



Rich feature hierarchies for accurate object detection and semantic segmentation

Girshick, Donahue, Darrell, Malik
CVPR 2014

Total citations	2870
Citations per year	957
Equations per page	0
Oldest citation*	1980
Self citations	6
Dataset size (log ₁₀)	4.1

*earliest year wins



ImageNet large scale visual recognition challenge

Russakovsky, Deng, Su, Krause, Satheesh, Ma, Huang, Karpathy, Khosla, Bernstein, Berg, Fei-Fei
IJCV 2015

Total citations	2427
Citations per year	1214
Equations per page	0.16
Oldest citation*	1995
Self citations	11
Dataset size (log ₁₀)	6.2

*earliest year wins



Very deep convolutional networks for large-scale image recognition

Simonyan, Zisserman
ICLR 2015

Total citations	4261
Citations per year	2131
Equations per page	0
Oldest citation*	2004
Self citations	3
Dataset size (log ₁₀)	6.2

*earliest year wins



Deep residual learning for image recognition

He, Zhang, Ren, Sun
CVPR 2016

Total citations	1919
Citations per year	1919
Equations per page	0.16
Oldest citation*	1989
Self citations	5
Dataset size (log ₁₀)	6.2

*earliest year wins



A taxonomy and evaluation of dense two-frame stereo correspondence algorithms

Scharstein, Szeliski, Zabih
SMBV 2001

Total citations	6585
Citations per year	412
Equations per page	0.8
Oldest citation*	1974
Self citations	15
Dataset size (log ₁₀)	1.8

*earliest year wins



Mean shift: A robust approach toward feature space analysis

Comaniciu, Meer
PAMI 2002

Total citations	10485
Citations per year	699
Equations per page	3.06
Oldest citation*	1973
Self citations	6
Dataset size (log ₁₀)	1.2

*earliest year wins



Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection

Belhumeur, Hespanha, Kriegman
PAMI 1997

Total citations	12387
Citations per year	619
Equations per page	1.1
Oldest citation*	1936
Self citations	2
Dataset size (log ₁₀)	2.7

*earliest year wins



Paper Legend: The Rules

To start the game, shuffle and deal out the cards face down. Each player holds their cards so that they can see the top card only.

The player to the dealer's left starts by reading out a category from the top card (e.g. self citations, value 1). The other players then read out the same category from their cards. The player with the highest value wins, except in the 'oldest citation' category, where the earliest year wins. The winner of the round collects everyone's top card and moves them to the bottom of their pile, along with their own winning card. It is then their turn again to choose a category from the next card at the top of their pile.

If two or more cards share the top value then all of the cards are placed in the middle and the same player chooses again from their next card. The winner of the hand then takes the cards in the middle as well.

If one of the players has authored the paper on their card they automatically win the round. As above, if two or more players have authored papers on the same round, all cards go into the middle.

The person with all of the cards at the end is the winner.

Citations correct as of May 2017.





 Microsoft



 Microsoft



 Microsoft



 Microsoft



 Microsoft



 Microsoft



 Microsoft



 Microsoft



Fast approximate energy minimization via graph cuts

Boykov, Veksler, Zabih
PAMI 2001

Total citations	7045
Citations per year	440
Equations per page	2.39
Oldest citation*	1952
Self citations	6
Dataset size (log ₁₀)	0.8

*earliest year wins



Shape matching and object recognition using shape contexts

Belongie, Malik, Puzicha
PAMI 2002

Total citations	6515
Citations per year	434
Equations per page	1.64
Oldest citation*	1917
Self citations	3
Dataset size (log ₁₀)	4.9

*earliest year wins



Robust wide-baseline stereo from maximally stable extremal regions

Matas, Chum, Urban, Pajdla
BMVC 2002

Total citations	4021
Citations per year	268
Equations per page	0
Oldest citation*	1988
Self citations	0
Dataset size (log ₁₀)	1.1

*earliest year wins



Video google: A text retrieval approach to object matching in videos

Sivic, Zisserman
ICCV 2003

Total citations	5504
Citations per year	393
Equations per page	0.25
Oldest citation*	1994
Self citations	2
Dataset size (log ₁₀)	3.9

*earliest year wins



A performance evaluation of local descriptors

Schmid, Mikolajczyk
ICPR 2003

Total citations	7794
Citations per year	557
Equations per page	0.29
Oldest citation*	1946
Self citations	4
Dataset size (log ₁₀)	3.0

*earliest year wins



Learning generative visual models from few training examples

Fei-Fei, Fergus, Perona
GMBV 2004

Total citations	2767
Citations per year	213
Equations per page	1.3
Oldest citation*	1987
Self citations	4
Dataset size (log ₁₀)	4.0

*earliest year wins



Distinctive image features from scale-invariant keypoints

Lowe
IJCV 2004

Total citations	41477
Citations per year	3191
Equations per page	0.75
Oldest citation*	1962
Self citations	8
Dataset size (log ₁₀)	2.1

*earliest year wins



Histograms of oriented gradients for human detection

Dalal, Triggs
CVPR 2005

Total citations	18739
Citations per year	1562
Equations per page	0
Oldest citation*	1977
Self citations	1
Dataset size (log ₁₀)	3.3

*earliest year wins





 Microsoft



 Microsoft



 Microsoft



 Microsoft



 Microsoft



 Microsoft



 Microsoft



 Microsoft