Graph Cuts, MRFs and Graphcut Textures

CS448V — Computational Video Manipulation

April 2019


What makes this a “texture”?

What makes this a “texture”?


stationary
What makes this a “texture”?

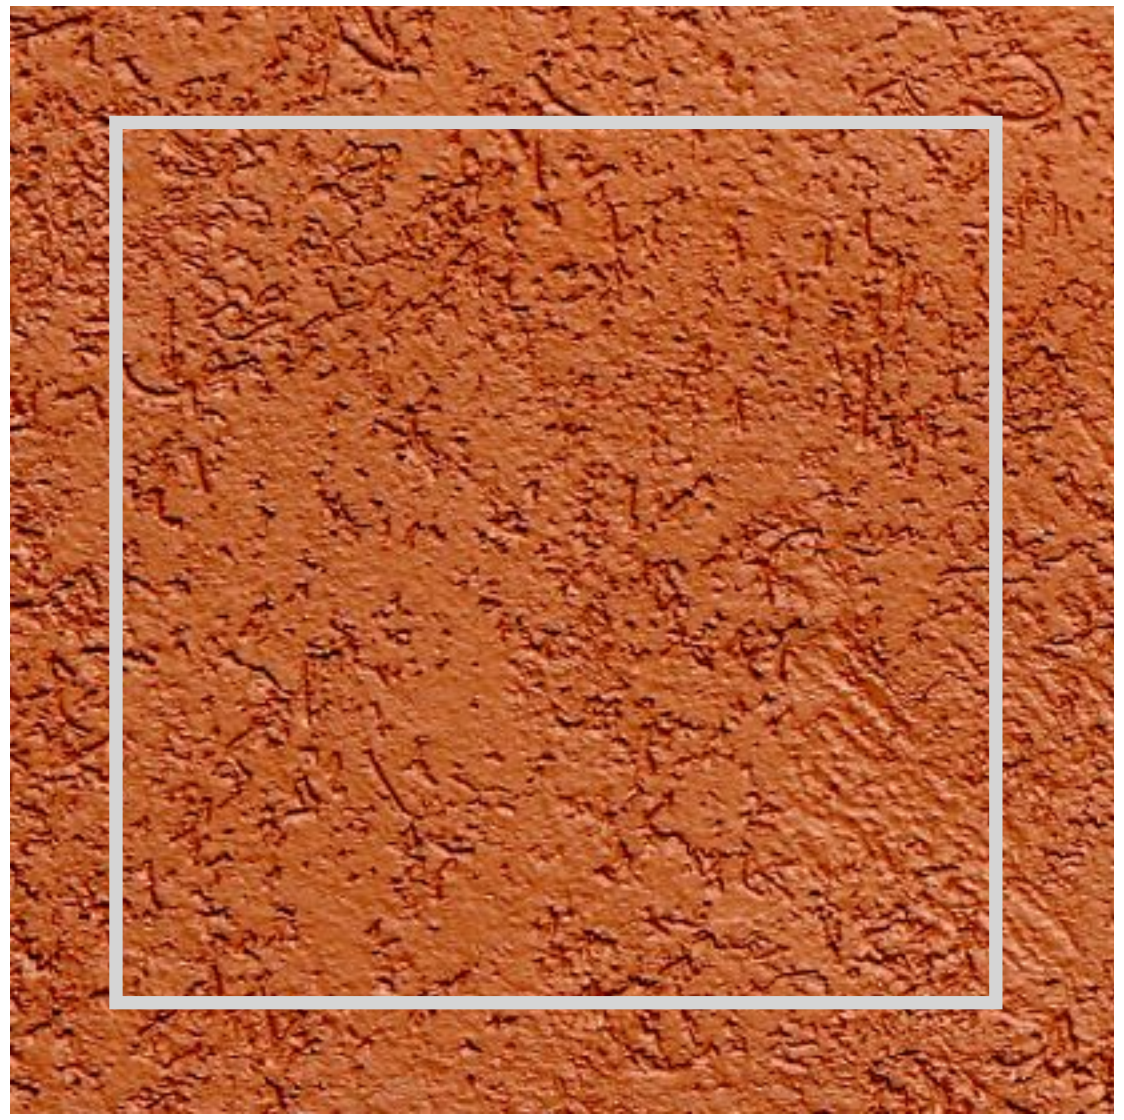

Local

What makes this a "texture"?

Stochastic

Regular

Stochastic

Regular

Texture?

Textures are everywhere!

105
AMMO
90%
HEALTH
2 3 4
5 6 7
ARMS
167%
ARMOR
BULL 108 / 400
SHEL 100 / 100
RCKT 15 / 100
CELL 105 / 600


Lives Pool
2
Time
00m 07s
Par Time
11m 00s
100
40

[Kwatra et al. 2003]


[Wexler et al. 2004]


[Wexler et al. 2004]

[Wexler et al. 2004]


[Wexler et al. 2004]

Graphcut Textures: Image and Video Synthesis Using Graph Cuts

Kwatra et al. 2003

Graphcut Textures: Image and Video Synthesis Using Graph Cuts

Kwatra et al. 2003

Graph Cuts

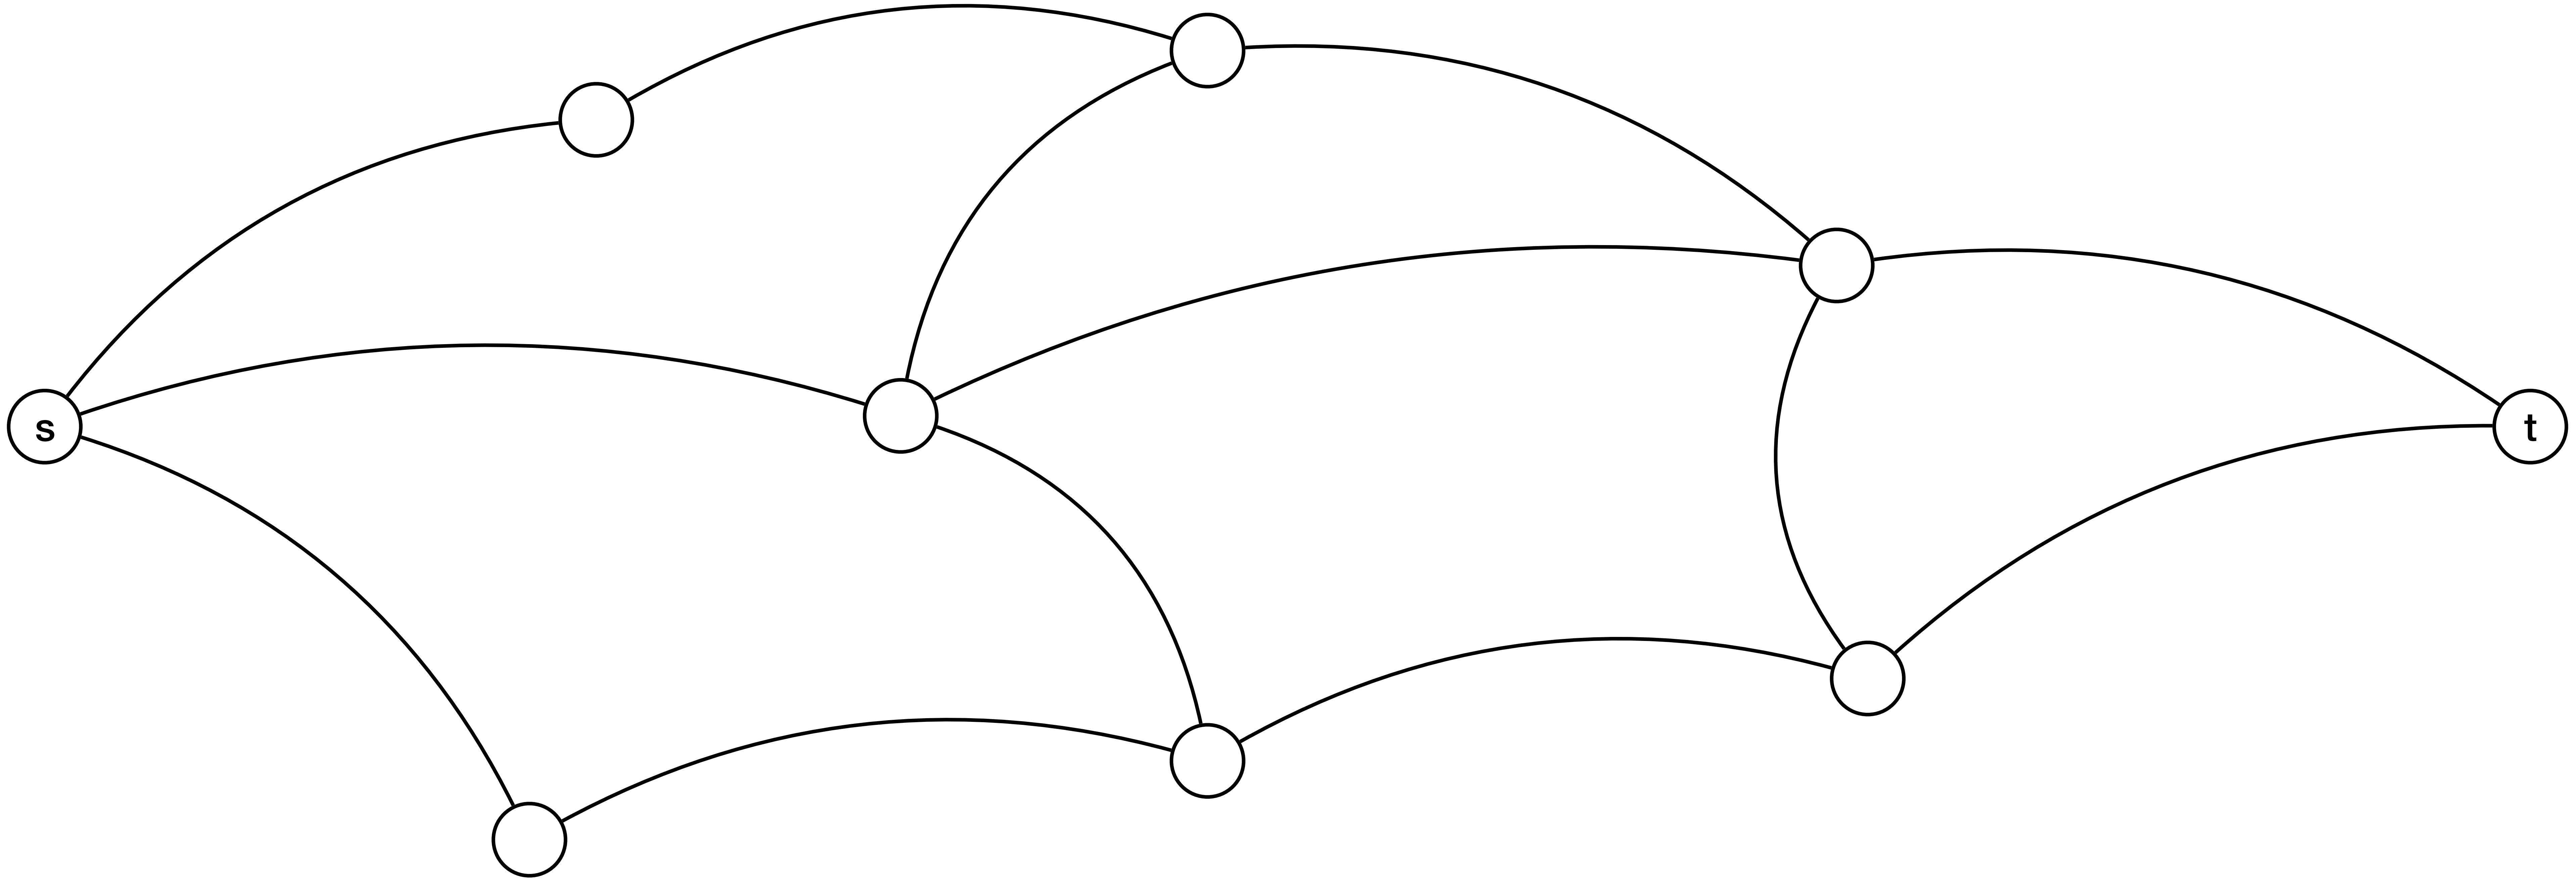

Graph Cuts

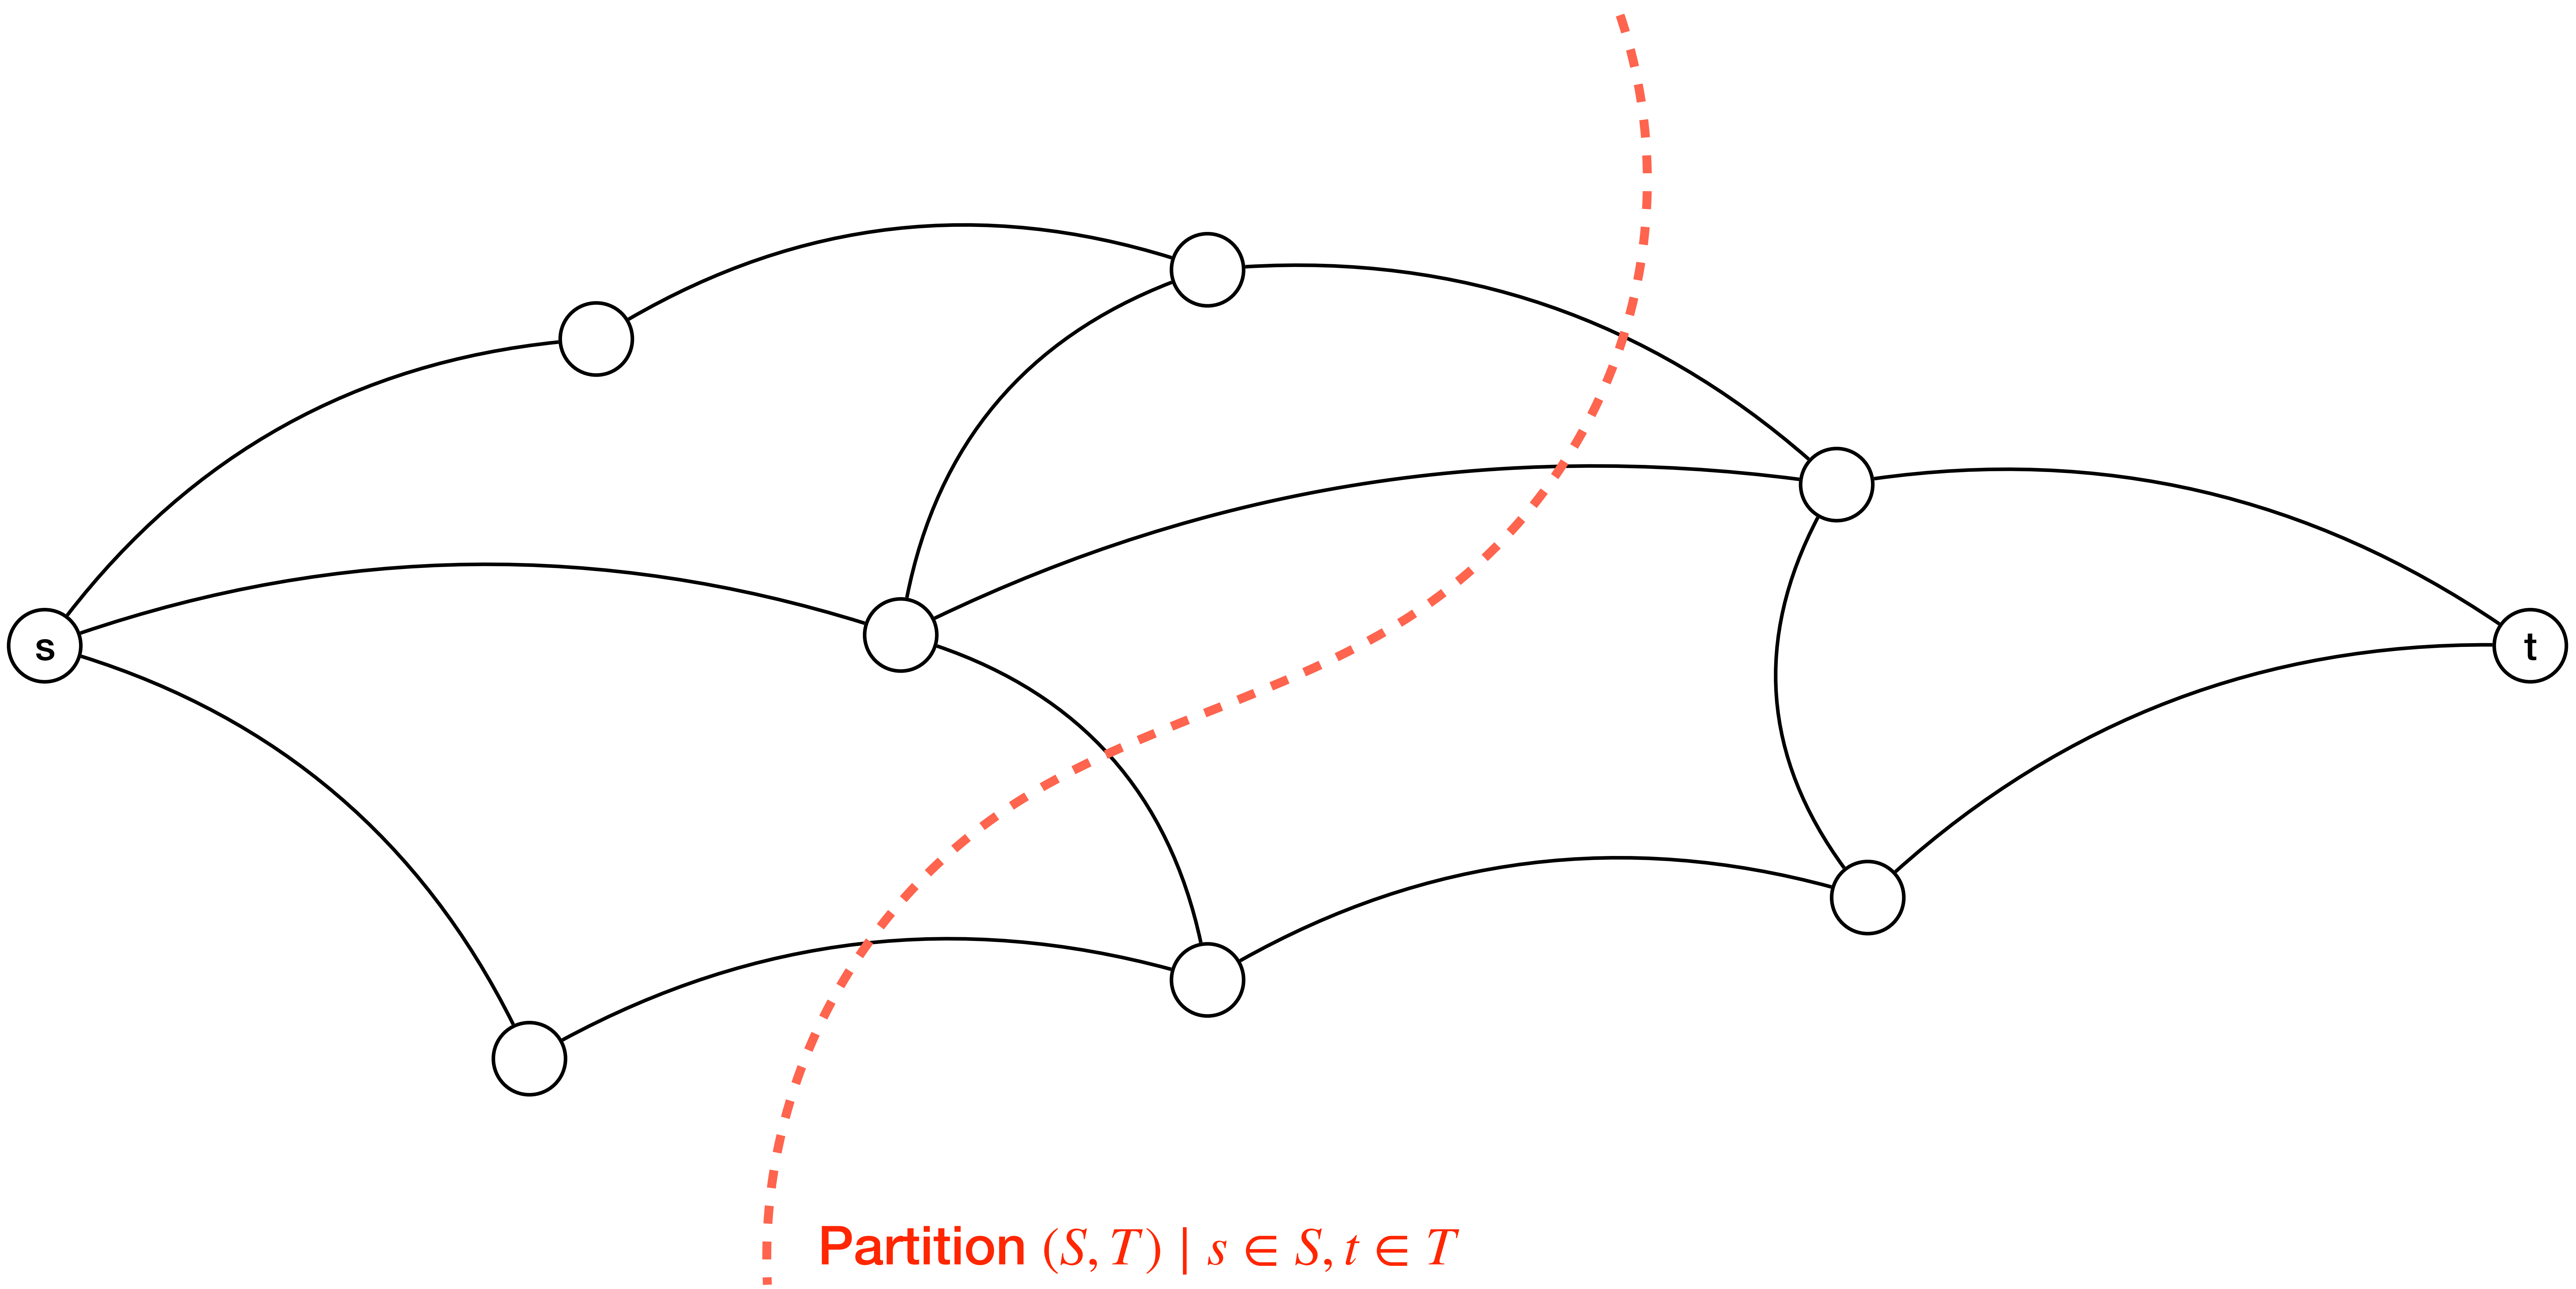

Partition $(S, T) \mid s \in S, t \in T$

Graph Cuts

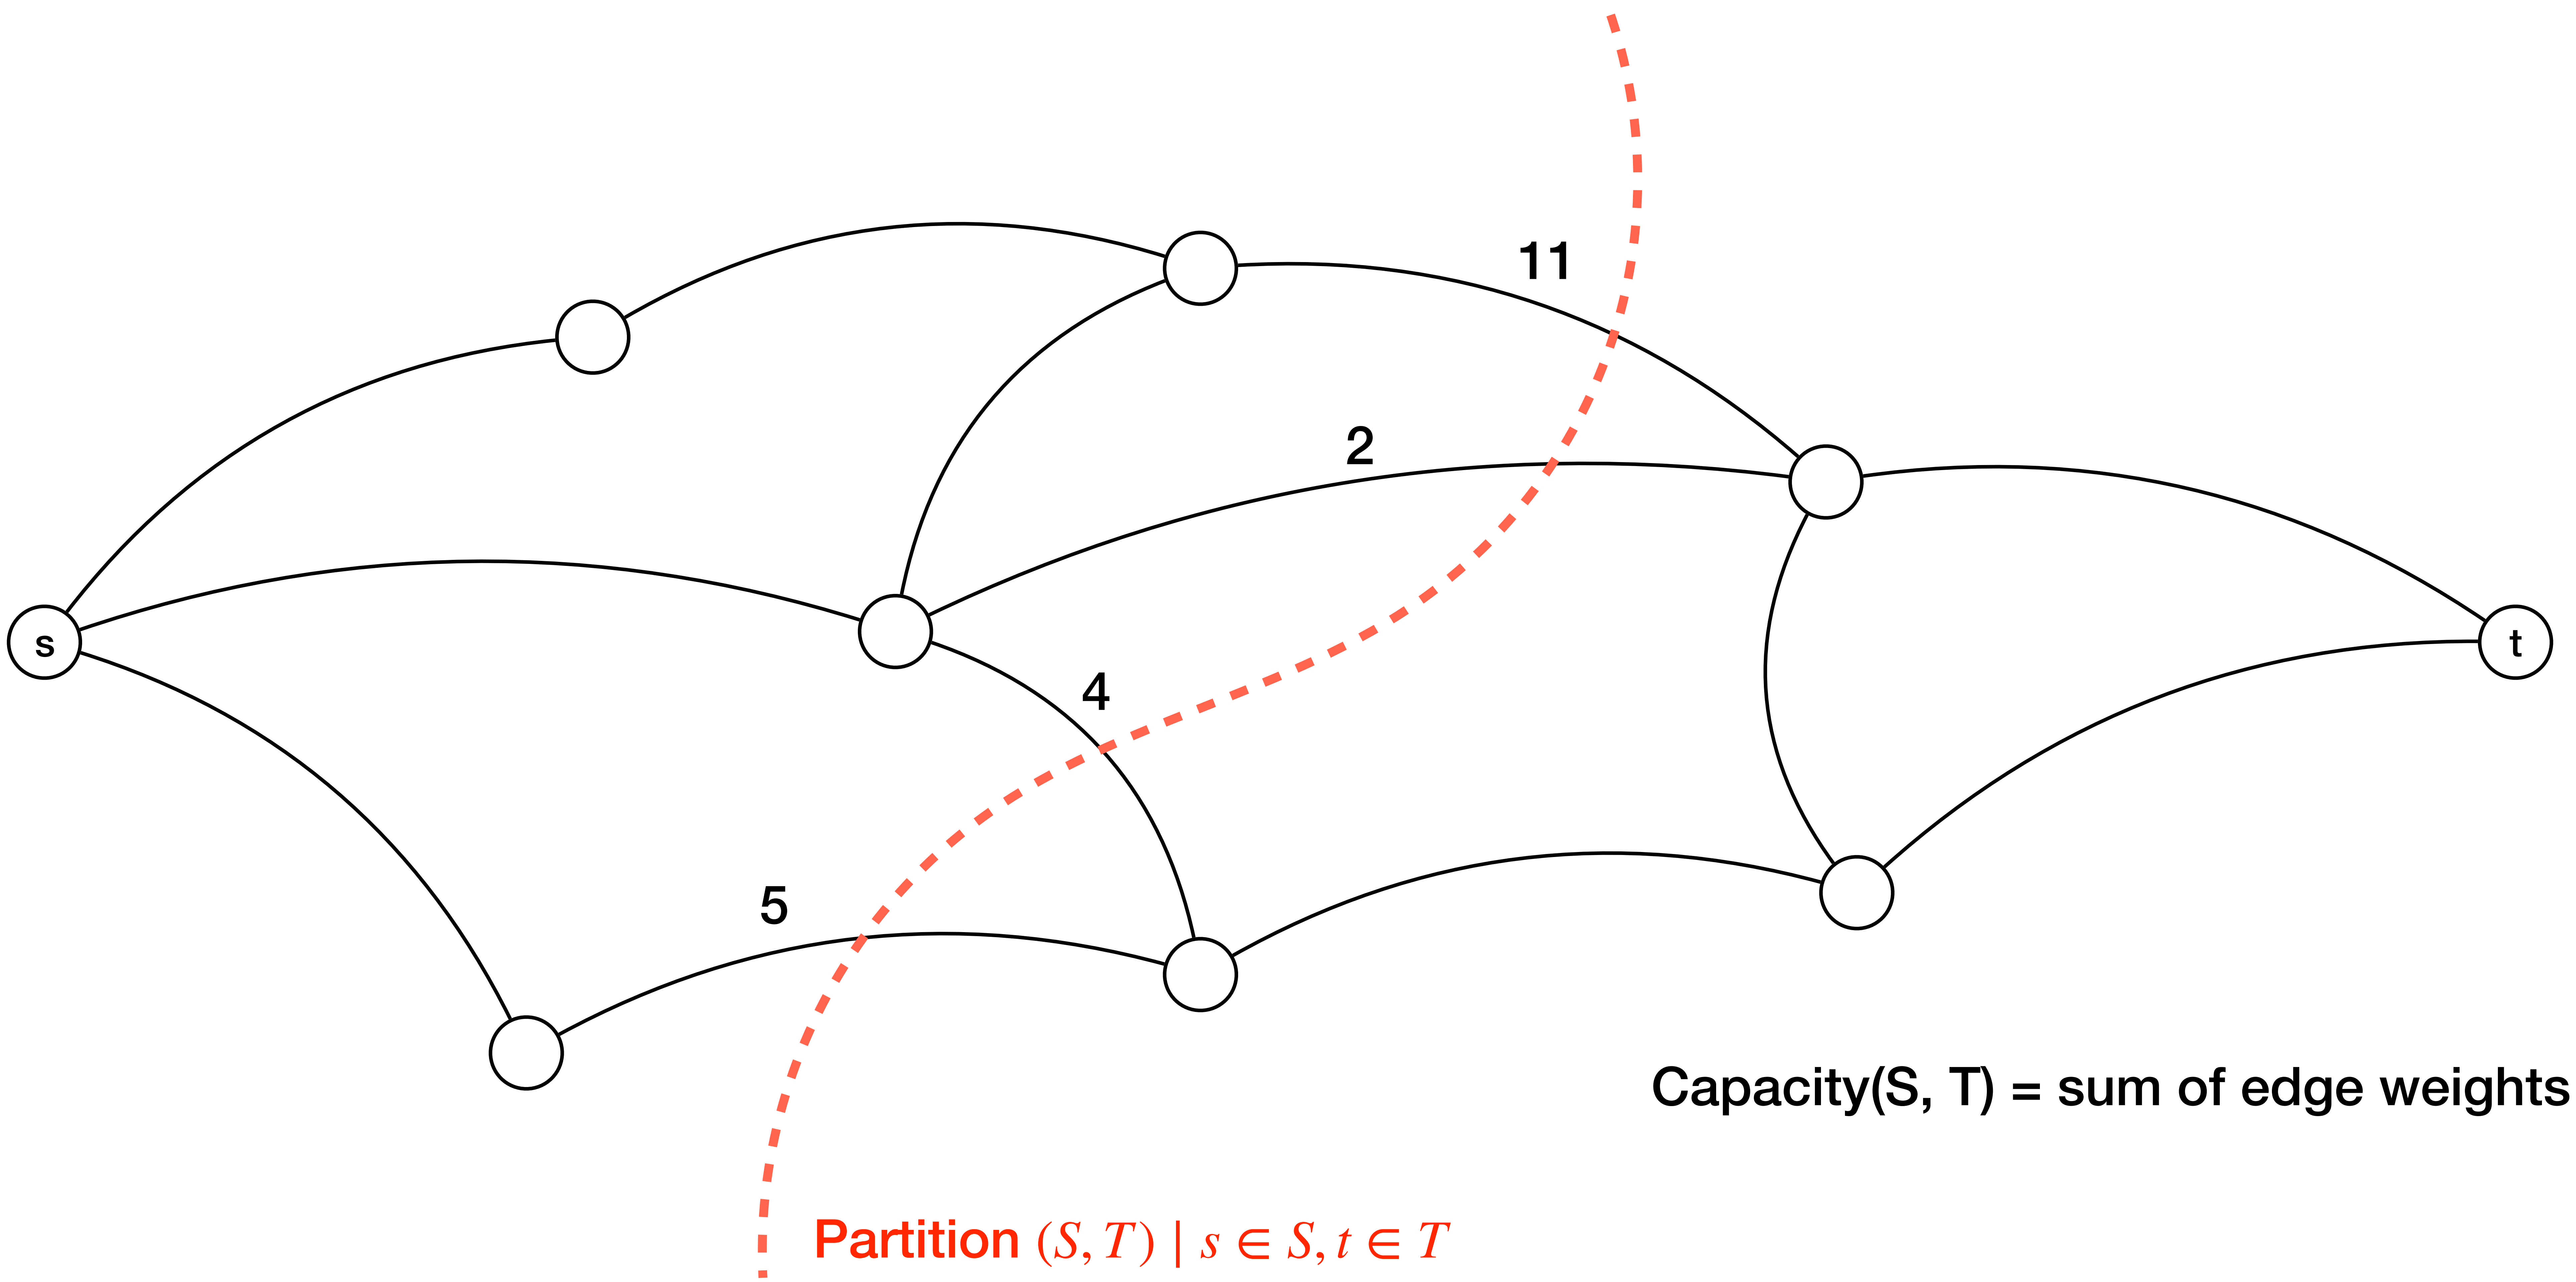

Capacity(S, T) = sum of edge weights (leaving) S

Max-flow Min-cut theorem

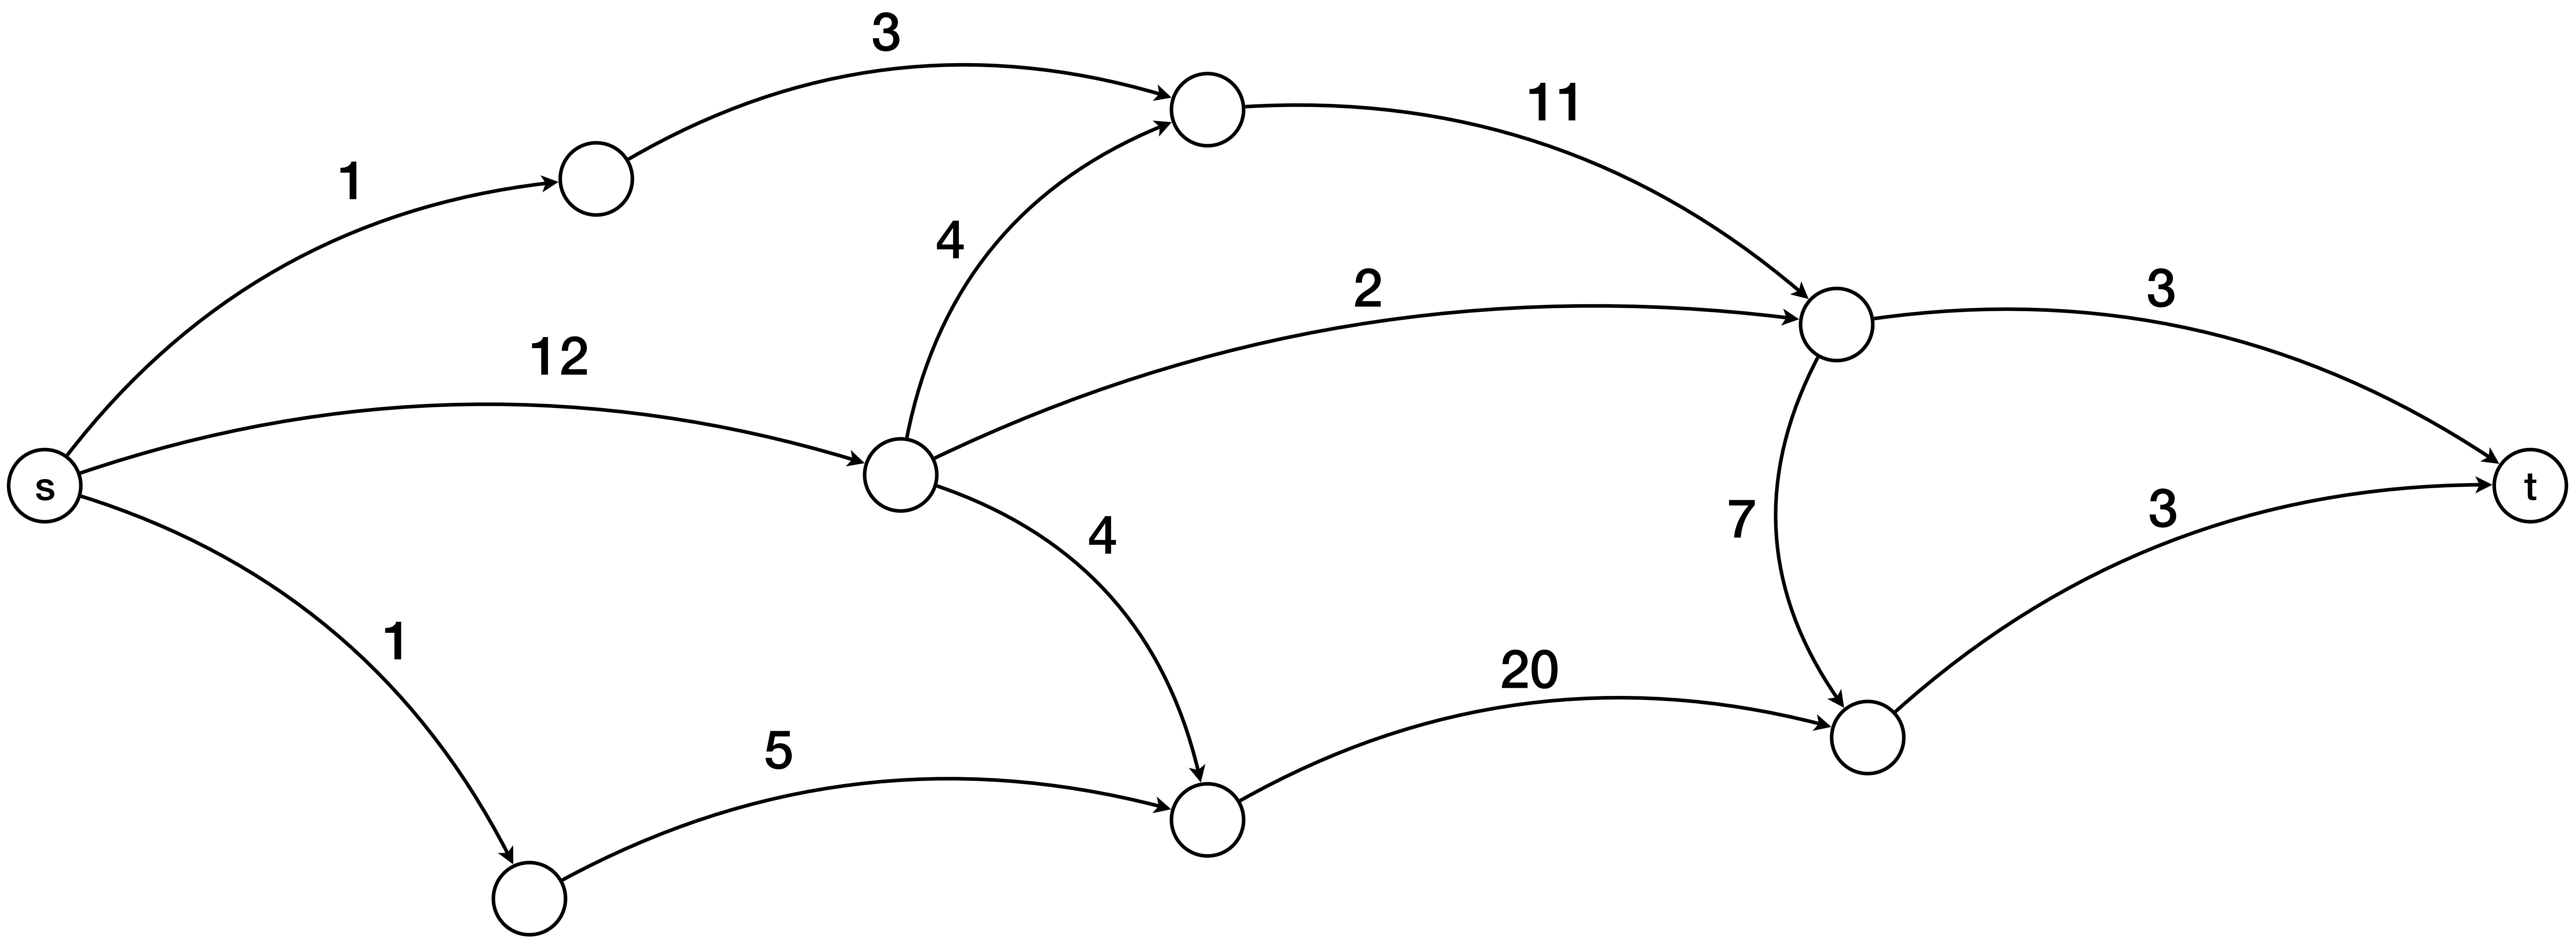

Max-flow Min-cut theorem

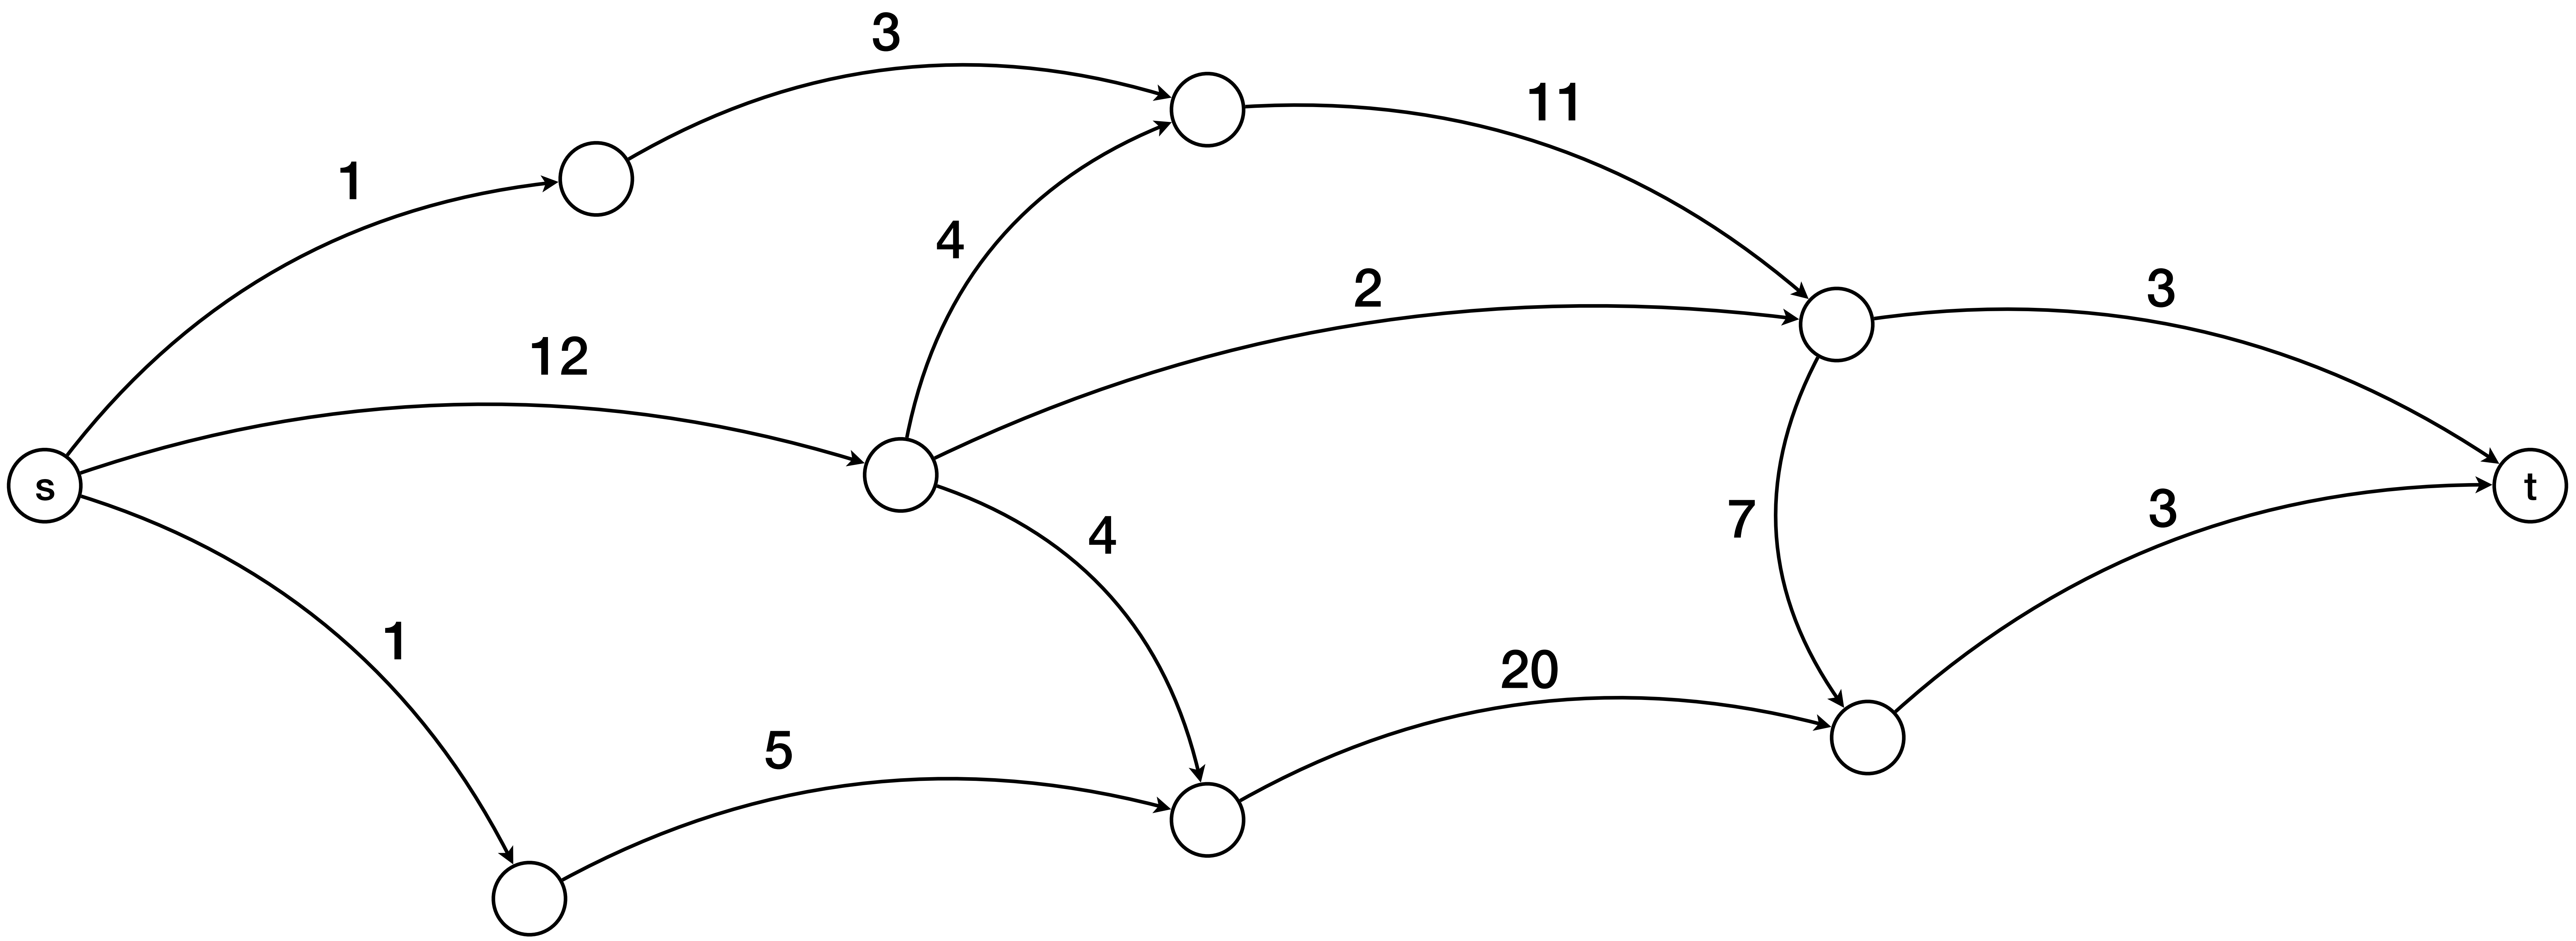

What is a flow?…

Max-flow Min-cut theorem

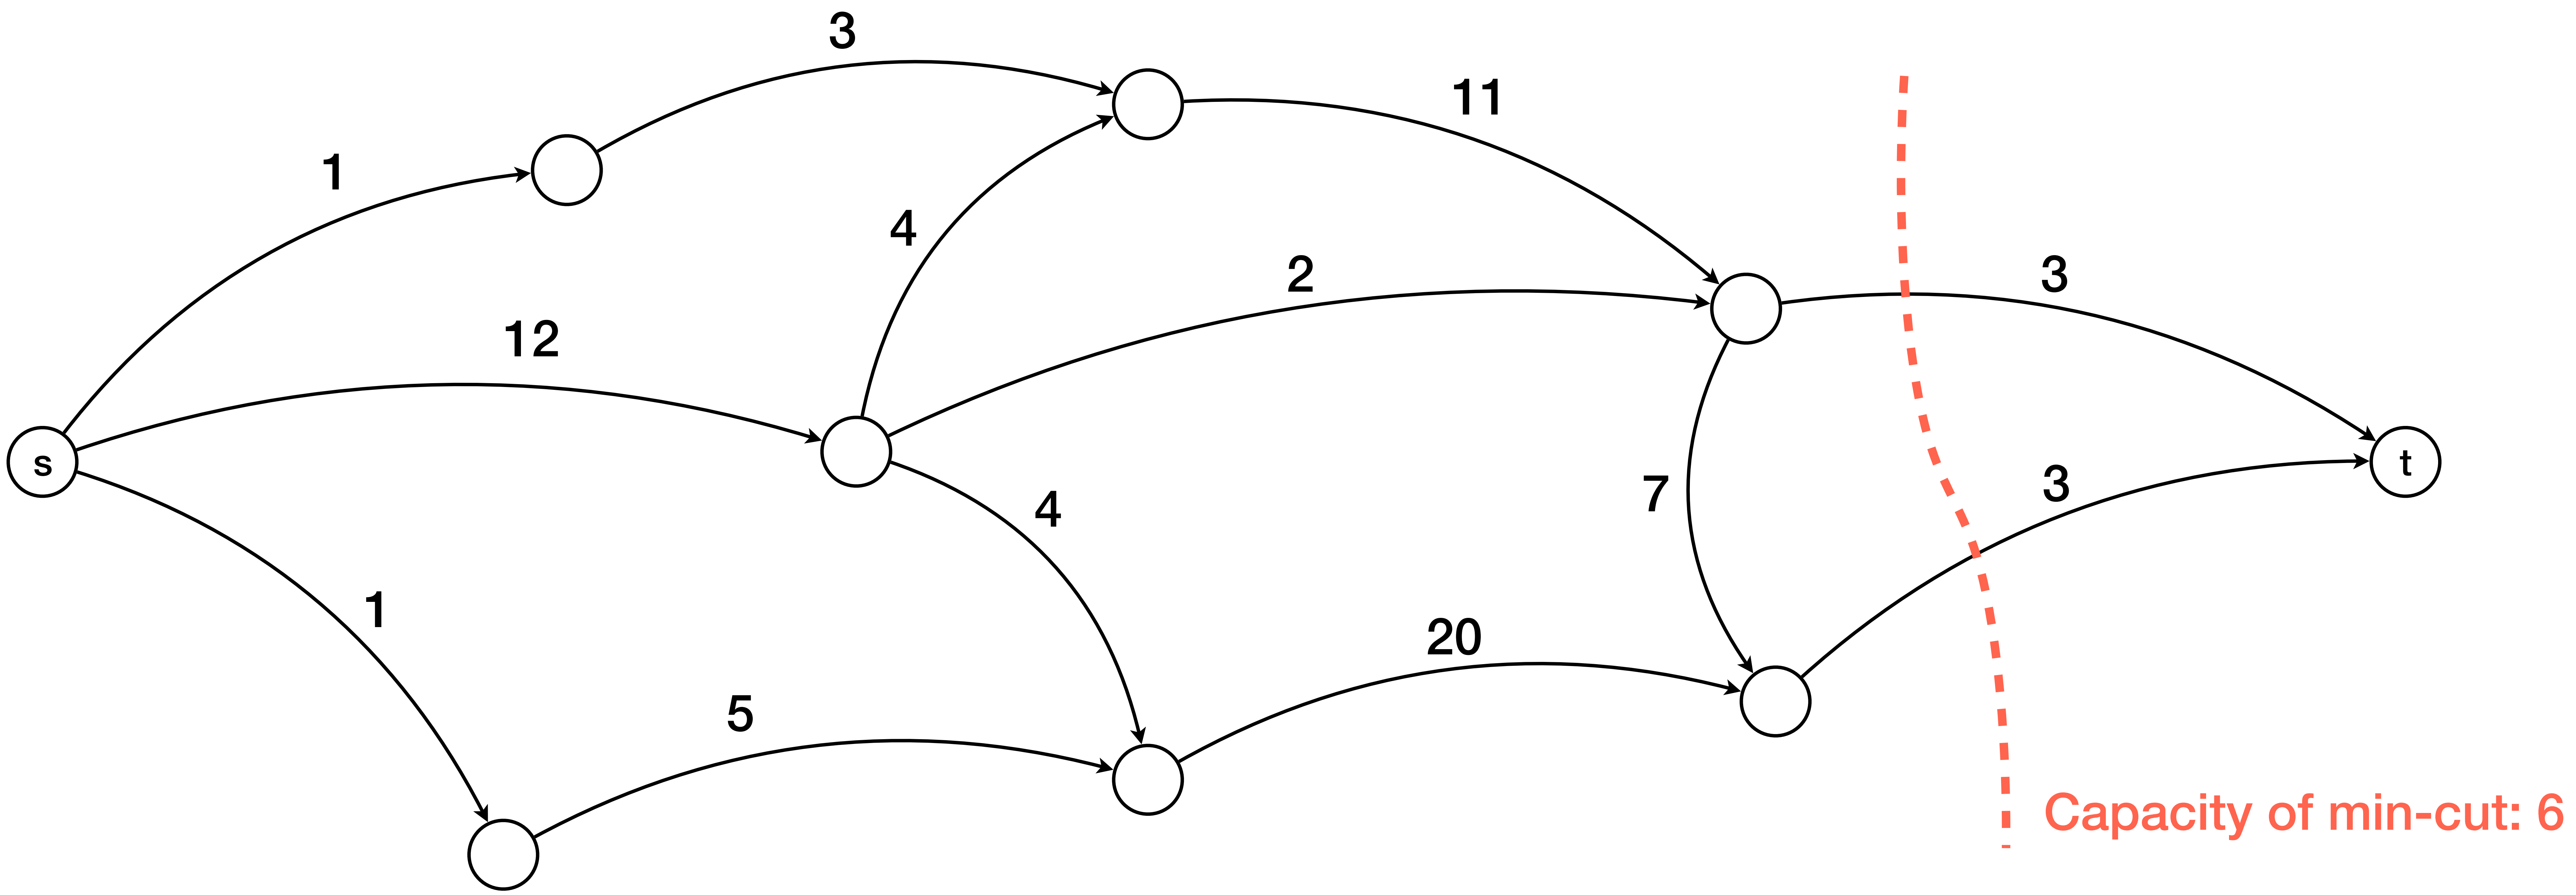

Max-flow Min-cut theorem

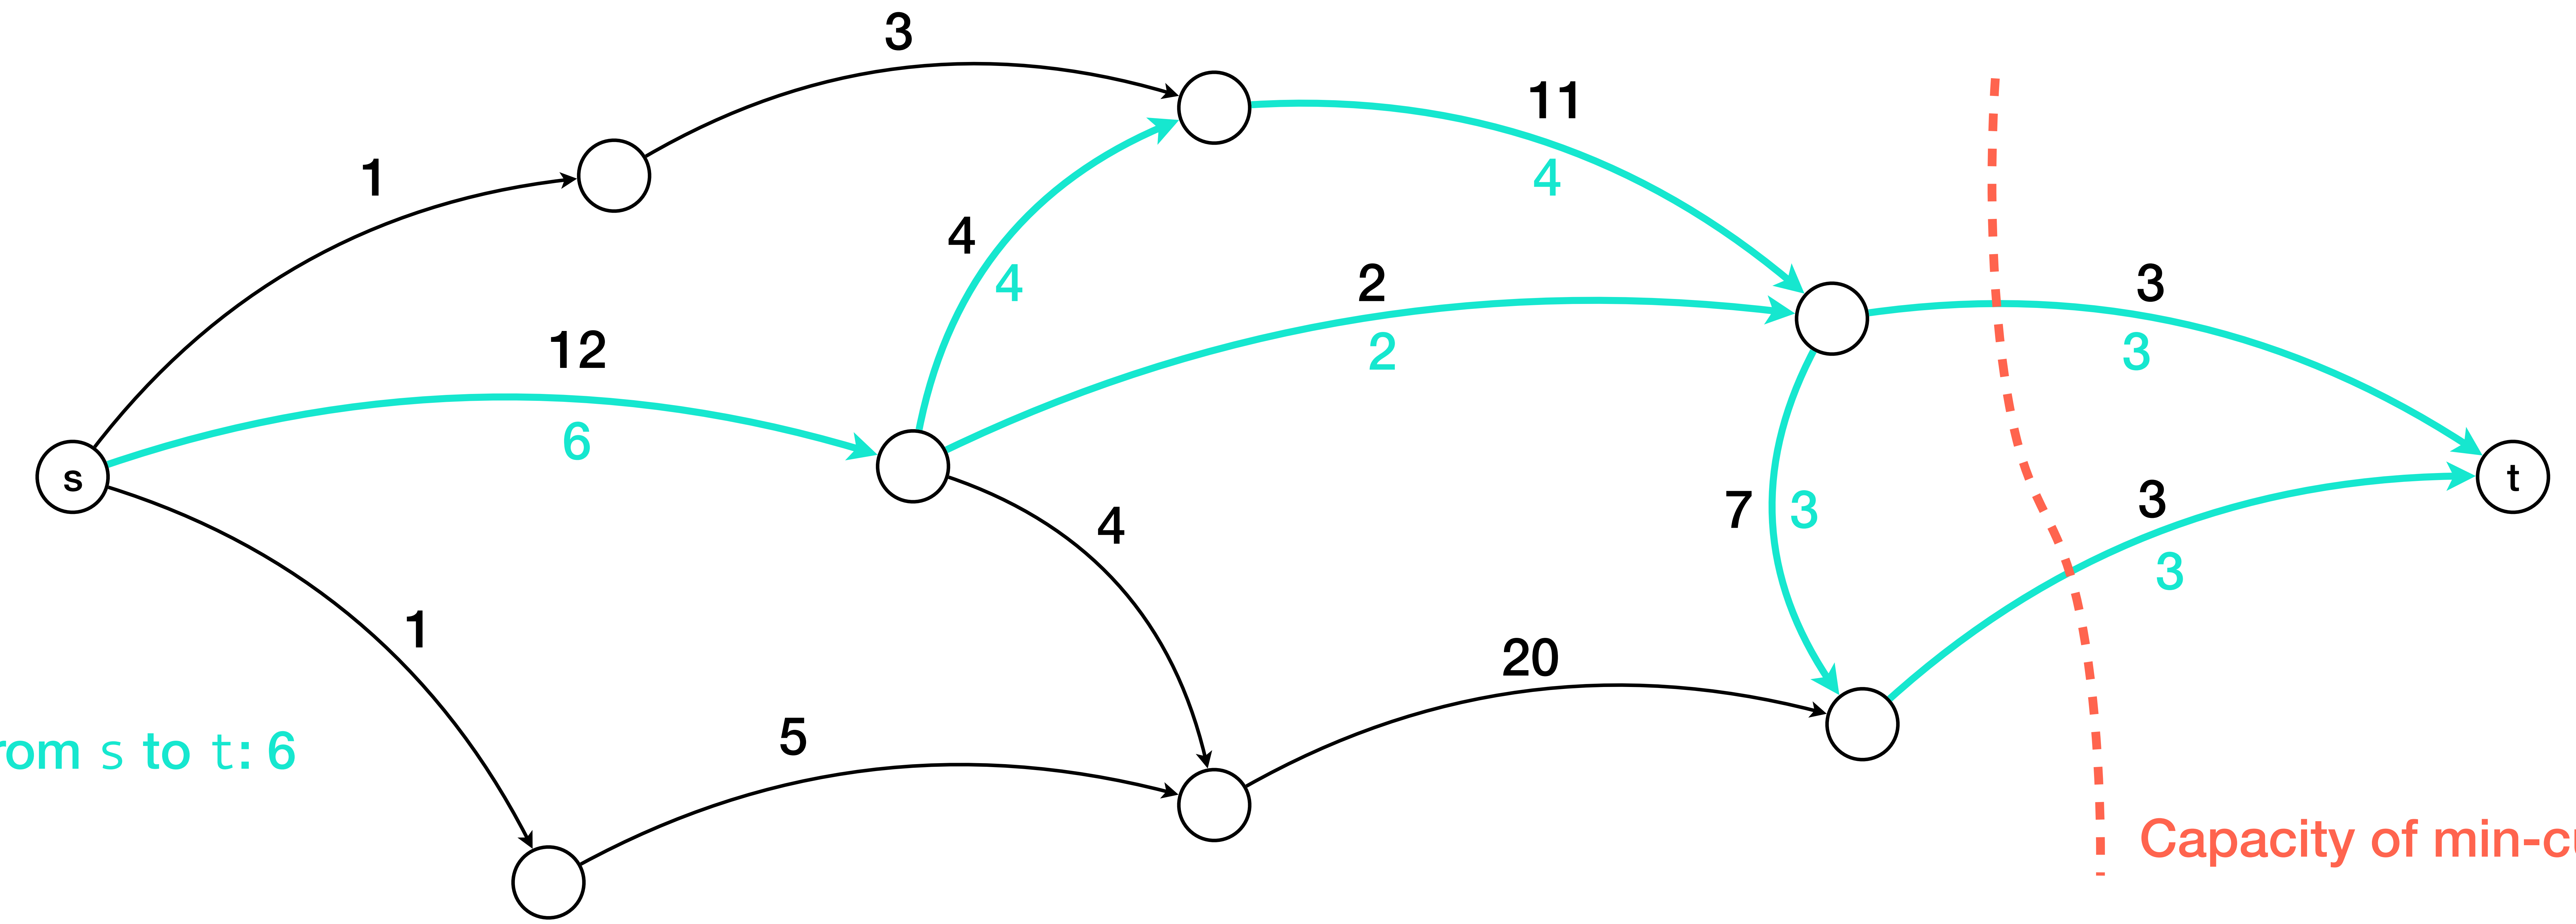

Max-flow from `s` to `t`: 6

Capacity of min-cut: 6

Cuts & Flows

Cuts & Flows

Many variants:

- directed/undirected
- with/without terminals
- multi-cut
- non integer weights
- negative weights
- …

Cuts & Flows

Many variants:

- directed/undirected
- with/without terminals
- multi-cut
- non integer weights
- negative weights
- …

Many many applications!

Max Bipartite Match

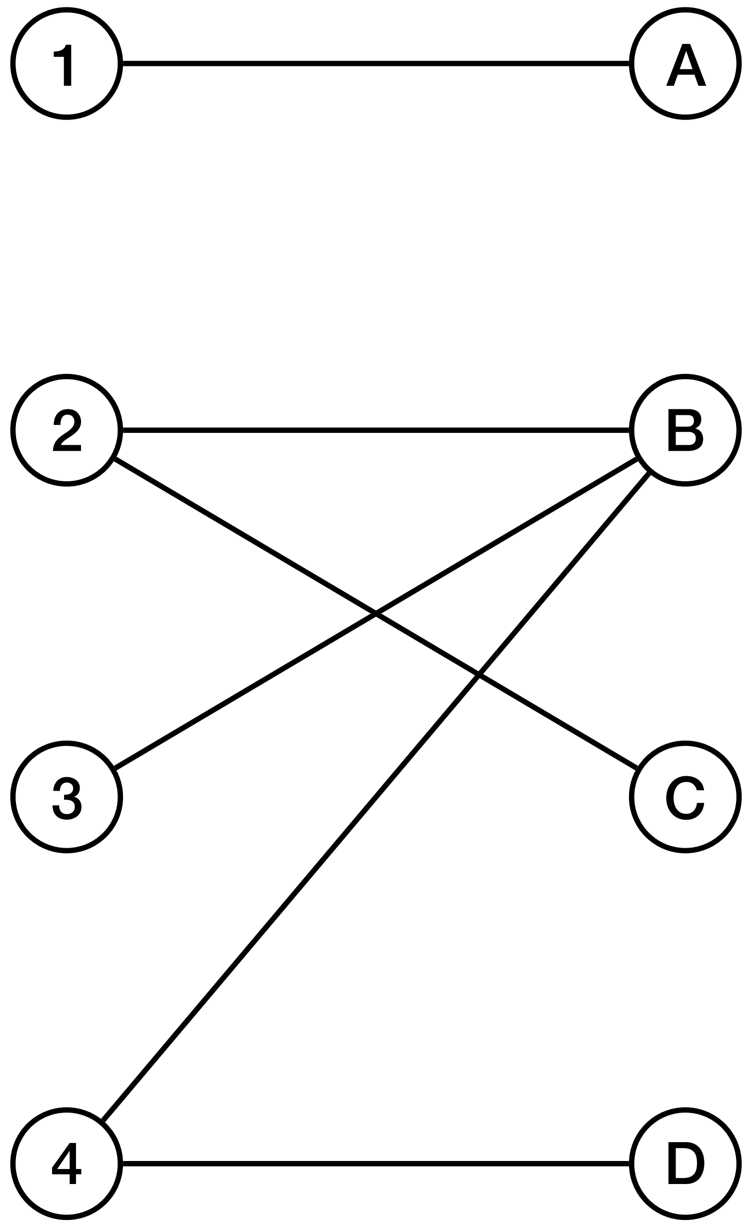

Max Bipartite Match

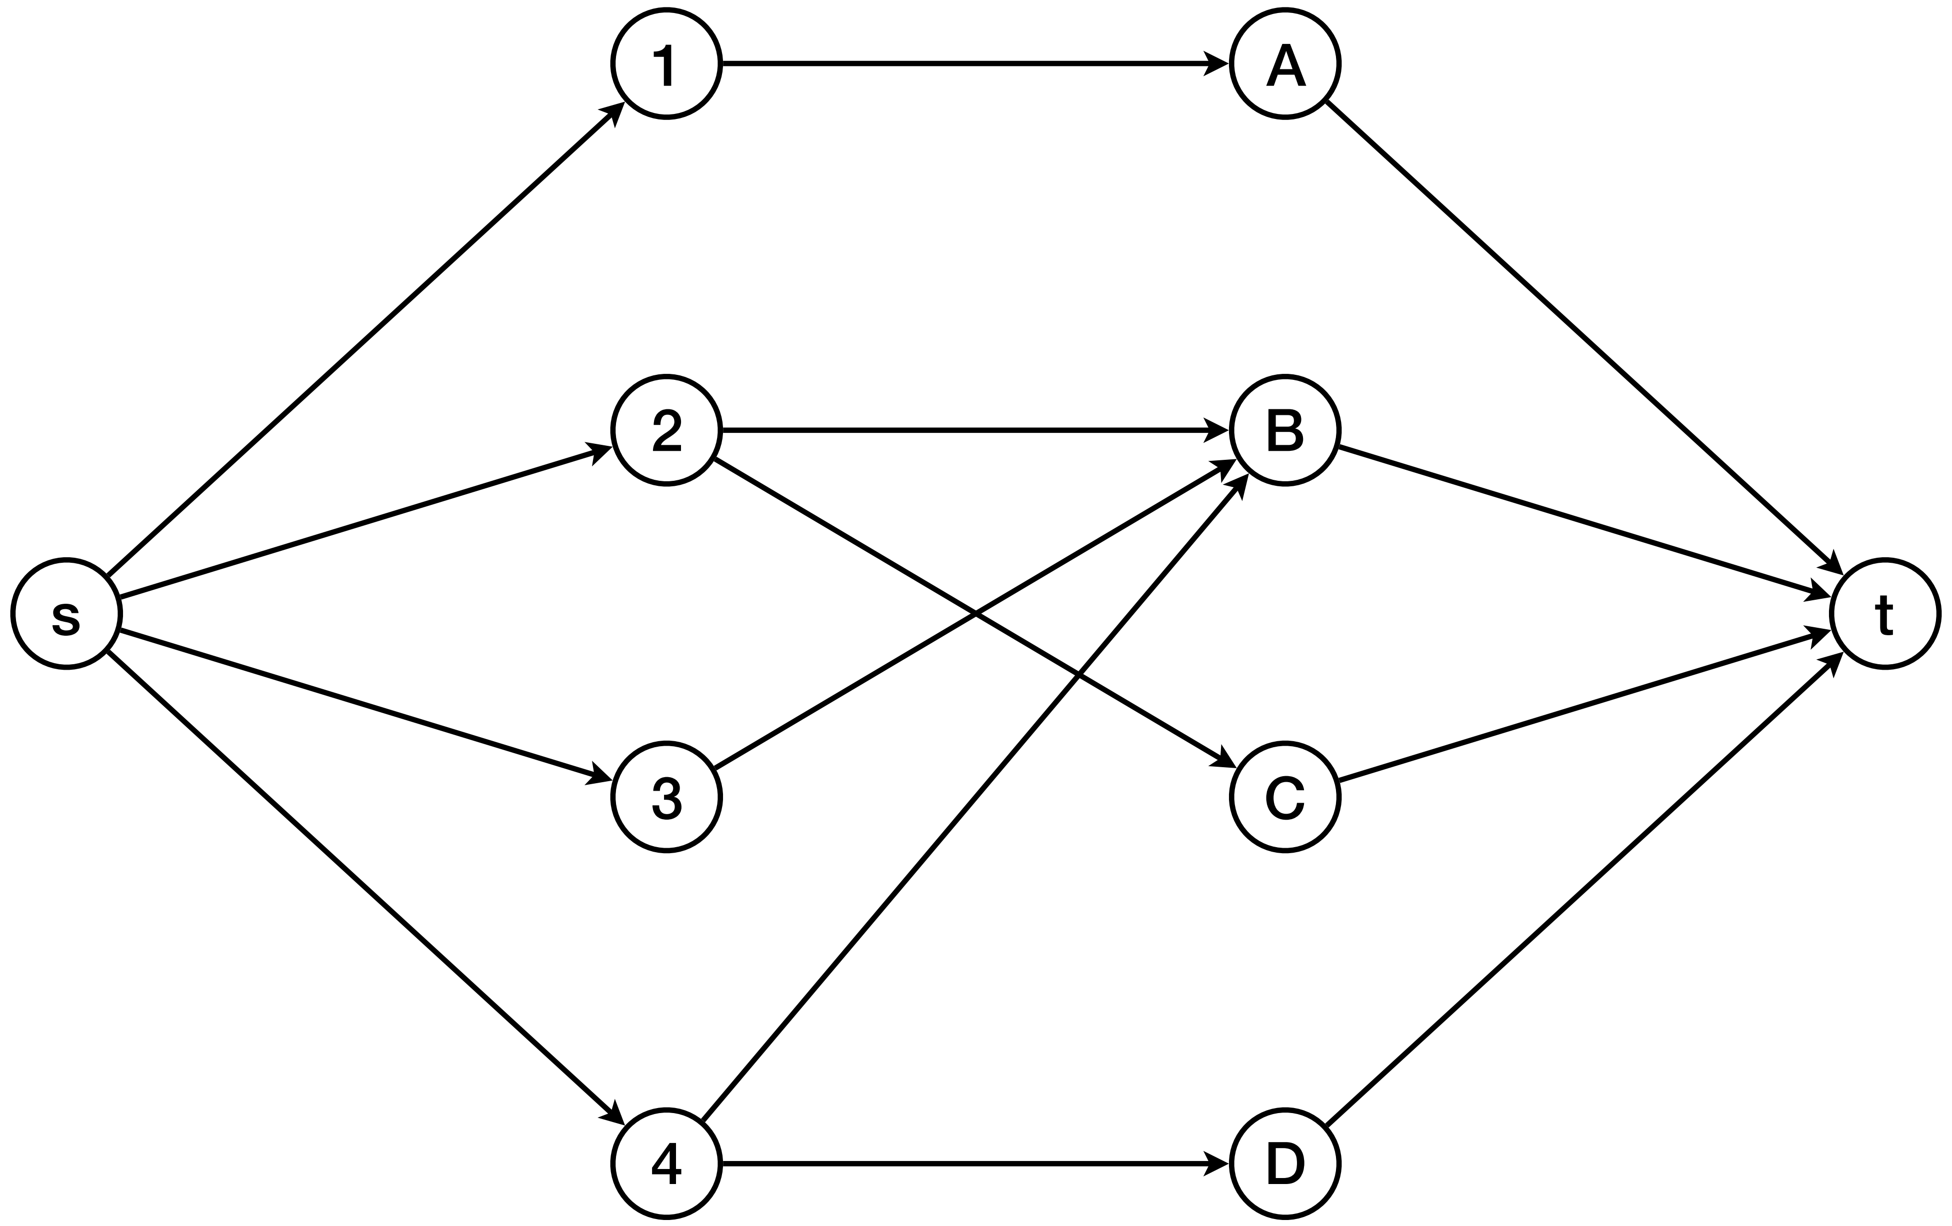

Max Bipartite Match

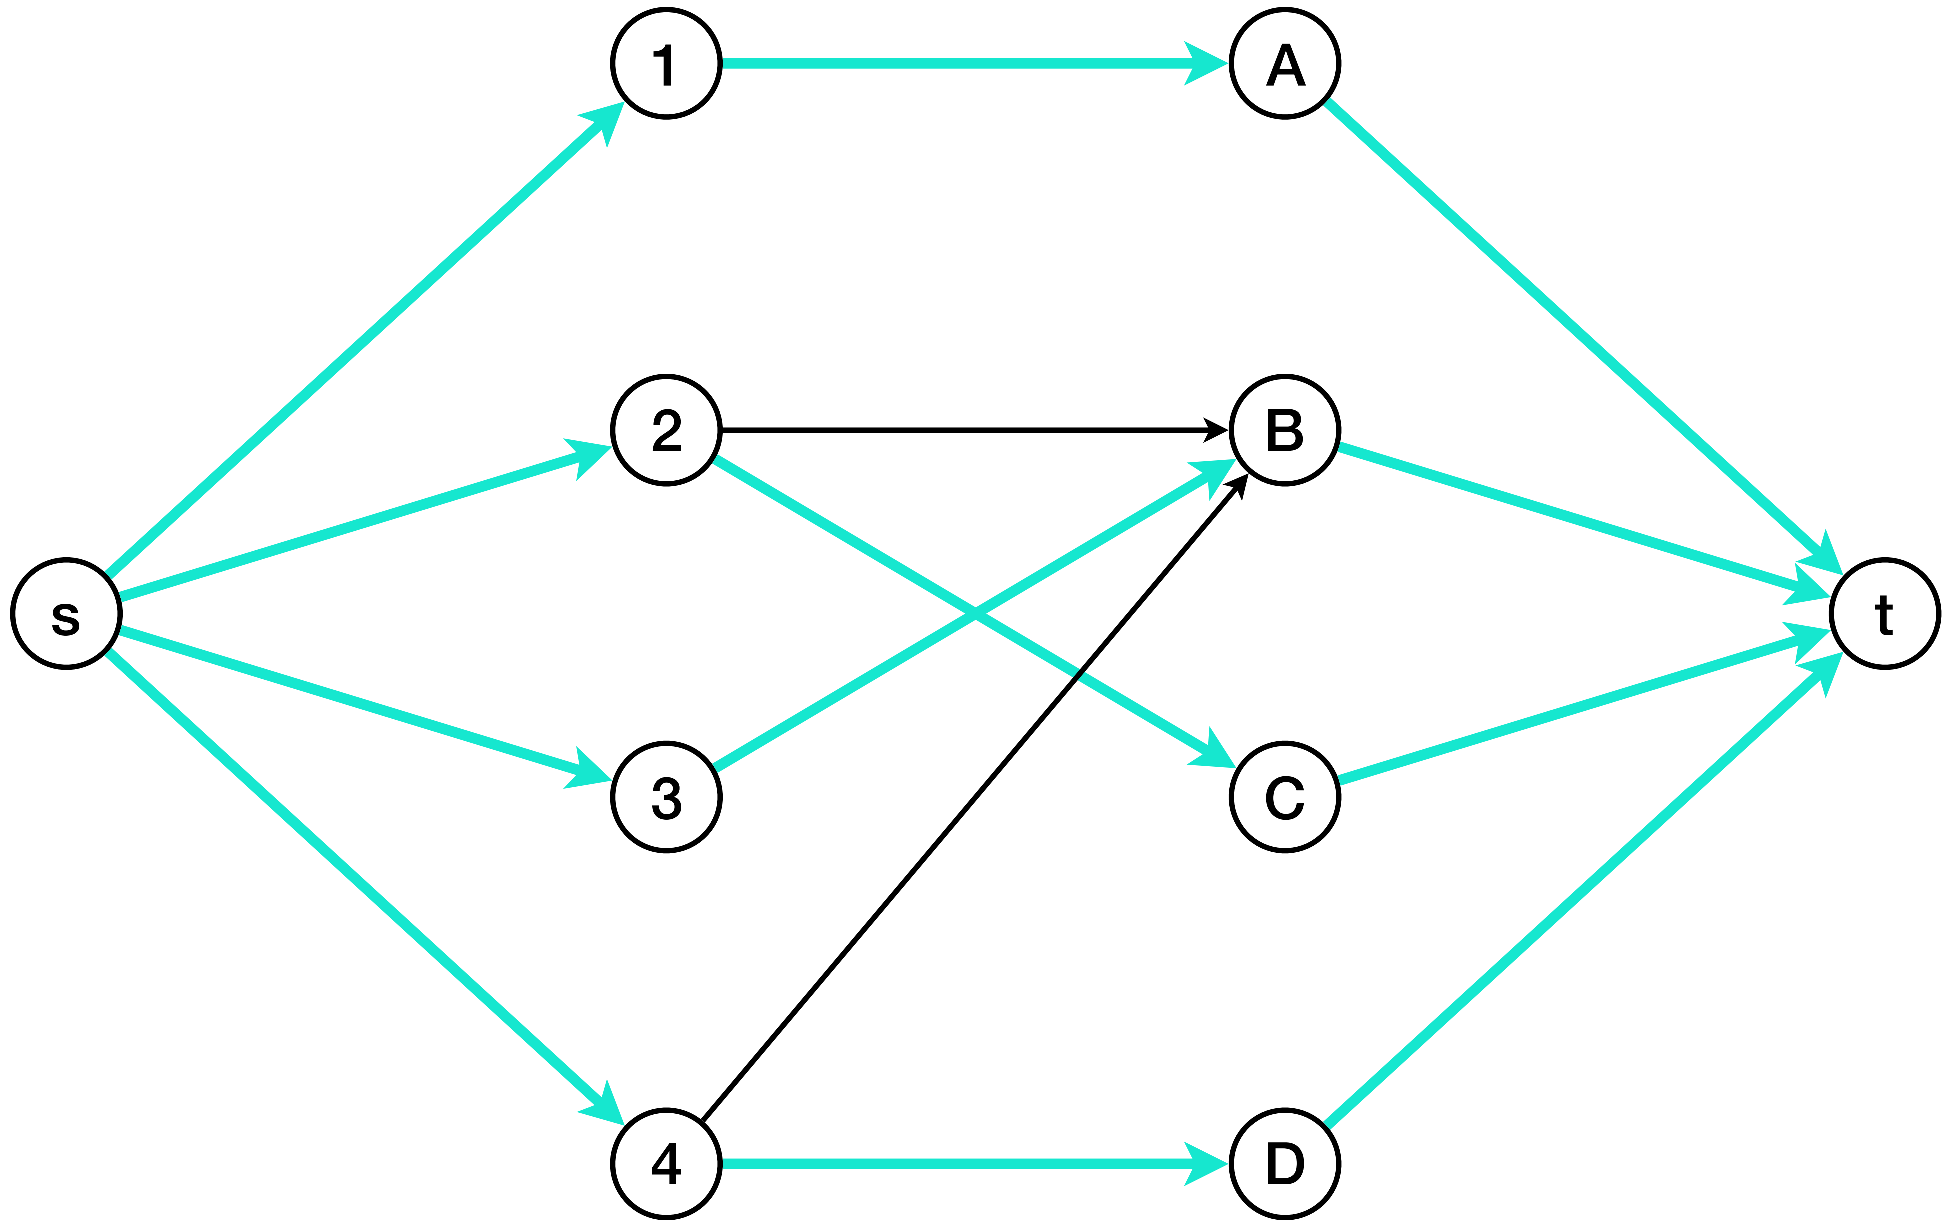

Max Bipartite Match

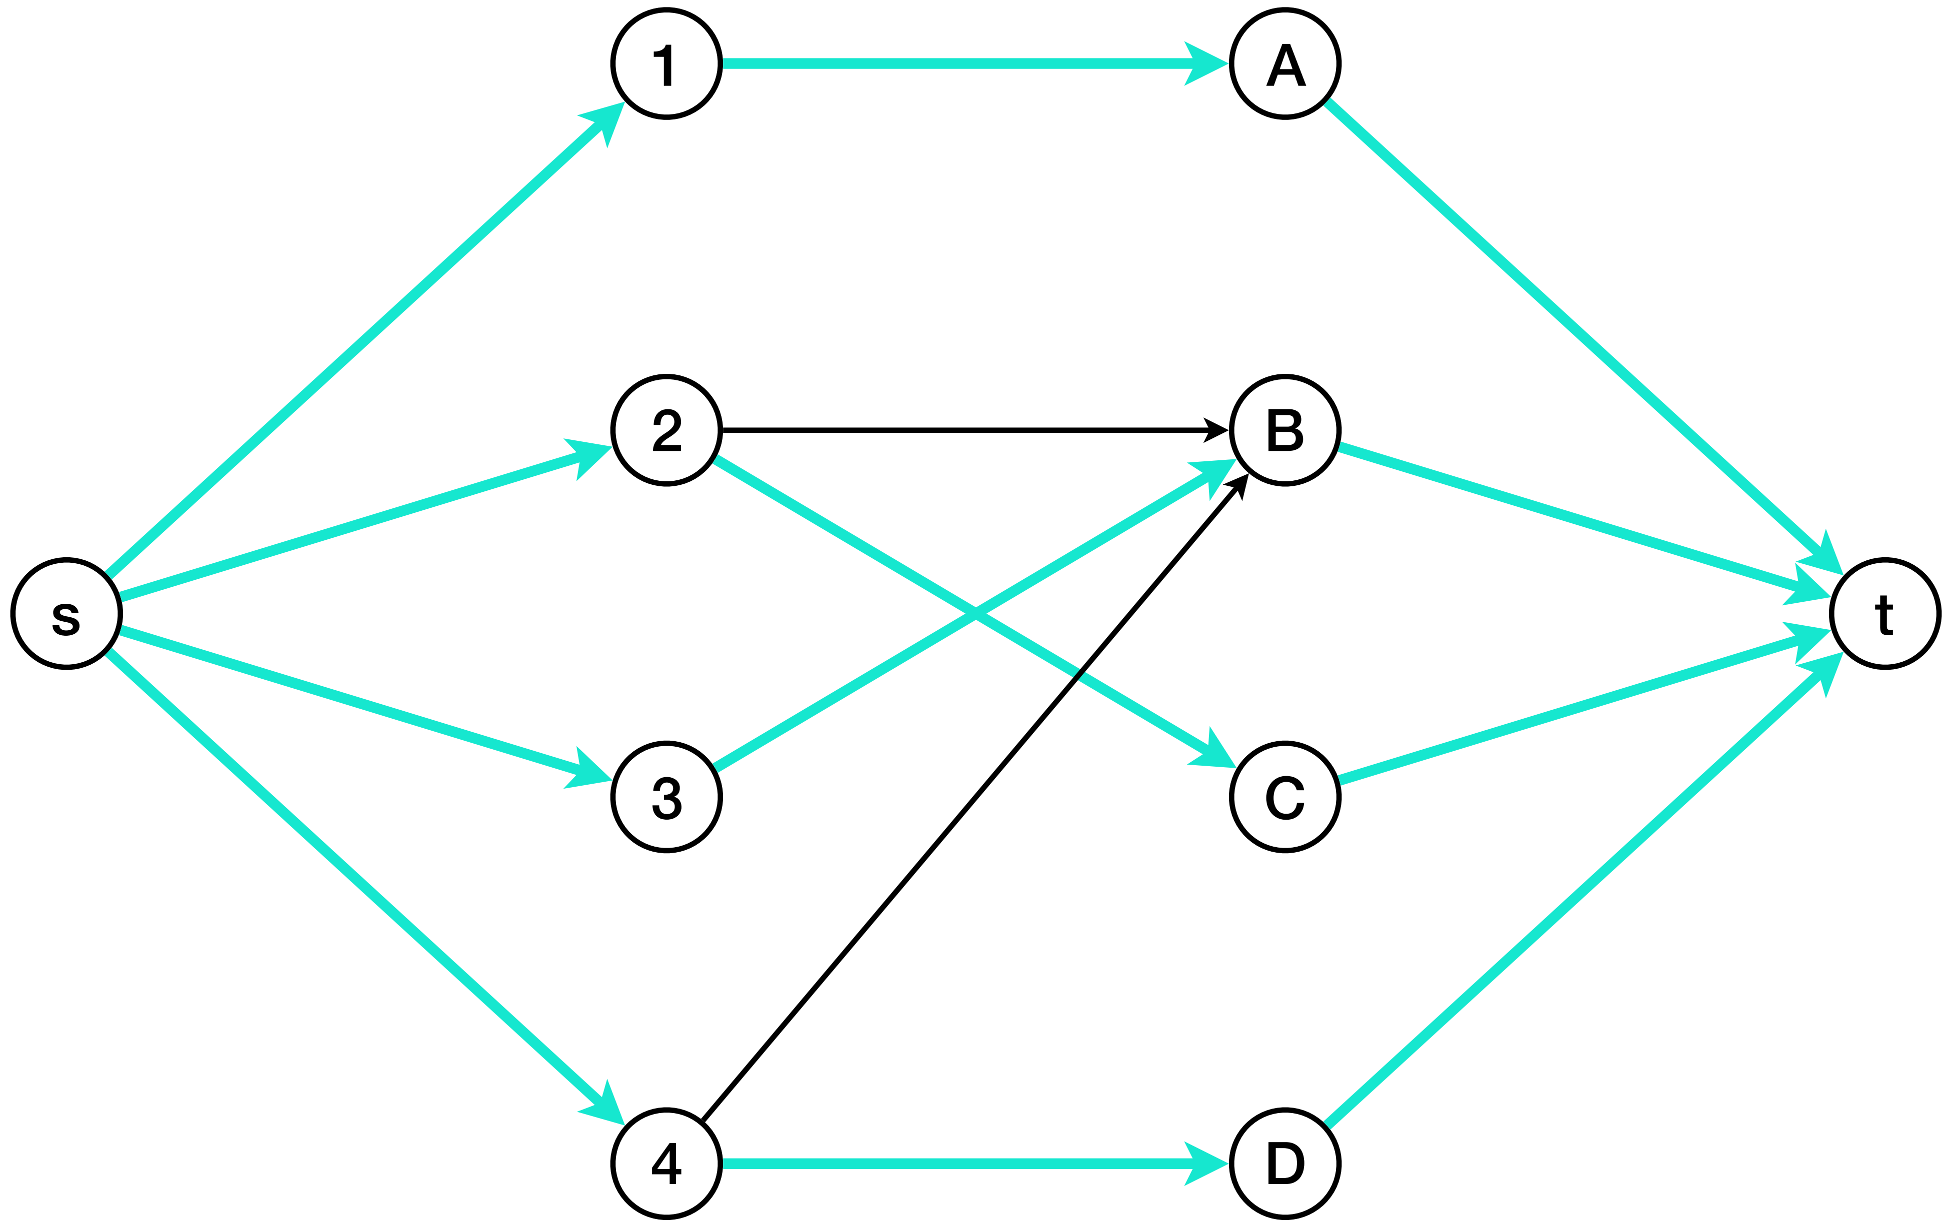

we will use a similar trick…

Back to Graphcut Textures…

“Chernobyl harvest”

Where to place next patch?

Which pixels to use?

Which pixels to use?

Which pixels to use?

Graph cuts to the rescue

Which pixels to use?

Graph cuts to the rescue

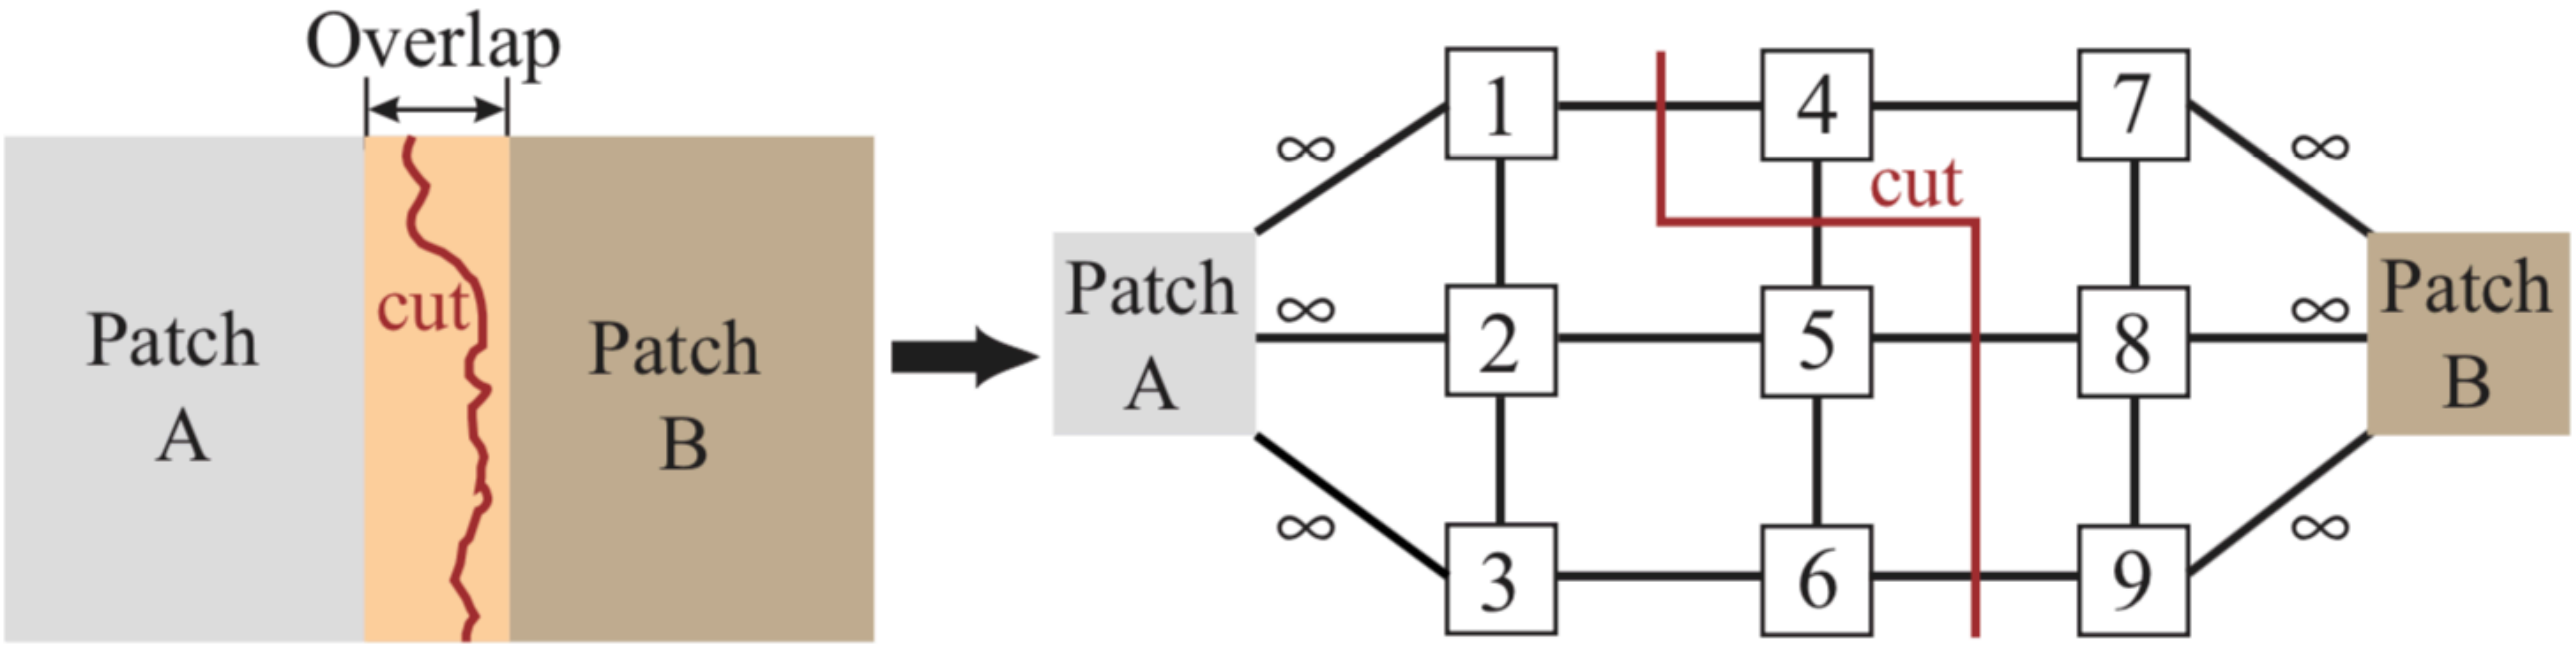

Which pixels to use?

Graph cuts to the rescue

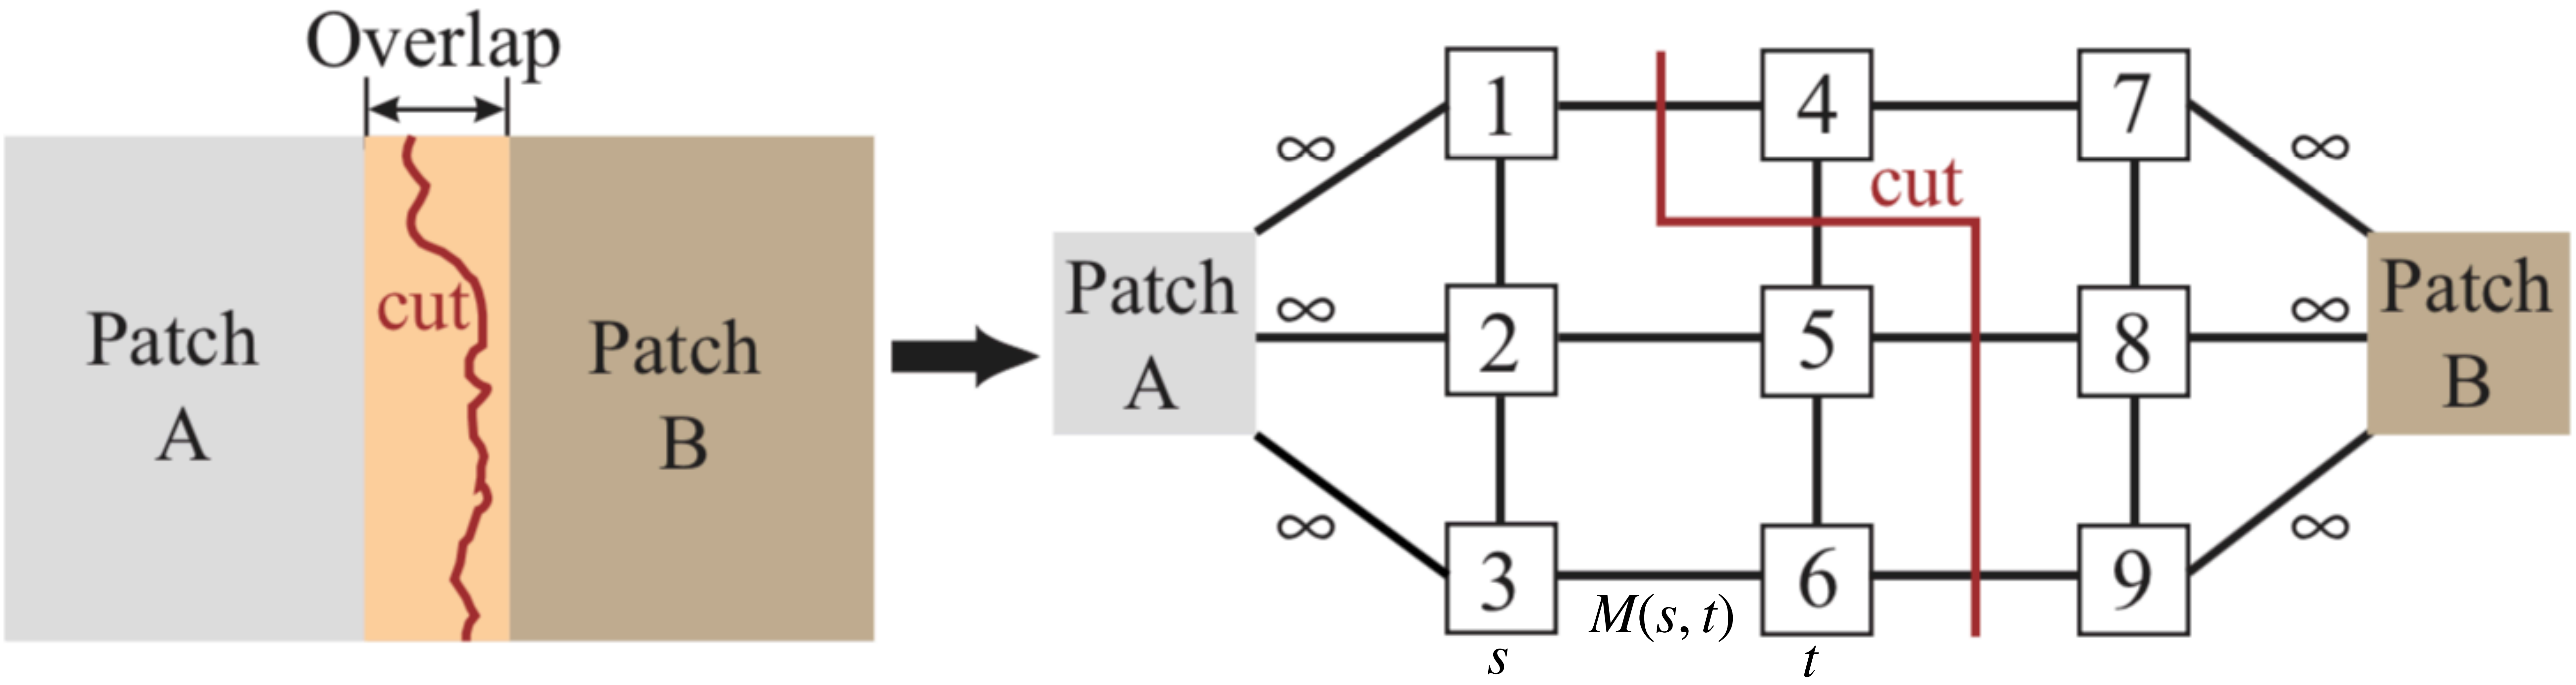

Which pixels to use?

Graph cuts to the rescue

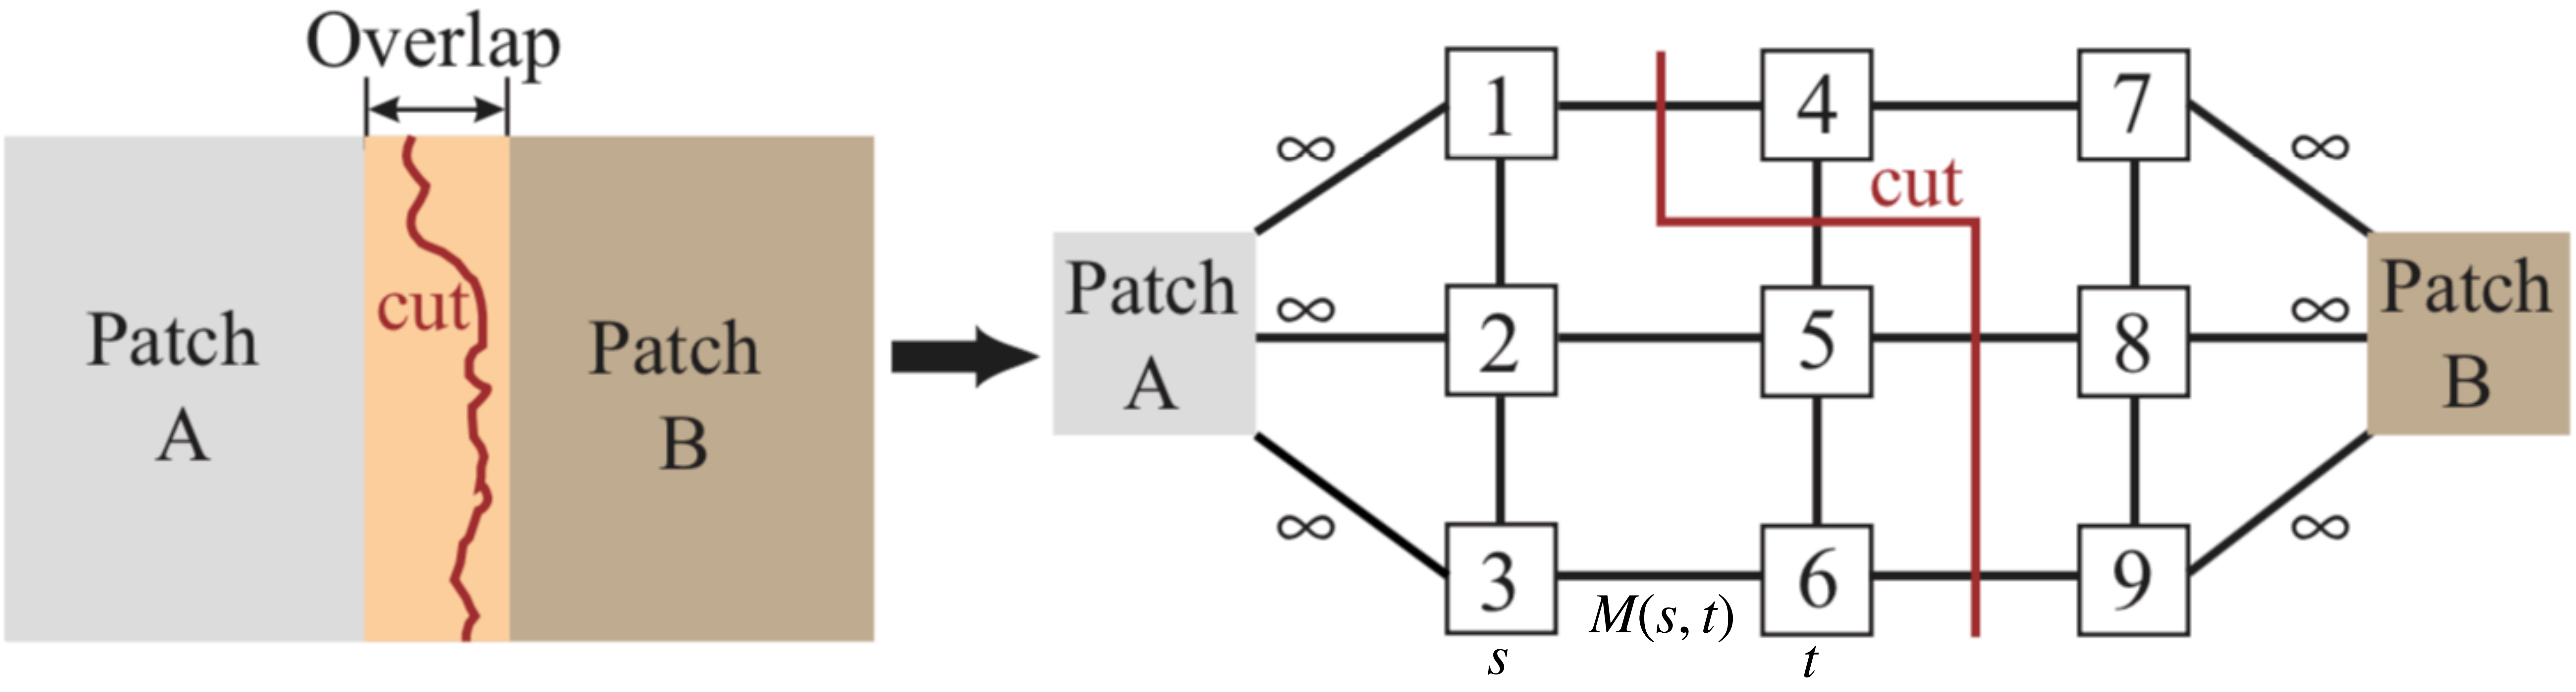

$$M(s,t) = \|A(s) - B(s)\| + \|A(t) - B(t)\|$$

Which pixels to use?

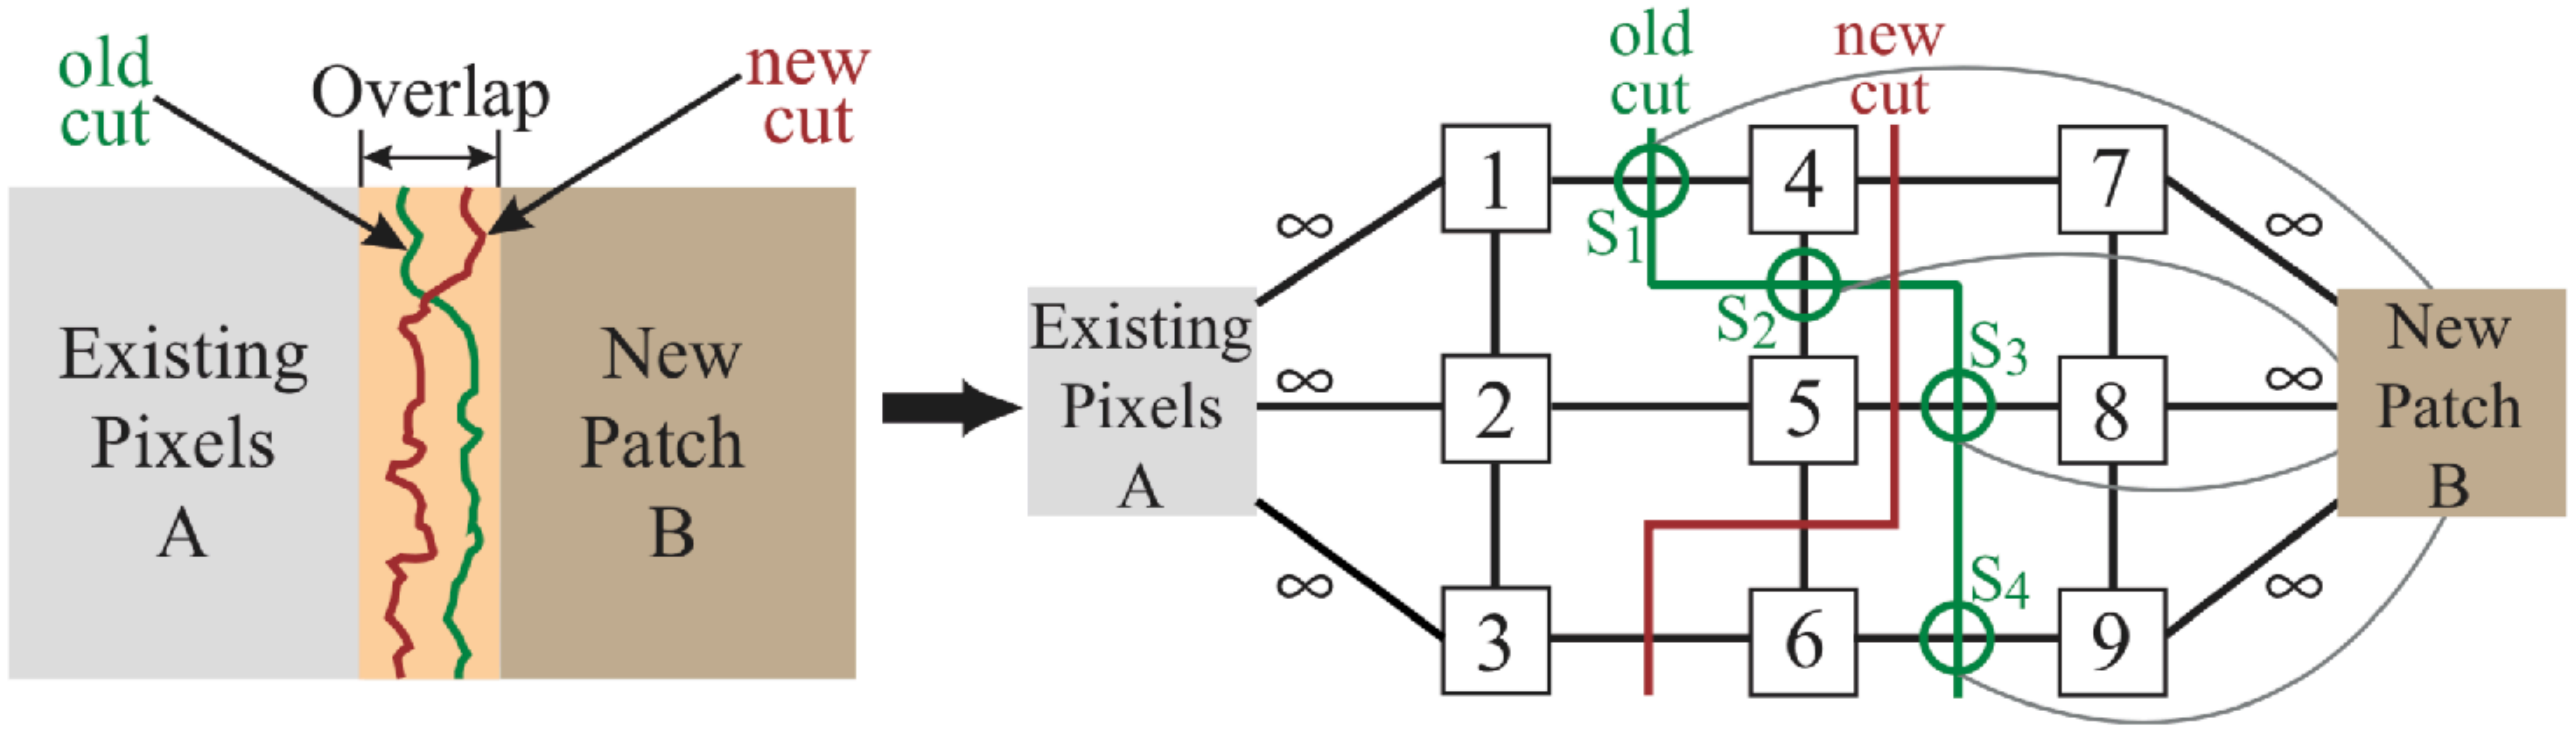

Which pixels to use?

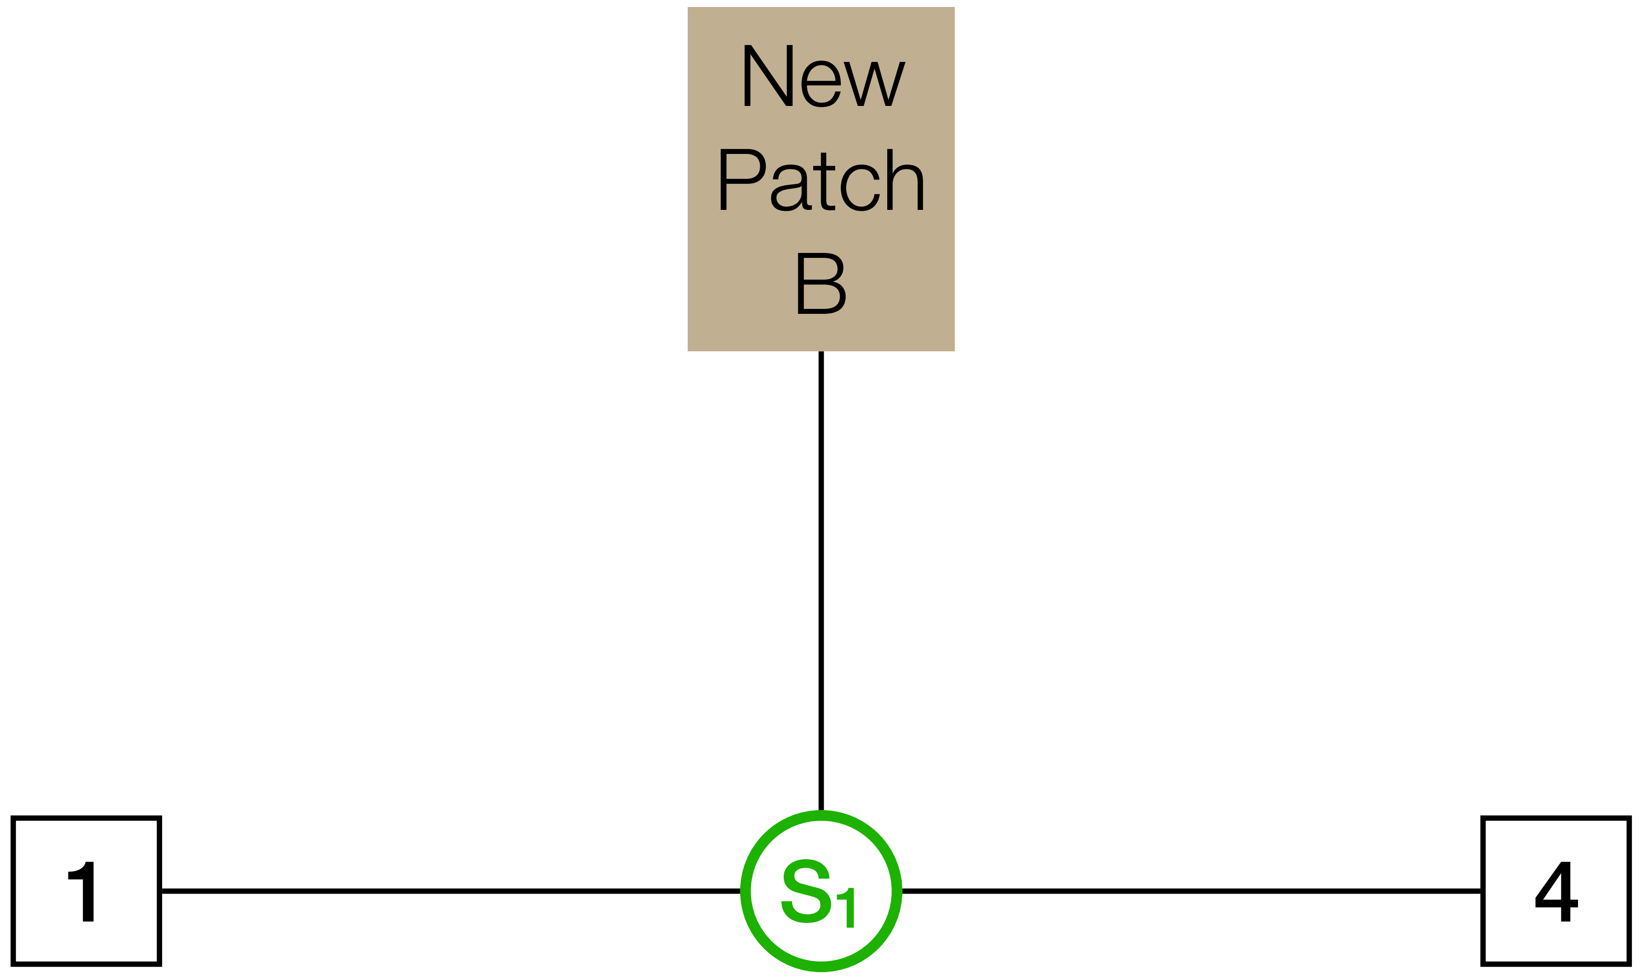

Which pixels to use?

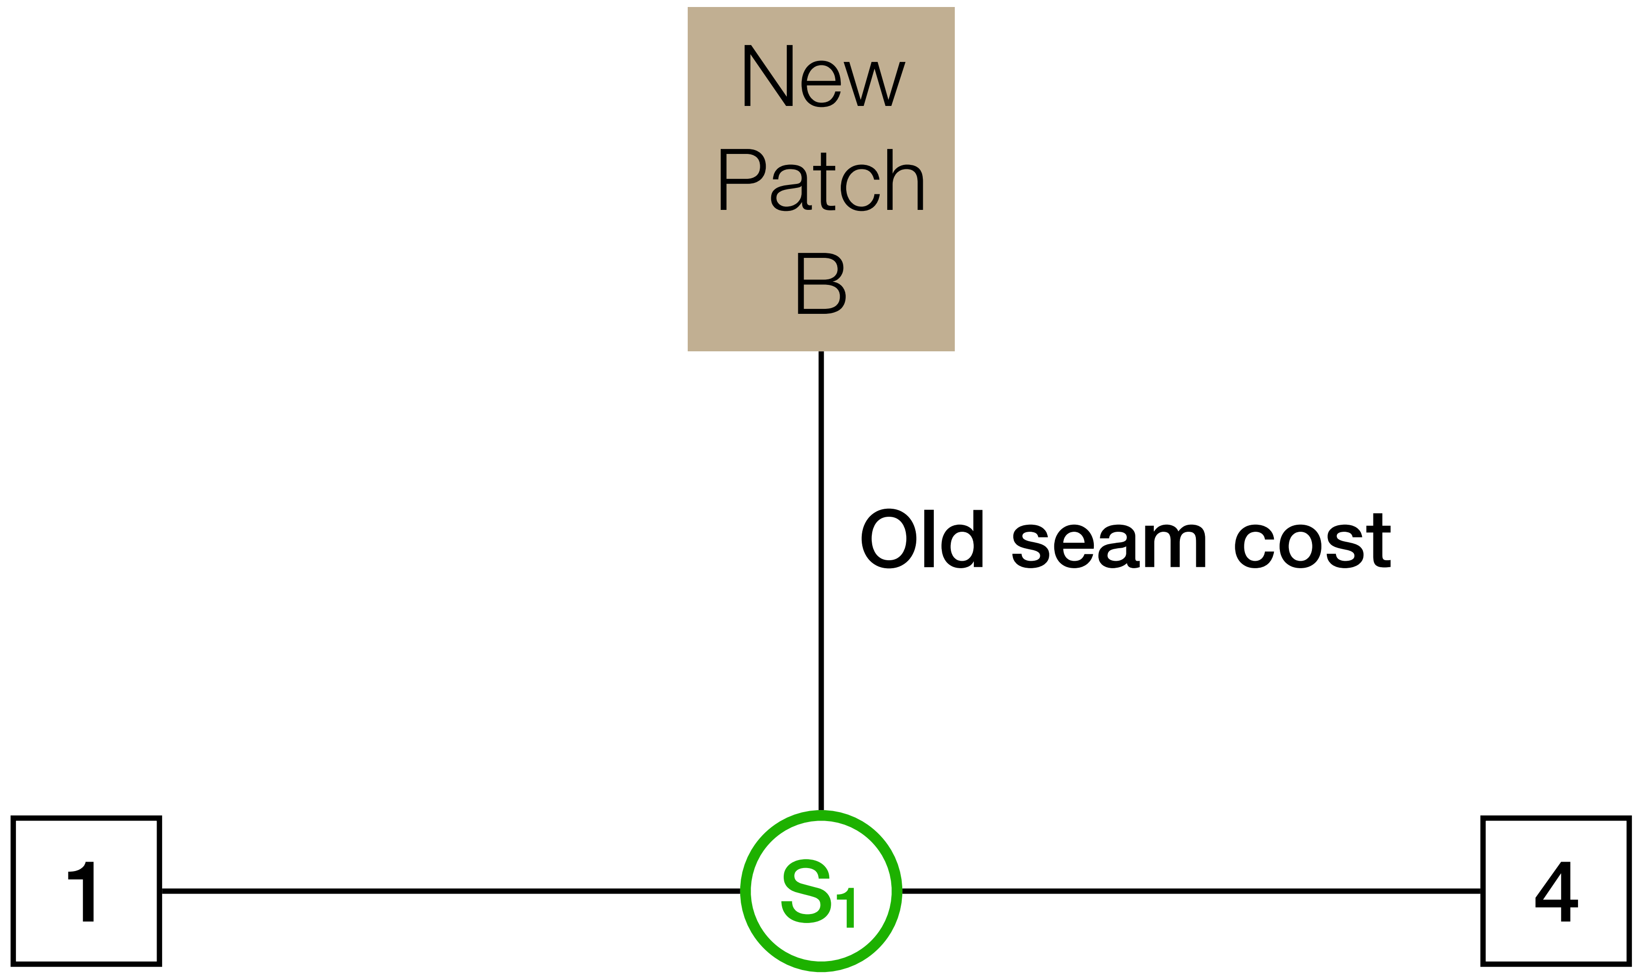

Which pixels to use?

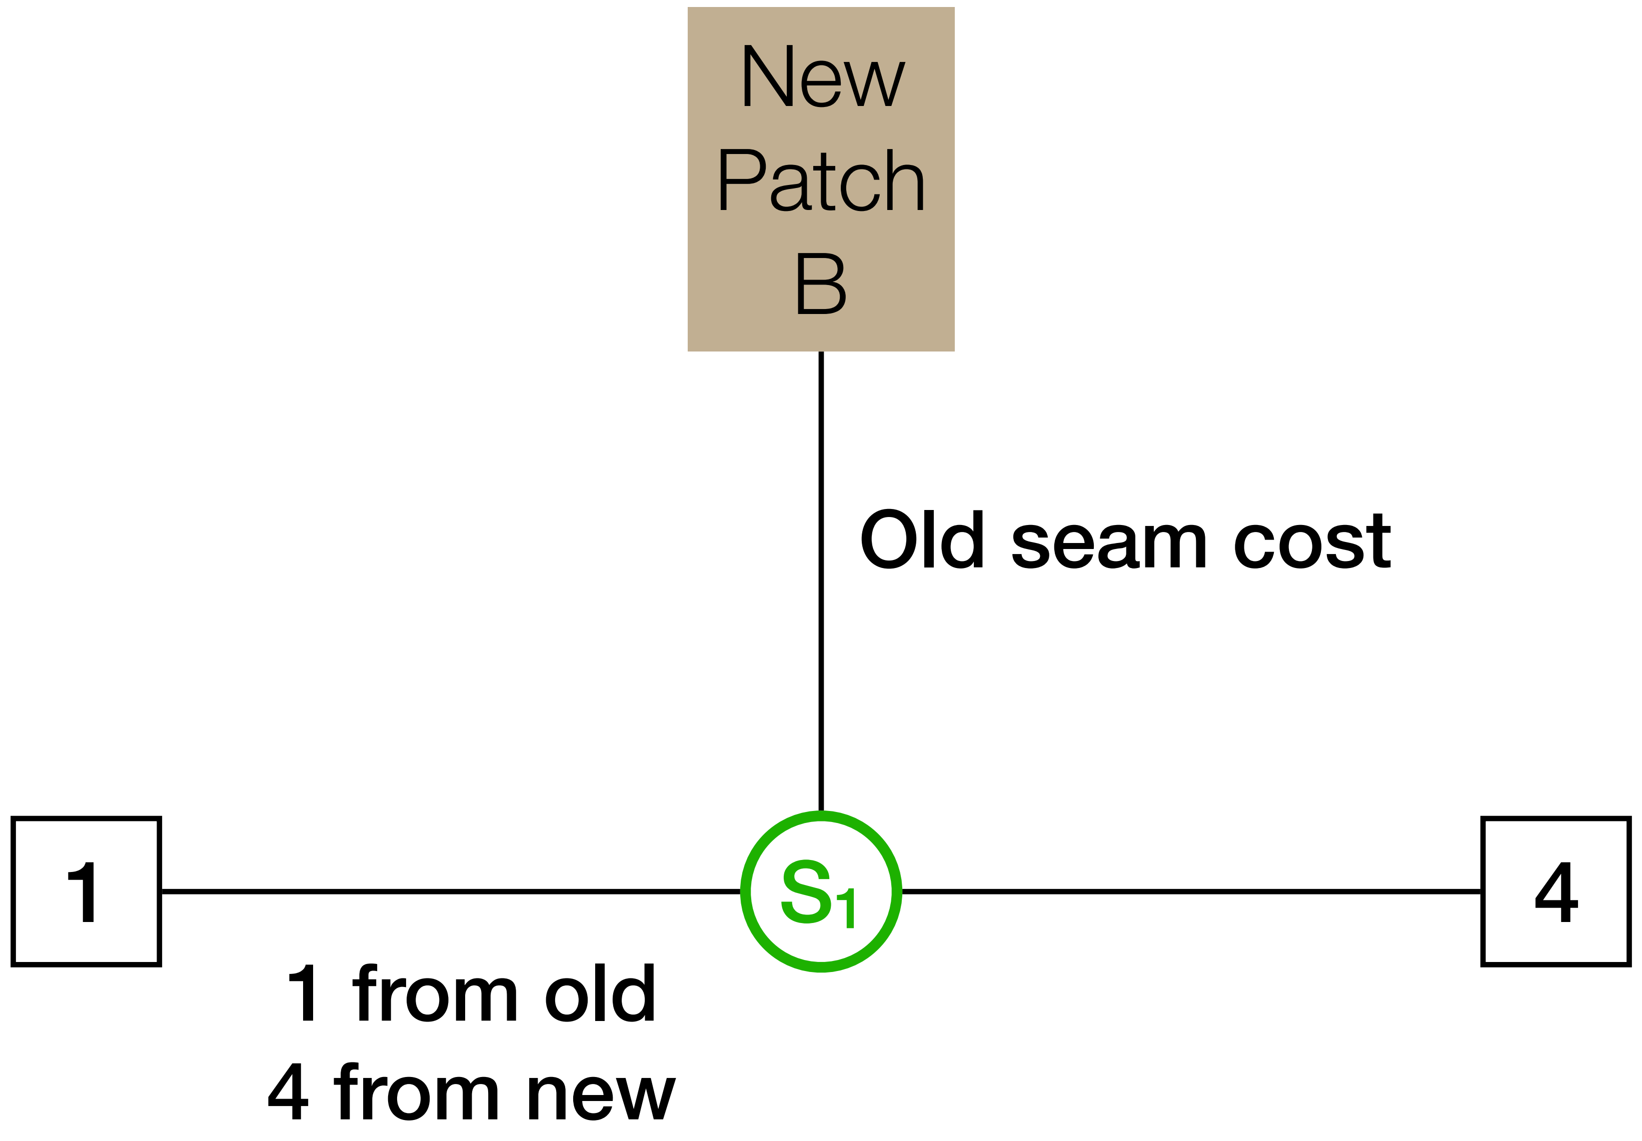

Which pixels to use?

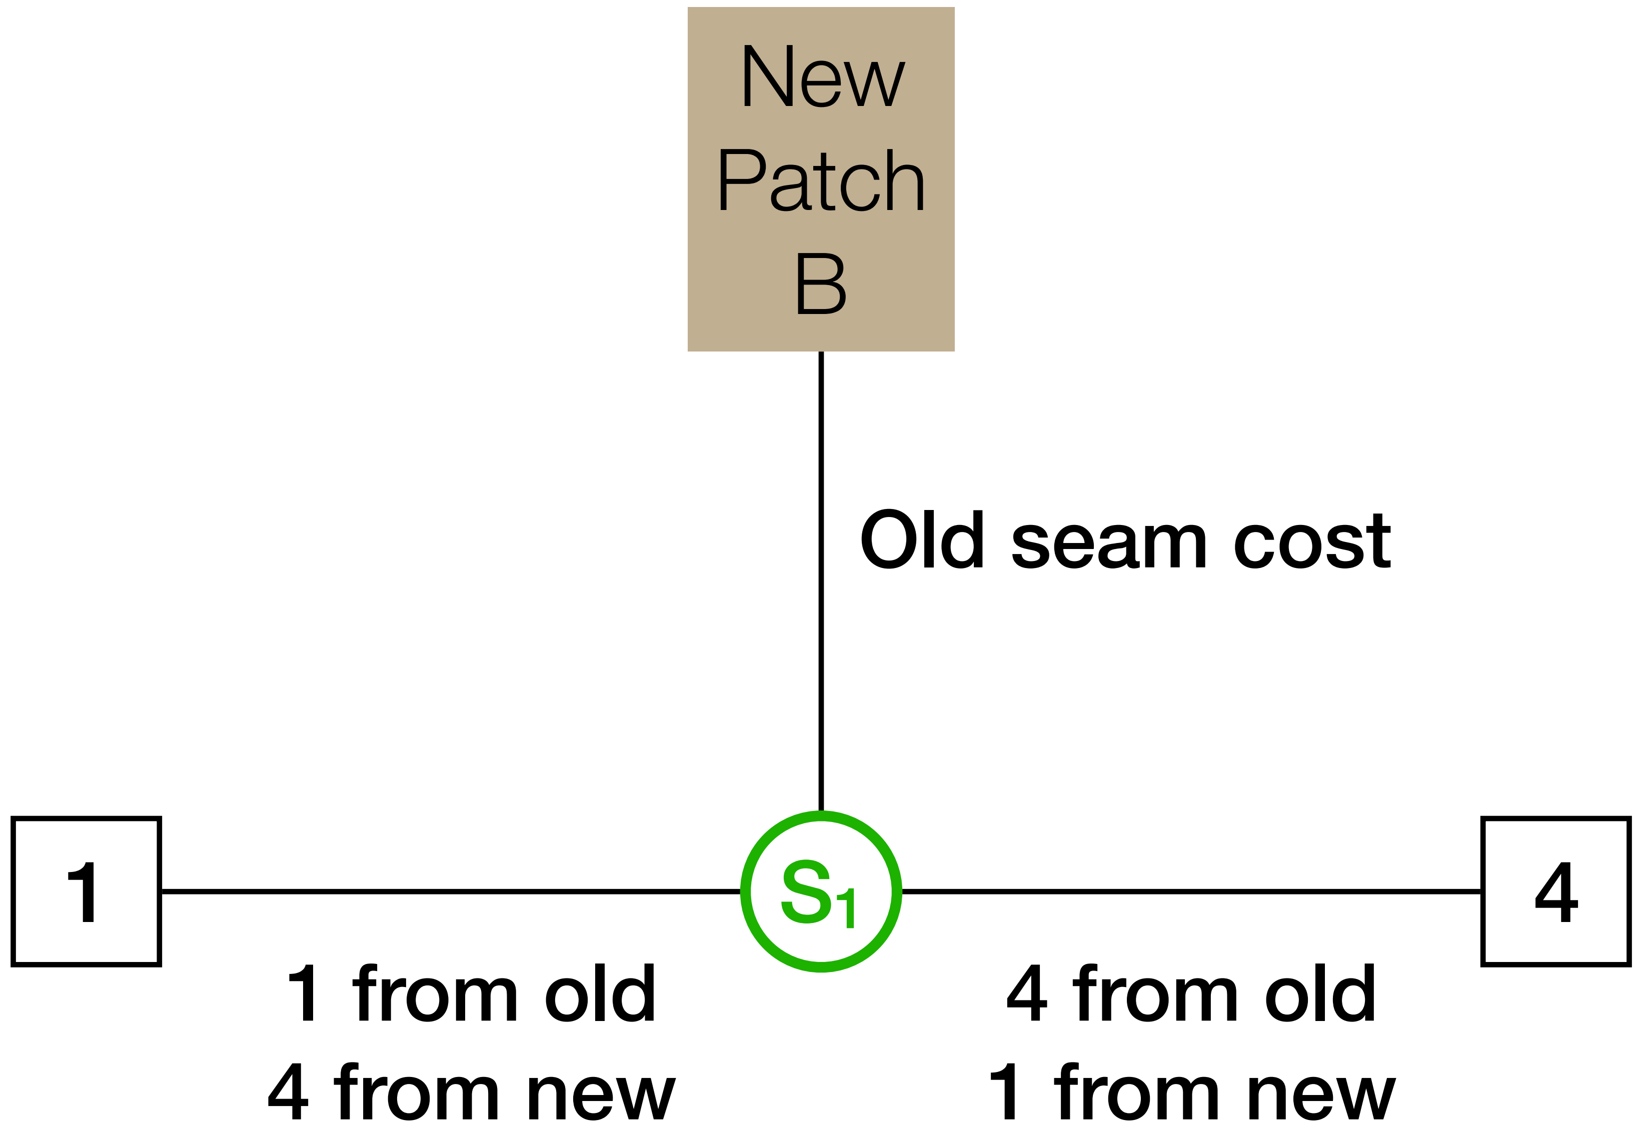

Which pixels to use?

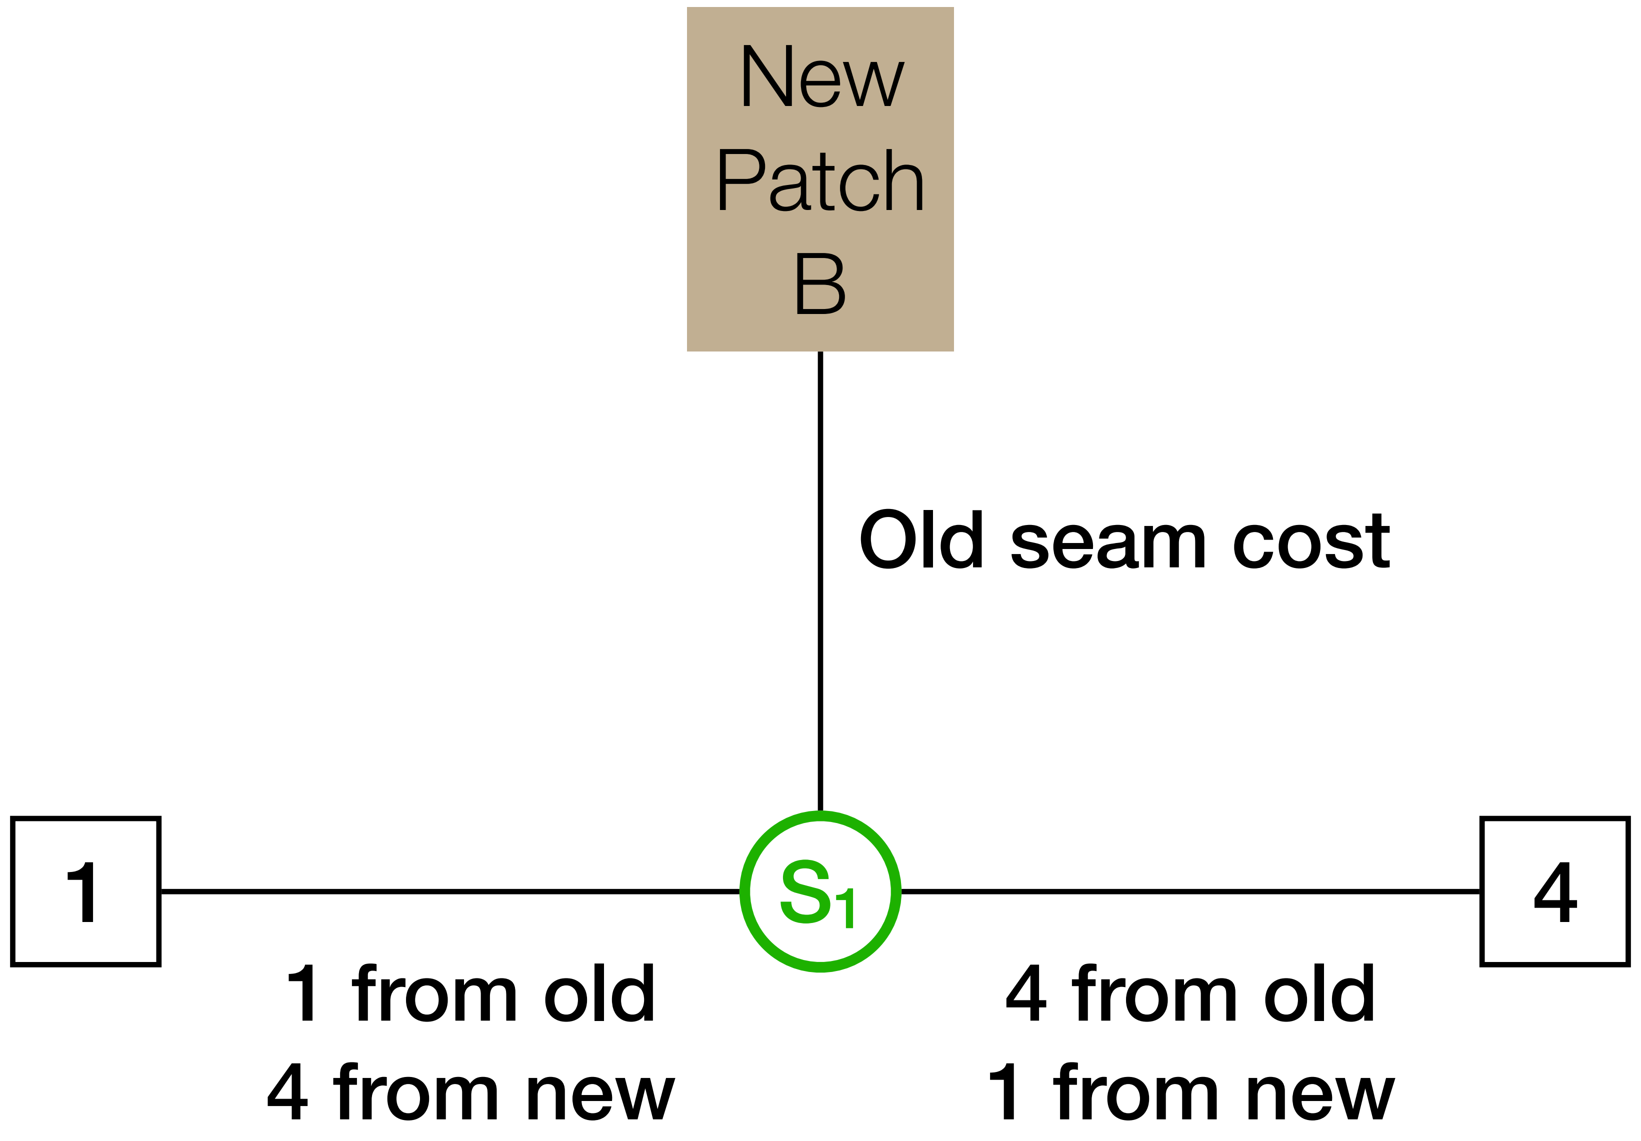

Cut at most one edge!

Which pixels to use?

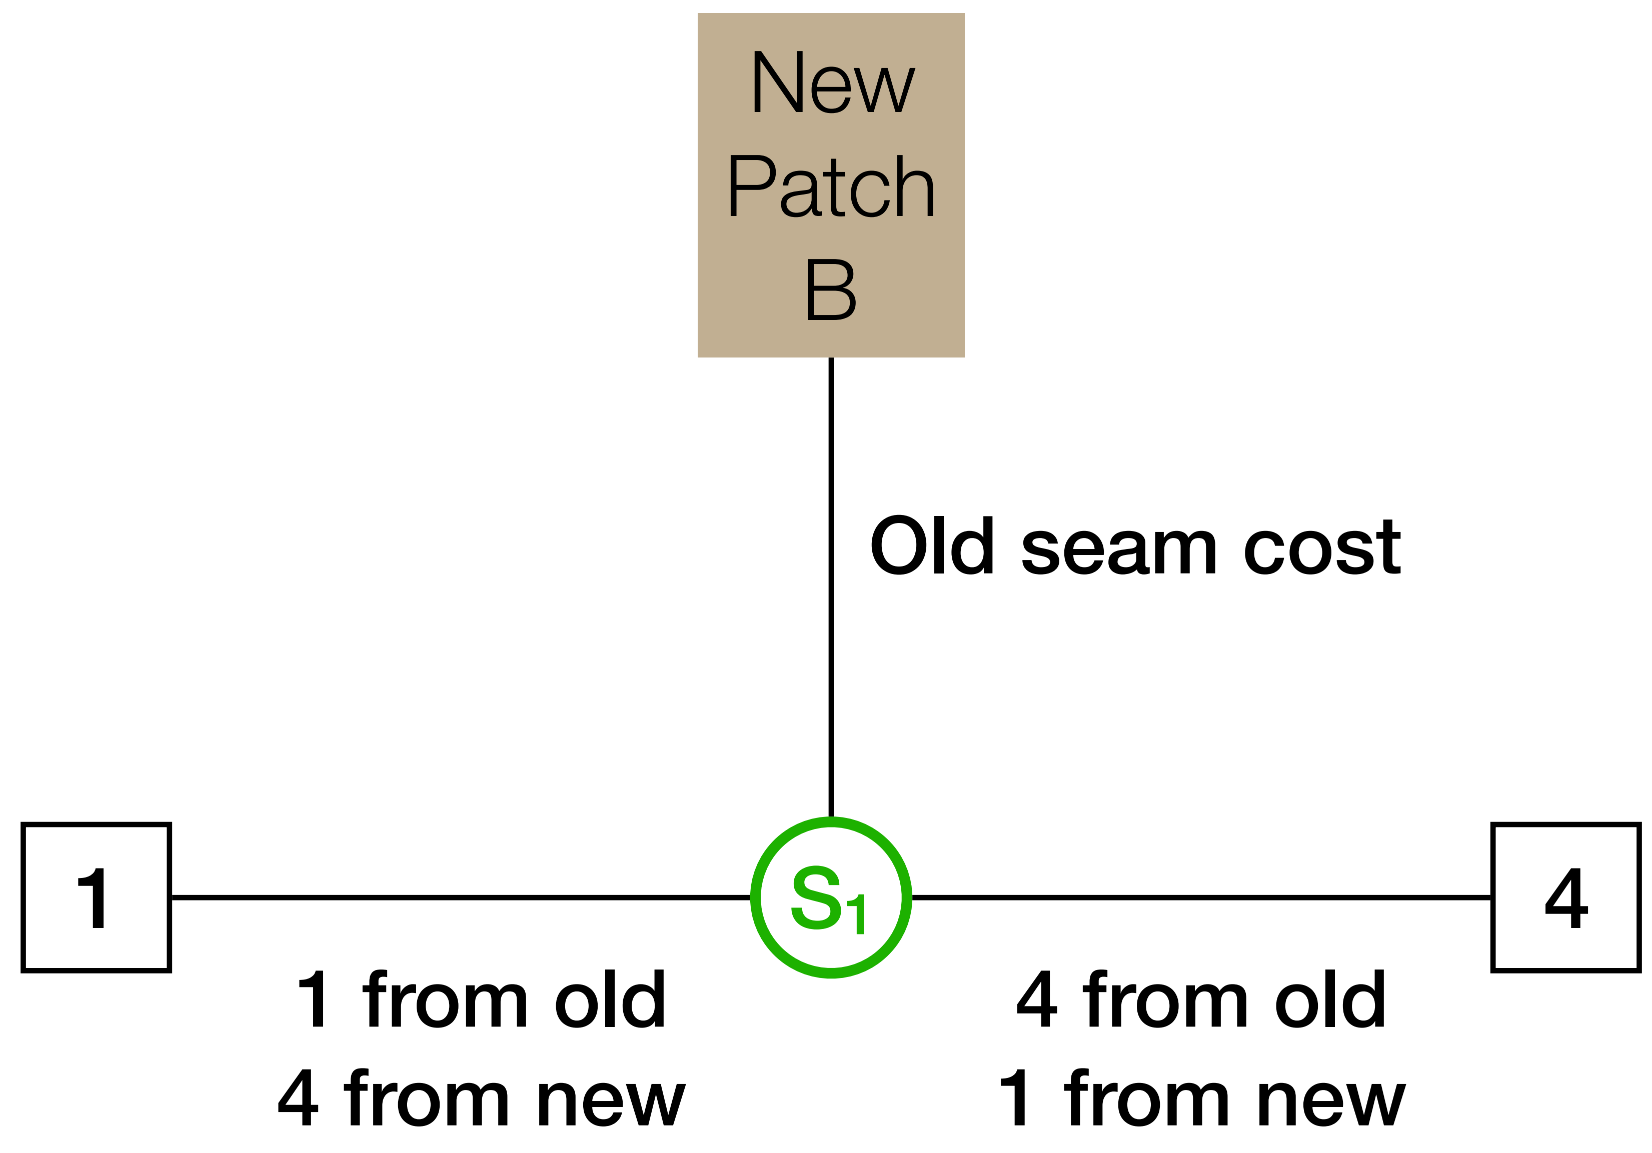

Cut at most one edge! M should be a metric

Which pixels to use?

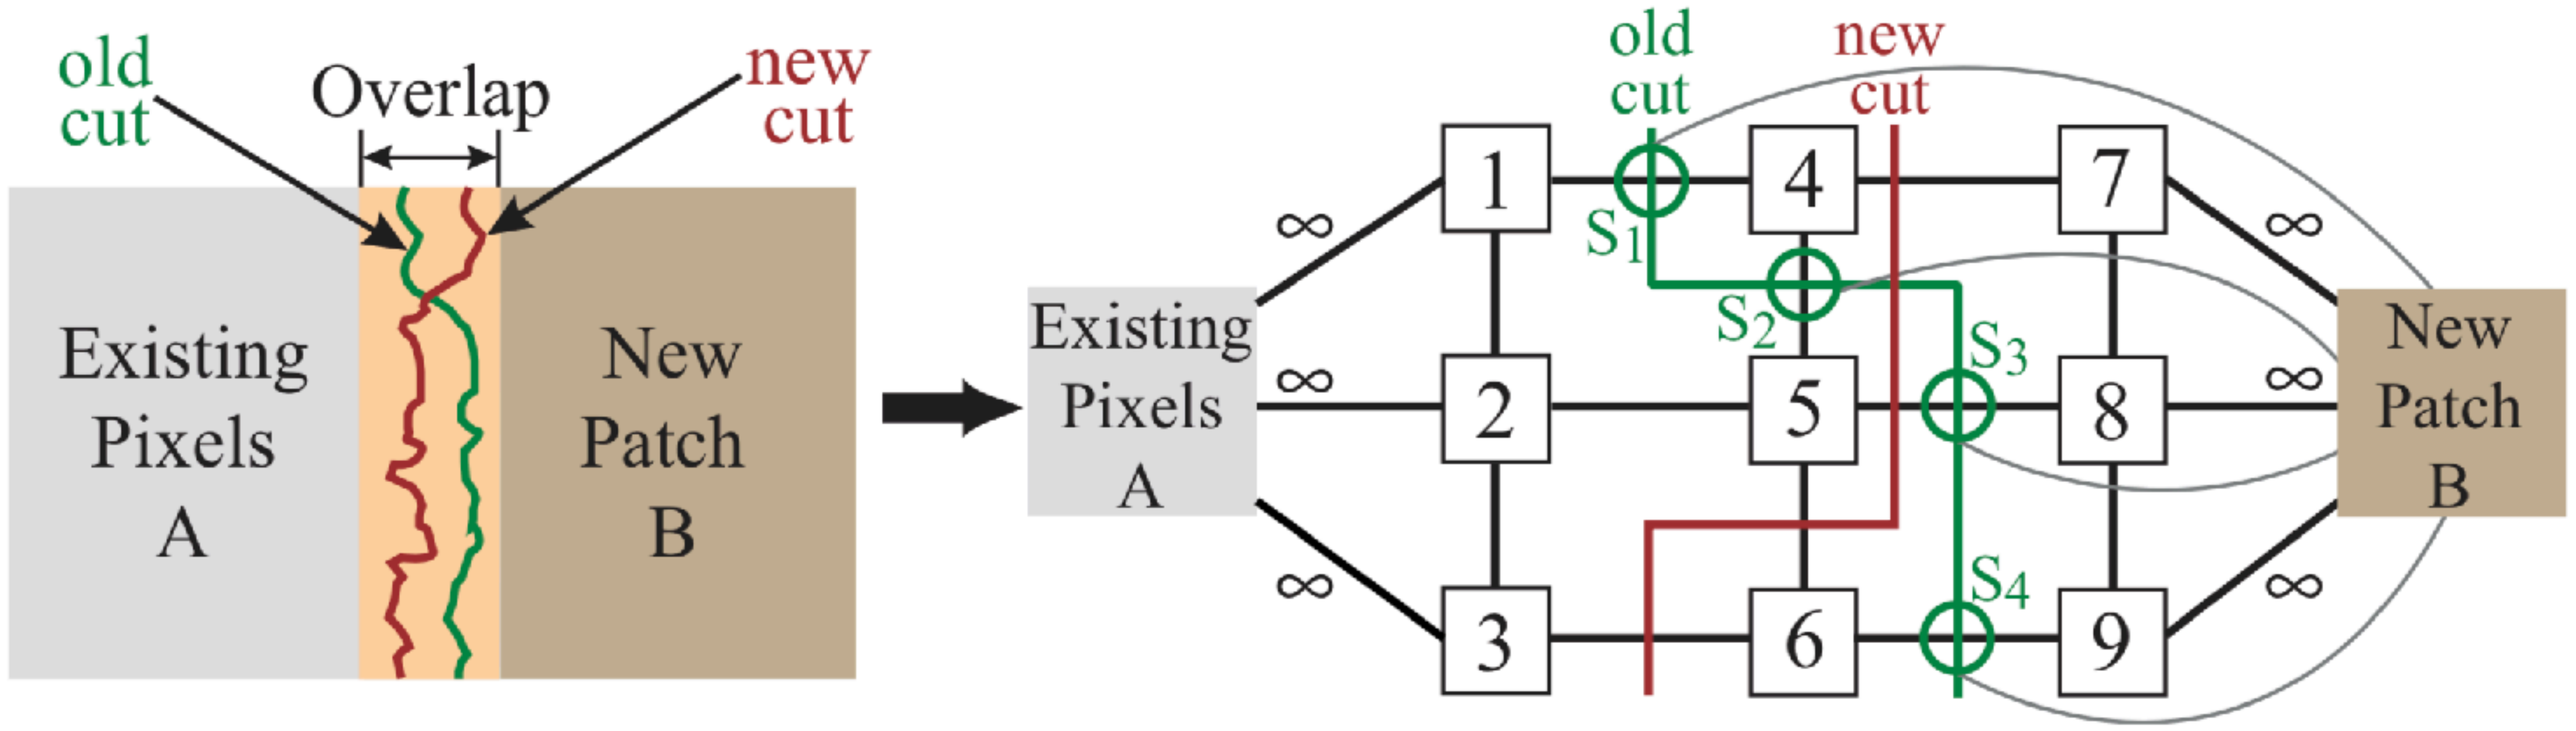

Which pixels to use?

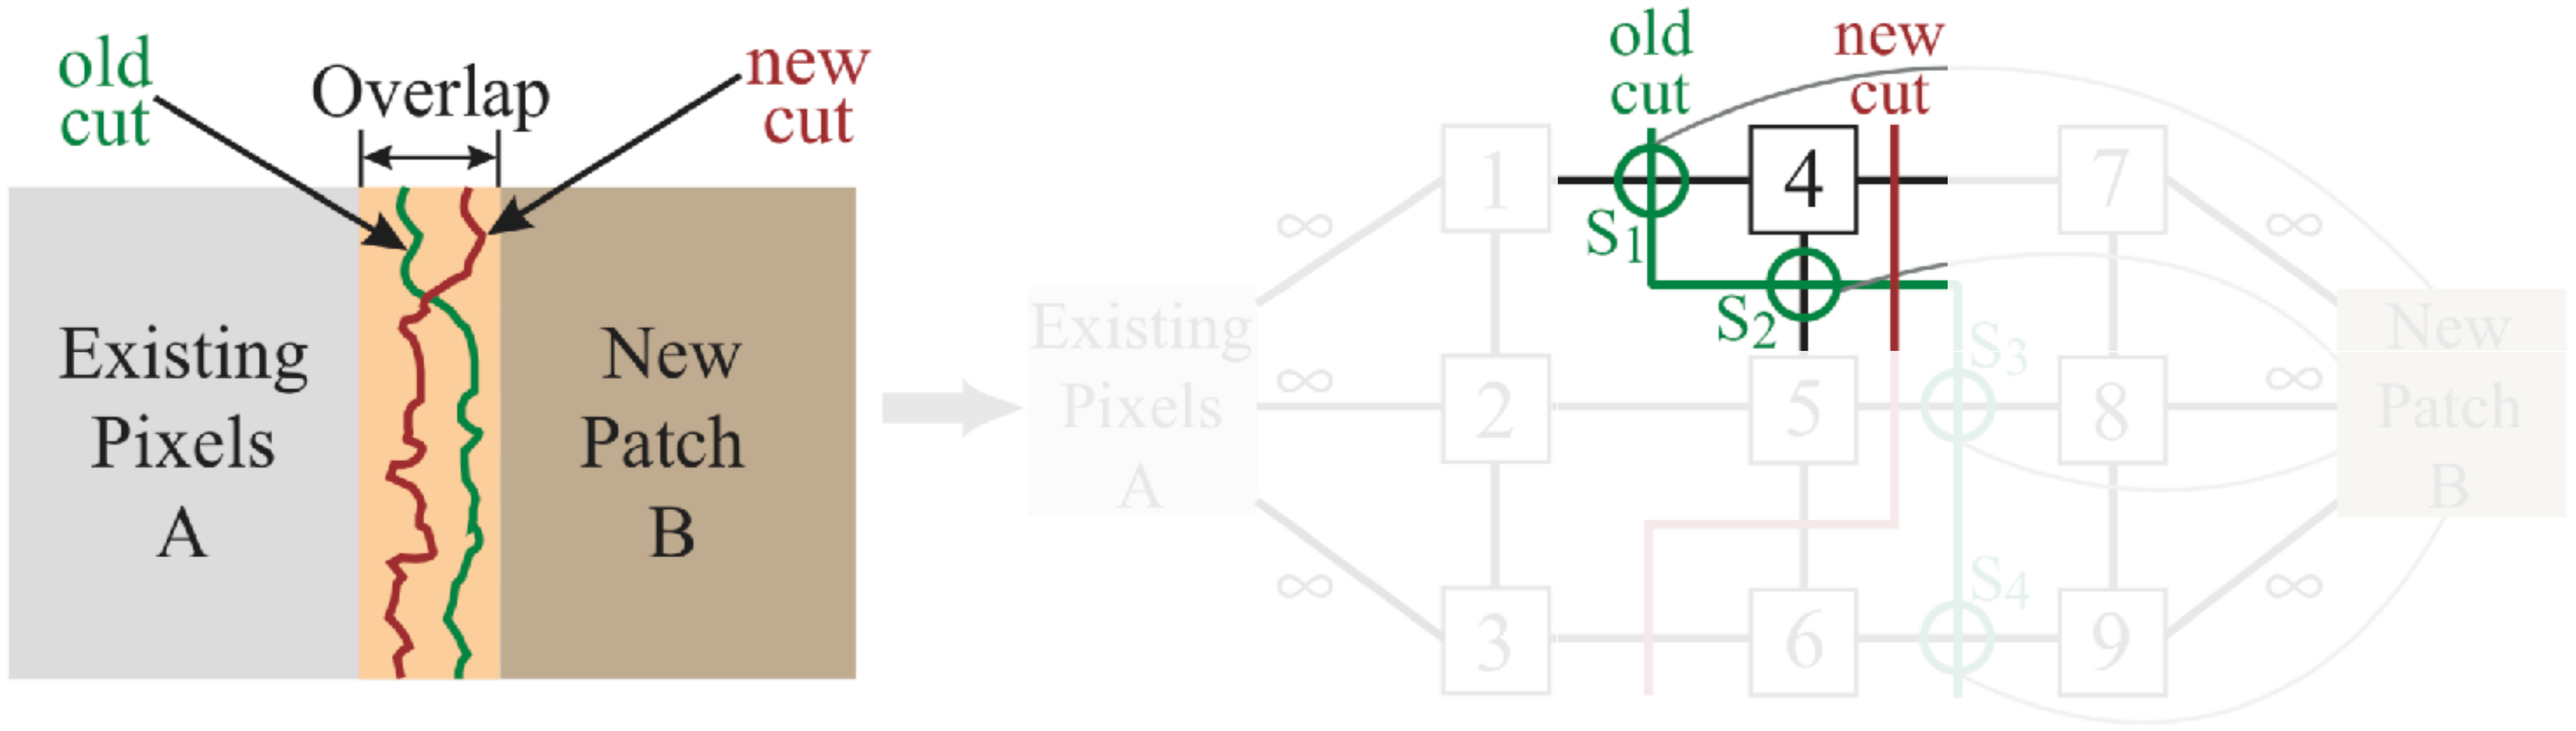

Kept old seam

Which pixels to use?

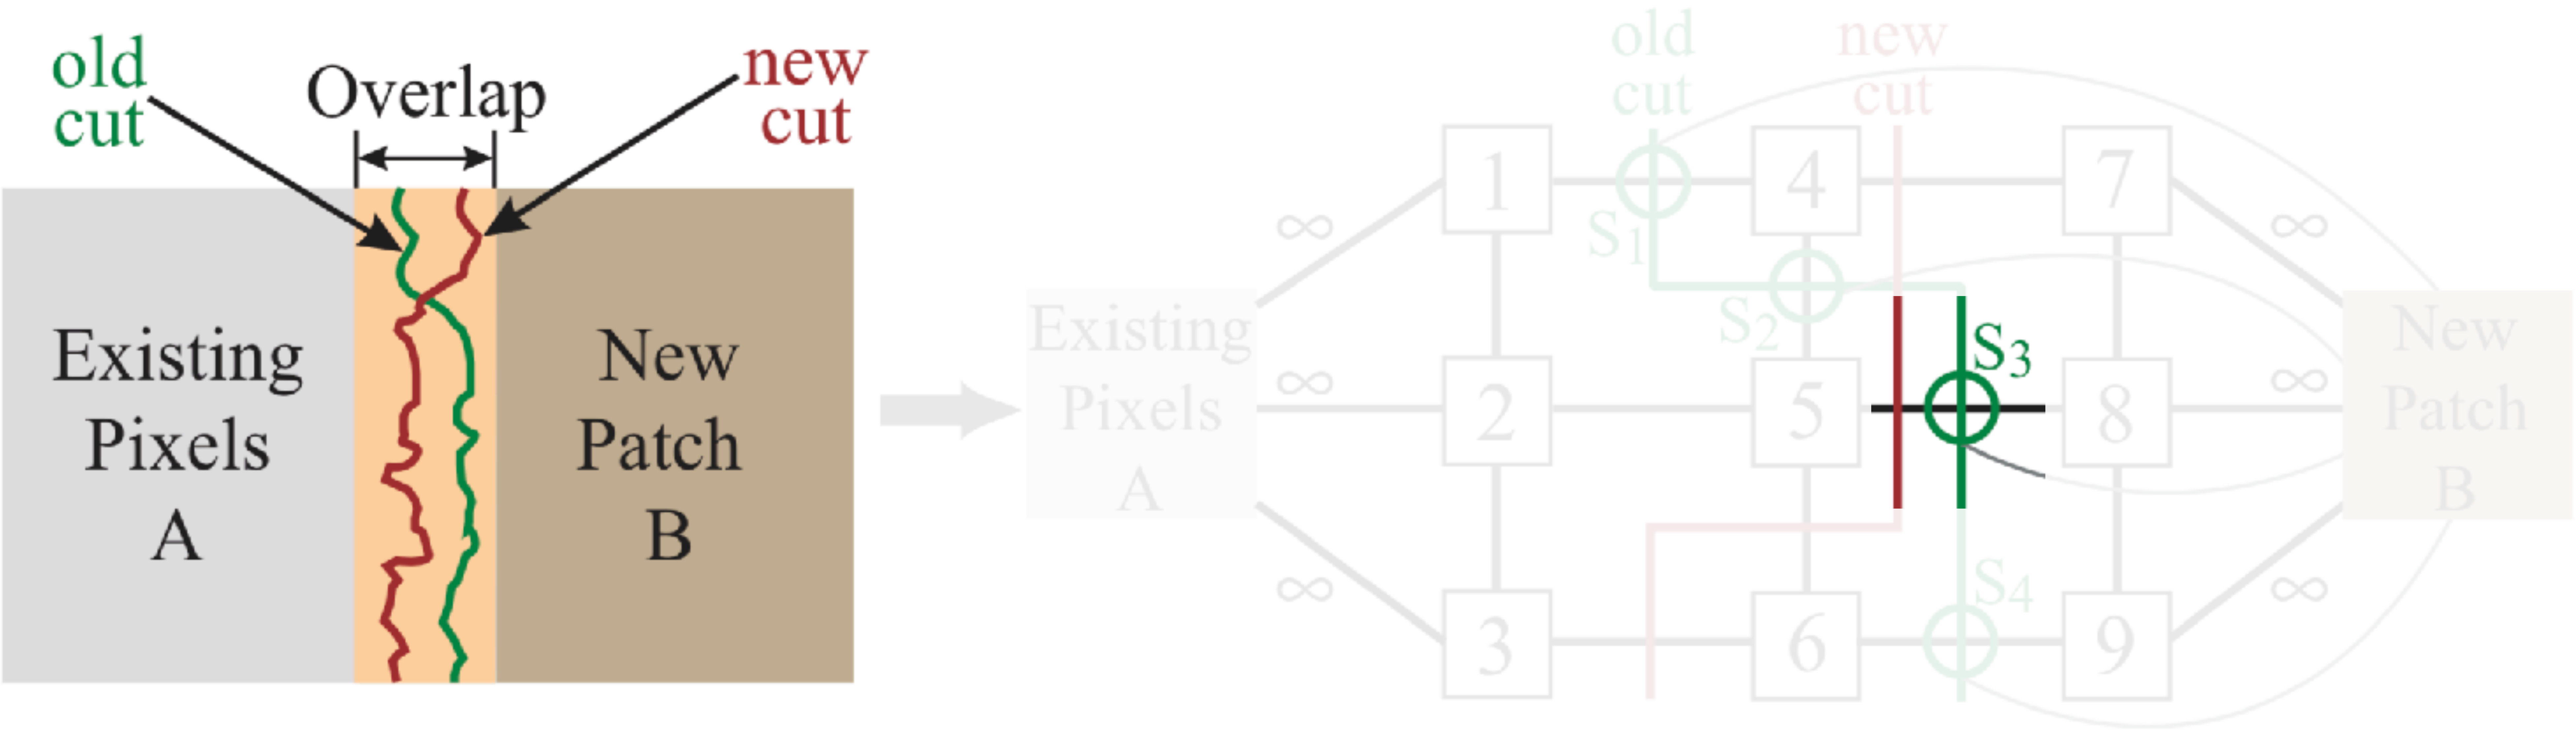

Updated seam

Which pixels to use?

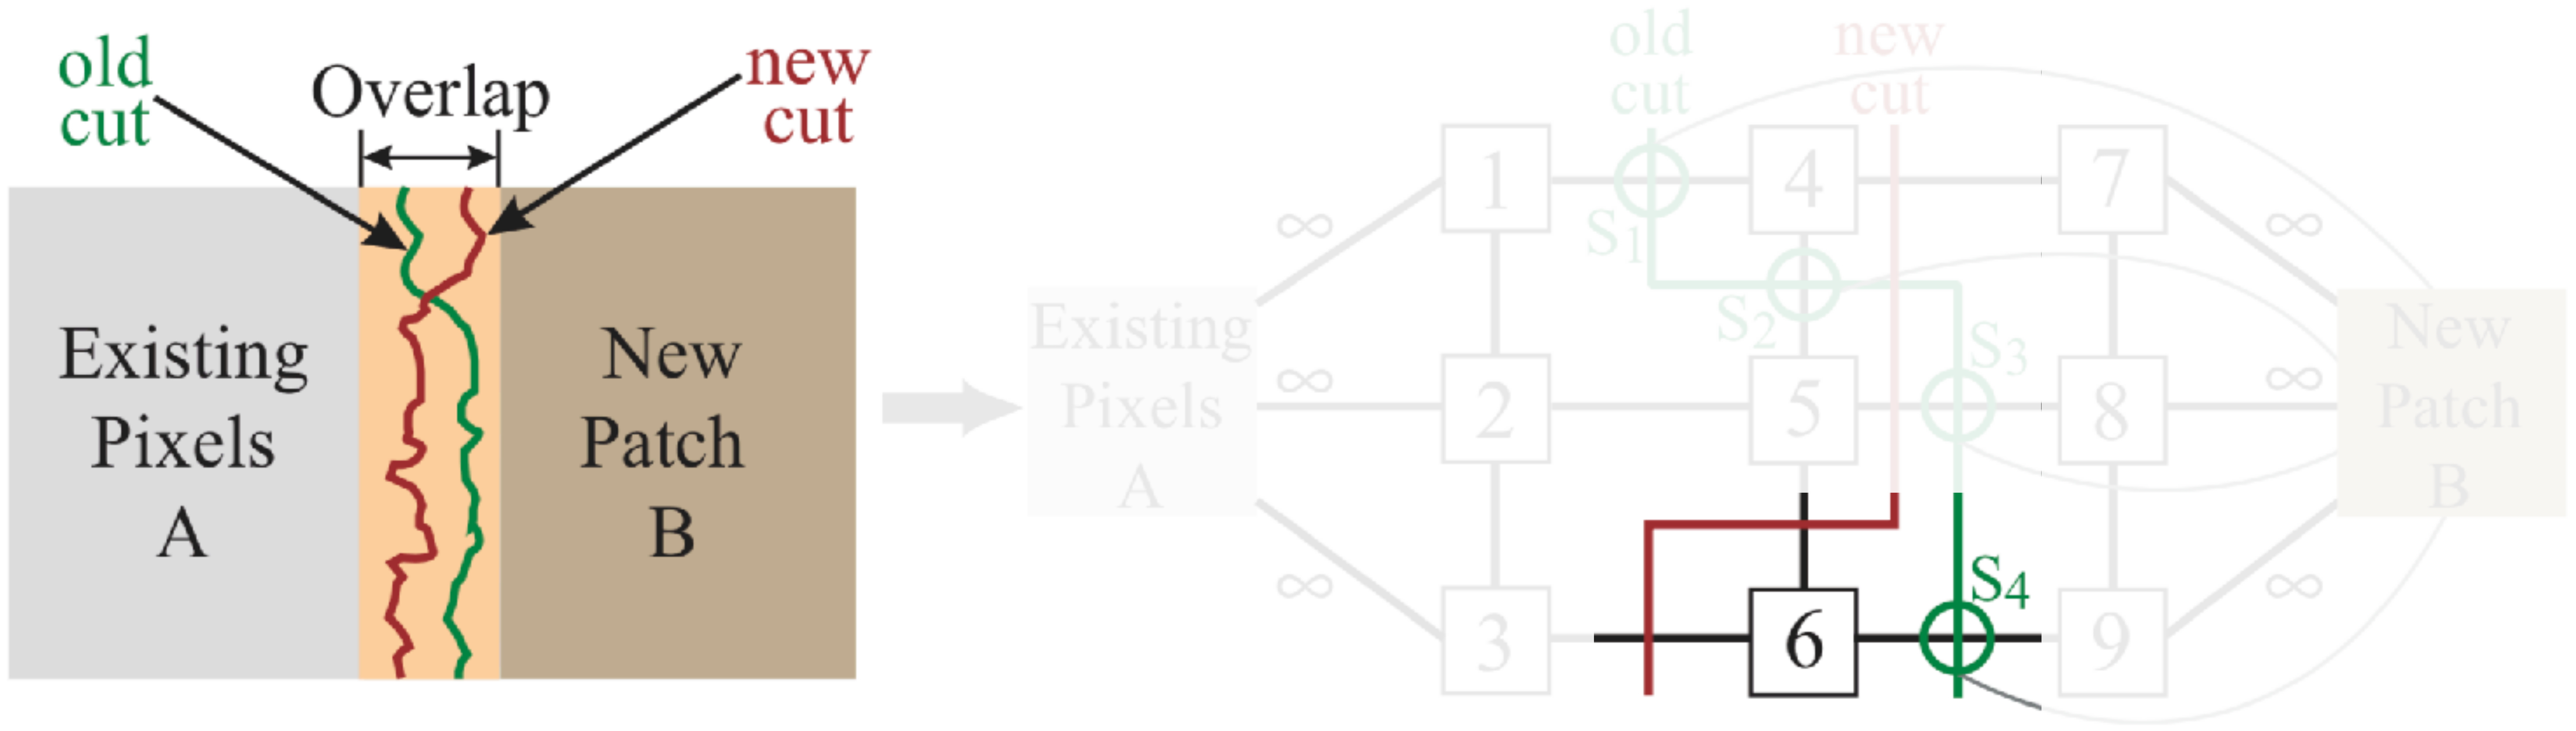

Removed seam

Which pixels to use?

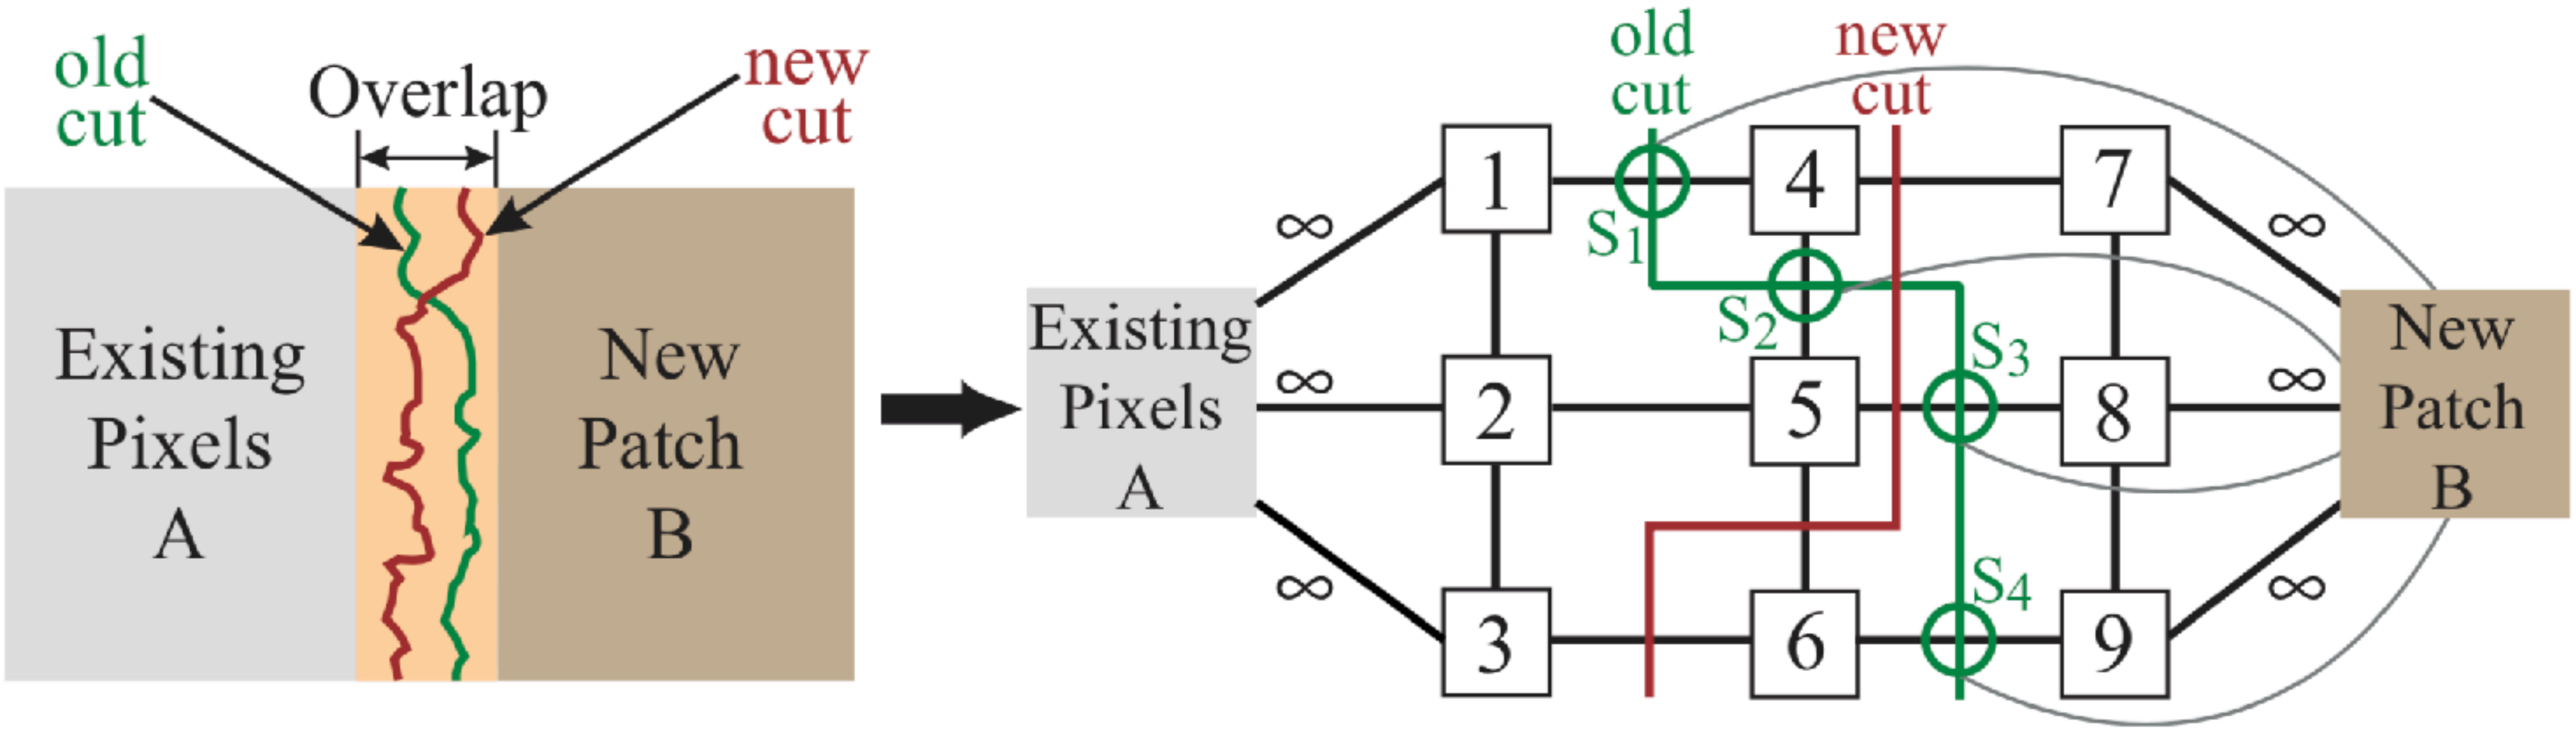

Which pixels to use?

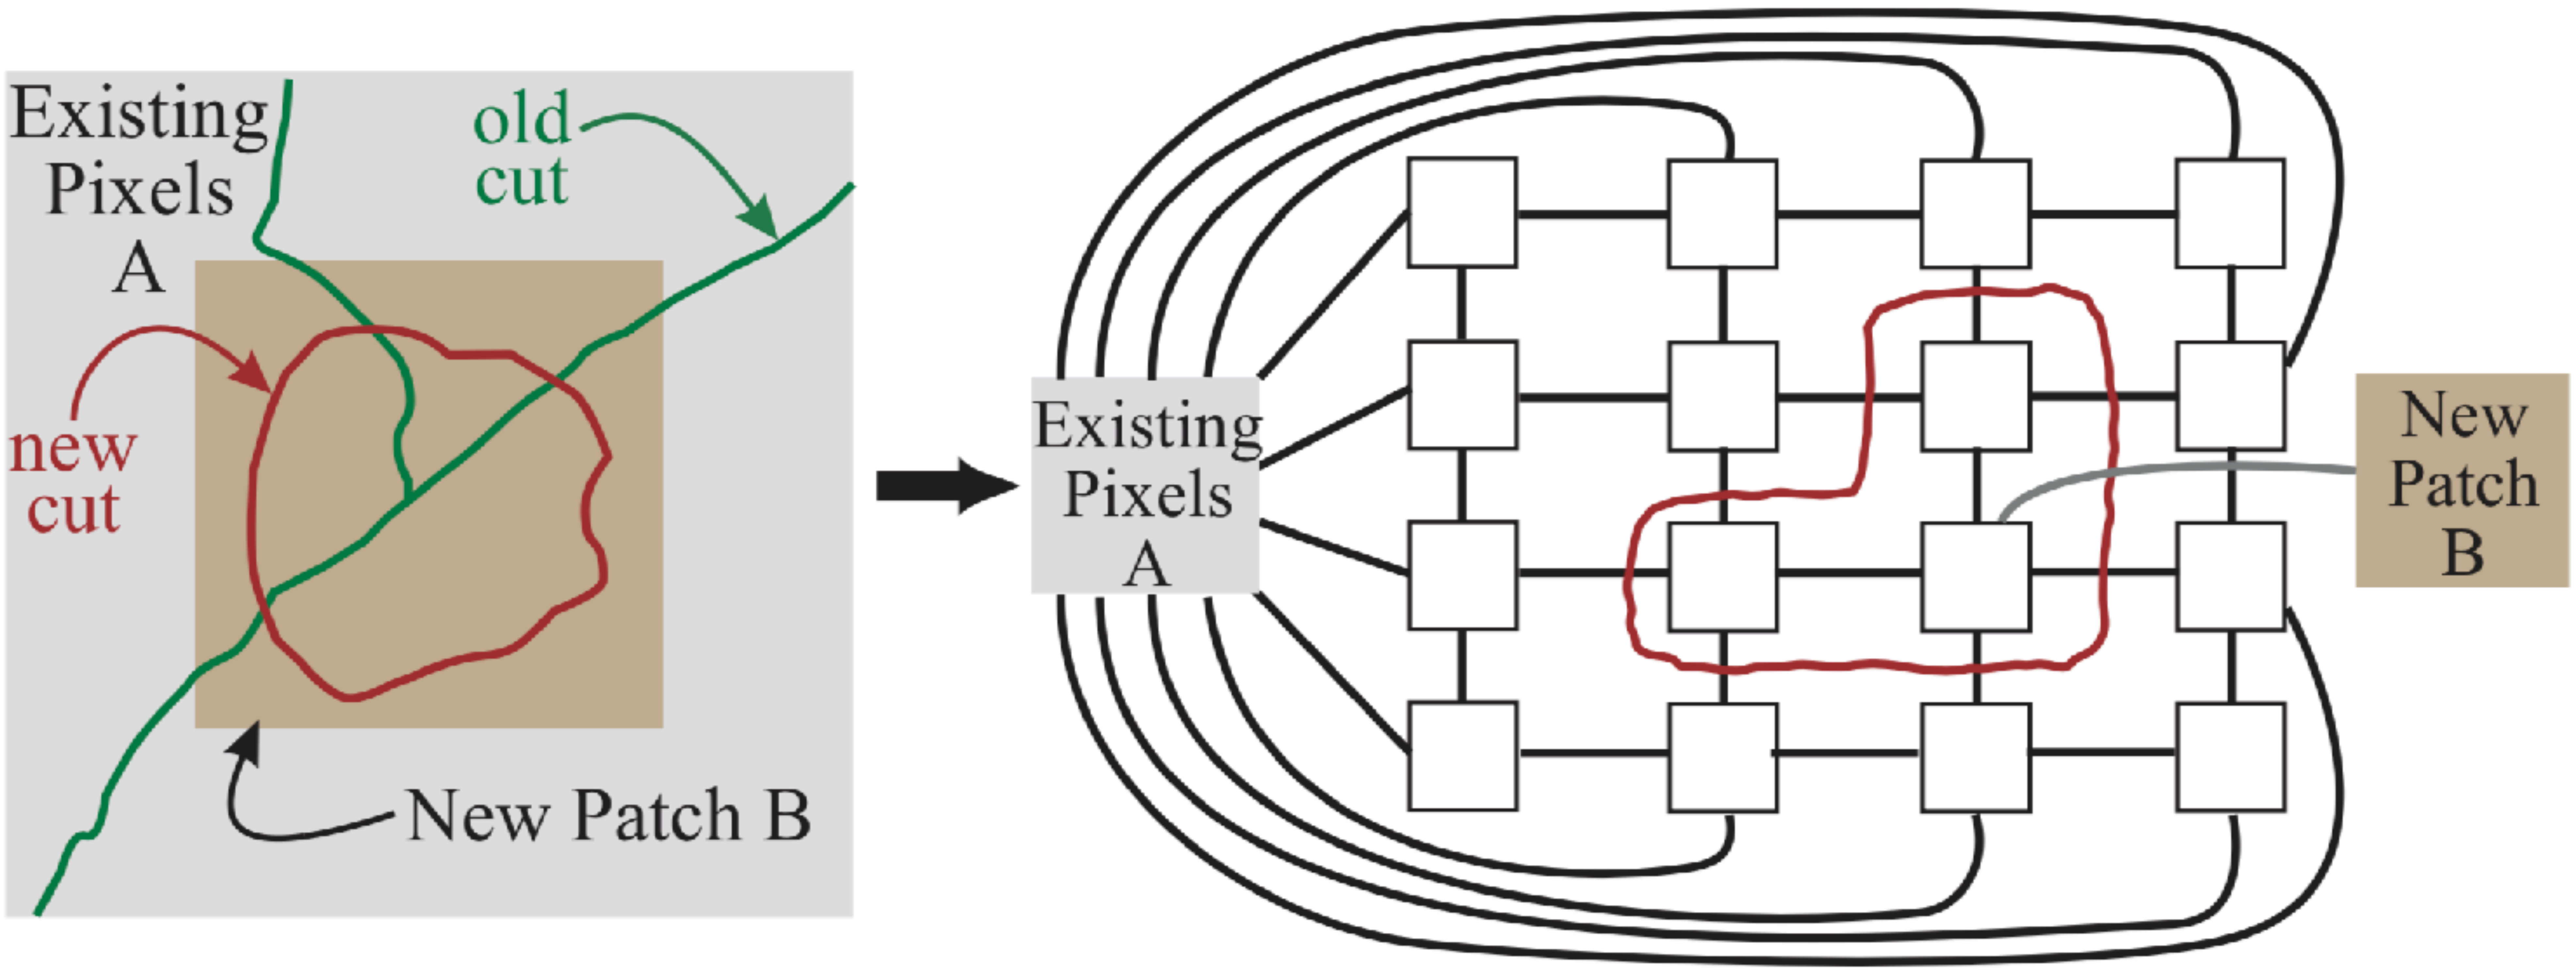

Which pixels to use?

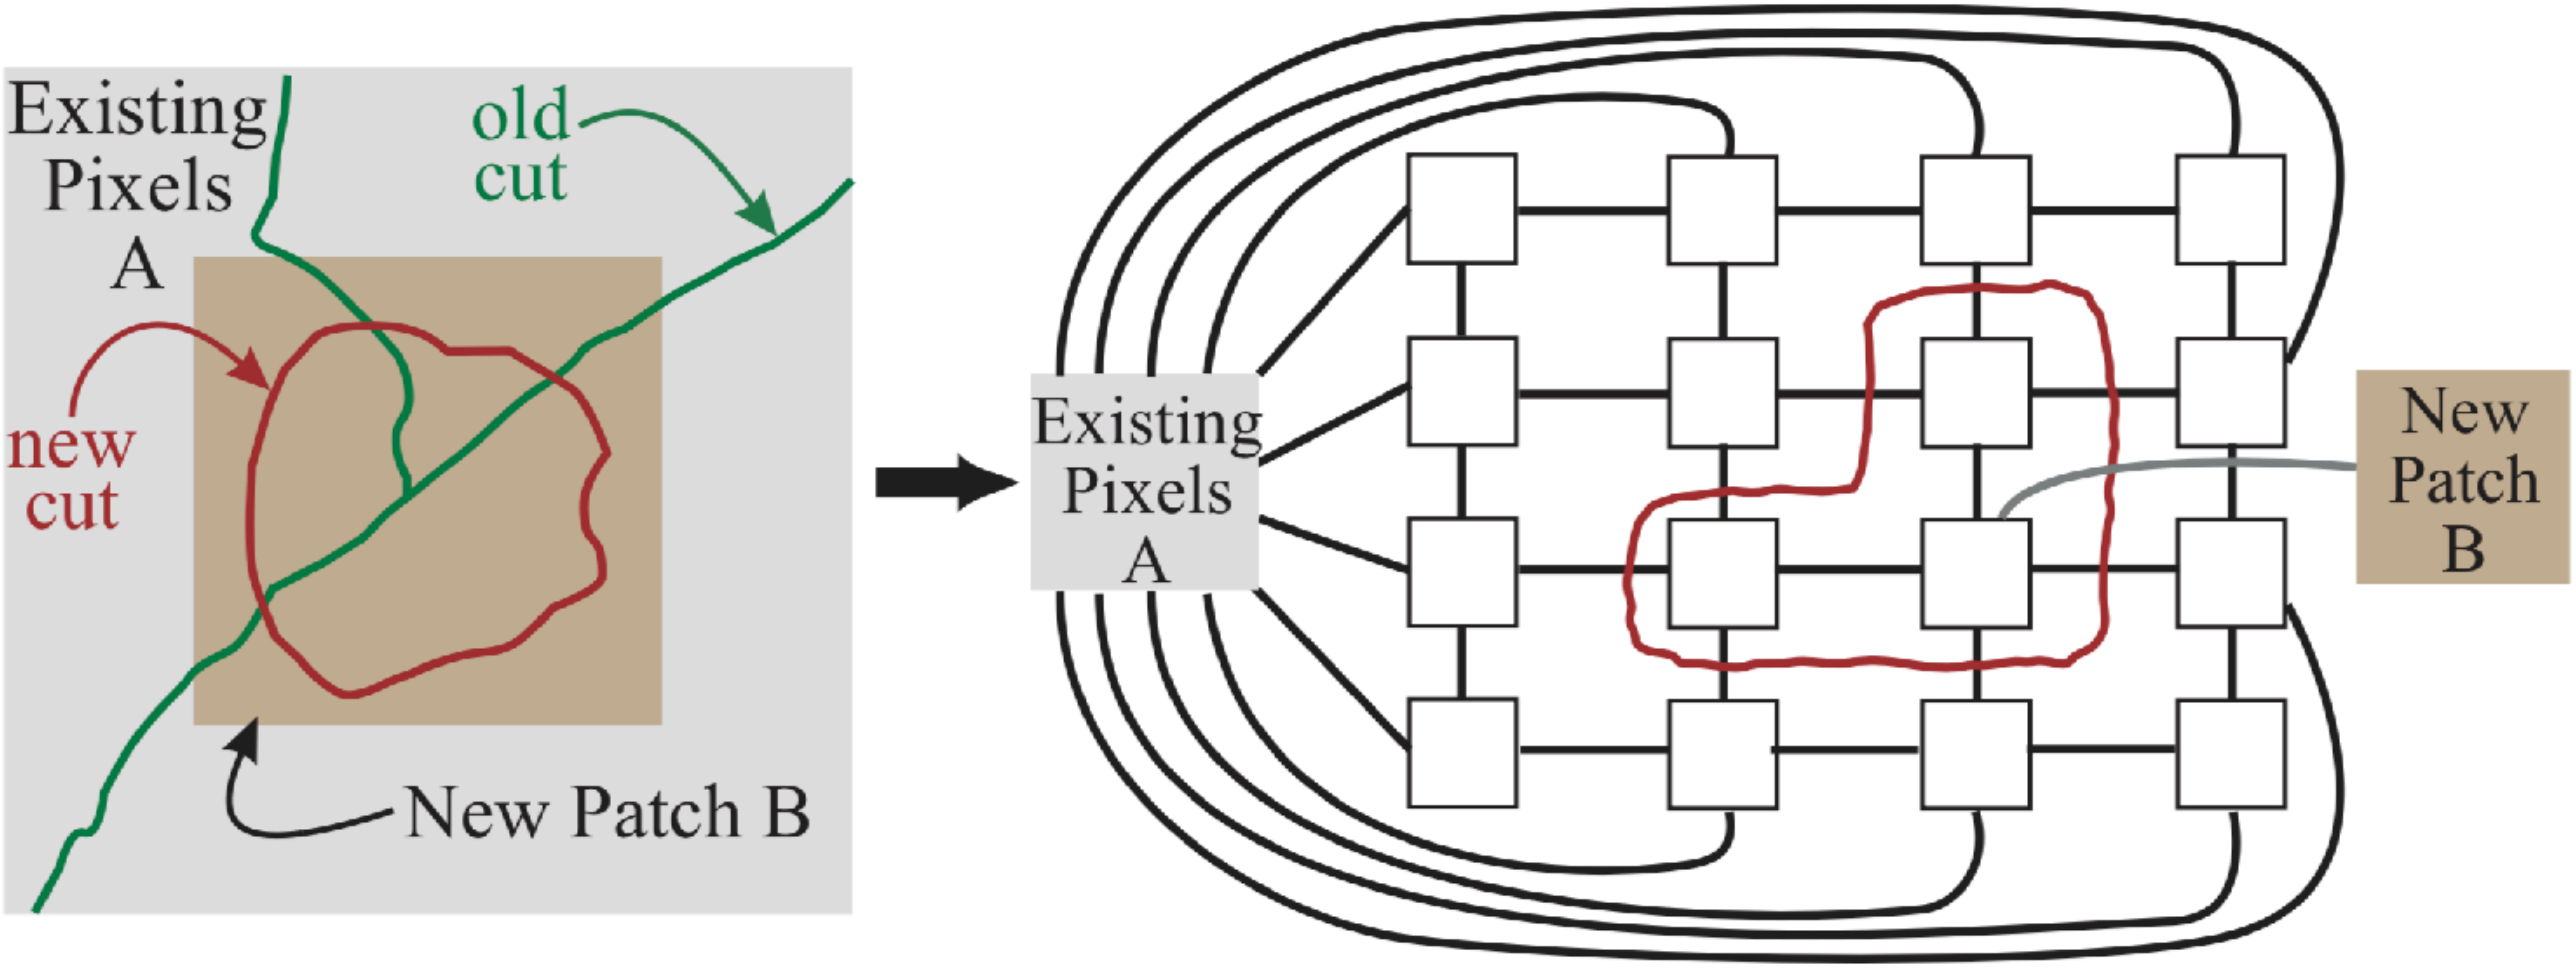

What might happen if we only connect a few pixels to B?

Minor detour: MRFs

Markov Random Field

Markov Random Field

Reminder: Markov property

Markov Random Field

Reminder: Markov property

“memoryless”

Markov Random Field

Reminder: Markov property

"memoryless"

For a discrete process: $P(X_n = x_n | X_{n-1} = x_{n-1}, \ldots, X_0 = x_0) = P(X_n = x_n | X_{n-1} = x_{n-1})$

Markov Random Field

Reminder: Markov property

"memoryless"

For a discrete process: $P(X_n = x_n | X_{n-1} = x_{n-1}, \ldots, X_0 = x_0) = P(X_n = x_n | X_{n-1} = x_{n-1})$

What about fields?

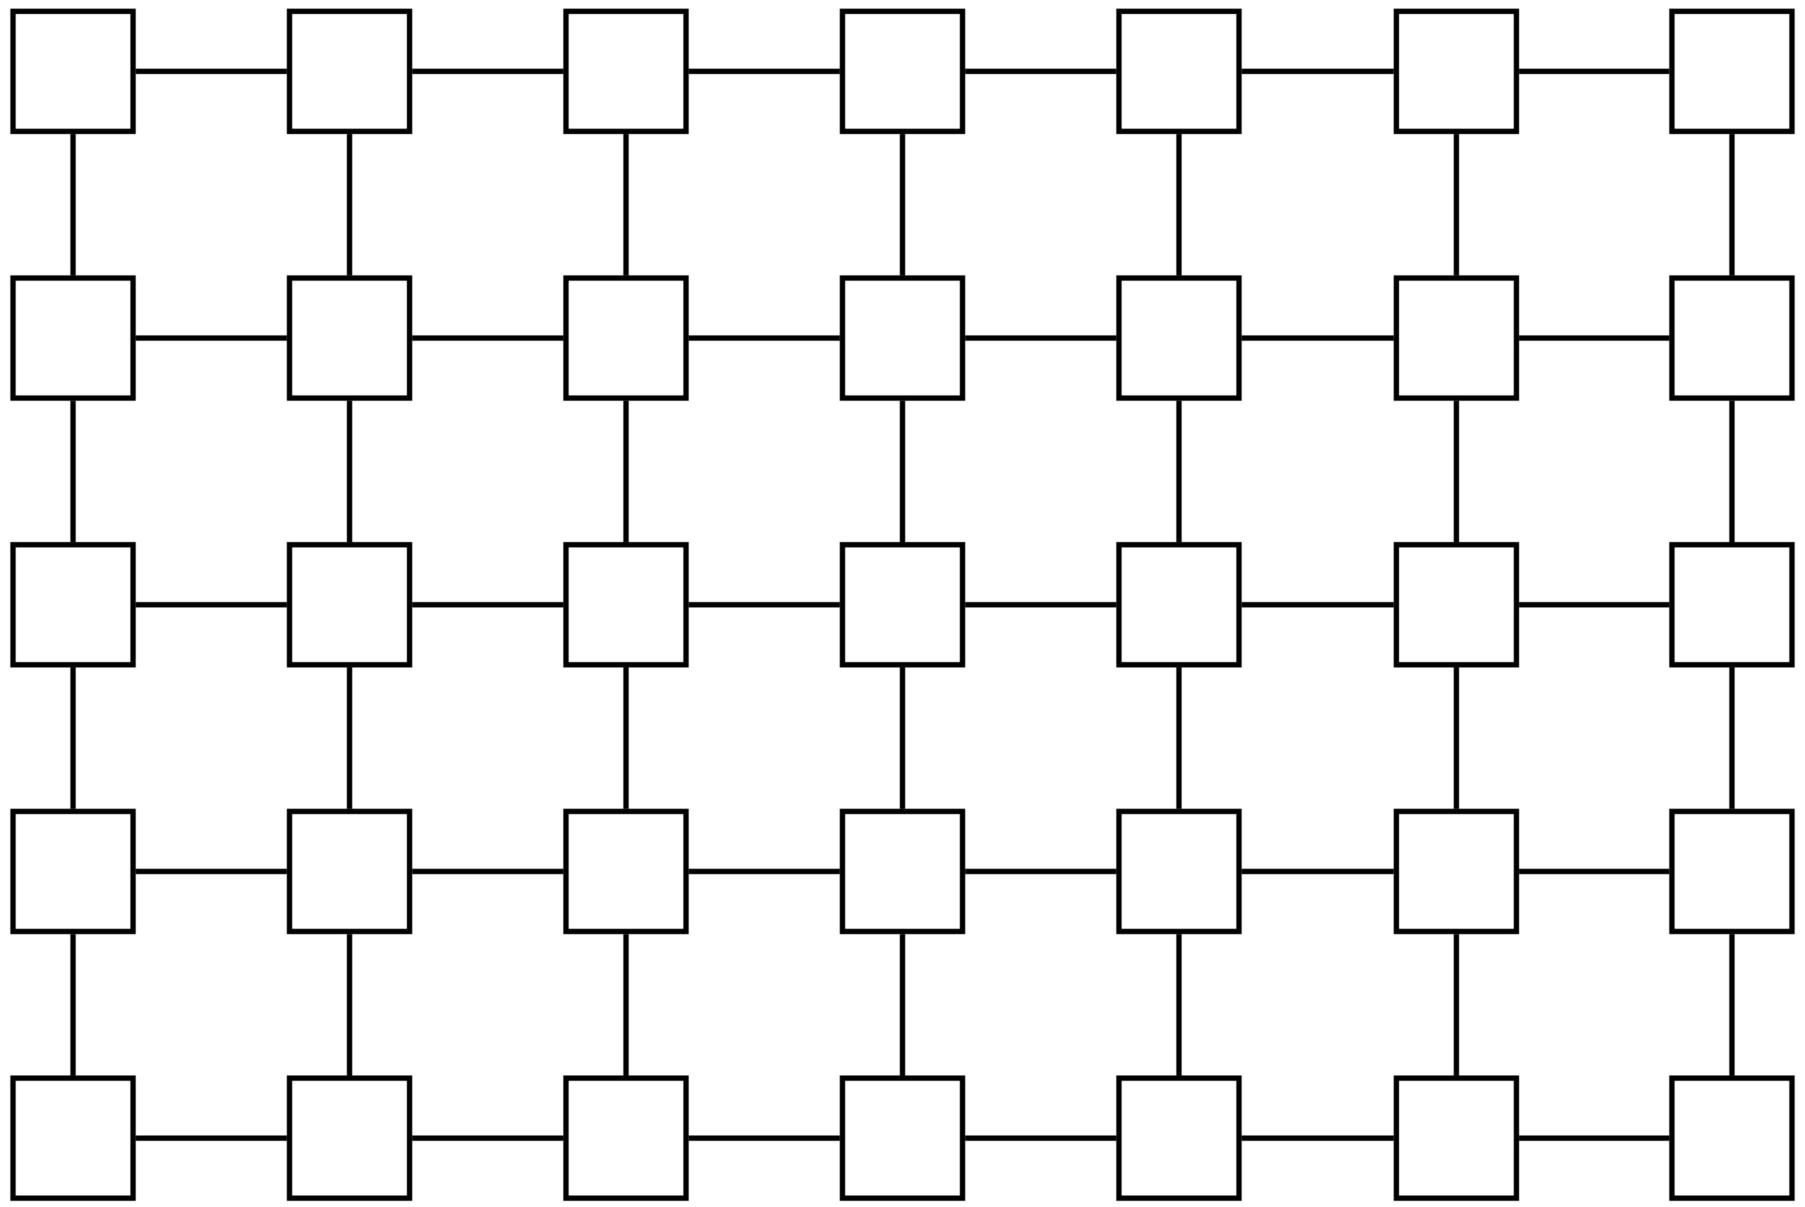

Markov Random Field

Reminder: Markov property

"memoryless"

For a discrete process: $P(X_n = x_n \,|\, X_{n-1} = x_{n-1}, \ldots, X_0 = x_0) = P(X_n = x_n \,|\, X_{n-1} = x_{n-1})$

What about fields?

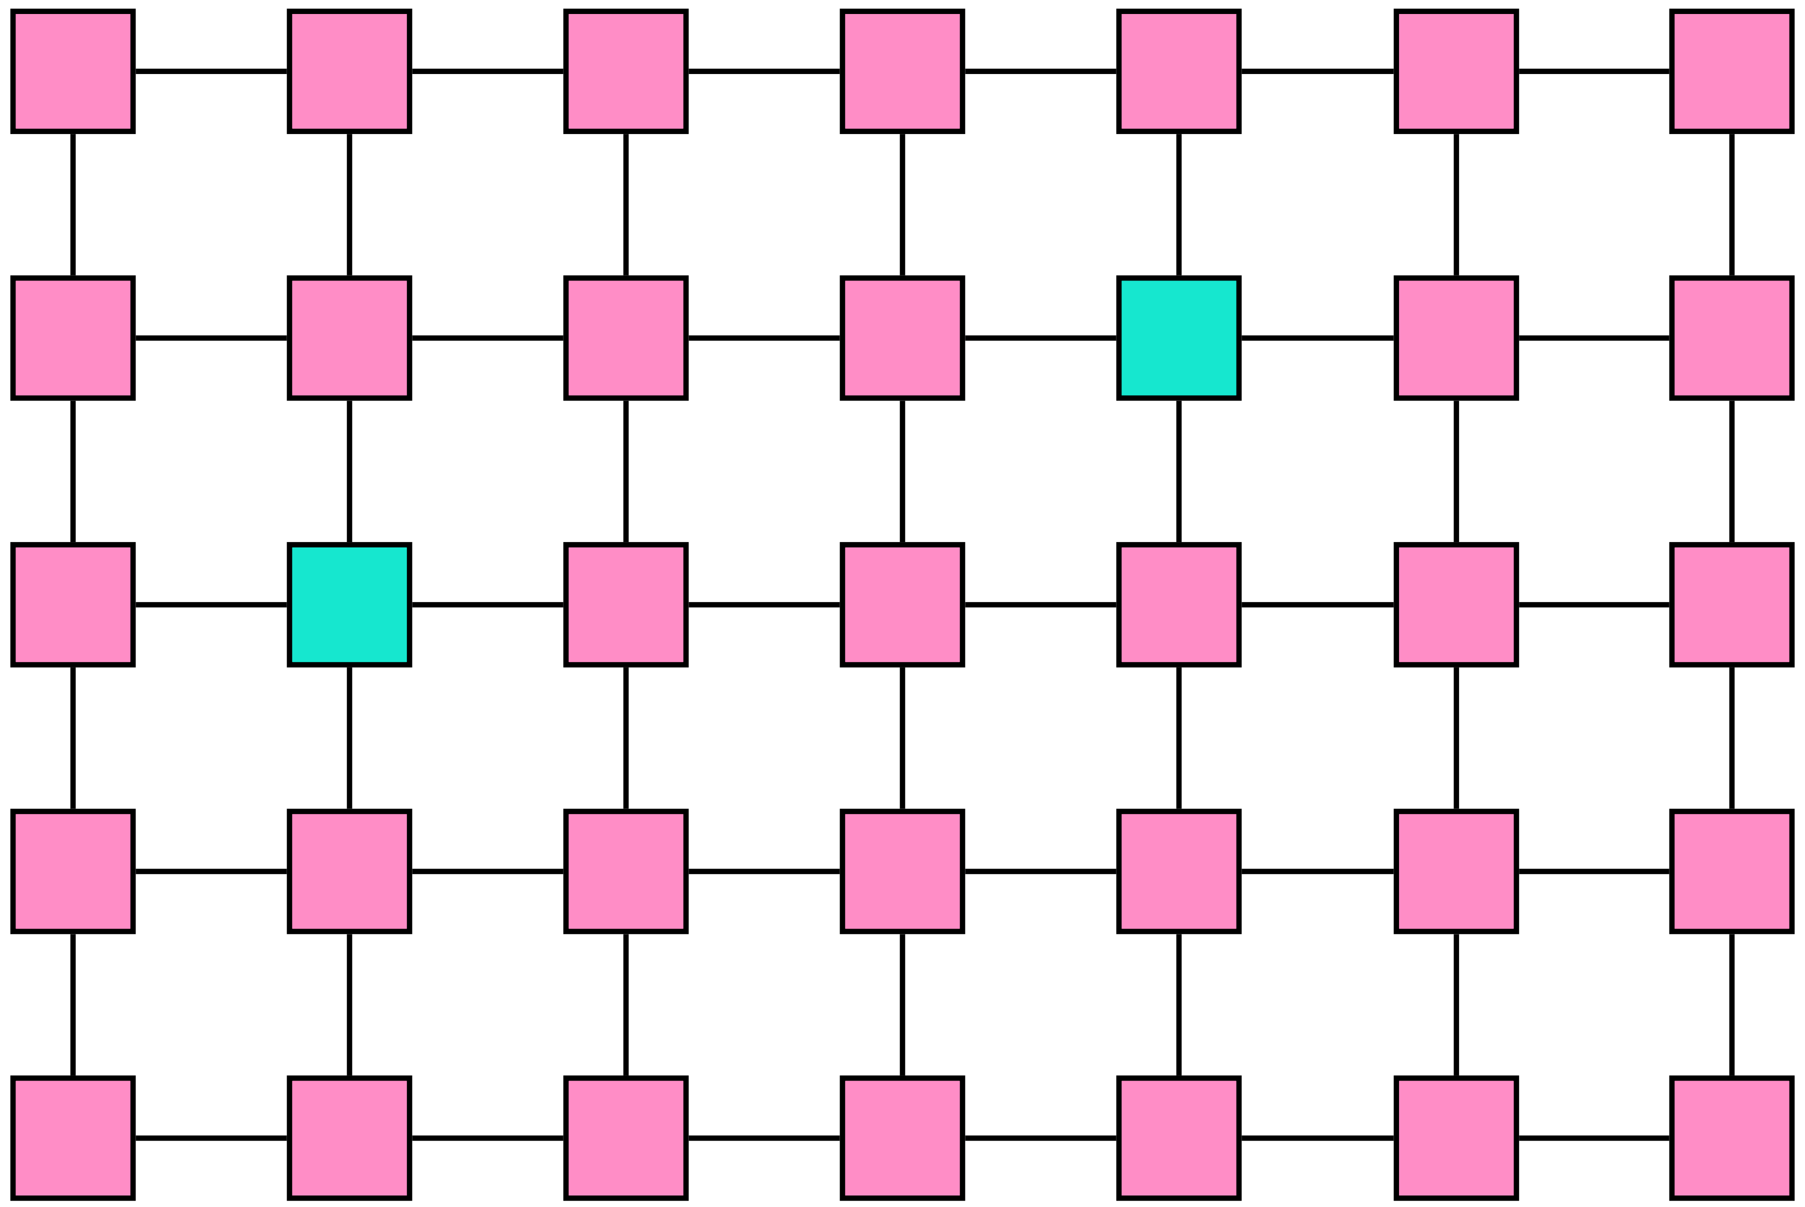

two non-adjacent variables are conditionally independent given all other variables

Markov Random Field

Reminder: Markov property

“memoryless”

For a discrete process: $P(X_n = x_n \mid X_{n-1} = x_{n-1}, \ldots, X_0 = x_0) = P(X_n = x_n \mid X_{n-1} = x_{n-1})$

What about fields?

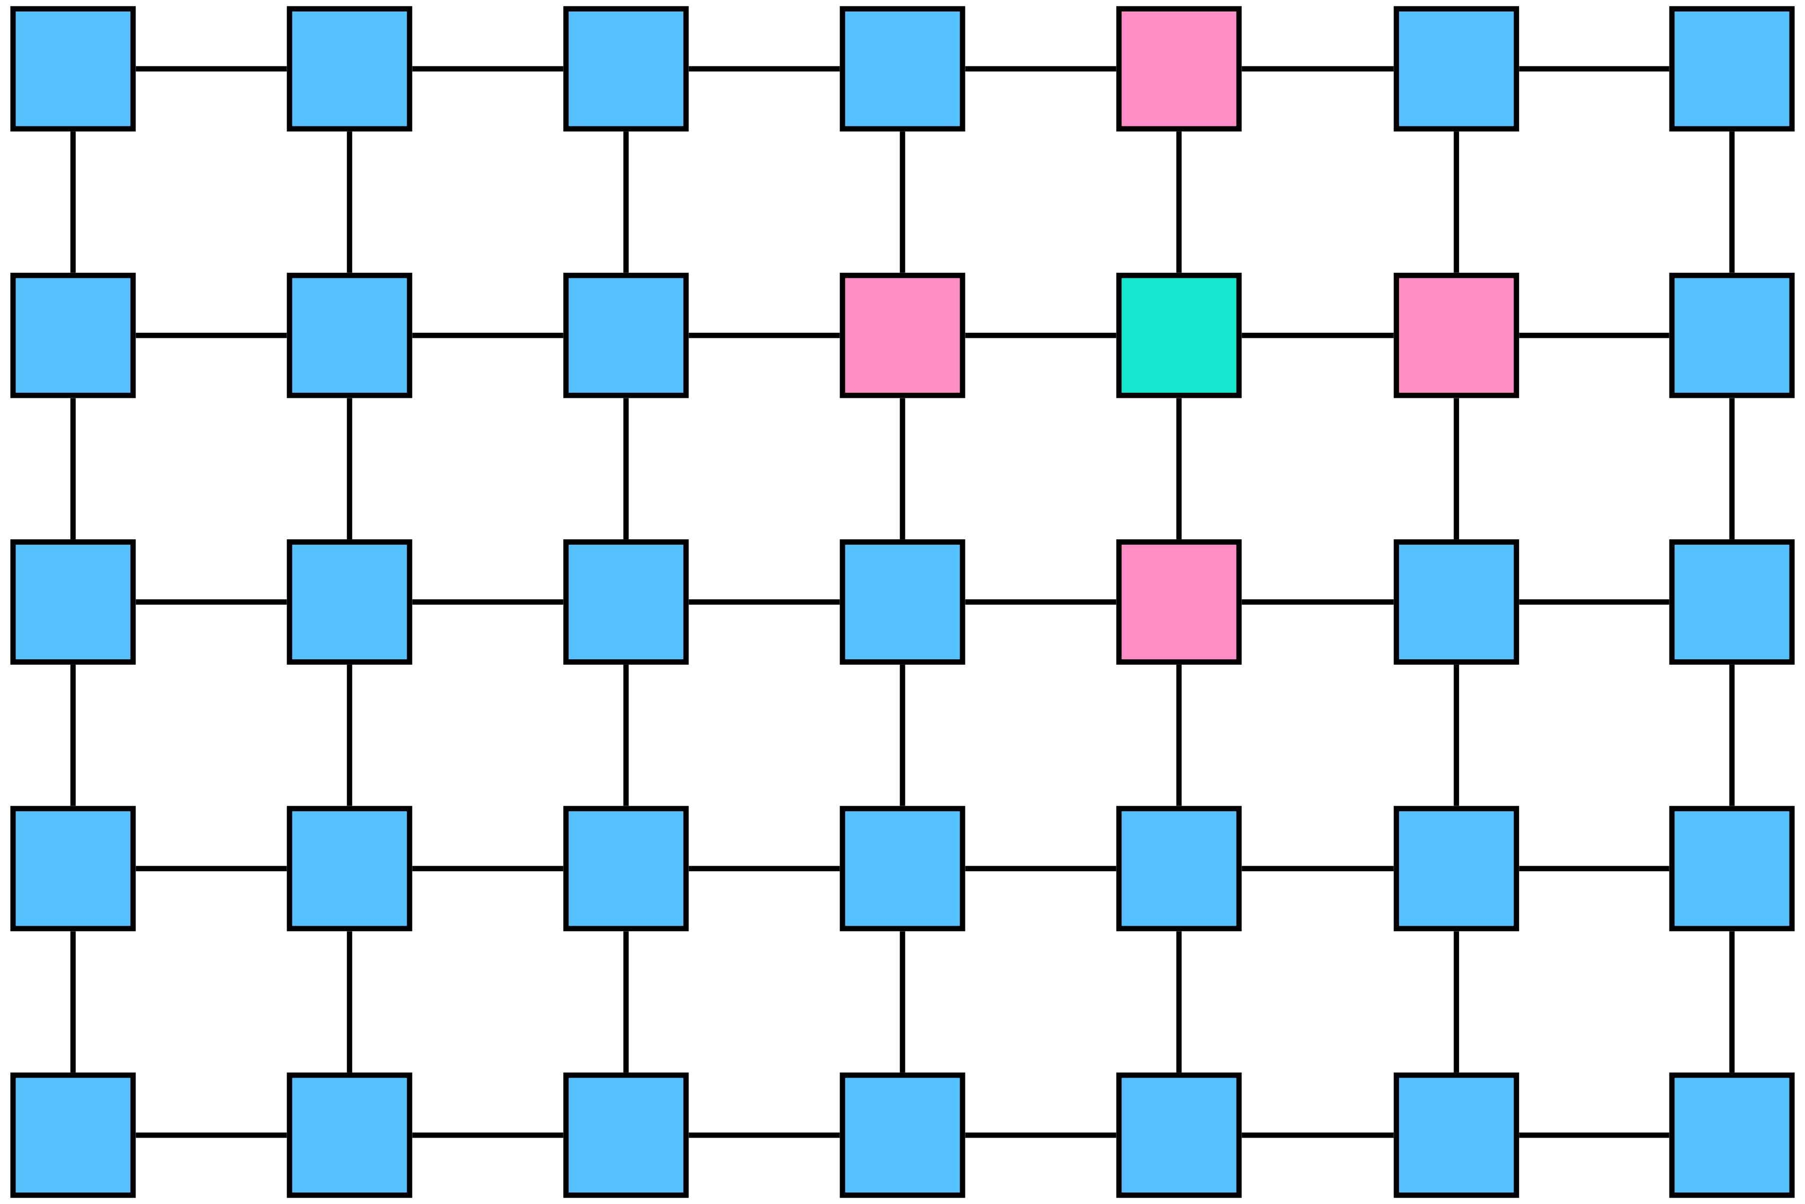

A variable
is conditionally independent
of all other variables
given its neighbors

Markov Random Field

Reminder: Markov property

“memoryless”

For a discrete process: $P(X_n = x_n \mid X_{n-1} = x_{n-1}, \ldots, X_0 = x_0) = P(X_n = x_n \mid X_{n-1} = x_{n-1})$

What about fields?

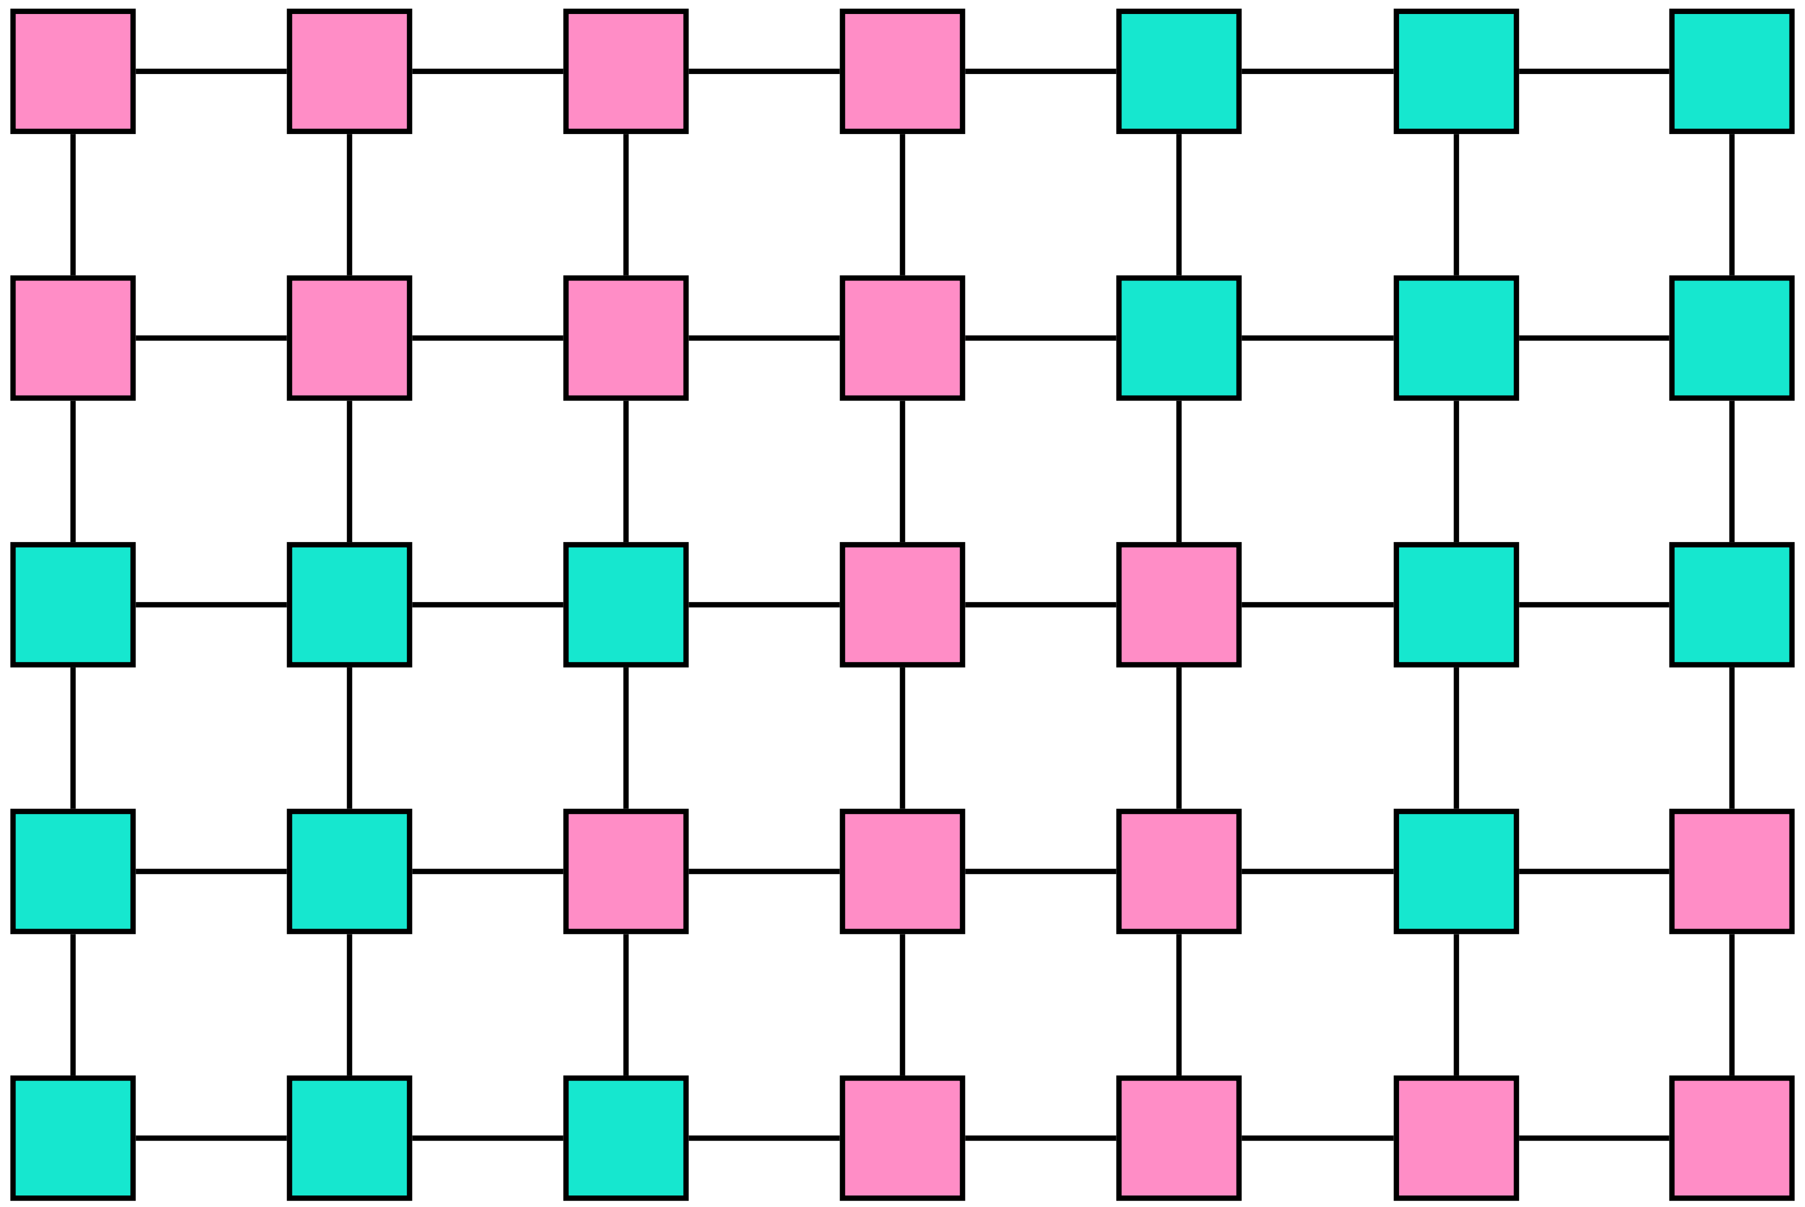

two subsets
are conditionally independent
given a separating subset

Markov Random Field

What does it mean in our setting?

Where to place next patch?

Which pixels to use?

Placing the next patch

Placing the next patch

- Random placement

Placing the next patch

- Random placement
- Entire patch matching

Placing the next patch

- Random placement
- Entire patch matching
- Sub-patch matching

Placing the next patch

- Random placement
- Entire patch matching
- Sub-patch matching

What would be the “right” thing to do, assuming no runtime constraints?

Results


describing the response of that neuro
ht as a function of position—is perhap
functional description of that neuron.
seek a single conceptual and mathem
scribe the wealth of simple-cell recep
d neurophysiologically[1–3] and inferred
especially if such a framework has the
it helps us to understand the functio
leeper way. Whereas no generic mo
ussians (DOG), difference of offset
rivative of a Gaussian, higher derivati
function, and so on—can be expecte
imple-cell receptive field, we noneth
describing the response of that neurophysiologically[1–3] and
ht as a function of position—is perhally if such a framework
functional description of that neuron. us to understand the
seek a single conceptual and mathr way. Whereas no gen
scribe the wealth of simple-cell ians (DOG), difference of
d neurophysiologically[1–3] and ivative of a Ga response of tha
especially if such a framework functionnction of position—i
t helps us to understand the funeional description of that
eeper way. Whereas no generick a single conceptual and
ussians (DOG), difference of a function of position—is per
rivative of a Gaussian, higher donal description of that neur
he response od so on—can be a single conceptual and math
scribing the response of that ne the wealth of simple-cell r
as a function of position—is perbphysiologically[1–3] and infe
ctional description of that neurony if such a framework has
ek a single conceptual and mathems to understand the fun
ribe the wealth of simple-onceptual Whereas no generic
neurophysiologically[1–3] and th of simple), difference of offs
pecially if such a frameworlogically[1–3] Gaussian, higher deri
helps us to understand such a framewor so on—can be exp
per way. Whereas us to understand the fun field, we nor

Input

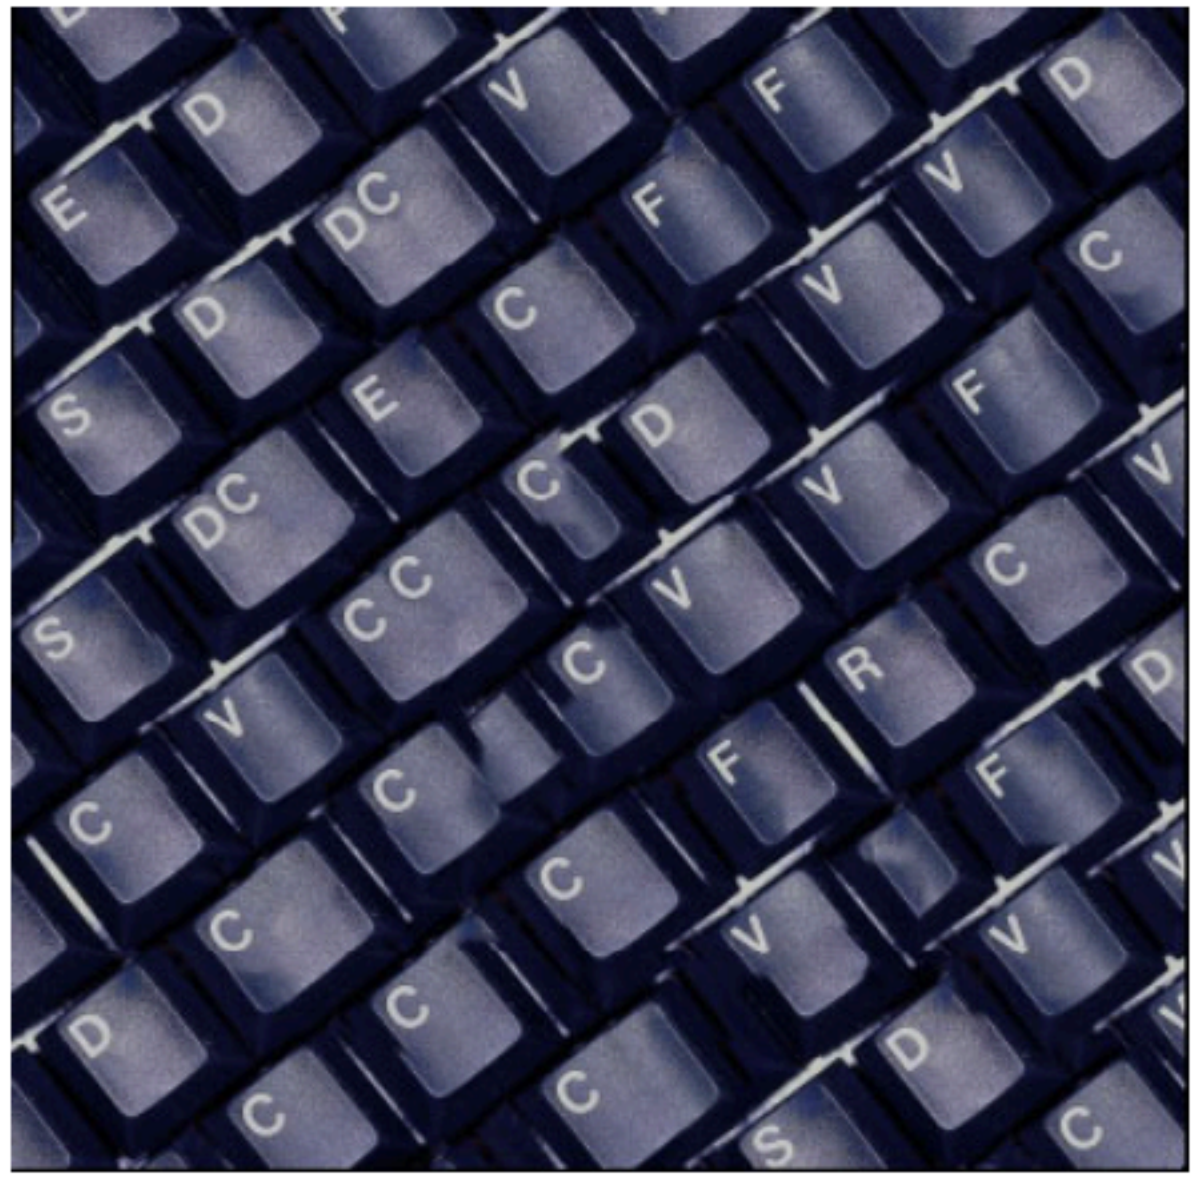

Image Quilting

Graph cut

Input

Image Quilting

Graph cut

Rotation & Mirroring

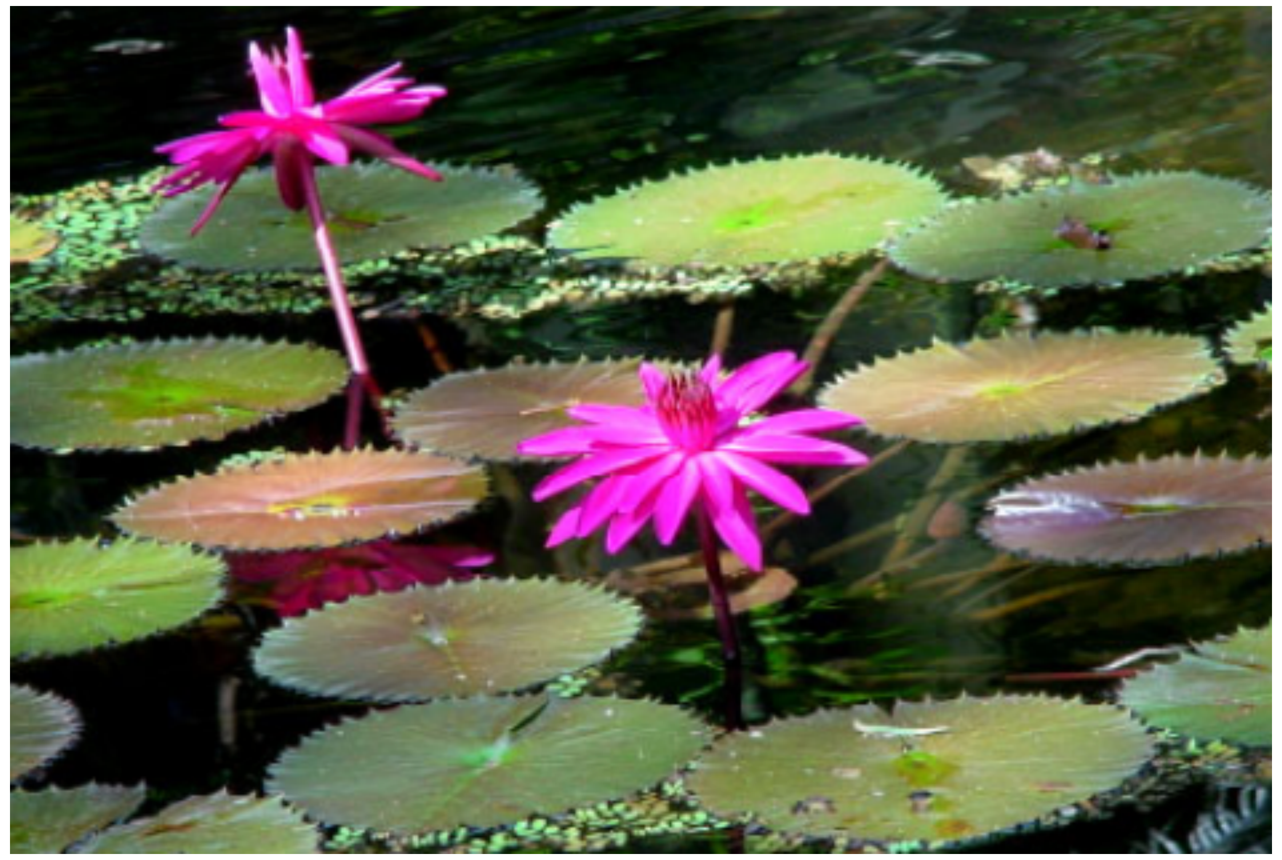

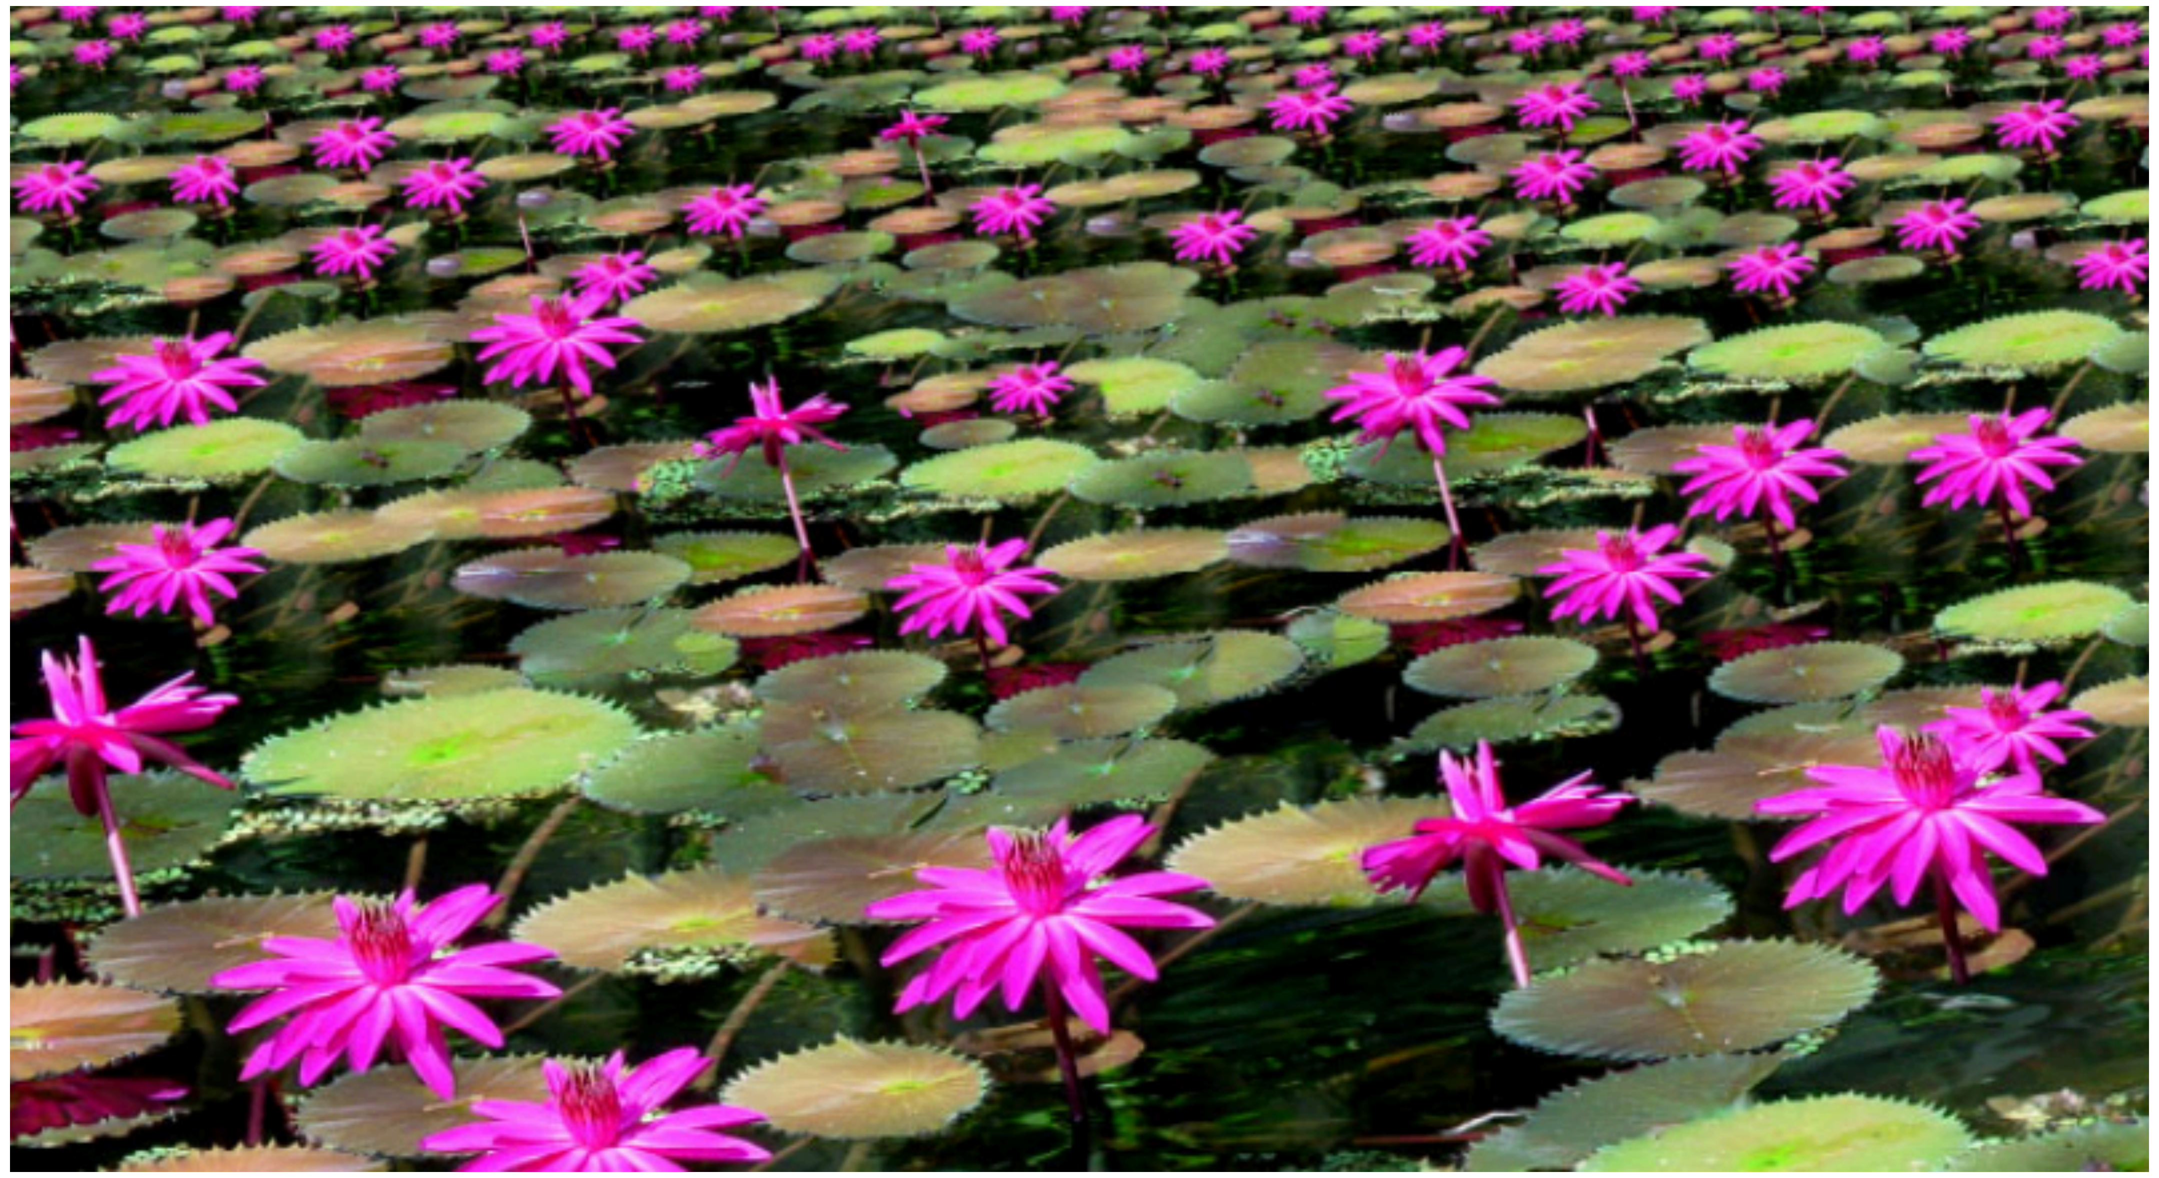

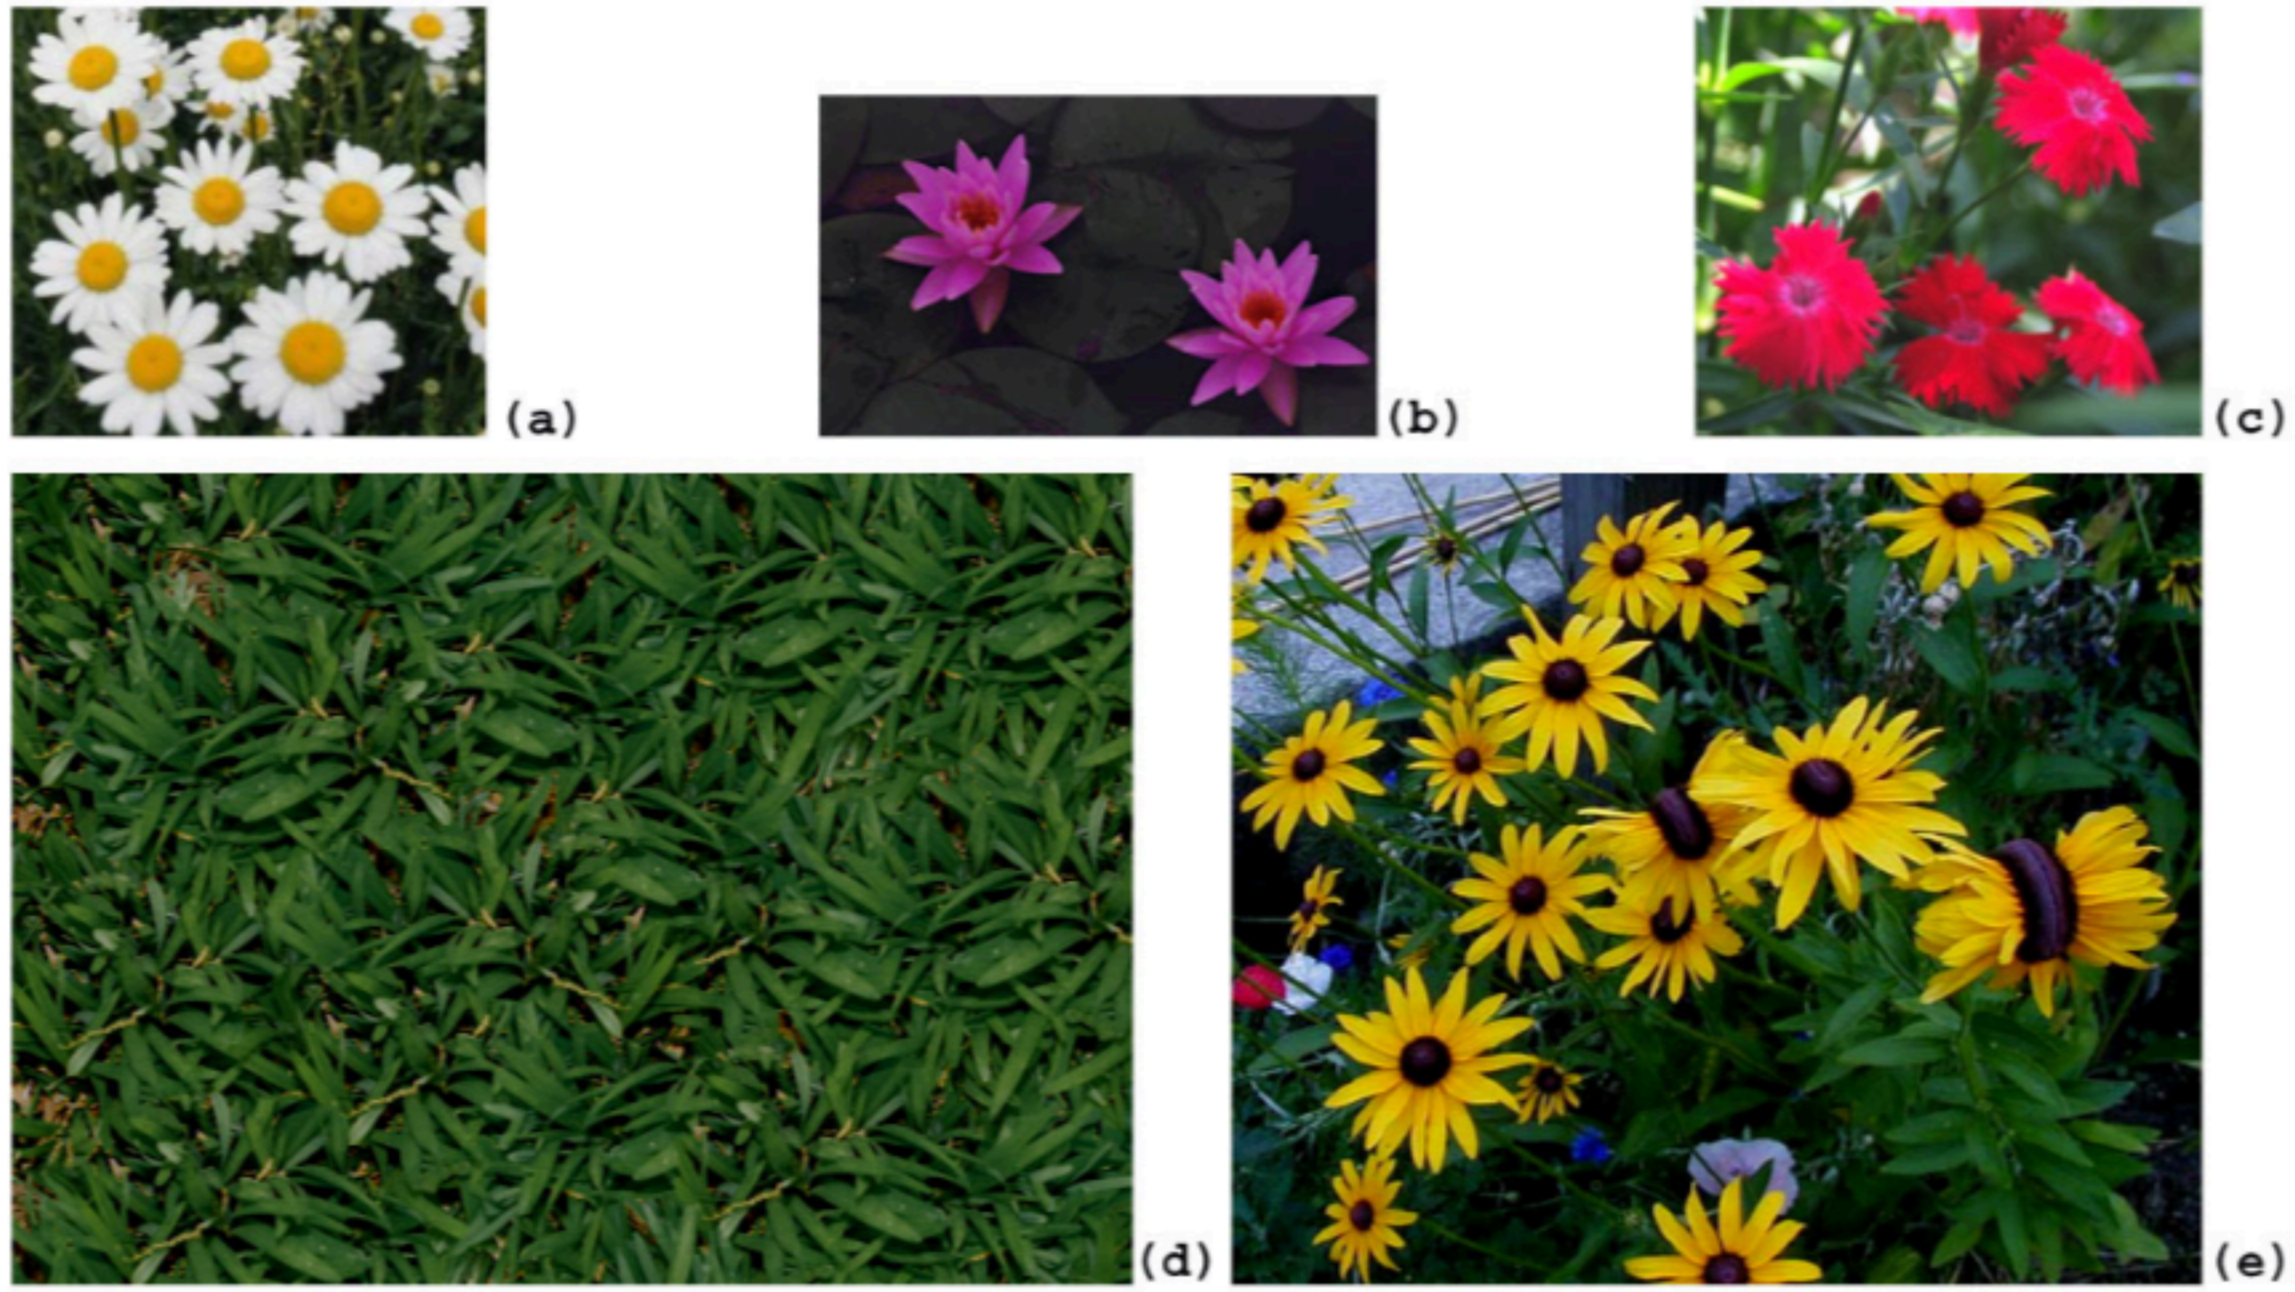
(a)
(b)
(c)
(d)
(e)


SIGGRAPH

Video synthesis

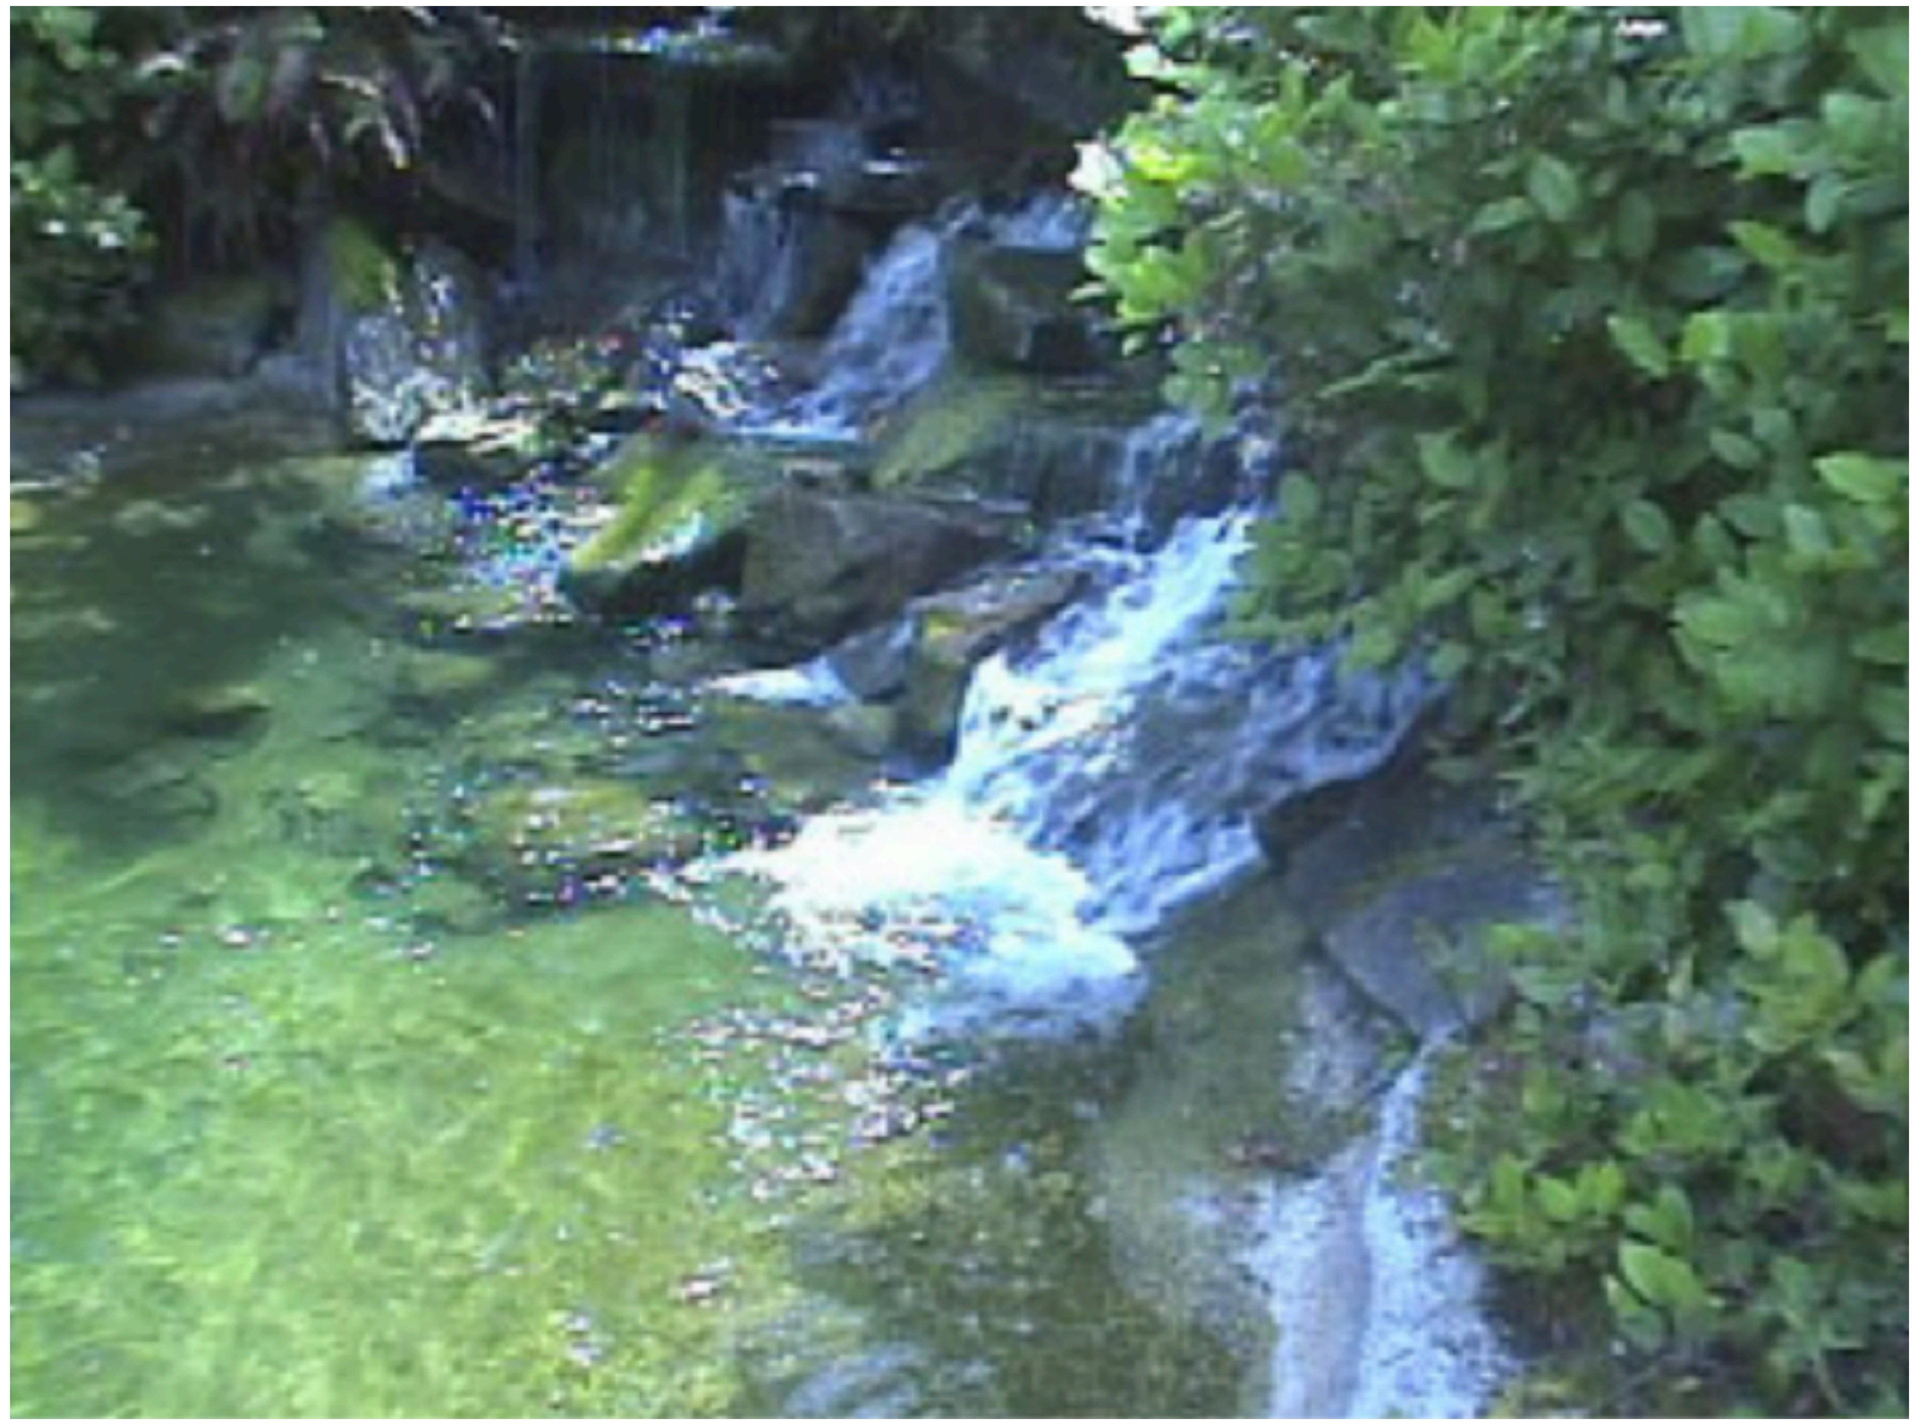

Temporally stationary

Spatio-temporally stationary

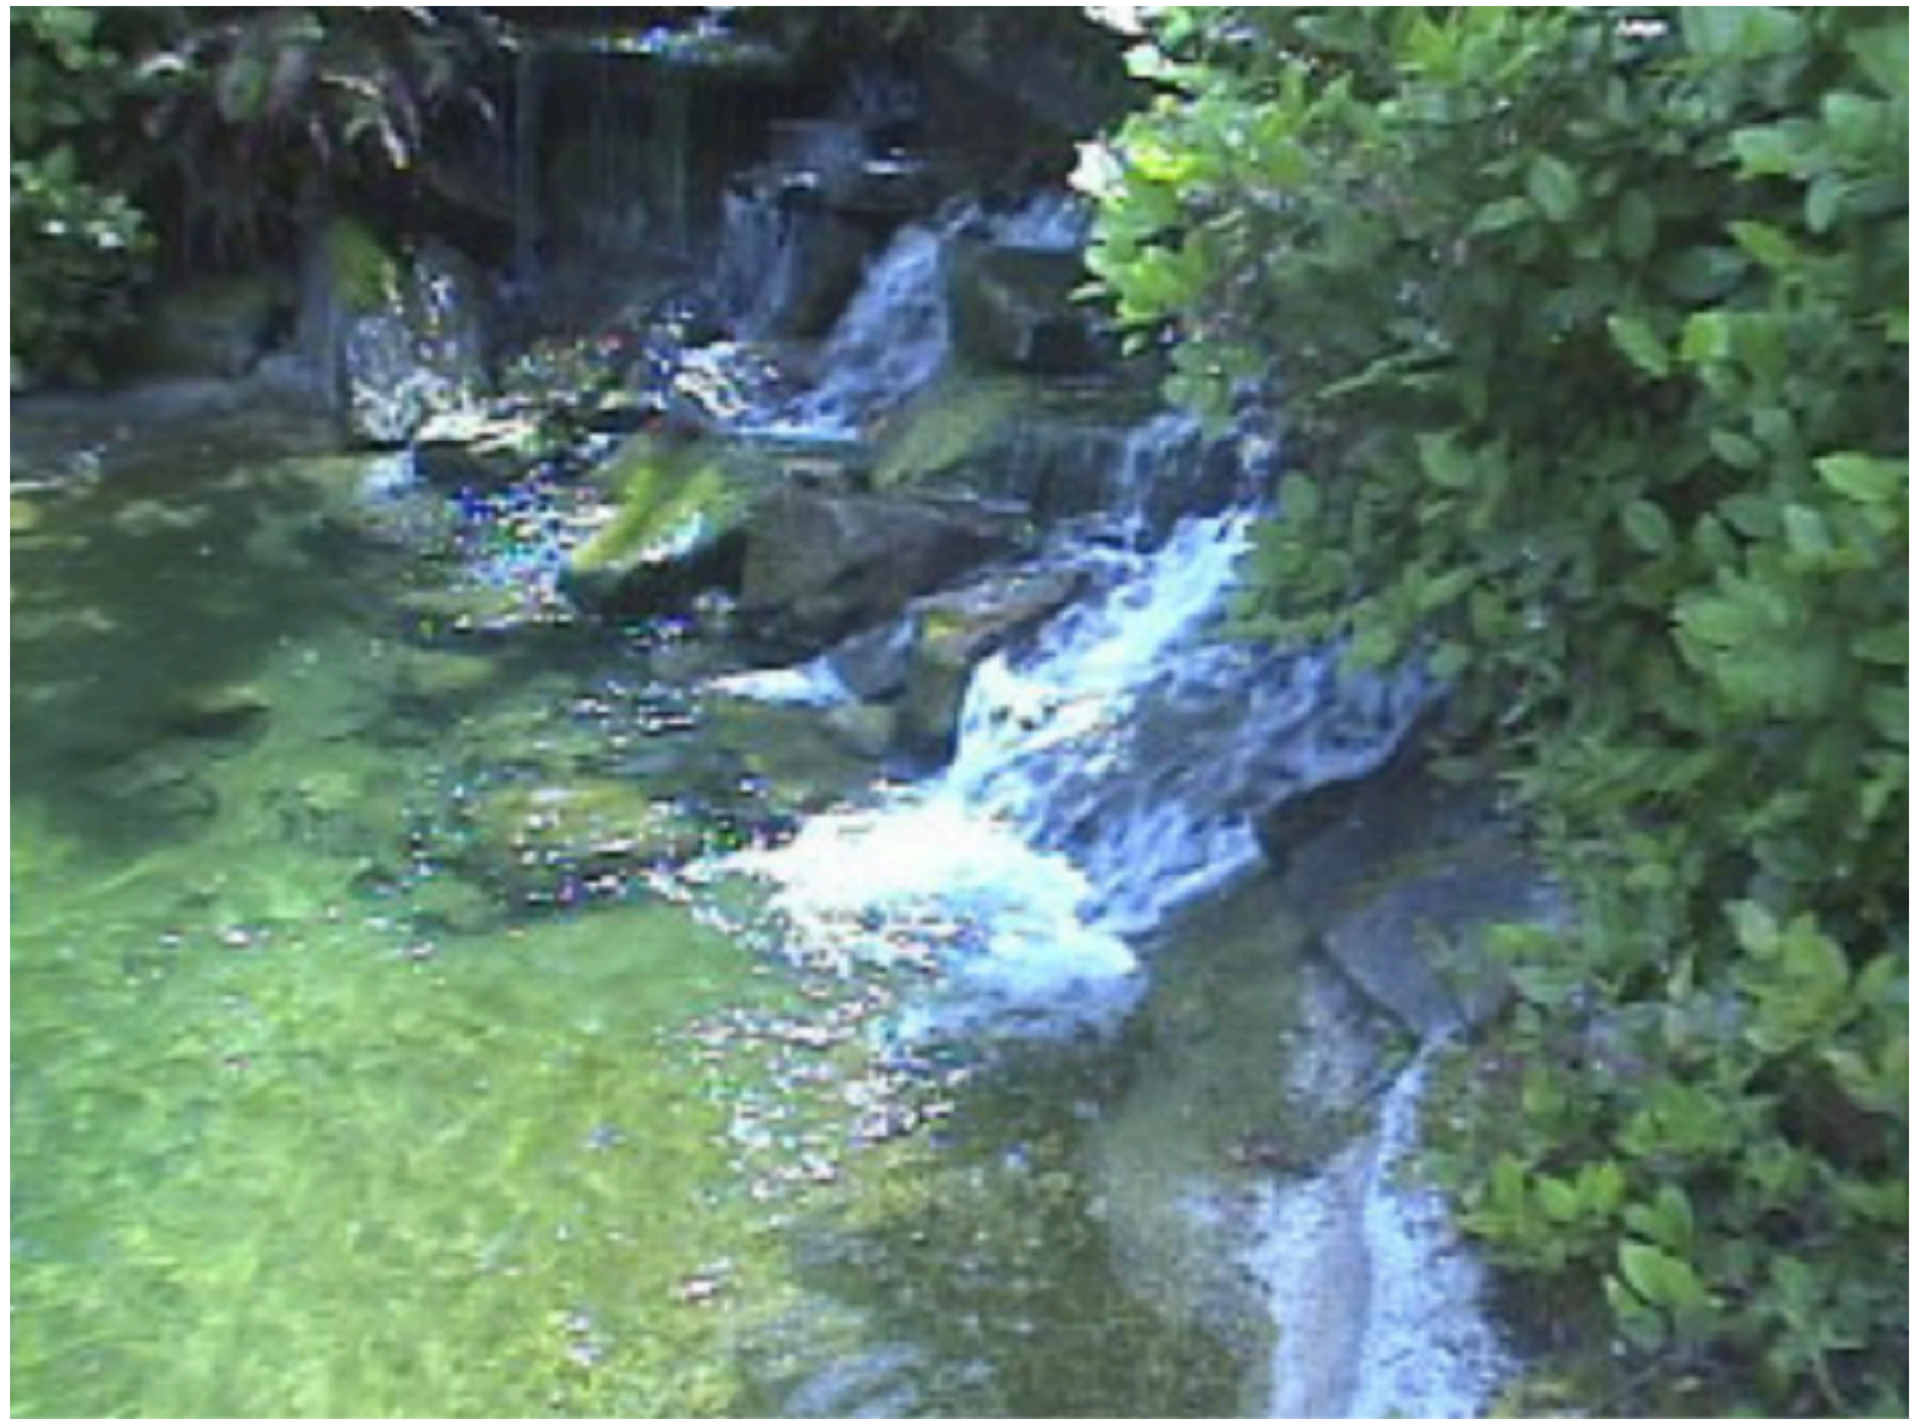

Temporally stationary

Spatio-temporally stationary

How should this affect patch search strategy?

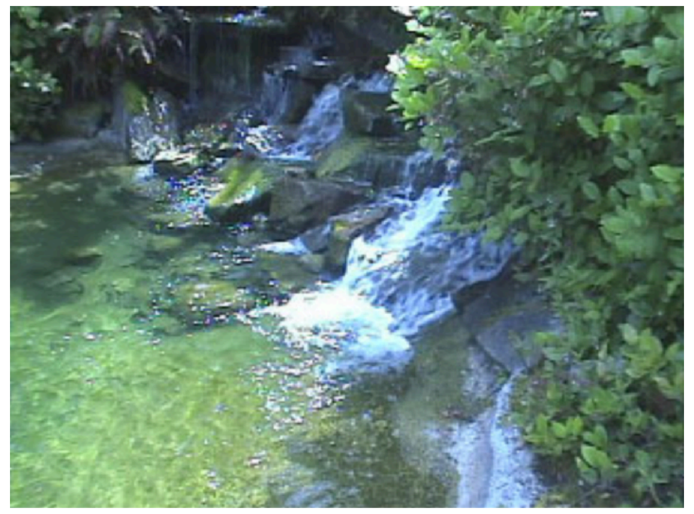

Temporally stationary

Video Textures

+

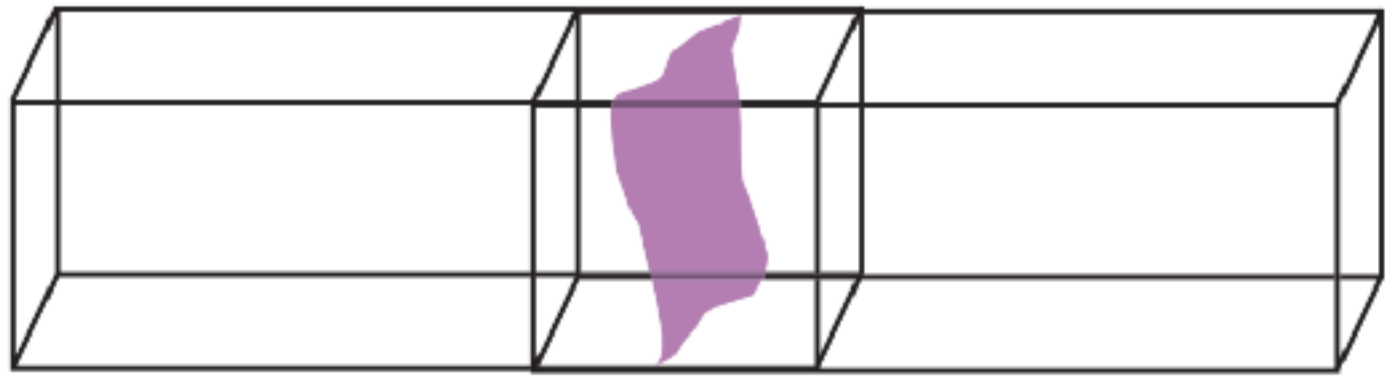

Per-pixel
transition
timing

Seam optimization

Original

Original

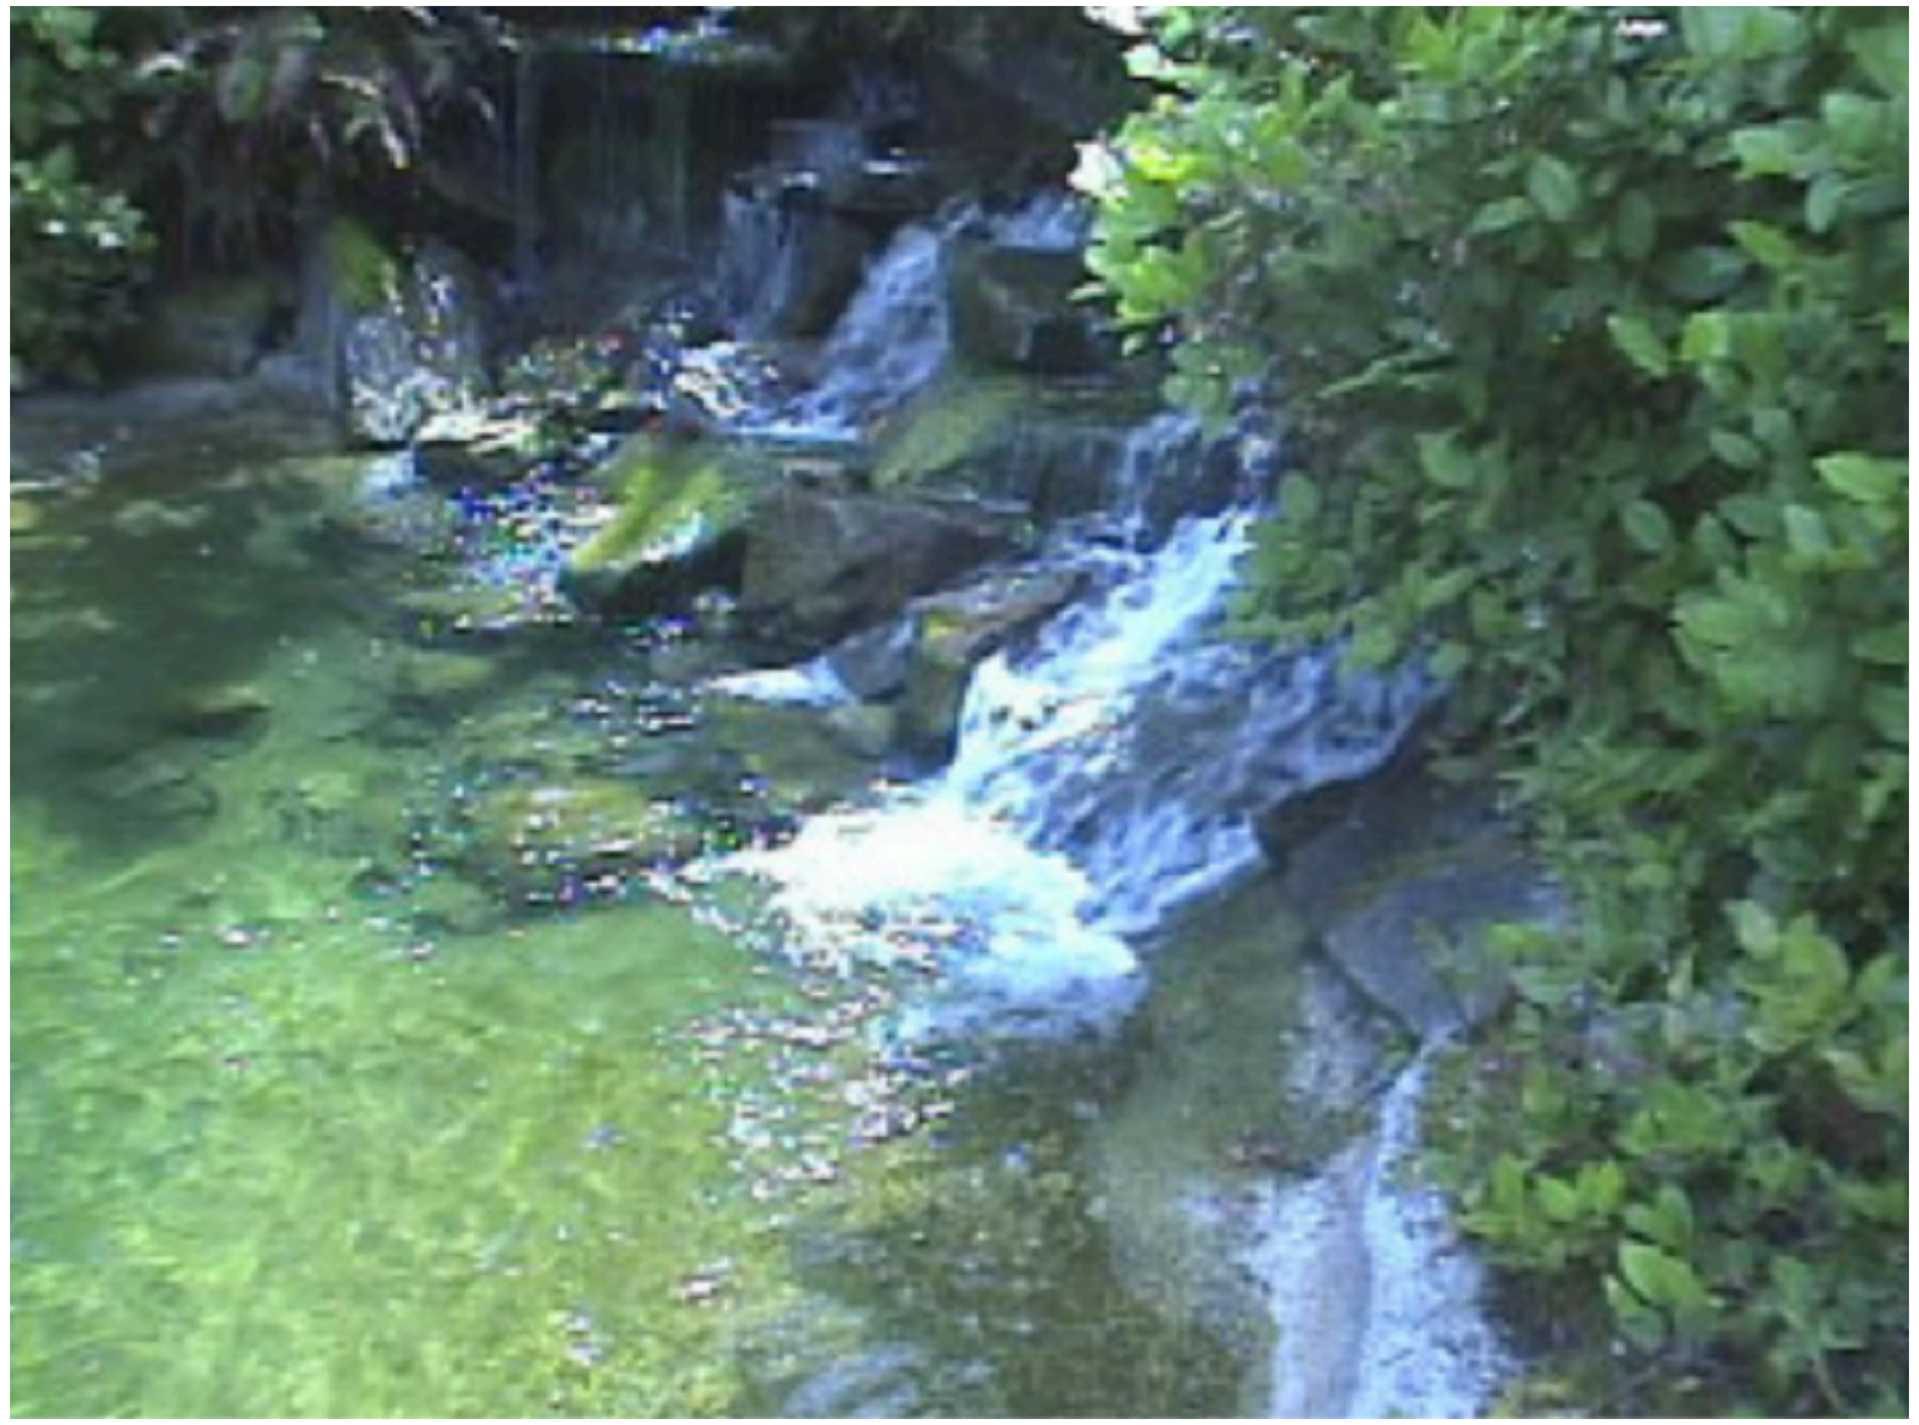

Temporally stationary

Spatio-temporally stationary

Spatio-temporally stationary

Can search patches in time and space!

Spatio-temporally stationary

Can search patches in time and space!

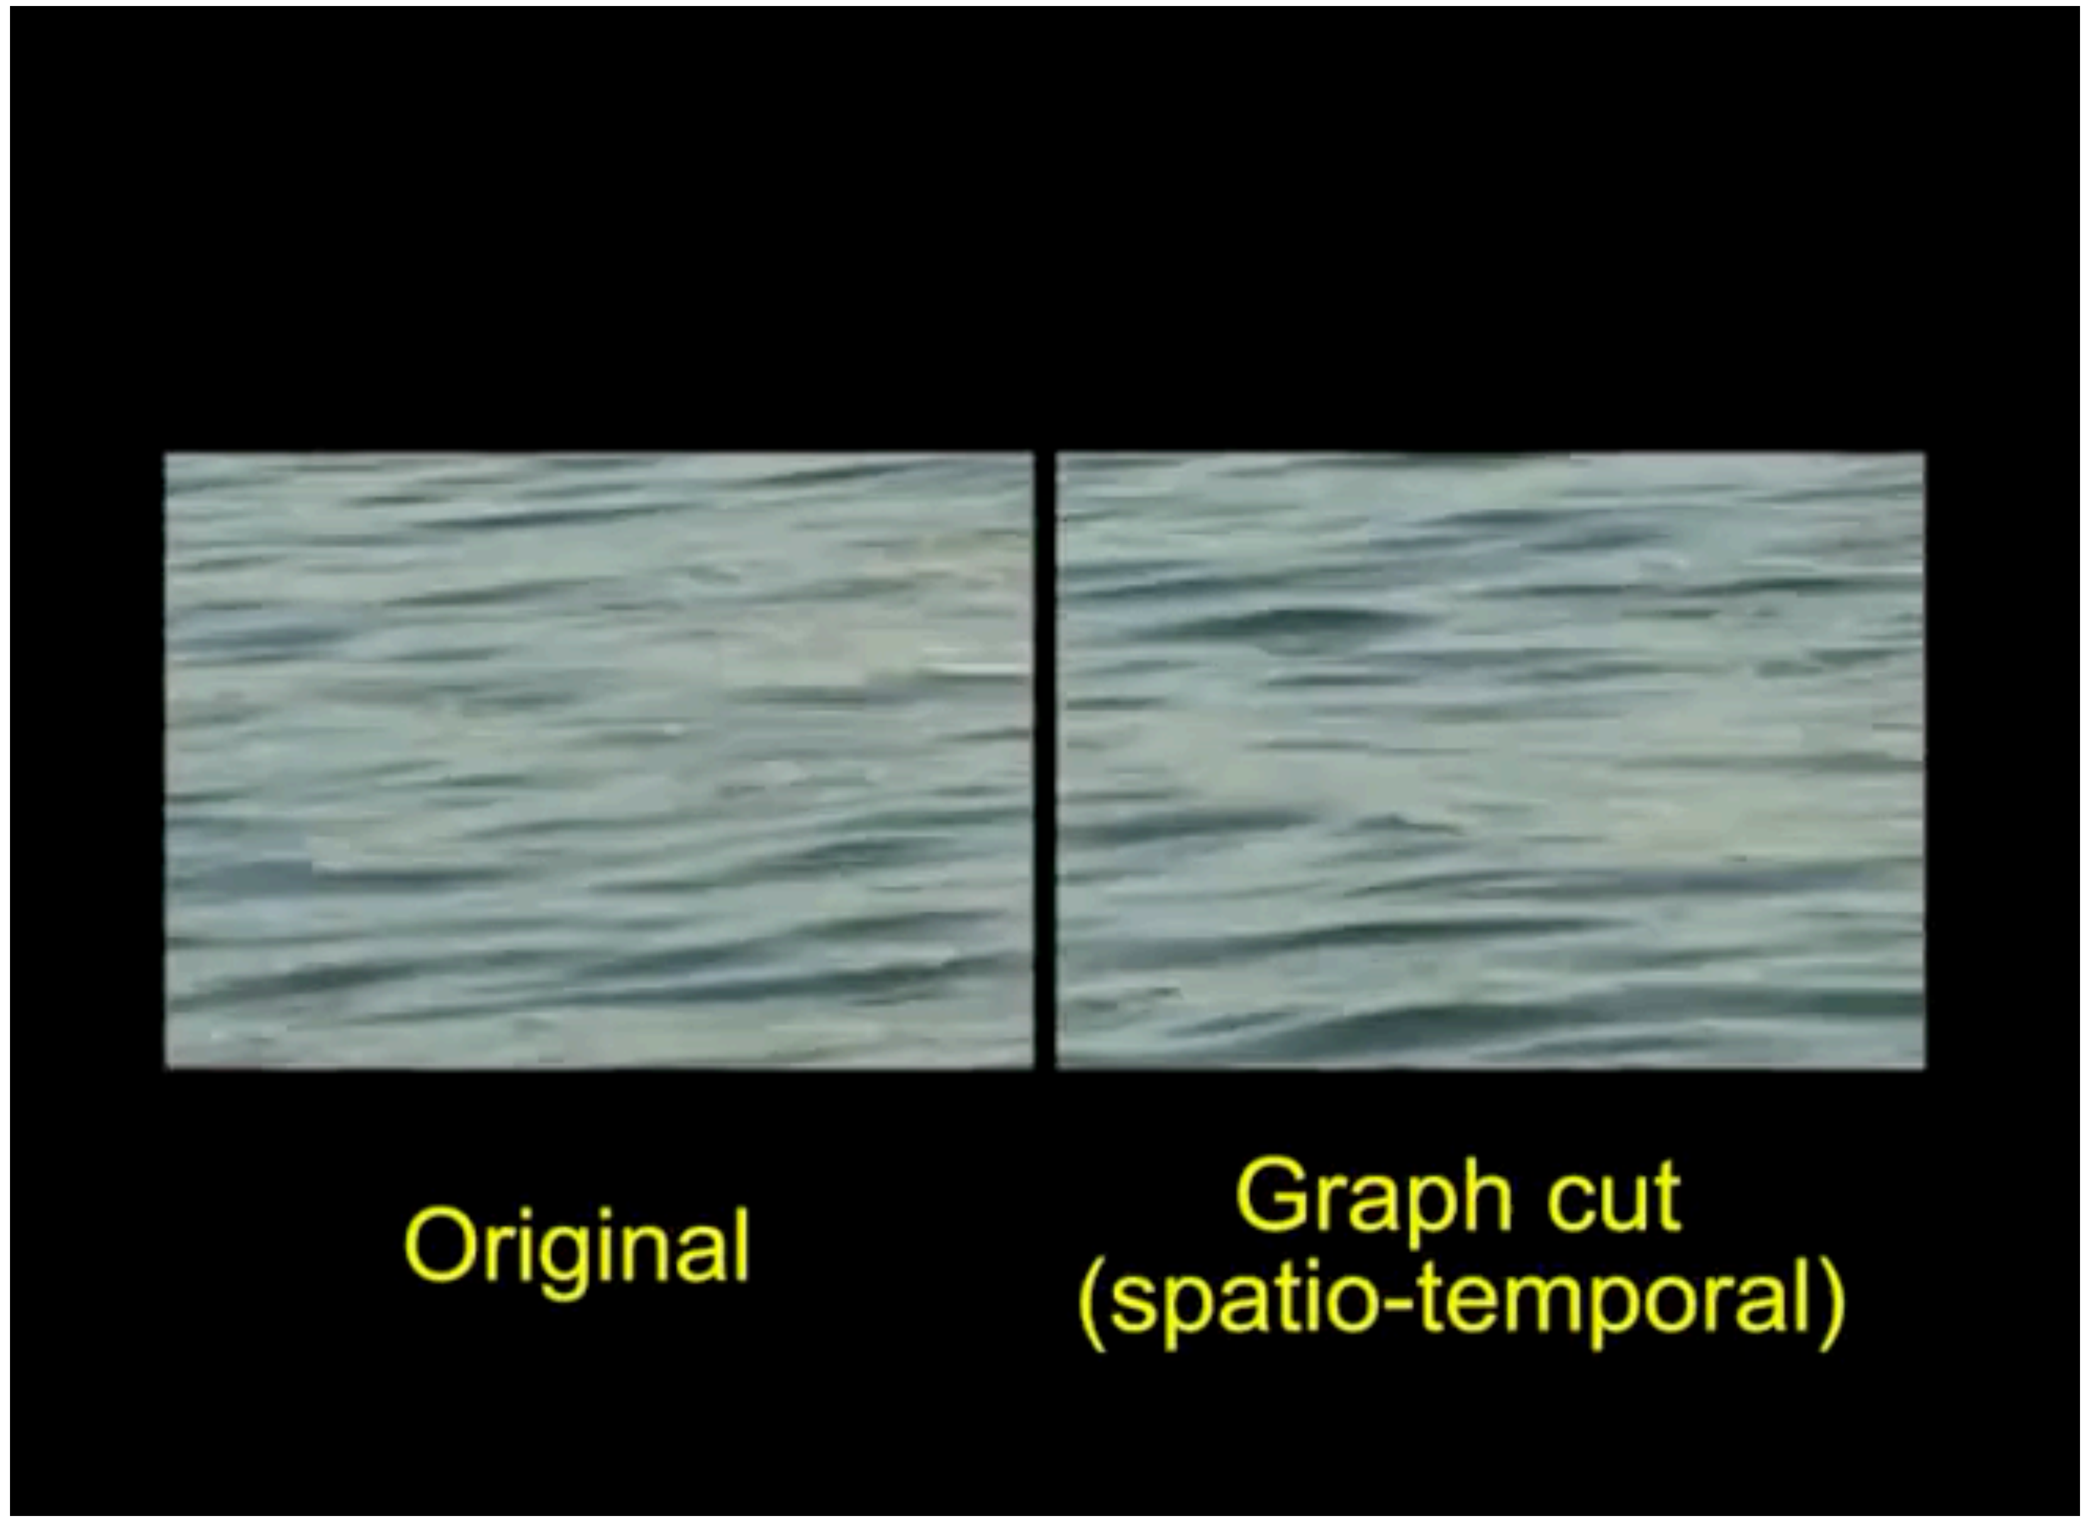

Robust results even for short sequences

Spatio-temporally stationary

Can search patches in time and space!

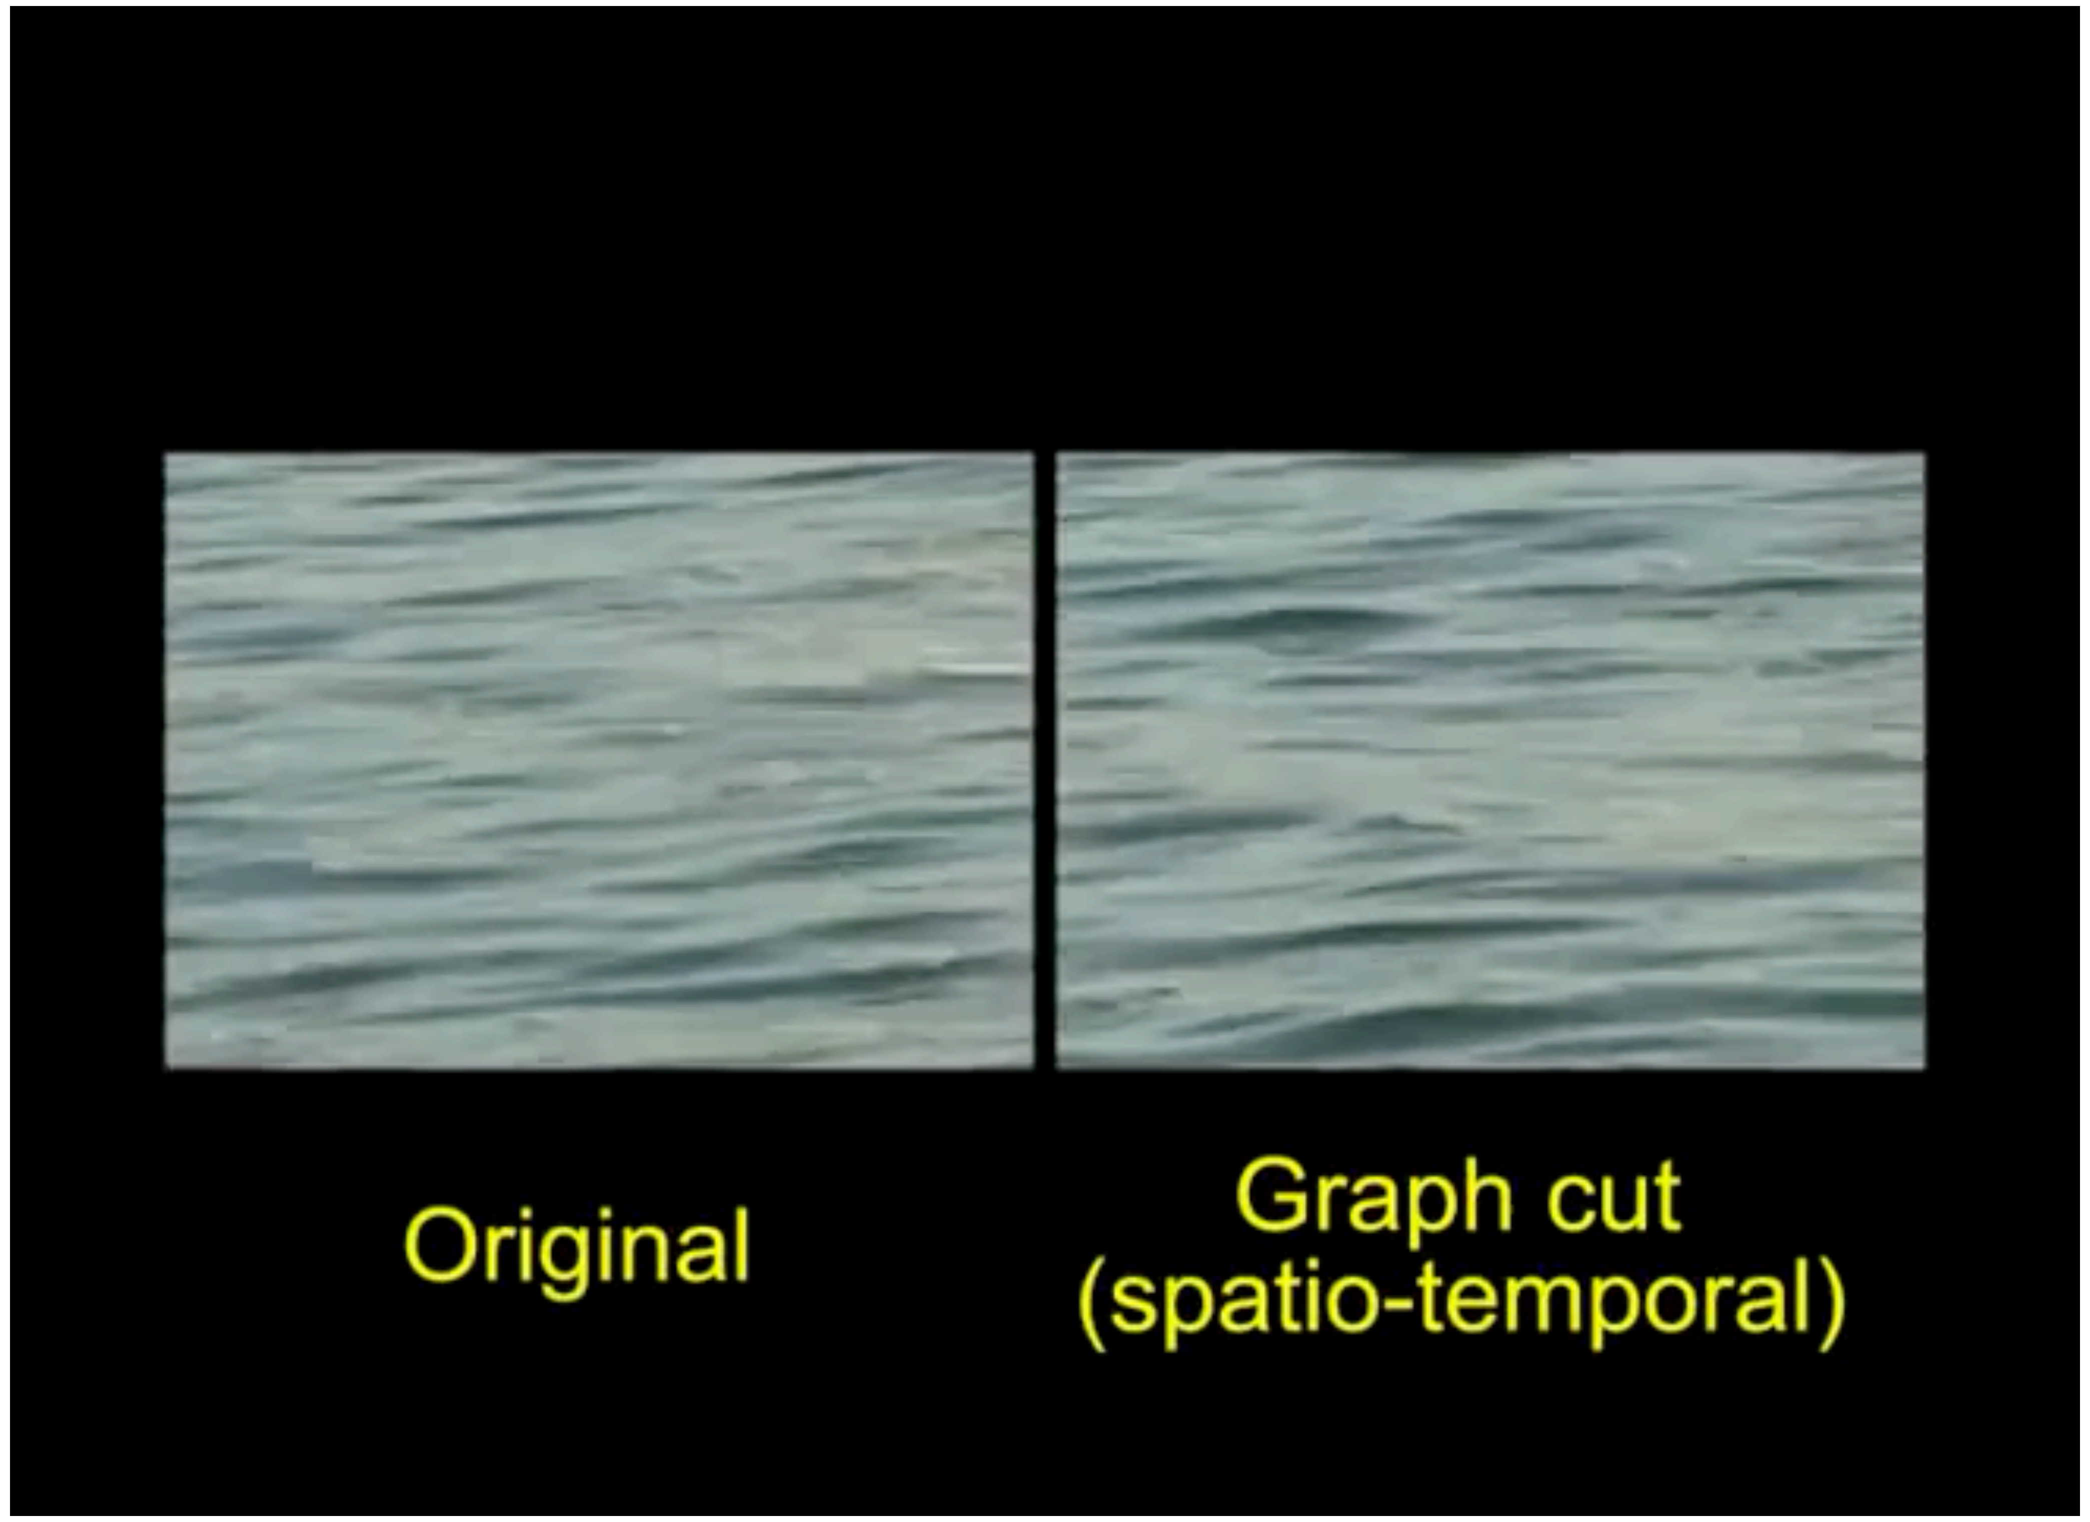

Robust results even for short sequences

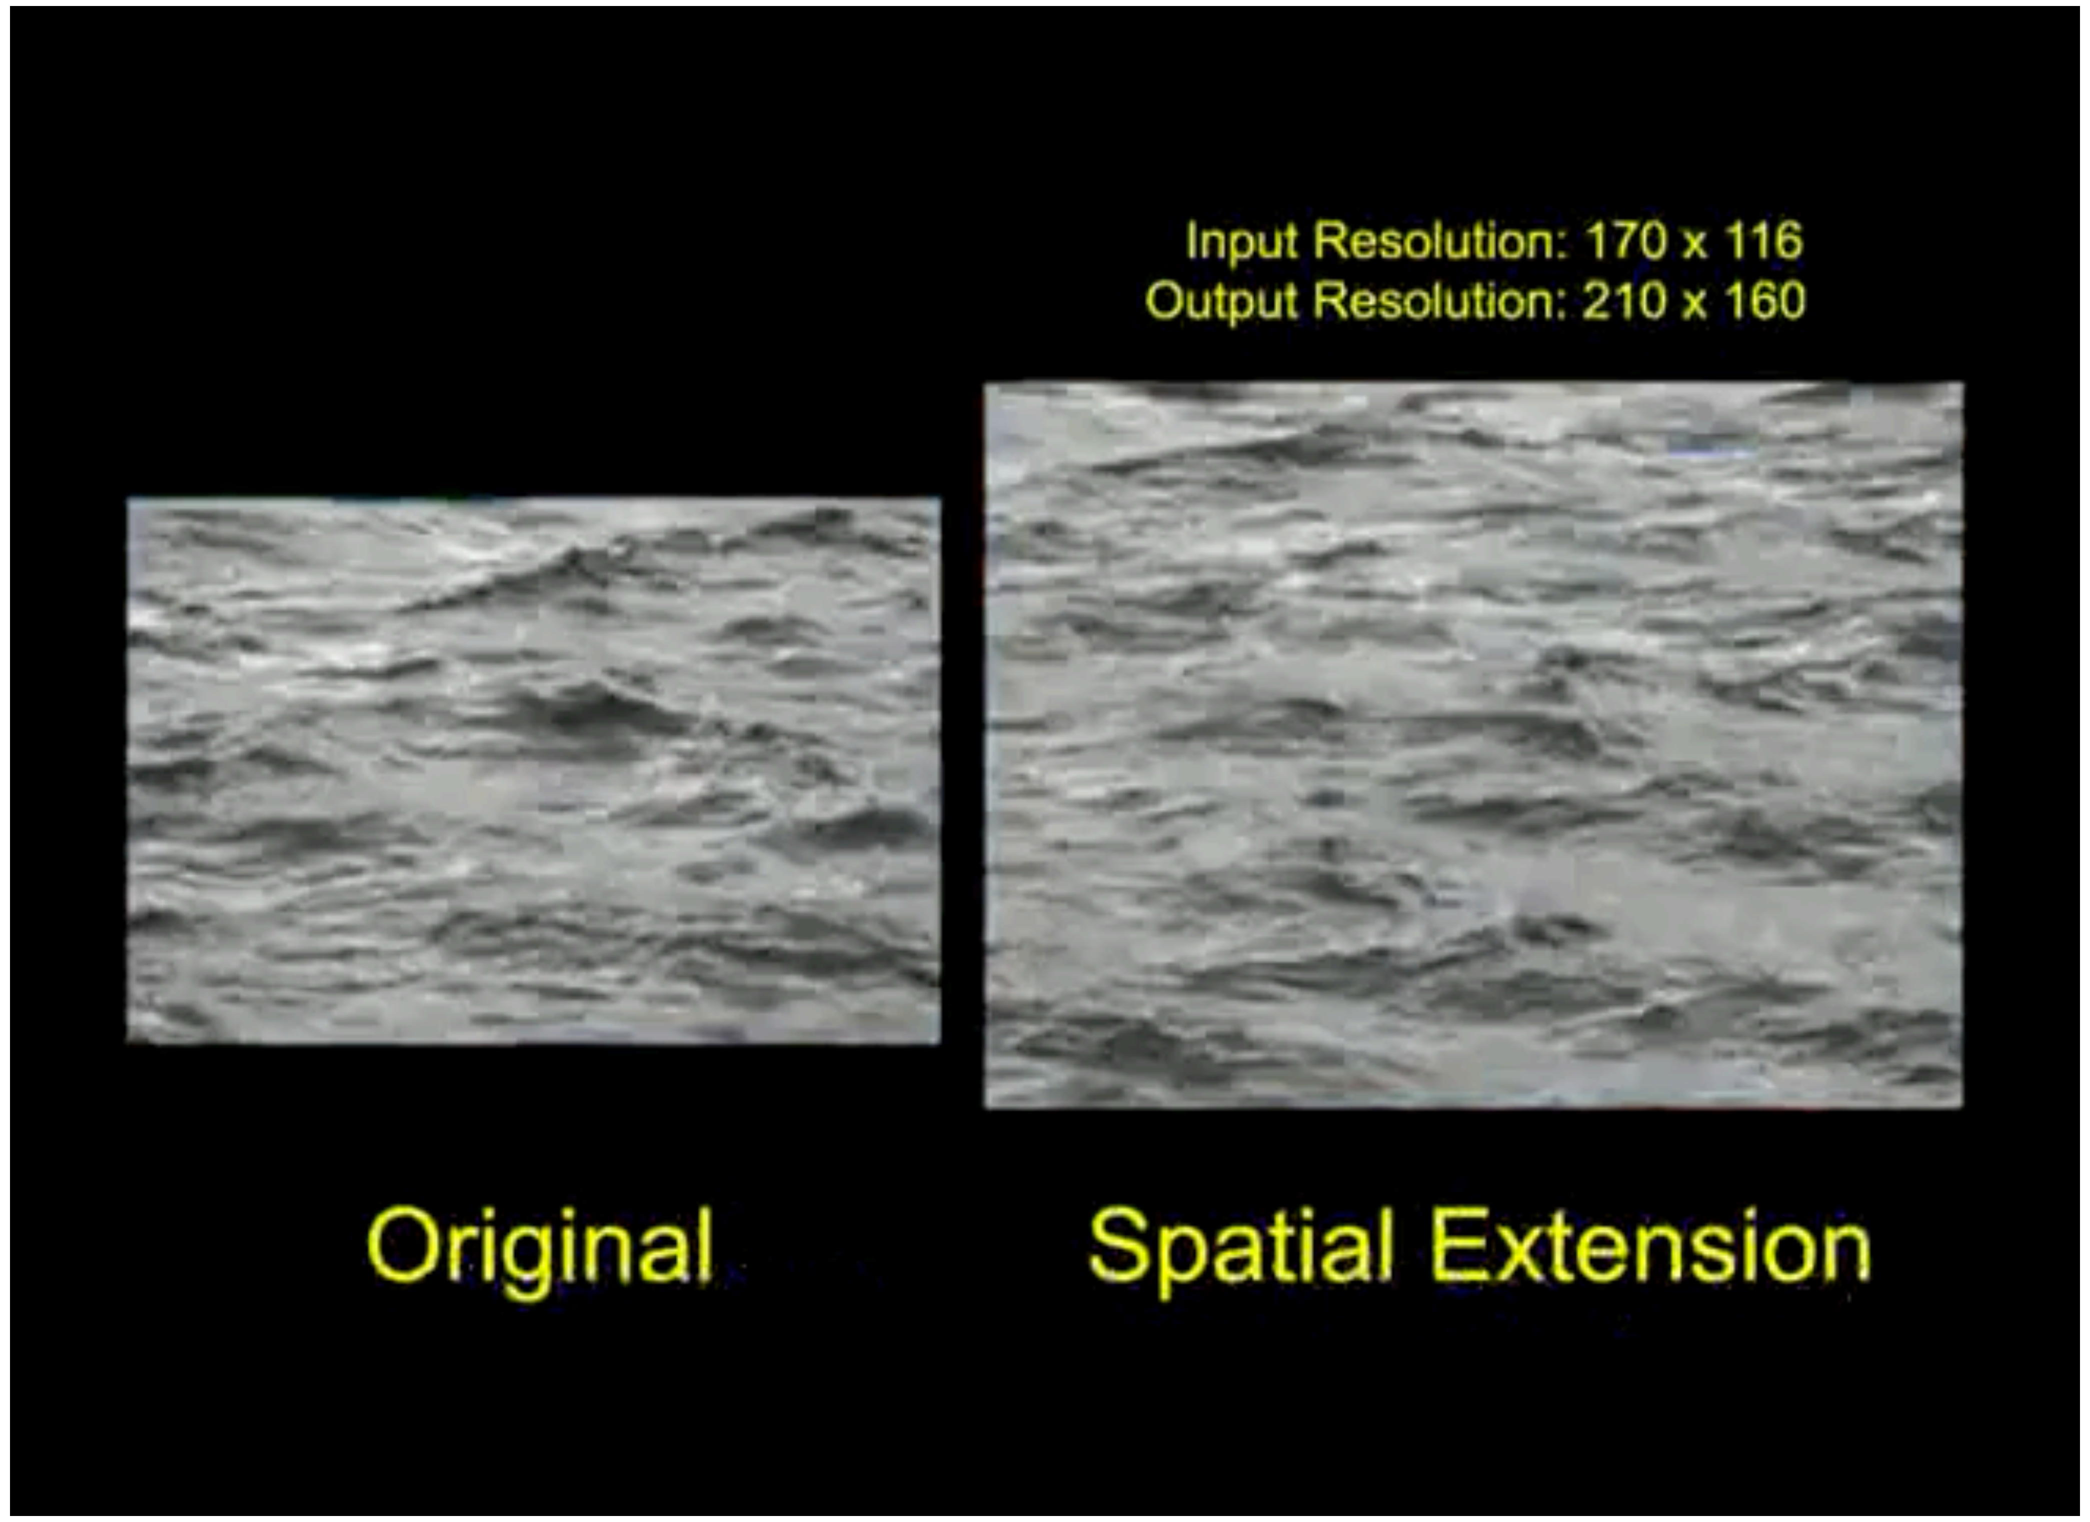

Can make videos larger

Harder to create loops. Why?

Harder to create loops. Why?

Harder to create loops. Why?

Solution: explicitly force beginning and end to match

Harder to create loops. Why?

Solution: explicitly force beginning and end to match

Recap

Recap

- Textures are everywhere!

Recap

- Textures are everywhere!

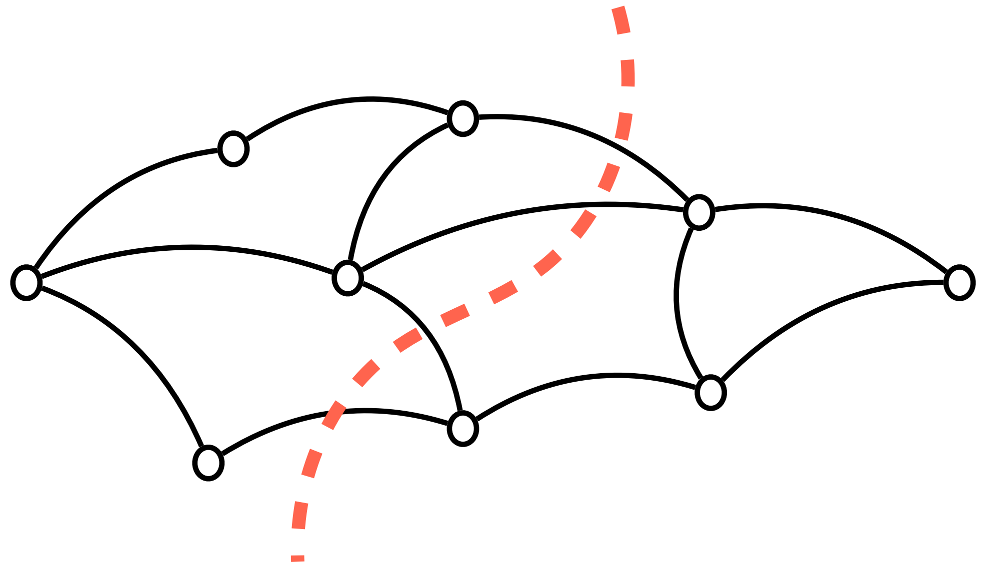

- Add to your tool belt: Graph Cuts

Recap

- Textures are everywhere!

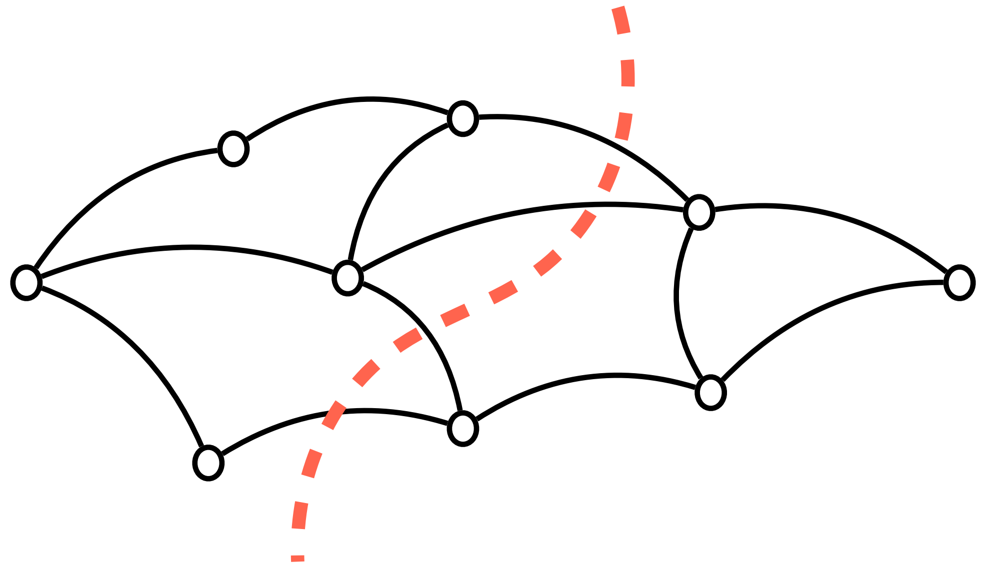

- Add to your tool belt: Graph Cuts

- Graphcut Textures

What didn’t we cover?

What didn’t we cover?

- Many things!

What didn’t we cover?

- Many things!
- E.g., image analogies

What didn’t we cover?

- Many things!
- E.g., image analogies
    - Convert between different representations of an image

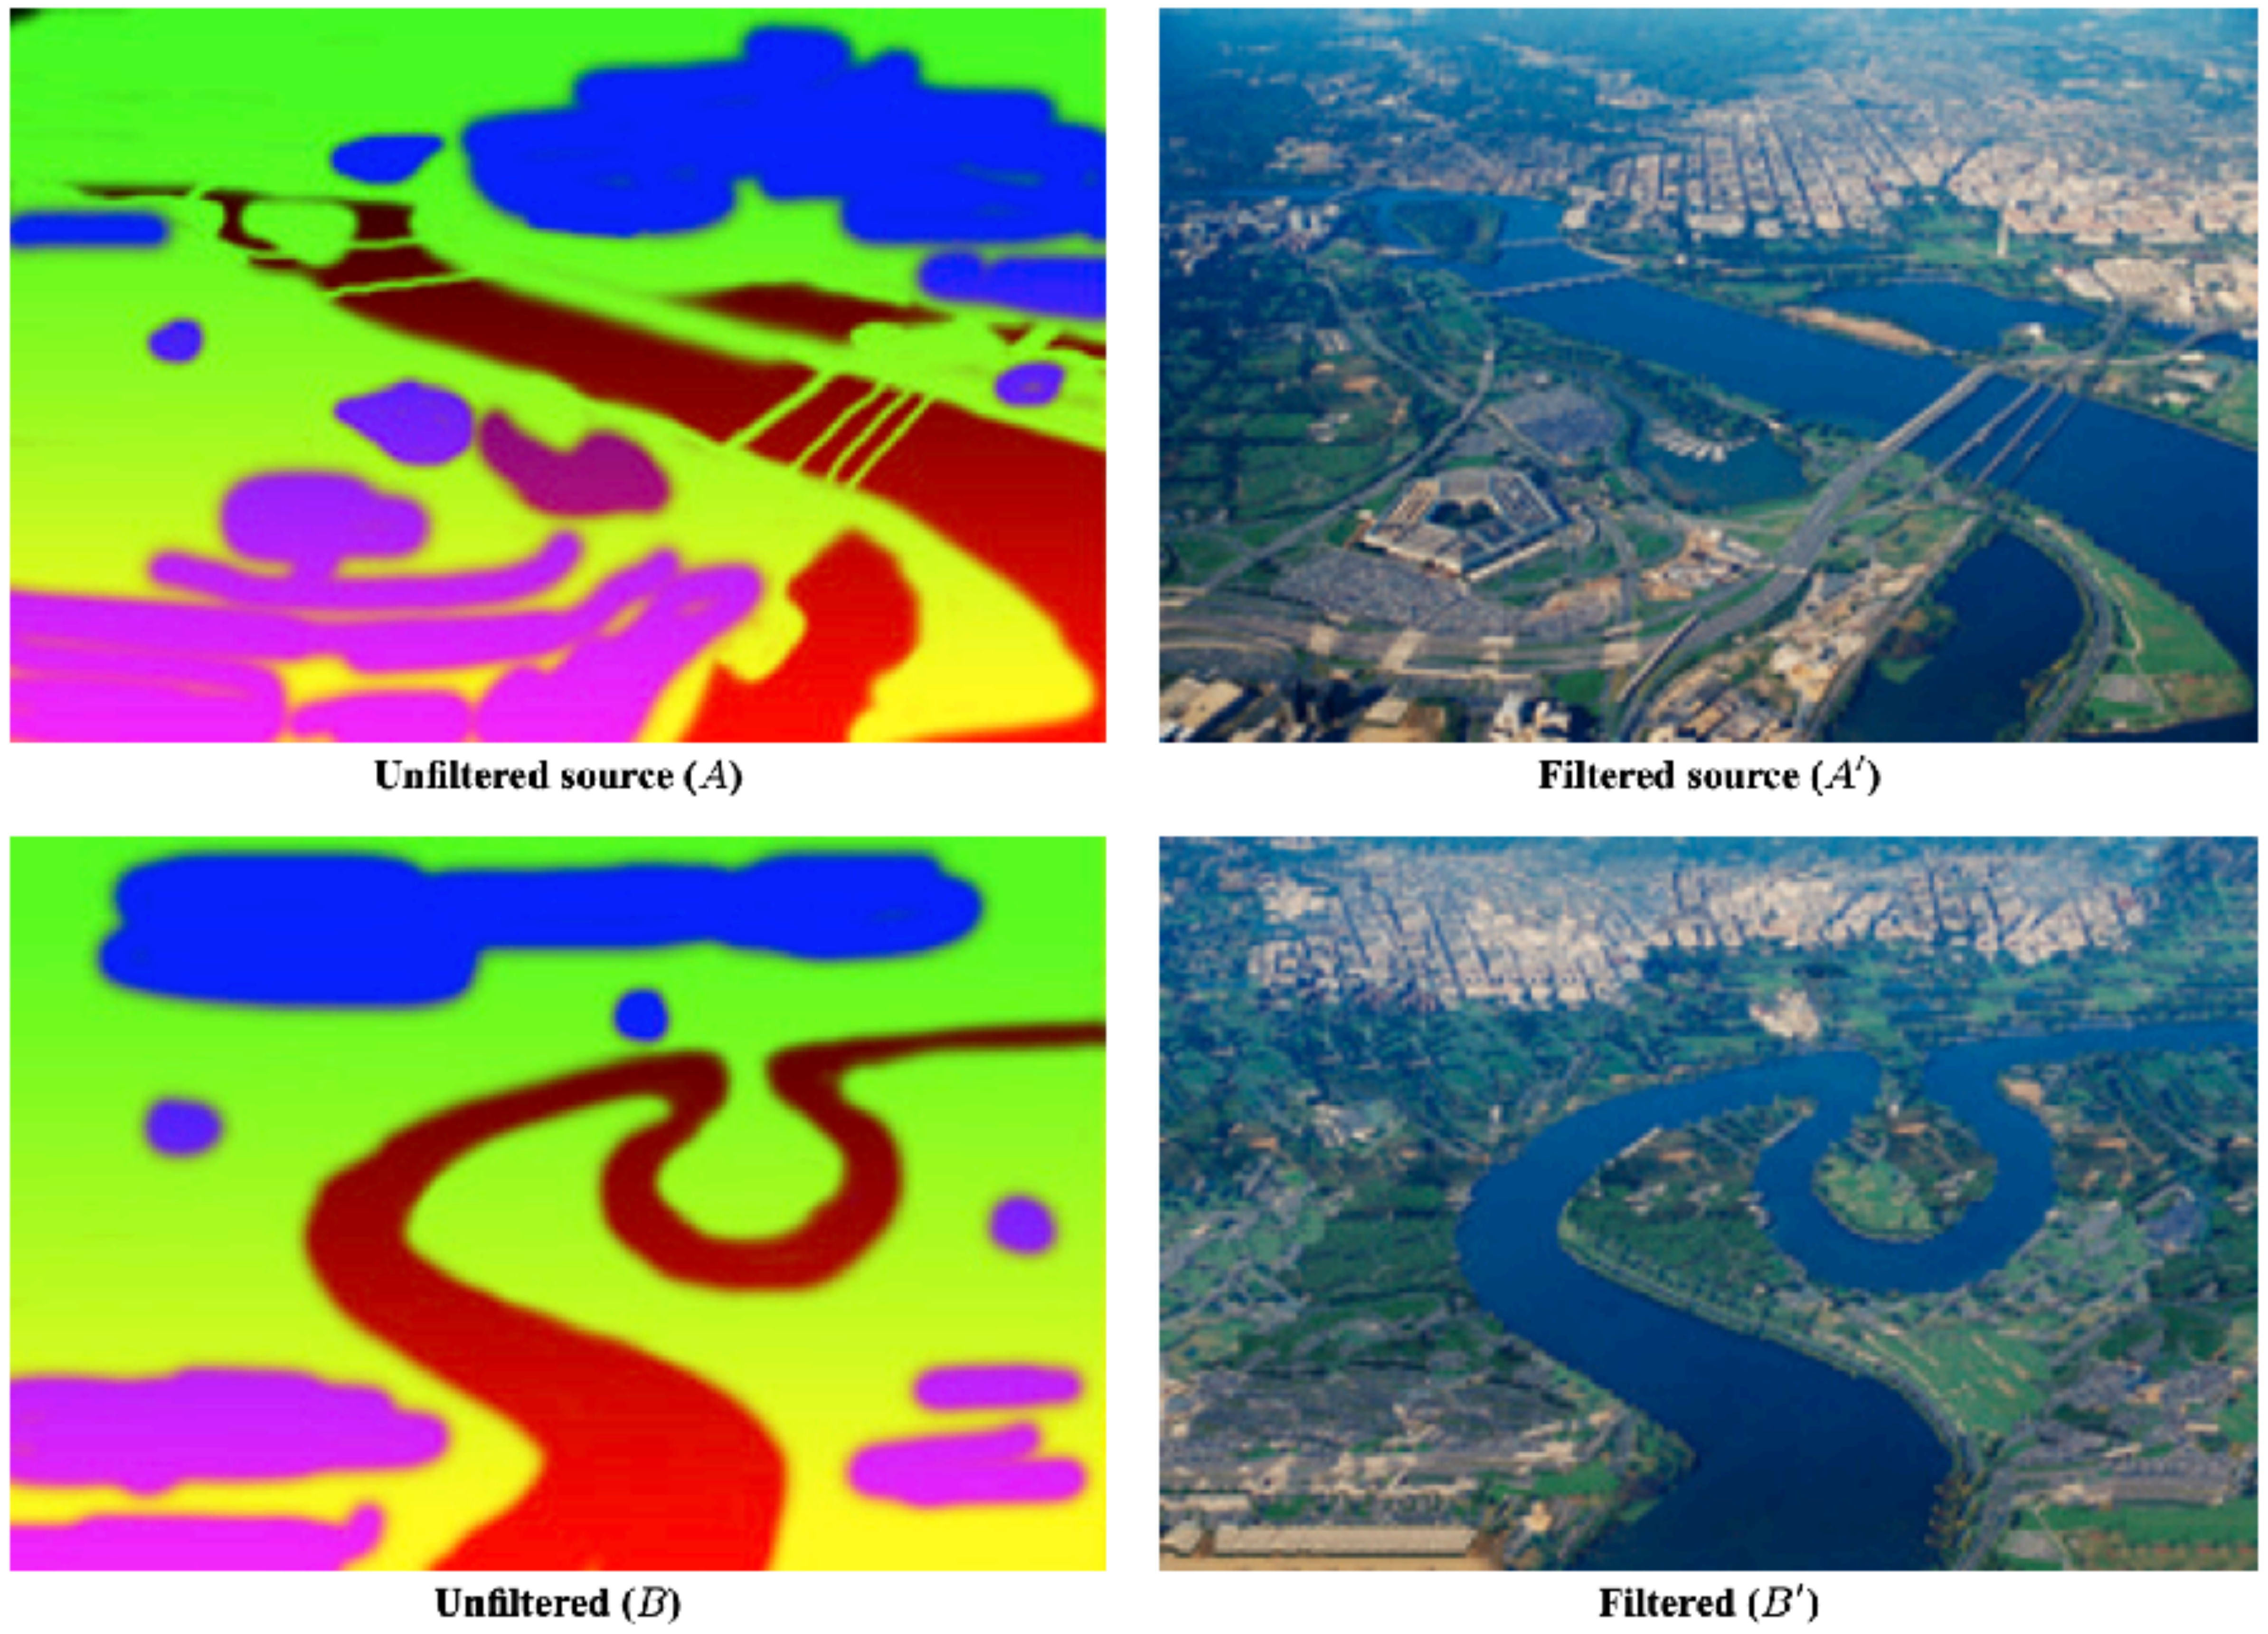

What didn’t we cover?

- Many things!
- E.g., image analogies
    - Convert between different representations of an image
    - Stylization

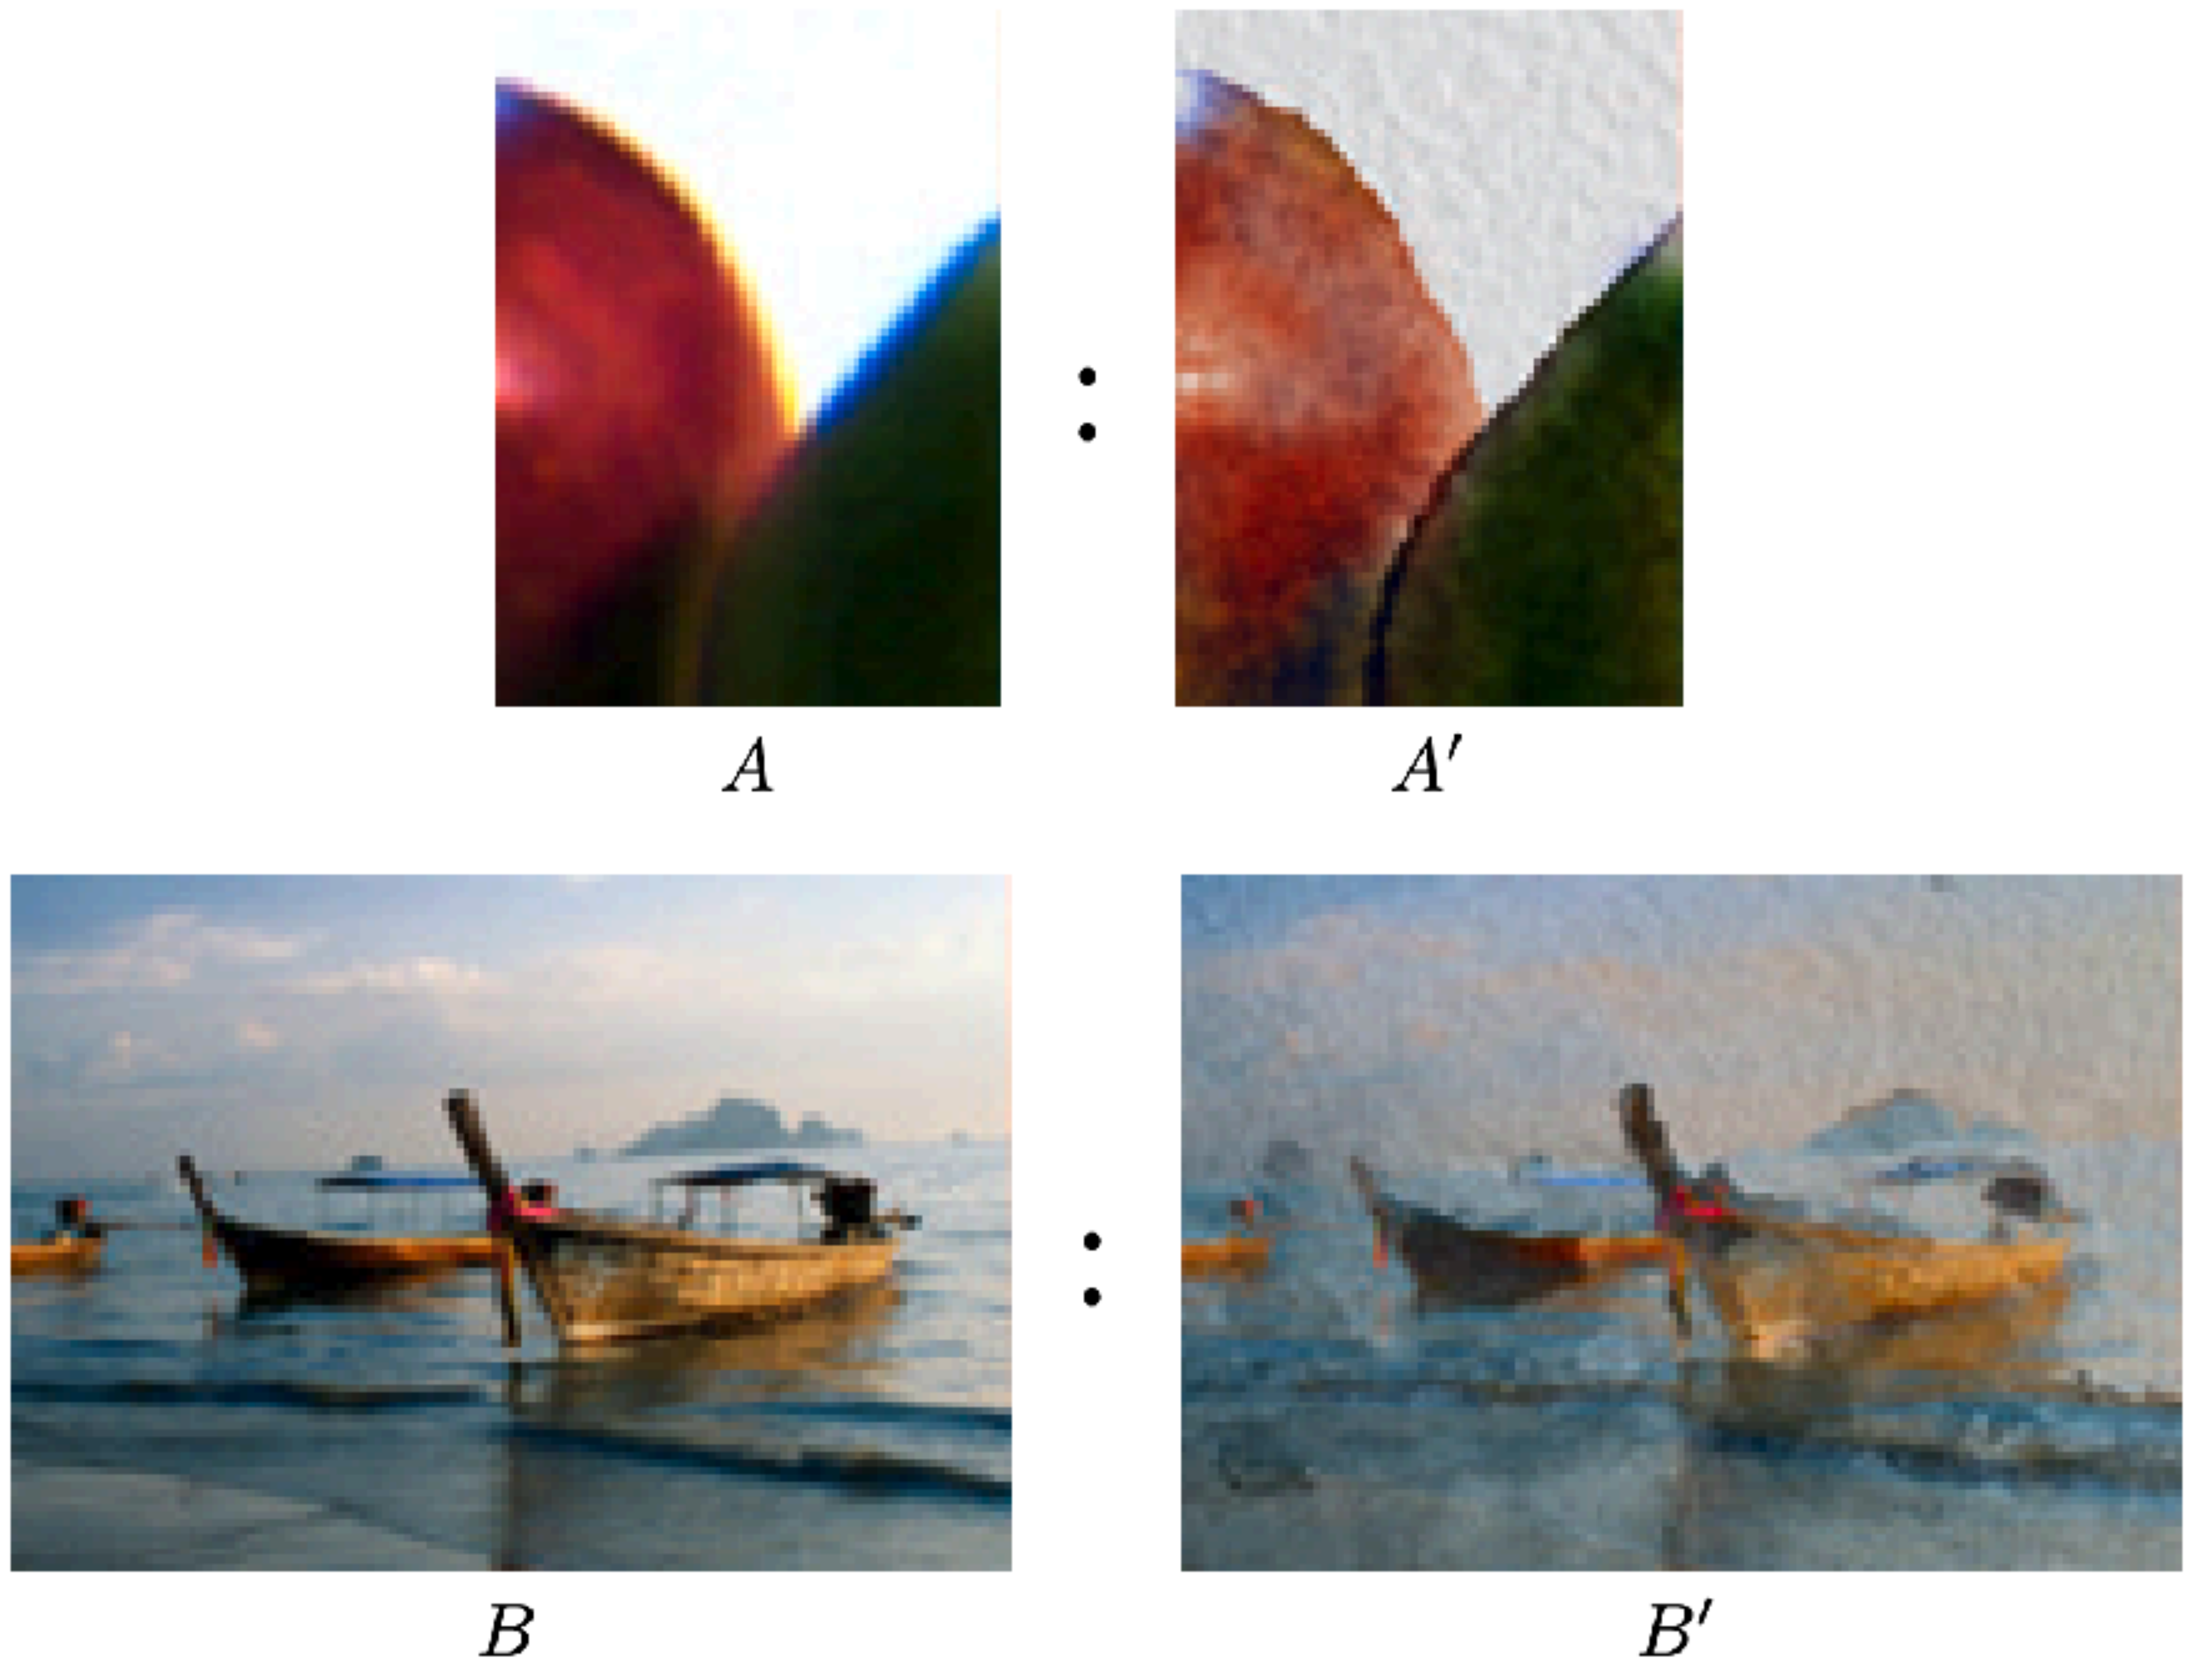

What didn't we cover?

- Many things!
- E.g., image analogies
  - Convert between different representations of an image
  - Stylization
- We will discuss these applications later in the course (using more recent methods)