Graph Geodesics to Find Progressively Similar Skin Lesions

Jeremy Kawahara, Kathleen Moriarty, Ghassan Hamarneh

Presented by Kathleen Moriarty

Database of skin lesion images

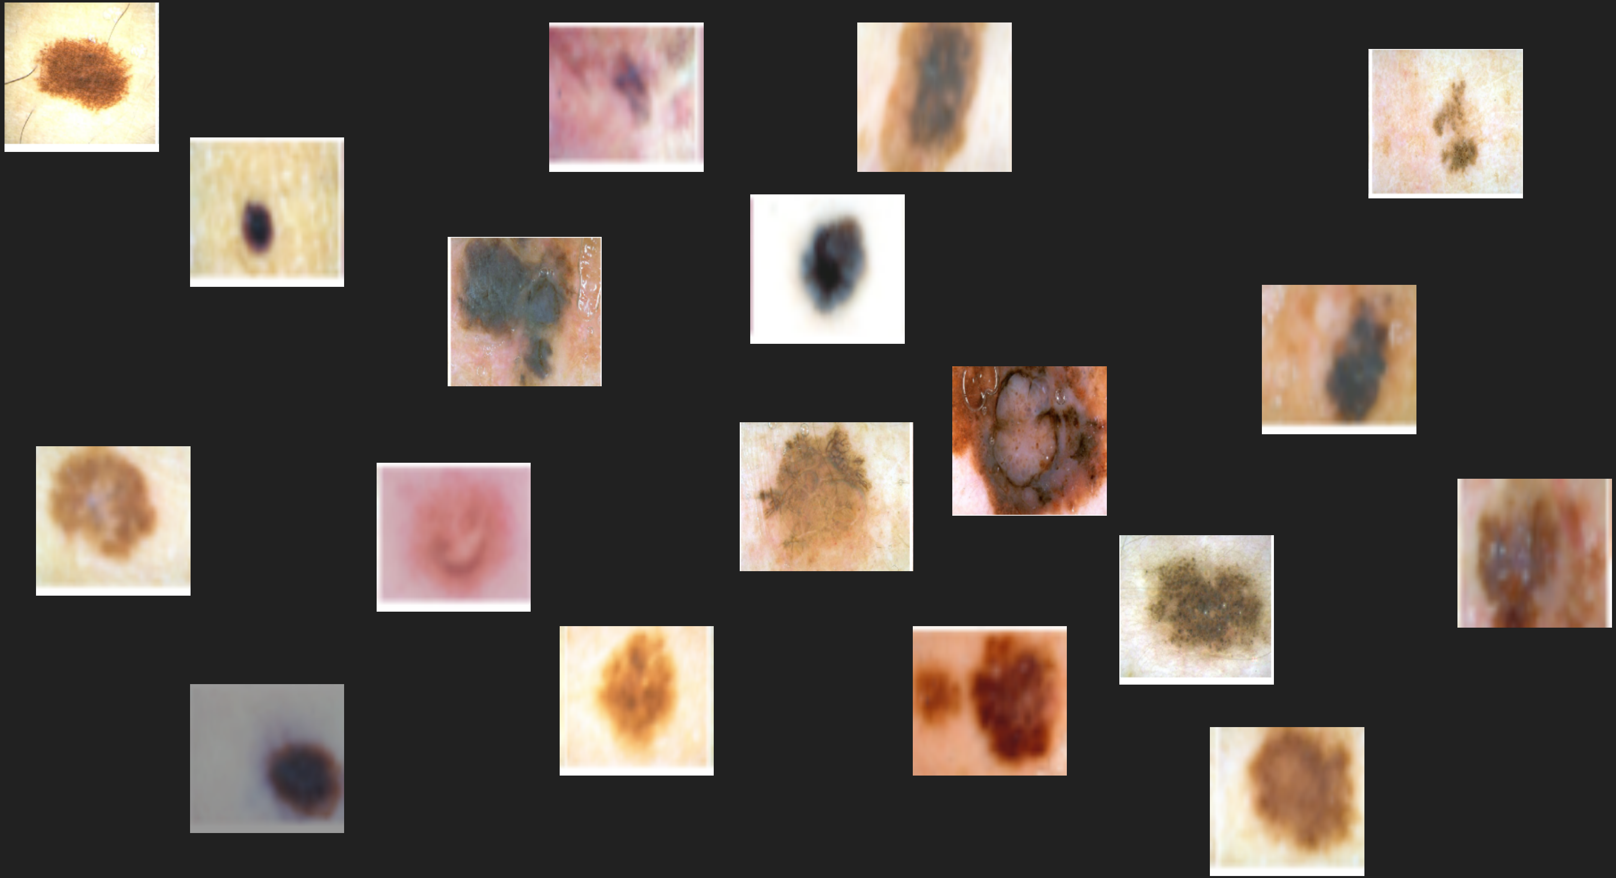

Image retrieval

Query
Image

Query
Image

Image retrieval

Query Image | Most Similar

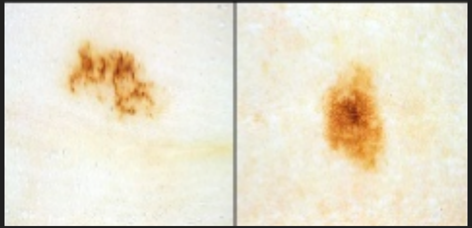

CN

Most Similar | Query Image

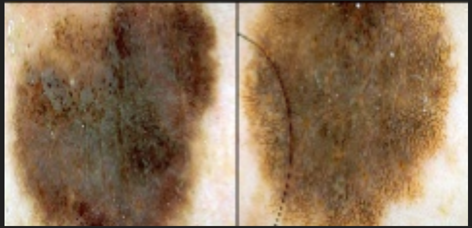

MEL

Find similar images from database of known skin images
Similar images can infer a diagnosis

CN = Clark Nevus (Benign)

MEL = Melanoma (Cancerous)

Image retrieval

Query Image

Most Similar

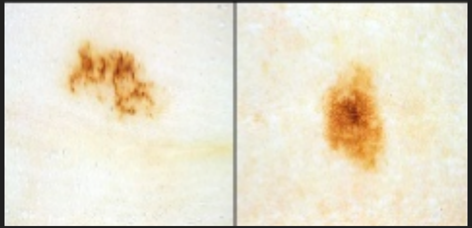

CN

CN

Most Similar

Query Image

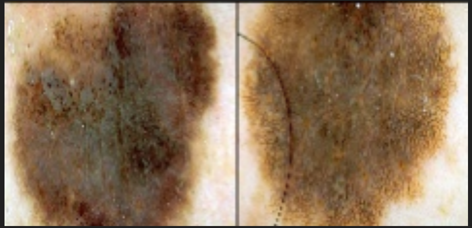

MEL

MEL

Find similar images from database of known skin images
Similar images can infer a diagnosis

CN = Clark Nevus (Benign)

MEL = Melanoma (Cancerous)

Image path retrieval

Query Image | Most Similar

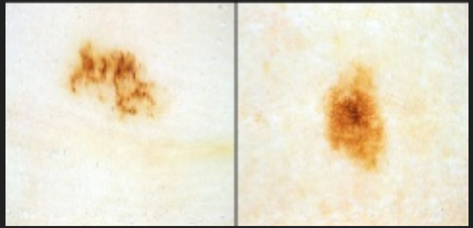

CN | CN

Most Similar | Query Image

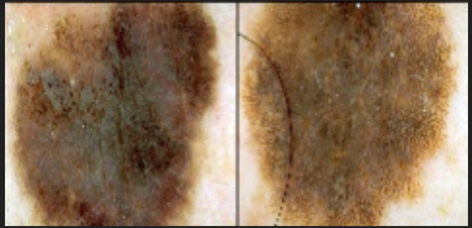

MEL | MEL

CN = Clark Nevus (Benign)

MEL = Melanoma (Cancerous)

Image path retrieval

Source Image

CN

Most Similar

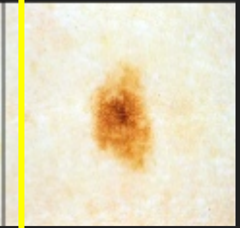

CN

Most Similar

MEL

Query Image

MEL

CN = Clark Nevus (Benign)

MEL = Melanoma (Cancerous)

Image path retrieval

Source Image

CN

Most Similar

CN

Most Similar

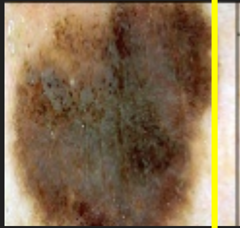

MEL

Target Image

MEL

CN = Clark Nevus (Benign)

MEL = Melanoma (Cancerous)

Image path retrieval

Source Image

Target Image

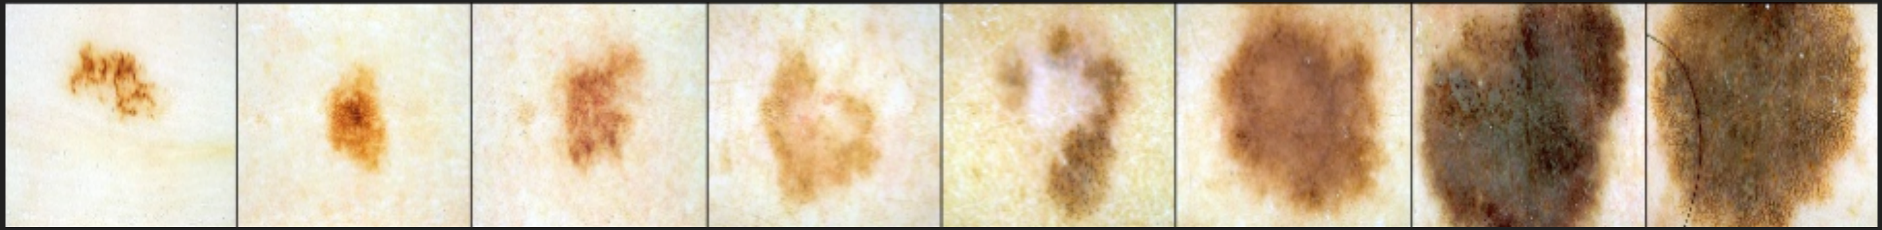

CN = Clark Nevus (Benign)

MEL = Melanoma (Cancerous)

Image path retrieval

Source Image

Target Image

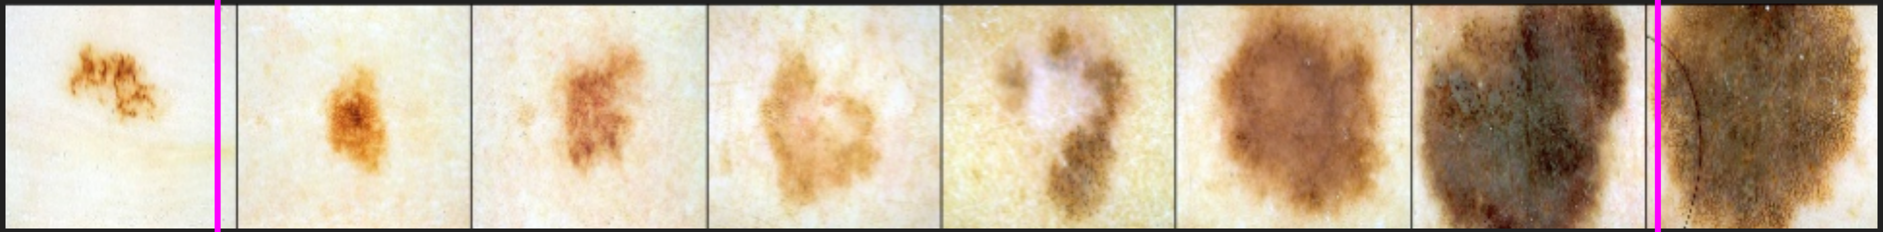

Retrieved path of images

CN = Clark Nevus (Benign)

MEL = Melanoma (Cancerous)

Image path retrieval

Source Image | **Target Image**

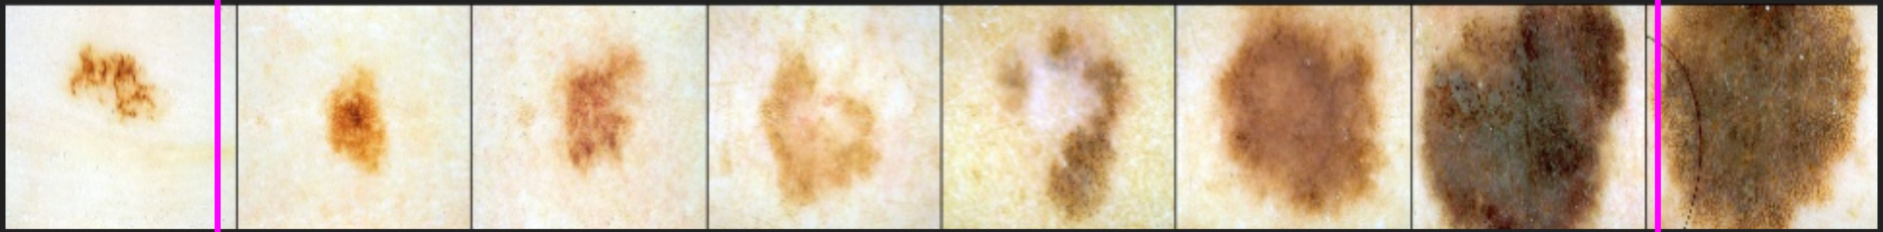

Retrieved path of images

A visual progression from a **source** to a **target** image

CN = Clark Nevus (Benign)

MEL = Melanoma (Cancerous)

Image path retrieval

Source Image

Target Image

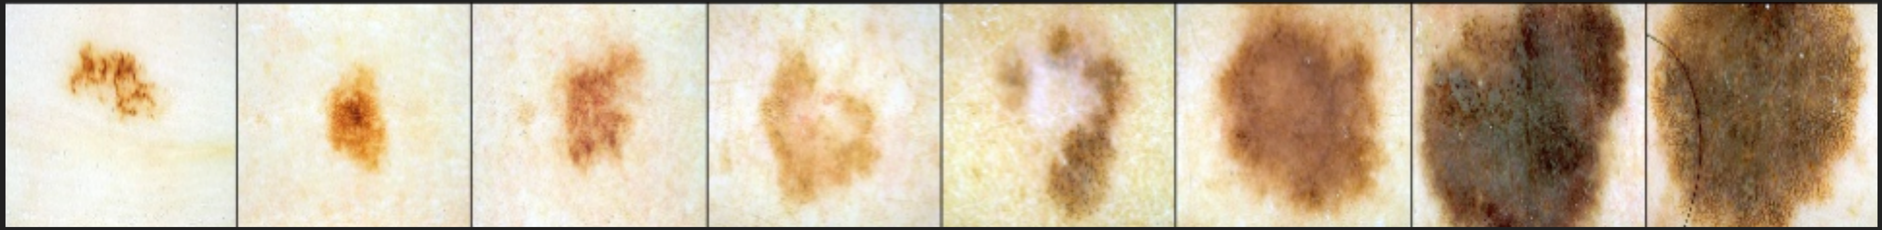

CN CN CN CN MEL CN MEL MEL

Example Applications

- Visualize a progression of diseases (e.g., from benign to melanoma)

CN = Clark Nevus (Benign)

MEL = Melanoma (Cancerous)

Image path retrieval

Source Image

Target Image

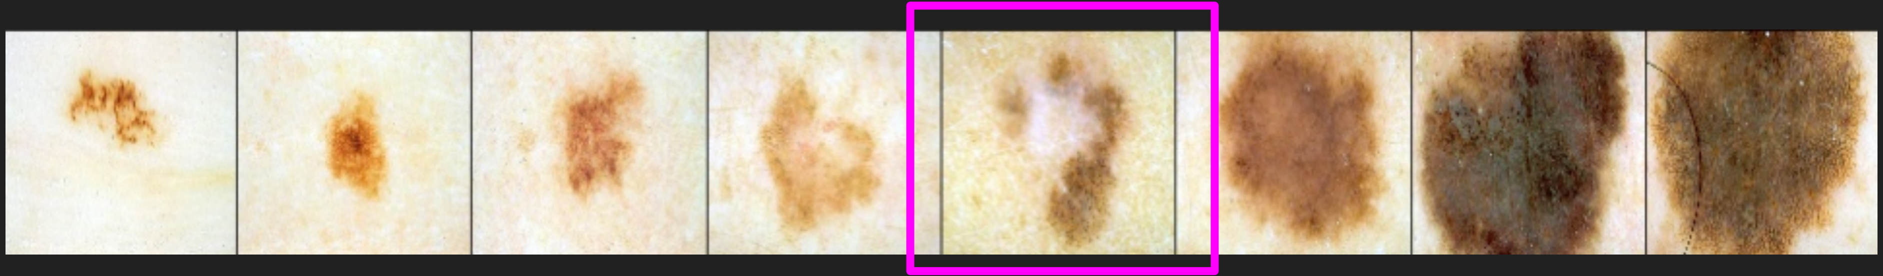

CN CN CN CN MEL CN MEL MEL

Example Applications

- Visualize a progression of diseases (e.g., from benign to melanoma)
- Find images of **borderline cases**

CN = Clark Nevus (Benign)

MEL = Melanoma (Cancerous)

Skin dataset encoded as a graph

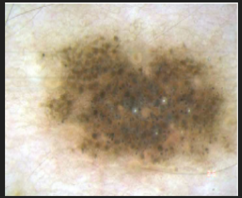

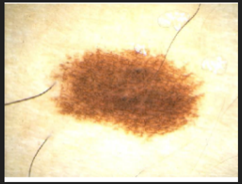

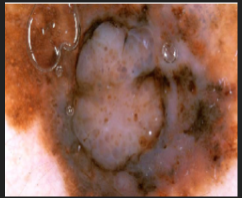

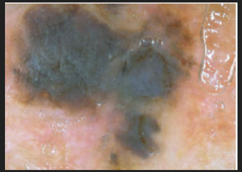

Skin dataset encoded as a graph

Images are encoded as nodes

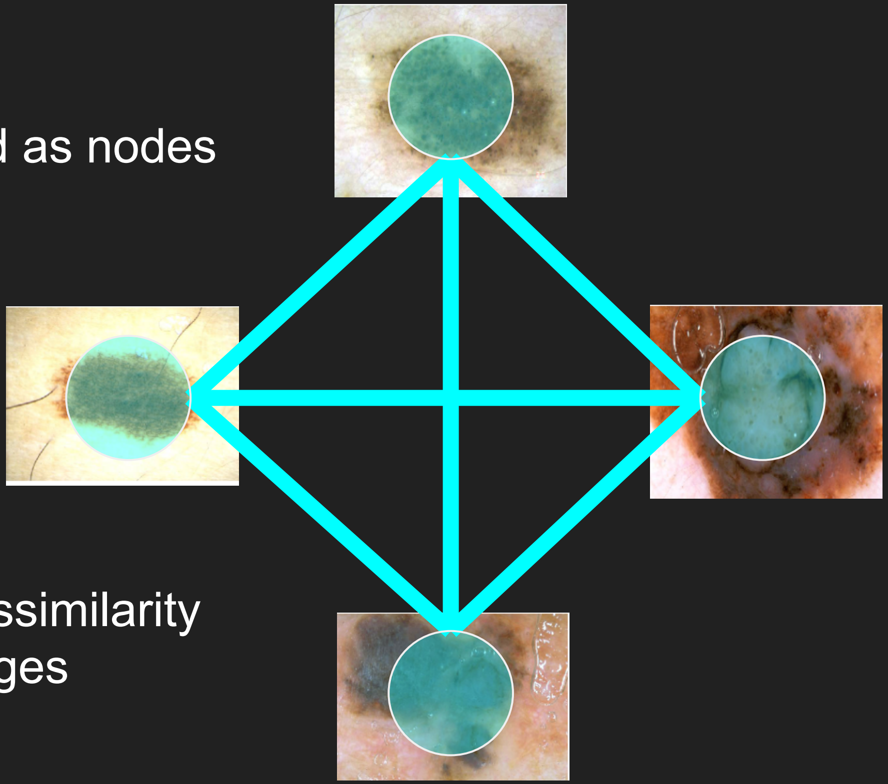

Edges represent dissimilarity between images

Skin dataset encoded as a graph

Images are encoded as nodes

Given a source and target node

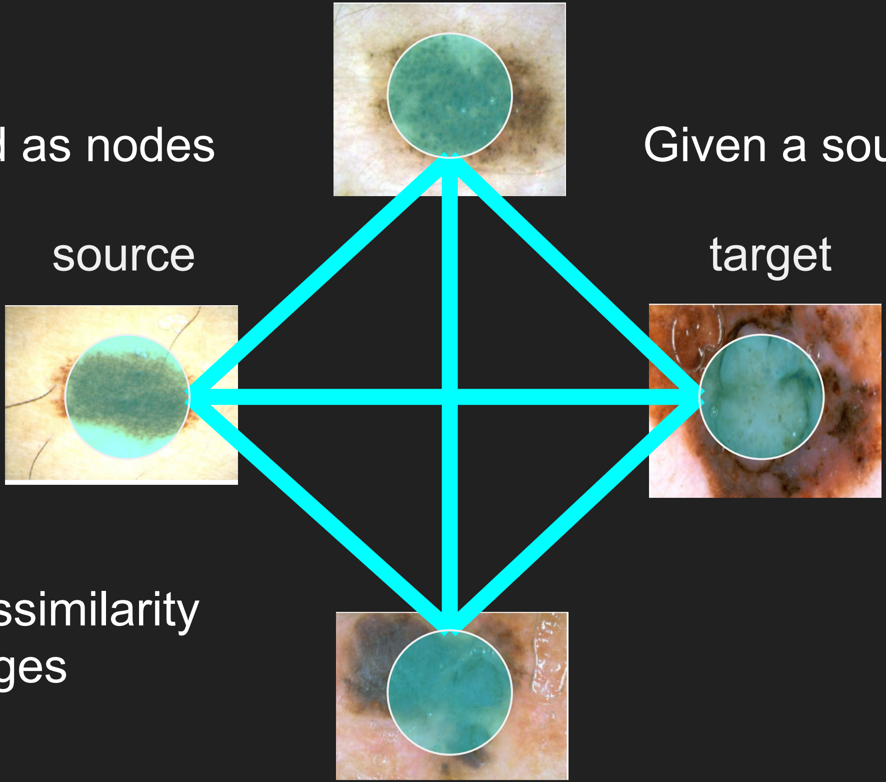

Edges represent dissimilarity between images

Skin dataset encoded as a graph

Images are encoded as nodes

Given a source and target node

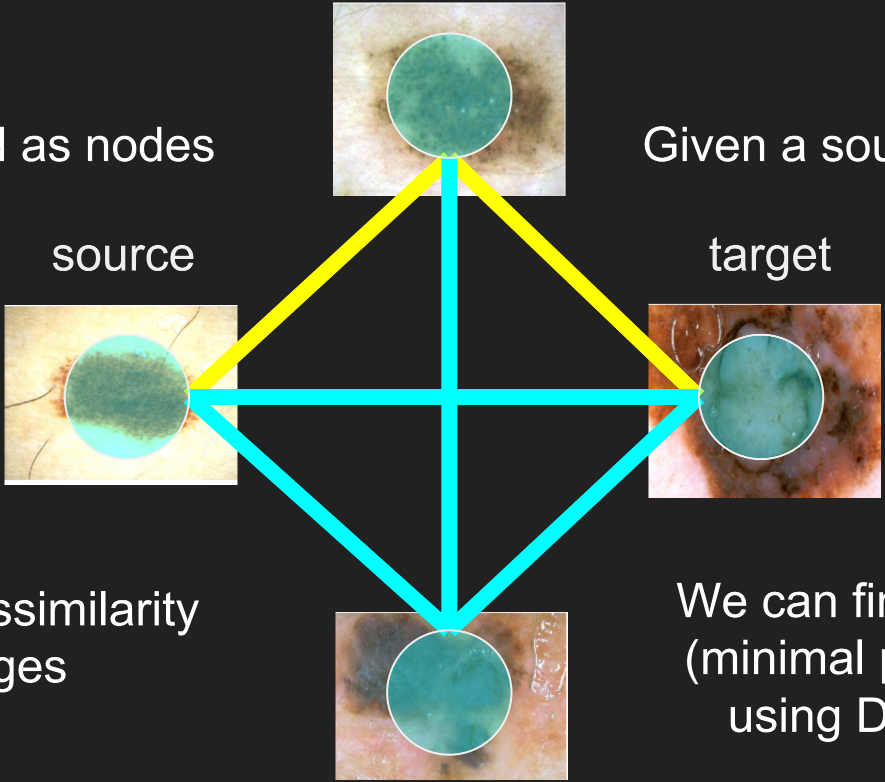

Edges represent dissimilarity between images

We can find a graph geodesic (minimal path) between them using Dijkstra's algorithm

Overview of our talk

1) How to encode images?

3) How to quantitatively evaluate the quality of paths?

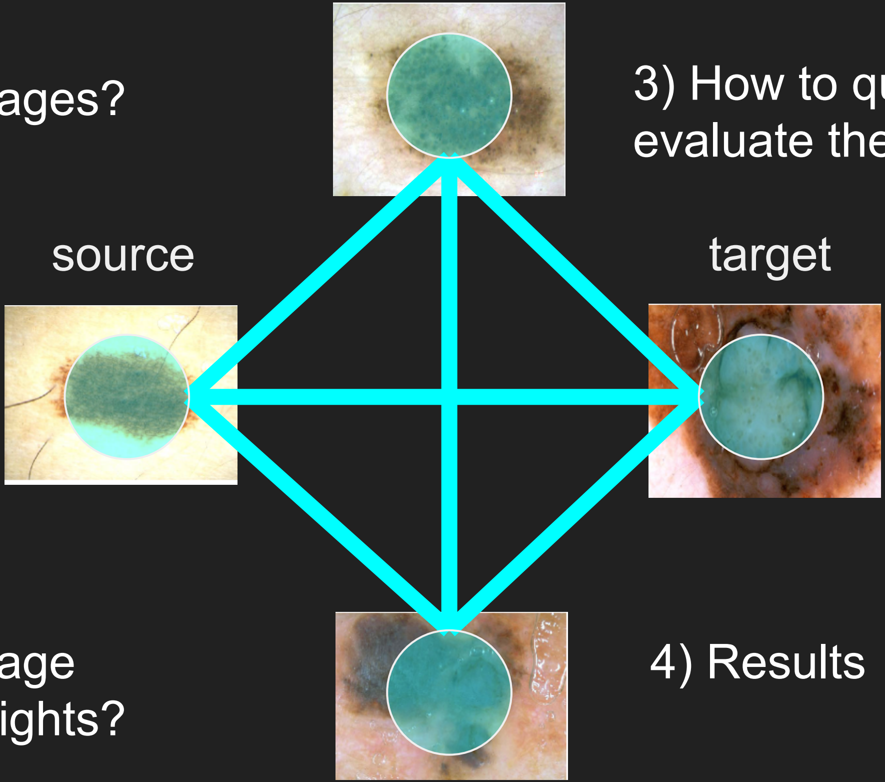

2) How to encode image dissimilarity/edge weights?

4) Results

Overview of our talk

1) How to encode images?

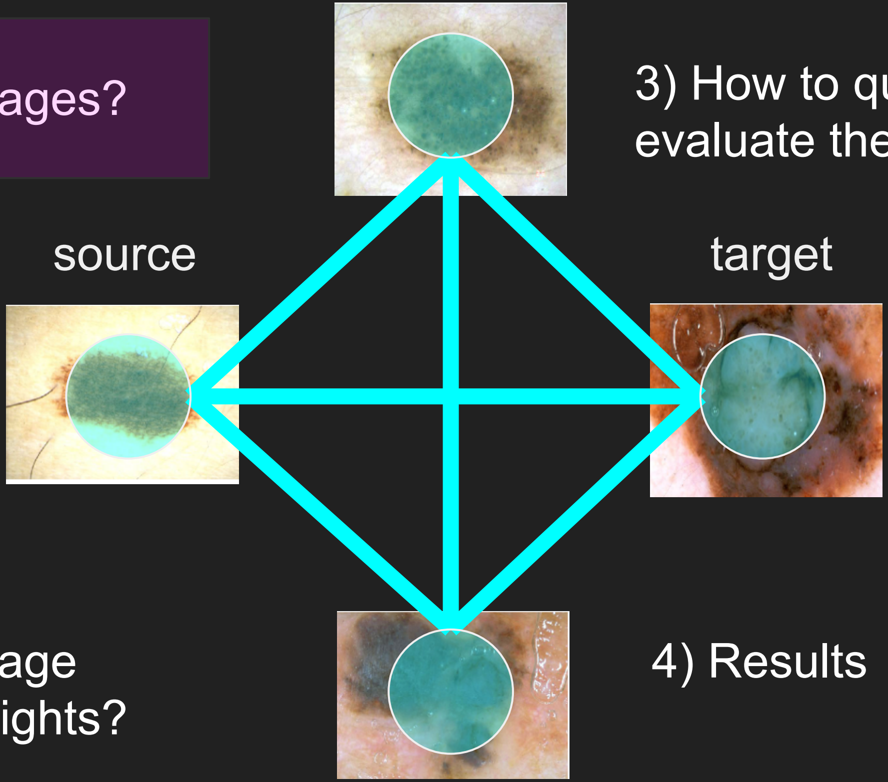

2) How to encode image dissimilarity/edge weights?

3) How to quantitatively evaluate the quality of paths?

4) Results

VGG 16 Pretrained over ImageNet

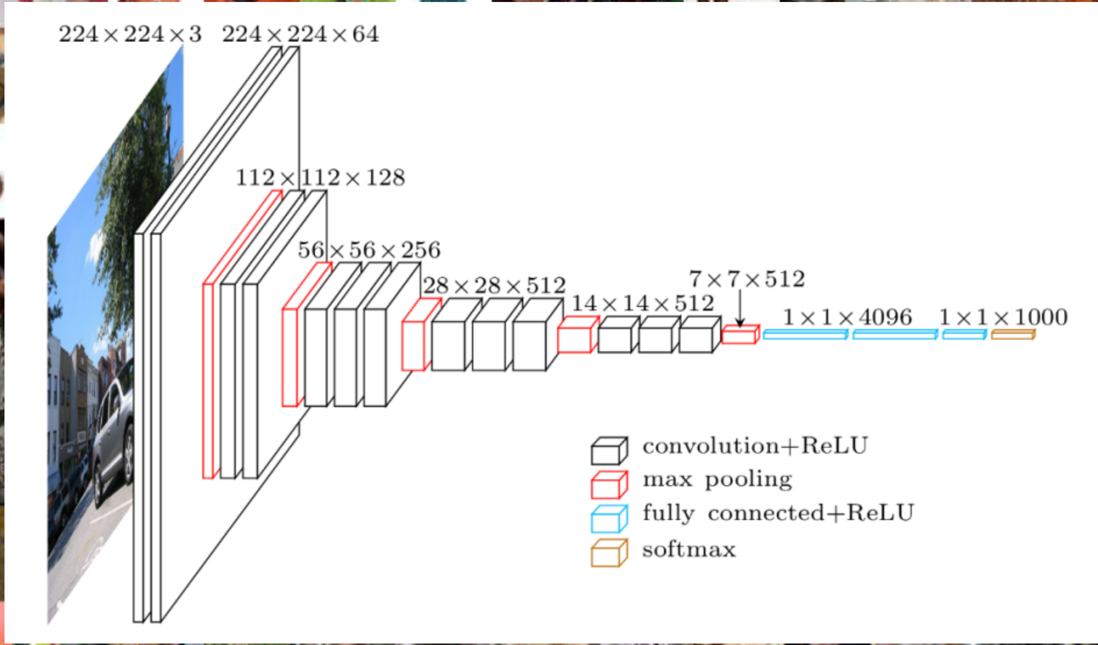

https://blog.heuritech.com/2016/02/29/a-brief-report-of-the-heuritech-deep-learning-meetup-5/

Skin images as deep neural network responses

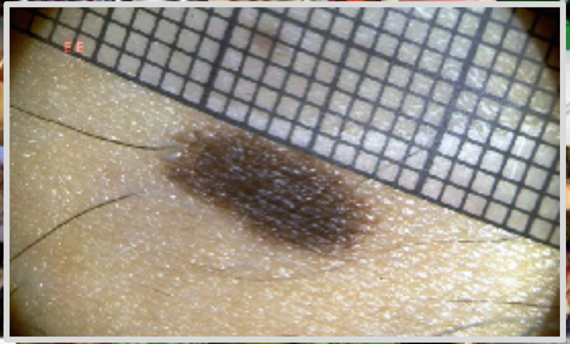

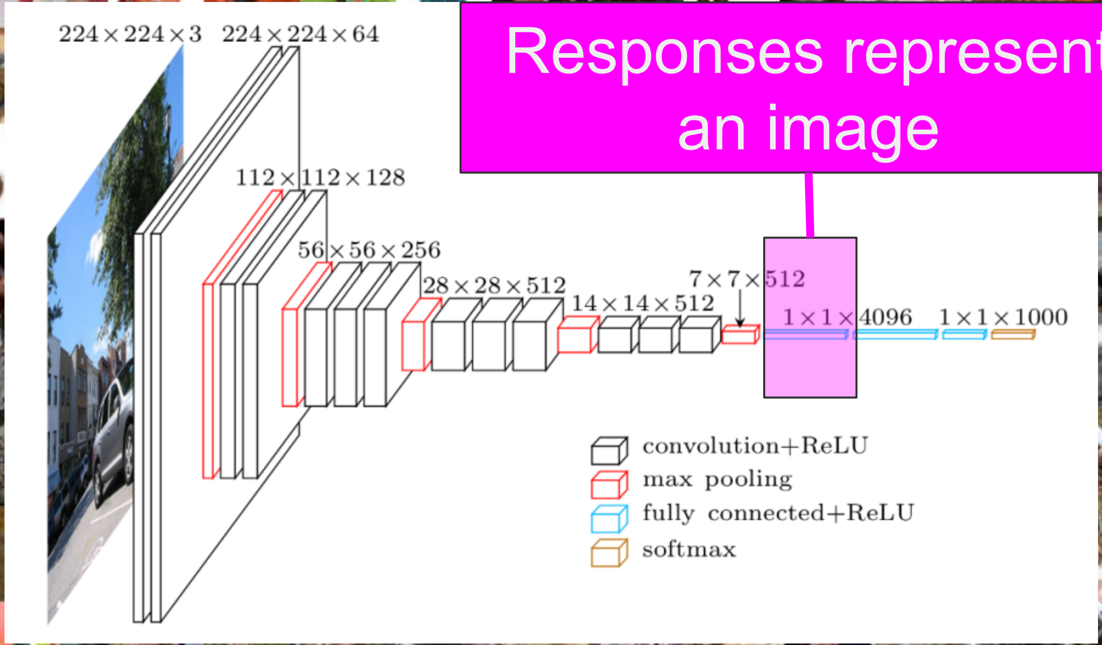

Responses represent an image

https://blog.heuritech.com/2016/02/29/a-brief-report-of-the-heuritech-deep-learning-meetup-5/

Skin images as deep neural network responses

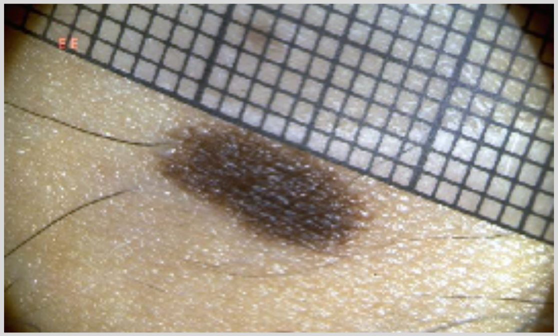

Left-right fip

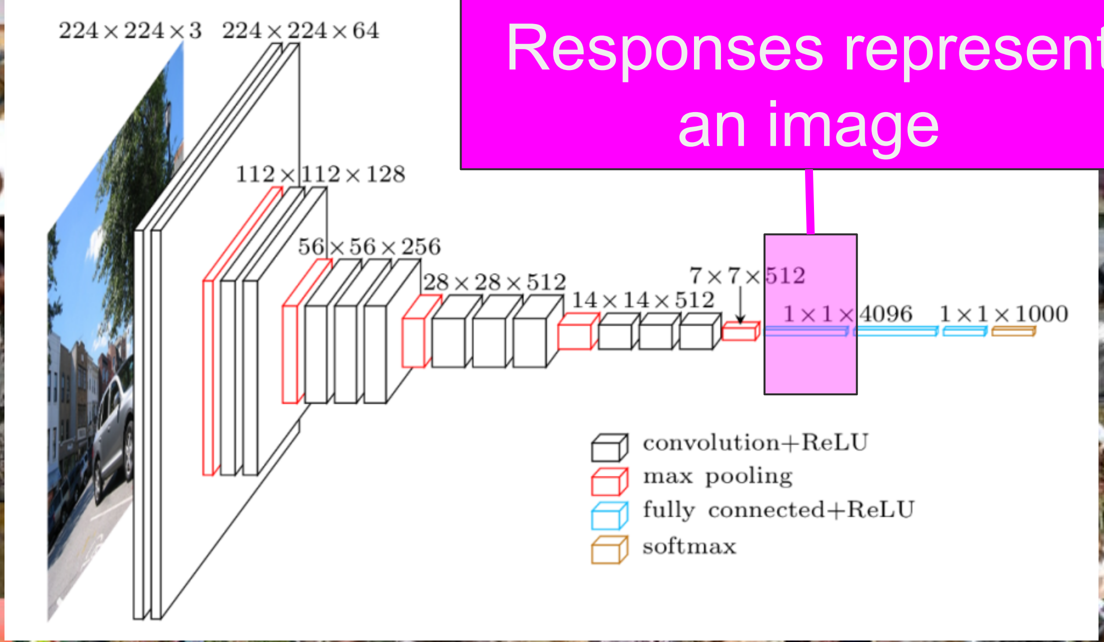

Responses represent an image

https://blog.heuritech.com/2016/02/29/a-brief-report-of-the-heuritech-deep-learning-meetup-5/

Skin images as deep neural network responses

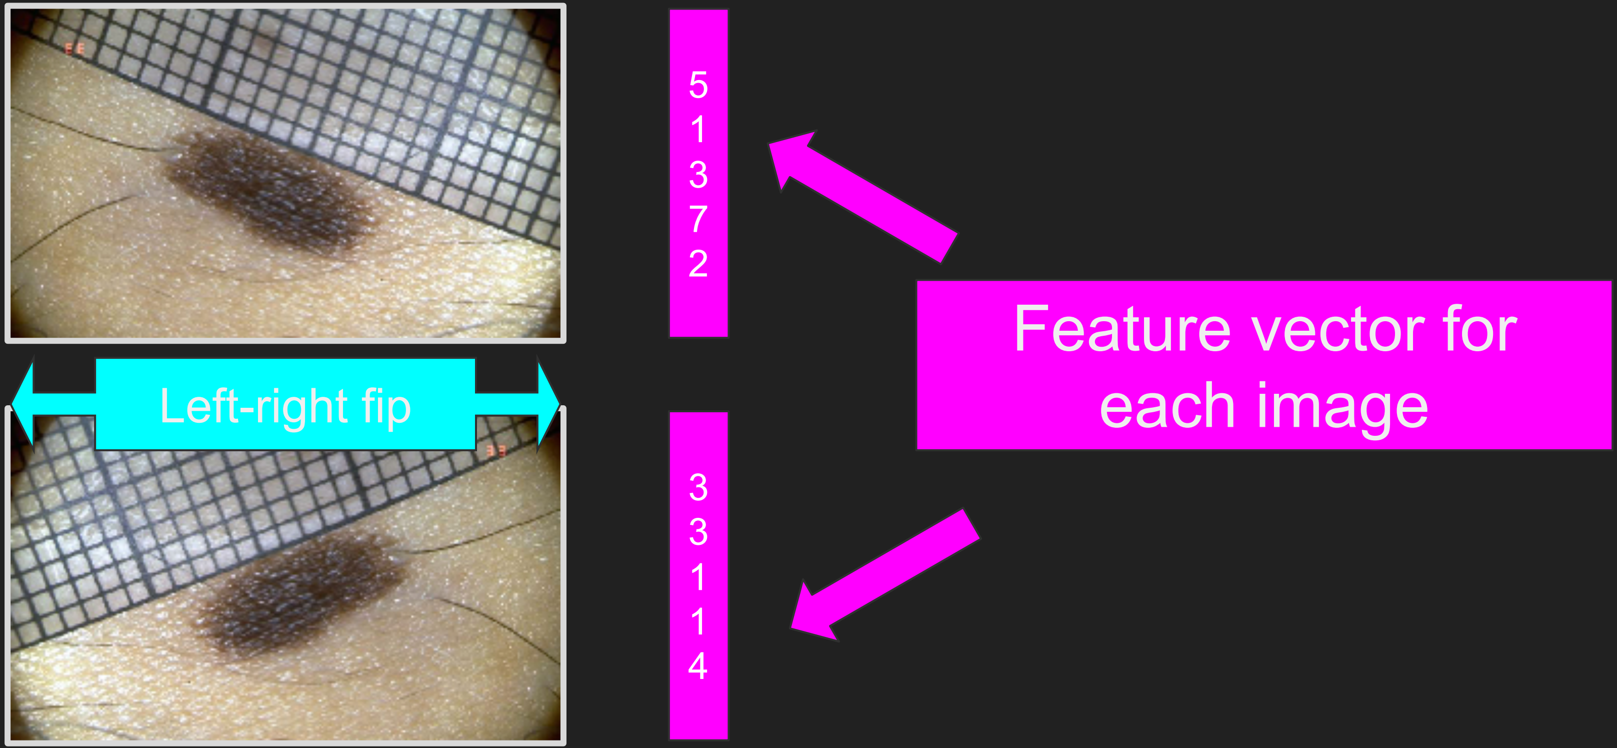

Skin images as deep neural network responses

5
1
3
7
2

3
3
1
1
4

4
2
2
4
3

Average responses to form a single vector that represents the image

Overview of our talk

1) How to encode images?

3) How to quantitatively evaluate the quality of paths?

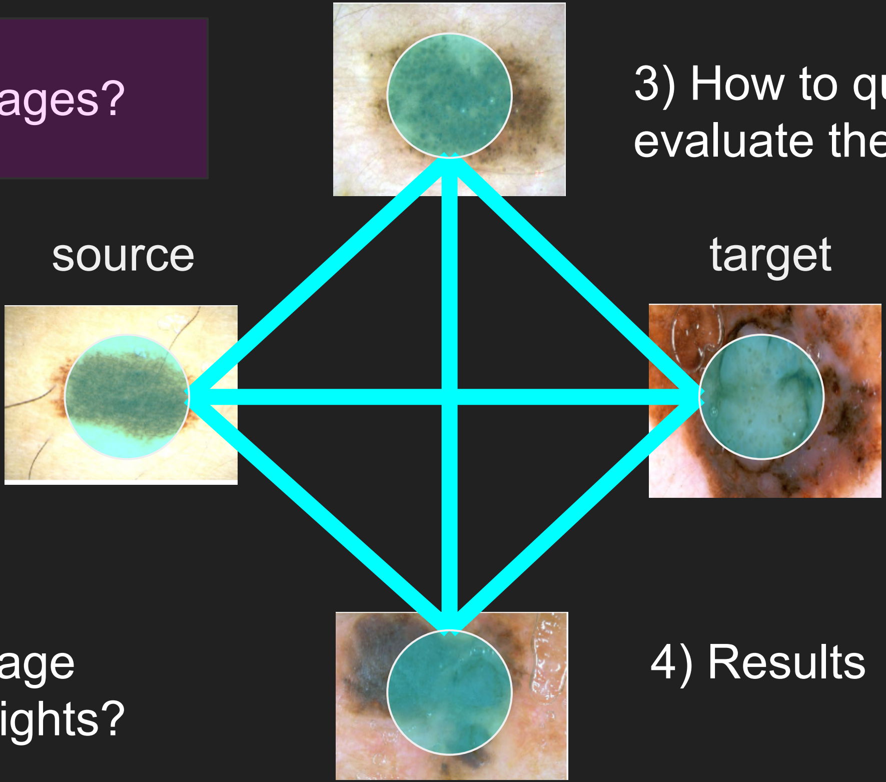

2) How to encode image dissimilarity/edge weights?

4) Results

Overview of our talk

1) How to encode images?

3) How to quantitatively evaluate the quality of paths?

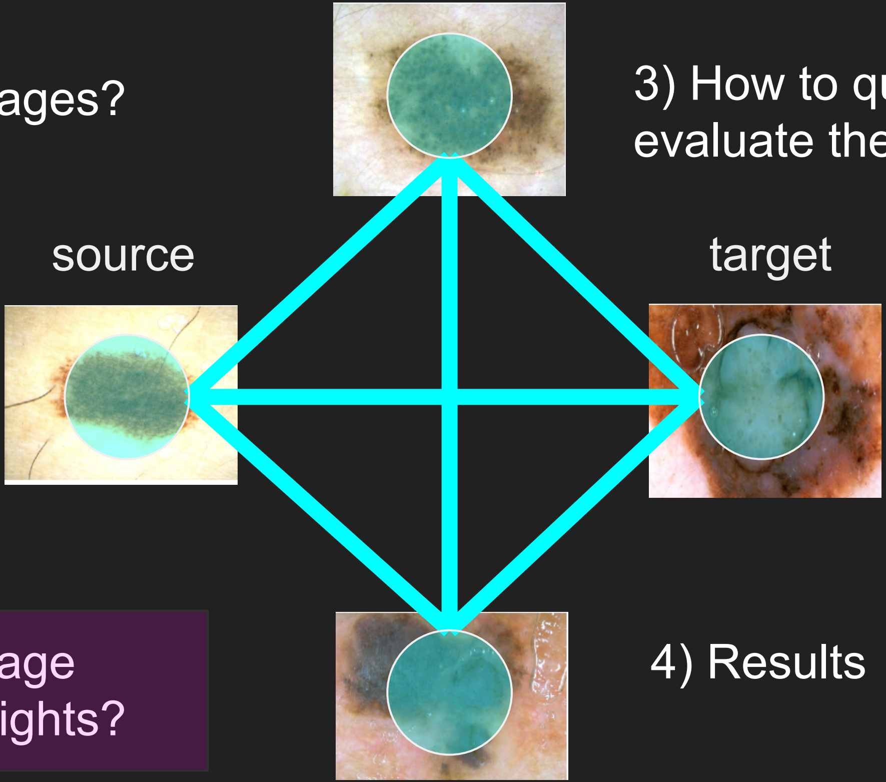

2) How to encode image dissimilarity/edge weights?

4) Results

Dissimilarity between pairs of images

$\mathcal{D}(x^{(i)}, x^{(j)})$ = the **cosine** distance

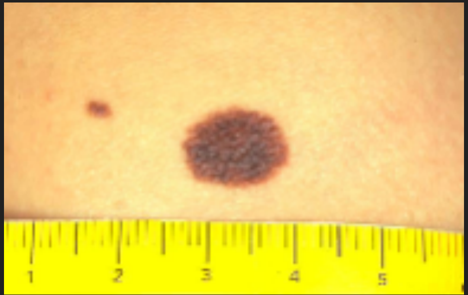

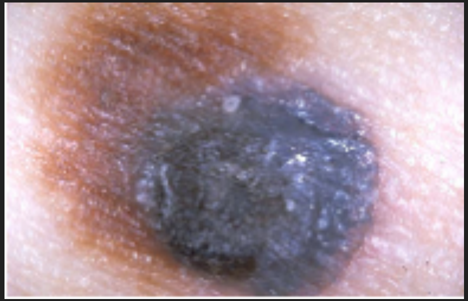

Dissimilarity between pairs of images

$\mathcal{D}(x^{(i)}, x^{(j)})$ = the **cosine** distance

Magenta indicates the dissimilarity between images

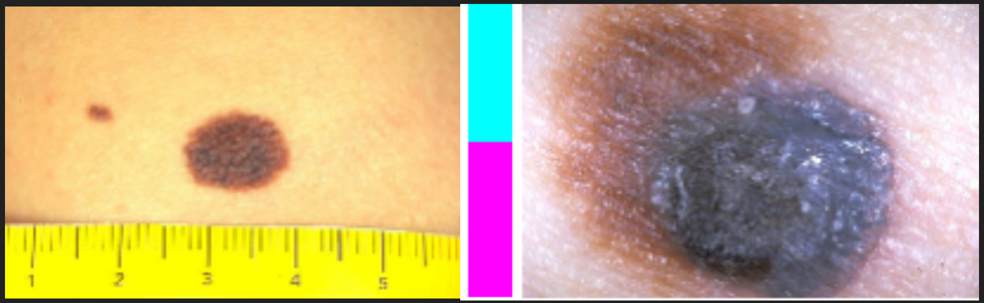

Dissimilarity between pairs of images

$\mathcal{D}(x^{(i)}, x^{(j)})$ = the **cosine** distance

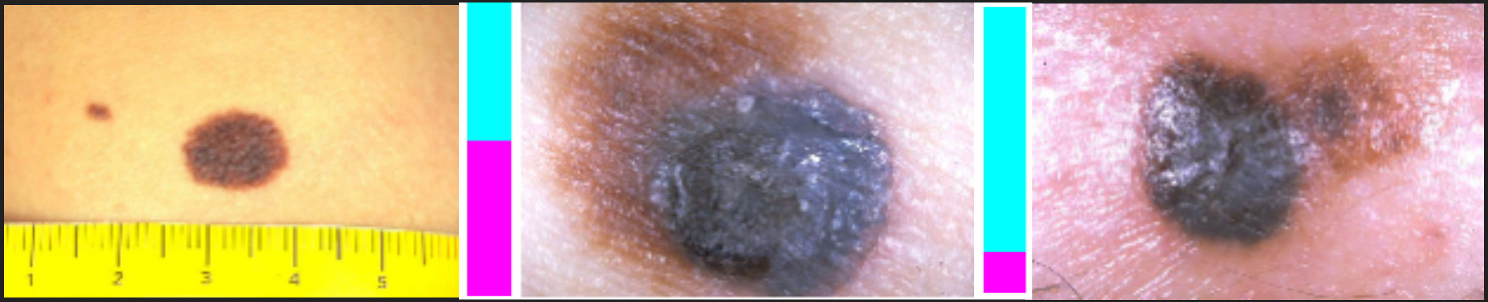

High dissimilarity between visually **dissimilar** images

Low dissimilarity between visually **similar** images

Dissimilarity encoded as graph edges

$\mathcal{D}(x^{(i)}, x^{(j)})$ = the **cosine** distance

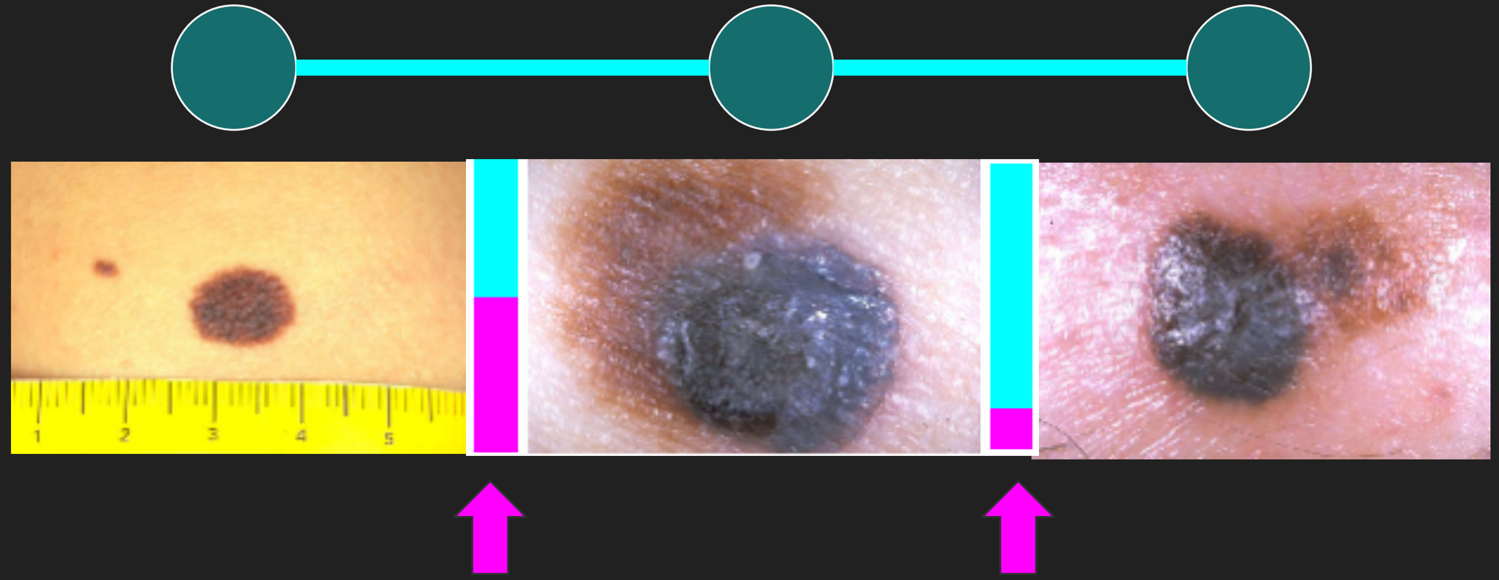

High dissimilarity between visually **dissimilar** images

Low dissimilarity between visually **similar** images

Dissimilarity encoded as graph edges

$$e^{(ij)} = \mathcal{D}(x^{(i)}, x^{(j)})$$ = the **cosine** distance

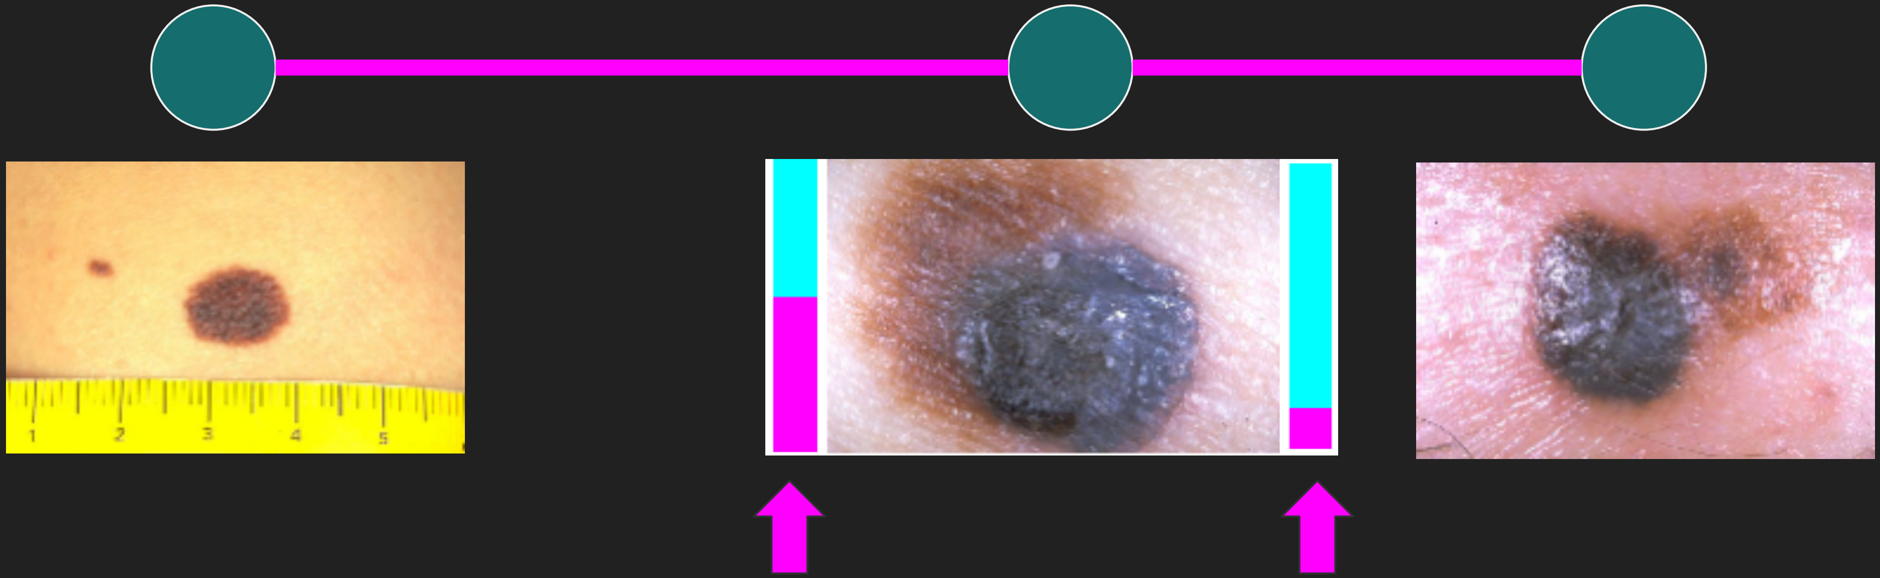

High dissimilarity between visually **dissimilar** images

Low dissimilarity between visually **similar** images

Equidistant edge weights

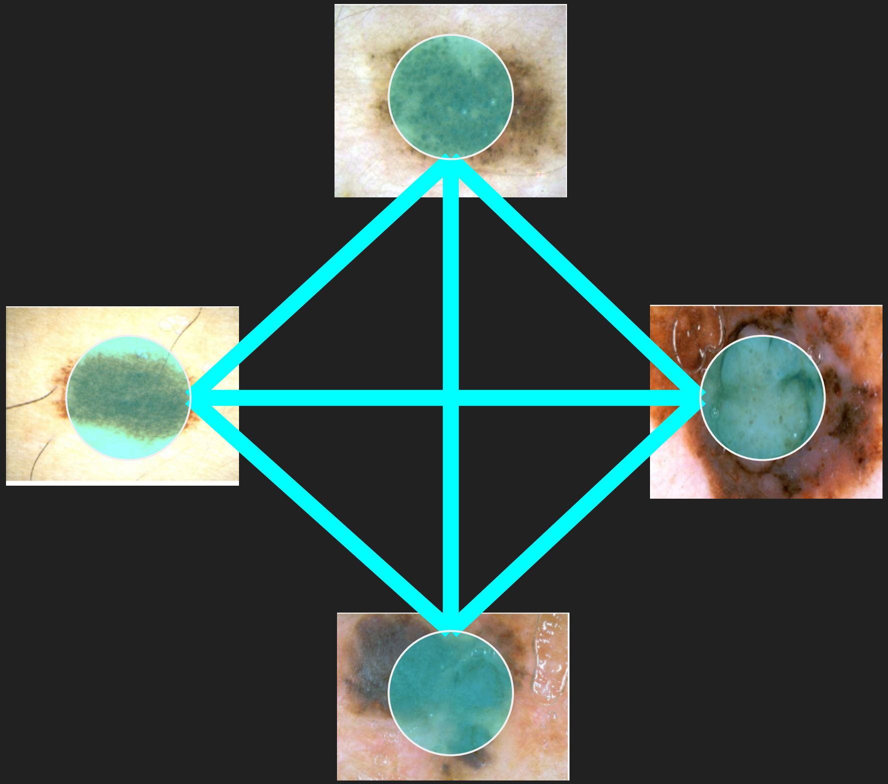

Image dissimilarity as edge weights

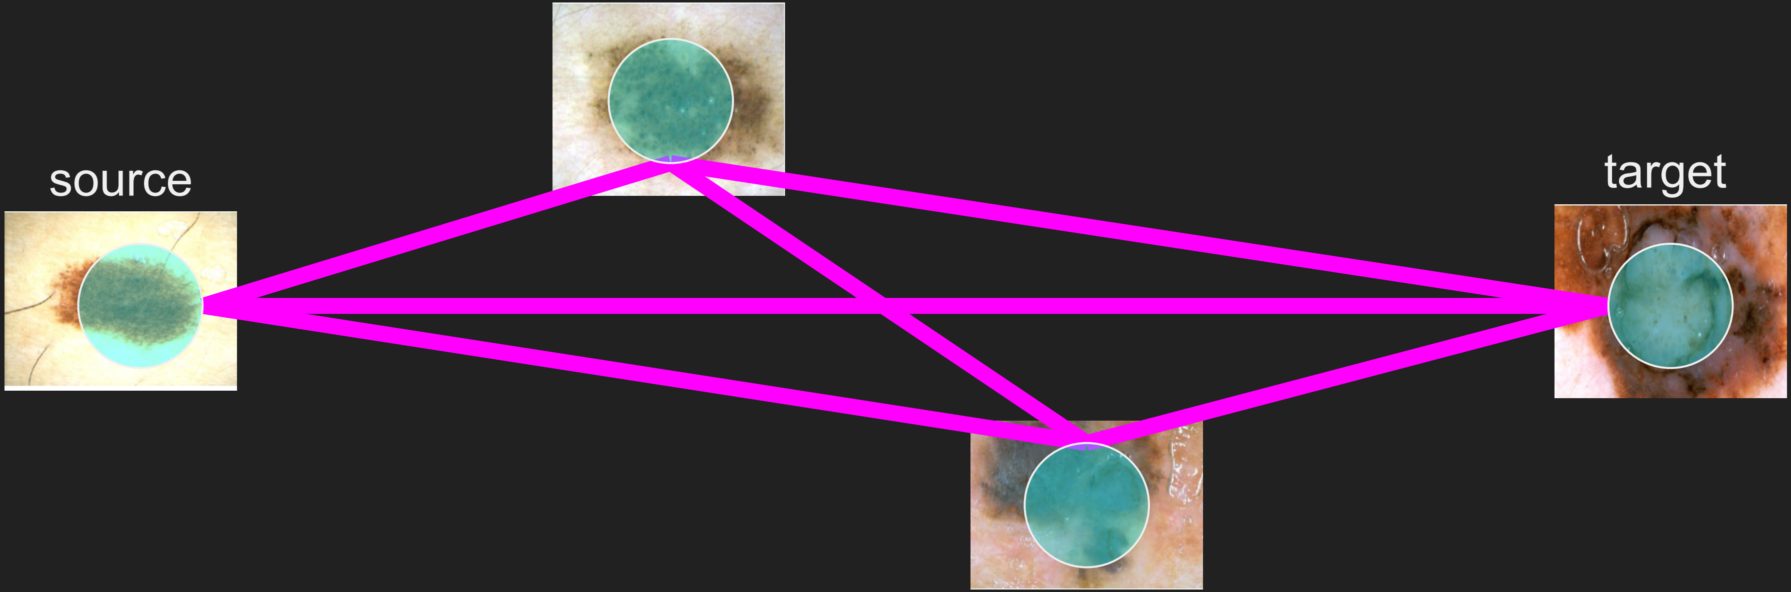

Edges represent the dissimilarity between two images

Problem: very short paths

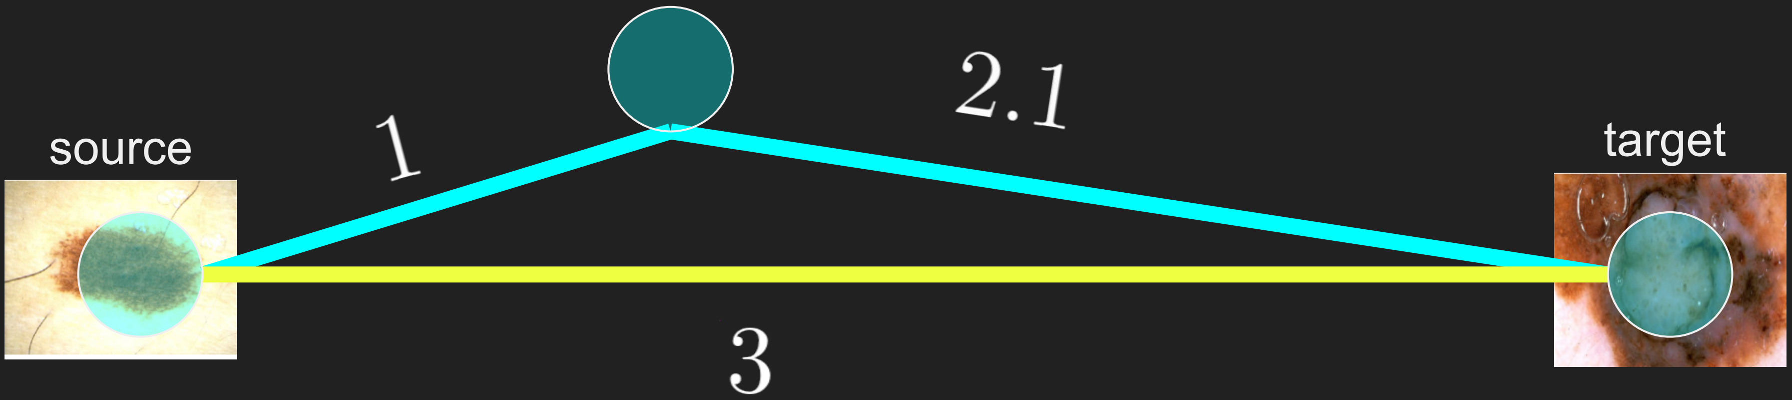

3 < 1+2.1

In a complete graph, the direct edge will almost always be chosen

Common solution: Prune edges

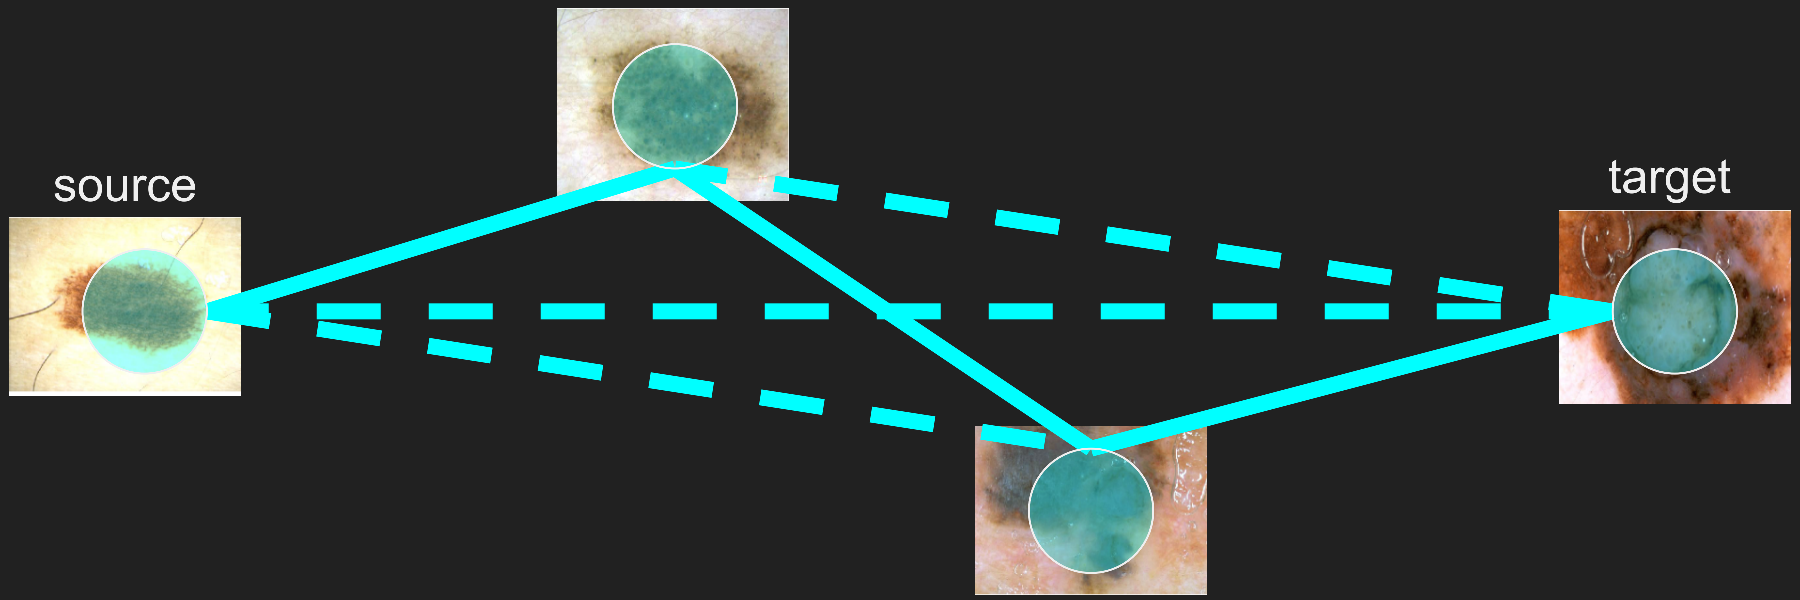

Common solution: Prune edges

Potential problem: how many edges to prune?

Can lead to disconnected graphs with no path

Problem: very short paths

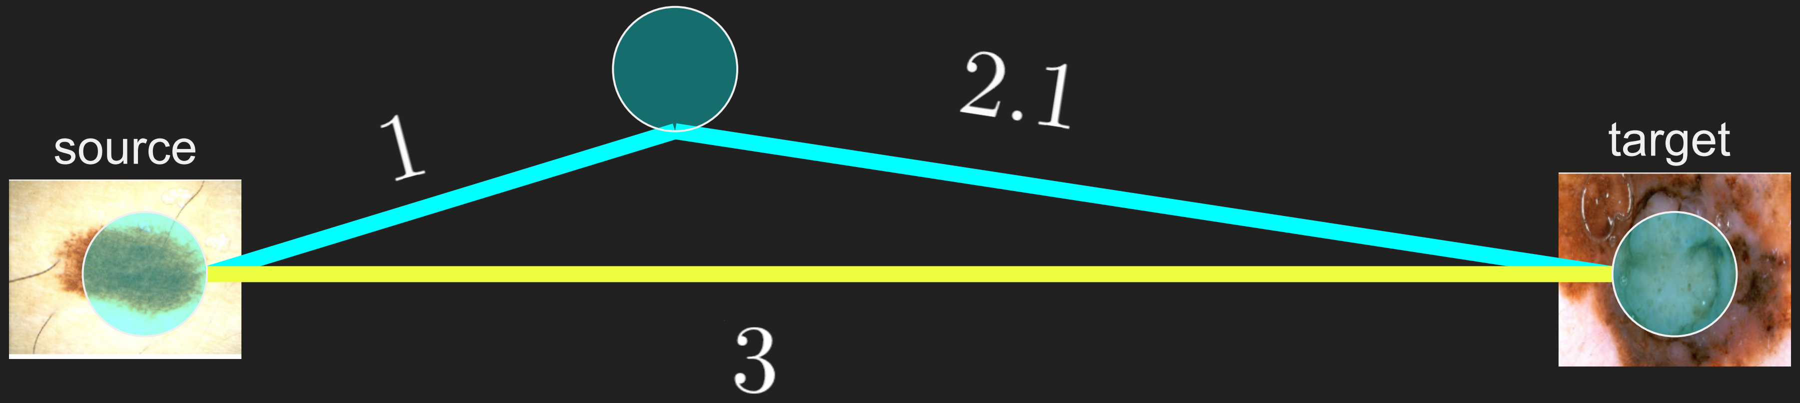

3 < 1+2.1

In a complete graph, the direct edge will almost always be chosen

Problem: very short paths

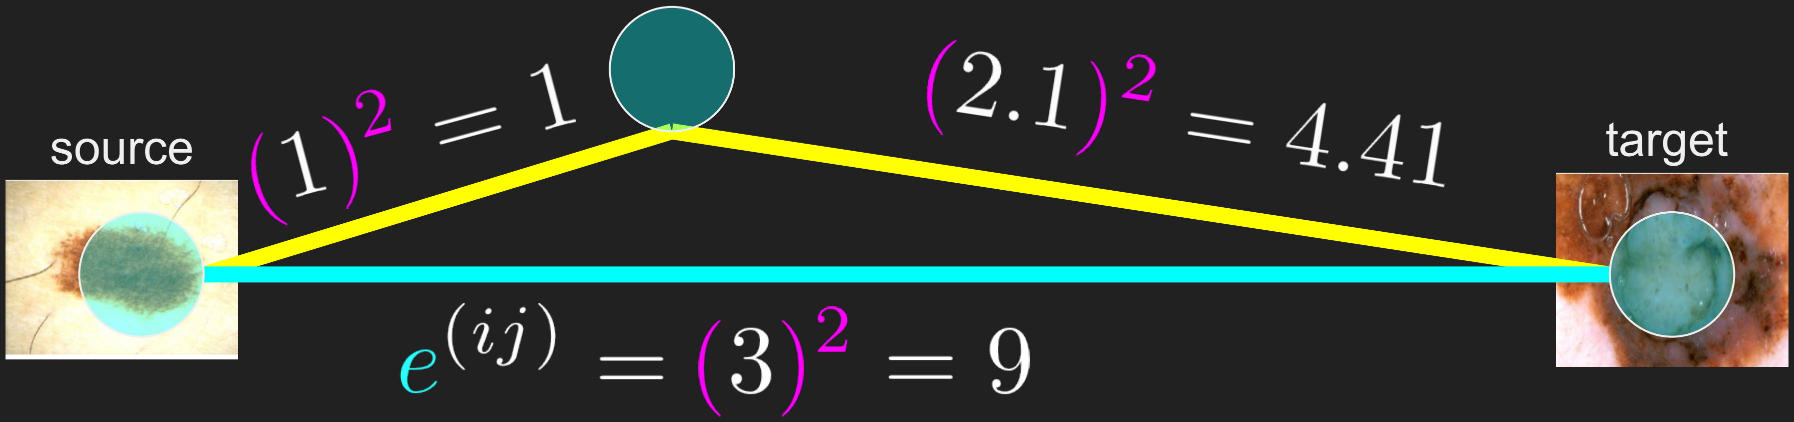

$$9 > 1 + 4.41$$

Solution: exponential dissimilarity

Multi-Modal Edge Weights

${x_d}^{(i)}$

dermoscopic

${x_d}^{(j)}$

dermoscopic

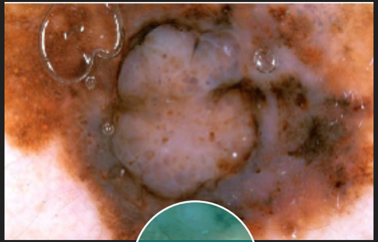

Computes dissimilarity between images

$$e^{(ij)} = \mathcal{D}({x_d}^{(i)}, {x_d}^{(j)})$$

So far, we have only looked at the
dermoscopic images (captured by a dermatoscope)

Multi-Modal Edge Weights

$x_d{}^{(i)}$

dermoscopic

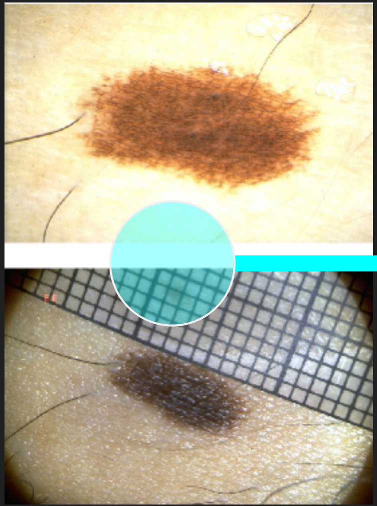

clinical

$x_c{}^{(i)}$

$x_d{}^{(j)}$

dermoscopic

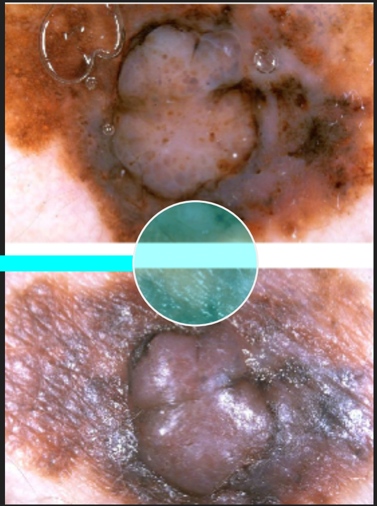

clinical

$x_c{}^{(j)}$

Computes dissimilarity between images

$$e^{(ij)} = \mathcal{D}(x_d{}^{(i)}, x_d{}^{(j)})$$

So far, we have only looked at the **dermoscopic images** (captured by a dermatoscope)

Each lesion also has a **clinical image** (can contain background)

Multi-Modal Edge Weights

$x_d^{(i)}$

dermoscopic

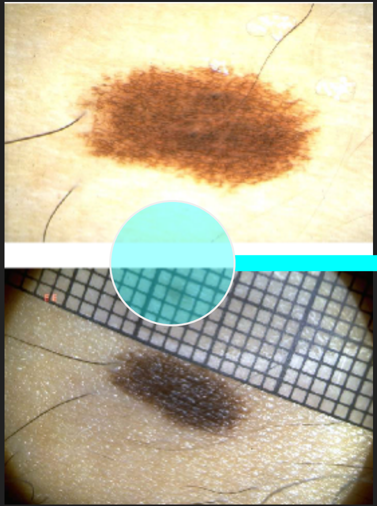

clinical

$x_c^{(i)}$

$x_d^{(j)}$

dermoscopic

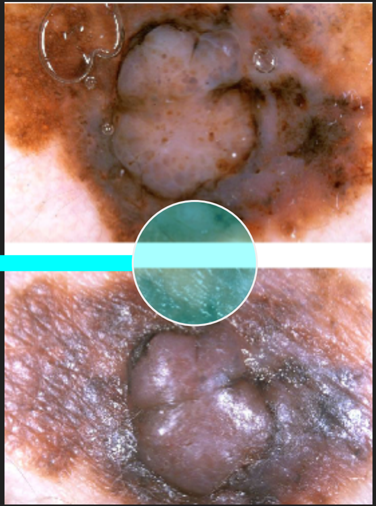

clinical

$x_c^{(j)}$

Computes dissimilarity between images

$$e^{(ij)} = \mathcal{D}({x_d}^{(i)}, {x_d}^{(j)}) + \mathcal{D}({x_c}^{(i)}, {x_c}^{(j)})$$

So far, we have only looked at the **dermoscopic images** (captured by a dermatoscope)

Each lesion also has a **clinical image**
(can contain background)

Multi-Modal Edge Weights

$x_d^{(i)}$

dermoscopic

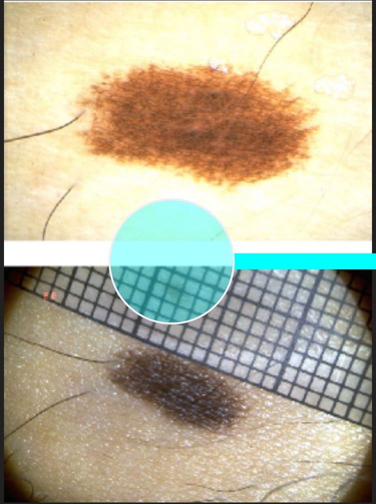

clinical

$x_c^{(i)}$

Computes dissimilarity between images

$$e^{(ij)} = \alpha\mathcal{D}(x_d^{(i)}, x_d^{(j)}) + (1-\alpha)\mathcal{D}(x_c^{(i)}, x_c^{(j)})$$

$x_d^{(j)}$

dermoscopic

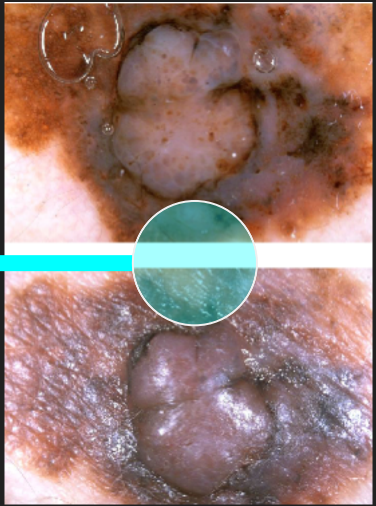

clinical

$x_c^{(j)}$

α = 0.8 gives heavier weight to dermoscopic images

(as clinical images can contain background artefacts)

Overview of our talk

1) How to encode images?

3) How to quantitatively evaluate the quality of paths?

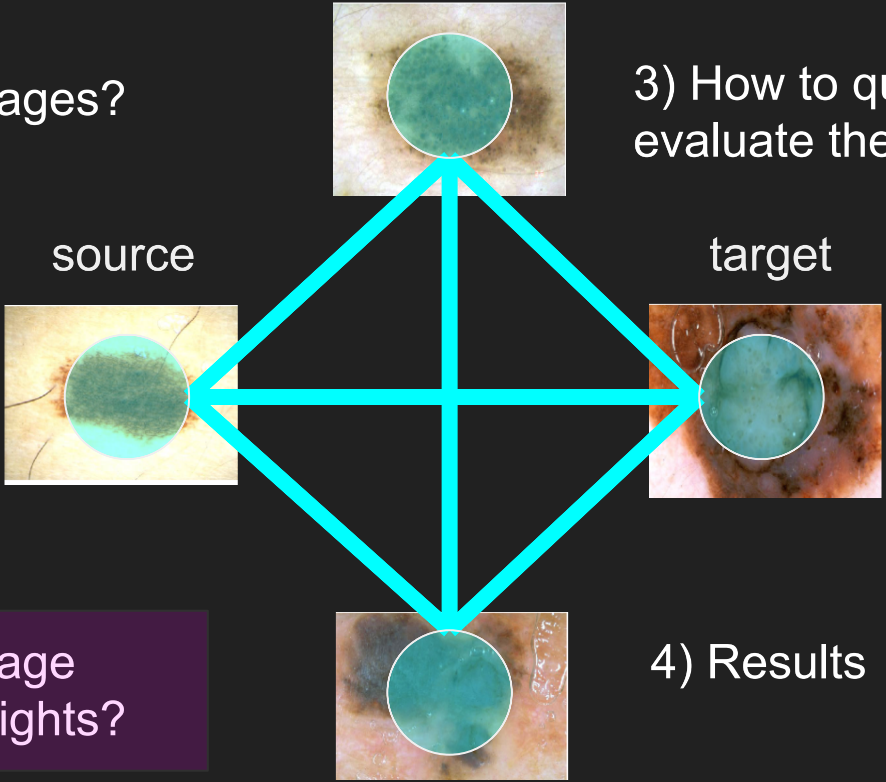

2) How to encode image dissimilarity/edge weights?

4) Results

Overview of our talk

1) How to encode images?

3) How to quantitatively evaluate the quality of paths?

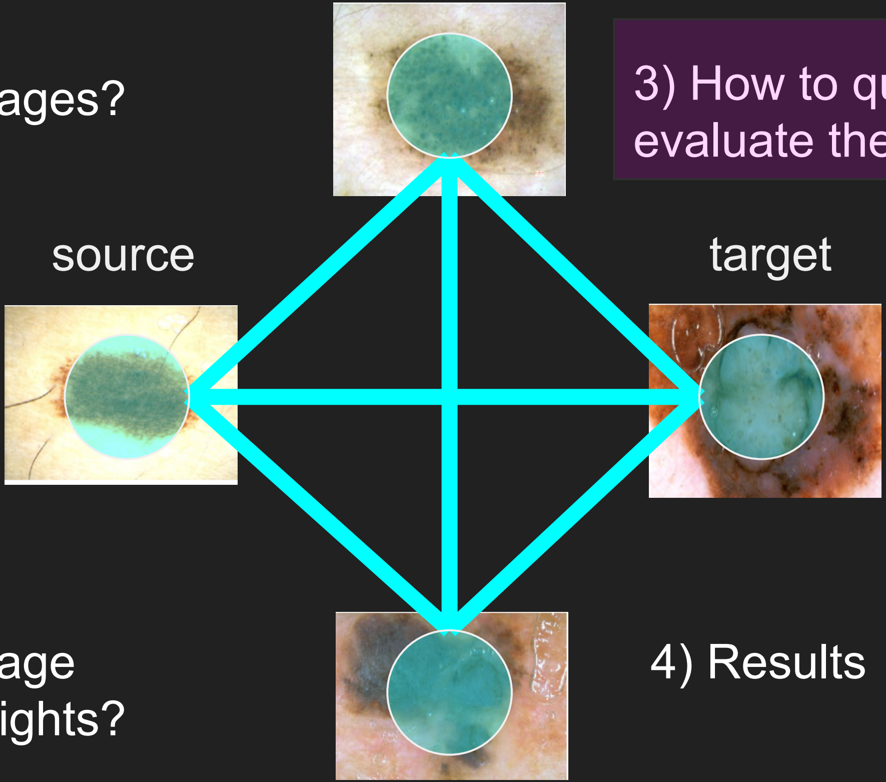

2) How to encode image dissimilarity/edge weights?

4) Results

Evaluate Path Quality How to quantify different paths?

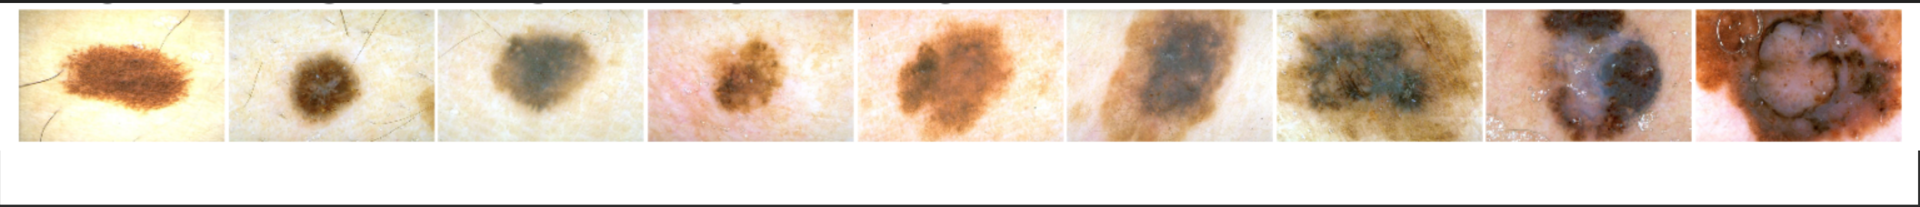

Evaluate Path Quality How to quantify different paths?

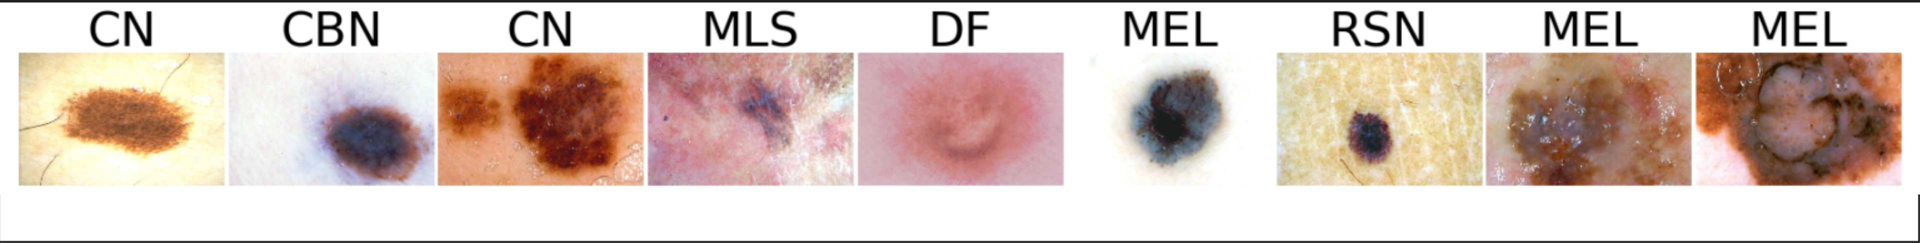

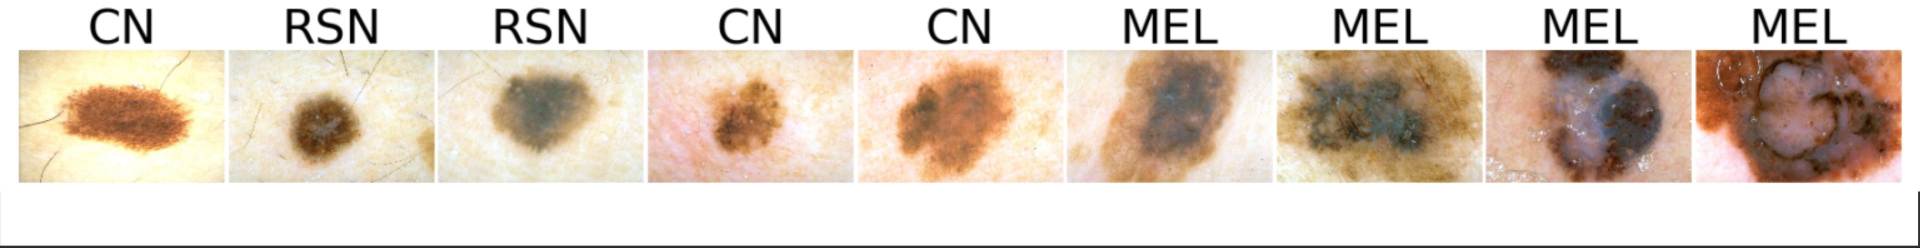

Images have class labels associated with them.

Evaluate Path Quality How to quantify different paths?

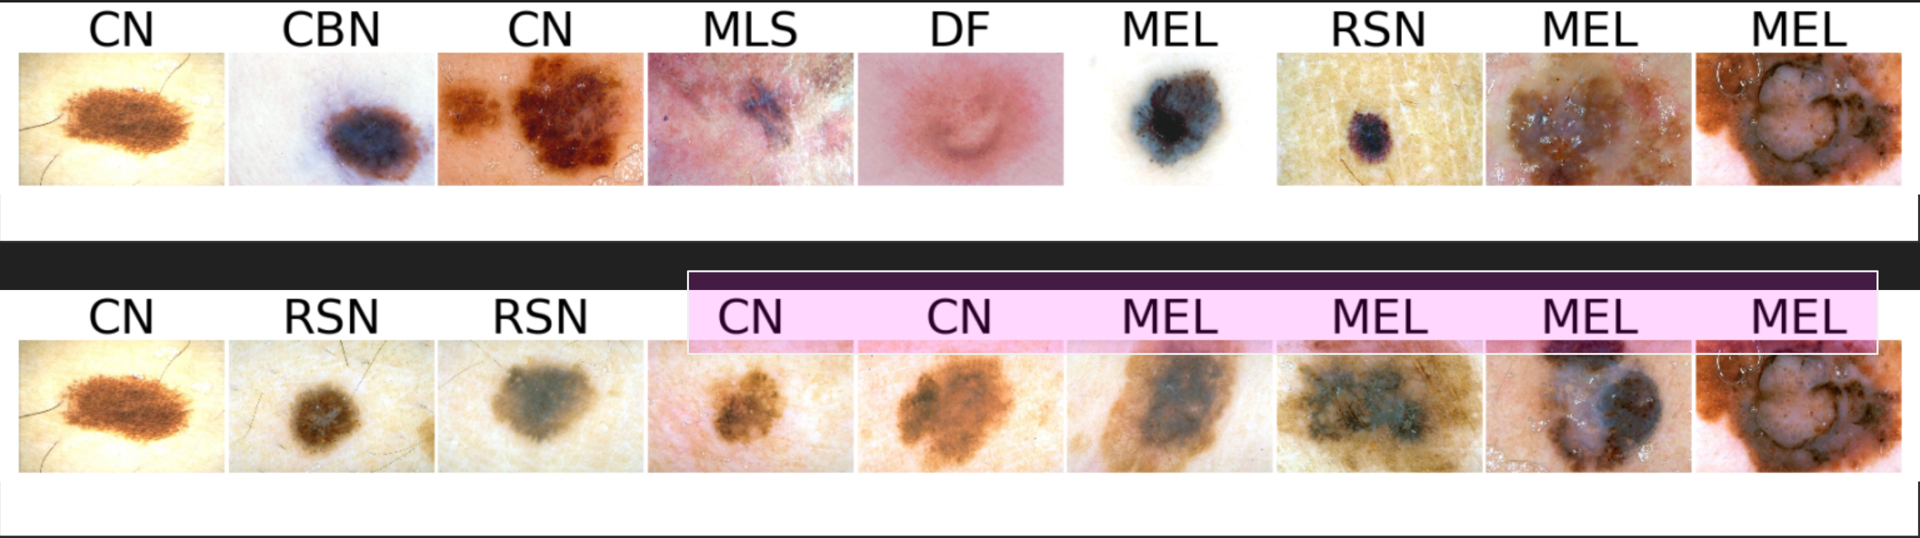

Images have class labels associated with them

Transition cost = a path should have a smooth transition of class labels

Evaluate Path Quality How to quantify different paths?

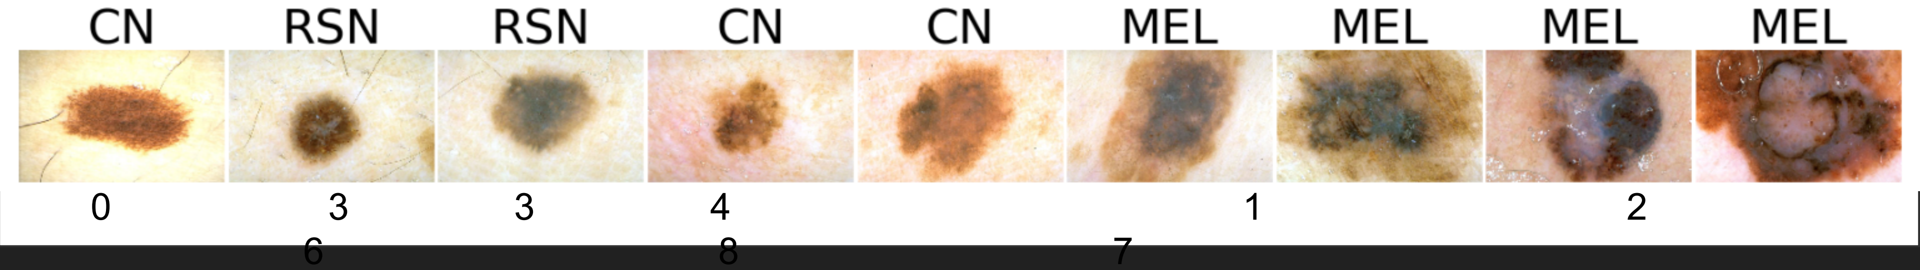

Images have class labels associated with them

Transition cost = a path should have a smooth transition of class labels

Images have a 7-point score, where higher values indicate melanoma

Evaluate Path Quality How to quantify different paths?

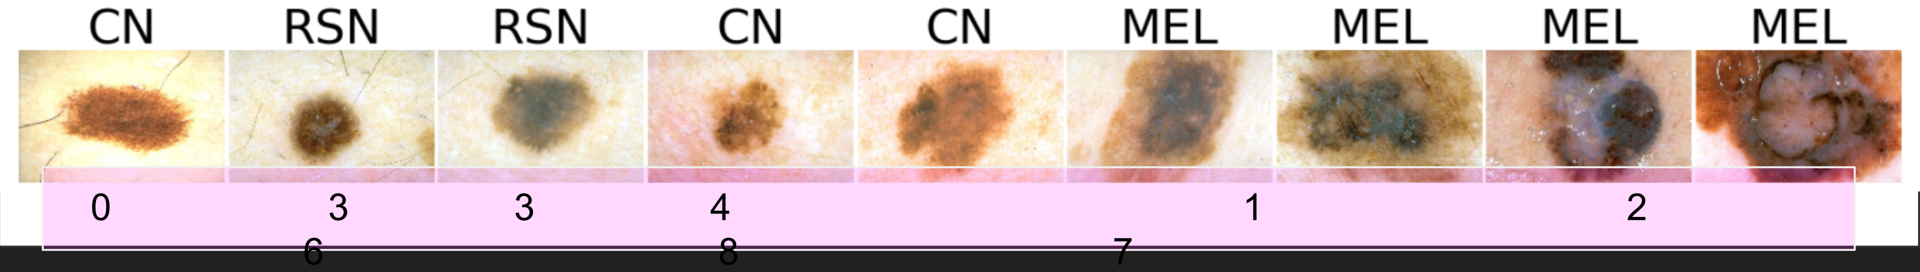

Images have class labels associated with them

Transition cost = a path should have a smooth transition of class labels

Images have a 7-point score, where higher values indicate melanoma

Progression cost = the **7-point scores** in a path should consistently increase/decrease between the source and target

Transition cost

smooth transition of class labels

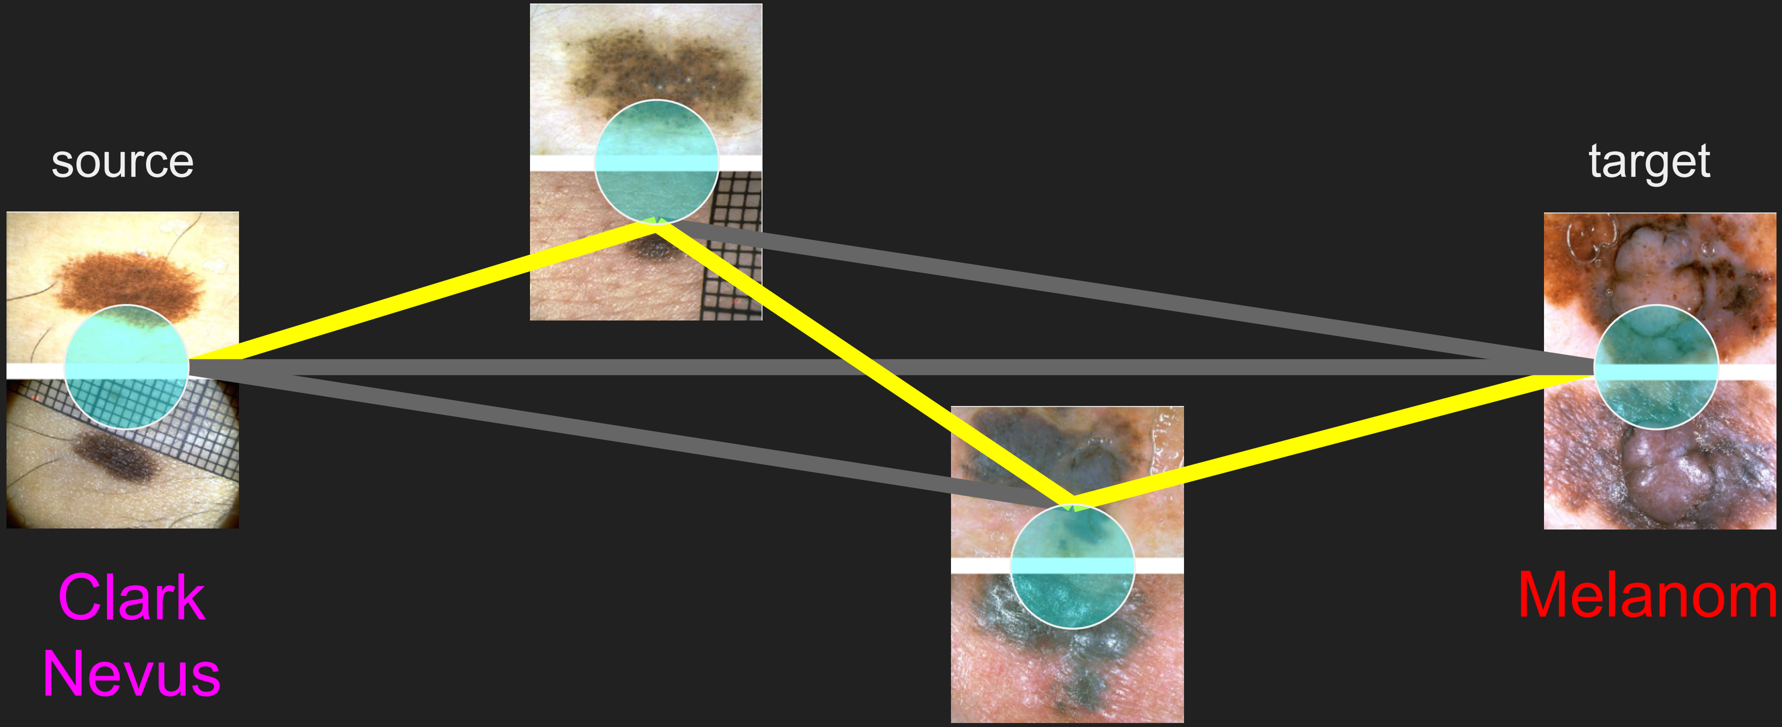

Transition cost

smooth transition of class labels

$$C = \begin{cases} 0, & \text{if same labels} \\ 1, & \text{otherwise} \end{cases}$$

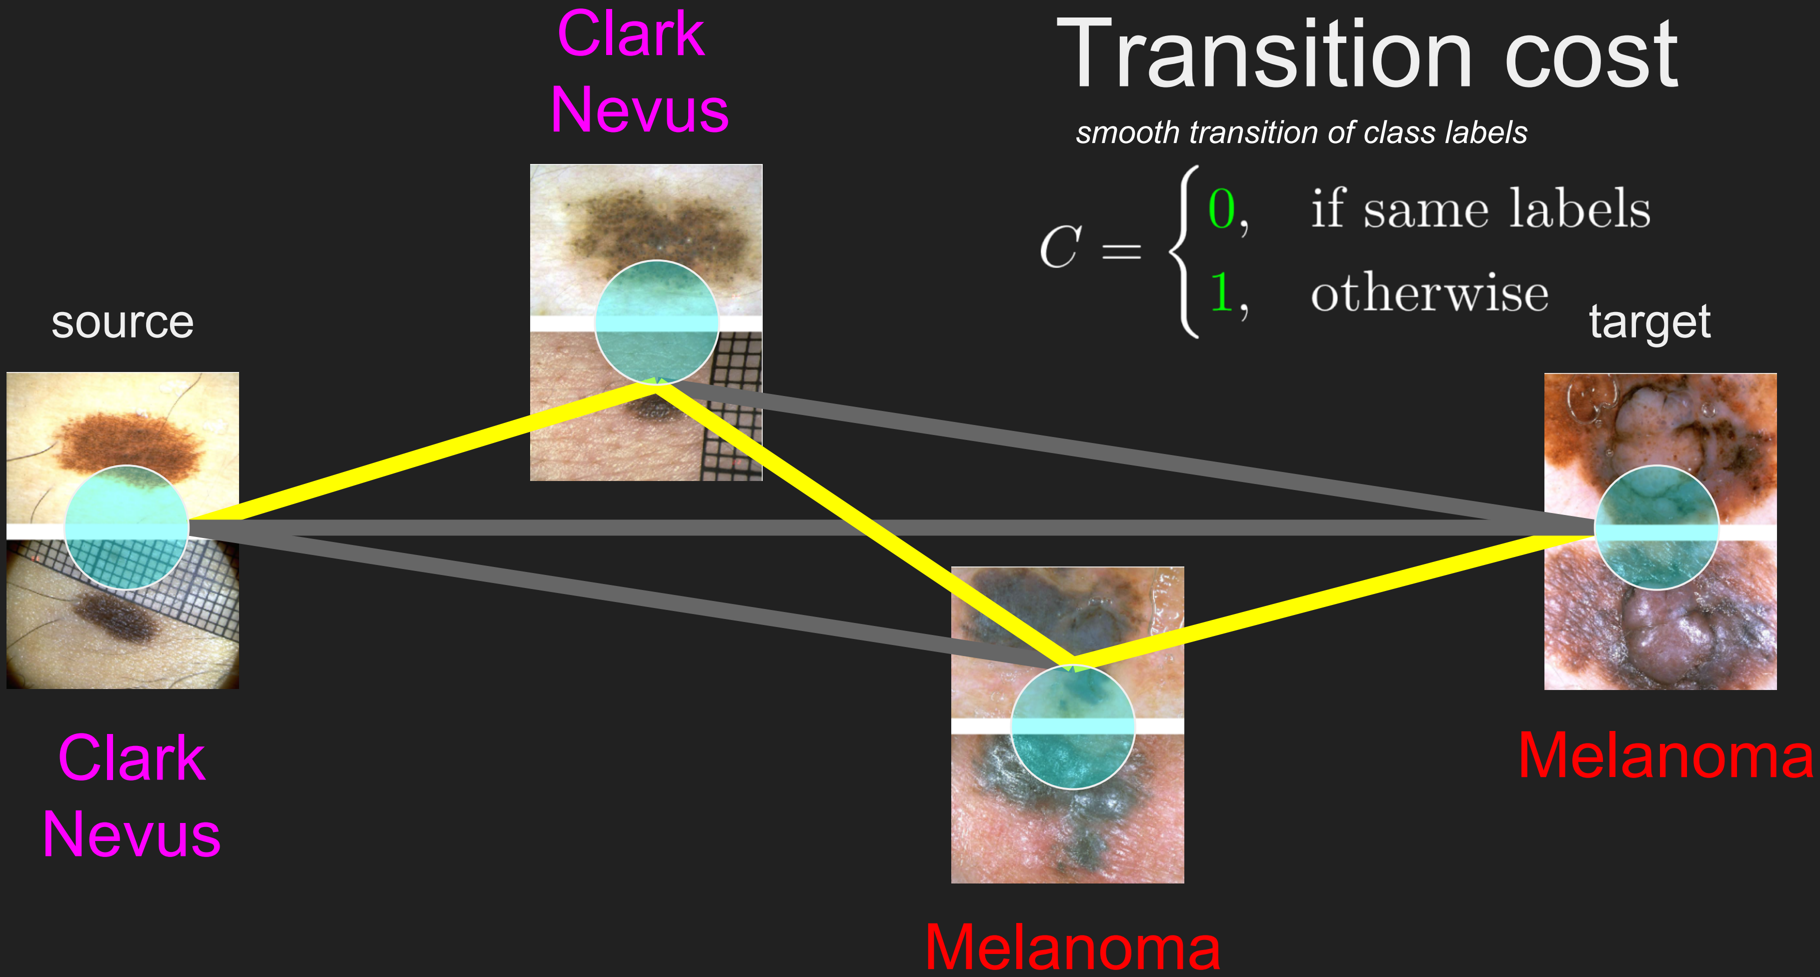

Transition cost

smooth transition of class labels

$$C = \begin{cases} 0, & \text{if same labels} \\ 1, & \text{otherwise} \end{cases}$$

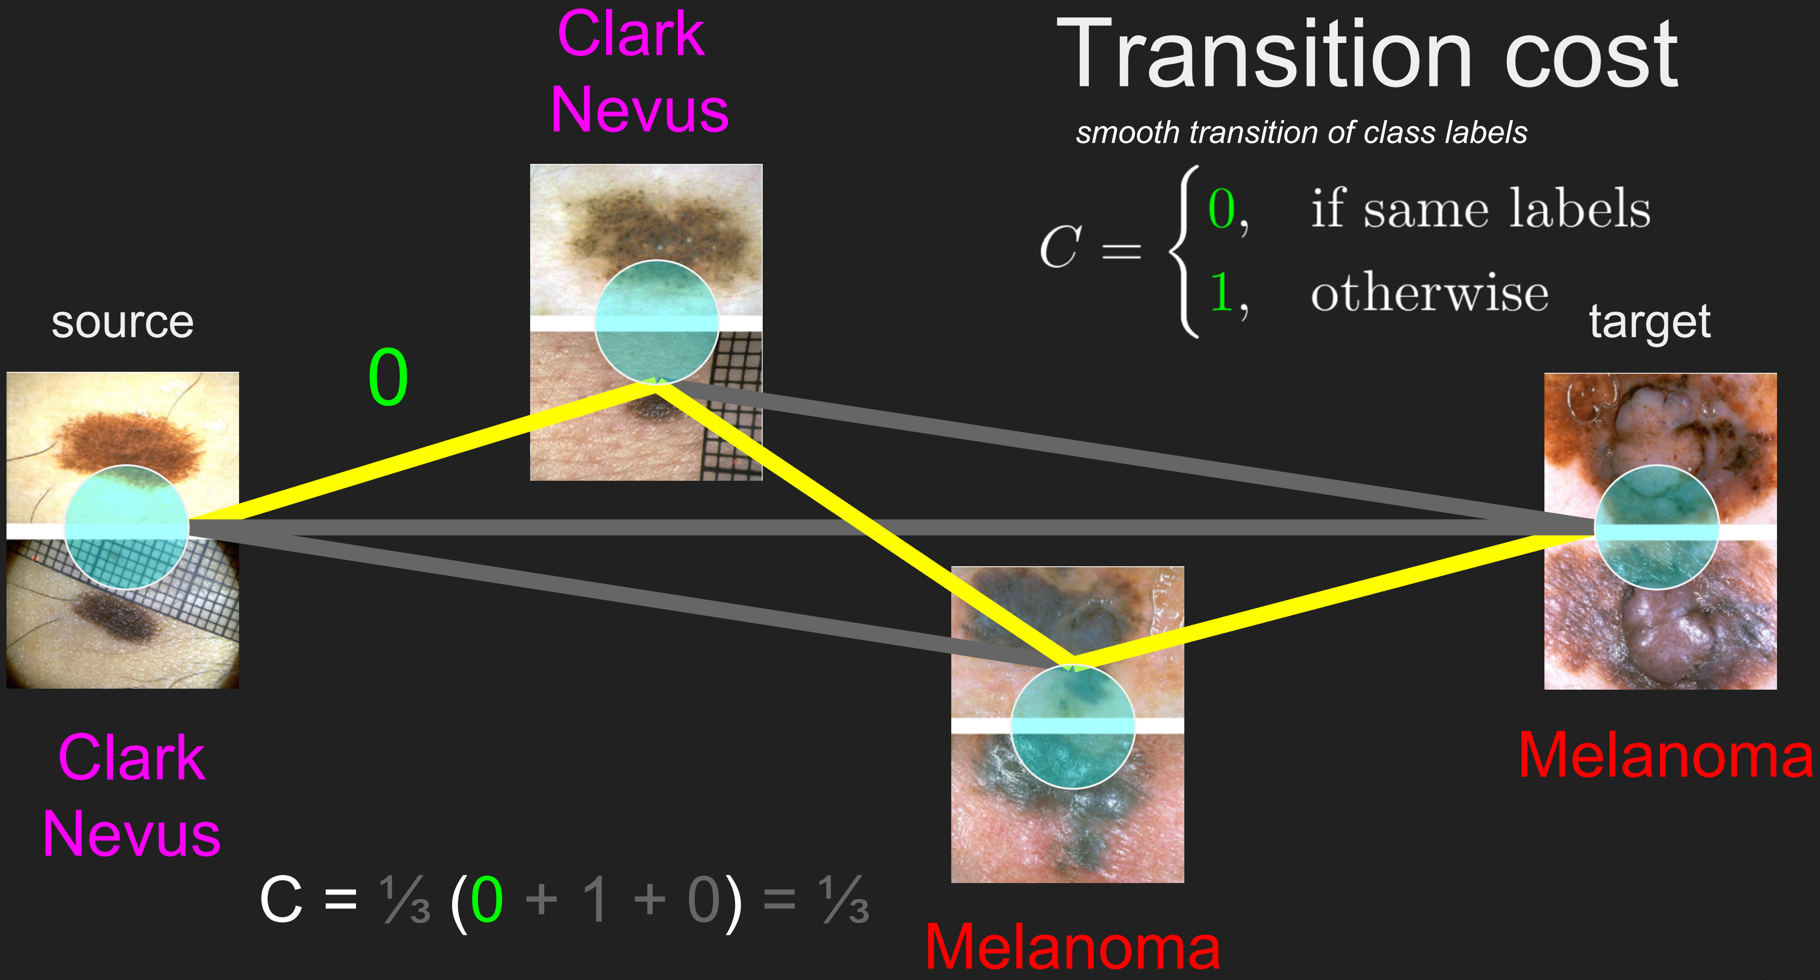

$C = ⅓ (0 + 1 + 0) = ⅓$

Transition cost

smooth transition of class labels

$$C = \begin{cases} 0, & \text{if same labels} \\ 1, & \text{otherwise} \end{cases}$$

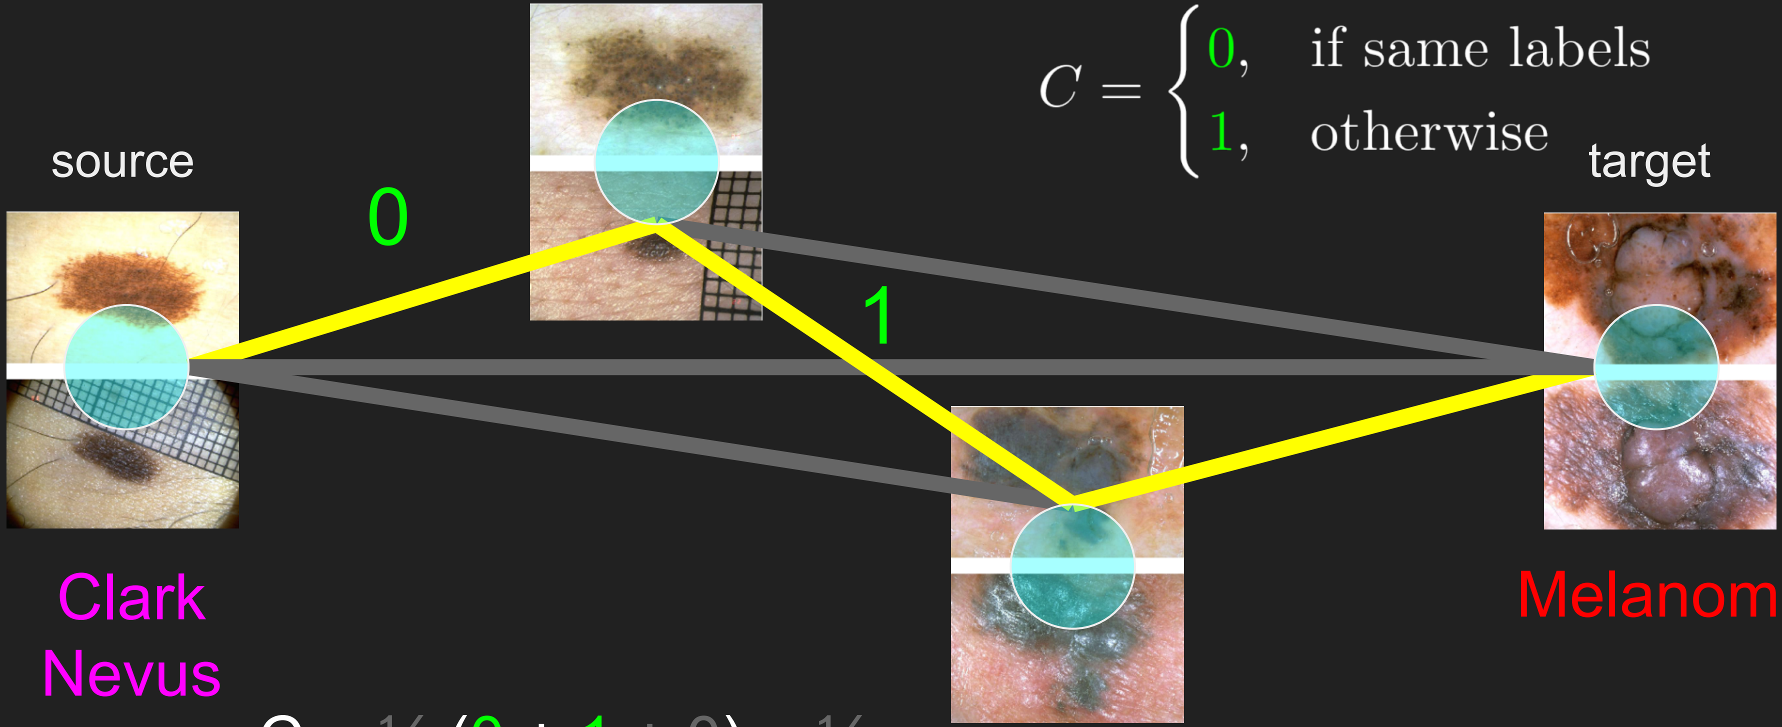

$$C = \tfrac{1}{3}\,(0 + 1 + 0) = \tfrac{1}{3}$$

Transition cost

smooth transition of class labels

$$C = \begin{cases} 0, & \text{if same labels} \\ 1, & \text{otherwise} \end{cases}$$

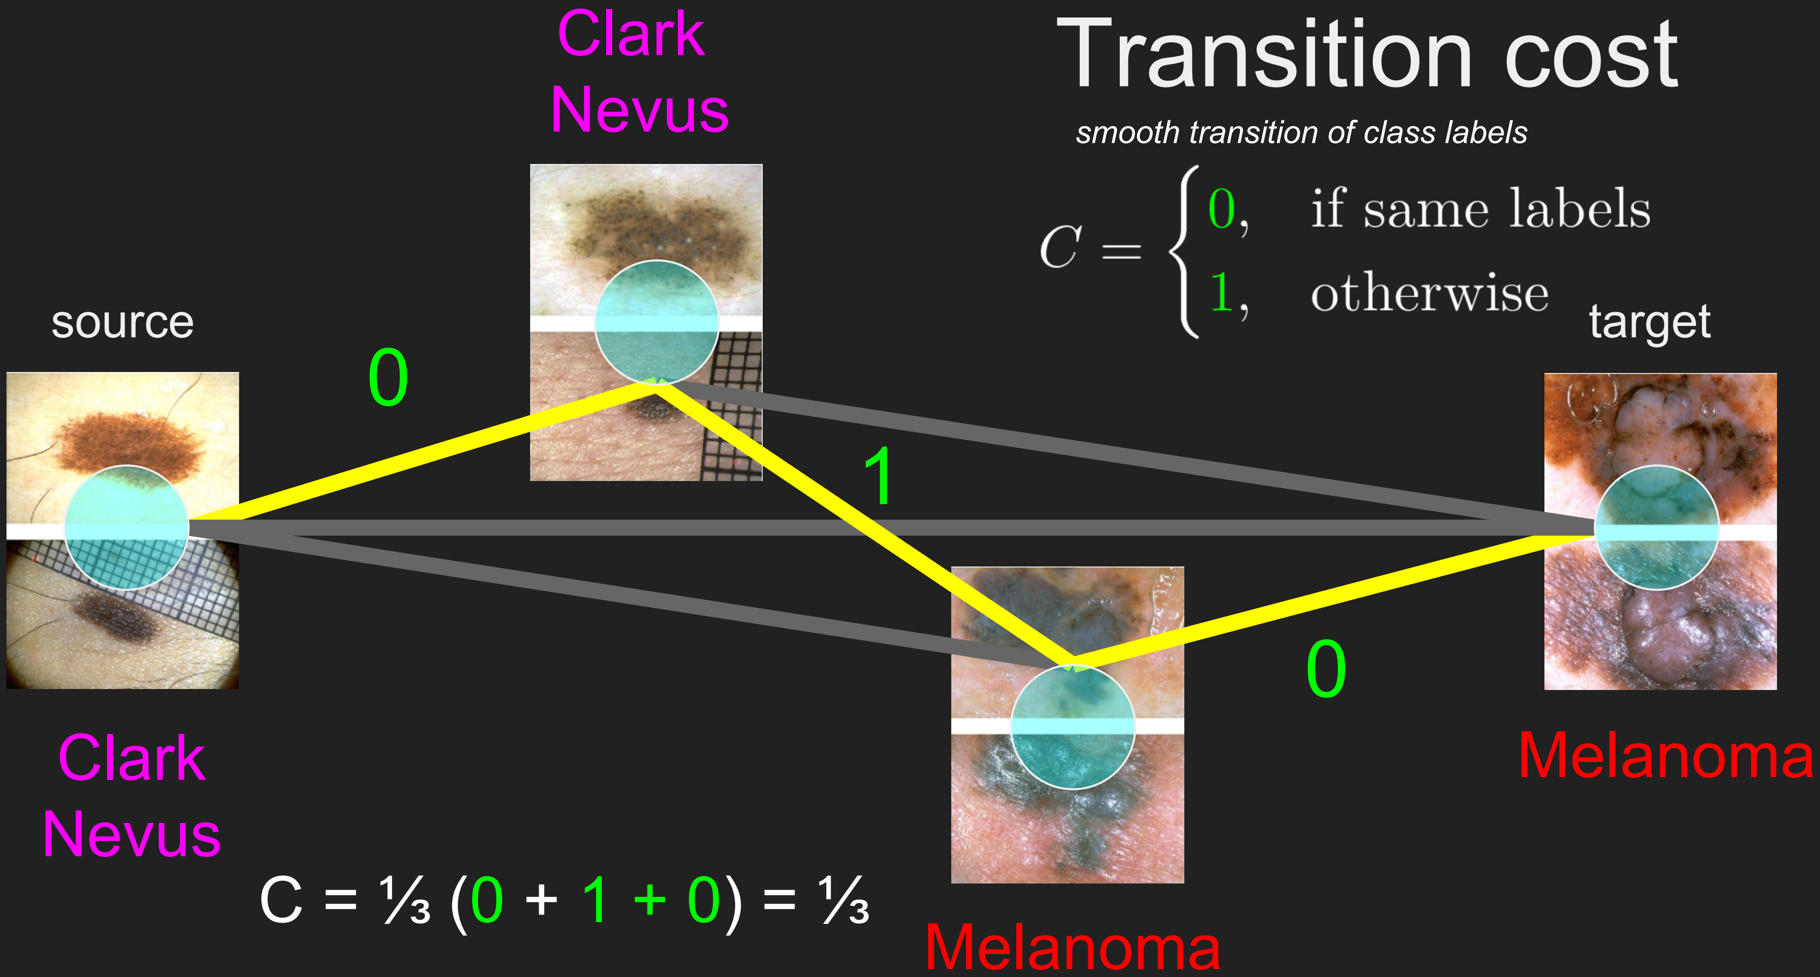

C = ⅓ (0 + 1 + 0) = ⅓

Progression cost

7-point score should consistently increase/decrease

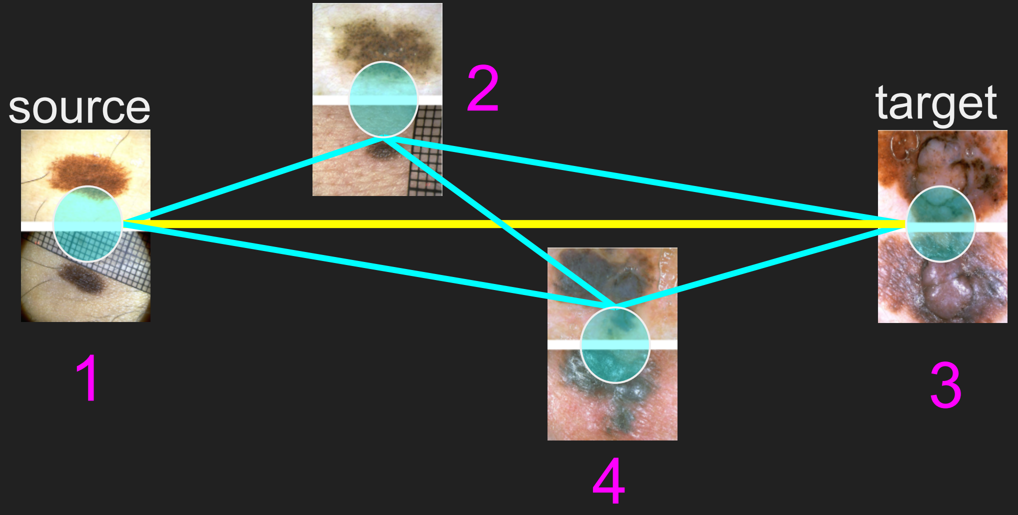

7-point scores

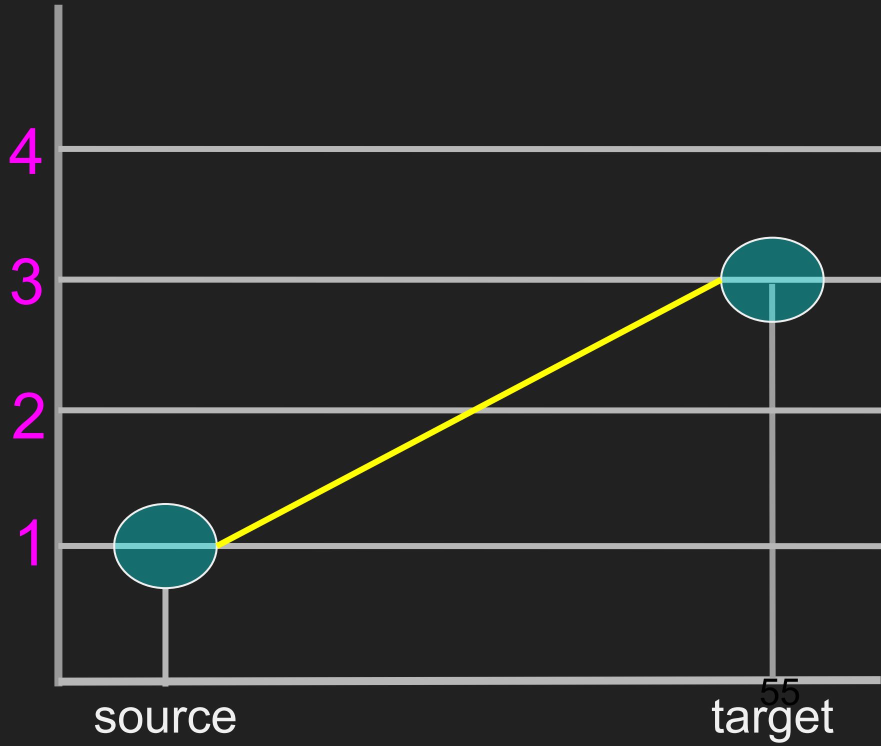

Progression cost

7-point score should consistently increase/decrease

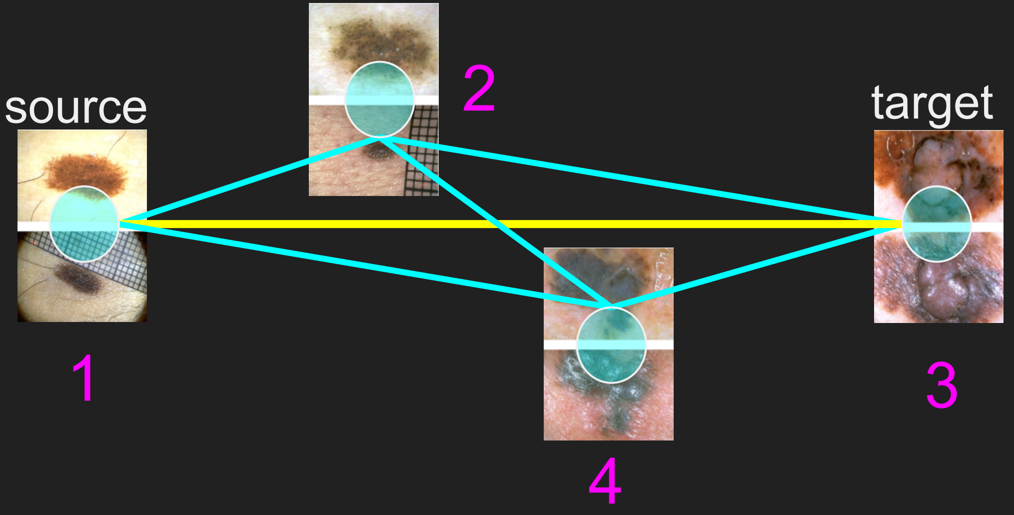

7-point scores

Increasing progression
between source and target

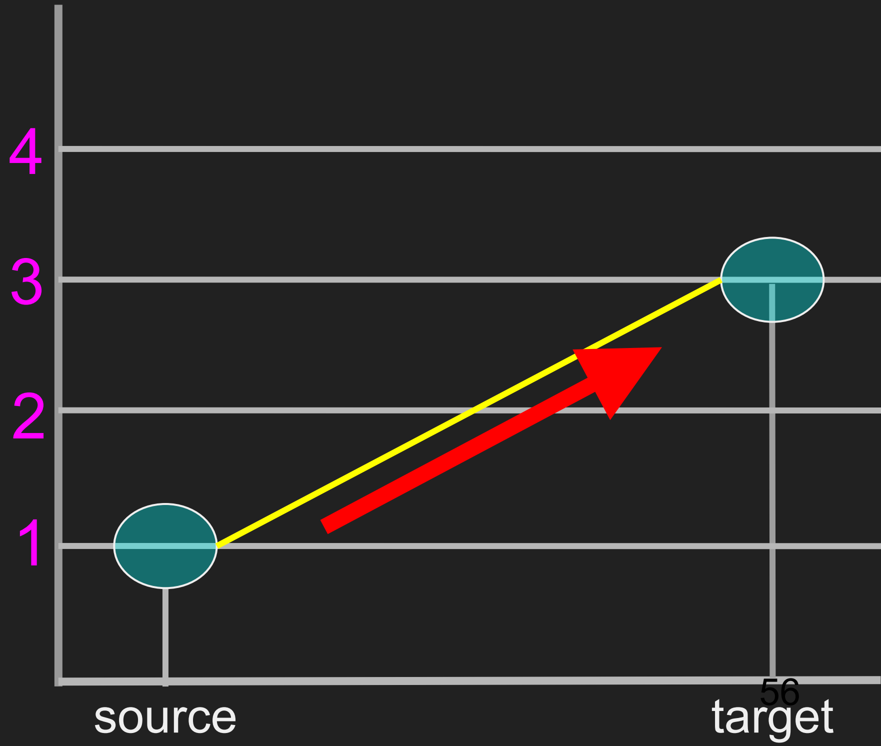

Progression cost

7-point score should consistently increase/decrease

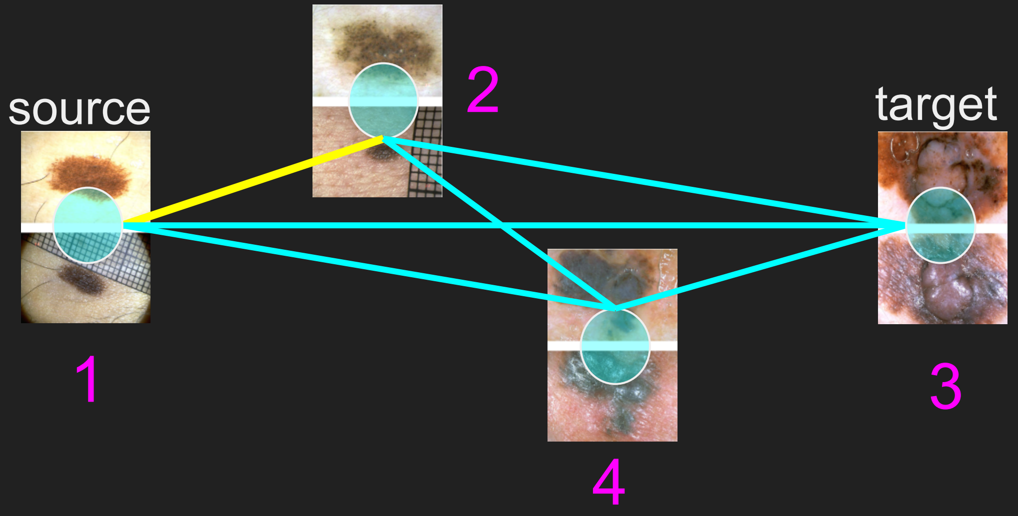

7-point scores

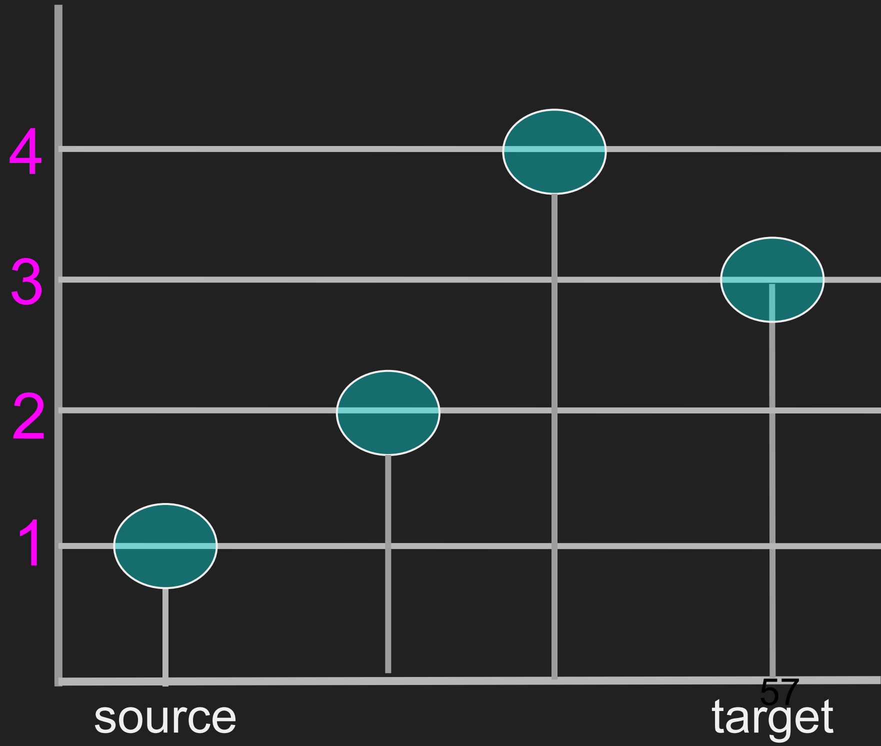

C =

Progression cost

7-point score should consistently increase/decrease

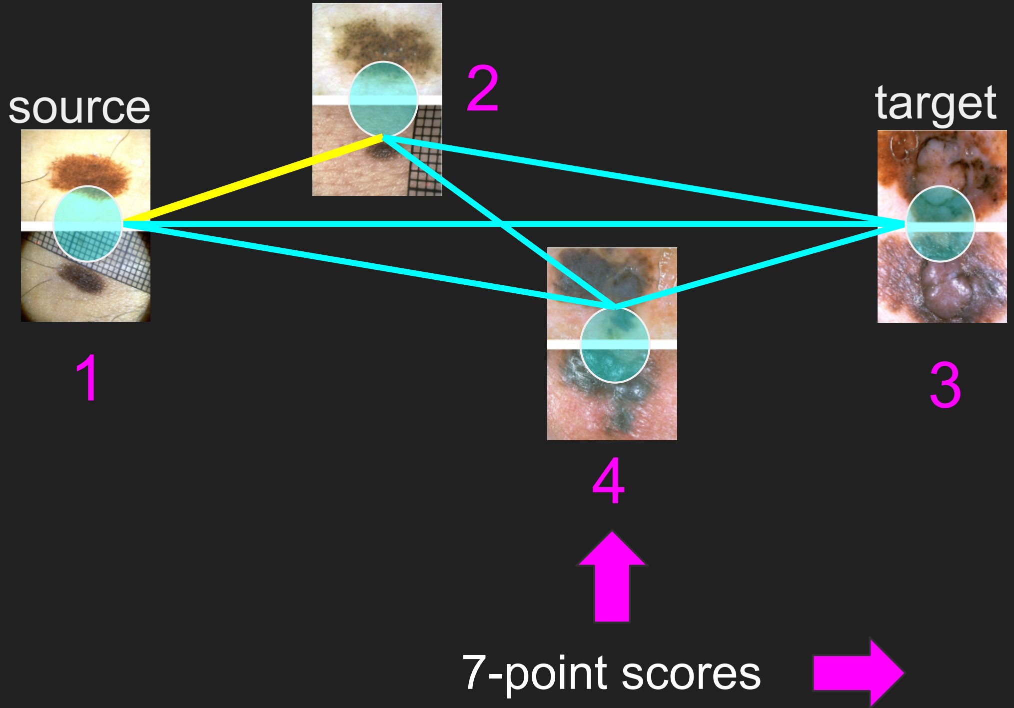

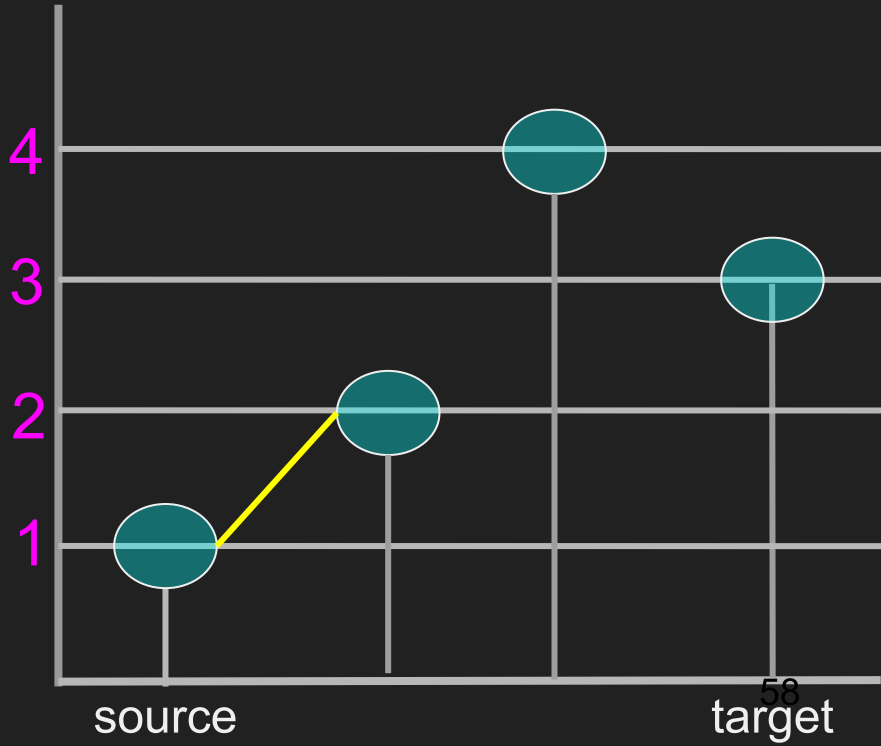

C = (0)

Progression cost

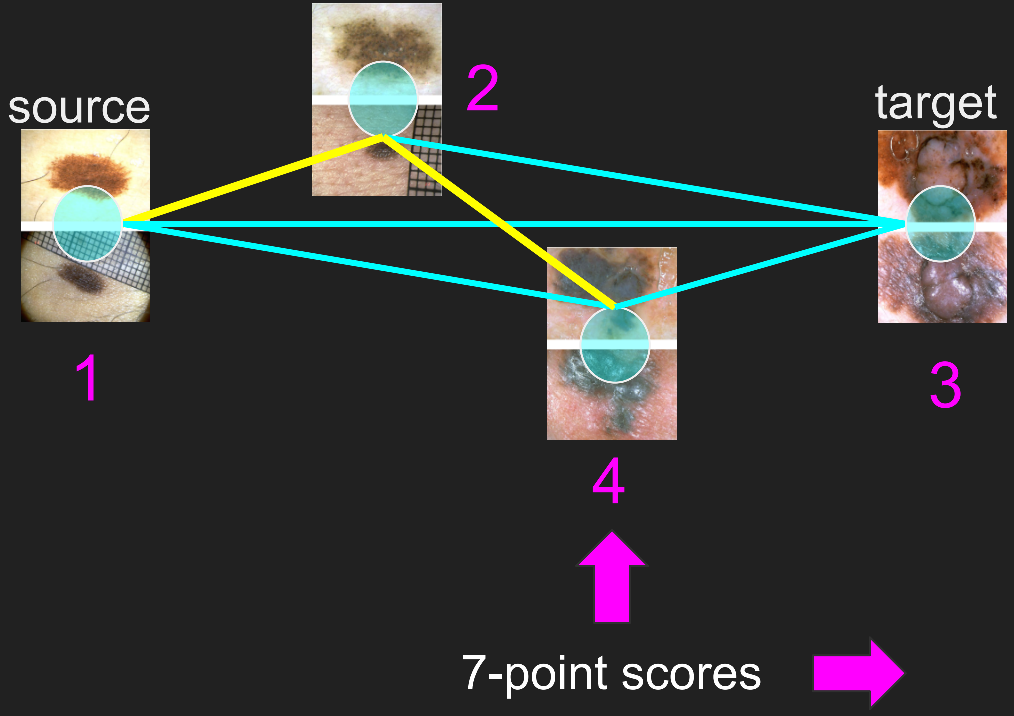

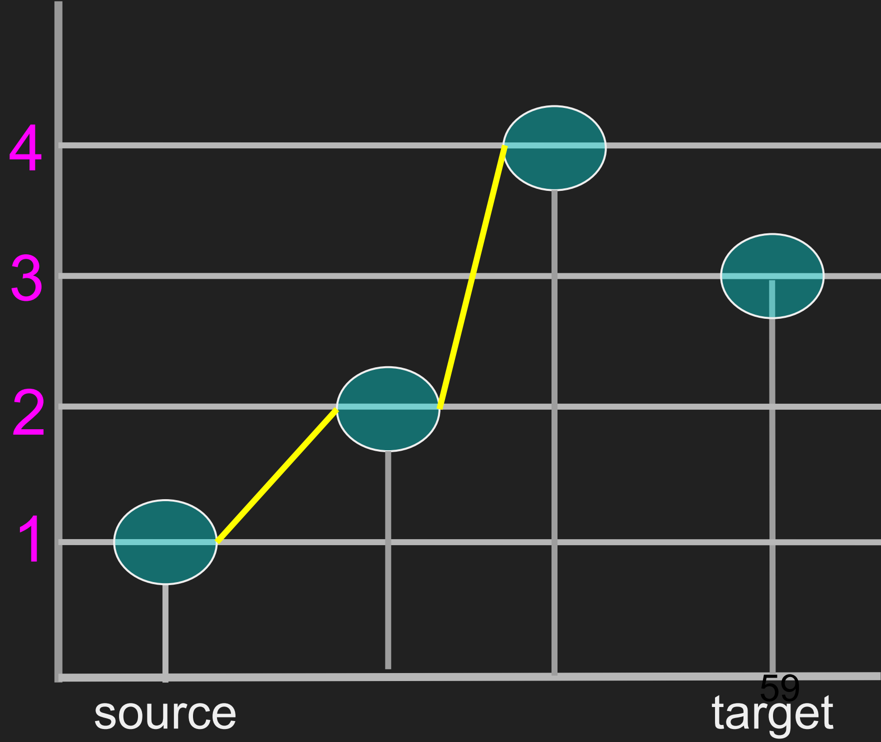

C = (0 + 0)

Progression cost

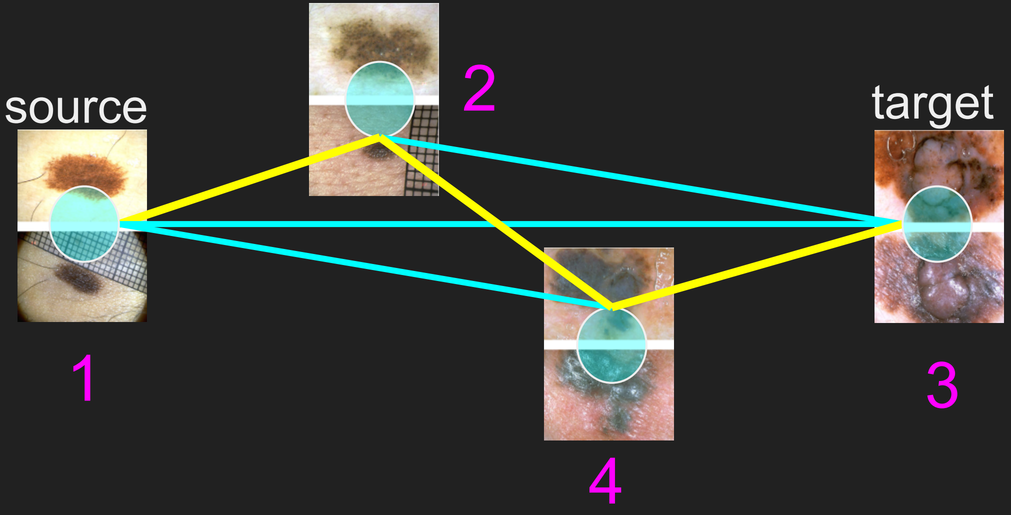

7-point scores

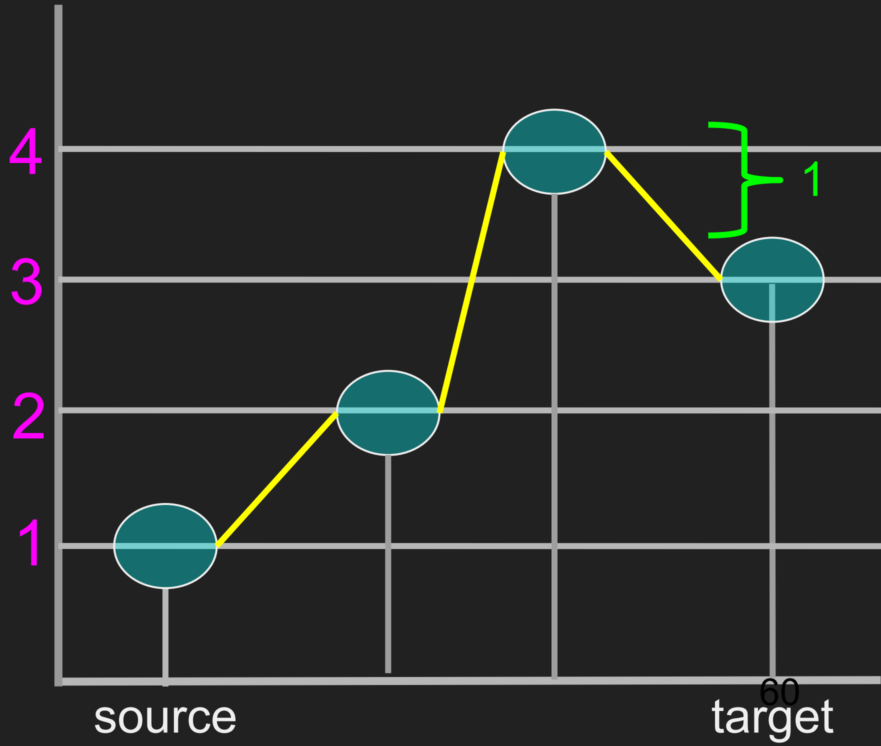

$C = \frac{1}{4}(0 + 0 + 1) = \frac{1}{4}$

Overview of our talk

1) How to encode images?

3) How to quantitatively evaluate the quality of paths?

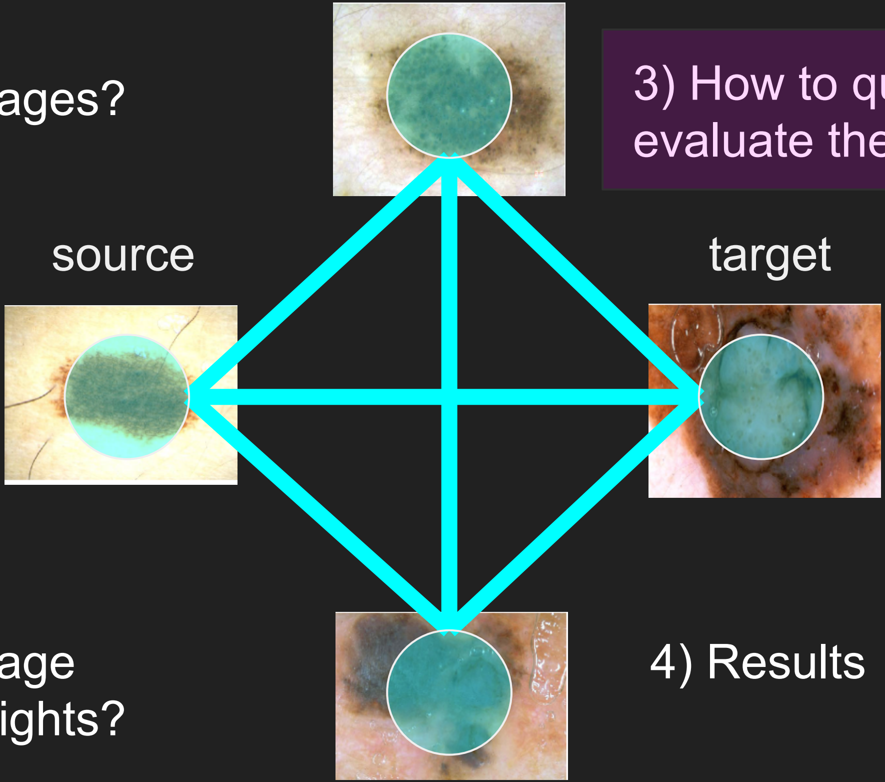

2) How to encode image dissimilarity/edge weights?

4) Results

Overview of our talk

1) How to encode images?

3) How to quantitatively evaluate the quality of paths?

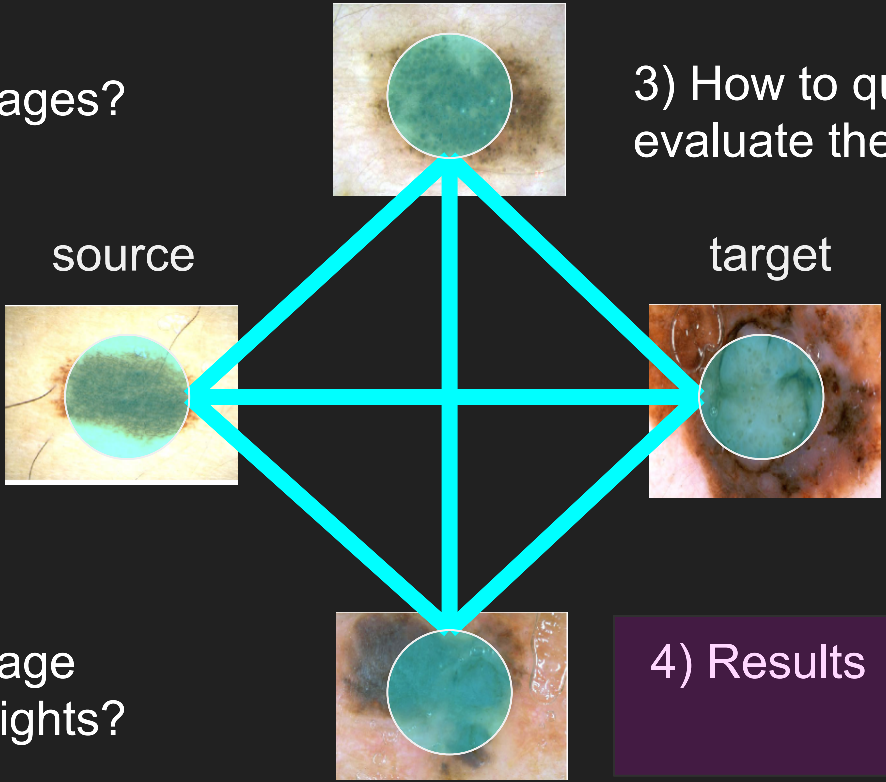

2) How to encode image dissimilarity/edge weights?

4) Results

Results (quantitative)

$$(2.1)^4 = 19.4$$

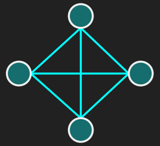

$$C = \begin{cases} 0, & \text{if same labels} \\ 1, & \text{otherwise} \end{cases}$$

Lower is better

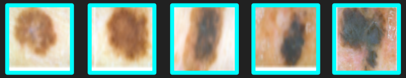

Image type	Exponential edge weights	Graph connectivity	Transition cost Mean (std. dev.)	Num. of path nodes Mean (std. dev.)
derm	no	complete	0.76 (0.42)	2.02 (0.13)
derm	no	30	0.64 (0.34)	3.59 (0.85)
derm	yes	complete	0.56 (0.26)	8.11 (2.87)
clinic	yes	30	0.65 (0.18)	10.64 (5.08)
derm/clinic	yes	30	**0.45 (0.24)**	7.90 (3.27)

Results (quantitative)

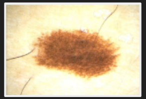

2.1

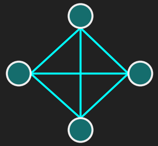

$$C = \begin{cases} 0, & \text{if same labels} \\ 1, & \text{otherwise} \end{cases}$$

Lower is better

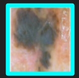

Image type	Exponential edge weights	Graph connectivity	Transition cost Mean (std. dev.)	Num. of path nodes Mean (std. dev.)
derm	no	complete	0.76 (0.42)	2.02 (0.13)
derm	no	30	0.64 (0.34)	3.59 (0.85)
derm	yes	complete	0.56 (0.26)	8.11 (2.87)
clinic	yes	30	0.65 (0.18)	10.64 (5.08)
derm/clinic	yes	30	**0.45 (0.24)**	7.90 (3.27)

A complete graph without exponential edge weights has a very short path

Results (quantitative)

$$C = \begin{cases} 0, & \text{if same labels} \\ 1, & \text{otherwise} \end{cases}$$

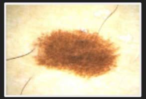

2.1

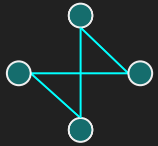

Lower is better

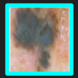

Image type	Exponential edge weights	Graph connectivity	Transition cost Mean (std. dev.)	Num. of path nodes Mean (std. dev.)
derm	no	complete	0.76 (0.42)	2.02 (0.13)
derm	no	30	0.64 (0.34)	3.59 (0.85)
derm	yes	complete	0.56 (0.26)	8.11 (2.87)
clinic	yes	30	0.65 (0.18)	10.64 (5.08)
derm/clinic	yes	30	**0.45 (0.24)**	7.90 (3.27)

Restricting graph connectivity results in slightly longer paths

Results (quantitative)

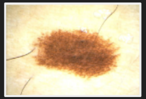

$$(2.1)^4 = 19.4$$

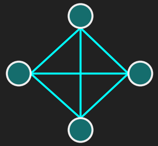

$$C = \begin{cases} 0, & \text{if same labels} \\ 1, & \text{otherwise} \end{cases}$$

Lower is better

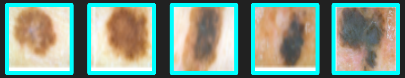

Image type	Exponential edge weights	Graph connectivity	Transition cost Mean (std. dev.)	Num. of path nodes Mean (std. dev.)
derm	no	complete	0.76 (0.42)	2.02 (0.13)
derm	no	30	0.64 (0.34)	3.59 (0.85)
derm	yes	complete	0.56 (0.26)	8.11 (2.87)
clinic	yes	30	0.65 (0.18)	10.64 (5.08)
derm/clinic	yes	30	**0.45 (0.24)**	7.90 (3.27)

Exponential edge weights yields longer paths even with a complete graph

Results (quantitative)

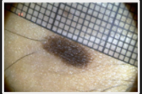

$$(2.1)^4 = 19.4$$

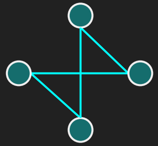

$$C = \begin{cases} 0, & \text{if same labels} \\ 1, & \text{otherwise} \end{cases}$$

Lower is better

Image type	Exponential edge weights	Graph connectivity	Transition cost Mean (std. dev.)	Num. of path nodes Mean (std. dev.)
derm	no	complete	0.76 (0.42)	2.02 (0.13)
derm	no	30	0.64 (0.34)	3.59 (0.85)
derm	yes	complete	0.56 (0.26)	8.11 (2.87)
clinic	yes	30	0.65 (0.18)	10.64 (5.08)
derm/clinic	yes	30	**0.45 (0.24)**	7.90 (3.27)

Clinical images score poorly on the transition costs

(as expected since clinical images contain more background clutter than dermoscopic images)

Results (quantitative)

$$(2.1)^4 = 19.4$$

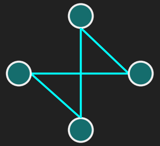

$$C = \begin{cases} 0, & \text{if same labels} \\ 1, & \text{otherwise} \end{cases}$$

Lower is better

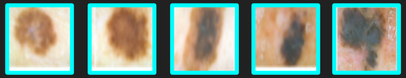

Image type	Exponential edge weights	Graph connectivity	Transition cost Mean (std. dev.)	Num. of path nodes Mean (std. dev.)
derm	no	complete	0.76 (0.42)	2.02 (0.13)
derm	no	30	0.64 (0.34)	3.59 (0.85)
derm	yes	complete	0.56 (0.26)	8.11 (2.87)
clinic	yes	30	0.65 (0.18)	10.64 (5.08)
derm/clinic	yes	30	**0.45 (0.24)**	7.90 (3.27)

Combined dermoscopic and clinical images have a low **transition cost**

Summary

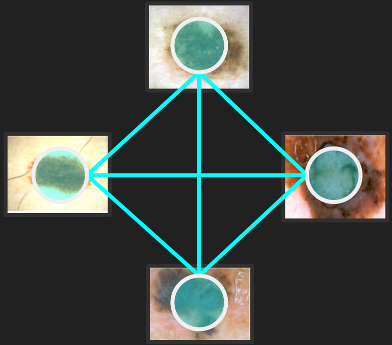

Graph geodesics (minimal path) to visualize skin lesions

Exponential multi-modal edge weights based on responses from a pretrained neural network

Proposed metrics to quantify the path quality

Progression cost

Transition cost

$$(2.1)^4 = 19.4$$

$$C = \begin{cases} 0, & \text{if same labels} \\ 1, & \text{otherwise} \end{cases}$$

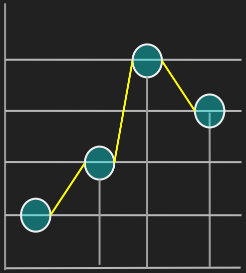

Kathleen: kmoriart@sfu.ca

Thank you!

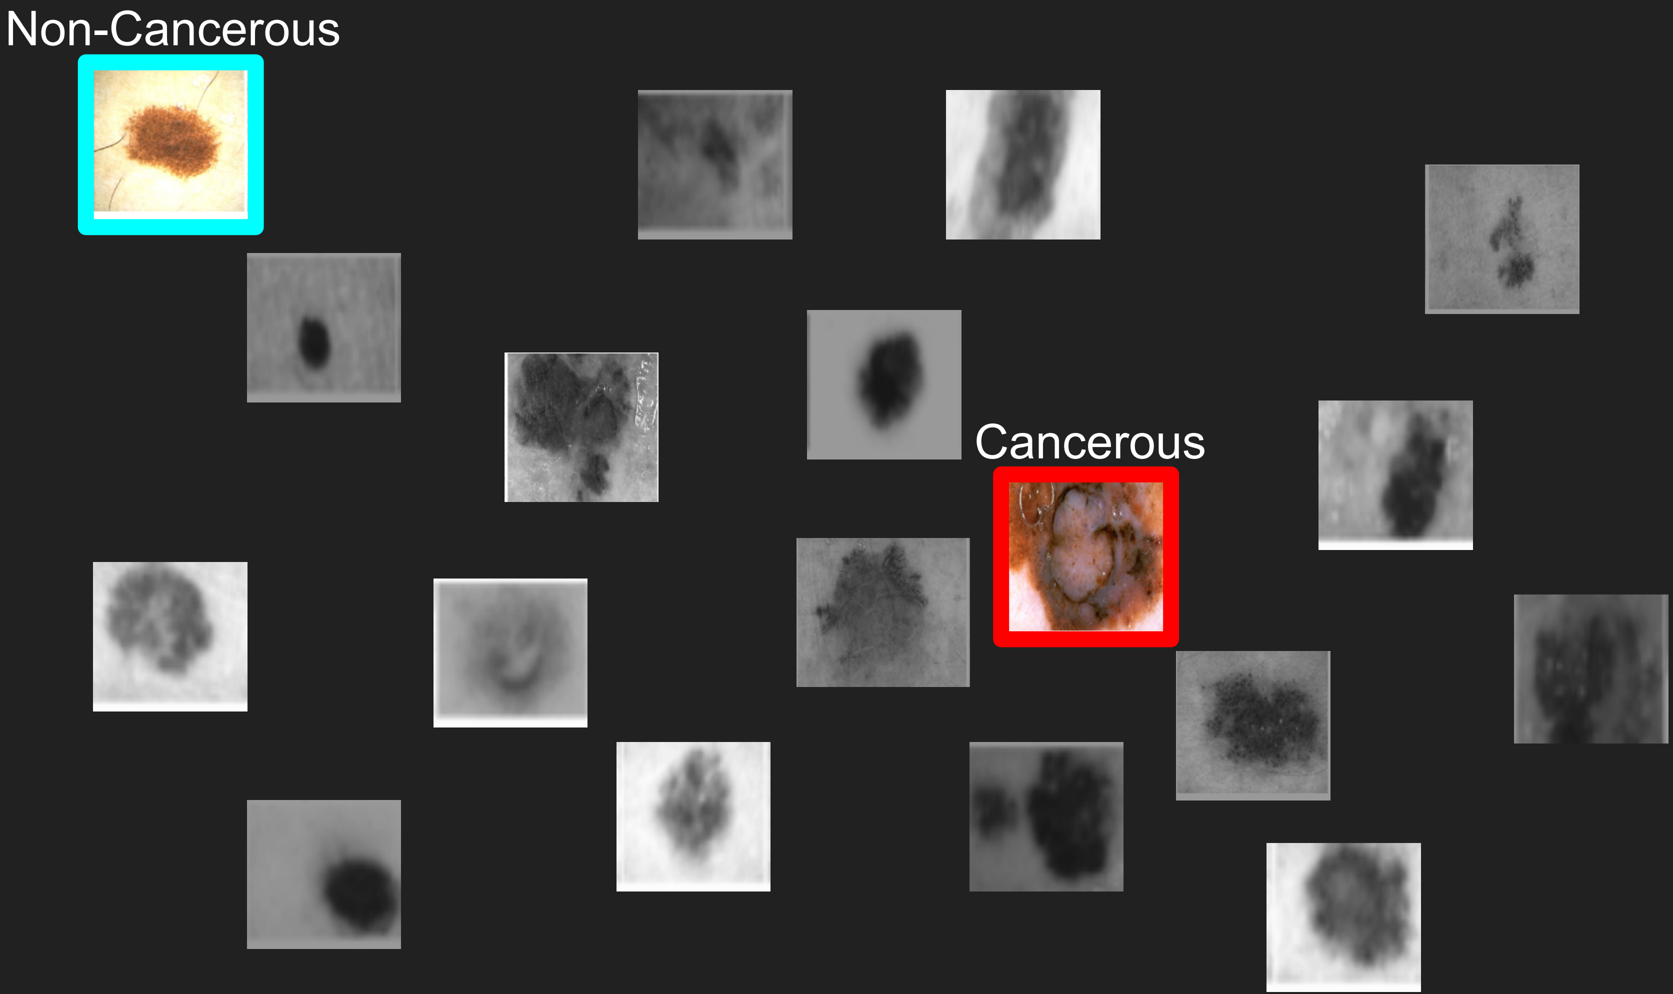
Non-Cancerous
Cancerous

Applying machine learning to skin lesion diagnosis

Image Retrieval K-Nearest Neighbours

Diagnosis can be **inferred** by inspecting the appearance of **similarly diseased images**.

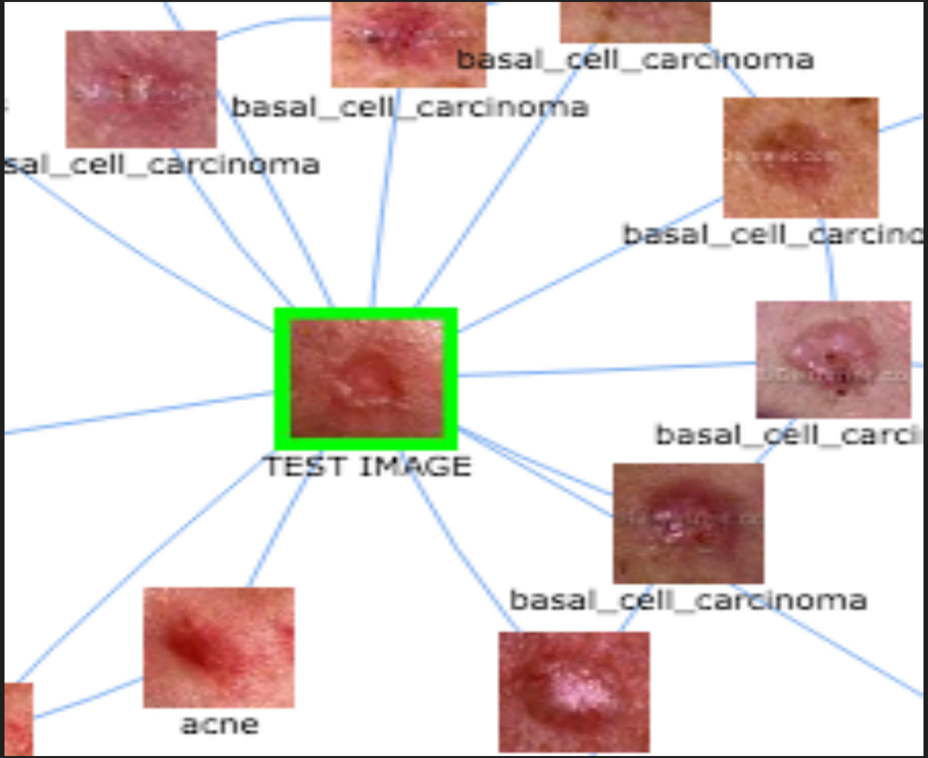

Results (quantitative)

$(2.1)^2 = 4.41$

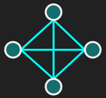

$$C = \begin{cases} 0, & \text{if same labels} \\ 1, & \text{otherwise} \end{cases}$$

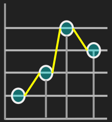

Lower is better

Lower is better

Image type	Exponential edge weights	Graph connectivity	Transition cost Mean (std. dev.)	Progress cost Mean (std. dev.)	Num. Nodes Mean (std. dev.)
derm	no	complete	0.76 (0.42)	0.10 (0.19)	2.02 (0.13)
derm	no	30	0.64 (0.34)	0.23 (0.26)	3.59 (0.85)
derm	yes	complete	0.56 (0.26)	0.37 (0.20)	8.11 (2.87)
derm	yes	30	0.56 (0.26)	0.37 (0.20)	8.12 (2.87)
clinic	yes	30	0.65 (0.18)	0.46 (0.20)	10.64 (5.08)
derm/clinic	yes	30	0.45 (0.24)	0.34 (0.19)	7.90 (3.27)

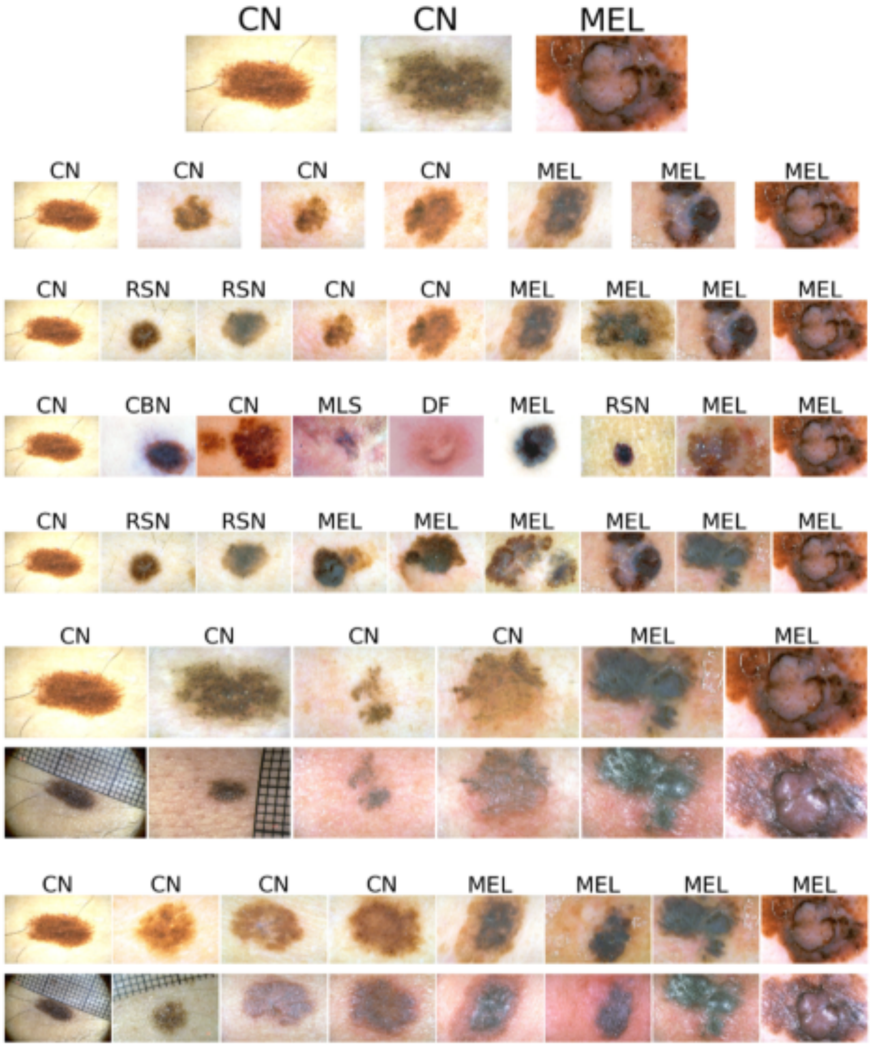
CN CN MEL
CN CN CN CN MEL MEL MEL
CN RSN RSN CN CN MEL MEL MEL MEL
CN CBN CN MLS DF MEL RSN MEL MEL
CN RSN RSN MEL MEL MEL MEL MEL MEL
CN CN CN CN MEL MEL
CN CN CN CN MEL MEL MEL MEL

Synthetic Examples