


Object Detection



CSC420

David Lindell

University of Toronto

cs.toronto.edu/~lindell/teaching/420

Slide credit: Babak Taati ← Ahmed Ashraf ← Sanja Fidler


Logistics

- A4 is out. Due date is **March 28**
- Final exam April 17th WB116/119 7pm–10pm
 - multiple choice, short answer, long answer
- Office hours (TA + instructor) will continue until the exam, held over zoom
- please submit course evals!

Improving stereo matching




Block matching




Ground truth



What are some problems with the result?



How can we improve depth estimation?




How can we improve depth estimation?

Too many discontinuities.
We expect disparity values to change slowly.

Let's make an assumption:
depth should change smoothly

Stereo matching as ...

Energy Minimization



What defines a good stereo correspondence?

1. Match quality

- Want each pixel to find a good match in the other image

2. Smoothness

- If two pixels are adjacent, they should (usually) move about the same amount

energy function
(for one pixel)

$$E(d) = \underbrace{E_d(d)}_{\text{data term}} + \lambda \underbrace{E_s(d)}_{\text{smoothness term}}$$



Want each pixel to find a good match in
the other image
(block matching result)



Adjacent pixels should (usually) move
about the same amount
(smoothness function)

$$E(d) = E_d(d) + \lambda E_s(d)$$

$$E_d(d) = \sum_{(x,y) \in I} C(x, y, d(x, y))$$

data term

SSD distance between windows centered
at $I(x, y)$ and $J(x + d(x, y), y)$

$$E(d) = E_d(d) + \lambda E_s(d)$$


$$E_d(d) = \sum_{(x,y) \in I} C(x, y, d(x, y))$$

SSD distance between windows centered at $I(x, y)$ and $J(x+d(x,y), y)$


$$E_s(d) = \sum_{(p,q) \in \mathcal{E}} V(d_p, d_q)$$

smoothness term

\mathcal{E} : set of neighboring pixels



4-connected neighborhood




8-connected neighborhood

$$E_s(d) = \sum_{(p,q) \in \mathcal{E}} V(d_p, d_q)$$

smoothness term


$$V(d_p, d_q) = |d_p - d_q|$$

L₁ distance



$$V(d_p, d_q) = \begin{cases} 0 & \text{if } d_p = d_q \\ 1 & \text{if } d_p \neq d_q \end{cases}$$

"Potts model"



One possible solution...

Dynamic Programming

$$E(d) = E_d(d) + \lambda E_s(d)$$

Can minimize this independently per scanline
using dynamic programming (DP) ●.....●.....●

$D(x, y, d)$: minimum cost of solution such that $d(x, y) = d$

$$D(x, y, d) = C(x, y, d) + \min_{d'} \{D(x - 1, y, d') + \lambda |d - d'|\}$$


One possible solution...

Dynamic Programming

Left Image




Right Image



Dissimilarity Values
(1-NCC) or SSD


One possible solution...

Dynamic Programming




One possible solution...

Dynamic Programming



One possible solution...

Dynamic Programming






Three cases:

- Sequential – add cost of match (small if intensities agree)
- Occluded – add cost of no match (large cost)
- Disoccluded – add cost of no match (large cost)

One possible solution...

Dynamic Programming




Three cases:

- Matching patches. Cost = dissimilarity score
- Occluded from right. Cost is some constant value.
- Occluded from left. Cost is some constant value.

$$C(i,j) = \min([C(i-1,j-1) + \text{dissimilarity}(i,j) \\ C(i-1,j) + \text{occlusionConstant}, \\ C(i,j-1) + \text{occlusionConstant}]);$$

One possible solution...


Dynamic Programming



Dynamic programming yields the optimal path through grid. This is the best set of matches that satisfy the ordering constraint

One possible solution...


Dynamic Programming



Disparity Space Image

One possible solution...

Dynamic Programming



Disparity Space Image



Y. Boykov, O. Veksler, and R. Zabih, [Fast Approximate Energy Minimization via Graph Cuts](#), PAMI 2001

Where are we in the Vision Landscape

- Template Detection, Normalized Correlation
- Linear Filters, Convolutions, Gradients
- Edges ... Non-Max Suppression
- Interest Points – Corners – Harris Corner Detector
- SIFT – Scale Invariant Feature Transform
- Feature Descriptor around Interest Points (Remember 128D descriptor)
- Feature Matching and RANSAC, Homography
- Camera Models, Perspective Projections, Stereo
- Deep Learning – Neural Nets
- Automatic Differentiation – Training Neural Nets
- **Object Detection**

Fast retrieval

Recognizing or Retrieving Specific Objects

- Example: Visual search in feature films

Visually defined query

“Groundhog Day” [Rammis, 1993]

“Find this
clock”



“Find this
place”





Demo: <http://www.robots.ox.ac.uk/~vgg/research/vgoogle/>

[Source: J. Sivic, slide credit: R. Urtasun]

Recognizing or Retrieving Specific Objects

- Example: Search photos on the web for particular places




Find these landmarks


...in these images and 1M more

Why is it Difficult?


- Objects can have possibly large changes in scale, viewpoint, lighting and partial occlusion.




Scale



Viewpoint



Lighting



Occlusion


Why is it Difficult?

- There is tons of data.




Our Case: Matching with Local Features

- For each image in our database we extracted local descriptors (e.g., SIFT)




Our Case: Matching with Local Features

- For each image in our database we extracted local descriptors (e.g., SIFT)



Our Case: Matching with Local Features

- Let's focus on descriptors only (vectors of e.g. 128 dim for SIFT)



Our Case: Matching with Local Features

Database of images




image 1

$$f_1^1 = [0.1, 0.2, \dots, 0.15]^T$$

$$f_2^1 = [0.23, 0.12, \dots, 0.1]^T$$

$$f_3^1 = [0.12, 0.15, \dots, 0.05]^T$$

⋮

$$f_n^1 = [0.05, 0.18, \dots, 0.09]^T$$




image 2

$$f_1^2 = [0.05, 0.11, \dots, 0.2]^T$$

$$f_2^2 = [0.09, 0.01, \dots, 0.18]^T$$

$$f_3^2 = [0.0, 0.08, \dots, 0.1]^T$$

⋮

$$f_m^2 = [0.1, 0.15, \dots, 0.14]^T$$




image 3

...




image hugeN

$$f_1^{hugeN} = [0.12, 0.15, \dots, 0.19]^T$$

$$f_2^{hugeN} = [0.1, 0.2, \dots, 0.2]^T$$

$$f_3^{hugeN} = [0.12, 0.22, \dots, 0.18]^T$$

⋮

$$f_k^{hugeN} = [0.15, 0.02, \dots, 0.08]^T$$

descriptors (vectors)

Now I get a reference (query) image of an object. I want to retrieve all images from the database that contain the object. **How?**


$$f_1^{ref} = [0.1, 0.2, \dots, 0.16]^T$$

$$f_2^{ref} = [0.15, 0.02, \dots, 0.06]^T$$

$$f_3^{ref} = [0.14, 0.22, \dots, 0.09]^T$$

⋮

$$f_p^{ref} = [0.17, 0.18, \dots, 0.2]^T$$



reference (query) image

Our Case: Matching with Local Features

Database of images




image 1




image 2




image 3

...




image hugeN

$$f_1^1 = [0.1, 0.2, \dots, 0.15]^T$$

$$f_2^1 = [0.23, 0.12, \dots, 0.1]^T$$

$$f_3^1 = [0.12, 0.15, \dots, 0.05]^T$$

⋮

$$f_n^1 = [0.05, 0.18, \dots, 0.09]^T$$

$$f_1^2 = [0.05, 0.11, \dots, 0.2]^T$$

$$f_2^2 = [0.09, 0.01, \dots, 0.18]^T$$

$$f_3^2 = [0.0, 0.08, \dots, 0.1]^T$$

⋮

$$f_m^2 = [0.1, 0.15, \dots, 0.14]^T$$

descriptors (vectors)

$$f_1^{hugeN} = [0.12, 0.15, \dots, 0.19]^T$$

$$f_2^{hugeN} = [0.1, 0.2, \dots, 0.2]^T$$

$$f_3^{hugeN} = [0.12, 0.22, \dots, 0.18]^T$$

⋮

$$f_k^{hugeN} = [0.15, 0.02, \dots, 0.08]^T$$

SLOW

Before (Assignment 3) we were matching **all** reference descriptors to **all** descriptors in **each** database image. Not very efficient.


$$f_1^{ref} = [0.1, 0.2, \dots, 0.16]^T$$

$$f_2^{ref} = [0.15, 0.02, \dots, 0.06]^T$$

$$f_3^{ref} = [0.14, 0.22, \dots, 0.09]^T$$

⋮

$$f_p^{ref} = [0.17, 0.18, \dots, 0.2]^T$$



reference (query) image

Our Case: Matching with Local Features

Database of images




image 1




image 2




image 3

...




image hugeN

$$f_1^1 = [0.1, 0.2, \dots, 0.15]^T$$

$$f_2^1 = [0.23, 0.12, \dots, 0.1]^T$$

$$f_3^1 = [0.12, 0.15, \dots, 0.05]^T$$

⋮

$$f_n^1 = [0.05, 0.18, \dots, 0.09]^T$$

$$f_1^2 = [0.05, 0.11, \dots, 0.2]^T$$

$$f_2^2 = [0.09, 0.01, \dots, 0.18]^T$$

$$f_3^2 = [0.0, 0.08, \dots, 0.1]^T$$

⋮

$$f_m^2 = [0.1, 0.15, \dots, 0.14]^T$$

descriptors (vectors)

$$f_1^{hugeN} = [0.12, 0.15, \dots, 0.19]^T$$

$$f_2^{hugeN} = [0.1, 0.2, \dots, 0.2]^T$$

$$f_3^{hugeN} = [0.12, 0.22, \dots, 0.18]^T$$

⋮

$$f_k^{hugeN} = [0.15, 0.02, \dots, 0.08]^T$$

What can we do to speed-up?

Before (Assignment 3) we were matching **all** reference descriptors to **all** descriptors in **each** database image. Not very efficient.


$$f_1^{ref} = [0.1, 0.2, \dots, 0.16]^T$$

$$f_2^{ref} = [0.15, 0.02, \dots, 0.06]^T$$

$$f_3^{ref} = [0.14, 0.22, \dots, 0.09]^T$$


⋮

$$f_p^{ref} = [0.17, 0.18, \dots, 0.2]^T$$



reference (query) image

Indexing!



Indexing Local Features: Inverted File Index

- For text documents, an efficient way to find all pages on which a word occurs is to use an index.

Index		
Along I-75, From Detroit to Florida; <i>inside back cover</i>	Butterfly Center, McGuire; 134	Driving Lanes; 85
Drive I-95, From Boston to Florida; <i>inside back cover</i>	CAA (see AAA)	Duval County; 163
1929 Spanish Trail Roadway; 101-102,104	CCC, The; 111,113,115,135,142	Eau Gallie; 175
511 Traffic Information; 83	Ca d'Zan; 147	Edison, Thomas; 152
A1A (Barrier Is) - I-95 Access; 86	Caloosahatchee River; 152	Eglin AFB; 116-118
AAA (and CAA); 83	Name; 150	Eight Reals; 176
AAA National Office; 88	Canaveral Natnl Seashore; 173	Ellenton; 144-145
Abbreviations,	Cannon Creek Airpark; 130	Emanuel Point Wreck; 120
Colored 25 mile Maps; cover	Canopy Road; 106,169	Emergency Callboxes; 83
Exit Services; 196	Cape Canaveral; 174	Epiphytes; 142,148,157,159
Travelogue; 85	Castillo San Marcos; 169	Escambia Bay; 119
Africa; 177	Cave Diving; 131	Bridge (I-10); 119
Agricultural Inspection Strs; 126	Cayo Costa, Name; 150	County; 120
Ah-Tah-Thi-Ki Museum; 160	Celebration; 93	Esterco; 153
Air Conditioning, First; 112	Charlotte County; 149	Everglade,90,95,139-140,154-160
Alabama; 124	Charlotte Harbor; 150	Drainage of; 156,181
Alachua; 132	Chautauqua; 116	Wildlife MA; 160
County; 131	Chipley; 114	Wonder Gardens; 154
Alafia River; 143	Name; 115	Falling Waters SP; 115
Alapaha, Name; 126	Choctawatchee, Name; 115	Fantasy of Flight; 95
Alfred B Maclay Gardens; 106	Circus Museum, Ringling; 147	Fayer Dykes SP; 171
Alligator Alley; 154-155	Citrus; 88,97,130,136,140,180	Fires, Forest; 166
Alligator Farm, St Augustine; 169	CityPlace, W Palm Beach; 180	Fires, Prescribed ; 148
Alligator Hole (definition); 157	City Maps,	Fisherman's Village; 151
Alligator, Buddy; 155	Fl Lauderdale Expwy; 194-195	Flagler County; 171
Alligators; 100,135,138,147,156	Jacksonville; 163	Flagler, Henry; 97,165,167,171
Anastasia Island; 170	Kissimmee Expwy; 192-193	Florida Aquarium; 186
Anhaica; 109-109,146	Miami Expressways; 194-195	Florida,
Apalachicola River; 112	Orlando Expressways; 192-193	12,000 years ago; 167
Appleton Mus of Art; 136	Pensacola; 26	Cavern SP; 114
Aquifer; 102	Tallahassee; 191	Map of all Expressways; 2-3
Arabic Nights; 94	Tampa-St. Petersburg; 63	Mus of Natural History; 134
Art Museum, Ringling; 147	St. Augustine; 191	National Cemetery ; 141
Aruba Beach Cafe; 183	Civil War; 100,108,127,138,141	Part of Africa; 177
Aucilla River Project; 106	Cleanwater Marine Aquarium; 167	Platform; 187
Babcock-Web WMA; 151	Collier County; 154	Sheriff's Boys Camp; 126
Bahia Mar Marina; 184	Collier, Barrow; 152	Sports Hall of Fame; 130
Baker County; 99	Colonial Spanish Quarters; 168	Sun 'n Fun Museum; 97
Barfoot Mailmen; 182	Columbia County; 101,128	Supreme Court; 107
Barge Canal; 137	Coquina Building Material; 165	Florida's Turnpike (FTP); 178,189
Bee Line Expy; 80	Corkscrew Swamp, Name; 154	25 mile Strip Maps; 66
Beltz Outlet Mall; 89	Cowboys; 95	Administration; 189
Bernard Castro; 136	Crab Trap II; 144	Coin System; 190
Big T; 165	Cracker, Florida; 88,95,132	Exit Services; 189
Big Cypress; 155,158	Crosslow Expy; 11,35,98,143	HEFT; 76,161,190
	Cuban Bread; 184	History; 189
	Dade Battlefield; 140	Names; 189
	Dade, Maj. Francis; 139-140,161	Service Plazas; 190

[Source: K. Grauman, slide credit: R. Urtasun]

Indexing Local Features: Inverted File Index

- For text documents, an efficient way to find all pages on which a word occurs is to use an index.
- We want to find all images in which a feature occurs.

Index		
Along I-75, From Detroit to Florida; <i>inside back cover</i>	Butterfly Center, McGuire; 134	Driving Lanes; 85
Drive I-95, From Boston to Florida; <i>inside back cover</i>	CAA (see AAA)	Duval County; 163
1929 Spanish Trail Roadway; Name; 150	CCC, The; 111,113,115,135,142	Eau Gallie; 175
511 Traffic Information; 83	Ca d'Zan; 147	Edison, Thomas; 152
A1A (Barrier Is) - I-95 Access; 86	Caloosahatchee River; 152	Eglin AFB; 116-118
AAA (and CAA); 83	Canaveral Natnl Seashore; 173	Eight Reals; 176
AAA National Office; 88	Cannon Creek Airpark; 130	Ellenton; 144-145
Abbreviations, Colored 25 mile Maps; cover	Canopy Road; 106,169	Emanuel Point Wreck; 120
Exit Services; 196	Cape Canaveral; 174	Emergency Callboxes; 83
Travelogue; 85	Castillo San Marcos; 169	Escambia Bay; 119
Africa; 177	Cave Diving; 131	Bridge (I-10); 119
Agricultural Inspection Strs; 126	Cayo Costa, Name; 150	County; 120
Ah-Tah-Thi-Ki Museum; 160	Celebration; 93	Esterco; 153
Air Conditioning, First; 112	Charlotte County; 149	Everglade,90,95,139-140,154-160
Alabama; 124	Charlotte Harbor; 150	Drainage of; 156,181
Alachua; 132	Chautauqua; 116	Wildlife MA; 160
County; 131	Chipley; 114	Wonder Gardens; 154
Alafia River; 143	Name; 115	Falling Waters SP; 115
Alapaha, Name; 126	Choctawatchee, Name; 115	Fantasy of Flight; 95
Alfred B MacLay Gardens; 106	Circus Museum, Ringling; 147	Fayer Dykes SP; 171
Alligator Alley; 154-155	Citrus; 88,97,130,136,140,180	Fires, Forest; 166
Alligator Farm, St Augustine; 169	CityPlace, W Palm Beach; 180	Fires, Prescribed ; 148
Alligator Hole (definition); 157	City Maps,	Fisherman's Village; 151
Alligator, Buddy; 155	Fl Lauderdale Expwys; 194-195	Flagler County; 171
Alligators; 100,135,138,147,156	Jacksonville; 163	Flagler, Henry; 97,165,167,171
Anastasia Island; 170	Kissimmee Expwys; 192-193	Florida Aquarium; 186
Anhaica; 109-109,146	Miami Expressways; 194-195	Florida,
Apalachicola River; 112	Orlando Expressways; 192-193	12,000 years ago; 167
Appleton Mus of Art; 136	Pensacola; 26	Cavern SP; 114
Aquifer; 102	Tallahassee; 191	Map of all Expressways; 2-3
Arabic Nights; 94	Tampa-St. Petersburg; 83	Mus of Natural History; 134
Art Museum, Ringling; 147	St. Augustine; 191	National Cemetery ; 141
Aruba Beach Cafe; 183	Civil War; 100,108,127,138,141	Part of Africa; 177
Aucilla River Project; 106	Cleanwater Marine Aquarium; 167	Platform; 187
Babcock-Web WMA; 151	Collier County; 154	Sheriff's Boys Camp; 126
Bahia Mar Marina; 184	Collier, Barrow; 152	Sports Hall of Fame; 130
Baker County; 99	Colonial Spanish Quarters; 168	Sun 'n Fun Museum; 97
Barfoot Mailmen; 182	Columbia County; 101,128	Supreme Court; 107
Barge Canal; 137	Coquina Building Material; 165	Florida's Turnpike (FTP); 178,189
Bee Line Expy; 80	Corkscrew Swamp, Name; 154	25 mile Strip Maps; 66
Beltz Outlet Mall; 89	Cowboys; 95	Administration; 189
Bernard Castro; 136	Crab Trap II; 144	Coin System; 190
Big T; 165	Cracker, Florida; 88,95,132	Exit Services; 189
Big Cypress; 155,158	Crosslow Expy; 11,35,98,143	HEFT; 76,161,190
	Cuban Bread; 184	History; 189
	Dade Battlefield; 140	Names; 189
	Dade, Maj. Francis; 139-140,161	Service Plazas; 190

[Source: K. Grauman, slide credit: R. Urtasun]

Indexing Local Features: Inverted File Index

- For text documents, an efficient way to find all pages on which a word occurs is to use an index.
- We want to find all images in which a feature occurs.
- To use this idea, we'll need to map our features to "visual words".

Index		
Along I-75, From Detroit to Florida; <i>inside back cover</i>	Butterfly Center, McGuire; 134	Driving Lanes; 85
Drive I-95, From Boston to Florida; <i>inside back cover</i>	CAA (see AAA)	Duval County; 163
1929 Spanish Trail Roadway; Name; 150	CCC, The; 111,113,115,135,142	Eau Gallie; 175
511 Traffic Information; 83	Ca d'Zan; 147	Edison, Thomas; 152
A1A (Barrier Is) - I-95 Access; 86	Caloosahatchee River; 152	Eglin AFB; 116-118
AAA (and CAA); 83	Canaveral Natrl Seashore; 173	Eight Reals; 176
AAA National Office; 88	Cannon Creek Airpark; 130	Ellenton; 144-145
Abbreviations, Colored 25 mile Maps; cover	Canopy Road; 106,169	Emanuel Point Wreck; 120
Exit Services; 196	Cape Canaveral; 174	Emergency Callboxes; 83
Travelogue; 85	Castillo San Marcos; 169	Escambia Bay; 119
Africa; 177	Cave Diving; 131	Bridge (I-10); 119
Agricultural Inspection Strs; 126	Cayo Costa, Name; 150	County; 120
Ah-Tah-Thi-Ki Museum; 160	Celebration; 93	Esterco; 153
Air Conditioning, First; 112	Charlotte County; 149	Everglade; 90,95,139-140,154-160
Alabama; 124	Charlotte Harbor; 150	Drainage of; 156,181
Alachua; 132	Chautauqua; 116	Wildlife MA; 160
County; 131	Chipley; 114	Wonder Gardens; 154
Alafia River; 143	Name; 115	Falling Waters SP; 115
Alapaha, Name; 126	Choctawatchee, Name; 115	Fantasy of Flight; 95
Alfred B MacIay Gardens; 106	Circus Museum, Ringling; 147	Fayer Dykes SP; 171
Alligator Alley; 154-155	Citrus; 88,97,130,136,140,180	Fires, Forest; 166
Alligator Farm, St Augustine; 169	CityPlace, W Palm Beach; 180	Fires, Prescribed ; 148
Alligator Hole (definition); 157	City Maps,	Fisherman's Village; 151
Alligator, Buddy; 155	Fl Lauderdale Expwys; 194-195	Flagler County; 171
Alligators; 100,135,138,147,156	Jacksonville; 163	Flagler, Henry; 97,165,167,171
Anastasia Island; 170	Kissimmee Expwys; 192-193	Florida Aquarium; 186
Anhaica; 109-109,146	Miami Expressways; 194-195	Florida,
Apalachicola River; 112	Orlando Expressways; 192-193	12,000 years ago; 167
Appleton Mus of Art; 136	Pensacola; 26	Cavern SP; 114
Aquifer; 102	Tallahassee; 191	Map of all Expressways; 2-3
Arabic Nights; 94	Tampa-St. Petersburg; 83	Mus of Natural History; 134
Art Museum, Ringling; 147	St. Augustine; 191	National Cemetery ; 141
Aruba Beach Cafe; 183	Civil War; 100,108,127,138,141	Part of Africa; 177
Aucilla River Project; 106	Cleanwater Marine Aquarium; 167	Platform; 187
Babcock-Web WMA; 151	Collier County; 154	Sheriff's Boys Camp; 126
Bahia Mar Marina; 184	Collier, Barrow; 152	Sports Hall of Fame; 130
Baker County; 99	Colonial Spanish Quarters; 168	Sun 'n Fun Museum; 97
Barfoot Mailmen; 182	Columbia County; 101,128	Supreme Court; 107
Barge Canal; 137	Coquina Building Material; 165	Florida's Turnpike (FTP); 178,189
Bee Line Expy; 80	Corkscrew Swamp, Name; 154	25 mile Strip Maps; 66
Beltz Outlet Mall; 89	Cowboys; 95	Administration; 189
Bernard Castro; 136	Crab Trap II; 144	Coin System; 190
Big T; 165	Cracker, Florida; 88,95,132	Exit Services; 189
Big Cypress; 155,158	Crosslow Expy; 11,35,98,143	HEFT; 76,161,190
	Cuban Bread; 184	History; 189
	Dade Battlefield; 140	Names; 189
	Dade, Maj. Francis; 139-140,161	Service Plazas; 190

[Source: K. Grauman, slide credit: R. Urtasun]

How would “visual words” help us?

Database of images




image 1

W1
W5
W4
⋮
W1




image 2

W2
W3
W6
⋮
W7




image 3

W7
W9
W1
⋮
W9

...




image hugeN

W6
W2
W7
⋮
W8

words

Imagine that I am somehow able to “name” my descriptors with a set of “words”.

How can this help me?

How would "visual words" help us?

Database of images




image 1

W1

W5

W4

⋮

W1




image 2

W2

W3

W6

⋮

W7




image 3

W7

W9

W1

⋮

W9

...




image hugeN

W6

W2

W7

⋮

W8

words

Visual word	Image
1	1,3
2	2,hugeN
3	2
4	1
5	1
6	2,hugeN
7	2,3,hugeN
...	

We can now build an **inverted file index**

This is like an Index of a book

How would "visual words" help us?

Database of images




image 1

W_1

W_5

W_4

\vdots

W_1




image 2

W_2

W_3

W_6

\vdots

W_7




image 3

W_7

W_9

W_1

\vdots

W_9

...




image hugeN

W_6

W_2

W_7


\vdots

W_8

words

Visual word	Image
1	1,3
2	2,hugeN
3	2
4	1
5	1
6	2,hugeN
7	2,3,hugeN
...	

We can also assign the descriptors in the reference image to the visual words



reference (query) image

How would "visual words" help us?

Database of images




image 1

W1
W5
W4
⋮
W2




image 2

W2
W3
W6
⋮
W7




image 3

W7
W1
W9
⋮
W91

...





image hugeN

W6
W2
W7
⋮
W8

words

Visual word	Image
1	1,3
2	2,hugeN
3	2
4	1
5	1
6	2,hugeN
7	2,3,hugeN
...	

And for each word in the reference image, we lookup our inverted file and check which images contain it.
We only need to match our reference image to the retrieved set of images.



reference (query) image

How would "visual words" help us?

Database of images




image 1




image 2




image 3

...




image hugeN

$$f_1^1 = [0.1, 0.2, \dots, 0.15]^T$$

$$f_2^1 = [0.23, 0.12, \dots, 0.1]^T$$

$$f_3^1 = [0.12, 0.15, \dots, 0.05]^T$$

⋮

$$f_n^1 = [0.05, 0.18, \dots, 0.09]^T$$

$$f_1^2 = [0.05, 0.11, \dots, 0.2]^T$$

$$f_2^2 = [0.09, 0.01, \dots, 0.18]^T$$

$$f_3^2 = [0.0, 0.08, \dots, 0.1]^T$$

⋮

$$f_m^2 = [0.1, 0.15, \dots, 0.14]^T$$

descriptors (vectors)

$$f_1^{hugeN} = [0.12, 0.15, \dots, 0.19]^T$$

$$f_2^{hugeN} = [0.1, 0.2, \dots, 0.2]^T$$

$$f_3^{hugeN} = [0.12, 0.22, \dots, 0.18]^T$$

⋮

$$f_k^{hugeN} = [0.15, 0.02, \dots, 0.08]^T$$

What are our visual "words"?


$$f_1^{ref} = [0.1, 0.2, \dots, 0.16]^T$$

$$f_2^{ref} = [0.15, 0.02, \dots, 0.06]^T$$

$$f_3^{ref} = [0.14, 0.22, \dots, 0.09]^T$$

⋮

$$f_p^{ref} = [0.17, 0.18, \dots, 0.2]^T$$



reference (query) image

How would "visual words" help us?

Database of images




image 1




image 2




image 3

...




image hugeN

$$f_1^1 = [0.1, 0.2, \dots, 0.15]^T$$

$$f_2^1 = [0.23, 0.12, \dots, 0.1]^T$$

$$f_3^1 = [0.12, 0.15, \dots, 0.05]^T$$

⋮

$$f_n^1 = [0.05, 0.18, \dots, 0.09]^T$$

$$f_1^2 = [0.05, 0.11, \dots, 0.2]^T$$

$$f_2^2 = [0.09, 0.01, \dots, 0.18]^T$$

$$f_3^2 = [0.0, 0.08, \dots, 0.1]^T$$

⋮

$$f_m^2 = [0.1, 0.15, \dots, 0.14]^T$$

descriptors (vectors)

$$f_1^{\text{hugeN}} = [0.12, 0.15, \dots, 0.19]^T$$

$$f_2^{\text{hugeN}} = [0.1, 0.2, \dots, 0.2]^T$$

$$f_3^{\text{hugeN}} = [0.12, 0.22, \dots, 0.18]^T$$

⋮

$$f_k^{\text{hugeN}} = [0.15, 0.02, \dots, 0.08]^T$$

The quest for visual words

We could do something like:

If all coordinates of vector smaller than 0.1, then call this vector word 1

If first n-1 coordinates < 0.1, but last coordinate is > 0.1, call this vector word 2

If first n-2 and last coordinate < 0.1, but n-1 coordinate > 0.1, call this vector word 3

...

Why is this not a very good choice? How can we do this better?

How would "visual words" help us?

Database of images




image 1




image 2




image 3

...




image hugeN

$$f_1^1 = [0.1, 0.2, \dots, 0.15]^T$$

$$f_2^1 = [0.23, 0.12, \dots, 0.1]^T$$

$$f_3^1 = [0.12, 0.15, \dots, 0.05]^T$$

⋮

$$f_n^1 = [0.05, 0.18, \dots, 0.09]^T$$

$$f_1^2 = [0.05, 0.11, \dots, 0.2]^T$$

$$f_2^2 = [0.09, 0.01, \dots, 0.18]^T$$

$$f_3^2 = [0.0, 0.08, \dots, 0.1]^T$$

⋮

$$f_m^2 = [0.1, 0.15, \dots, 0.14]^T$$

descriptors (vectors)

$$f_1^{hugeN} = [0.12, 0.15, \dots, 0.19]^T$$

$$f_2^{hugeN} = [0.1, 0.2, \dots, 0.2]^T$$

$$f_3^{hugeN} = [0.12, 0.22, \dots, 0.18]^T$$


⋮

$$f_k^{hugeN} = [0.15, 0.02, \dots, 0.08]^T$$

The quest for visual words

You can imagine each descriptor vector as a point in a high-dimensional space (128-dim for SIFT).

Disclaimer: This is only for the purpose of easier visualization of the solution.



How would "visual words" help us?

Database of images




image 1




image 2




image 3

...




image hugeN

$$f_1^1 = [0.1, 0.2, \dots, 0.15]^T$$

$$f_2^1 = [0.23, 0.12, \dots, 0.1]^T$$

$$f_3^1 = [0.12, 0.15, \dots, 0.05]^T$$

⋮

$$f_n^1 = [0.05, 0.18, \dots, 0.09]^T$$

$$f_1^2 = [0.05, 0.11, \dots, 0.2]^T$$

$$f_2^2 = [0.09, 0.01, \dots, 0.18]^T$$

$$f_3^2 = [0.0, 0.08, \dots, 0.1]^T$$

⋮

$$f_m^2 = [0.1, 0.15, \dots, 0.14]^T$$

descriptors (vectors)

$$f_1^{hugeN} = [0.12, 0.15, \dots, 0.19]^T$$

$$f_2^{hugeN} = [0.1, 0.2, \dots, 0.2]^T$$


$$f_3^{hugeN} = [0.12, 0.22, \dots, 0.18]^T$$

⋮

$$f_k^{hugeN} = [0.15, 0.02, \dots, 0.08]^T$$

The quest for visual words

- We can choose our visual words as "representative" vectors in this space
- We can perform **clustering** (for example **k-means**)



How would "visual words" help us?

Database of images




image 1




image 2




image 3

...




image hugeN

$$f_1^1 = [0.1, 0.2, \dots, 0.15]^T$$

$$f_2^1 = [0.23, 0.12, \dots, 0.1]^T$$

$$f_3^1 = [0.12, 0.15, \dots, 0.05]^T$$

⋮

$$f_n^1 = [0.05, 0.18, \dots, 0.09]^T$$

$$f_1^2 = [0.05, 0.11, \dots, 0.2]^T$$

$$f_2^2 = [0.09, 0.01, \dots, 0.18]^T$$

$$f_3^2 = [0.0, 0.08, \dots, 0.1]^T$$

⋮

$$f_m^2 = [0.1, 0.15, \dots, 0.14]^T$$

descriptors (vectors)

$$f_1^{hugeN} = [0.12, 0.15, \dots, 0.19]^T$$

$$f_2^{hugeN} = [0.1, 0.2, \dots, 0.2]^T$$

$$f_3^{hugeN} = [0.12, 0.22, \dots, 0.18]^T$$

⋮

$$f_k^{hugeN} = [0.15, 0.02, \dots, 0.08]^T$$

Visual words: cluster centers


$$\times W1 = [0.1, 0.15, \dots, 0.8]^T$$

$$\times W2 = [0.15, 0.01, \dots, 0.09]^T$$

$$\times W3 = [0.01, 0.09, \dots, 0.1]^T$$

$$\times W4 = [0.2, 0.02, \dots, 0.14]^T$$

⋮



How would "visual words" help us?

Database of images




image 1




image 2




image 3

...




image hugeN

$$\begin{aligned} f_1^1 &= [0.1, 0.2, \dots, 0.15]^T \\ f_2^1 &= [0.23, 0.12, \dots, 0.1]^T \\ f_3^1 &= [0.12, 0.15, \dots, 0.05]^T \\ &\vdots \\ f_n^1 &= [0.05, 0.18, \dots, 0.09]^T \end{aligned}$$

$$\begin{aligned} f_1^2 &= [0.05, 0.11, \dots, 0.2]^T \\ f_2^2 &= [0.09, 0.01, \dots, 0.18]^T \\ f_3^2 &= [0.0, 0.08, \dots, 0.1]^T \\ &\vdots \\ f_m^2 &= [0.1, 0.15, \dots, 0.14]^T \end{aligned}$$


descriptors (vectors)

$$\begin{aligned} f_1^{hugeN} &= [0.12, 0.15, \dots, 0.19]^T \\ f_2^{hugeN} &= [0.1, 0.2, \dots, 0.2]^T \\ f_3^{hugeN} &= [0.12, 0.22, \dots, 0.18]^T \\ &\vdots \\ f_k^{hugeN} &= [0.15, 0.02, \dots, 0.08]^T \end{aligned}$$

Visual words

- ✘ $W1 = [0.1, 0.15, \dots, 0.8]^T$
- ✘ $W2 = [0.15, 0.01, \dots, 0.09]^T$
- ✘ $W3 = [0.01, 0.09, \dots, 0.1]^T$
- ✘ $W4 = [0.2, 0.02, \dots, 0.14]^T$
- ⋮

How do we map this vector to a visual word?



How would "visual words" help us?

Database of images




image 1




image 2




image 3

...




image hugeN

$W1$

$$f_2^1 = [0.23, 0.12, \dots, 0.1]^T$$

$$f_3^1 = [0.12, 0.15, \dots, 0.05]^T$$

\vdots

$$f_n^1 = [0.05, 0.18, \dots, 0.09]^T$$

$$f_1^2 = [0.05, 0.11, \dots, 0.2]^T$$

$$f_2^2 = [0.09, 0.01, \dots, 0.18]^T$$

$$f_3^2 = [0.0, 0.08, \dots, 0.1]^T$$

\vdots

$$f_m^2 = [0.1, 0.15, \dots, 0.14]^T$$

descriptors (vectors)

$$f_1^{hugeN} = [0.12, 0.15, \dots, 0.19]^T$$

$$f_2^{hugeN} = [0.1, 0.2, \dots, 0.2]^T$$

$$f_3^{hugeN} = [0.12, 0.22, \dots, 0.18]^T$$

\vdots

$$f_k^{hugeN} = [0.15, 0.02, \dots, 0.08]^T$$

Visual words

$$\times W1 = [0.1, 0.15, \dots, 0.8]^T$$

$$\times W2 = [0.15, 0.01, \dots, 0.09]^T$$


$$\times W3 = [0.01, 0.09, \dots, 0.1]^T$$

$$\times W4 = [0.2, 0.02, \dots, 0.14]^T$$

\vdots



We find the closest visual word (Euclidean distance)

$$\arg \min_i \|f - Wi\|$$



Visual Words

- All example patches on the right belong to the same visual word.



Now We Can do Our Fast Matching

Database of images




image 1

W1
W5
W4
⋮
W2




image 2

W2
W3
W6
⋮
W7




image 3

W7
W1
W9
⋮
W91

...





image hugeN

W6
W2
W7
⋮
W8

words

Visual word	Image
1	1,3
2	2,hugeN
3	2
4	1
5	1
6	2,hugeN
7	2,3,hugeN
...	

And for each word in the reference image, we lookup our inverted file and check which images contain it.
We only need to match our reference image to the retrieved set of images.




W5
W1
W4
⋮
W1

reference (query) image

Inverted File Index

- Now we found all images in the database that have at least one visual word in common with the query image
- But this can still give us lots of images... What can we do?



Inverted File Index


- Now we found all images in the database that have at least one visual word in common with the query image
- But this can still give us lots of images... What can we do?
- Idea: Compute meaningful similarity (efficiently) between query image and retrieved images. Then just match query to top K most similar images and forget about the rest.

Inverted File Index

- Now we found all images in the database that have at least one visual word in common with the query image
- But this can still give us lots of images... What can we do?
- Idea: Compute meaningful similarity (efficiently) between query image and retrieved images. Then just match query to top K most similar images and forget about the rest.
- How can we do compute a meaningful similarity, and do it fast?


Relation to Documents

Of all the sensory impressions proceeding to the brain, the visual experiences are the dominant ones. Our perception of the world around us is based essentially on the messages that reach our eyes. For a long time, the retinal image was considered as a visual centers in the brain as a movie screen. The image is discovered through the perceptual process more complex following the path to the various cortical areas. Hubel and Wiesel have demonstrated that the message about the image falling on the retina undergoes a point-by-point analysis in a system of nerve cells stored in columns. In this system each cell has its specific function and is responsible for a specific detail in the pattern of the retinal image.



**sensory, brain,
visual, perception,
retinal, cerebral cortex,
eye, cell, optical
nerve, image
Hubel, Wiesel**


China is forecasting a trade surplus of \$90bn (£51bn) to \$100bn this year, a threefold increase on 2004's \$32bn. The Commerce Ministry said the surplus would be created by a predicted 30% increase in exports to \$750bn, compared with \$560bn in 2004. The government also needs to increase demand so that the country. China has permitted it to trade within a narrow band but the US wants the yuan to be allowed to trade freely. However, Beijing has made it clear that it will take its time and tread carefully before allowing the yuan to rise further in value.



**China, trade,
surplus, commerce,
exports, imports, US,
yuan, bank, domestic,
foreign, increase,
trade, value**

Bags of Visual Words

- Summarize entire image based on its distribution (histogram) of word occurrences.
- Analogous to bag of words representation commonly used for documents.



[Slide credit: R. Urtasun]

Compute a Bag-of-Words Description

Database of images




image 1




image 2




image 3

...




image hugeN

W1
W5
W4
⋮
W1

W2
W3
W6
⋮
W7

W7
W9
W1
⋮
W9

words

W6
W2
W7
⋮
W8

How many times a word repeats in image (frequency)




image 1 representation

[2 6 3 1 5 2 1 ...]

Compute a Bag-of-Words Description

Database of images




image 1

W1
W5
W4
⋮
W1




image 2

W2
W3
W6
⋮
W7




image 3

W7
W9
W1
⋮
W9

...





image hugeN


W6
W2
W7
⋮
W8

words



[2 6 3 1 5 2 1 ...]

We can do the same for the reference image



[1 2 4 5 1 2 2 ...]

Compute a Bag-of-Words Description

Database of images




image 1

W1
W5
W4
⋮
W1




image 2

W2
W3
W6
⋮
W7




image 3

W7
W9
W1
⋮
W9

...






image hugeN

W6
W2
W7
⋮
W8

words



How do we compare?

[2 6 3 1 5 2 1 ...] ← → [1 2 4 5 1 2 2 ...]

Comparing Images

- Compute the similarity by normalized dot product between their representations (vectors)

$$\text{sim}(\mathbf{t}_i, \mathbf{q}) = \frac{\langle \mathbf{t}_i, \mathbf{q} \rangle}{\|\mathbf{t}_i\| \cdot \|\mathbf{q}\|}$$

Comparing Images

- Compute the similarity by normalized dot product between their representations (vectors)

$$\text{sim}(\mathbf{t}_i, \mathbf{q}) = \frac{\langle \mathbf{t}_i, \mathbf{q} \rangle}{\|\mathbf{t}_i\| \cdot \|\mathbf{q}\|}$$

- Rank images in database based on the similarity score (the higher the better)
- Take top K best ranked images and do spatial verification (compute transformation and count inliers)

Compute a Better Bag-of-Words Description

Database of images




image 1




image 2




image 3

...




image hugeN


W1
W5
W4
⋮
W1

W2
W3
W6
⋮
W7

W7
W9
W1
⋮
W9

words

W6
W2
W7
⋮
W8



[2 6 3 1 5 2 1 ...]

Problem can quickly occur if one word appears in many many images and has a big count in each image (it dominates the vector)

This way any similarity based on this vector will be dominated with this very frequent, non-discriminative word.

Our similarity will not have much sense.

Compute a Better Bag-of-Words Description

Database of images




image 1

W1
W5
W4
⋮
W1




image 2

W2
W3
W6
⋮
W7




image 3

W7
W9
W1
⋮
W9

...





image hugeN

W6
W2
W7
⋮
W8

words



[2 6 3 1 5 2 1 ...]

Intuition:

Re-weight the entries such that words that appear in many images (documents) are down-weighted

This re-weighting is called **tf-idf**

Compute a Better Bag-of-Words Description

- Instead of a histogram, for retrieval it's better to re-weight the image description vector $t = [t_1, t_2, \dots, t_i, \dots]$ with **term frequency-inverse document frequency (tf-idf)**, a standard trick in document retrieval:

$$t_i = \frac{n_{id}}{n_d} \log \frac{N}{n_i}$$

Where:

n_{id} ...is the number of occurrences of word i in image d

n_d ...is the total number words in image d

n_i ...is the number of documents where the word i occurs

N ...is the number of documents in the whole database

Compute a Better Bag-of-Words Description

- Instead of a histogram, for retrieval it's better to re-weight the image description vector $t = [t_1, t_2, \dots, t_i, \dots]$ with **term frequency-inverse document frequency (tf-idf)**, a standard trick in document retrieval:

$$t_i = \frac{n_{id}}{n_d} \log \frac{N}{n_i}$$

Where:

n_{id} ...is the number of occurrences of word i in image d

n_d ...is the total number words in image d

n_i ...is the number of documents where the word i occurs

N ...is the number of documents in the whole database

- The weighting is a product of two terms: the word frequency $\frac{n_{id}}{n_d}$, and the inverse document frequency $\log \frac{N}{n_i}$

Compute a Better Bag-of-Words Description

- Instead of a histogram, for retrieval it's better to re-weight the image description vector $t = [t_1, t_2, \dots, t_i, \dots]$ with **term frequency-inverse document frequency (tf-idf)**, a standard trick in document retrieval:

$$t_i = \frac{n_{id}}{n_d} \log \frac{N}{n_i}$$

Where:

n_{id} ...is the number of occurrences of word i in image d

n_d ...is the total number words in image d

n_i ...is the number of documents where the word i occurs

N ...is the number of documents in the whole database

- **Intuition behind this:** word frequency weights words occurring often in a particular document, and thus describe it well, while the inverse document frequency describes how much information the word provides (is it common or rare across documents?)

Comparing Images


- Compute the similarity by normalized dot product between their tf-idf representations (vectors)

$$\text{sim}(\mathbf{t}_i, \mathbf{q}) = \frac{\langle \mathbf{t}_i, \mathbf{q} \rangle}{\|\mathbf{t}_i\| \cdot \|\mathbf{q}\|}$$

- Rank images in database based on the similarity score (the higher the better)
- Take top K best ranked images and do spatial verification (compute transformation and count inliers)

Spatial Verification

- Both image pairs have many visual words in common
- Only some of the matches are mutually consistent



Visual Words/Bags of Words

Good

- flexible to geometry / deformations / viewpoint
- compact summary of image content
- provides vector representation for sets
- good results in practice

Visual Words/Bags of Words

Good

- flexible to geometry / deformations / viewpoint
- compact summary of image content
- provides vector representation for sets
- good results in practice

Bad

- background and foreground mixed when bag covers whole image
- optimal vocabulary formation remains unclear
- basic model ignores geometry, must verify afterwards, or encode via features

Summary – Stuff You Need To Know

Fast image retrieval:

- Compute features in all images from database, and query image.

Summary – Stuff You Need To Know

Fast image retrieval:

- Compute features in all images from database, and query image.
- Cluster the descriptors from the images in the database (e.g., k-means) to get k clusters. These clusters are vectors that live in the same dimensional space as the descriptors. We call them visual words.

Summary – Stuff You Need To Know

Fast image retrieval:

- Compute features in all images from database, and query image.
- Cluster the descriptors from the images in the database (e.g., k-means) to get k clusters. These clusters are vectors that live in the same dimensional space as the descriptors. We call them visual words.
- Assign each descriptor in database and query image to the closest cluster.

Summary – Stuff You Need To Know

Fast image retrieval:

- Compute features in all images from database, and query image.
- Cluster the descriptors from the images in the database (e.g., k-means) to get k clusters. These clusters are vectors that live in the same dimensional space as the descriptors. We call them visual words.
- Assign each descriptor in database and query image to the closest cluster.
- Build an inverted file index

Summary – Stuff You Need To Know

Fast image retrieval:

- Compute features in all images from database, and query image.
- Cluster the descriptors from the images in the database (e.g., k-means) to get k clusters. These clusters are vectors that live in the same dimensional space as the descriptors. We call them visual words.
- Assign each descriptor in database and query image to the closest cluster.
- Build an inverted file index
- For a query image, lookup all the visual words in the inverted file index to get a list of images that share at least one visual word with the query

Summary – Stuff You Need To Know

Fast image retrieval:

- Compute features in all images from database, and query image.
- Cluster the descriptors from the images in the database (e.g., k-means) to get k clusters. These clusters are vectors that live in the same dimensional space as the descriptors. We call them visual words.
- Assign each descriptor in database and query image to the closest cluster.
- Build an inverted file index
- For a query image, lookup all the visual words in the inverted file index to get a list of images that share at least one visual word with the query
- Compute a bag-of-words (BoW) vector for each retrieved image and query. This vector just counts the number of occurrences of each word. It has as many dimensions as there are visual words. Weight the vector with tf-idf.

Summary – Stuff You Need To Know

Fast image retrieval:

- Compute features in all images from database, and query image.
- Cluster the descriptors from the images in the database (e.g., k-means) to get k clusters. These clusters are vectors that live in the same dimensional space as the descriptors. We call them visual words.
- Assign each descriptor in database and query image to the closest cluster.
- Build an inverted file index
- For a query image, lookup all the visual words in the inverted file index to get a list of images that share at least one visual word with the query
- Compute a bag-of-words (BoW) vector for each retrieved image and query. This vector just counts the number of occurrences of each word. It has as many dimensions as there are visual words. Weight the vector with tf-idf.
- Compute similarity between query BoW vector and all retrieved image BoW vectors. Sort (highest to lowest). Take top K most similar images (e.g, 100)

Summary – Stuff You Need To Know

Fast image retrieval:

- Compute features in all images from database, and query image.
- Cluster the descriptors from the images in the database (e.g., k-means) to get k clusters. These clusters are vectors that live in the same dimensional space as the descriptors. We call them visual words.
- Assign each descriptor in database and query image to the closest cluster.
- Build an inverted file index
- For a query image, lookup all the visual words in the inverted file index to get a list of images that share at least one visual word with the query
- Compute a bag-of-words (BoW) vector for each retrieved image and query. This vector just counts the number of occurrences of each word. It has as many dimensions as there are visual words. Weight the vector with tf-idf.
- Compute similarity between query BoW vector and all retrieved image BoW vectors. Sort (highest to lowest). Take top K most similar images (e.g, 100)
- Do spatial verification on all top K retrieved images (RANSAC + affine or homography + remove images with too few inliers)

Even Faster?

Can we make the retrieval process even more efficient?

Vocabulary Trees

- Hierarchical clustering for large vocabularies, [Nister et al., 06].

Vocabulary Trees

- Hierarchical clustering for large vocabularies, [Nister et al., 06].
- k defines the branch factor (number of children of each node) of the tree.

Vocabulary Trees

- Hierarchical clustering for large vocabularies, [Nister et al., 06].
- k defines the branch factor (number of children of each node) of the tree.
- First, an initial k -means process is run on the training data, defining k cluster centers (same as we did before).

Vocabulary Trees


- Hierarchical clustering for large vocabularies, [Nister et al., 06].
- k defines the branch factor (number of children of each node) of the tree.
- First, an initial k -means process is run on the training data, defining k cluster centers (same as we did before).
- The same process is then recursively applied to each group.

Vocabulary Trees

- Hierarchical clustering for large vocabularies, [Nister et al., 06].
- k defines the branch factor (number of children of each node) of the tree.
- First, an initial k -means process is run on the training data, defining k cluster centers (same as we did before).
- The same process is then recursively applied to each group.
- The tree is determined level by level, up to some maximum number of levels L .


Constructing the tree

- Offline phase: hierarchical clustering (e.g., k-means at each level).




Constructing the tree

- Offline phase: hierarchical clustering (e.g., k-means at each level).




Constructing the tree

- Offline phase: hierarchical clustering (e.g., k-means at each level).




Constructing the tree

- Offline phase: hierarchical clustering (e.g., k-means at each level).




Assigning Descriptors to Words




The words that I use to form the descriptor are the **leaves** of the tree

Assigning Descriptors to Words



$f_2^1 = [0.23, 0.12, \dots, 0.1]^T$
 $f_3^1 = [0.12, 0.15, \dots, 0.05]^T$
 \vdots
 $f_n^1 = [0.05, 0.18, \dots, 0.09]^T$


$\xrightarrow{?} W?$



How do I transform my (eg, SIFT) descriptors into such visual words?

Assigning Descriptors to Words


- Each descriptor vector is propagated down the tree by at each level comparing the descriptor vector to the k candidate cluster centers (represented by k children in the tree) and choosing the closest one.



Find the closest word at each level for a selected parent, starting from top

Assigning Descriptors to Words


- Each descriptor vector is propagated down the tree by at each level comparing the descriptor vector to the k candidate cluster centers (represented by k children in the tree) and choosing the closest one.



Find the closest word at each level for a selected parent, starting from top

Assigning Descriptors to Words

- The tree allows us to efficiently match a descriptor to a very large vocabulary



Assigning Descriptors to Words




image 1

W1
W5
W4
⋮
W1




image 2

W2
W3
W6
⋮
W7




image 3

W7
W9
W1
⋮
W9

...






image hugeN

W6
W2
W7
⋮
W8


words




[2 6 3 1 5 2 1 ...]



Querying Images




Querying Images



For all descriptors in the image calculate $q_i = n_i w_i$
where n_i is the number of descriptors with a path through node i

Querying Images




For all descriptors in the image calculate $q_i = n_i w_i$
where n_i is the number of descriptors with a path through node i

Do the same for all images in the database ($d_i = m_i w_i$) and retrieve images based on the similarity between d and q

Vocabulary Size

- Complexity: branching factor and number of levels
- Most important for the retrieval quality is to have a large vocabulary



Object Detection


Object Detection

- The goal of object detection is to localize objects in an image and tell their class
- Localization: place a tight bounding box around object



Object Detection

- The goal of object detection is to localize objects in an image and tell their class
- Localization: place a tight bounding box around object
- Can scale up to many classes using hierarchical tree of visual concepts



Type of Approaches

Different approaches tackle detection differently. They can roughly be categorized into three main types:

- Find interest points, followed by Hough voting

Type of Approaches

Different approaches tackle detection differently. They can roughly be categorized into three main types:

- Find interest points, followed by Hough voting
- Sliding windows: "slide" a box around image and classify each image crop inside a box (contains object or not?)

Type of Approaches

Different approaches tackle detection differently. They can roughly be categorized into three main types:

- Find interest points, followed by Hough voting
- Sliding windows: "slide" a box around image and classify each image crop inside a box (contains object or not?)
- Generate region (object) proposals, and classify each region

Type of Approaches

Different approaches tackle detection differently. They can roughly be categorized into three main types:

- Find interest points, followed by Hough voting
- Sliding windows: "slide" a box around image and classify each image crop inside a box (contains object or not?)
- Generate region (object) proposals, and classify each region

Interest Point Based Approaches


- Compute interest points (e.g., Harris corner detector is a popular choice)
- Vote for where the object could be given the content around interest points



Interest points

Interest Point Based Approaches


- Compute interest points (e.g., Harris corner detector is a popular choice)
- Vote for where the object could be given the content around interest points



Interest points

Interest Point Based Approaches

- Compute interest points (e.g., Harris corner detector is a popular choice)
- Vote for where the object could be given the content around interest points




- Is this part of cow?
- Where on cow have we see this?



Interest points

Interest Point Based Approaches


- Compute interest points (e.g., Harris corner detector is a popular choice)
- Vote for where the object could be given the content around interest points



Interest points

Interest Point Based Approaches

- Compute interest points (e.g., Harris corner detector is a popular choice)
- Vote for where the object could be given the content around interest points



Interest points

Type of Approaches

Different approaches tackle detection differently. They can roughly be categorized into three main types:

- Find interest points, followed by Hough voting
- Sliding windows: "slide" a box around image and classify each image crop inside a box (contains object or not?)
- Generate region (object) proposals, and classify each region

Sliding Window Approaches

- Slide window and ask a classifier: "Is sheep in window or not?"



0.1
confidence

Sliding Window Approaches

- Slide window and ask a classifier: "Is sheep in window or not?"



-0.2
confidence

Sliding Window Approaches

- Slide window and ask a classifier: "Is sheep in window or not?"



-0.1
confidence

Sliding Window Approaches

- Slide window and ask a classifier: "Is sheep in window or not?"



0.1
confidence

Sliding Window Approaches

- Slide window and ask a classifier: "Is sheep in window or not?"



1.5
confidence

Sliding Window Approaches

- Slide window and ask a classifier: "Is sheep in window or not?"



0.5
confidence

Sliding Window Approaches

- Slide window and ask a classifier: "Is sheep in window or not?"



0.4
confidence

Sliding Window Approaches

- Slide window and ask a classifier: "Is sheep in window or not?"



0.3
confidence

Sliding Window Approaches

- Slide window and ask a classifier: "Is sheep in window or not?"



Confidence

-0.1

0.2

-0.1

0.1

...

1.5

...

0.5

0.4

0.3

Type of Approaches

Different approaches tackle detection differently. They can roughly be categorized into three main types:

- Find interest points, followed by Hough voting
- Sliding windows: “slide” a box around image and classify each image crop inside a box (contains object or not?)
- Generate region (object) proposals, and classify each region

Region Proposal Based Approaches

- Group pixels into object-like regions



Region Proposal Based Approaches

- Group pixels into object-like regions




Region Proposal Based Approaches

- Generate many different regions



Region Proposal Based Approaches

- Generate many different regions



Region Proposal Based Approaches

- Generate many different regions



Region Proposal Based Approaches

- Generate many different regions



Region Proposal Based Approaches

- The hope is that at least a few will cover real objects



Region Proposal Based Approaches

- The hope is that at least a few will cover real objects



Region Proposal Based Approaches

- Select a region



Region Proposal Based Approaches

- Crop out an image patch around it, throw to classifier (e.g., Neural Net)



classifier
"dog" or not?

confidence: -2.5

Region Proposal Based Approaches

- Do this for every region



Region Proposal Based Approaches

- Do this for every region



Region Proposal Based Approaches

- Do this for every region



classifier
“dog” or not?

confidence: 1.5

Dog!!!



Type of Approaches

Different approaches tackle detection differently. They can roughly be categorized into three main types:

- Find interest points, followed by Hough voting ← Let's first look at one example method for this
- Sliding windows: "slide" a box around image and classify each image crop inside a box (contains object or not?)
- Generate region (object) proposals, and classify each region

Object Detection via Hough Voting: Implicit Shape Model

B. Leibe, A. Leonardis, B. Schiele

Robust Object Detection with Interleaved Categorization and
Segmentation


IJCV, 2008

Paper:

<http://www.vision.rwth-aachen.de/publications/pdf/leibe-interleaved-ijcv07final.pdf>

Start simple: line detection

- How can I find lines in this image?



[Source: K. Grauman]

Hough Transform


- Idea: Voting (Hough Transform)

Hough Transform

- Idea: Voting (Hough Transform)
- Voting is a general technique where we let the features vote for all models that are compatible with it.
 - Cycle through features, cast votes for model parameters.
 - Look for model parameters that receive a lot of votes.

Hough Transform: Line Detection


- Hough space: parameter space



- Connection between image (x, y) and Hough (m, b) spaces
 - A line in the image corresponds to a point in Hough space
 - What does a point (x_0, y_0) in the image space map to in Hough space?

Hough Transform: Line Detection



- Hough space: parameter space



- Connection between image (x, y) and Hough (m, b) spaces
 - A line in the image corresponds to a point in Hough space
 - A point in image space votes for all the lines that go through this point. These votes are a line in the Hough space.

Hough Transform: Line Detection

- Hough space: parameter space



- Two points: Each point corresponds to a line in the Hough space
- A point where these two lines meet defines a line in the image!

Hough Transform: Line Detection

- Hough space: parameter space





image space




Hough (parameter) space

- Vote with each image point
- Find peaks in Hough space. Each peak is a line in the image.

Hough Transform: Line Detection

- Issues with usual (m, b) parameter space: undefined for vertical lines
- A better representation is a polar representation of lines



d : perpendicular distance from line to origin



θ : angle the perpendicular makes with the x-axis

$$x \cos \theta - y \sin \theta = d$$

Point in image space \rightarrow sinusoid segment in Hough space

Example Hough Transform

- With the parameterization $x \cos \theta + y \sin \theta = d$
- Points in picture represent sinusoids in parameter space
- Points in parameter space represent lines in picture
- Example $0.6x + 0.4y = 2.4$, Sinusoids intersect at $d = 2.4$, $\theta = 0.9273$



[Source: M. Kazhdan, slide credit: R. Urtasun]

Hough Transform: Line Detection

- Hough Voting algorithm

Using the polar parameterization:

$$x \cos \theta - y \sin \theta = d$$


Basic Hough transform algorithm

1. Initialize $H[d, \theta] = 0$
2. for each edge point $I[x, y]$ in the image
for $\theta = [\theta_{\min} \text{ to } \theta_{\max}]$ // some quantization

$$d = x \cos \theta - y \sin \theta$$

$$H[d, \theta] += 1$$

3. Find the value(s) of (d, θ) where $H[d, \theta]$ is maximum
4. The detected line in the image is given by $d = x \cos \theta - y \sin \theta$



Hough Transform: Circle Detection


- What about circles? How can I fit circles around these coins?



Hough Transform: Circle Detection

Assume we are looking for a circle of known radius r


- Circle: $(x - a)^2 + (y - b)^2 = r^2$
- Hough space (a, b) : A point (x_0, y_0) maps to $(a - x_0)^2 + (b - y_0)^2 = r^2$
→ a circle around (x_0, y_0) with radius r
- Each image point votes for a circle in Hough space



Each point in geometric space (left) generates a circle in parameter space (right). The circles in parameter space intersect at the (a, b) that is the center in geometric space.


Hough Transform: Circle Detection

- What if we don't know r ?
 - Hough space: ?



Hough Transform: Circle Detection

- What if we don't know r ?
 - Hough space: conics




Hough Transform: Circle Detection

- Find the coins


Original



Edges



Votes: Penny



Hough Transform: Circle Detection

- Iris detection



Gradient+threshold




Hough space
(fixed radius)



Max detections

Generalized Hough Transform

Model:



ϕ -Table

Edge Direction	$\bar{\pi} = (\pi, \alpha)$
ϕ_1	$\bar{\pi}'_1, \bar{\pi}'_2, \bar{\pi}'_3$
ϕ_2	$\bar{\pi}^2_1, \bar{\pi}^2_2$
\vdots	\vdots
ϕ_i	$\bar{\pi}^i_1, \bar{\pi}^i_2$
\vdots	\vdots
ϕ_n	$\bar{\pi}^n_1, \bar{\pi}^n_2$

Generalized Hough Transform

Find Object Center (x_c, y_c) given edges (x_i, y_i, ϕ_i)

Create Accumulator Array $A(x_c, y_c)$

Initialize: $A(x_c, y_c) = 0 \quad \forall(x_c, y_c)$

For each edge point (x_i, y_i, ϕ_i)

For each entry \bar{r}_k^i in table, compute:

$$x_c = x_i + r_k^i \cos \alpha_k^i$$

$$y_c = y_i + r_k^i \sin \alpha_k^i$$

Increment Accumulator: $A(x_c, y_c) = A(x_c, y_c) + 1$

Find Local Maxima in $A(x_c, y_c)$


Results



Model




Model Detected



Hough Transform $A(x_c, y_c)$



Model



Model Detected




Hough Transform $A(x_c, y_c)$



Scale & Rotation:

Use Accumulator Array:

$$A[x_c, y_c, S, \theta]$$



Use:

$$x_c = x_i + r_k^i S \cos(\alpha_k^i + \theta)$$

$$y_c = y_i + r_k^i S \sin(\alpha_k^i + \theta)$$

$$A(x_c, y_c, S, \theta) = A(x_i, y_i, S, \theta) + 1.$$

A. Train phase:

1. Get features


2. Store all displacements of feature from center

B. Test phase:

1. Get features & lookup displacements

2. Vote for center location

Template



A. Train phase:



1. Get features

2. Store all displacements of feature from center

B. Test phase:

1. Get features & lookup displacements

2. Vote for center location



A. Train phase:



1. Get features

2. Store all displacements of feature from center

B. Test phase:

1. Get features & lookup displacements

2. Vote for center location




Implicit Shape Model

- Implicit Shape Model adopts the idea of voting
- Basic idea:
 - Find interest points in an image
 - Match patch around each interest point to a training patch
 - Vote for object center given that training instance

Implicit Shape Model: Basic Idea


- Vote for object center



vote for center of object

Implicit Shape Model: Basic Idea


- Vote for object center



vote for center of object

Implicit Shape Model: Basic Idea


- Vote for object center



vote for center of object

Implicit Shape Model: Basic Idea


- Vote for object center



of course some wrong votes are bound to happen...

Implicit Shape Model: Basic Idea


- Vote for object center



But that's ok. We want only **peaks** in voting space.

Implicit Shape Model: Basic Idea


- Find the patches that produced the peak



Find patches that voted for the peaks (back-projection).

Implicit Shape Model: Basic Idea


- Place a box around these patches → objects!



Find full objects based on the back-projected patches.

Implicit Shape Model: Basic Idea


- Really easy. Only one problem... Would be slow... How do we make it fast?



we need to match a patch around each yellow + to all patches in all training images → **SLOW**

Implicit Shape Model: Basic Idea

- Visual vocabulary (we saw this for retrieval)
- Compare each patch to a small set of visual words (clusters)



Visual words (visual codebook)!

Implicit Shape Model: Basic Idea

- Training: Getting the vocabulary


training image



Implicit Shape Model: Basic Idea

- Find interest points in each training image

training image




detect interest points (e.g. Harris)

Implicit Shape Model: Basic Idea

- Collect patches around each interest point
-

training image




extract an image patch around each interest point

Implicit Shape Model: Basic Idea

- Collect patches across all training examples

training images




collect all patches



Implicit Shape Model: Basic Idea

- Cluster the patches to get a small set of "representative" patches


training images



collect all patches



visual codebook




- cluster the patches to get a few "representative" patches
- each cluster represented as the average of all patches that belong to the cluster


clusters

Implicit Shape Model: Training

- Represent each training patch with the closest visual word.
- Record the displacement vectors for each word across all training examples.



Training image




Visual codeword with displacement vectors

Implicit Shape Model: Test

- At test times detect interest points
- Assign each patch around interest point to closes visual word
- Vote with all displacement vectors for that word



Recognition Pipeline



Recognition Summary

- Apply interest points and extract features around selected locations.
- Match those to the codebook.
- Collect consistent configurations using Generalized Hough Transform.
- Each entry votes for a set of possible positions and scales in continuous space.
- Extract maxima, localize in continuous space using Mean Shift.
- Refinement can be done by sampling more local features.

Example



Original Image

[Source: B. Leibe, credit: R. Urtasun]


Example



Interest points

[Source: B. Leibe, credit: R. Urtasun]

Example



Matched patches

[Source: B. Leibe, credit: R. Urtasun]


Example



Voting space

[Source: B. Leibe, credit: R. Urtasun]


Example



1st hypothesis

[Source: B. Leibe, credit: R. Urtasun]


Example



2nd hypothesis

[Source: B. Leibe, credit: R. Urtasun]

Example



3rd hypothesis

[Source: B. Leibe, credit: R. Urtasun]

Scale Invariant Voting

- Scale-invariant feature selection
 - Scale-invariant interest points
 - Based on patches around interest points, at training time a codebook of visual words is created.
 - Associated with each codebook entry, the displacements to object centre are stored along with the scale at which the interest point (mapped to the respective codebook entry) occurs... x_{occ} , y_{occ} , s_{occ} .

Scale Invariant Voting

- Generate scale votes (suppose an image feature was found at x_{img} , y_{img} , s_{img} , and gets mapped to a codebook entry that was observed while training at x_{occ} , y_{occ} , s_{occ} , then vote for the following location and scale:
 - Scale as 3rd dimension in voting space


$$x_{vote} = x_{img} - x_{occ}(s_{img}/s_{occ})$$

$$y_{vote} = y_{img} - y_{occ}(s_{img}/s_{occ})$$

$$s_{vote} = s_{img}/s_{occ}$$

- Search for maxima in 3D voting space


Scale Invariant Voting




[Slide credit: R. Urtasun]

Scale Voting: Efficient Computation


- Continuous Generalized Hough Transform
 - Binned accumulator array similar to standard Gen. Hough Transf.
 - Quickly identify candidate maxima locations
 - Refine locations by Mean-Shift search only around those points
 - Avoid quantization effects by keeping exact vote locations.




Scale votes



Binned accum. array





Candidate maxima



Refinement (Mean-Shift)

Extension: Rotation-Invariant Detection

- Polar instead of Cartesian voting scheme
- Recognize objects under image-plane rotations
- Possibility to share parts between articulations
- But also increases false positive detections



[Source: B. Leibe, credit: R. Urtasun]

Sometimes it's Necessary




[Figure from Mikolajczyk et al., CVPR'06]

Source: B. Leibe, credit: R. Urtasun]


Recognition and Segmentation

- Augment each visual word with meta-data: for example, segmentation mask




augment each cluster with a figure-ground mask

Recognition and Segmentation




Results



(a) detections

(b) p(figure)

(c) segmentation




(a) detections

(b) p(figure)

(c) segmentation

Results




[Source: B. Leibe]

Results




[Source: B. Leibe]

Results





[Source: B. Leibe]


Inferring Other Information: Part Labels



Inferring Other Information: Part Labels



Inferring Other Information: Depth



“Depth from a single image”

[Source: B. Leibe]

Deep learning

You Only Look Once: Unified, Real-Time Object Detection

Joseph Redmon*, Santosh Divvala*†, Ross Girshick†, Ali Farhadi*†

University of Washington*, Allen Institute for AI†, Facebook AI Research†

<http://pjreddie.com/yolo/>

Abstract

We present YOLO, a new approach to object detection. Prior work on object detection repurposes classifiers to perform detection. Instead, we frame object detection as a regression problem to spatially separated bounding boxes and associated class probabilities. A single neural network predicts bounding boxes and class probabilities directly from full images in one evaluation. Since the whole detection pipeline is a single network, it can be optimized end-to-end directly on detection performance.

Our unified architecture is extremely fast. Our base YOLO model processes images in real-time at 45 frames per second. A smaller version of the network, Fast YOLO, processes an astounding 155 frames per second while still achieving double the mAP of other real-time detectors. Compared to state-of-the-art detection systems, YOLO makes more localization errors but is less likely to predict false positives on background. Finally, YOLO learns very general representations of objects. It outperforms other detection methods, including DPM and R-CNN, when generalizing from natural images to other domains like artwork.

1. Introduction

Humans glance at an image and instantly know what objects are in the image, where they are, and how they interact. The human visual system is fast and accurate, allow-






Figure 1: The YOLO Detection System. Processing images with YOLO is simple and straightforward. Our system (1) resizes the input image to 448×448 , (2) runs a single convolutional network on the image, and (3) thresholds the resulting detections by the model's confidence.

methods to first generate potential bounding boxes in an image and then run a classifier on these proposed boxes. After classification, post-processing is used to refine the bounding boxes, eliminate duplicate detections, and rescore the boxes based on other objects in the scene [13]. These complex pipelines are slow and hard to optimize because each individual component must be trained separately.

We reframe object detection as a single regression problem, straight from image pixels to bounding box coordinates and class probabilities. Using our system, you only look once (YOLO) at an image to predict what objects are present and where they are.

YOLO is refreshingly simple: see Figure 1. A single convolutional network simultaneously predicts multiple bounding boxes and class probabilities for those boxes. YOLO trains on full images and directly optimizes detection performance. This unified model has several benefits






Some concluding thoughts...

What is computer vision?

- A field trying to develop automatic algorithms that can “see”



This course focused on standard techniques in vision and image processing


... But you have the skills to understand how state-of-the-art builds on these methods



Generate an image from a caption (stable diffusion)



"Dwayne Johnson side view"



"Dwayne Johnson top view"



The picture above is funny.

User What is funny about this image? Describe it panel by panel.



Source: [hmmm \(Reddit\)](#)

GPT-4 The image shows a package for a "Lightning Cable" adapter with three panels.

Panel 1: A smartphone with a VGA connector (a large, blue, 15-pin connector typically used for computer monitors) plugged into its charging port.

Panel 2: The package for the "Lightning Cable" adapter with a picture of a VGA connector on it.

Panel 3: A close-up of the VGA connector with a small Lightning connector (used for charging iPhones and other Apple devices) at the end.

The humor in this image comes from the absurdity of plugging a large, outdated VGA connector into a small, modern smartphone charging port.

The image features a central graphic consisting of several concentric circles. The outermost ring is a dark red, followed by a black ring, and then several more red rings of varying shades, creating a tunnel-like or target-like effect. In the center of these circles is a solid dark blue circle. Overlaid on this central blue circle is the text "That's all Folks!" written in a white, elegant cursive font. The text is slightly tilted and positioned across the middle of the blue circle.

That's all Folks!