


Nonlinear Component Analysis as a Kernel Eigenvalue Problem

Schölkopf, Smola, Müller
Neural Computation 1998

Citations per year	365
Equations per page	1.57
Oldest citation*	1953
Self citations	2
Hinton references	0

If you authored this paper you win this round!

*earliest year wins




Temporal Constraint Networks

Dechter, Meiri, Pearl
Artificial Intelligence 1991

Citations per year	83
Equations per page	2.50
Oldest citation*	1962
Self citations	4
Hinton references	0

If you authored this paper you win this round!

*earliest year wins




A Unifying Review of Linear Gaussian Models

Roweis, Ghahramani
Neural Computation 1999

Citations per year	49
Equations per page	1.51
Oldest citation*	1960
Self citations	5
Hinton references	8

If you authored this paper you win this round!

*earliest year wins




Exploiting Generative Models in Discriminative Classifiers

Jaakkola, Haussler
NIPS 1999

Citations per year	81
Equations per page	2.18
Oldest citation*	1990
Self citations	1
Hinton references	0

If you authored this paper you win this round!

*earliest year wins




An Information-maximization Approach to Blind Separation and Blind Deconvolution

Bell, Sejnowski
Neural Computation 1995

Citations per year	378
Equations per page	1.39
Oldest citation*	1961
Self citations	3
Hinton references	2

If you authored this paper you win this round!

*earliest year wins



Gradient-based Learning Applied to Document Recognition

LeCun, Bottou, Bengio, Haffner
Proc. IEEE 1998

Citations per year	328
Equations per page	0.63
Oldest citation*	1961
Self citations	39
Hinton references	4

If you authored this paper you win this round!

*earliest year wins




Divergence Measures and Message Passing

Minka
Techreport 2005

Citations per year	24
Equations per page	8.29
Oldest citation*	1966
Self citations	5
Hinton references	0

If you authored this paper you win this round!

*earliest year wins



Q-learning

Watkins, Dayan
Machine Learning 1992

Citations per year	251
Equations per page	2.57
Oldest citation*	1962
Self citations	1
Hinton references	0

If you authored this paper you win this round!

*earliest year wins





 Microsoft



 Microsoft



 Microsoft



 Microsoft



 Microsoft




 Microsoft



 Microsoft



 Microsoft




Learning Long-term Dependencies with Gradient Descent is Difficult
Bengio, Simard, Frasconi
 IEEE Trans. Neural Net. 94

Citations per year	53
Equations per page	2.00
Oldest citation*	1960
Self citations	6
Hinton references	2

If you authored this paper you win this round!

**earliest year wins*




Learning to Predict by the Methods of Temporal Differences
Sutton
 Machine Learning 1988

Citations per year	150
Equations per page	1.61
Oldest citation*	1959
Self citations	5
Hinton references	2

If you authored this paper you win this round!

**earliest year wins*



Policy Gradient Methods for Reinforcement Learning with Function Approximation
Sutton, McAllester, Singh, Mansour
 NIPS 1999

Citations per year	57
Equations per page	4.71
Oldest citation*	1983
Self citations	5
Hinton references	0

If you authored this paper you win this round!

**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 October 2016.





 Microsoft



 Microsoft



 Microsoft



 Microsoft