Master in Computer Vision *Barcelona*

M2: Optimisation and Inference for Computer Vision

Final Project Presentation

Team 1:

- Òscar Lorente Corominas
- Antoni Rodriguez Villegas
- Aditya Sangram Singh Rana

Contents

Optimisation in Computer Vision

Optimization in Computer Vision is about trying to solve complex ill-posed problems where we cannot find analytical solutions and hence we try to formulate it based on some criterions and look for a numerical solution that is acceptable upto some tolerance.

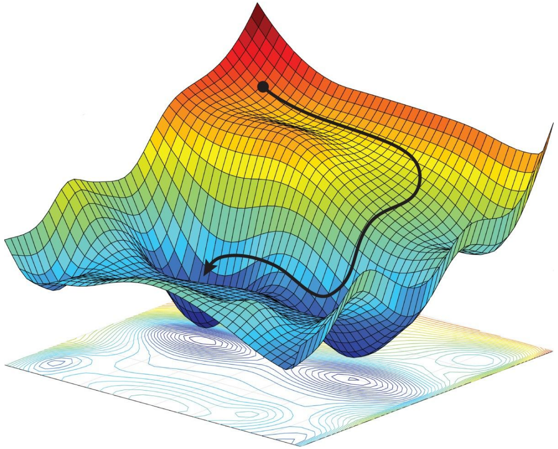

Research Paper

Determining Optical Flow

Horn, Berthold & Schunck, Brian. (1981). Artificial Intelligence. 17. 185-203. 10.1016/0004-3702(81)90024-2.

Link : http://image.diku.dk/imagecanon/material/HornSchunckOptical_Flow.pdf

Optical Flow: Problem

How to estimate pixel motion from Image $I(x, y, t) \rightarrow I(x, y, t+1)$?

Solve The Pixel Correspondence Problem: Given a pixel in $I(x, y, t)$, look for **nearby** pixels of the **same intensity** in $I(x, y, t+1)$.

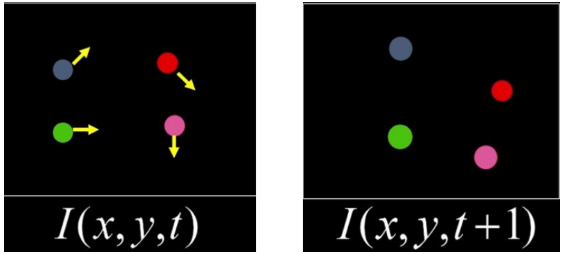

Optical Flow is the **apparent** motion of objects or surfaces.

Criterion	Implication	Visual
Small motion: points do not move very far in consecutive frames, meaning motion is smooth.	Displacement: $(\delta x, \delta y) = (u\delta t, v\delta t)$ Optical flow (velocities): (u, v)	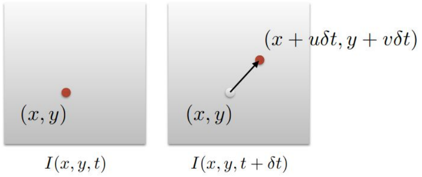
Brightness: brightness of point will remain the same	$I(x(t), y(t), t) = C$ constant	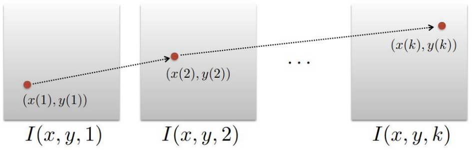

INSIGHT: If the time step is really small, we can linearize the intensity function using Multivariable Taylor Series Expansion (First order approximation with two variables)

SOLVE: For a really small space time-step, I(x+uδt, y+ vδt, t+δt) = I(x, y, t)…. The brightness between two consecutive frames is same

Brightness Constancy Equation

$$\frac{dI}{dt} = \frac{\partial I}{\partial x}\frac{dx}{dt} + \frac{\partial I}{\partial y}\frac{dy}{dt} + \frac{\partial I}{\partial t} = 0$$

total derivative

partial derivative

Optical Flow: Solution

$$\frac{\partial I}{\partial x}\frac{dx}{dt} + \frac{\partial I}{\partial y}\frac{dy}{dt} + \frac{\partial I}{\partial t} = 0$$

Brightness Constancy Equation

$$I_x u + I_y v + I_t = 0$$

shorthand notation

QUESTION: What do the term of the brightness constancy equation represent?

flow velocities

$$I_x u + I_y v + I_t = 0$$

Image gradients / temporal gradient

$$I_x = \frac{\partial I}{\partial x} \quad I_y = \frac{\partial I}{\partial y}$$

spatial derivative

Forward difference
Sobel filter
Scharr filter

QUESTION: How do we compute these terms?

$$I_t = \frac{\partial I}{\partial t}$$

temporal derivative

frame differencing

unknown

$$I_x u + I_y v + I_t = 0$$

known

NOTE: We need at least **2** equations to solve for **2** unknowns.

Where do we get more equations (constraints)?

Horn-Schunck Assumption: most objects in the world are rigid, deform elastically and move together coherently.

Implication: Enforce optical flow fields to be smooth

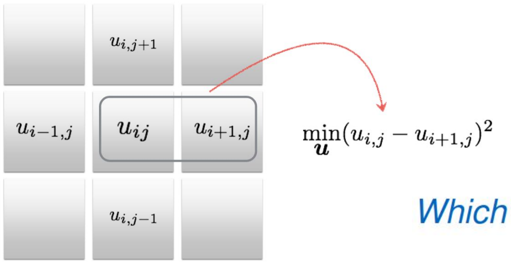

u-component of flow

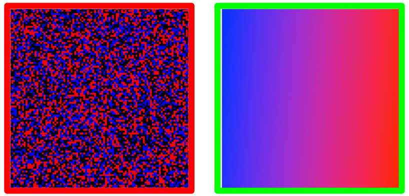

Which flow field optimizes the objective?

$$\min_{u}(u_{i,j} - u_{i+1,j})^2$$

$$\min_{u,v} \sum_{i,j} \left\{ E_s(i,j) + \lambda E_d(i,j) \right\}$$

smoothness / brightness constancy / weight

HOW TO SOLVE?
1. Compute partial derivative,
2 .Derive update equations
3. Use grad. descent

Optical Flow: Results

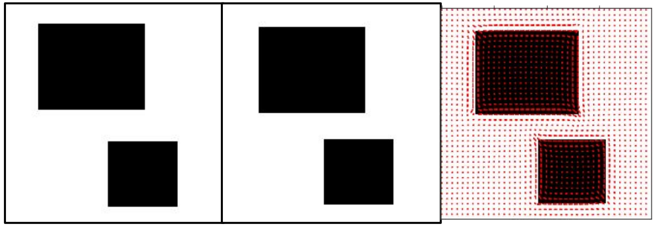

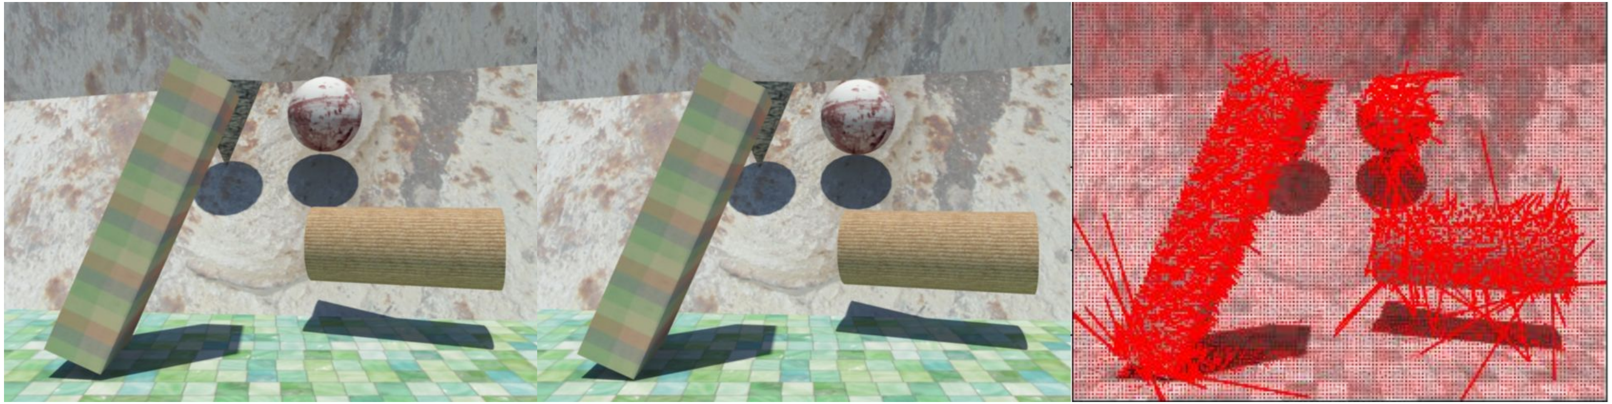

Poisson Editing: Problem

Add a region of a source image B to a target image A while keeping it **photorealistic**.

We **can't just crop the source region and paste it** into the destination image **D**.

Create a new image **H** that **keeps the colors of image A and the details (i.e. gradients) of image B**.

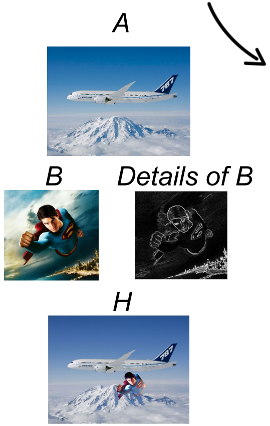

Minimisation problem

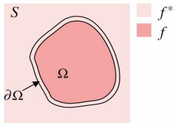

f^*: known image values

f: unknown values over region Ω

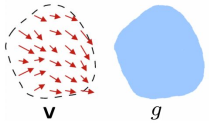

v: guided field

v may be gradient of a function g

Simplest interpolant -> Inpainting

$$\min_f \iint_\Omega |\nabla f|^2 \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$$

Edit would be too blurry!!

Guided interpolation

$$\min_f \iint_\Omega |\nabla f - \mathbf{v}|^2 \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$$

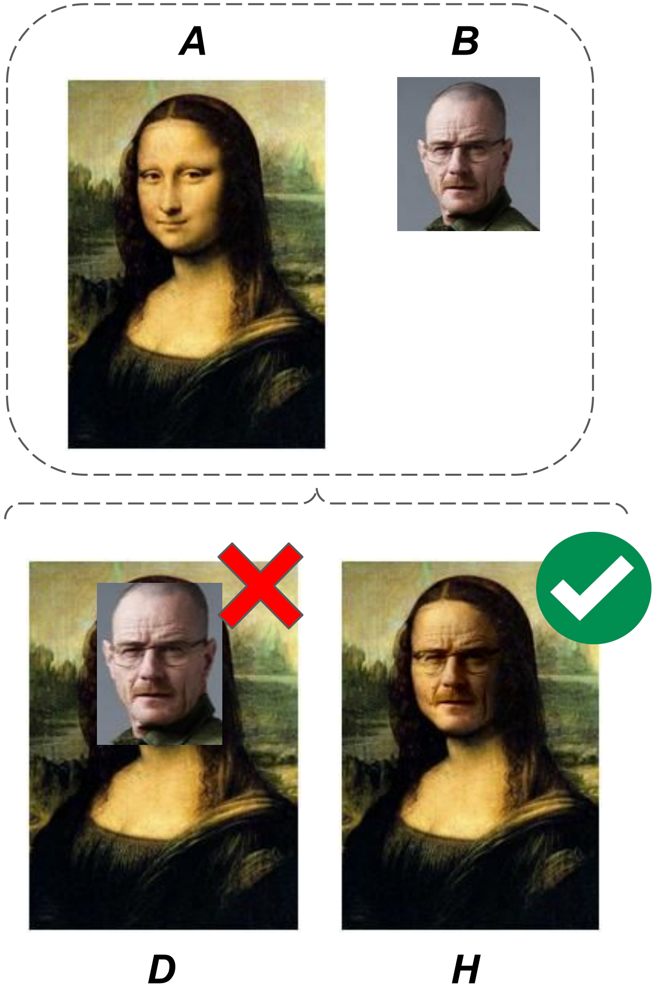

Source image g -> B
Target image f* -> A
Resulting image f -> H

Poisson Editing: Solution

Poisson equation with Dirichlet boundary conditions

$$\Delta f = \operatorname{div} \mathbf{v} \text{ over } \Omega, \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$$

Importing gradients

$\mathbf{v} = \nabla g$

$$\Delta f = \Delta g \text{ over } \Omega, \text{ with } f|_{\partial\Omega} = f^*|_{\partial\Omega}$$

Criterion	Implication
The transition between **A** and **B** must be **smooth**, that is **conserving the colours of image A**	**H(x,y) = A(x,y)** at each (x,y) of the **boundary of B**
The details of B must be kept, that is the **gradients must be conserved**	**∇H(x,y) = ∇B(x,y)** at each (x,y) **inside B** $4 * H_{i,j} - (H_{i+1,j} + H_{i-1,j} + H_{i,j+1} + H_{i,j-1})$ = $4 * B_{i,j} - (B_{i+1,j} + B_{i-1,j} + B_{i,j+1} + B_{i,j-1})$

Mixing gradients

- When **importing gradients**, **no trace of A remains inside the region** we wanted to fill
- There are **situations** where it is **desirable to combine properties of A with those of B**, (e.g. to add objects with holes, or partially transparent ones, on top of a textured or cluttered background)
- With the **mixing gradients** algorithm, we **retain the stronger of the variations in image A or in image B**, using the following guidance field:

$$\text{for all } \mathbf{x} \in \Omega, \ \mathbf{v}(\mathbf{x}) = \begin{cases} \nabla f^*(\mathbf{x}) & \text{if } |\nabla f^*(\mathbf{x})| > |\nabla g(\mathbf{x})|, \\ \nabla g(\mathbf{x}) & \text{otherwise.} \end{cases}$$

$$\mathbf{v(x)} = \begin{cases} \nabla A(x) & \text{if } |\nabla A(x)| > |\nabla B(x)| \\ \nabla B(x) & \text{otherwise} \end{cases}$$

Poisson Editing: Results

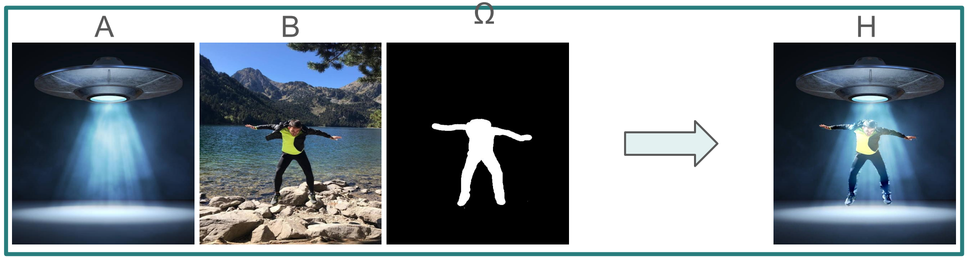

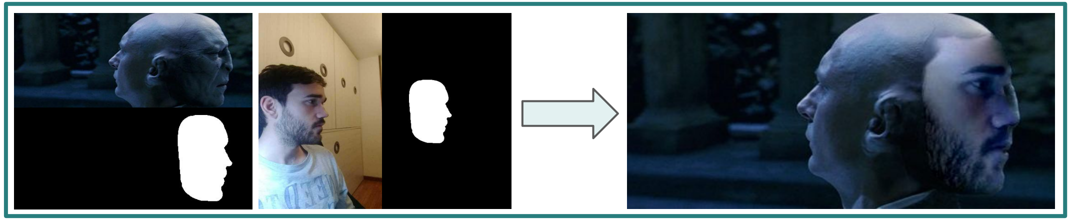

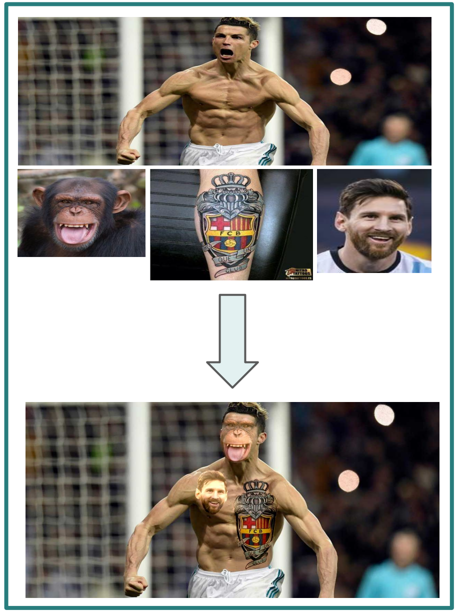

Importing gradients

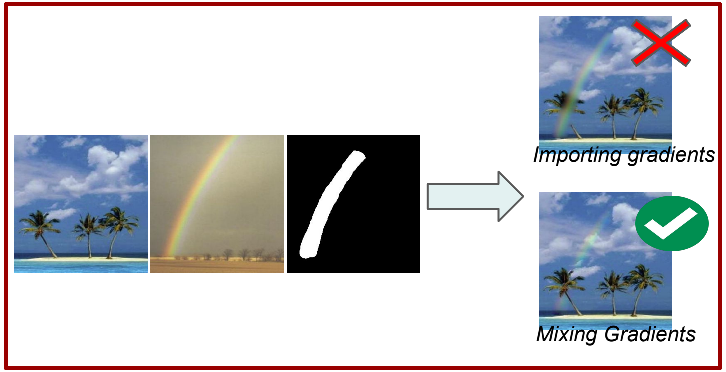

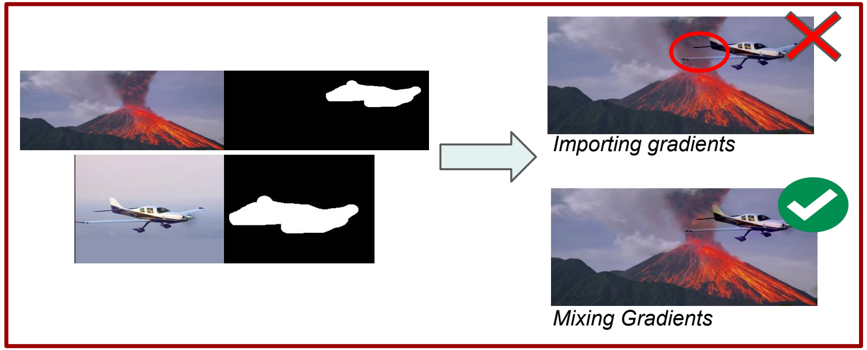

Mixing gradients

Chan-Vese Segmentation: Problem

Segment image into meaningful regions

Criterion	Implication
Segmentation result is an image	Mapping in the same domain $f \to u$
Segmentation image is similar to the original image	Minimise $\int_{\Omega} (f(x) - u(x))^2 \, dx$
Segmented regions have smooth boundaries	Minimise $\arg\min_{u,C} \; \mu \, \mathrm{Length}(C)$
Segmented regions are homogeneous	Minimise $\int_{\Omega \setminus C} \lvert \nabla u(x) \rvert^2 \, dx$
Segmentation image has two regions	Constrain values $u(x) = \{c1, c2\}$ 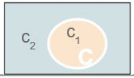

f

$u(x)$

Minimisation problem

$$\arg\min_{c_1,c_2,C} \; \mu \, \mathrm{Length}(C) + \nu \, \mathrm{Area}(inside(C)) + \lambda_1 \int_{inside(C)} \lvert f(x) - c_1 \rvert^2 \, dx + \lambda_2 \int_{outside(C)} \lvert f(x) - c_2 \rvert^2 \, dx$$

Chan-Vese Segmentation: Solution

First approach: **Active contours**

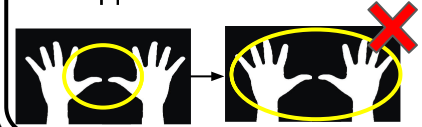

Limitation: Active contours cannot deal with disconnected regions

Solution: **Active surfaces**

$$\underset{c_1,c_2,\varphi}{\arg\min}\ \mu\int_\Omega \delta(\varphi(x))|\nabla\varphi(x)|\,dx + \nu\int_\Omega H(\varphi(x))\,dx$$
$$+\lambda_1\int_\Omega |f(x)-c_1|^2 H(\varphi(x))\,dx + \lambda_2\int_\Omega |f(x)-c_2|^2(1-H(\varphi(x)))\,dx$$

Contours that segment the regions: $C = \{x\in\Omega : \varphi(x)=0\}$

Heaviside step function

$$H(t)=\begin{cases}1 & t\geq 0,\\ 0 & t<0,\end{cases}\qquad \delta(t)=\frac{d}{dt}H(t)$$

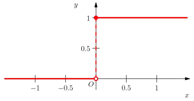

Regularised heaviside step function

$$H_\epsilon(t)=\frac{1}{2}\left(1+\frac{2}{\pi}\arctan\left(\frac{t}{\epsilon}\right)\right)$$

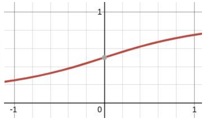

Dirac

$$\delta_\epsilon(t) := \frac{d}{dt}H_\epsilon(t)=\frac{\epsilon}{\pi(\epsilon^2+t^2)}$$

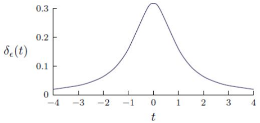

Level set function φ

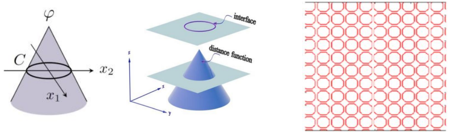

$$\varphi(x)=r-\sqrt{x_1^2+x_2^2}\qquad \varphi(x)=\sin(\tfrac{\pi}{5}x_1)\sin(\tfrac{\pi}{5}y)$$

Boundary conditions

$$\left.\begin{aligned}\varphi_{-1,j}&=\varphi_{0,j},\\ \varphi_{M,j}&=\varphi_{M-1,j},\\ \varphi_{i,-1}&=\varphi_{i,0},\\ \varphi_{i,M}&=\varphi_{i,M-1}\end{aligned}\right\}$$

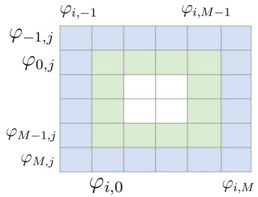

Average of each region

$$c_1=\frac{\int_\Omega f(x)H(\varphi(x))\,dx}{\int_\Omega H(\varphi(x))\,dx}\qquad c_2=\frac{\int_\Omega f(x)(1-H(\varphi(x)))\,dx}{\int_\Omega(1-H(\varphi(x)))\,dx}$$

Coefficients A and B

$$A_{i,j}=\frac{\mu}{\sqrt{\eta^2+(\nabla_x^+\varphi_{i,j})^2+(\nabla_y^0\varphi_{i,j})^2}}\quad B_{i,j}=\frac{\mu}{\sqrt{\eta^2+(\nabla_x^0\varphi_{i,j})^2+(\nabla_y^+\varphi_{i,j})^2}}$$

Gauss-Seidel

$$\begin{aligned}\varphi_{i,j}^{n+1}\leftarrow\Big[&\varphi_{i,j}^n+dt\,\delta_\epsilon(\varphi_{i,j}^n)\big(A_{i,j}\varphi_{i+1,j}^n+A_{i-1,j}\varphi_{i-1,j}^{n+1}+B_{i,j}\varphi_{i,j+1}^n+B_{i,j-1}\varphi_{i,j-1}^{n+1}\\ &-\nu-\lambda_1(f_{i,j}-c_1)^2+\lambda_2(f_{i,j}-c_2)^2\big)\Big]\\ &\Big/\ \big[1+dt\,\delta_\epsilon(\varphi_{i,j})(A_{i,j}+A_{i-1,j}+B_{i,j}+B_{i,j-1})\big].\end{aligned}$$

Chan-Vese Segmentation: Results

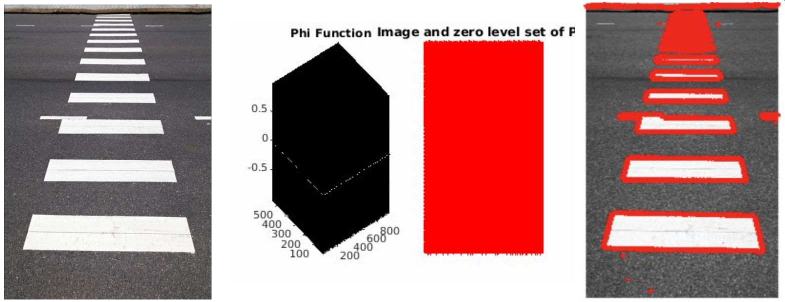

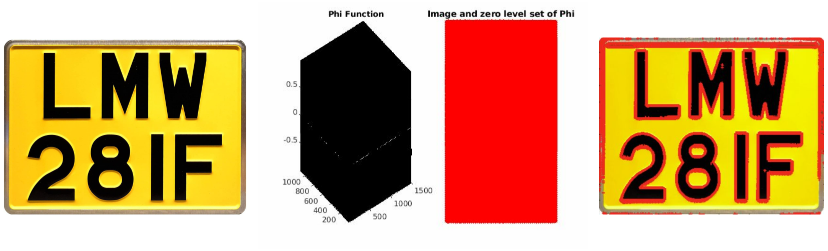

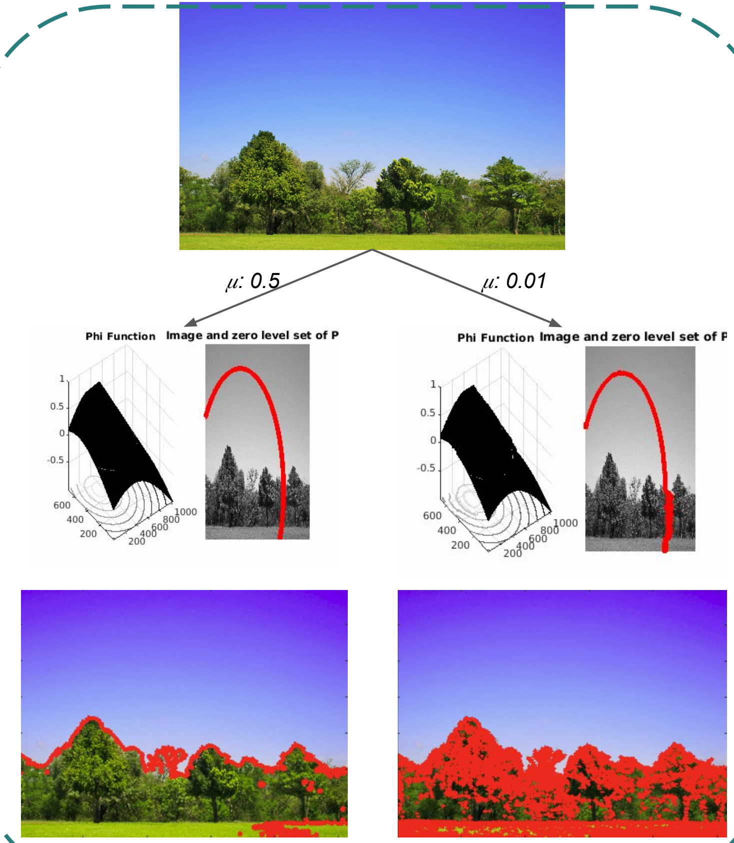

Markov Random Fields: Problem

Homogeneous segmentation of non homogeneous regions in an image

- Applications in Medical image segmentation, where there are intensity inhomogeneities in the regions of interest. Some examples of applications are:
    - Segmentation of polyp ulceration in optical biopsies
    - Breast cancer MRIs with radio-frequency induced noise.

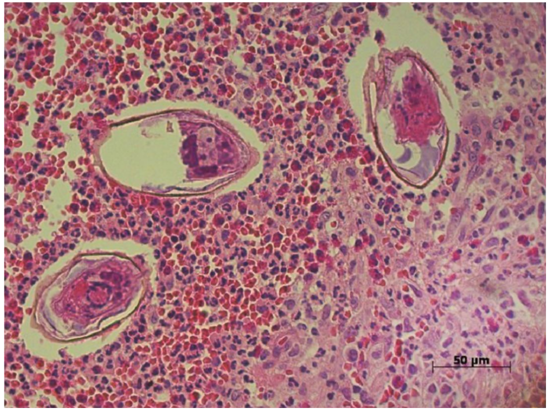

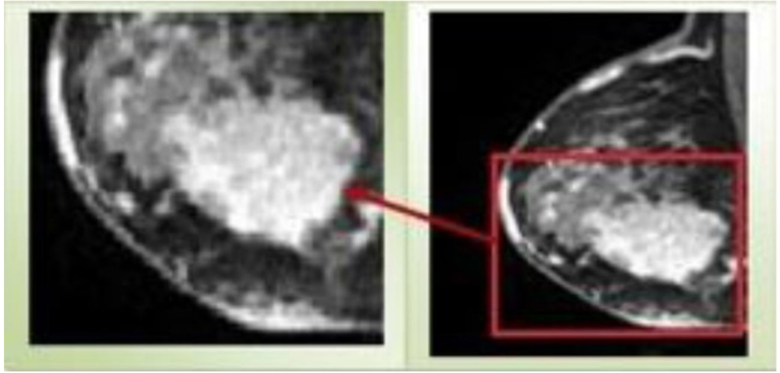

Markov Random Fields: Solution

Probabilistic approach, take region consistency into account in classification of pixels

Criterion	Implication
Each pixel should belong to the class that best corresponds its pixel value.	Maximise $P(y_i \| C_i)$ 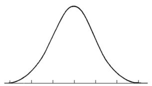
Neighbouring pixels should belong to the same class as often as possible.	Maximise $P(C_i \| C_j)$

Minimise

$$E(x,y) = \sum_i \underbrace{\varphi(x_i, y_i)}_{\text{unary term}} + \sum_{ij} \underbrace{\psi(x_i, x_j)}_{\text{pairwise term}} \quad \alpha$$

Using ICM

Markov Random Fields: Results

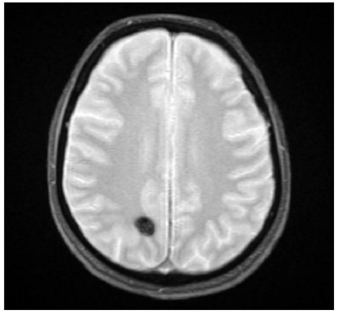

Original Image

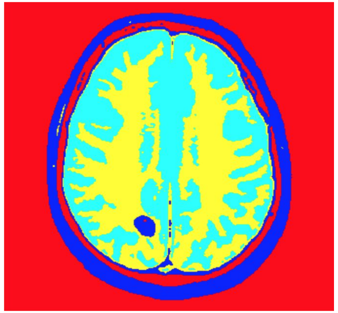

Output image
with alpha=0.1
30 iterations

Original Image

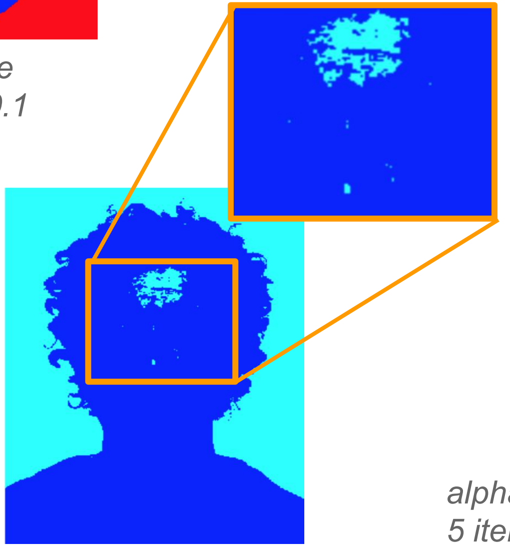

Kmeans initial labeling

alpha=5
5 iterations

Advantages and Limitations

	Advantages	Limitations
Poisson Editing	Not computationally complex. Good results in general cases.	Gradients of target image are lost in some cases if not using mixing gradients.
Chan-Vese Segmentation	Good segmentation of regions not separated by edges.	- Fine tuning of parameters for every specific image. - Only binary segmentation. - Pixels intensities are assumed to be statistically homogenous in each region.
Markov Random Fields	Segmentation of non homogeneous regions.	Results sensitive to initial parameters.

Future Work and Conclusions

Conclusions

- Optimization techniques for computer vision are not overthrown by deep learning.
- Low computational complexity. Good for embedded systems (applications in ADAS).

Future Work

- Work with other images and examples to find more limitations of each algorithm.
- Study of optimization techniques that are specific to medical implementations.

Master in Computer Vision Barcelona

Thank you!

Q&A

Team 1:

- Òscar Lorente Corominas
- Antoni Rodriguez Villegas
- Aditya Sangram Singh Rana