, \ldots, \sqrt{z_{n}} v_{n}\right)$.

## Convex Optimisation

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x) \\
\text { subject to } & f_{i}(x) \leq 0, \quad i=1, \ldots, p \\
& a_{i}^{T} x=b_{i}, \quad i=1, \ldots, q
\end{array}
$$

- $f_{0}, f_{1}, \ldots, f_{p}$ are convex
- $h_{i}(x) \triangleq a_{i}^{T} x-b_{i}$ are affine, often written as $A x=b$
minimise a convex objective over a convex feasible set


## Local Optima are Global Optima

any local minimum of a convex problem is (globally) optimal

## Local Optima are Global Optima

any local minimum of a convex problem is (globally) optimal

## Proof Sketch.

- towards contradiction, suppose $x$ is locally optimal, but there exists a feasible $y$ with lower objective
- since $x$ is locally optimally there exists a radius $R$ such that no other point within $R$ of $x$ has lower objective
- (so $y$ must be further than $R$ from $x$ )
- pick a point $z$ on the line segment between $x$ and $y$ and within $R$ of $x$
- so $z$ must be feasible and have objective no lower than $x$
- but, by the basic inequality of convex functions,

the objective value at $z$ must be between that at $x$ and $y$, i.e., lower than $f_{0}(x)$
- we have a contradiction


## Optimality Criterion for Differentiable $f_{0}$

$x$ is optimal if and only if it is feasible and $\nabla f_{0}(x)^{T}(y-x) \geq 0$ for all feasible $y$

if nonzero,

- $\nabla f_{0}(x)$ defines a supporting hyperplane to feasible set $\mathcal{X}$ at $x$
- $f_{0}$ cannot be improved by moving in a direction where $x$ stays feasible


## Lagrangian

Standard form problem (not necessarily convex),

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x) \\
\text { subject to } & f_{i}(x) \leq 0, \quad i=1, \ldots, p \\
& h_{i}(x)=0, \quad i=1, \ldots, q
\end{array}
$$

variable $x \in \mathbb{R}^{n}$, domain $\mathcal{D}$, optimal value $p^{\star}$

## Lagrangian

Standard form problem (not necessarily convex),

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x) \\
\text { subject to } & f_{i}(x) \leq 0, \quad i=1, \ldots, p \\
& h_{i}(x)=0, \quad i=1, \ldots, q
\end{array}
$$

variable $x \in \mathbb{R}^{n}$, domain $\mathcal{D}$, optimal value $p^{\star}$
Lagrangian: $\mathcal{L}: \mathbb{R}^{n} \times \mathbb{R}^{p} \times \mathbb{R}^{q} \rightarrow \mathbb{R}$, with $\operatorname{dom}(\mathcal{L})=\mathcal{D} \times \mathbb{R}^{p} \times \mathbb{R}^{q}$,

$$
\mathcal{L}(x, \lambda, \nu)=f_{0}(x)+\sum_{i=1}^{p} \lambda_{i} f_{i}(x)+\sum_{i=1}^{q} \nu_{i} h_{i}(x)
$$

- weighted sum of objective and constraint functions
- $\lambda_{i}$ is the Lagrange multiplier (dual variable) associated with $f_{i}(x) \leq 0$
- $\nu_{i}$ is the Lagrange multiplier (dual variable) associated with $h_{i}(x)=0$


## Karush-Kuhn-Tucker (KKT) Conditions

The following four conditions are called KKT conditions (for differentiable $f_{i}, h_{i}$ ):

- primal feasible: $\quad f_{i}(x) \leq 0, \quad i=1, \ldots, p$

$$
h_{i}(x)=0, \quad i=1, \ldots, q
$$

- dual feasible: $\lambda \succeq 0$
- complementary slackness: $\lambda_{i} f_{i}(x)=0$ for $i=1, \ldots, p$
- gradient of Lagrangian with respect to $x$ vanishes,

$$
\nabla f_{0}(x)+\sum_{i=1}^{p} \lambda_{i} \nabla f_{i}(x)+\sum_{i=1}^{q} \nu_{i} \nabla h_{i}(x)=0
$$

Generalizes optimality condition $\nabla f_{0}(x)=0$ for unconstrained problems.

## Gradient Descent

$$
\operatorname{minimize} \quad f_{0}(x)
$$

- $f_{0}$ convex, twice continuously differentiable
- we assume optimal value $p^{\star}=\inf _{x} f_{0}(x)$ is attained (and finite)


## Gradient Descent

$$
\operatorname{minimize} \quad f_{0}(x)
$$

- $f_{0}$ convex, twice continuously differentiable
- we assume optimal value $p^{\star}=\inf _{x} f_{0}(x)$ is attained (and finite)


## Gradient descent:

1. given a starting point $x \in \boldsymbol{\operatorname { d o m }}\left(f_{0}\right)$
2. repeat $x:=x-t \nabla f_{0}(x)$. (choose step size, $t$ )
3. until stopping criterion satisfied, e.g., $\left\|\nabla f_{0}(x)\right\|_{2} \leq \epsilon$.

- variants of gradient descent define step direction $\Delta x$ different to $-\nabla f_{0}(x)$


## Choosing Step Size

fixed schedule: set $t$ to a small constant or decay with each iteration exact line search: $t=\operatorname{argmin}_{t>0} f_{0}(x+t \Delta x)$
backtracking line search (with parameters $\alpha \in(0,1 / 2), \beta \in(0,1)$ )

- starting at $t=1$ with search direction $\Delta x$, repeat $t:=\beta t$ until

$$
f_{0}(x+t \Delta x)<f_{0}(x)+\alpha t \nabla f_{0}(x)^{T} \Delta x
$$

## Choosing Step Size

fixed schedule: set $t$ to a small constant or decay with each iteration exact line search: $t=\operatorname{argmin}_{t>0} f_{0}(x+t \Delta x)$
backtracking line search (with parameters $\alpha \in(0,1 / 2), \beta \in(0,1)$ )

- starting at $t=1$ with search direction $\Delta x$, repeat $t:=\beta t$ until

$$
f_{0}(x+t \Delta x)<f_{0}(x)+\alpha t \nabla f_{0}(x)^{T} \Delta x
$$



## Example

Gradient descent (even with exact line search) can be slow. E.g.,

$$
f_{0}(x)=x_{1}^{2}+\gamma x_{2}^{2}, \quad \gamma \gg 1
$$



## Newton's Method

$$
\Delta x_{\mathrm{nt}}=-\nabla^{2} f_{0}(x)^{-1} \nabla f_{0}(x)
$$

- $x+\Delta x_{\mathrm{nt}}$ minimizes the second-order approximation of $f_{0}$ at $x$,

$$
\hat{f}(x+v)=f_{0}(x)+\nabla f_{0}(x)^{T} v+\frac{1}{2} v^{T} \nabla^{2} f_{0}(x) v
$$

## Newton's method:

1. given a starting point $x \in \operatorname{dom}\left(f_{0}\right)$.
2. repeat $x:=x+t \Delta x_{\mathrm{nt}}$. (choose step size, $t$ )
3. until stopping criterion satisfied.

## Equality Constrained Methods

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x) \\
\text { subject to } & A x=b
\end{array}
$$

- $f_{0}$ convex, twice continuously differentiable
- $A \in \mathbb{R}^{q \times n}$ with $\boldsymbol{\operatorname { r a n k }}(A)=q$ (and $b \in \operatorname{range}(A)$ )
- we assume $p^{\star}$ is finite and attained


## Equality Constrained Methods

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x) \\
\text { subject to } & A x=b
\end{array}
$$

- $f_{0}$ convex, twice continuously differentiable
- $A \in \mathbb{R}^{q \times n}$ with $\boldsymbol{\operatorname { r a n k }}(A)=q$ (and $b \in \operatorname{range}(A)$ )
- we assume $p^{\star}$ is finite and attained
optimality condition: $x^{\star}$ is optimal iff there exists a $\nu^{\star}$ such that

$$
\nabla f_{0}\left(x^{\star}\right)+A^{T} \nu^{\star}=0, \quad A x^{\star}=b
$$

## Newton Step for Equality Constrained Optimisation

Newton step $\Delta x_{\text {nt }}$ of $f_{0}$ at feasible $x$ is given by solution $v$ of

$$
\left[\begin{array}{cc}
\nabla^{2} f_{0}(x) & A^{T} \\
A & 0
\end{array}\right]\left[\begin{array}{c}
v \\
w
\end{array}\right]=\left[\begin{array}{c}
-\nabla f_{0}(x) \\
0
\end{array}\right]
$$

- second row ensures that $x$ iterates stay feasible
- solves quadratic approximation of optimisation problem

$$
\begin{array}{ll}
\operatorname{minimize} & \hat{f}(x+v) \triangleq f_{0}(x)+\nabla f_{0}(x)^{T} v+\frac{1}{2} v^{T} \nabla^{2} f_{0}(x) v \\
\text { subject to } & A(x+v)=b
\end{array}
$$

- solves linear approximation of optimality condition


## The Barrier Method

For inequality constrained problems,

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x) \\
\text { subject to } & f_{i}(x) \leq 0, \quad i=1, \ldots, p \\
& A x=b
\end{array}
$$

## The Barrier Method

For inequality constrained problems,

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x) \\
\text { subject to } & f_{i}(x) \leq 0, \quad i=1, \ldots, p \\
& A x=b
\end{array}
$$

we reformulate using an indicator function,

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x)+\sum_{i=1}^{p} I_{\mathbb{R}_{-}}\left(f_{i}(x)\right) \\
\text { subject to } & A x=b
\end{array}
$$

where $I_{\mathbb{R}_{-}}(u)=0$ if $u \leq 0$ and $I_{\mathbb{R}_{-}}(u)=\infty$ otherwise,

## The Barrier Method

For inequality constrained problems,

$$
\begin{array}{ll}
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x) \\
\text { subject to } & f_{i}(x) \leq 0, \quad i=1, \ldots, p \\
& A x=b
\end{array}
\end{array}
$$

we reformulate using an indicator function,

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x)+\sum_{i=1}^{p} I_{\mathbb{R}_{-}}\left(f_{i}(x)\right) \\
\text { subject to } & A x=b
\end{array}
$$

where $I_{\mathbb{R}_{-}}(u)=0$ if $u \leq 0$ and $I_{\mathbb{R}_{-}}(u)=\infty$ otherwise, which we approximate with a logarithmic barrier

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x)-\frac{1}{t} \sum_{i=1}^{p} \log \left(-f_{i}(x)\right) \\
\text { subject to } & A x=b
\end{array}
$$

to get an equality constrained approximation.

## The Barrier Method

For inequality constrained problems,

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x) \\
\text { subject to } & f_{i}(x) \leq 0, \quad i=1, \ldots, p \\
& A x=b
\end{array}
$$

we reformulate using an indicator function,

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x)+\sum_{i=1}^{p} I_{\mathbb{R}_{-}}\left(f_{i}(x)\right) \\
\text { subject to } & A x=b
\end{array}
$$

where $I_{\mathbb{R}_{-}}(u)=0$ if $u \leq 0$ and $I_{\mathbb{R}_{-}}(u)=\infty$ otherwise, which we approximate with a logarithmic barrier

$$
\begin{array}{ll}
\operatorname{minimize} & f_{0}(x)-\frac{1}{t} \sum_{i=1}^{p} \log \left(-f_{i}(x)\right) \\
\text { subject to } & A x=b
\end{array}
$$

to get an equality constrained approximation.



## Algorithms for Large Scale Problems

- for large scale problems, e.g., deep learning, Newton's method is too expensive
- even computing the true gradient may be too expensive
- many loss functions in machine learning decompose over train data $\left\{\left(x_{i}, y_{i}\right)\right\}_{i=1}^{m}$,

$$
L(\theta)=\sum_{i=1}^{m} \ell\left(f\left(x_{i} ; \theta\right), y_{i}\right)
$$

- SGD approximates the gradient on mini-batches $\mathcal{I} \subseteq\{1, \ldots, m\}$

$$
\widehat{\nabla_{\theta} L}=\sum_{i \in \mathcal{I}} \nabla_{\theta} \ell\left(f\left(x_{i} ; \theta\right), y_{i}\right)
$$

- under mild assumptions $E\left[\widehat{\nabla_{\theta} L}\right]=\nabla_{\theta} L$
- for constrained problems can project back onto feasible set

Many, many other schemes and variations!

## lecture 2

## Lecture 2: Differentiable Optimisation and Deep Learning



| Springer Series in Operations Research <br> and finundal Eagineering |
| :---: |
| Asen L. Dontchey R. Tyrrell Rockafellar |
| Implicit Functions and Solution Mappings <br> A View from Variational Analysis Second Edition |


|  | (1) |
| :---: | :---: |
| On the salution of cowex bilevel aptimization penbiem <br> Q. Itwape" , 4- Freatry |  |
|  |  |
|  |  |
|  |  |
|  |  |
|  |  |
|  |  |
| 1 trimerater |  |
| \%tow | (a) |
| x+1.ancm |  |
|  |  |


| On Diffetentiating Parameterizad Argmin and Argmax Problems with Application to Bi-level Opeimeation | Differentiable Optimization-Based Modeling for Machine Learning |
| :---: | :---: |
|  | Brandon Aimas |
|  | Cx-Cstame |
|  | May 2179 |
|  |  |
| tre. |  |
|  | Camer Mdia Pemery |
| Hemerex |  |
| $\underline{\square}$ | Thusk Commitae |
| - - =. |  |
| 1 Iteruanction |  |
| Ausmen |  |
| , enter | Seite |
|  |  |
|  | Capmetio ans bristo Amm |
|  |  |
| Nut |  |

## Machine Learning from 10,000ft



## Machine Learning from 10,000ft



$$
\text { minimize (over } \theta) \quad \sum_{(x, y) \sim \mathcal{X} \times \mathcal{Y}} L\left(f_{\theta}(x), y\right)
$$

- loss $L$ - what to do
- model $f_{\theta}$ - how to do it
- optimised by gradient descent


## Deep Learning as an End-to-end Computation Graph

Deep learning does this by defining a function (equiv. computation graph) composed of many simple parametrized functions (equiv. computation nodes).


## Backward Pass



## Example 1.

$$
\frac{\partial L}{\partial \theta_{7}}=\frac{\partial L}{\partial y} \frac{\partial y}{\partial z_{7}} \frac{\partial z_{7}}{\partial \theta_{7}}
$$

## Backward Pass



Example 2.

$$
\frac{\partial L}{\partial \theta_{1}}=\frac{\partial L}{\partial y}\left(\frac{\partial y}{\partial z_{4}} \frac{\partial z_{4}}{\partial z_{3}} \frac{\partial z_{3}}{\partial z_{2}} \frac{\partial z_{2}}{\partial z_{1}}+\frac{\partial y}{\partial z_{7}} \frac{\partial z_{7}}{\partial z_{6}} \frac{\partial z_{6}}{\partial z_{5}} \frac{\partial z_{5}}{\partial z_{4}}\right) \frac{\partial z_{1}}{\partial \theta_{1}}
$$

## Deep Learning Node



- Forward pass: compute output $y$ as a function of the input $x$ (and model parameters $\theta$ ).
- Backward pass: compute the derivative of the loss with respect to the input $x$ (and model parameters $\theta$ ) given the derivative of the loss with respect to the output $y$.


## Notational Aside (Often Sloppy)

For scalar-valued functions:

$$
\text { total derivative: } \frac{\mathrm{d} f}{\mathrm{~d} x} \quad \text { partial derivative: } \frac{\partial f}{\partial x}
$$

## Notational Aside (Often Sloppy)

For scalar-valued functions:

$$
\text { total derivative: } \frac{\mathrm{d} f}{\mathrm{~d} x} \quad \text { partial derivative: } \quad \frac{\partial f}{\partial x}
$$

For multi-dimensional scalar-valued functions, $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ :

$$
\nabla f(x)=\left(\frac{\mathrm{d} f}{\mathrm{~d} x_{1}}, \ldots, \frac{\mathrm{~d} f}{\mathrm{~d} x_{n}}\right) \in \mathbb{R}^{n}
$$

## Notational Aside (Often Sloppy)

For scalar-valued functions:

$$
\text { total derivative: } \frac{\mathrm{d} f}{\mathrm{~d} x} \quad \text { partial derivative: } \quad \frac{\partial f}{\partial x}
$$

For multi-dimensional scalar-valued functions, $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ :

$$
\nabla f(x)=\left(\frac{\mathrm{d} f}{\mathrm{~d} x_{1}}, \ldots, \frac{\mathrm{~d} f}{\mathrm{~d} x_{n}}\right) \in \mathbb{R}^{n}
$$

For multi-dimensional vector-valued functions, $f: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ :

$$
\frac{\mathrm{d}}{\mathrm{~d} x} f(x)=\left[\begin{array}{ccc}
\frac{\mathrm{d} f_{1}}{\mathrm{~d} x_{1}} & \cdots & \frac{\mathrm{~d} f_{1}}{\mathrm{~d} x_{n}} \\
\vdots & \ddots & \vdots \\
\frac{\mathrm{~d} f_{m}}{\mathrm{~d} x_{1}} & \cdots & \frac{\mathrm{~d} f_{m}}{\mathrm{~d} x_{n}}
\end{array}\right] \in \mathbb{R}^{m \times n} \quad\left(\frac{\partial}{\partial x} f(x, y) \text { for partial }\right)
$$

Sometimes D and $\mathrm{D}_{X}$ for $\frac{\mathrm{d}}{\mathrm{d} x}$ and $\frac{\partial}{\partial x}$, respectively.

## Automatic Differentiation (AD)

- algorithmic procedure that produces code for computing exact derivatives
- assumes numeric computations are composed of a small set of elementary operations that we know how to differentiate
- arithmetic, exp, log, trigonometric
- workhorse of modern machine learning that greatly reduces development effort


## Automatic Differentiation (AD)

- algorithmic procedure that produces code for computing exact derivatives
- assumes numeric computations are composed of a small set of elementary operations that we know how to differentiate
- arithmetic, exp, log, trigonometric
- workhorse of modern machine learning that greatly reduces development effort
- two flavours
- (forward mode) propagates results on the first-order approximation $x+\Delta x$ forward through the computations
- (reverse mode) builds a program to compute derivative based on the chain rule re-using computation where applicable

$$
\frac{\mathrm{d} L}{\mathrm{~d} x}=\frac{\mathrm{d} L}{\mathrm{~d} y} \frac{\mathrm{~d} y}{\mathrm{~d} x}
$$

- different deep learning frameworks use slightly different approaches (explicit graph construction versus implicit operator tracking)


## Computing $1 / \sqrt{x}$

```
float Q_rsqrt( float number )
{
    long i;
    float x2, y;
    const float threehalfs = 1.5F;
    x2 = number * 0.5F;
    y = number;
    i = * ( long * ) &y; // evil floating point bit level hacking
    i = 0x5f3759df - ( i >> 1 ); // what the f**k?
    y = * ( float * ) &i;
    y = y * ( threehalfs - ( x2 * y * y ) ); // 1st iter
    // y = y * ( threehalfs - ( x2 * y * y ) ); // 2nd iter, can be removed
    return y;
}
```


## Separate Forward and Backward Operations



## Imperative vs Declarative Nodes



- imperative node
- input-output relationship explicit,

$$
y=\tilde{f}(x ; \theta)
$$

## Imperative vs Declarative Nodes



- imperative node
- input-output relationship explicit,

$$
y=\tilde{f}(x ; \theta)
$$



- declarative node
- input-output relationship specified as solution to an optimisation problem,

$$
y \in \underset{u \in C(x)}{\arg \min } f(x, u ; \theta)
$$

## Imperative vs Declarative Nodes



- imperative node
- input-output relationship explicit,

$$
y=\tilde{f}(x ; \theta)
$$



- declarative node
- input-output relationship specified as solution to an optimisation problem,

$$
y \in \underset{u \in C(x)}{\arg \min } f(x, u ; \theta)
$$

can co-exist in the same computation graph (network)

## Average Pooling Example

$$
\left\{x_{i} \in \mathbb{R}^{m} \mid i=1, \ldots, n\right\} \rightarrow \mathbb{R}^{m}
$$

- imperative specification
- declarative specification

$$
y=\frac{1}{n} \sum_{i=1}^{n} x_{i}
$$

$$
y=\operatorname{argmin}_{u \in \mathbb{R}^{m}} \sum_{i=1}^{n}\left\|u-x_{i}\right\|^{2}
$$

## Average Pooling Example

$$
\left\{x_{i} \in \mathbb{R}^{m} \mid i=1, \ldots, n\right\} \rightarrow \mathbb{R}^{m}
$$

- imperative specification

$$
y=\frac{1}{n} \sum_{i=1}^{n} x_{i}
$$

- declarative specification

$$
y=\operatorname{argmin}_{u \in \mathbb{R}^{m}} \sum_{i=1}^{n}\left\|u-x_{i}\right\|^{2}
$$

- can be easily varied, e.g., made robust

$$
y=\operatorname{argmin}_{u \in \mathbb{R}^{m}} \sum_{i=1}^{n} \phi\left(u-x_{i}\right)
$$

for some penalty function $\phi$

## Average Pooling Example

$$
\left\{x_{i} \in \mathbb{R}^{m} \mid i=1, \ldots, n\right\} \rightarrow \mathbb{R}^{m}
$$



- declarative specification

$$
y=\operatorname{argmin}_{u \in \mathbb{R}^{m}} \sum_{i=1}^{n}\left\|u-x_{i}\right\|^{2}
$$

- can be easily varied, e.g., made robust

$$
y=\operatorname{argmin}_{u \in \mathbb{R}^{m}} \sum_{i=1}^{n} \phi\left(u-x_{i}\right)
$$

for some penalty function $\phi$

## Bi-level Optimisation: Stackelberg Games

Consider two players, a leader and a follower

- the market dictates the price it's willing to pay for some goods based on supply, i.e., quantity produced by both players, $P\left(q_{1}+q_{2}\right)$
- each player has a cost structure associated with producing goods, $C_{i}\left(q_{i}\right)$ and wants to maximize profits, $q_{i} P\left(q_{1}+q_{2}\right)-C_{i}\left(q_{i}\right)$
- the leader picks a quantity of goods to produce knowing that the follower will respond optimally. In other words, the leader solves

$$
\begin{array}{ll}
\operatorname{maximize}\left(\text { over } q_{1}\right) & q_{1} P\left(q_{1}+q_{2}\right)-C_{1}\left(q_{1}\right) \\
\text { subject to } & q_{2} \in \operatorname{argmax}_{q} q P\left(q_{1}+q\right)-C_{2}(q)
\end{array}
$$

## Solving Bi-level Optimisation Problems

$$
\begin{array}{ll}
\operatorname{minimize}(\text { over } x) & L(x, y) \\
\text { subject to } & y \in \operatorname{argmin}_{u \in C(x)} f(x, u)
\end{array}
$$

## Solving Bi-level Optimisation Problems

$$
\begin{array}{ll}
\operatorname{minimize}(\text { over } x) & L(x, y) \\
\text { subject to } & y \in \operatorname{argmin}_{u \in C(x)} f(x, u)
\end{array}
$$

- closed-form solution: substitute for $y$ in upper-level problem (if possible)

$$
\text { minimize (over } x) \quad L(x, y(x))
$$

## Solving Bi-level Optimisation Problems

```
minimize (over x) L(x,y)
subject to }\quady\in\mp@subsup{\operatorname{argmin}}{u\inC(x)}{}f(x,u
```

- closed-form solution: substitute for $y$ in upper-level problem (if possible)

$$
\text { minimize (over } x) \quad L(x, y(x))
$$

- convex lower-level problem: replace lower-level problem with sufficient optimality conditions (e.g., KKT conditions),

$$
\begin{array}{ll}
\operatorname{minimize}(\text { over } x, y) & L(x, y) \\
\text { subject to } & h(x, y)=0
\end{array}
$$

## Solving Bi-level Optimisation Problems

$$
\begin{array}{ll}
\operatorname{minimize}(\text { over } x) & L(x, y) \\
\text { subject to } & y \in \operatorname{argmin}_{u \in C(x)} f(x, u)
\end{array}
$$

- closed-form solution: substitute for $y$ in upper-level problem (if possible)

$$
\text { minimize (over } x) \quad L(x, y(x))
$$

- convex lower-level problem: replace lower-level problem with sufficient optimality conditions (e.g., KKT conditions),

$$
\begin{array}{ll}
\operatorname{minimize}(\text { over } x, y) & L(x, y) \\
\text { subject to } & h(x, y)=0
\end{array}
$$

- gradient descent: compute gradient of lower-level solution $y$ with respect to $x$, and use the chain rule to get the total derivative,

$$
x \leftarrow x-\eta\left(\frac{\partial L(x, y)}{\partial x}+\frac{\partial L(x, y)}{\partial y} \frac{\mathrm{~d} y}{\mathrm{~d} x}\right)
$$

## Solving Bi-level Optimisation Problems

```
minimize (over x) L(x,y)
subject to }\quady\in\mp@subsup{\operatorname{argmin}}{u\inC(x)}{}f(x,u
```

- closed-form solution: substitute for $y$ in upper-level problem (if possible)

$$
\text { minimize (over } x) \quad L(x, y(x))
$$

- convex lower-level problem: replace lower-level problem with sufficient optimality conditions (e.g., KKT conditions),

$$
\begin{array}{ll}
\operatorname{minimize}(\text { over } x, y) & L(x, y) \\
\text { subject to } & h(x, y)=0
\end{array}
$$

- gradient descent: compute gradient of lower-level solution $y$ with respect to $x$, and use the chain rule to get the total derivative,

$$
x \leftarrow x-\eta\left(\frac{\partial L(x, y)}{\partial x}+\frac{\partial L(x, y)}{\partial y} \frac{\mathrm{~d} y}{\mathrm{~d} x}\right)
$$

- by back-propagating through optimisation procedure or implicit differentiation


## Parametrized Optimisation

In the context of deep learning the upper-level Stackelberg problem is the learning problem and the lower-level Stackelberg problem is the inference problem.

A declarative node defines a family of problems indexed by continuous variable $x \in \mathbb{R}^{n}$,

$$
\left\{\begin{array}{ll}
\operatorname{minimize}\left(\text { over } u \in \mathbb{R}^{m}\right) & f_{0}(x, u) \\
\text { subject to } & f_{i}(x, u) \leq 0, \quad i=1, \ldots, p \\
& h_{i}(x, u)=0, \quad i=1, \ldots, q
\end{array}\right\}_{x \in \mathbb{R}^{n}}
$$

## Parametrized Optimisation Example


input/parameter, $x$


## Parametrized Optimisation Example



Main question: How do we compute $\frac{\mathrm{d}}{\mathrm{d} x} \operatorname{argmin}_{u} f(x, u)$ ?

## Dini's Implicit Function Theorem

Consider the solution mapping associated with the equation $f(x, u)=0$,

$$
Y: x \mapsto\left\{u \in \mathbb{R}^{m} \mid f(x, u)=0\right\} \text { for } x \in \mathbb{R}^{n}
$$

We are interested in how elements of $Y(x)$ change as a function of $x$.

## Dini's Implicit Function Theorem

Consider the solution mapping associated with the equation $f(x, u)=0$,

$$
Y: x \mapsto\left\{u \in \mathbb{R}^{m} \mid f(x, u)=0\right\} \text { for } x \in \mathbb{R}^{n}
$$

We are interested in how elements of $Y(x)$ change as a function of $x$.
Theorem
Let $f: \mathbb{R}^{n} \times \mathbb{R}^{m} \rightarrow \mathbb{R}^{m}$ be differentiable in a neighbourhood of $(x, u)$ and such that $f(x, u)=0$, and let $\frac{\partial}{\partial u} f(x, u)$ be nonsingular. Then the solution mapping $Y$ has a single-valued localization $y$ around $x$ for $u$ which is differentiable in a neighbourhood $\mathcal{X}$ of $x$ with Jacobian satisfying

$$
\frac{d y(x)}{d x}=-\left(\frac{\partial f(x, y(x))}{\partial y}\right)^{-1} \frac{\partial f(x, y(x))}{\partial x}
$$

for every $x \in \mathcal{X}$.

## Unit Circle Example



$$
\begin{aligned}
y & = \pm \sqrt{1-x^{2}} \\
\frac{\mathrm{~d} y}{\mathrm{~d} x} & =\frac{\mp 2 x}{2 \sqrt{1-x^{2}}}=-\frac{x}{y}
\end{aligned}
$$

$$
\begin{aligned}
f(x, y) & =x^{2}+y^{2}-1 \\
\frac{\mathrm{~d} y}{\mathrm{~d} x} & =-\left(\frac{\partial f}{\partial y}\right)^{-1}\left(\frac{\partial f}{\partial x}\right) \\
& =-\left(\frac{1}{2 y}\right)(2 x)=-\frac{x}{y}
\end{aligned}
$$

## Differentiating Unconstrained Optimisation Problems

Let $f: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be twice differentiable and let

$$
y(x) \in \operatorname{argmin}_{u} f(x, u)
$$

then for non-zero Hessian

$$
\frac{\mathrm{d} y(x)}{\mathrm{d} x}=-\left(\frac{\partial^{2} f}{\partial y^{2}}\right)^{-1} \frac{\partial^{2} f}{\partial x \partial y}
$$

## Differentiating Unconstrained Optimisation Problems

Let $f: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be twice differentiable and let

$$
y(x) \in \operatorname{argmin}_{u} f(x, u)
$$

then for non-zero Hessian

$$
\frac{\mathrm{d} y(x)}{\mathrm{d} x}=-\left(\frac{\partial^{2} f}{\partial y^{2}}\right)^{-1} \frac{\partial^{2} f}{\partial x \partial y}
$$



Proof. The derivative of $f$ vanishes at $(x, y)$, i.e., $y \in \operatorname{argmin}_{u} f(x, u) \Longrightarrow \frac{\partial f(x, y)}{\partial y}=0$.

$$
\begin{aligned}
\text { LHS : } & \frac{\mathrm{d}}{\mathrm{~d} x} \frac{\partial f(x, y)}{\partial y} & =\frac{\partial^{2} f(x, y)}{\partial x \partial y}+\frac{\partial^{2} f(x, y)}{\partial y^{2}} \frac{\mathrm{~d} y}{\mathrm{~d} x} \\
\text { RHS : } & \frac{\mathrm{d}}{\mathrm{~d} x} 0 & =0
\end{aligned}
$$

Equating and rearranging gives the result.

## Differentiable Optimisation: Big Picture Idea



## Differentiating Equality Constrained Optimisation Problems

Consider functions $f: \mathbb{R}^{n} \times \mathbb{R}^{m} \rightarrow \mathbb{R}$ and $h: \mathbb{R}^{n} \times \mathbb{R}^{m} \rightarrow \mathbb{R}^{q}$. Let

$$
\begin{aligned}
& y(x) \in \underset{\text { subject to }}{\arg \min _{u \in \mathbb{R}^{m}}} \quad h(x, u) \\
& h(x, u)=0_{q}
\end{aligned}
$$

Assume that $y(x)$ exists, that $f$ and $h$ are twice differentiable in the neighbourhood of $(x, y(x))$, and that $\operatorname{rank}\left(\frac{\partial h(x, y)}{\partial y}\right)=q$.

## Differentiating Equality Constrained Optimisation Problems

Consider functions $f: \mathbb{R}^{n} \times \mathbb{R}^{m} \rightarrow \mathbb{R}$ and $h: \mathbb{R}^{n} \times \mathbb{R}^{m} \rightarrow \mathbb{R}^{q}$. Let

$$
\begin{array}{rl}
y(x) \in \underset{u}{ } \underset{ }{\arg \min _{u \in \mathbb{R}^{m}}} & f(x, u) \\
\text { subject to } & h(x, u)=0_{q}
\end{array}
$$

Assume that $y(x)$ exists, that $f$ and $h$ are twice differentiable in the neighbourhood of $(x, y(x))$, and that $\operatorname{rank}\left(\frac{\partial h(x, y)}{\partial y}\right)=q$. Then for $H$ non-singular

$$
\frac{\mathrm{d} y(x)}{\mathrm{d} x}=H^{-1} A^{T}\left(A H^{-1} A^{T}\right)^{-1}\left(A H^{-1} B-C\right)-H^{-1} B
$$

where

$$
\begin{array}{ll}
A=\frac{\partial h(x, y)}{\partial y} \in \mathbb{R}^{q \times m} & B=\frac{\partial^{2} f(x, y)}{\partial x \partial y}-\sum_{i=1}^{q} \nu_{i} \frac{\partial^{2} h_{i}(x, y)}{\partial x \partial y} \\
C=\frac{\partial h(x, y)}{\partial x} \in \mathbb{R}^{q \times n} & H=\frac{\partial^{2} f(x, y)}{\partial y^{2}}-\sum_{i=1}^{q} \nu_{i} \frac{\partial^{2} h_{i}(x, y)}{\partial y^{2}}
\end{array}
$$

and $\nu \in \mathbb{R}^{q}$ satisfies $\nu^{T} A=\frac{\partial f(x, y)}{\partial y}$.

## Dealing with Inequality Constraints

$$
\begin{array}{ll}
y(x) \in \underset{u \in \mathbb{R}^{m}}{\arg \min _{u}(x, u)} & f_{0}(x, \ldots, p
\end{array}
$$

- Replace inequality constraints with log-barrier approximation (see last lecture)
- Treat as equality constraints if active ( $y_{2}$ or $y_{3}$ ) and ignore otherwise ( $y_{1}$ or $y_{3}$ )

- may lead to one-sided gradients since $\lambda \succeq 0$


## Automatic Differentiation for Differentiable Optimisation

- At one extreme we can try back propagate through the optimisation algorithm (i.e., unrolling the optimisation procedure using automatic differentiation)
- At the other extreme we can use the implicit differentiation result to hand-craft efficient backward pass code
- There are two options in between:
- Use automatic differentiation to obtain quantities $A, B, C$ and $H$ from software implementations of the objective and (active) constraint functions
- Implement the optimality condition $\nabla \mathcal{L}=0$ in software and automatically differentiate that
(in the next lecture we will see examples of the first two)


## Vector-Jacobian Product

For brevity consider the unconstrained optimisation case. The backward pass computes

$$
\begin{aligned}
\frac{\mathrm{d} L}{\mathrm{~d} x} & =\frac{\mathrm{d} L}{\mathrm{~d} y} \frac{\mathrm{~d} y}{\mathrm{~d} x} \\
& =\underbrace{\left(v^{T}\right)}_{\mathbb{R}^{1 \times m}} \underbrace{\left(-H^{-1} B\right)}_{\mathbb{R}^{m \times n}}
\end{aligned}
$$

evaluation order: $\quad-v^{T}\left(H^{-1} B\right) \quad\left(-v^{T} H^{-1}\right) B$

$$
\text { cost }^{\dagger}: O\left(m^{2} n+m n\right) \quad O\left(m^{2}+m n\right)
$$

${ }^{\dagger}$ assumes $H^{-1}$ is already factored (in $O\left(m^{3}\right)$ if unstructured, less if structured)

## Summary and Open Questions

- optimisation problems can be embedded inside deep learning models
- back-propagation by either unrolling the optimisation algorithm or implicit differentiation of the optimality conditions
- the former is easy to implement using automatic differentiation but memory intensive
- the latter requires that solution be strongly convex locally (i.e., invertible $H$ )
- but does not need to know how the problem was solved, nor store intermediate forward-pass calculations
- computing $H^{-1}$ may be costly


## Summary and Open Questions

- optimisation problems can be embedded inside deep learning models
- back-propagation by either unrolling the optimisation algorithm or implicit differentiation of the optimality conditions
- the former is easy to implement using automatic differentiation but memory intensive
- the latter requires that solution be strongly convex locally (i.e., invertible $H$ )
- but does not need to know how the problem was solved, nor store intermediate forward-pass calculations
- computing $H^{-1}$ may be costly
- active area of research and many open questions
- Are declarative nodes slower?
- Do declarative nodes give theoretical guarantees?
- How best to handle non-smooth or discrete optimization problems?
- What about problems with multiple solutions?
- What if the forward pass solution is suboptimal?
- Can problems become infeasible during learning?
- ...


## lecture 3

## Lecture 3: Examples and Applications

| = menm |  |
| :---: | :---: |
| Deep Declarative Networks |  |
|  |  |
|  |  |
| numung coose |  |
| $\cdots$ |  |
|  |  |
|  |  |
| - |  |
|  |  |
| Ucense |  |


https://deepdeclarativenetworks.com

## Common Theme



## Differentiable Least Squares

Consider our old friend, the least-squares problem,

$$
\operatorname{minimize} \quad\|A x-b\|_{2}^{2}
$$

parameterized by $A$ and $b$ and with closed-form solution $x^{\star}=\left(A^{T} A\right)^{-1} A^{T} b$.

## Differentiable Least Squares

Consider our old friend, the least-squares problem,

$$
\operatorname{minimize}\|A x-b\|_{2}^{2}
$$

parameterized by $A$ and $b$ and with closed-form solution $x^{\star}=\left(A^{T} A\right)^{-1} A^{T} b$.

We are interested in derivatives of the solution with respect to the elements of $A$,

$$
\frac{\mathrm{d} x^{\star}}{\mathrm{d} A_{i j}}=\frac{\mathrm{d}}{\mathrm{~d} A_{i j}}\left(A^{T} A\right)^{-1} A^{T} b
$$

We could also compute derivatives with respect to elements of $b$ (but not here).

## Least Squares Backward Pass

The backward pass combines $\frac{\mathrm{d} x^{\star}}{\mathrm{d} A_{i j}}$ with $v^{T}=\frac{\mathrm{d} L}{\mathrm{~d} x^{\star}}$ via the vector-Jacobian product. After some algebraic manipulation (see lecture notes) we get

$$
\left(\frac{\mathrm{d} L}{\mathrm{~d} A}\right)^{T}=w r^{T}-x^{\star}(A w)^{T}
$$

where $w^{T}=v^{T}\left(A^{T} A\right)^{-1}$.

## Least Squares Backward Pass

The backward pass combines $\frac{\mathrm{d} x^{\star}}{\mathrm{d} A_{i j}}$ with $v^{T}=\frac{\mathrm{d} L}{\mathrm{~d} x^{\star}}$ via the vector-Jacobian product. After some algebraic manipulation (see lecture notes) we get

$$
\left(\frac{\mathrm{d} L}{\mathrm{~d} A}\right)^{T}=w r^{T}-x^{\star}(A w)^{T}
$$

$\in \mathbb{R}^{m \times n}$
where $w^{T}=v^{T}\left(A^{T} A\right)^{-1}$.

- $\left(A^{T} A\right)^{-1}$ is used in both the forward and backward pass
- factored once to solve for $x$, e.g., into $A=Q R$
- cache $R$ and re-use when computing gradients


## Aside: PyTorch and Batched Data

Deep learning frameworks process data in batches, passed as tensors, for stochastic gradient descent. The first dimension of the tensor is the batch dimension.

Example. For the operation $y=A x+b$ we might have

$$
\begin{align*}
X & =\left\{x^{(1)}, \ldots, x^{(K)}\right\}  \tag{input}\\
Y & =\left\{A x^{(1)}+b, \ldots, A x^{(K)}+b\right\} \tag{output}
\end{align*}
$$

Many PyTorch functions are batch-aware, e.g., torch.bmm. For many operations the einsum function and broadcasting are particularly useful, e.g.,

```
y = torch.einsum("ij,kj->ki", A, x) + b
```

computes $y=A x^{(k)}+b$ on each element $k=1, \ldots, K$ of the batch.

## PyTorch Implementation: Forward Pass

```
class LeastSquaresFcn(torch.autograd.Function):
    """PyTorch autograd function for least squares."""
    @staticmethod
    def forward(ctx, A, b):
        B, M, N = A.shape
        assert b.shape == (B, M, 1)
        with torch.no_grad():
            Q, R = torch.linalg.qr(A, mode='reduced')
            x = torch.linalg.solve_triangular(R,
                torch.bmm(b.view(B, 1, M), Q).view(B, N, 1), upper=True)
        # save state for backward pass
        ctx.save_for_backward(A, b, x, R)
        # return solution
        return x
```

$$
\begin{aligned}
A & =Q R \\
x & =R^{-1}\left(Q^{T} b\right)
\end{aligned}
$$

(solves $R x=Q^{T} b$ )

## PyTorch Implementation: Backward Pass

```
@staticmethod
def backward(ctx, dx):
    \# check for None tensors
    if \(d x\) is None:
        return None, None
    \# unpack cached tensors
    A, \(b, x, R=c t x . s a v e d \_t e n s o r s\)
    B, \(M, N=\) A. shape
    \(\mathrm{dA}, \mathrm{db}=\) None, None
    \(\mathrm{w}=\) torch.linalg.solve_triangular (R,
        torch.linalg.solve_triangular (torch.transpose (R, 2, 1),
        dx, upper=False), upper=True)
    Aw = torch.bmm(A, w)
    if ctx. needs_input_grad[0]:
        \(r=b-t o r c h . b m m(A, x)\)
        \(d A=\) torch.einsum("bi,bj->bij", r.view (B,M), w.view(B,N)) - \}
            torch.einsum("bi,bj->bij", Aw.view (B, M), x.view (B,N))
    if ctx. needs_input_grad[1]:
        \(\mathrm{db}=\mathrm{Aw}\)
    \# return gradients
    return \(d A\), \(d b\)
```


## Example

Bi-level optimisation problem with lower-level least squares:

$$
\begin{array}{ll}
\operatorname{minimize} & \frac{1}{2}\left\|x^{\star}-x^{\text {target }}\right\|_{2}^{2} \\
\text { subject to } & x^{\star}=\operatorname{argmin}_{x}\|A x-b\|_{2}^{2}
\end{array}
$$


with upper-level variable $A \in \mathbb{R}^{m \times n}$.

## Profiling


(problems with $m=2 n$; run for 1000 iterations on CPU using PyTorch 1.13.0)

## Profiling


(problems with $m=2 n$; run for 1000 iterations on CPU using PyTorch 1.13.0)

## Optimal Transport

One view of optimal transport is as a matching problem

- from an $m$-by- $n$ cost matrix $M$
- to an $m$-by- $n$ probability matrix $P$,
often formulated with an entropic regularisation term,

$$
\begin{array}{ll}
\operatorname{minimize} & \langle M, P\rangle+\frac{1}{\gamma}\langle P, \log P\rangle \\
\text { subject to } & P \mathbf{1}=r \\
& P^{T} \mathbf{1}=c
\end{array}
$$

with $\mathbf{1}^{T} r=\mathbf{1}^{T} c=1$.
The row and column sum constraints ensure that $P$ is a
 doubly stochastic matrix (lies within the convex hull of permutation matrices).

## Solving Entropic Optimal Transport

Solution takes the form

$$
P_{i j}=\alpha_{i} \beta_{j} e^{-\gamma M_{i j}}
$$

and can be found using the Sinkhorn algorithm,

- Set $K_{i j}=e^{-\gamma M_{i j}}$ and $\alpha, \beta \in \mathbb{R}_{++}^{n}$
- Iterate until convergence,

$$
\begin{aligned}
& \alpha \leftarrow r \oslash K \beta \\
& \beta \leftarrow c \oslash K^{T} \alpha
\end{aligned}
$$

where $\oslash$ denotes componentwise division

- Return $P=\boldsymbol{\operatorname { d i a g }}(\alpha) K \boldsymbol{\operatorname { d i a g }}(\beta)$


## Differentiable Optimal Transport

- Option 1: back-propagate through Sinkhorn algorithm


## Differentiable Optimal Transport

- Option 1: back-propagate through Sinkhorn algorithm
- Option 2: use the implicit differentiation result

$$
\underbrace{\frac{\mathrm{d} L}{\mathrm{~d} M}}_{m \text {-by-n }}=\underbrace{\frac{\mathrm{d} L}{\mathrm{~d} P}}_{m \text {-by-n }} \overbrace{\frac{\mathrm{d} P}{\mathrm{~d} M}}^{\mathrm{d} M}
$$

## Differentiable Optimal Transport

- Option 1: back-propagate through Sinkhorn algorithm
- Option 2: use the implicit differentiation result

$$
\underbrace{\frac{\mathrm{d} L}{\mathrm{~d} M}}_{1 \text {-by-mn }}=\underbrace{\frac{\mathrm{d} L}{\mathrm{~d} P}}_{1 \text {-by- } m n} \overbrace{\frac{\mathrm{~d} P}{\frac{\mathrm{~d} P}{\mathrm{~d} M}}}^{m n \text {-by- } m n} \quad \text { (think of vectorising } M \text { and } P \text { ) }
$$

## Optimal Transport Gradient

Derivation of the optimal transport gradient is quite tedious (see notes). The result:

$$
\begin{aligned}
\frac{\mathrm{d} L}{\mathrm{~d} M} & =\frac{\mathrm{d} L}{\mathrm{~d} P}\left(H^{-1} A^{T}\left(A H^{-1} A^{T}\right)^{-1} A H^{-1}-H^{-1}\right) B \\
& =\gamma \frac{\mathrm{d} L}{\mathrm{~d} P} \boldsymbol{\operatorname { d i a g }}(P)\left[\begin{array}{l}
A_{1} \\
A_{2}
\end{array}\right]^{T}\left[\begin{array}{ll}
\Lambda_{11} & \Lambda_{12} \\
\Lambda_{12}^{T} & \Lambda_{22}
\end{array}\right]\left[\begin{array}{c}
A_{1} \\
A_{2}
\end{array}\right] \operatorname{diag}(P)-\gamma \frac{\mathrm{d} L}{\mathrm{~d} P} \mathbf{d i a g}(P)
\end{aligned}
$$

where

$$
\left[\begin{array}{l}
A_{1} \\
A_{2}
\end{array}\right]=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \cdots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \cdots & I_{n \times n}
\end{array}\right] \quad\left(A H^{-1} A^{T}\right)^{-1}=\frac{1}{\gamma}\left[\begin{array}{cc}
\Lambda_{11} & \Lambda_{12} \\
\Lambda_{12}^{T} & \Lambda_{22}
\end{array}\right] \quad \begin{array}{ll} 
& =\frac{1}{\gamma}\left[\begin{array}{cc}
\operatorname{diag}\left(r_{2: m}\right) & P_{2: m, 1: n} \\
P_{2: m, 1: n}^{T} & \operatorname{diag}(c)
\end{array}\right]^{-1}
\end{array}
$$

## Implementation

```
@staticmethod
def backward(ctx, dJdP)
    # unpacked cached tensors
    M, r, c, P = ctx.saved_tensors
    batches, m, n = P.shape
    # initialize backward gradients (-v^T H^{-1} B)
    dLdM = -1.0 * gamma * P * dLdP
    # compute [vHAt1, vHAt2] = -v^T H^{-1} A^T
    vHAt1, vHAt2 = sum(dJdM[:, 1:m, 0:n], dim=2), sum(dJdM, dim=1)
    # compute [v1, v2] = -v^T H^{-1} A^T (A H^{-1] A^ T)^{-1}
    P_over_c = P[:, 1:m, 0:n] / c.view(batches, 1, n)
    lmd_11 = cholesky(diag_embed(r[:, 1:m]) - einsum("bij,bkj->bik", P[:, 1:m, 0:n], P_over_c))
    lmd_12 = cholesky_solve(P_over_c, lmd_11)
    lmd_22 = diag_embed(1.0 / c) + einsum("bji,bjk->bik", lmd_12, P_over_c)
    v1 = cholesky_solve(vHAt1.view(batches, m-1, 1), lmd_11).view(batches, m-1) -
        einsum("bi,bji->bj", vHAt2, lmd_12)
    v2 = einsum("bi,bij->bj", vHAt2, lmd_22) - einsum("bi,bij->bj", vHAt1, lmd_12)
    # compute V^T H^{-1} A^T (A H^{-1] A^^T)^{-1} A H
    dLdM[:, 1:m, 0:n] -= v1.view(batches, m-1, 1) * P[:, 1:m, 0:n]
    dJdM -= v2.view(batches, 1, n) * P
    # return gradients
    return dJdM
```


## Experiment

Bi-level optimisation problem with lower-level optimal transport problem:
minimize $\quad \frac{1}{2}\left\|P-P^{\text {target }}\right\|_{F}^{2}$
subject to minimize $\langle M, P\rangle+\frac{1}{\gamma}\langle P, \log P\rangle$
subject to $\quad P \mathbf{1}=\frac{1}{n} \mathbf{1}$
$P^{T} \mathbf{1}=\frac{1}{m} \mathbf{1}$

with upper-level variable $M \in \mathbb{R}^{m \times n}$.

## Results: Running Time

Running time on cpu with batch size 1


Running time on cuda with batch size 16


## Results: Memory Usage




## Application to Blind Perspective-n-Point


find the location where the photograph was taken

## Coupled Problem



- if we knew correspondences then determining camera pose would be easy
- if we knew camera pose then determining correspondences would be easy


## Blind Perspective-n-Point Network Architecture



## Blind Perspective-n-Point Results



## Further Resources

Where to from here?

- Deep declarative networks (http://deepdeclarativenetworks.com)
- lots of small code examples and tutorials
- CVXPyLayers (https://github.com/cvxgrp/cvxpylayers)
- Theseus (https://sites.google.com/view/theseus-ai)
- JAXopt (https://github.com/google/jaxopt)
lecture notes available at https://users.cecs.anu.edu.au/~sgould


## break-out

## Local Optima are Global Optima Proof mack

## any local minimum of a convex problem is (globally) optimal

Proof. Suppose that $x$ is locally optimal, but there exists a feasible $y$ with lower objective, i.e., $f_{0}(y)<f_{0}(x)$. Local optimality of $x$ means there must be an $R>0$ such that

$$
z \text { feasible and }\|z-x\|_{2} \leq R \Longrightarrow f_{0}(z) \geq f_{0}(x)
$$

Consider $z=\theta y+(1-\theta) x$ with $\theta=\frac{R}{2\|y-x\|_{2}}$. We have that $\|y-x\|_{2}>R$ since we assumed $f_{0}(y)<f_{0}(x)$, so $0<\theta<1 / 2<1$. Therefore $z$ is a convex combination of two feasible points, hence also feasible. Moreover, $\|z-x\|_{2}=R / 2$ (from our choice of $\theta$ ) and therefore $f_{0}(z) \geq f_{0}(x)$ by our assumption that $x$ is locally optimal. But

$$
\begin{aligned}
f_{0}(z) & \leq \theta f_{0}(y)+(1-\theta) f_{0}(x) \\
& <\theta f_{0}(x)+(1-\theta) f_{0}(x) \\
& =f_{0}(x)
\end{aligned}
$$

where the first inequality is by the definition of convex function and the second inequality is from our assumption that $f_{0}(y)<f_{0}(x)$. We have a contradiction. Therefore every locally optimal point is globally optimal.
automatic differentiation

## Toy Example: Babylonian Algorithm

Consider the following implementation for a forward operation:

```
procedure \(\operatorname{FwdFcn}(x)\)
    \(y_{0} \leftarrow \frac{1}{2} x\)
    for \(t=1, \ldots, T\) do
        \(y_{t} \leftarrow \frac{1}{2}\left(y_{t-1}+\frac{x}{y_{t-1}}\right)\)
    end for
    return \(y_{T}\)
end procedure
```


## Toy Example: Babylonian Algorithm

Consider the following implementation for a forward operation:

```
procedure FwdFcn( \(x\) )
    \(y_{0} \leftarrow \frac{1}{2} x\)
    for \(t=1, \ldots, T\) do
        \(y_{t} \leftarrow \frac{1}{2}\left(y_{t-1}+\frac{x}{y_{t-1}}\right)\)
    end for
    return \(y_{T}\)
end procedure
```

Automatic differentiation algorithmically generates the backward code:

```
procedure \(\operatorname{BckFcN}\left(x, y_{T}, \frac{\mathrm{~d} L}{\mathrm{~d} y_{T}}\right)\)
    \(\frac{\mathrm{d} L}{\mathrm{~d} x} \leftarrow 0\)
    for \(t=T, \ldots, 1\) do
        \(\frac{\mathrm{d} L}{\mathrm{~d} x} \leftarrow \frac{\mathrm{~d} L}{\mathrm{~d} x}+\frac{\mathrm{d} L}{\mathrm{~d} y_{t}} \overbrace{\left(\frac{1}{2 y_{t-1}}\right)}^{\partial y_{t} / \partial x}\)
        \(\frac{\mathrm{d} L}{\mathrm{~d} y_{t-1}} \leftarrow \frac{\mathrm{~d} L}{\mathrm{~d} y y_{t}} \underbrace{\left(\frac{1}{2}-\frac{x}{2 y_{t-1}^{2}}\right)}_{\partial y_{t} / \partial y_{t-1}}\)
    end for
    \(\frac{\mathrm{d} L}{\mathrm{~d} x} \leftarrow \frac{\mathrm{~d} L}{\mathrm{~d} x}+\frac{\mathrm{d} L}{\mathrm{~d} y_{0}} \frac{1}{2}\)
    return \(\frac{\mathrm{d} L}{\mathrm{~d} x}\)
end procedure
```


## Toy Example: Babylonian Algorithm mack

Consider the following implementation for a forward operation:

```
procedure FwdFcn(x)
    y0}\leftarrow\frac{1}{2}
    for }t=1,\ldots,T\mathrm{ do
        yt}\leftarrow\frac{1}{2}(\mp@subsup{y}{t-1}{}+\frac{x}{\mp@subsup{y}{t-1}{}}
    end for
    return }\mp@subsup{y}{T}{
end procedure
```

- computes $y=\sqrt{x}$
- derivative computed directly is

$$
\frac{\mathrm{d} y}{\mathrm{~d} x}=\frac{1}{2 \sqrt{x}}=\frac{1}{2 y}
$$

Automatic differentiation algorithmically generates the backward code:

```
procedure \(\operatorname{BCKFCN}\left(x, y_{T}, \frac{\mathrm{~d} L}{\mathrm{~d} y_{T}}\right)\)
    \(\frac{\mathrm{d} L}{\mathrm{~d} x} \leftarrow 0\)
    for \(t=T, \ldots, 1\) do
        \(\frac{\mathrm{d} L}{\mathrm{~d} x} \leftarrow \frac{\mathrm{~d} L}{\mathrm{~d} x}+\frac{\mathrm{d} L}{\mathrm{~d} y_{t}} \overbrace{\left(\frac{1}{2 y_{t-1}}\right)}^{\partial y_{t} / \partial x}\)
        \(\frac{\mathrm{d} L}{\mathrm{~d} y_{t-1}} \leftarrow \frac{\mathrm{~d} L}{\mathrm{~d} y_{t}} \underbrace{\left(\frac{1}{2}-\frac{x}{2 y_{t-1}^{2}}\right)}_{\partial y_{t} / \partial y_{t-1}}\)
    end for
    \(\frac{\mathrm{d} L}{\mathrm{~d} x} \leftarrow \frac{\mathrm{~d} L}{\mathrm{~d} x}+\frac{\mathrm{d} L}{\mathrm{~d} y_{0}} \frac{1}{2}\)
    return \(\frac{\mathrm{d} L}{\mathrm{~d} x}\)
end procedure
```


## Computation Graph for Babylonian Algorithm



## duality

## Lagrange Dual Function

Define Lagrange dual function, $g: \mathbb{R}^{p} \times \mathbb{R}^{q} \rightarrow \mathbb{R}$, as

$$
\begin{aligned}
g(\lambda, \nu) & =\inf _{x \in \mathcal{D}} \mathcal{L}(x, \lambda, \nu) \\
& =\inf _{x \in \mathcal{D}}\left(f_{0}(x)+\sum_{i=1}^{p} \lambda_{i} f_{i}(x)+\sum_{i=1}^{q} \nu_{i} h_{i}(x)\right)
\end{aligned}
$$

- $g$ is concave (always), can be $-\infty$ for some $\lambda, \nu$
- lower bound property: if $\lambda \succeq 0$, then $g(\lambda, \nu) \leq p^{\star}$ (since for feasible $x$ we have $f_{i}(x) \leq 0$ and $h_{i}(x)=0$ )


## The Dual Problem

The Lagrange dual problem is to maximise the dual function

$$
\begin{array}{ll}
\text { maximize } & g(\lambda, \nu) \\
\text { subject to } & \lambda \succeq 0
\end{array}
$$

- finds the best lower bound on $p^{\star}$, obtained from Lagrange dual function
- a convex optimisation problem with optimal value denoted by $d^{\star}$
- $\lambda, \nu$ are dual feasible if $\lambda \succeq 0$ and $(\lambda, \nu) \in \operatorname{dom}(g)$
- original problem is known as the primal problem


## Weak and Strong Duality © back

weak duality: $d^{\star} \leq p^{\star}$

- always holds (for convex and nonconvex problems)
- can be used to find nontrivial lower bounds for difficult problems
strong duality: $d^{\star}=p^{\star}$
- does not hold in general
- (usually) holds for convex problems
- conditions that guarantee strong duality on convex problems are called constraint qualifications
differentiating equality constrained problems


## Abridged Derivation

Forming the Lagrangian at optimal $y$ for fixed $x$ we have

$$
\mathcal{L}(x, y, \nu)=f(x, y)-\sum_{i=1}^{q} \nu_{i} h_{i}(x, y) .
$$

Since $\frac{\partial h(x, y)}{\partial y}$ is full rank we have that $y$ is a regular point. Then there exists a $\nu$ such that the Lagrangian is stationary at the point $(y, \nu)$. Thus

$$
\left[\begin{array}{c}
\frac{\partial \mathcal{L}^{T}}{\partial Y} \\
\frac{\partial \mathcal{L}}{\partial \nu}
\end{array}\right]=\left[\begin{array}{c}
\left(\frac{\partial f(x, y)}{\partial y}-\sum_{i=1}^{q} \nu_{i} \frac{\partial h_{i}(x, y)}{\partial y}\right)^{T} \\
h(x, y)
\end{array}\right]=\mathbf{0}_{m+q}
$$

which we can differentiate with respect to $x$,

$$
\frac{\mathrm{d}}{\mathrm{~d} x}\left[\left(\frac{\partial f(x, y)}{\partial y}\right)^{T}-\sum_{h(x, 1}^{q} \nu_{i}\left(\frac{\partial h_{i}(x, y)}{\partial y}\right)^{T}\right]=\mathbf{0}_{(m+q) \times n}
$$

to get (after some re-arranging in matrix form)

$$
\left[\begin{array}{cc}
\frac{\partial^{2} f(x, y)}{\partial y^{2}}-\sum_{i=1}^{q} \nu_{i} \frac{\partial^{2} h_{i}(x, y)}{\partial y^{2}} & -\left(\frac{\partial h(x, y)}{\partial y}\right)^{T} \\
\frac{\partial h(x, y)}{\partial y} & \mathbf{0}_{q \times q}
\end{array}\right]\left[\begin{array}{c}
\frac{\mathrm{d} y(x)}{\mathrm{d} x} \\
\frac{\mathrm{~d} \nu(x)}{\mathrm{d} x}
\end{array}\right]=-\left[\begin{array}{c}
\frac{\partial^{2} f(x, y)}{\partial x \partial y}-\sum_{i=1}^{q} \nu_{i} \frac{\partial^{2} h_{i}(x, y)}{\partial x \partial y} \\
\frac{\partial}{\partial x} h(x, y)
\end{array}\right] .
$$

## Abridged Derivation

Forming the Lagrangian at optimal $y$ for fixed $x$ we have

$$
\mathcal{L}(x, y, \nu)=f(x, y)-\sum_{i=1}^{q} \nu_{i} h_{i}(x, y) .
$$

Since $\frac{\partial h(x, y)}{\partial y}$ is full rank we have that $y$ is a regular point. Then there exists a $\nu$ such that the Lagrangian is stationary at the point $(y, \nu)$. Thus

$$
\left[\begin{array}{c}
\frac{\partial \mathcal{L}}{}^{T} \\
\frac{\partial \mathcal{L}}{\partial \nu}^{T}
\end{array}\right]=\left[\begin{array}{c}
\left(\frac{\partial f(x, y)}{\partial y}-\sum_{i=1}^{q} \nu_{i} \frac{\partial h_{i}(x, y)}{\partial y}\right)^{T} \\
h(x, y)
\end{array}\right]=\mathbf{0}_{m+q}
$$

which we can differentiate with respect to $x$,

$$
\frac{\mathrm{d}}{\mathrm{~d} x}\left[\left(\frac{\partial f(x, y)}{\partial y}\right)^{T}-\sum_{h(x=1}^{q} \nu_{i}\left(\frac{\partial h_{i}(x, y)}{\partial y}\right)^{T}\right]=\mathbf{0}_{(m+q) \times n}
$$

to get (after some re-arranging in matrix form)

$$
\left[\begin{array}{cc}
H & -A^{T} \\
A & \mathbf{0}_{q \times q}
\end{array}\right]\left[\begin{array}{c}
\frac{\mathrm{d} y(x)}{\mathrm{d} x} \\
\frac{\mathrm{~d} \nu(x)}{\mathrm{d} x}
\end{array}\right]=-\left[\begin{array}{l}
B \\
C
\end{array}\right] .
$$

## Abridged Derivation (cont.)

(from last slide:)

$$
\left[\begin{array}{cc}
H & -A^{T} \\
A & \mathbf{0}_{q \times q}
\end{array}\right]\left[\begin{array}{l}
\frac{\mathrm{d} y(x)}{\mathrm{d} x} \\
\frac{\mathrm{~d} \nu(x)}{\mathrm{d} x}
\end{array}\right]=-\left[\begin{array}{l}
B \\
C
\end{array}\right]
$$

We can solve this system of equations directly or solve by variable elimination. Multiplying out we have

$$
\begin{align*}
H \frac{\mathrm{~d} y(x)}{\mathrm{d} x}-A^{T} \frac{\mathrm{~d} \nu(x)}{\mathrm{d} x} & =-B  \tag{1}\\
A \frac{\mathrm{~d} y(x)}{\mathrm{d} x} & =-C \tag{2}
\end{align*}
$$

Substituting $\frac{d y(x)}{d x}$ from (1) into (2) gives,

$$
\begin{aligned}
A H^{-1}\left(A^{T} \frac{\mathrm{~d} \nu(x)}{\mathrm{d} x}-B\right) & =-C \\
\therefore \frac{\mathrm{~d} \nu(x)}{\mathrm{d} x} & =\left(A H^{-1} A^{T}\right)^{-1}\left(A H^{-1} B-C\right)
\end{aligned}
$$

Then substituting back into (1) we get the result

$$
\frac{\mathrm{d} y(x)}{\mathrm{d} x}=H^{-1} A^{T}\left(A H^{-1} A^{T}\right)^{-1}\left(A H^{-1} B-C\right)-H^{-1} B
$$

## least squares

## Least Squares Backward Pass Derivation © back

Differentiating $x^{\star}$ with respect to single element $A_{i j}$, we have

$$
\begin{aligned}
\frac{\mathrm{d}}{\mathrm{~d} A_{i j}} x^{\star} & =\frac{\mathrm{d}}{\mathrm{~d} A_{i j}}\left(A^{T} A\right)^{-1} A^{T} b \\
& =\left(\frac{\mathrm{d}}{\mathrm{~d} A_{i j}}\left(A^{T} A\right)^{-1}\right) A^{T} b+\left(A^{T} A\right)^{-1}\left(\frac{\mathrm{~d}}{\mathrm{~d} A_{i j}} A^{T} b\right)
\end{aligned}
$$

Using the identity $\frac{\mathrm{d}}{\mathrm{d} z} Z^{-1}=-Z^{-1}\left(\frac{\mathrm{~d}}{\mathrm{~d} z} Z\right) Z^{-1}$ we get, for the first term,

$$
\begin{aligned}
\frac{\mathrm{d}}{\mathrm{~d} A_{i j}}\left(A^{T} A\right)^{-1} & =-\left(A^{T} A\right)^{-1}\left(\frac{\mathrm{~d}}{\mathrm{~d} A_{i j}}\left(A^{T} A\right)\right)\left(A^{T} A\right)^{-1} \\
& =-\left(A^{T} A\right)^{-1}\left(E_{i j}^{T} A+A^{T} E_{i j}\right)\left(A^{T} A\right)^{-1}
\end{aligned}
$$

where $E_{i j}$ is a matrix with one in the $(i, j)$-th element and zeros elsewhere. Furthermore, for the second term,

$$
\frac{\mathrm{d}}{\mathrm{~d} A_{i j}} A^{T} b=E_{i j}^{T} b
$$

## Least Squares Backward Pass Derivation (cont.)

Plugging these back into parent equation we have

$$
\begin{aligned}
\frac{\mathrm{d}}{\mathrm{~d} A_{i j}} x^{\star} & =-\left(A^{T} A\right)^{-1}\left(E_{i j}^{T} A+A^{T} E_{i j}\right)\left(A^{T} A\right)^{-1} A^{T} b+\left(A^{T} A\right)^{-1} E_{i j}^{T} b \\
& =-\left(A^{T} A\right)^{-1}\left(E_{i j}^{T} A+A^{T} E_{i j}\right) x^{\star}+\left(A^{T} A\right)^{-1} E_{i j}^{T} b \\
& =-\left(A^{T} A\right)^{-1}\left(E_{i j}^{T}\left(A x^{\star}-b\right)+A^{T} E_{i j} x^{\star}\right) \\
& =-\left(A^{T} A\right)^{-1}\left(\left(a_{i}^{T} x^{\star}-b_{i}\right) e_{j}+x_{j}^{\star} a_{i}\right)
\end{aligned}
$$

where $e_{j}=(0,0, \ldots, 1,0, \ldots) \in \mathbb{R}^{n}$ is the $j$-th canonical vector, i.e., vector with a one in the $j$-th component and zeros everywhere else, and $a_{i}^{T} \in \mathbb{R}^{1 \times n}$ is the $i$-th row of matrix $A$.

## Least Squares Backward Pass Derivation (cont.)

Let $r=b-A x^{\star}$ and let $v^{T}$ denote the backward coming gradient $\frac{\mathrm{d}}{\mathrm{d} x^{\star}} L$. Then

$$
\begin{aligned}
\frac{\mathrm{d} L}{\mathrm{~d} A_{i j}} & =v^{T} \frac{\mathrm{~d} x^{\star}}{\mathrm{d} A_{i j}} \\
& =v^{T}\left(A^{T} A\right)^{-1}\left(r_{i} e_{j}-x_{j}^{\star} a_{i}\right) \\
& =w^{T}\left(r_{i} e_{j}-x_{j}^{\star} a_{i}\right) \\
& =r_{i} w_{j}-w^{T} a_{i} x_{j}^{\star}
\end{aligned}
$$

where $w=\left(A^{T} A\right)^{-1} v$. We can compute the entire matrix of $m \times n$ derivatives efficiently as the sum of outer products

$$
\left(\frac{\mathrm{d} L}{\mathrm{~d} A}\right)^{T}=\left[\frac{\mathrm{d} L}{\mathrm{~d} A_{i j}}\right]_{\substack{i=1, \ldots, m \\ j=1, \ldots, n}}=w r^{T}-x^{\star}(A w)^{T}
$$

optimal transport

## Objective and Constraint Functions

$$
f(M, P)=\sum_{i=1}^{m} \sum_{j=1}^{n} M_{i j} P_{i j}+\frac{1}{\gamma} \sum_{i=1}^{m} \sum_{j=1}^{n} P_{i j} \log P_{i j}
$$

$$
h(M, P)=\left[\begin{array}{cccc}
\mathbf{1}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right]\left[\begin{array}{c}
P_{11} \\
P_{12} \\
\vdots \\
P_{1 n} \\
P_{21} \\
\vdots \\
P_{m n}
\end{array}\right]-\left[\begin{array}{c}
r_{1} \\
r_{2} \\
\vdots \\
r_{m} \\
c_{1} \\
\vdots \\
c_{n}
\end{array}\right]
$$

(one constraint is redundant-a linear combination of the others-and removed to ensure $\operatorname{rank}(A)=q$ )

## Deriving the Gradient © back

$$
f(M, P)=\sum_{i=1}^{m} \sum_{j=1}^{n} M_{i j} P_{i j}+\frac{1}{\gamma} \sum_{i=1}^{m} \sum_{j=1}^{n} P_{i j} \log P_{i j} \quad h(M, P)=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right] \vec{P}-\left[\begin{array}{c}
r_{2} \\
\vdots \\
r_{m} \\
c
\end{array}\right]
$$

## Deriving the Gradient © back

$$
\begin{gathered}
f(M, P)=\sum_{i=1}^{m} \sum_{j=1}^{n} M_{i j} P_{i j}+\frac{1}{\gamma} \sum_{i=1}^{m} \sum_{j=1}^{n} P_{i j} \log P_{i j} \quad h(M, P)=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \cdots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \cdots & I_{n \times n}
\end{array}\right] \vec{P}-\left[\begin{array}{c}
r_{2} \\
\vdots \\
r_{m} \\
c
\end{array}\right] \\
\frac{\mathrm{d} P}{\mathrm{~d} M}=\left(H^{-1} A^{T}\left(A H^{-1} A^{T}\right)^{-1} A H^{-1}-H^{-1}\right) B \\
A=\frac{\mathrm{d}}{\mathrm{~d} P} h \in \mathbb{R}^{(m+n-1) \times m n} \quad B=\frac{\mathrm{d}^{2}}{\mathrm{~d} M \partial P} f \in \mathbb{R}^{m n \times n n} \quad H=\frac{\mathrm{d}^{2}}{\mathrm{~d} P^{2}} f \in \mathbb{R}^{m n \times m n}
\end{gathered}
$$

## Deriving the Gradient © back

$$
\begin{aligned}
& f(M, P)=\sum_{i=1}^{m} \sum_{j=1}^{n} M_{i j} P_{i j}+\frac{1}{\gamma} \sum_{i=1}^{m} \sum_{j=1}^{n} P_{i j} \log P_{i j} \quad h(M, P)=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right] \vec{P}-\left[\begin{array}{c}
r_{2} \\
\vdots \\
r_{m} \\
c
\end{array}\right] \\
& \frac{\mathrm{d} P}{\mathrm{~d} M}=\left(H^{-1} A^{T}\left(A H^{-1} A^{T}\right)^{-1} A H^{-1}-H^{-1}\right) B \\
& A=\frac{\mathrm{d}}{\mathrm{~d} P} h \in \mathbb{R}^{(m+n-1) \times m n} \\
&=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right] \quad B=\frac{\mathrm{d}^{2}}{\mathrm{dM} \mathrm{\partial P}} f \in \mathbb{R}^{m n \times n n} \quad H=\frac{\mathrm{d}^{2}}{\mathrm{~d} P^{2}} f \in \mathbb{R}^{m n \times m n}
\end{aligned}
$$

## Deriving the Gradient © back

$$
\begin{aligned}
& f(M, P)=\sum_{i=1}^{m} \sum_{j=1}^{n} M_{i j} P_{i j}+\frac{1}{\gamma} \sum_{i=1}^{m} \sum_{j=1}^{n} P_{i j} \log P_{i j} \quad h(M, P)=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right] \vec{P}-\left[\begin{array}{c}
r_{2} \\
\vdots \\
r_{m} \\
c
\end{array}\right] \\
& \frac{\mathrm{d} P}{\mathrm{~d} M}=\left(H^{-1} A^{T}\left(A H^{-1} A^{T}\right)^{-1} A H^{-1}-H^{-1}\right) B \\
& A=\frac{\mathrm{d}}{\mathrm{~d} P} h \in \mathbb{R}^{(m+n-1) \times m n} \\
& =\left[\begin{array}{cccc}
\frac{\mathrm{d} P}{\mathbf{0}_{n}^{T}} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right] \\
& B=\frac{\mathrm{d}^{2}}{\mathrm{~d} M \partial P} f \in \mathbb{R}^{m n \times n n} \quad H=\frac{\mathrm{d}^{2}}{\mathrm{dP} P^{2}} f \\
& B_{i j, k l}= \begin{cases}1 & \text { if } i j=k l \\
0 & \text { otherwise }\end{cases}
\end{aligned}
$$

## Deriving the Gradient © back

$$
\begin{aligned}
& f(M, P)=\sum_{i=1}^{m} \sum_{j=1}^{n} M_{i j} P_{i j}+\frac{1}{\gamma} \sum_{i=1}^{m} \sum_{j=1}^{n} P_{i j} \log P_{i j} \quad h(M, P)=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right] \vec{P}-\left[\begin{array}{c}
r_{2} \\
\vdots \\
r_{m} \\
c
\end{array}\right] \\
& \frac{\mathrm{d} P}{\mathrm{~d} M}=\left(H^{-1} A^{T}\left(A H^{-1} A^{T}\right)^{-1} A H^{-1}-H^{-1}\right) B \\
& A=\frac{\mathrm{d}}{\mathrm{~d} P} h \in \mathbb{R}^{(m+n-1)} \mathbf{( m ) m n} \\
&=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right] \quad B=\frac{\mathrm{d}^{2}}{\mathrm{dM} \mathrm{\partial P}} f \in \mathbb{R}^{m n \times n n} \quad H=\frac{\mathrm{d}^{2}}{\mathrm{~d} P^{2}} f \in \mathbb{R}^{m n \times m n}
\end{aligned}
$$

## Deriving the Gradient © back

$$
\begin{gathered}
f(M, P)=\sum_{i=1}^{m} \sum_{j=1}^{n} M_{i j} P_{i j}+\frac{1}{\gamma} \sum_{i=1}^{m} \sum_{j=1}^{n} P_{i j} \log P_{i j} \quad h(M, P)=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \cdots & I_{n \times n}
\end{array}\right] \vec{P}-\left[\begin{array}{c}
r_{2} \\
\vdots \\
r_{m} \\
c
\end{array}\right] \\
\frac{\mathrm{d} P}{\mathrm{~d} M}=\left(H^{-1} A^{T}\left(A H^{-1} A^{T}\right)^{-1} A H^{-1}-H^{-1}\right) B \\
A=\frac{\mathrm{d}}{\mathrm{~d} P} h \in \mathbb{R}^{(m+n-1) \times m n} \begin{array}{l}
\mathrm{d} P
\end{array} \\
=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right] \quad B=\frac{\mathrm{d}^{2}}{\mathrm{dM} \mathrm{\partial P}} f \in \mathbb{R}^{m n \times m n} \quad H=\frac{\mathrm{d}^{2}}{\mathrm{~d} P^{2}} f \in \mathbb{R}^{m n \times m n} \\
I_{m n \times m n}
\end{gathered}
$$

## Deriving the Gradient © back

$$
\begin{aligned}
& f(M, P)=\sum_{i=1}^{m} \sum_{j=1}^{n} M_{i j} P_{i j}+\frac{1}{\gamma} \sum_{i=1}^{m} \sum_{j=1}^{n} P_{i j} \log P_{i j} \quad h(M, P)=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right] \vec{P}-\left[\begin{array}{c}
r_{2} \\
\vdots \\
r_{m} \\
c
\end{array}\right] \\
& \frac{\mathrm{d} P}{\mathrm{~d} M}=\left(H^{-1} A^{T}\left(A H^{-1} A^{T}\right)^{-1} A H^{-1}-H^{-1}\right) B \\
& A=\frac{\mathrm{d}}{\mathrm{~d} P} h \in \mathbb{R}^{(m+n-1) \times m n} \\
& =\left[\begin{array}{cccc}
\mathrm{d} P \\
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T}
\end{array}\right] \quad=I_{m n \times m n} \quad H=\frac{\mathrm{d} P^{2}}{} \mathrm{~d} M \partial P \quad H^{-1}=\gamma \boldsymbol{d i a g}(\vec{P})
\end{aligned}
$$

## Computing $\left(A H^{-1} A^{T}\right)^{-1}$

$$
\begin{gathered}
H^{-1}=\gamma \operatorname{diag}(\vec{P}) \quad A=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \cdots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \cdots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \cdots & I_{n \times n}
\end{array}\right] \\
\\
\frac{\mathrm{d} P}{\mathrm{~d} M}=\left(H^{-1} A^{T}\left(A H^{-1} A^{T}\right)^{-1} A H^{-1}\right. \\
\end{gathered}
$$

## Computing $\left(A H^{-1} A^{T}\right)^{-1}$

$$
\left.\left.\begin{array}{c}
H^{-1}=\gamma \operatorname{diag}(\vec{P}) \quad A=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \cdots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \cdots & I_{n \times n}
\end{array}\right] \\
\\
\frac{\mathrm{d} P}{\mathrm{~d} M}=\left(H^{-1} A^{T}\left(A H^{-1} A^{T}\right)^{-1} A H^{-1}\right.
\end{array}-H^{-1}\right) B\right)
$$

The $(k, l)$-th entry of $A H^{-1} A^{T}$ for $k, l \in 1, \ldots, m+n-1$ is

$$
\left(A H^{-1} A^{T}\right)_{k l}=\sum_{i=1}^{m} \sum_{j=1}^{n} \frac{A_{k, i j} A_{l, i j}}{H_{i j, i j}}=\gamma \sum_{i=1}^{m} \sum_{j=1}^{n} A_{k, i j} A_{l, i j} P_{i j}
$$

Interpreting $A_{k, i j} A_{l, i j}$

$$
{ }_{k}\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right] \quad k\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right]
$$

$$
{ }_{l}\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right] \quad{ }_{l}^{k}\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \ldots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \ldots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \ldots & I_{n \times n}
\end{array}\right]
$$

Evaluating $\left(A H^{-1} A^{T}\right)_{k l}=\gamma \sum_{i=1}^{m} \sum_{j=1}^{n} A_{k, i j} A_{l, i j} P_{i j}$

|  | $0 \leq l \leq m-1$ | $m \leq l \leq m+n-1$ |
| :---: | :---: | :---: |
| $0 \leq k \leq m-1$ | $\begin{cases}\gamma \sum_{j=1}^{n} P_{k+1, j} & \text { if } k=l \\ 0 & \text { otherwise }\end{cases}$ | $\gamma P_{k+1, l-m+1}$ |

## Computing $\left(A H^{-1} A^{T}\right)^{-1}$

$$
\begin{gathered}
H^{-1}=\gamma \operatorname{diag}(\vec{P}) \quad A=\left[\begin{array}{cccc}
\mathbf{0}_{n}^{T} & \mathbf{1}_{n}^{T} & \cdots & \mathbf{0}_{n}^{T} \\
\vdots & \vdots & \ddots & \vdots \\
\mathbf{0}_{n}^{T} & \mathbf{0}_{n}^{T} & \cdots & \mathbf{1}_{n}^{T} \\
I_{n \times n} & I_{n \times n} & \cdots & I_{n \times n}
\end{array}\right] \\
\\
\frac{\mathrm{d} P}{\mathrm{~d} M}=\left(H^{-1} A^{T}\left(A H^{-1} A^{T}\right)^{-1} A H^{-1}\right. \\
\end{gathered}
$$

$$
\begin{gathered}
A H^{-1} A^{T}=\gamma\left[\begin{array}{cc}
\boldsymbol{\operatorname { d i a g }}\left(r_{2: m}\right) & P_{2: m, 1: n} \\
P_{2: m, 1: n}^{T} & \mathbf{d i a g}(c)
\end{array}\right] \quad\left(A H^{-1} A^{T}\right)^{-1}=\frac{1}{\gamma}\left[\begin{array}{ll}
\Lambda_{11} & \Lambda_{12} \\
\Lambda_{12}^{T} & \Lambda_{22}
\end{array}\right] \\
\Lambda_{11}=\left(\operatorname{diag}\left(r_{2: m}-P_{2: m, 1: n} \operatorname{diag}(c)^{-1} P_{2: m, 1: n}^{T}\right)\right)^{-1} \\
\Lambda_{12}=-\Lambda_{11} P_{2: m, 1: n} \operatorname{diag}(c)^{-1} \\
\Lambda_{22}=\operatorname{diag}(c)^{-1}-\operatorname{diag}(c)^{-1} P_{2: m, 1: n}^{T} \Lambda_{12}
\end{gathered}
$$


[^0]:    ${ }^{1}$ In these lectures we will be concerned with continuous-valued variables

[^1]:    ${ }^{1}$ Warning: notation clash between $p$ and $p^{\star}$ !

[^2]:    ${ }^{1}$ Warning: notation clash between $p$ and $p^{\star}$ !

