UNIVERSITY OF AMSTERDAM

Blendle

Multileave Gradient Descent for Fast Online Learning to Rank

Anne Schuth[1,2], Harrie Oosterhuis[1], Shimon Whiteson[3] and Maarten de Rijke[1]

[1]University of Amsterdam, [2]Blendle, [3]University of Oxford

Search Engines that Learn

- Search engines are complex machines that base their ranking hundreds of signals

Search Engines that Learn

- Search engines are complex machines that base their ranking hundreds of signals
- Learn to combine these signals

Search Engines that Learn

- Search engines are complex machines that base their ranking hundreds of signals
- Learn to combine these signals
  - Offline: using labeled, static datasets

Search Engines that Learn

- Search engines are complex machines that base their ranking hundreds of signals
- Learn to combine these signals
  - Offline: using labeled, static datasets
  - **Online: directly from users**

Existing work

Dueling Bandit Gradient Descent (DBGD)

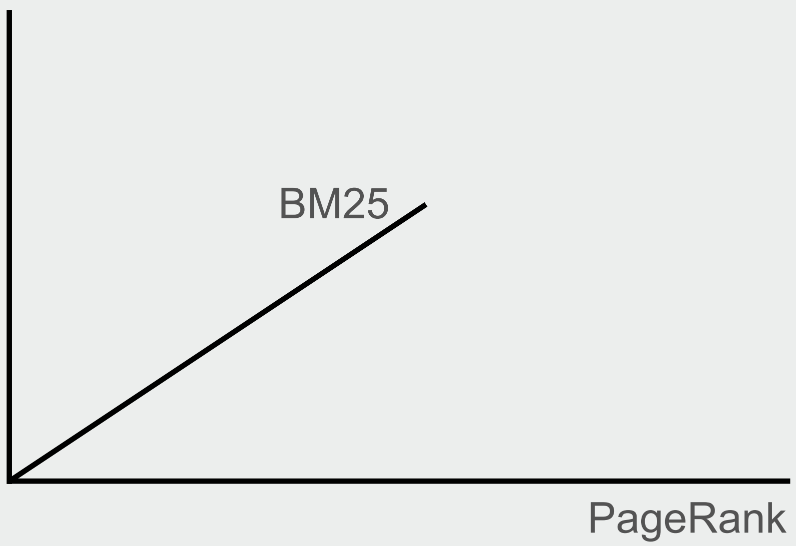

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

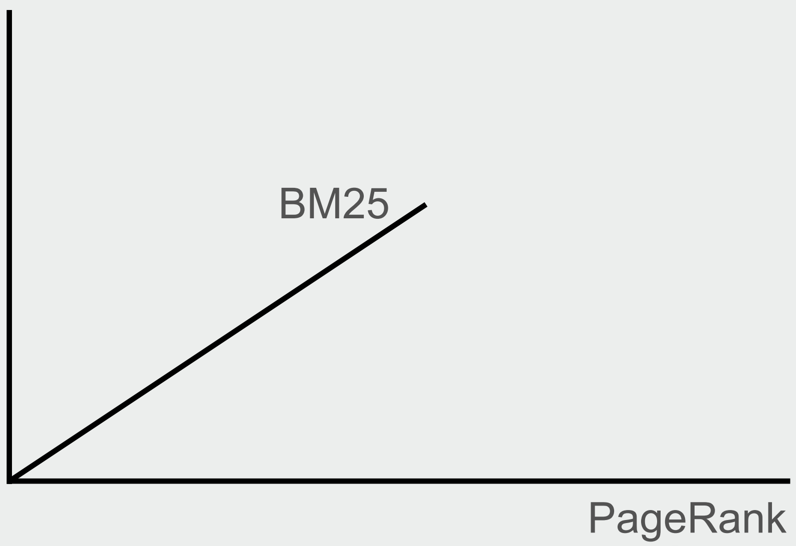

Weights for features

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

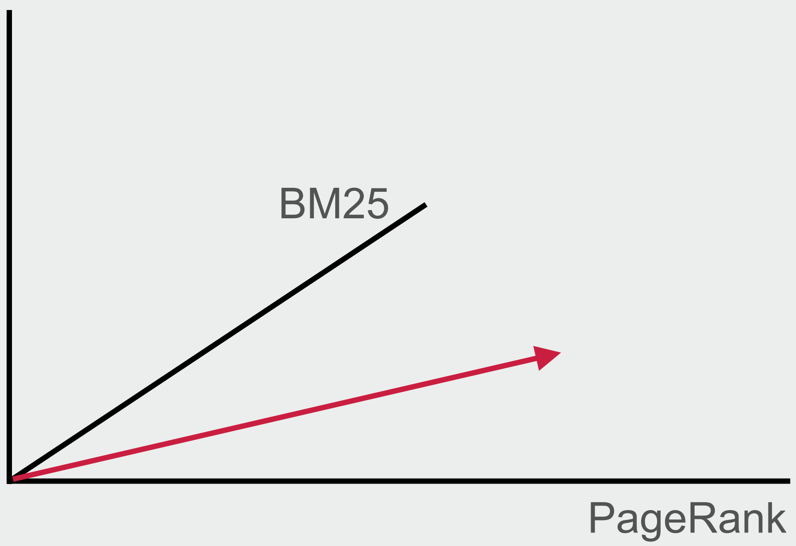

Weights for features

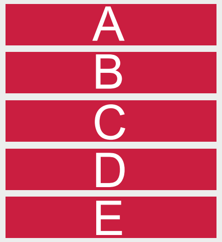

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

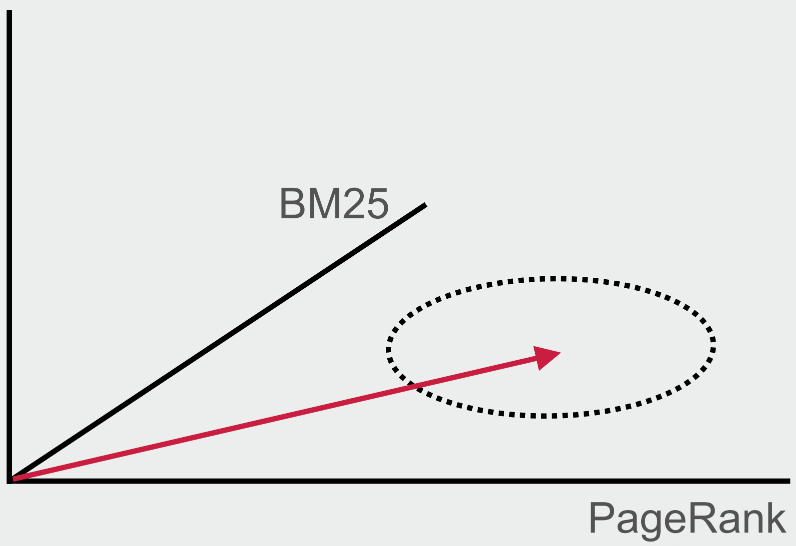

Weights for features

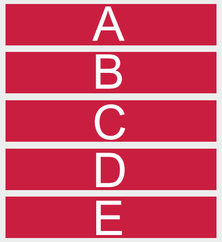

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

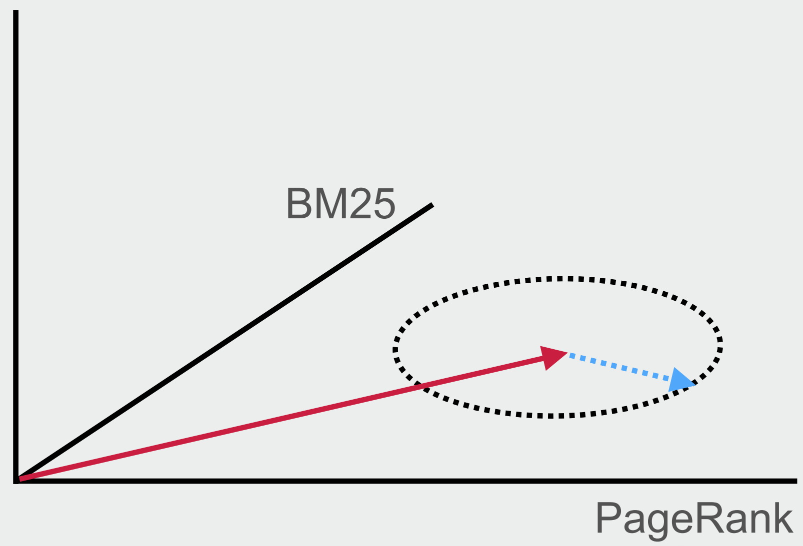

Weights for features

A
B
C
D
E

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

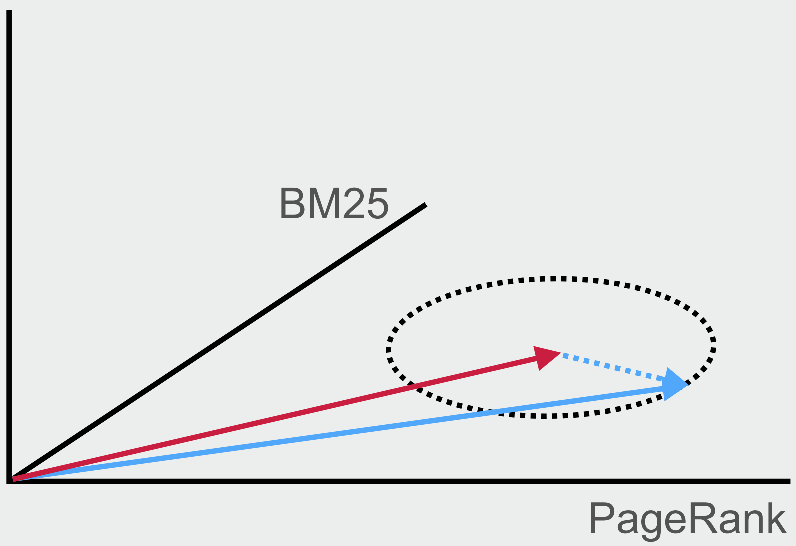

Weights for features

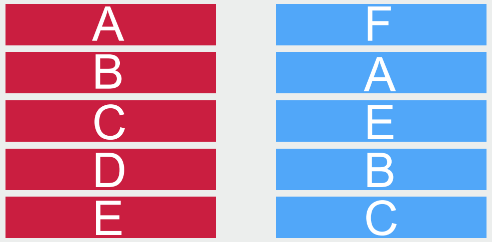

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

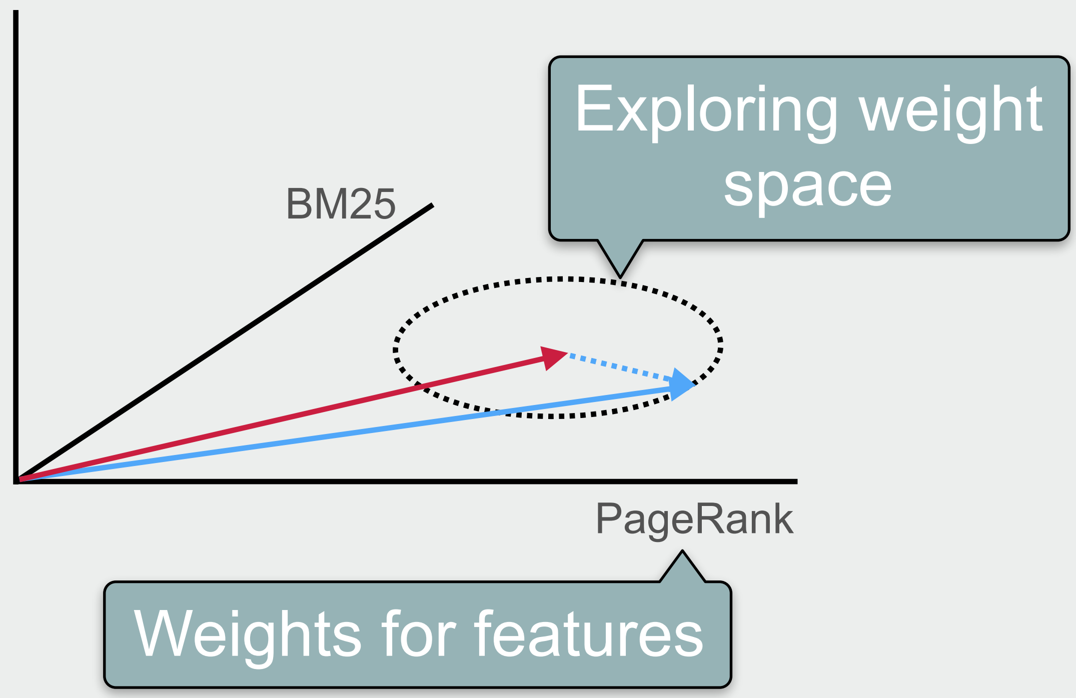

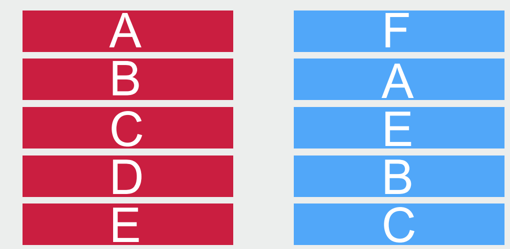

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

Preferences through **interleaved** comparisons

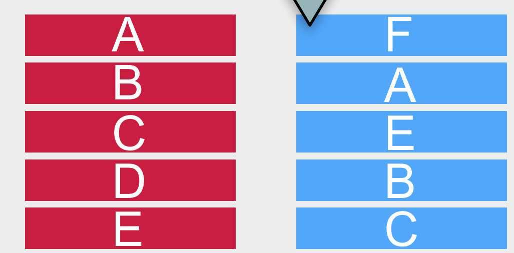

Exploring weight space

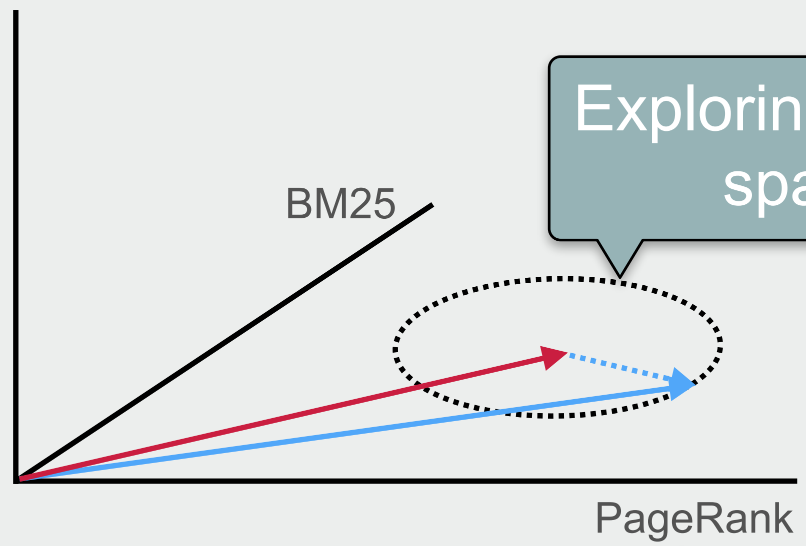

Weights for features

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

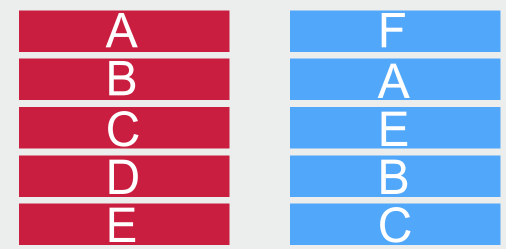

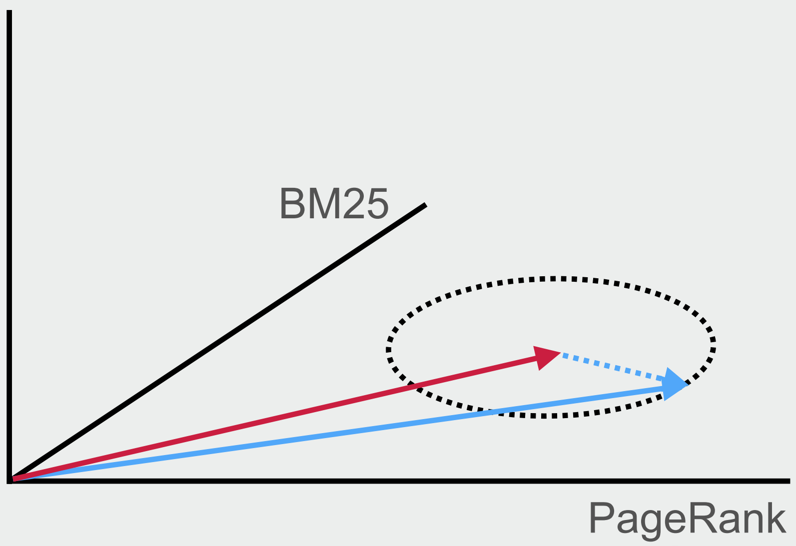

Weights for features

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

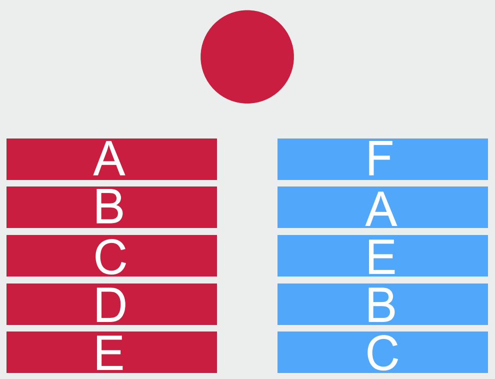

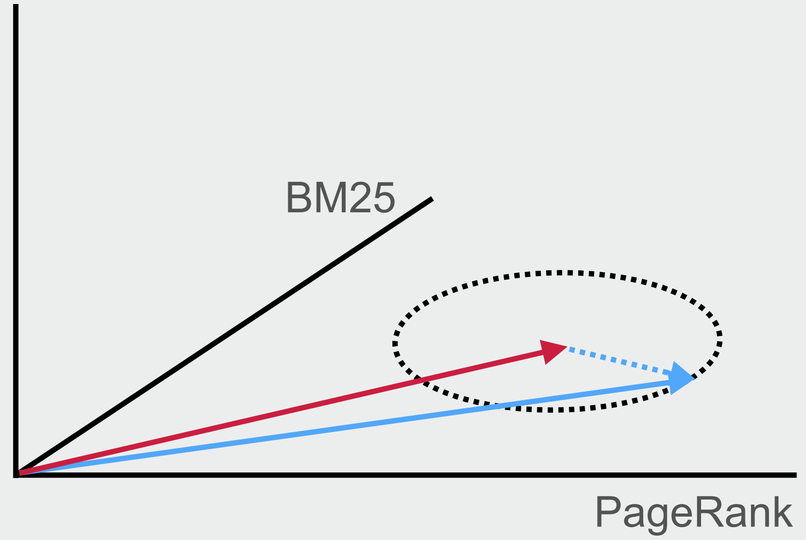

Weights for features

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

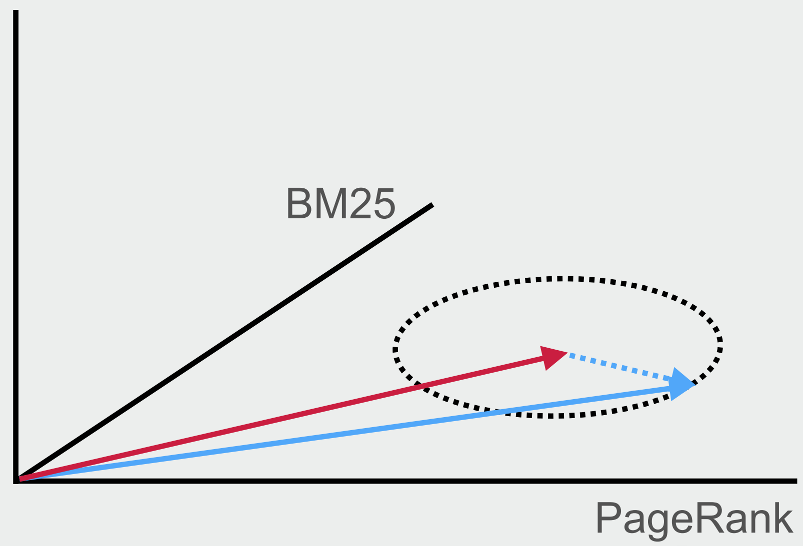

Weights for features

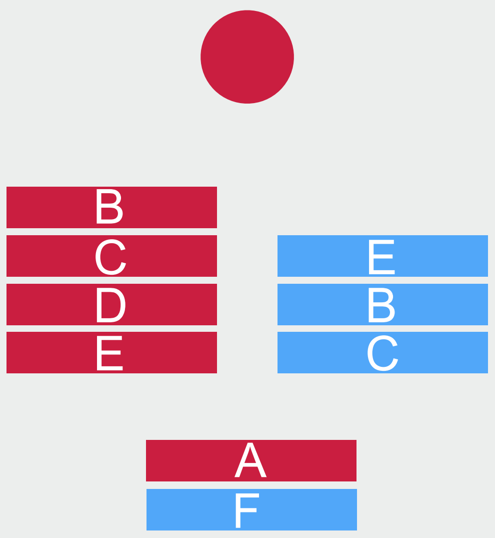

[Yue et al., 2009; Hofmann et al., 2011;
Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

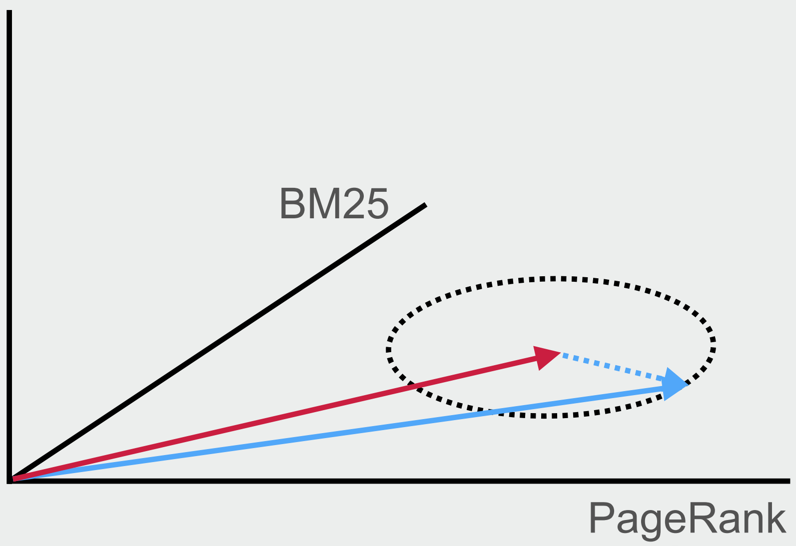

Weights for features

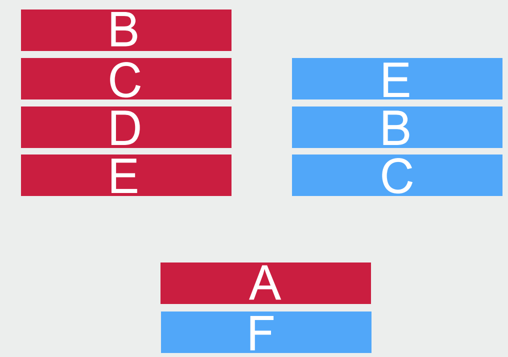

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

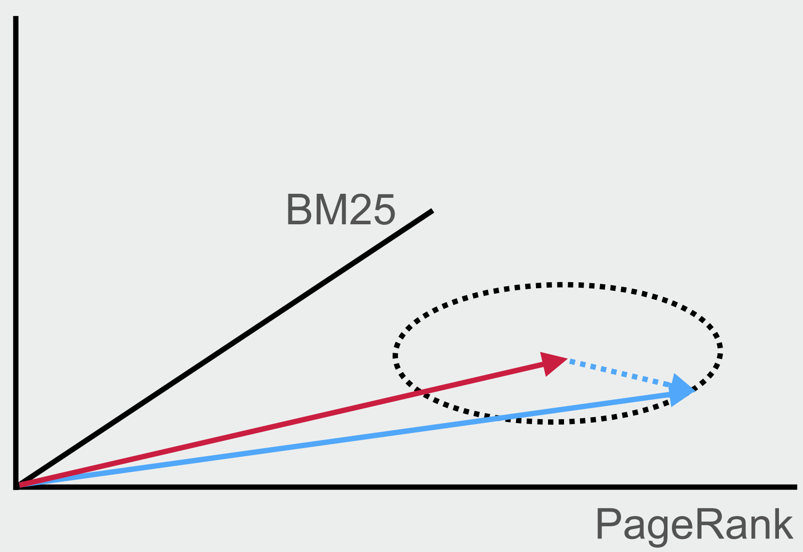

Weights for features

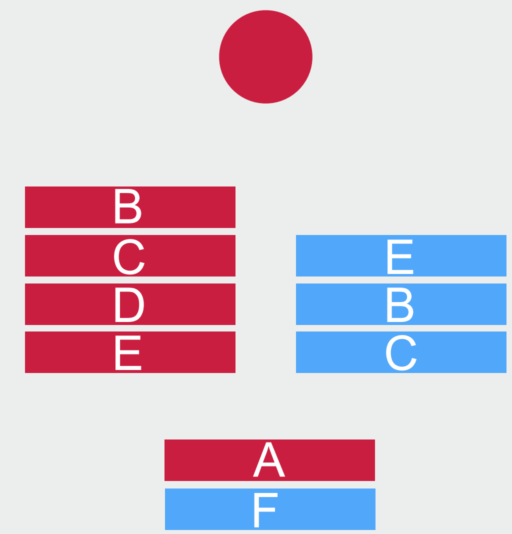

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

University of Amsterdam

Existing work

Dueling Bandit Gradient Descent (DBGD)

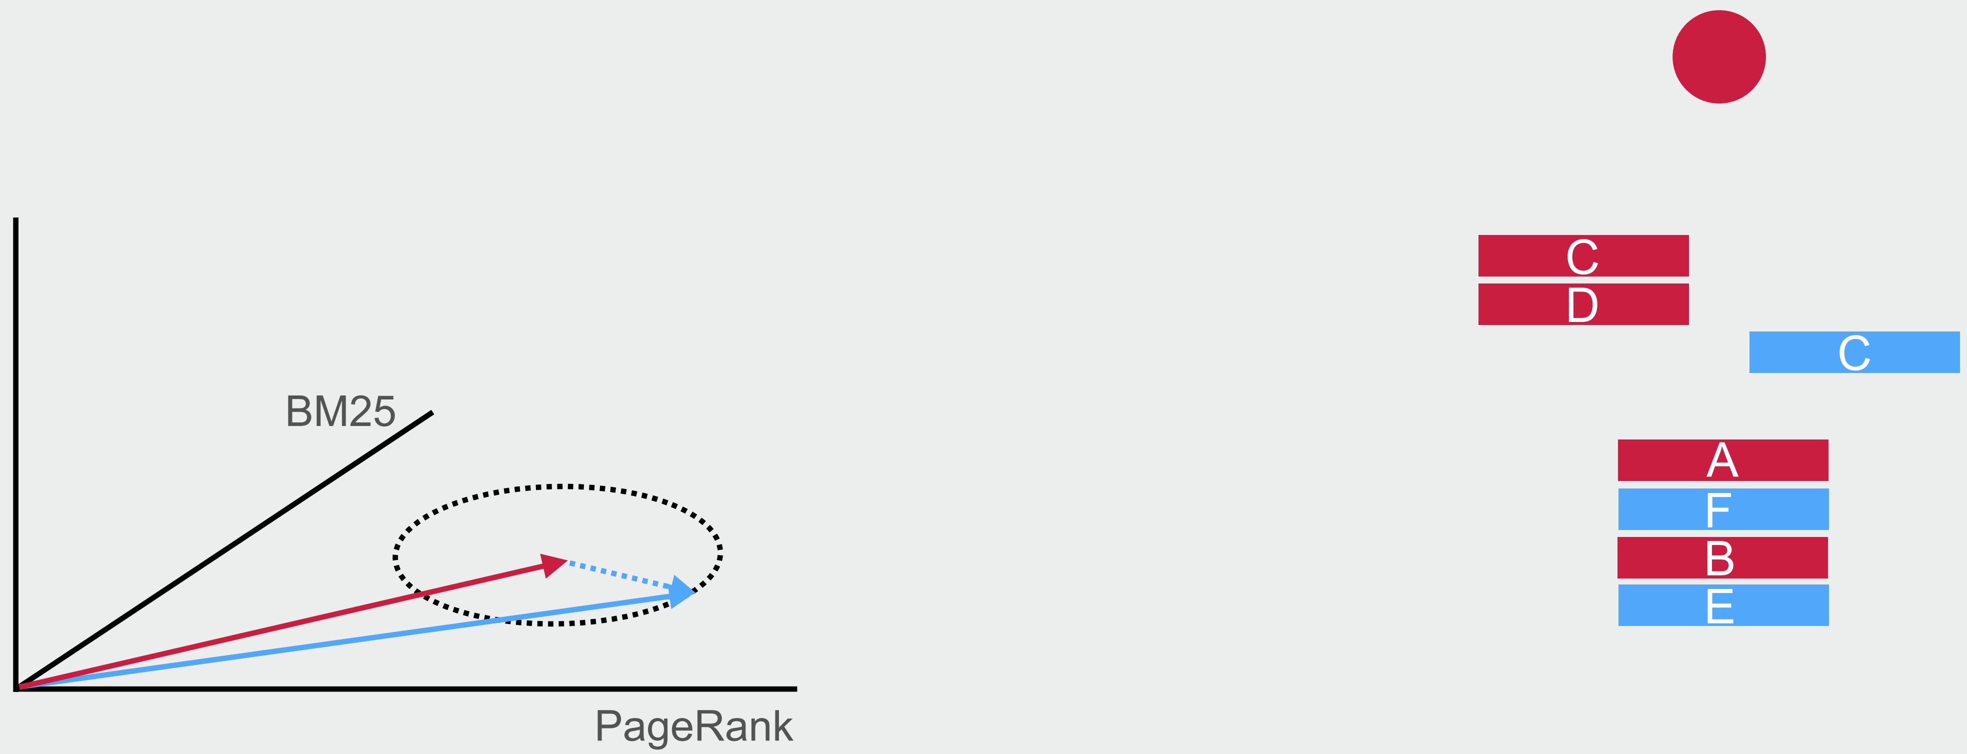

Weights for features

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

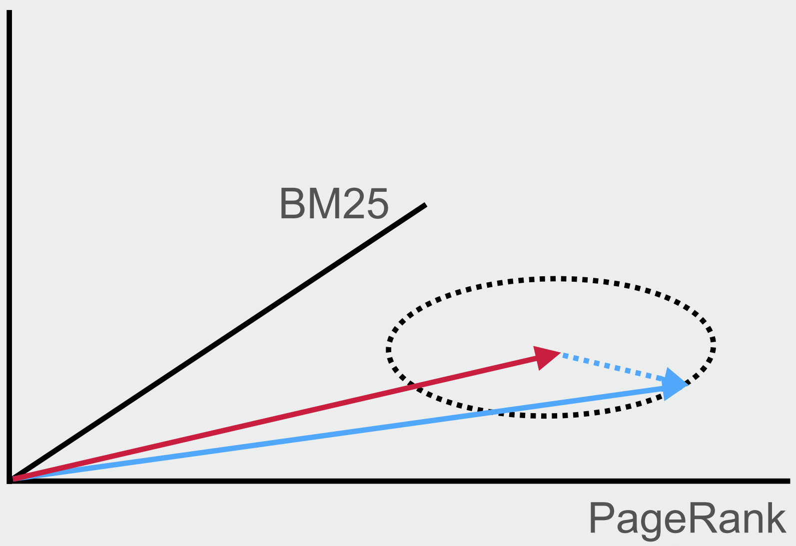

Weights for features

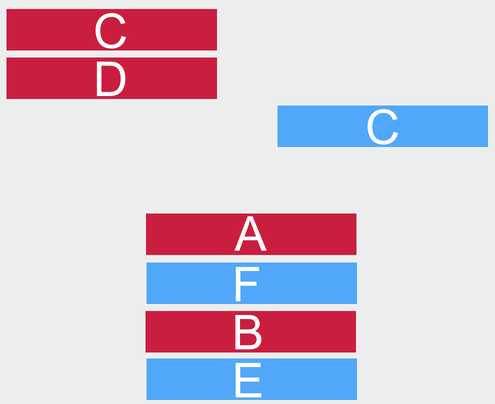

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

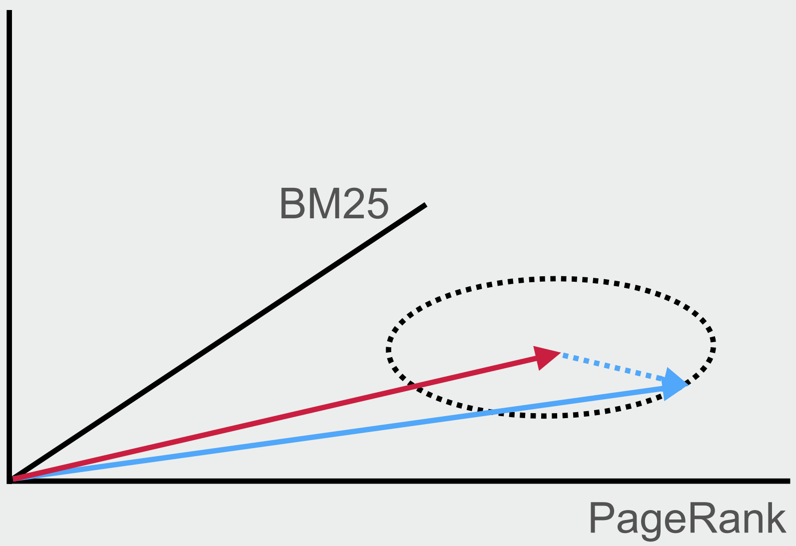

Weights for features

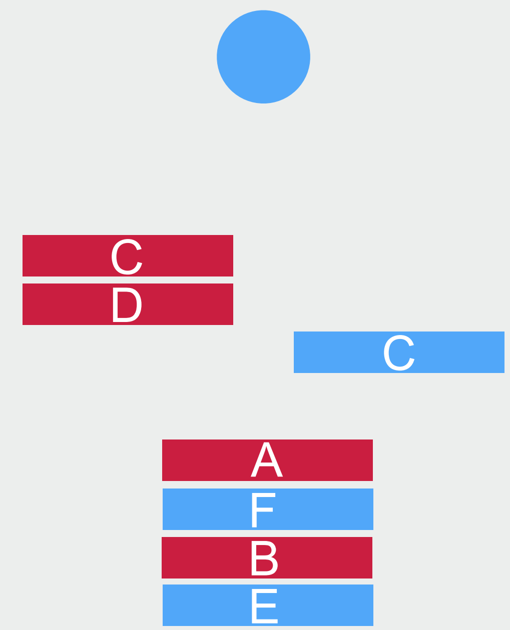

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

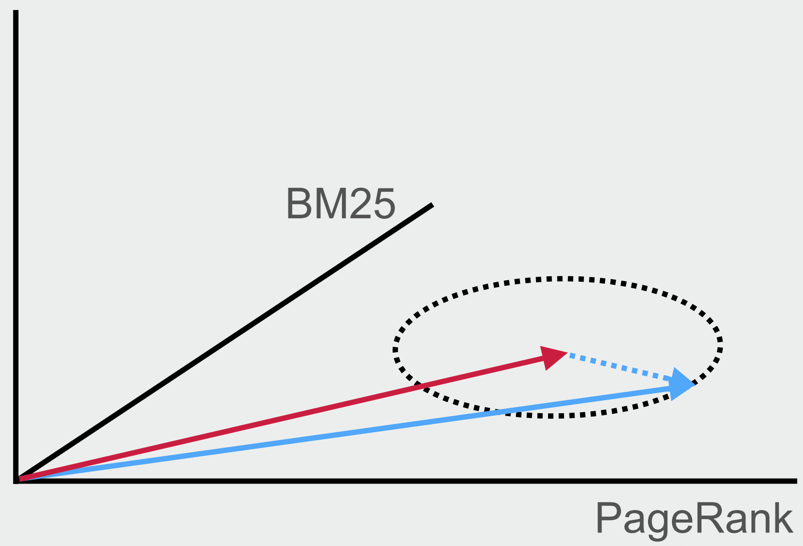

Weights for features

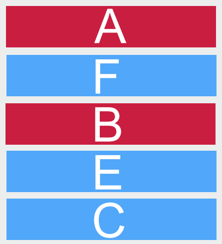

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

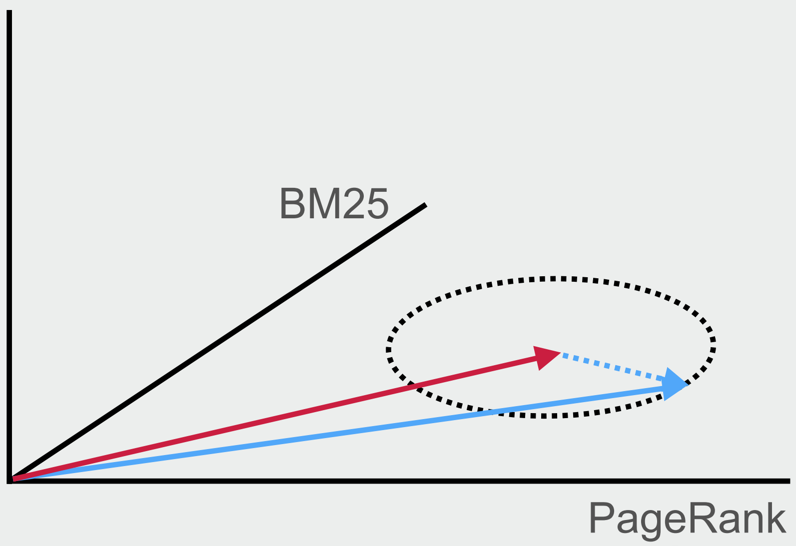

Weights for features

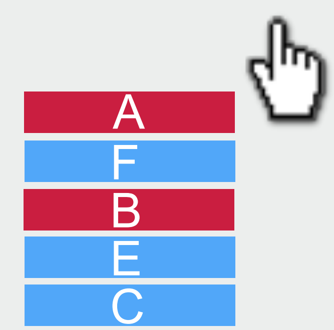

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

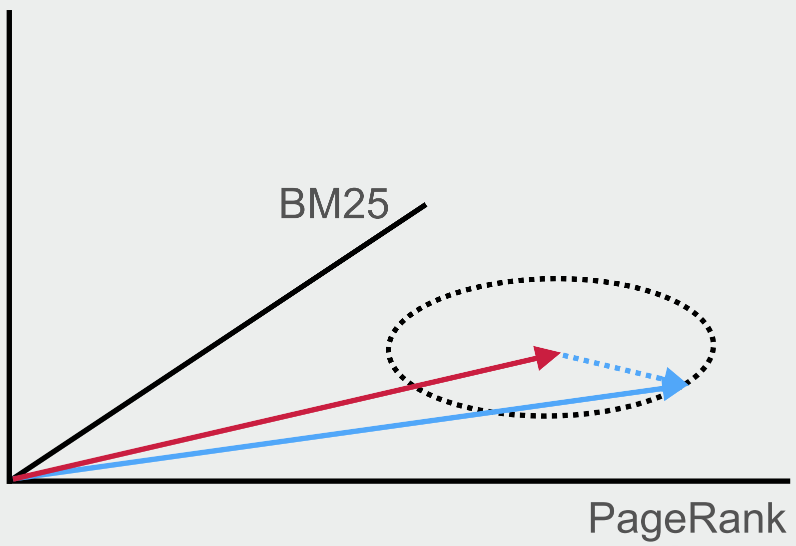

Weights for features

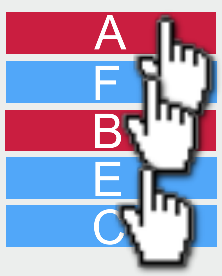

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

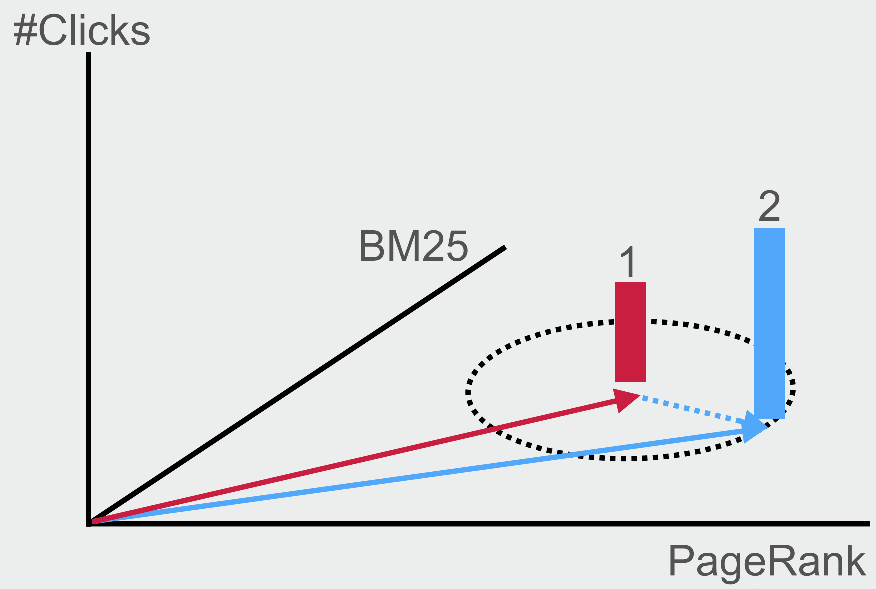

Weights for features

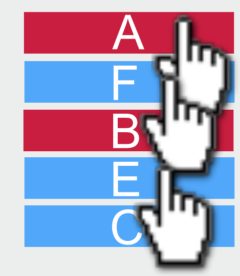

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

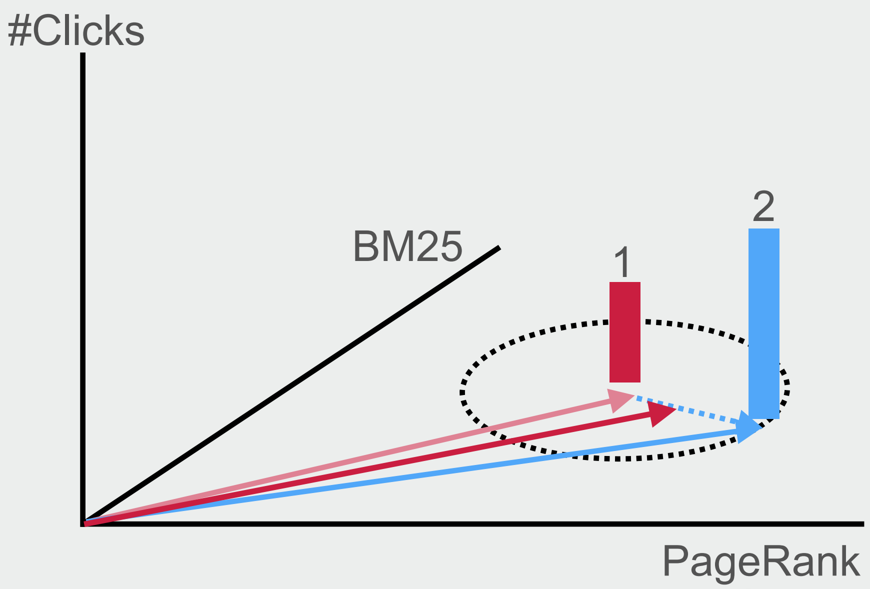

Weights for features

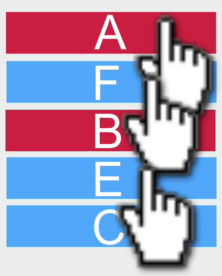

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

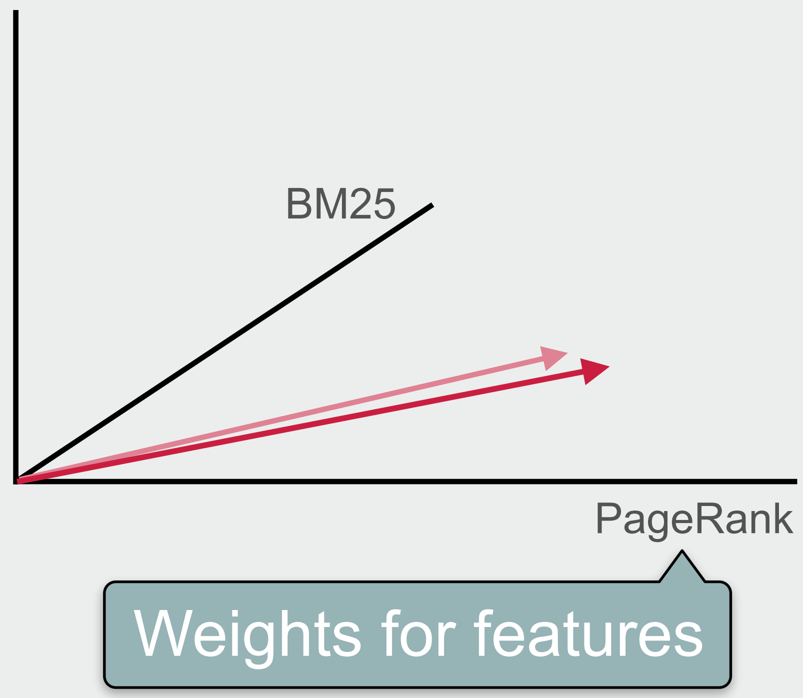

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Existing work

Dueling Bandit Gradient Descent (DBGD)

- DBGD updates after exploring only a **single** direction

Existing work

Dueling Bandit Gradient Descent (DBGD)

- DBGD updates after exploring only a **single** direction
- Exploring **multiple** directions would lead to a better ranker with less updates

Existing work

Dueling Bandit Gradient Descent (DBGD)

- DBGD updates after exploring only a **single** direction
- Exploring **multiple** directions would lead to a better ranker with less updates
- But would be **expensive** when interleaving was used

Existing work

Dueling Bandit Gradient Descent (DBGD)

- DBGD updates after exploring only a **single** direction
- Exploring **multiple** directions would lead to a better ranker with less updates
- But would be **expensive** when interleaving was used
  - Requires **pairwise** comparisons between all pairs of directions

Existing work

Dueling Bandit Gradient Descent (DBGD)

- DBGD updates after exploring only a **single** direction
- Exploring **multiple** directions would lead to a better ranker with less updates
- But would be **expensive** when interleaving was used
  - Requires **pairwise** comparisons between all pairs of directions
- **Multileaved comparisons can avoid this**

Our work

Multileave Gradient Descent (MGD)

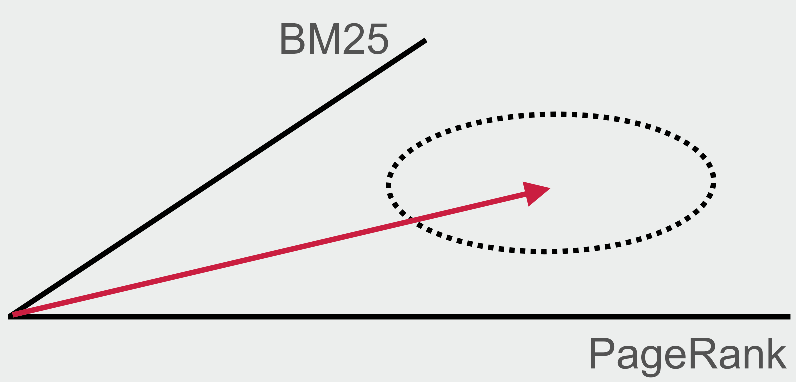

[Yue et al., 2009; Hofmann et al., 2011;
Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

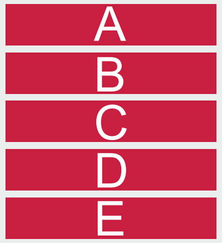

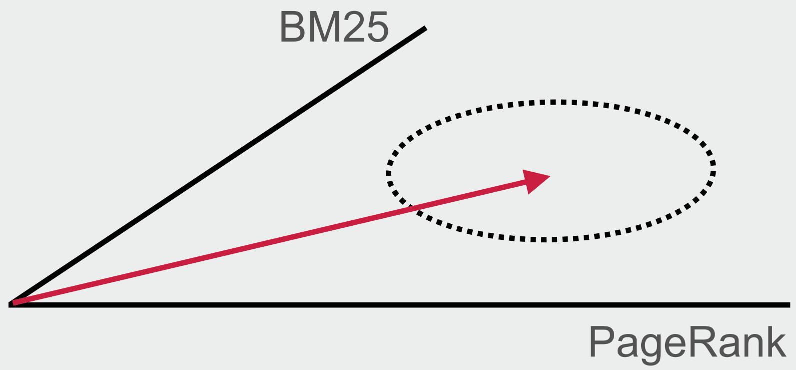

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

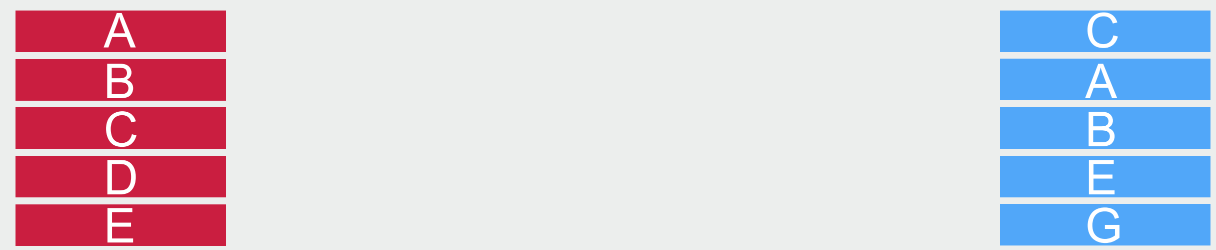

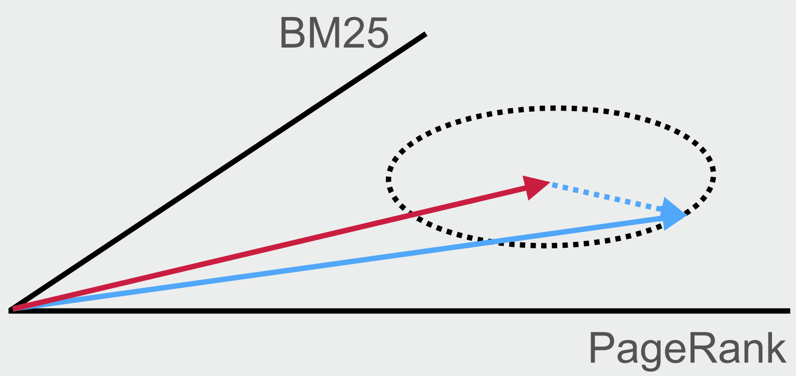

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

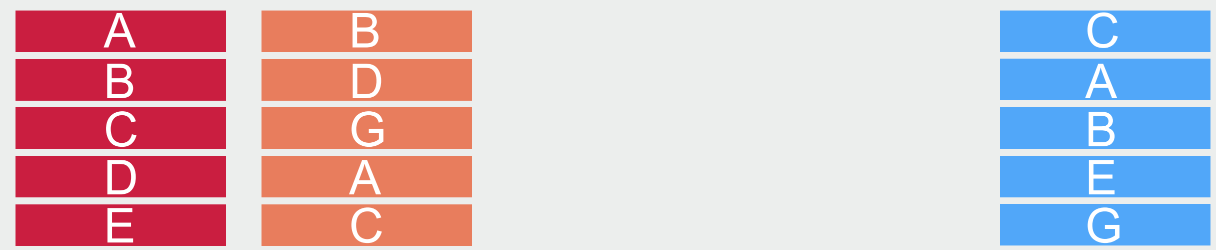

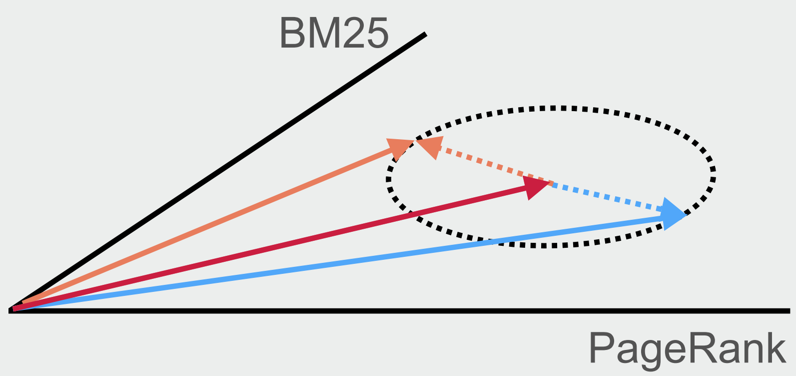

[Yue et al., 2009; Hofmann et al., 2011;
Radlinski et al., 2008; Schuth et al., 2014]

UNIVERSITY OF AMSTERDAM

Our work

Multileave Gradient Descent (MGD)

A	B	A	C
B	D	B	A
C	G	H	B
D	A	C	E
E	C	I	G

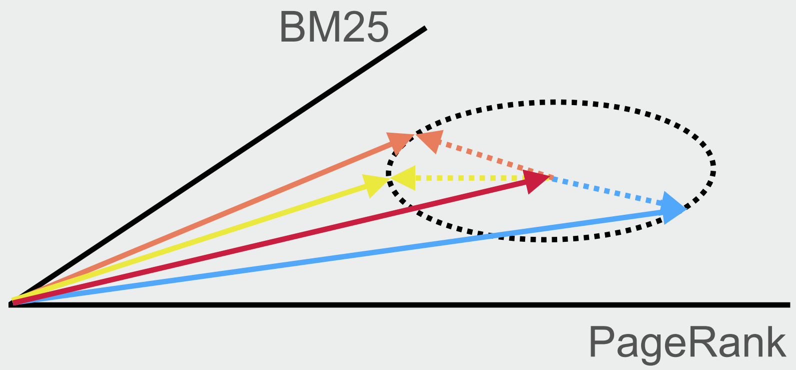

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

A	B	A	D	C
B	D	B	B	A
C	G	H	A	B
D	A	C	I	E
E	C	I	E	G

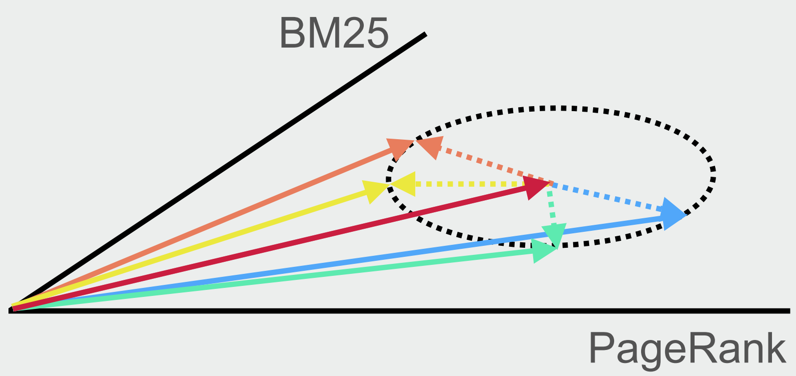

[Yue et al., 2009; Hofmann et al., 2011;
Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

A	B	A	D	C
B	D	B	B	A
C	G	H	A	B
D	A	C	I	E
E	C	I	E	G

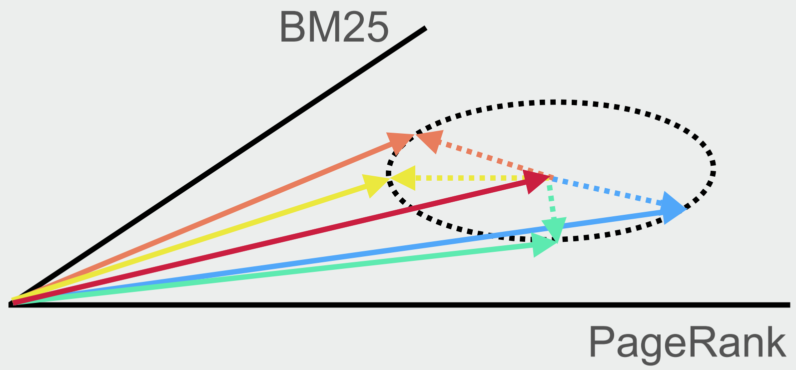

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

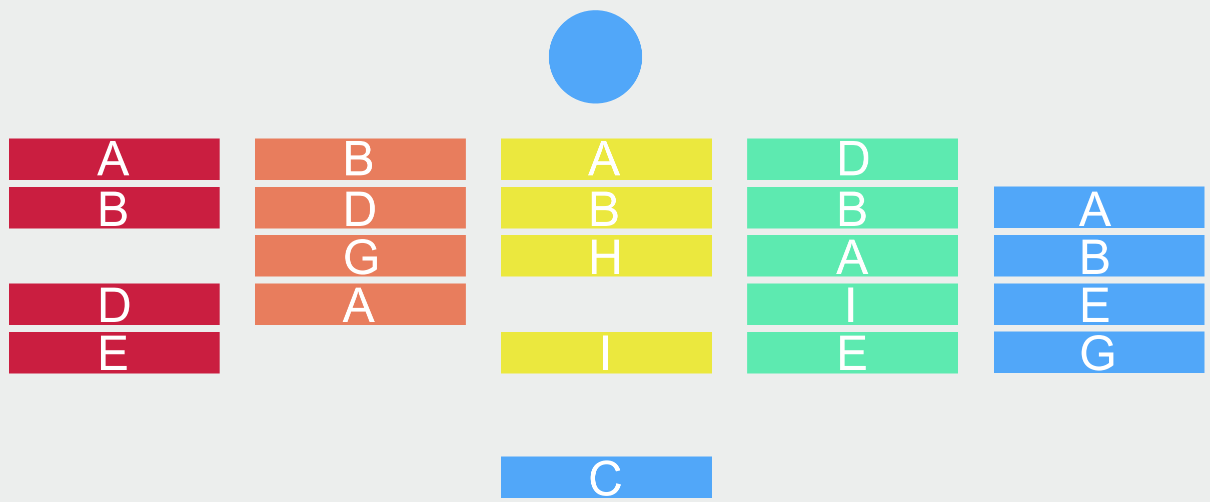

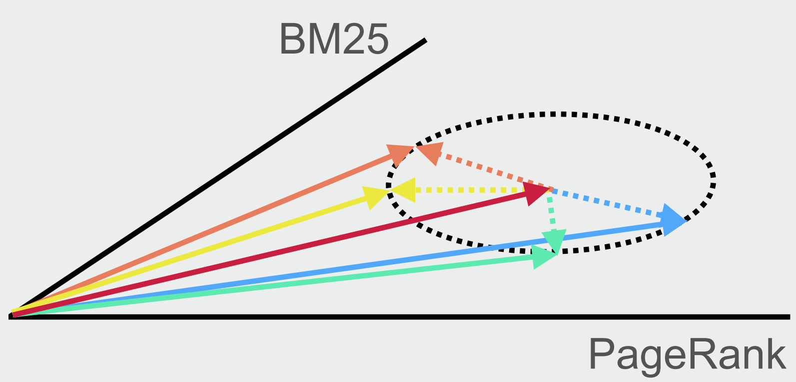

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

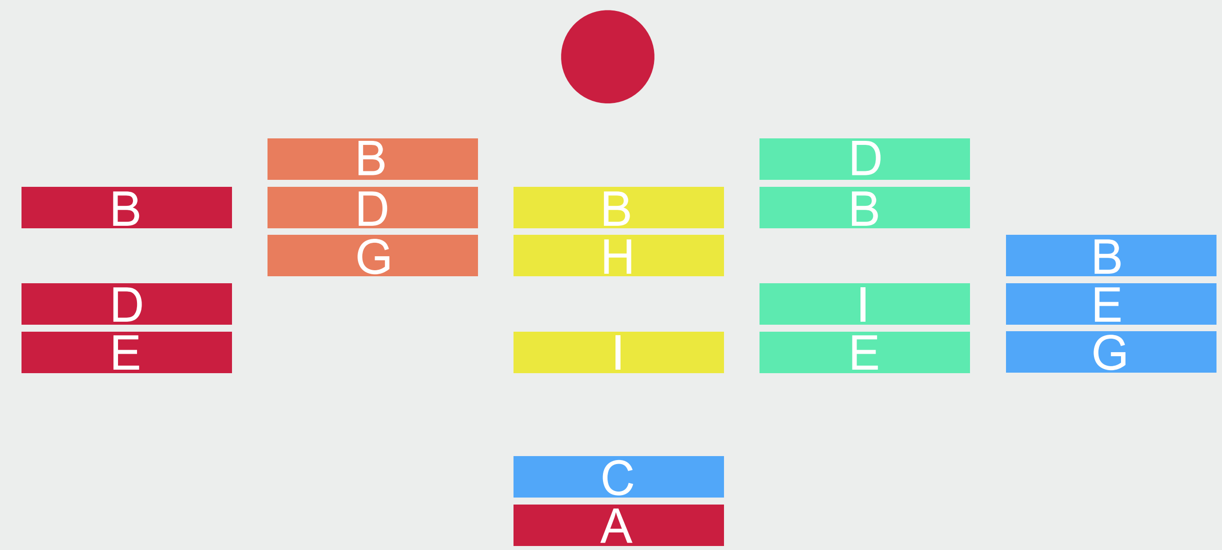

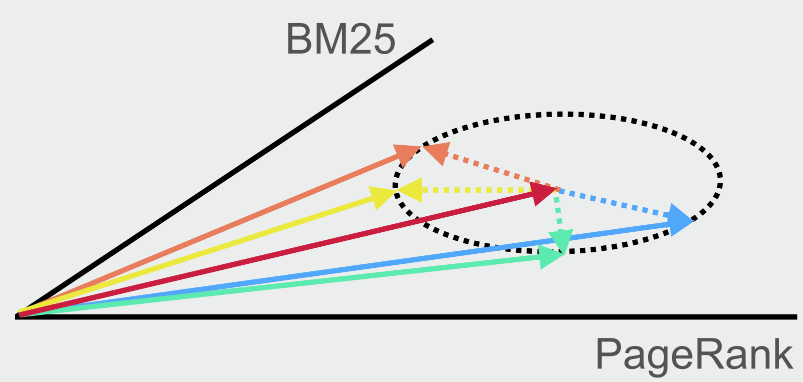

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

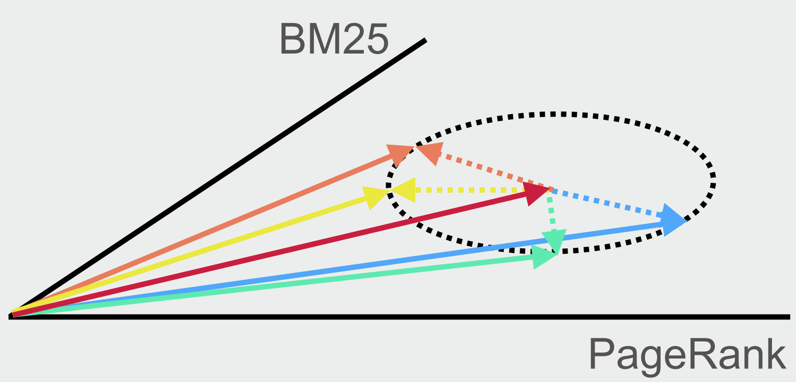

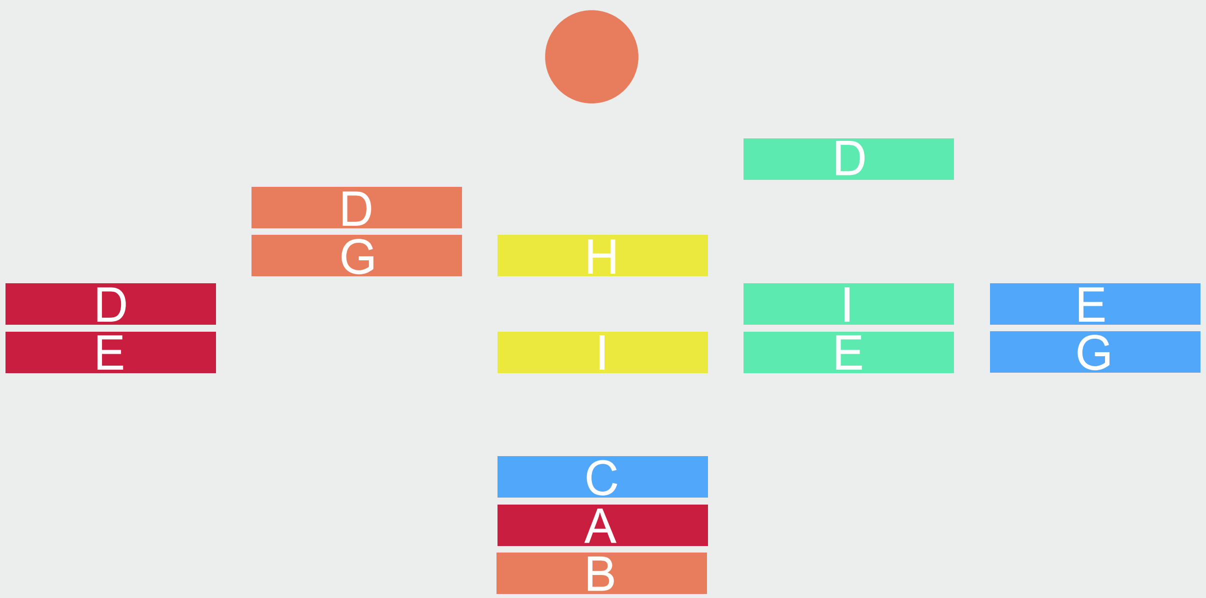

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

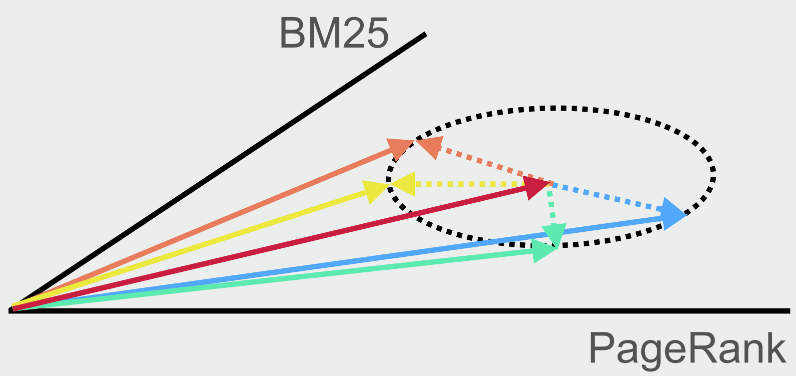

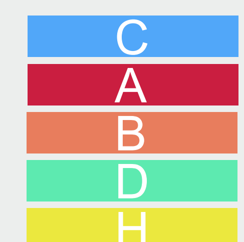

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

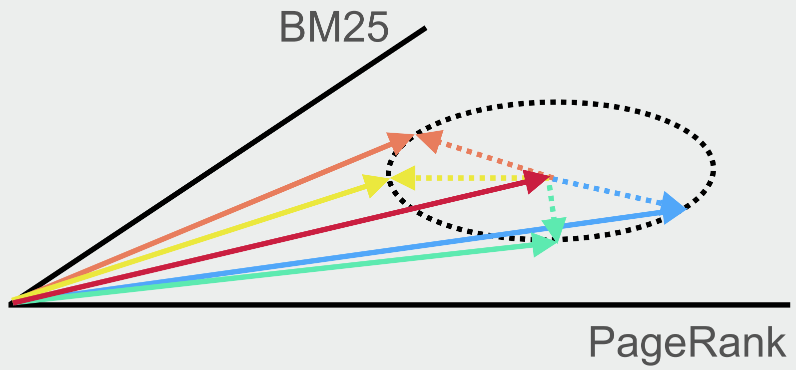

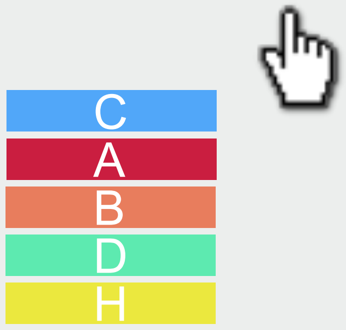

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

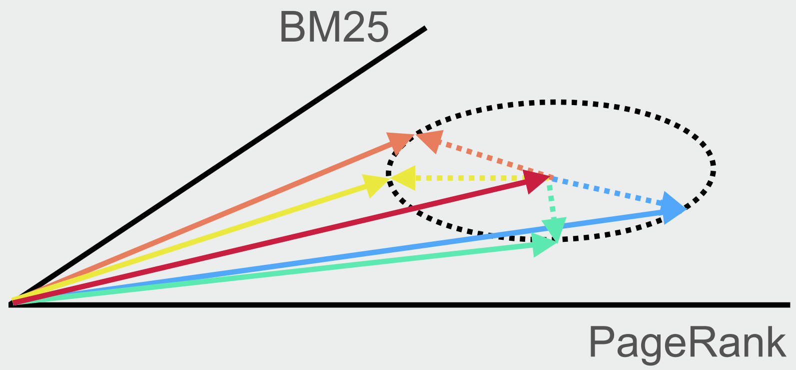

[Yue et al., 2009; Hofmann et al., 2011;
Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

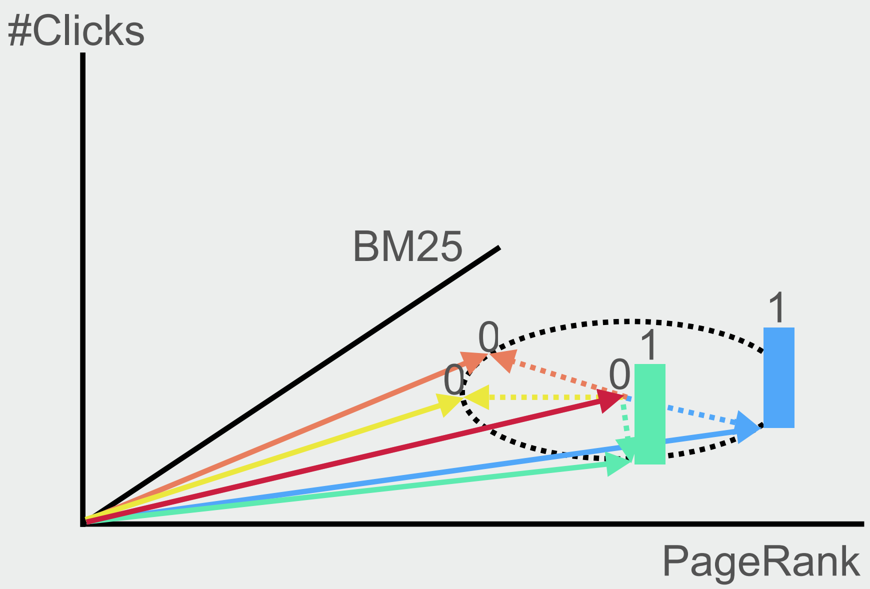

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Our work

Multileave Gradient Descent (MGD)

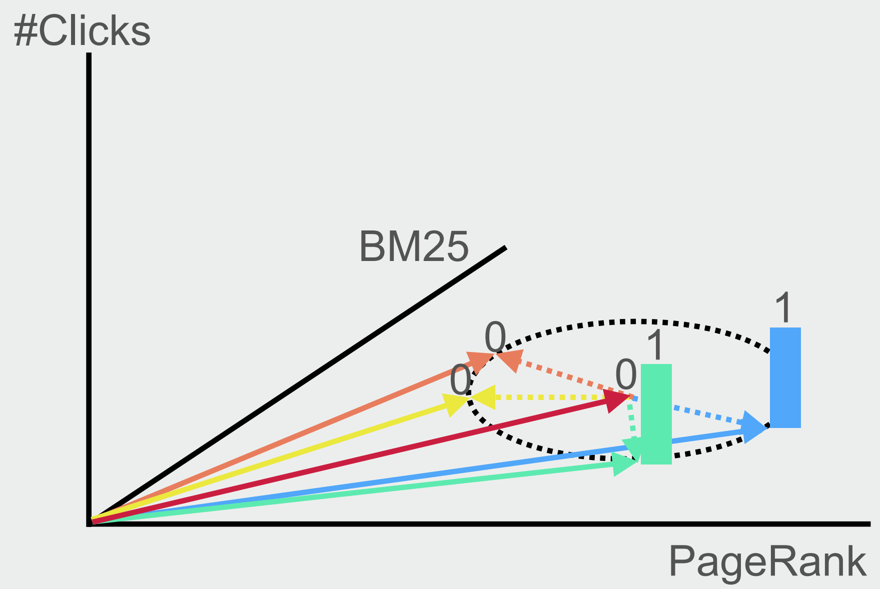

Our work

Multileave Gradient Descent (MGD)

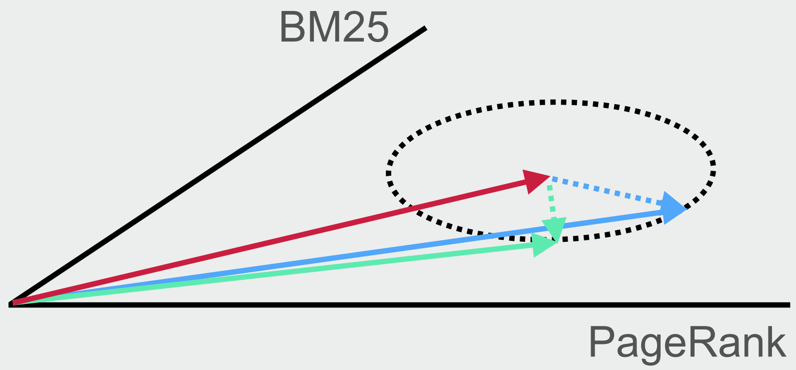

Our work

Multileave Gradient Descent (MGD)

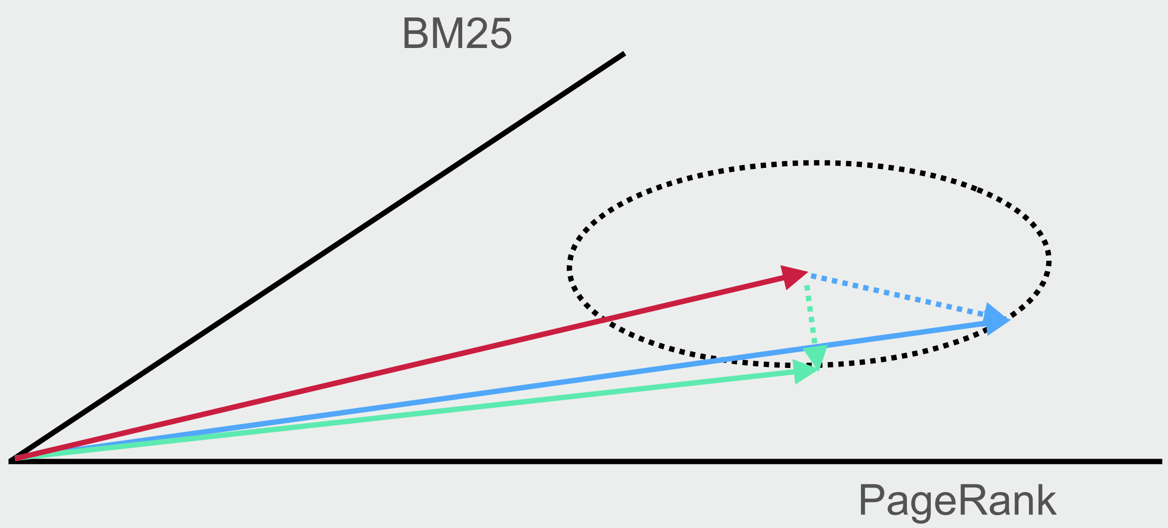

Our work

Multileave Gradient Descent (MGD)

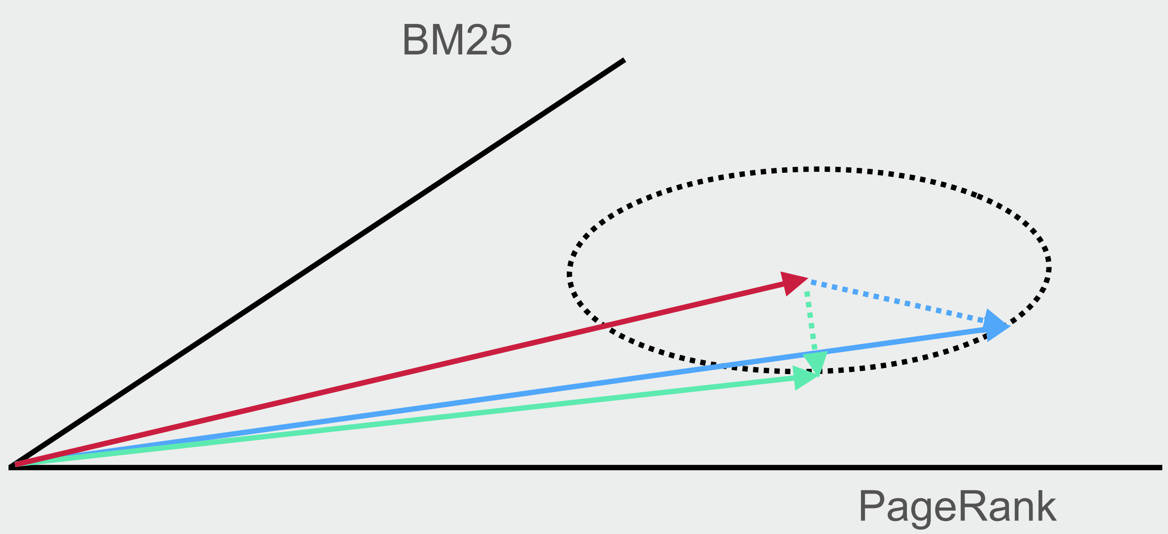

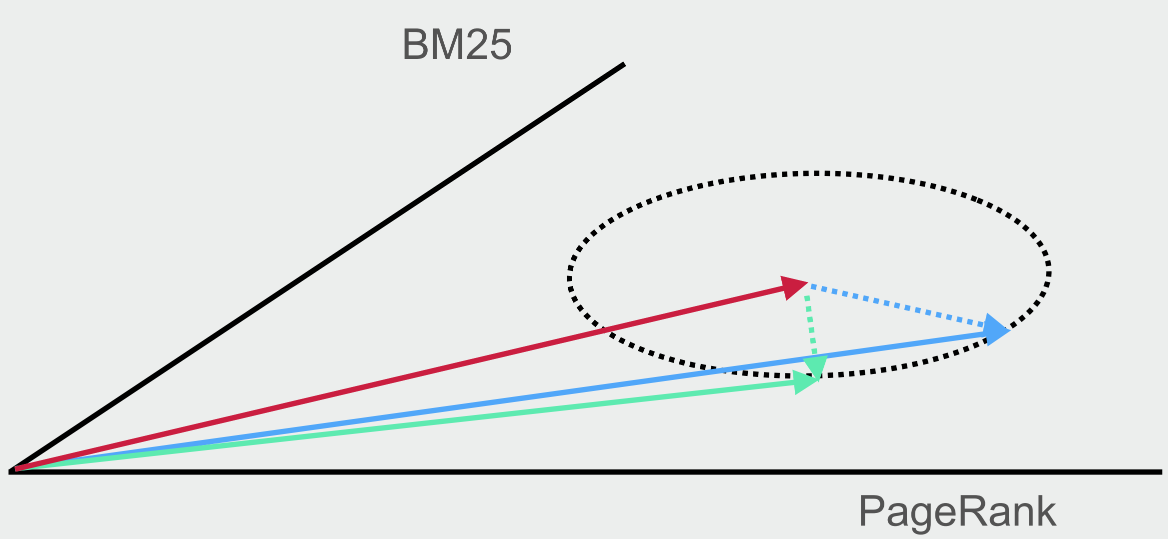

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

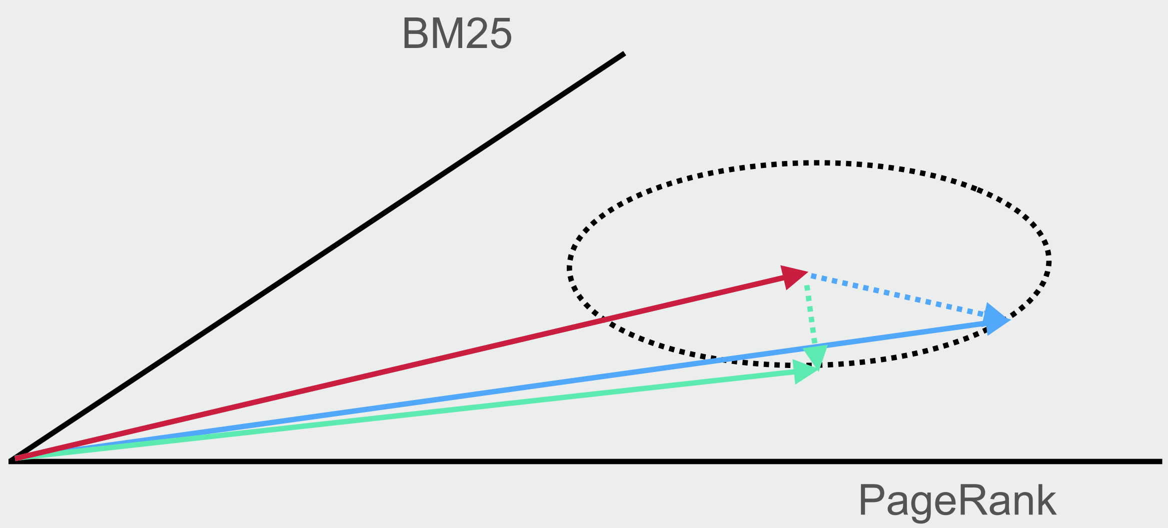

Mean winner (MGD-M)

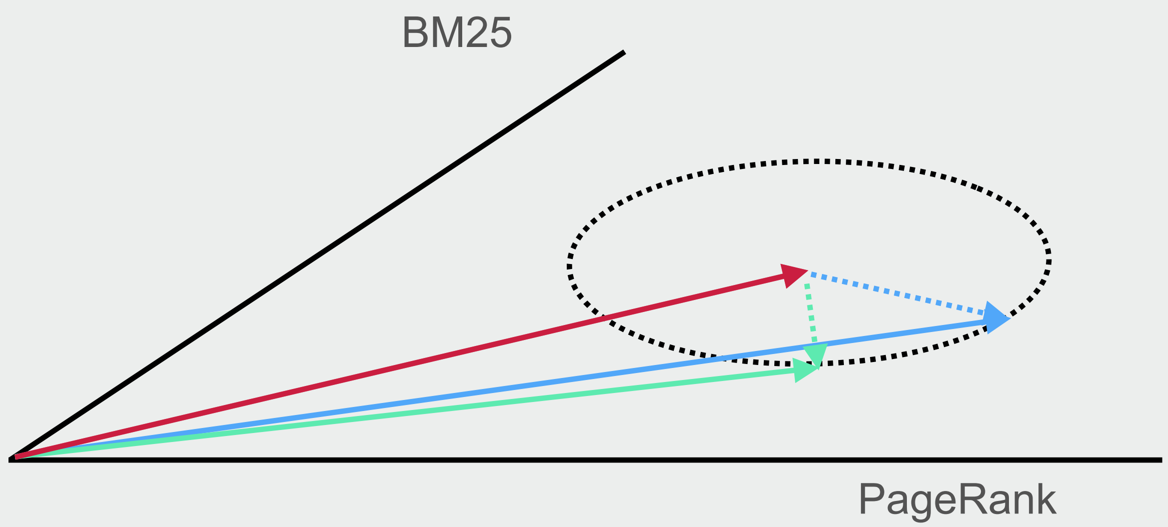

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

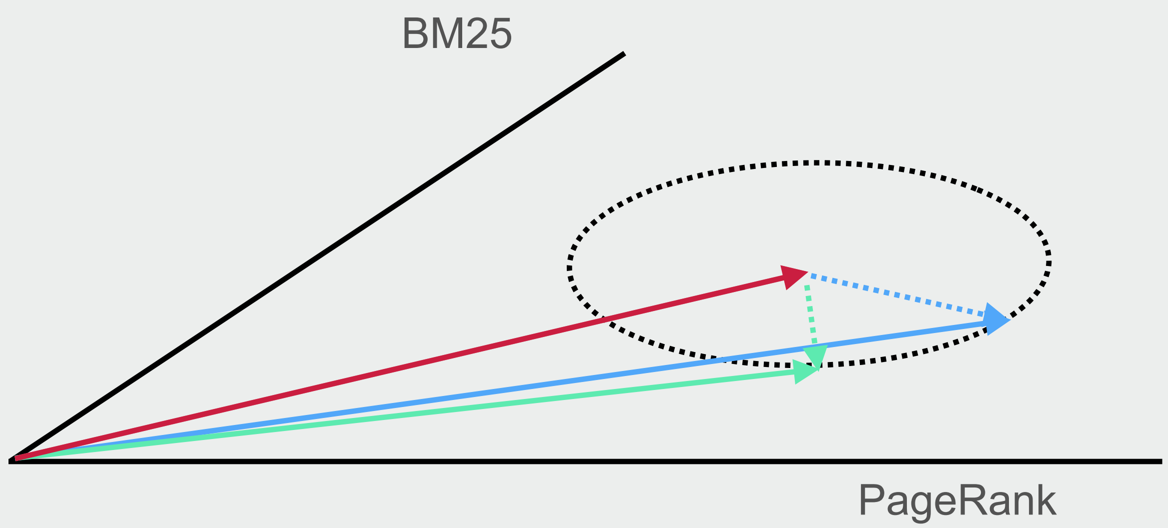

Mean winner (MGD-M)

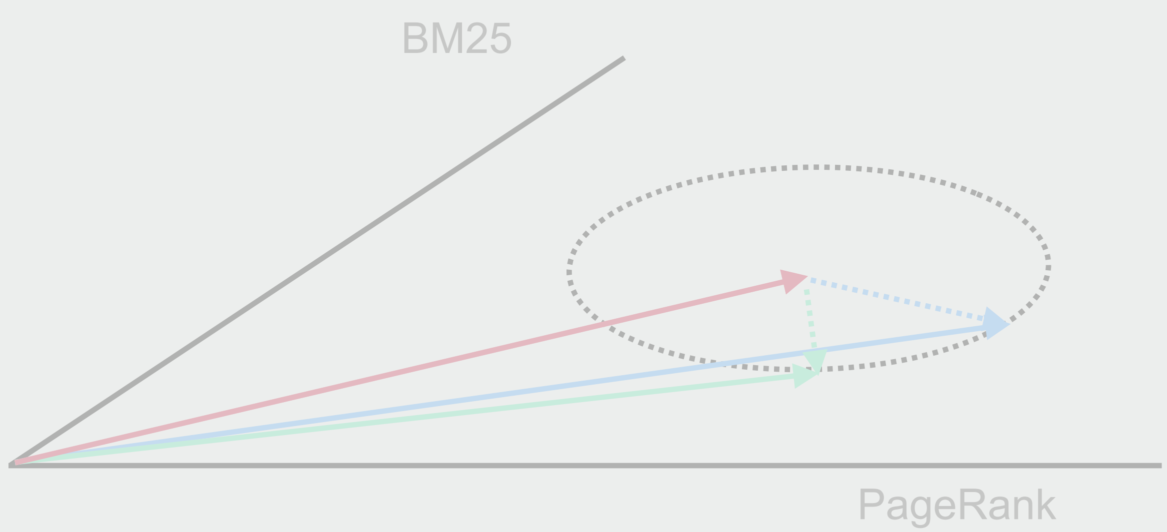

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

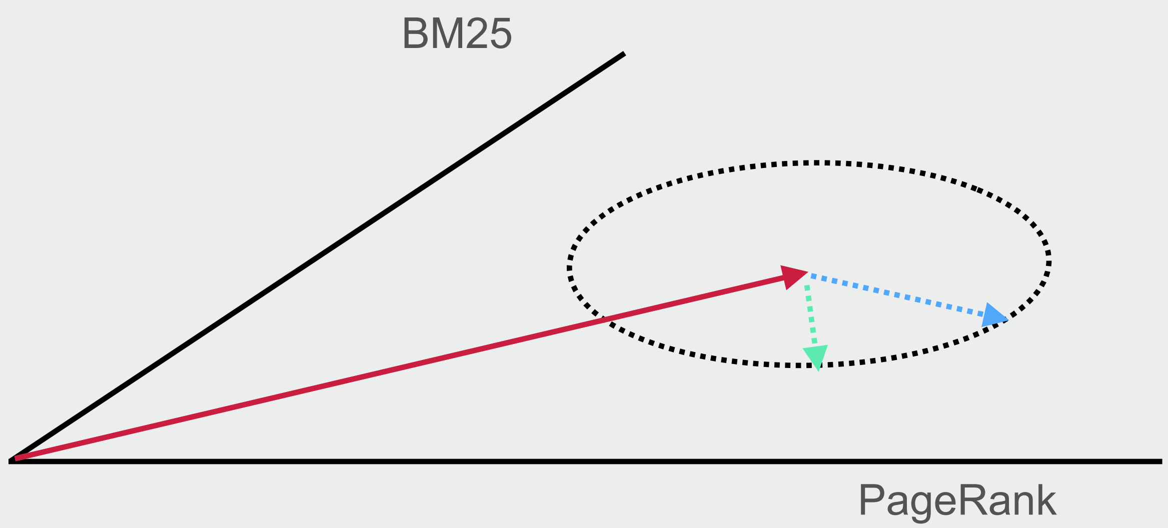

Mean winner (MGD-M)

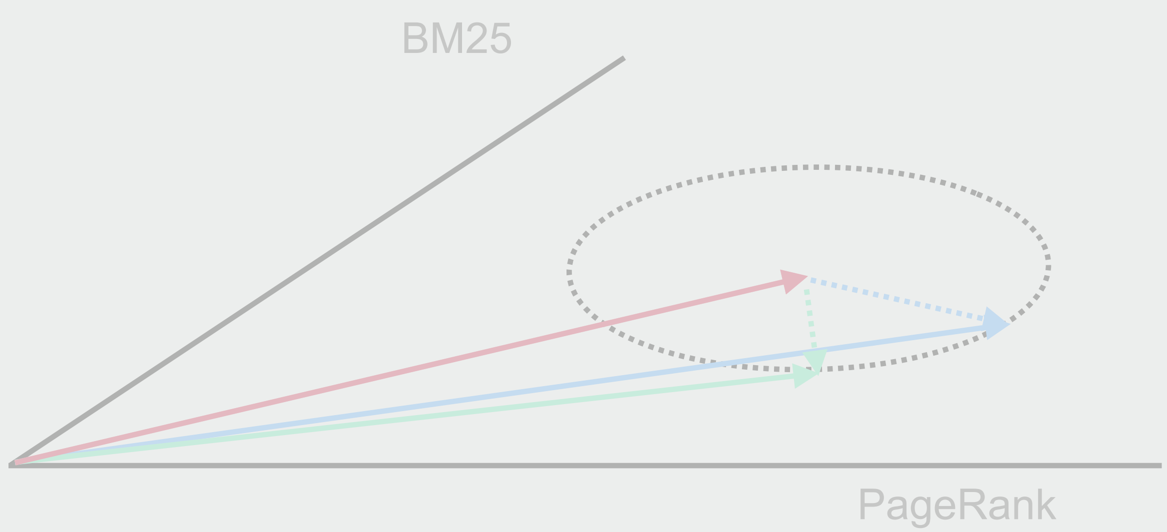

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

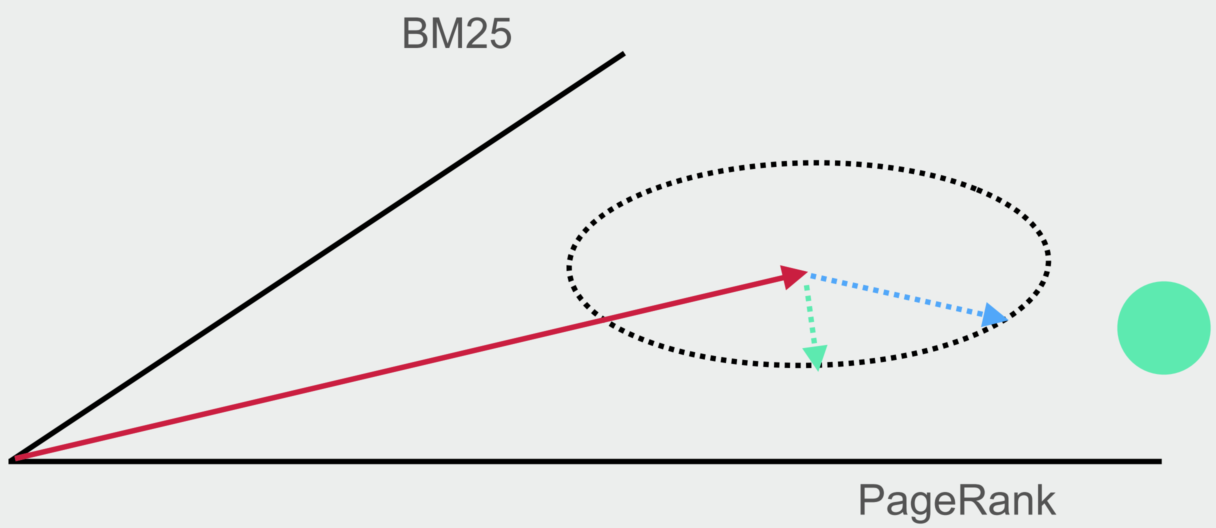

Mean winner (MGD-M)

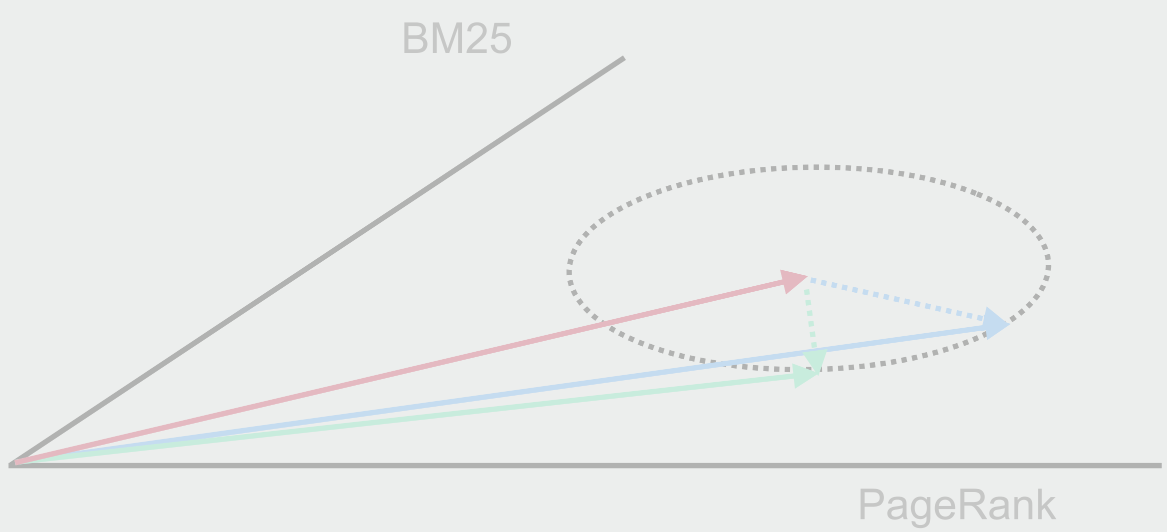

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

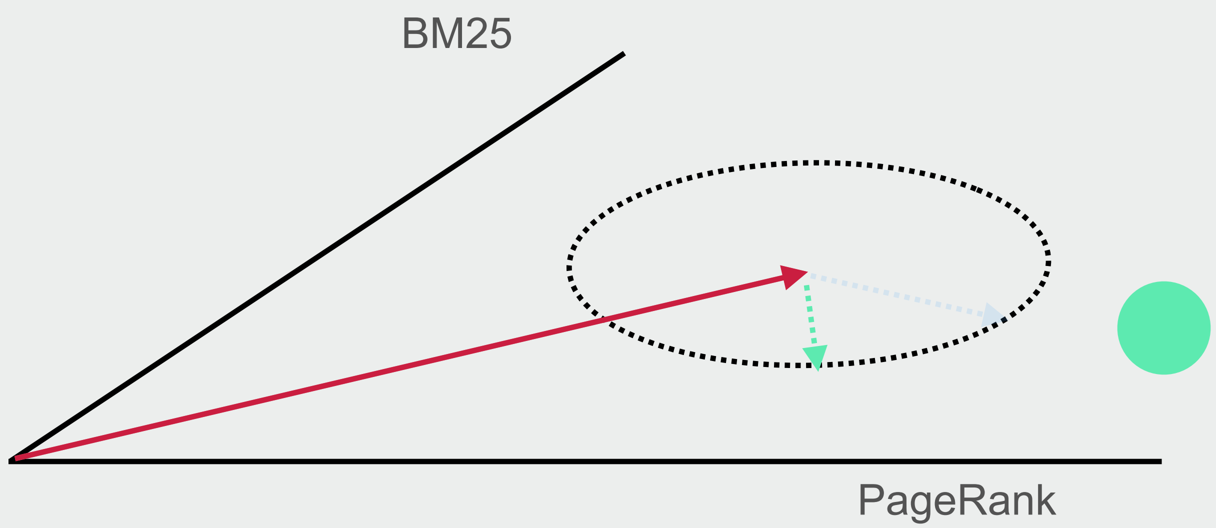

Mean winner (MGD-M)

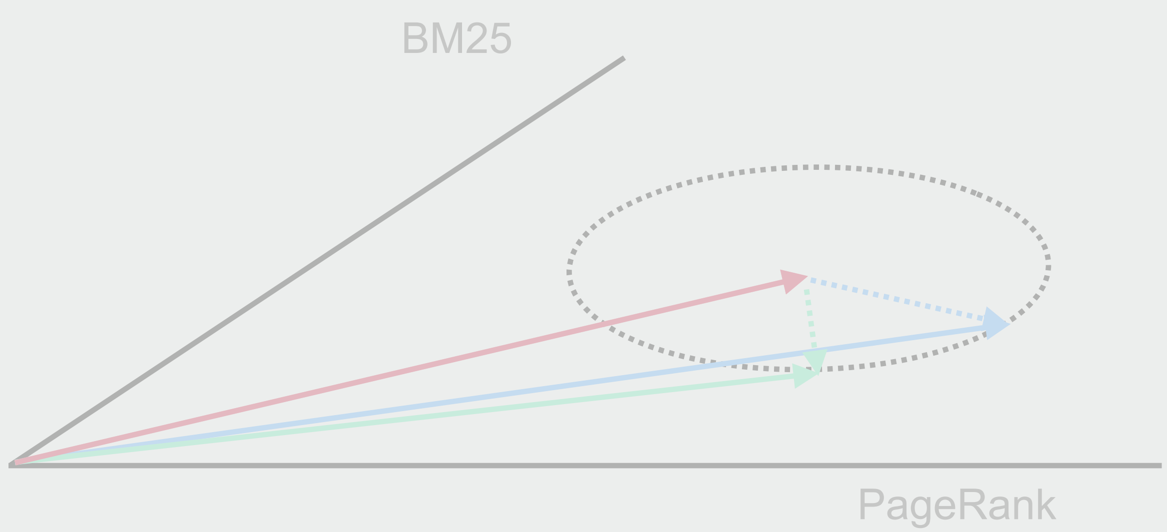

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

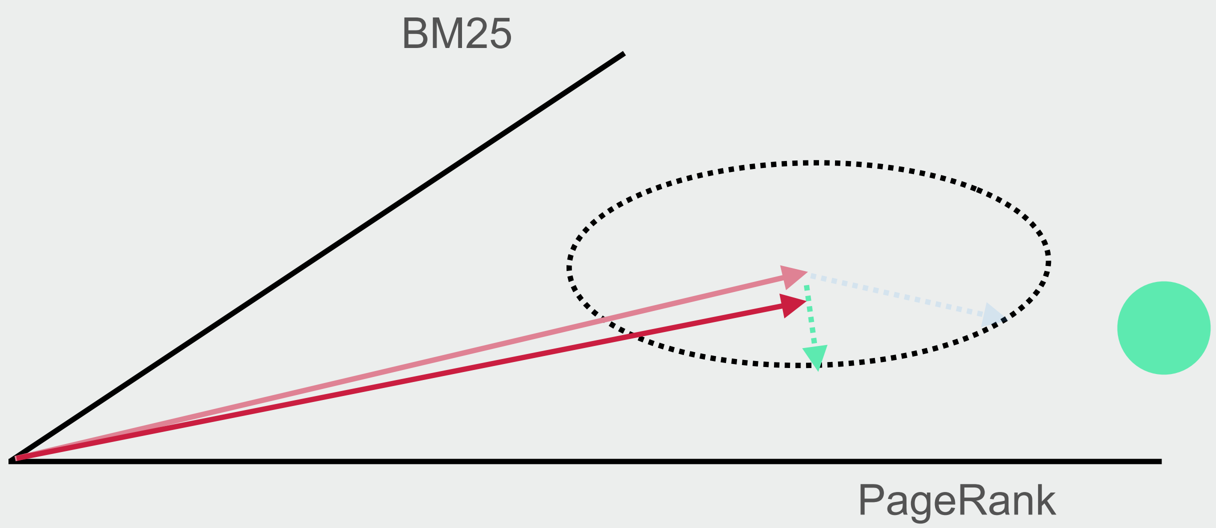

Mean winner (MGD-M)

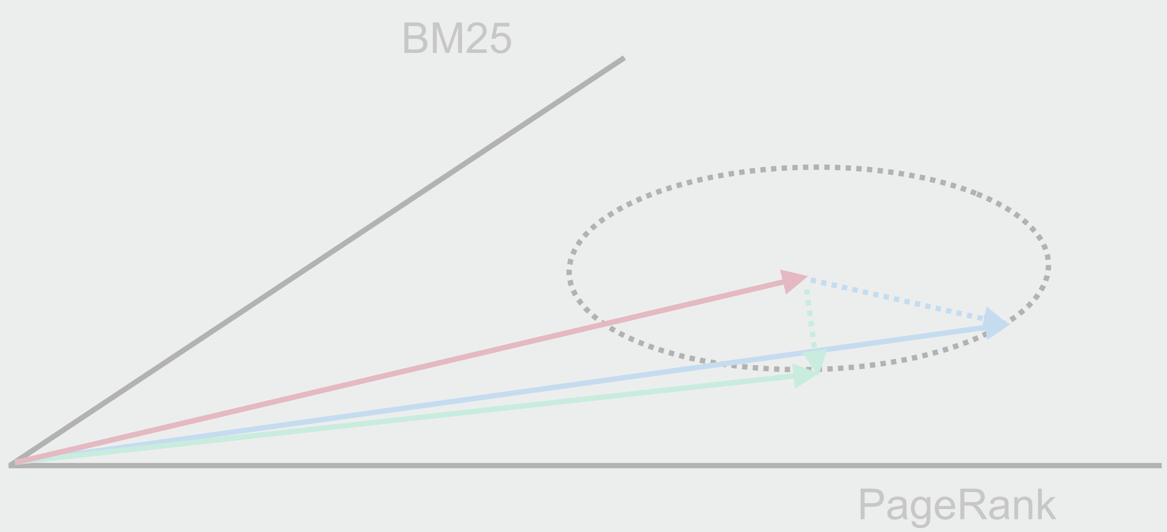

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

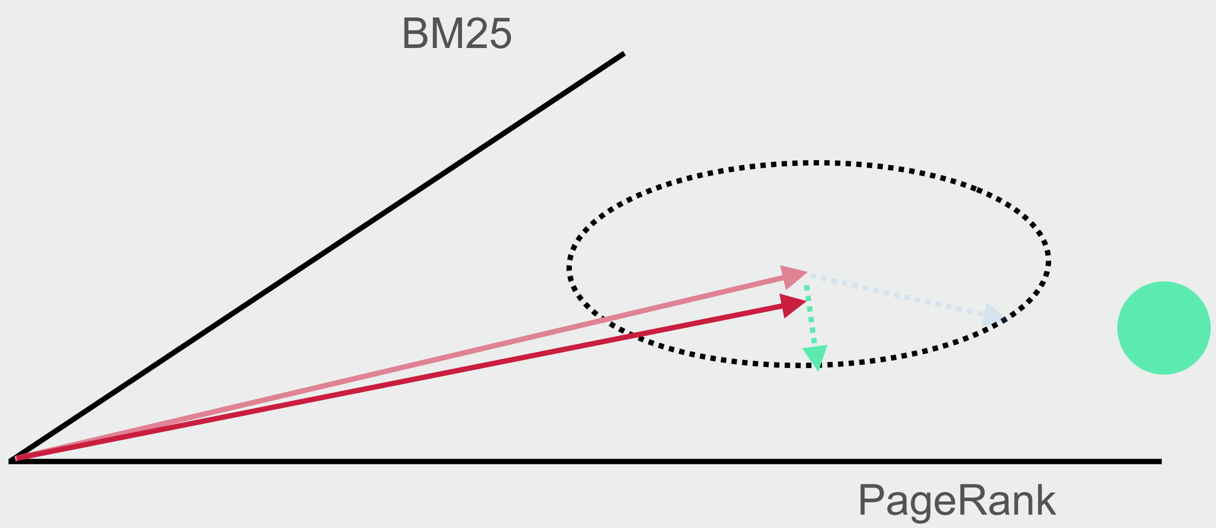

Mean winner (MGD-M)

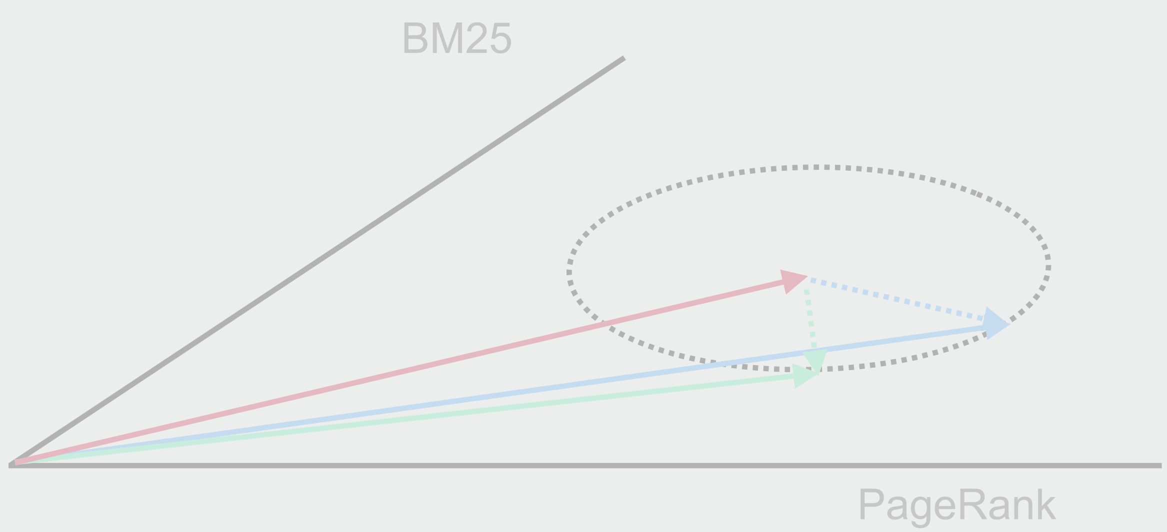

- Pick one of the winners
- Update with an alpha step

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

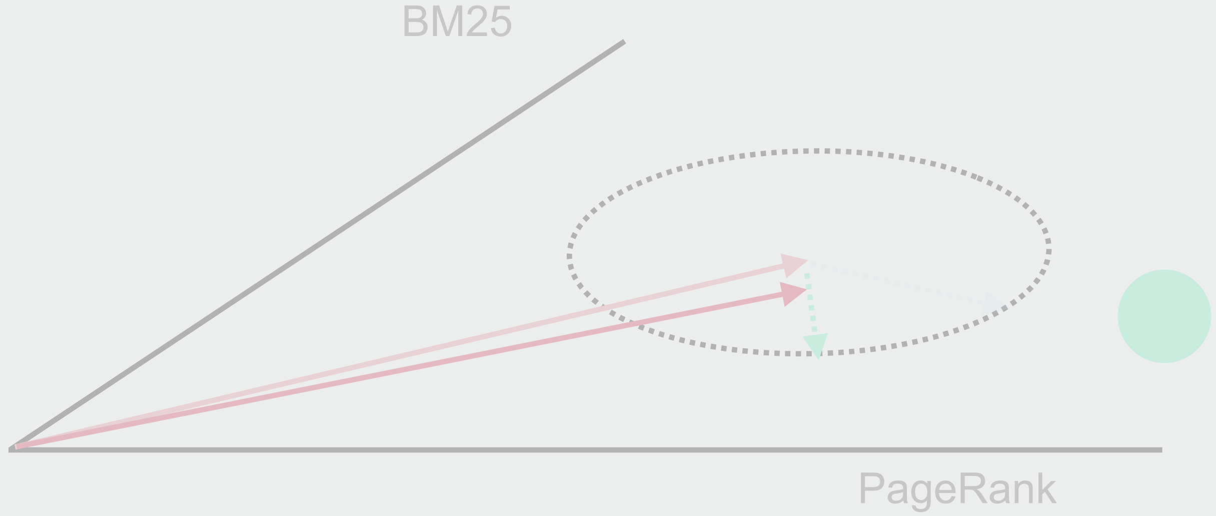

- Pick one of the winners
- Update with an alpha step

Mean winner (MGD-M)

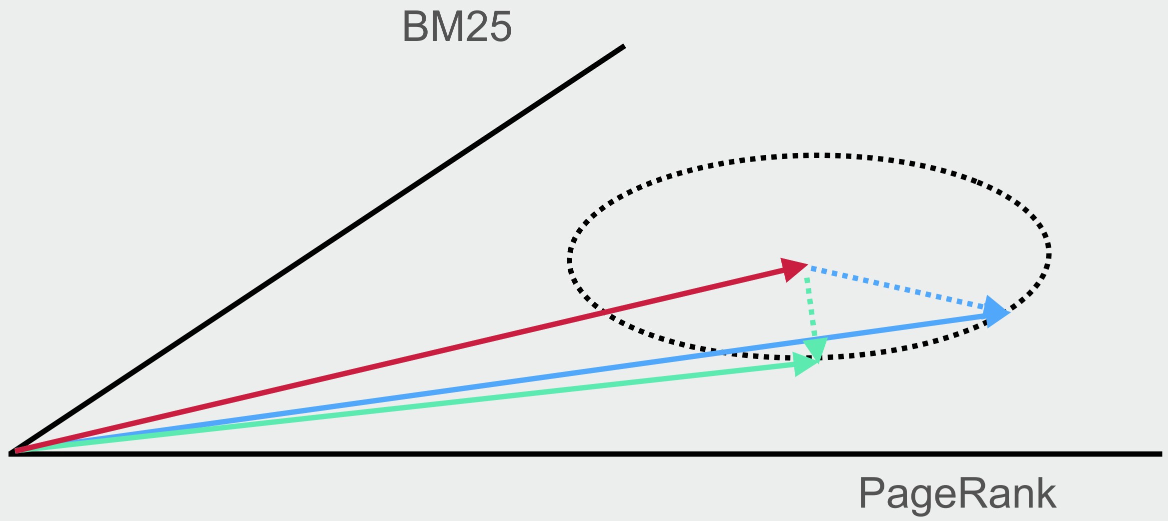

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

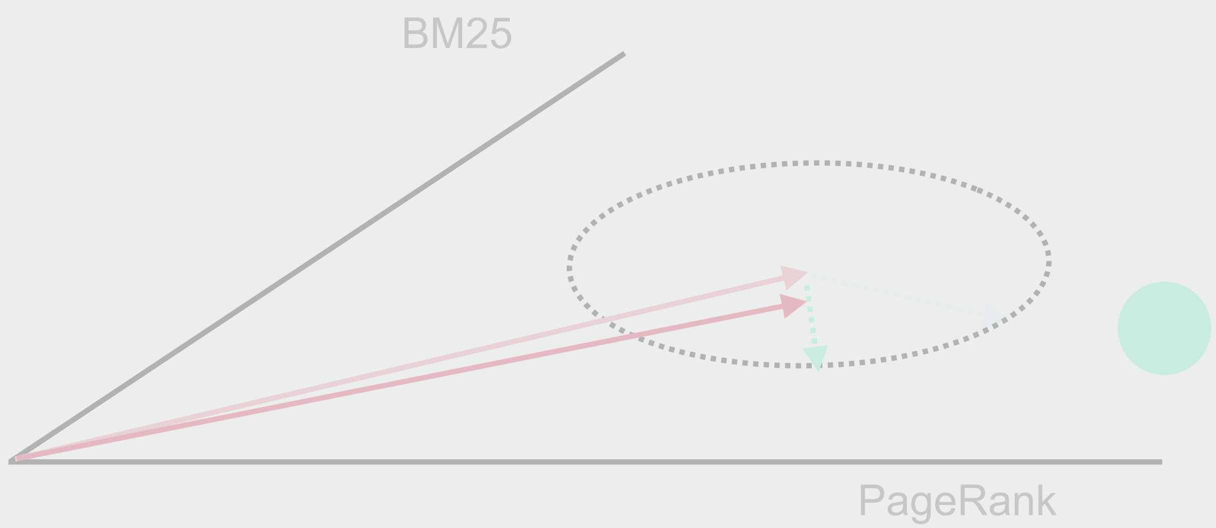

Mean winner (MGD-M)

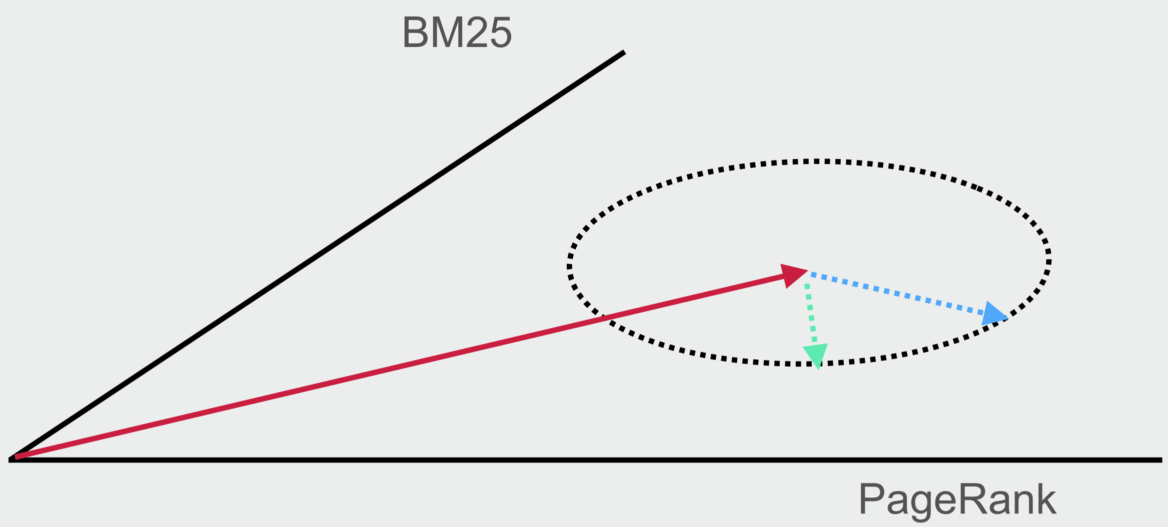

- Pick one of the winners
- Update with an alpha step

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

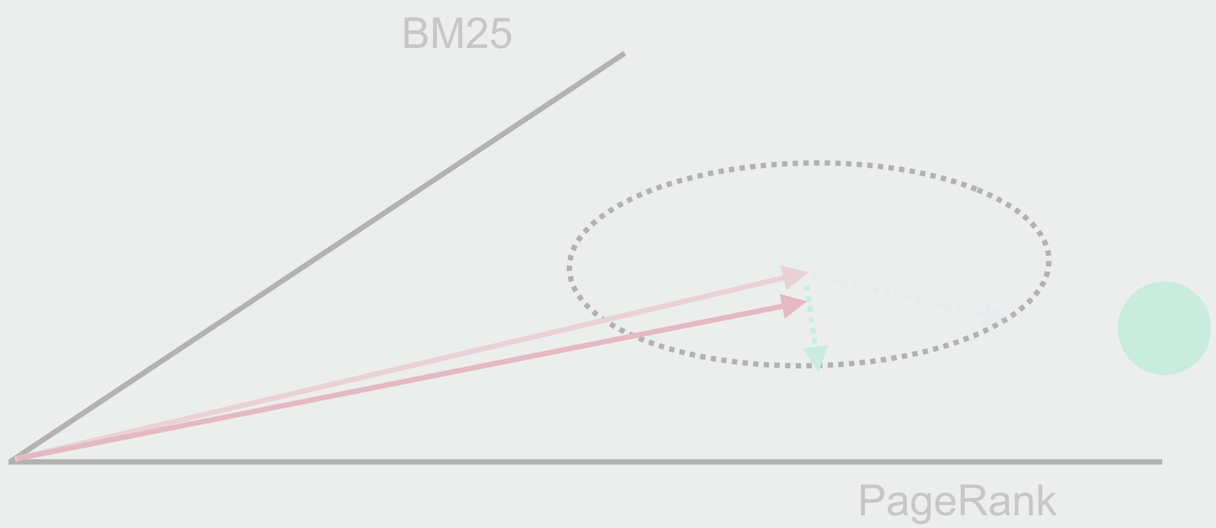

- Pick one of the winners
- Update with an alpha step

Mean winner (MGD-M)

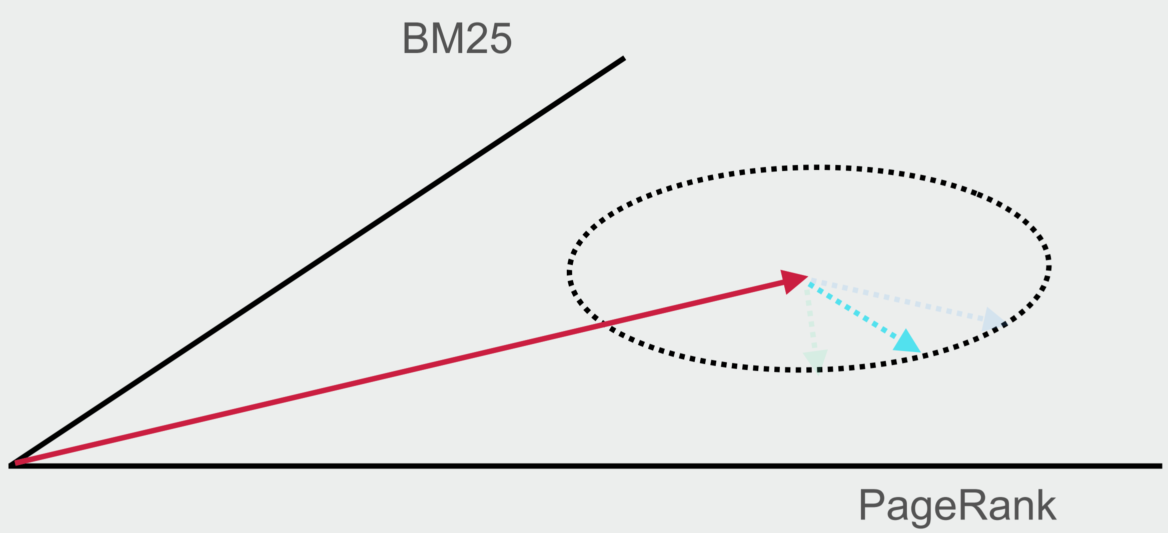

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

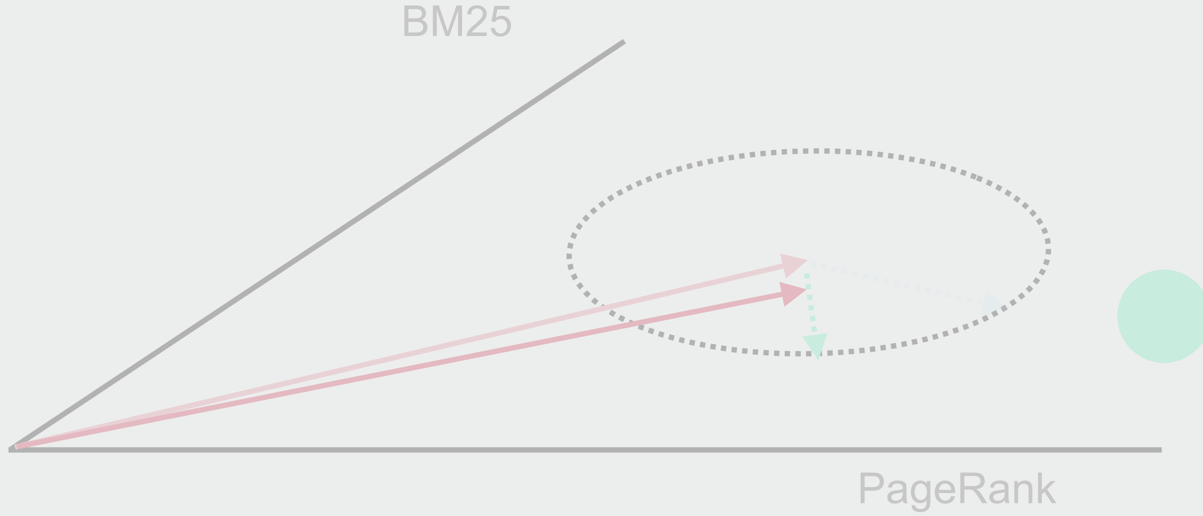

Mean winner (MGD-M)

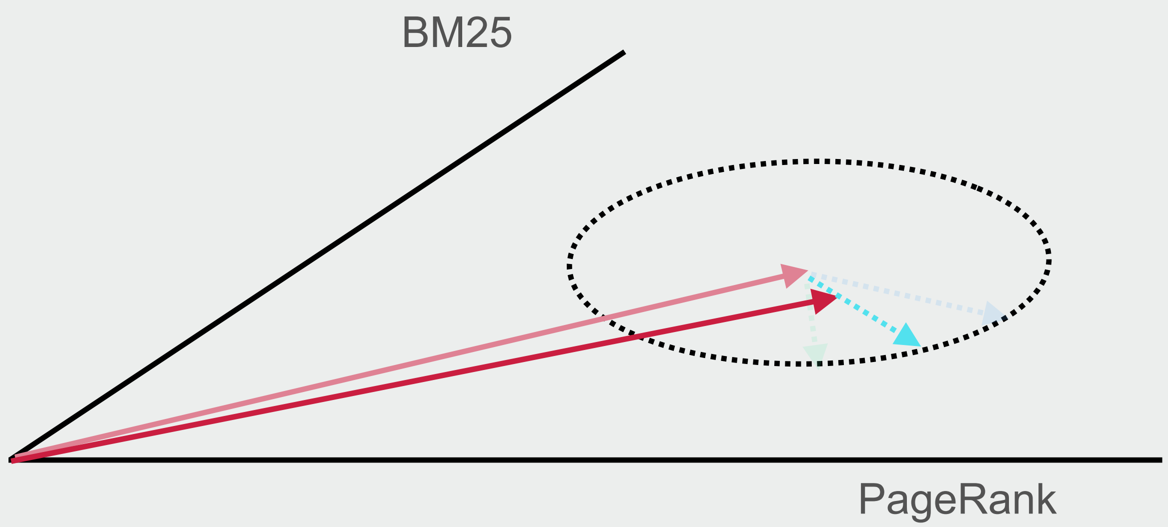

- Pick one of the winners
- Update with an alpha step

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

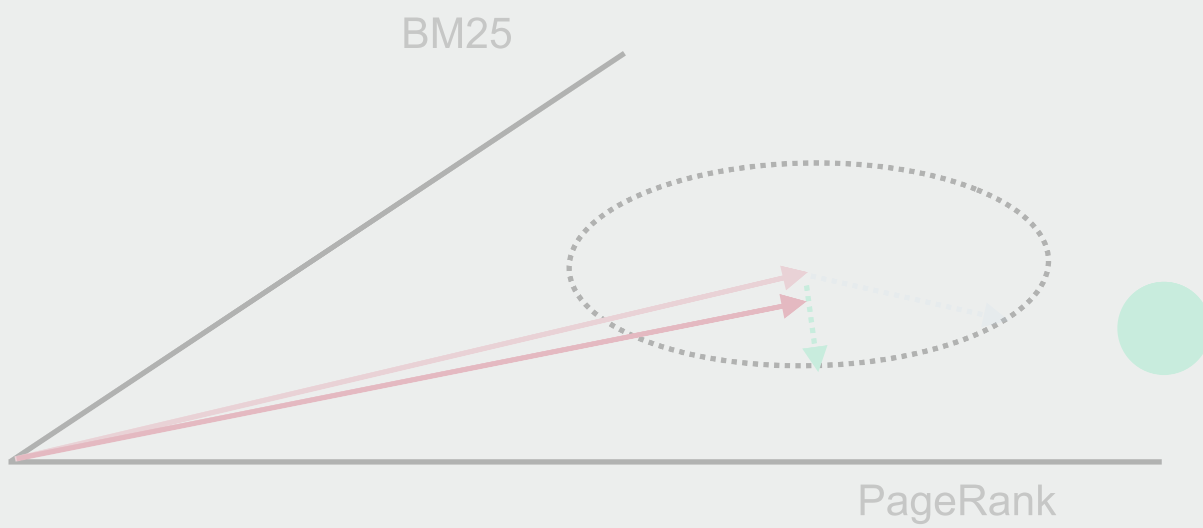

- Pick one of the winners
- Update with an alpha step

Mean winner (MGD-M)

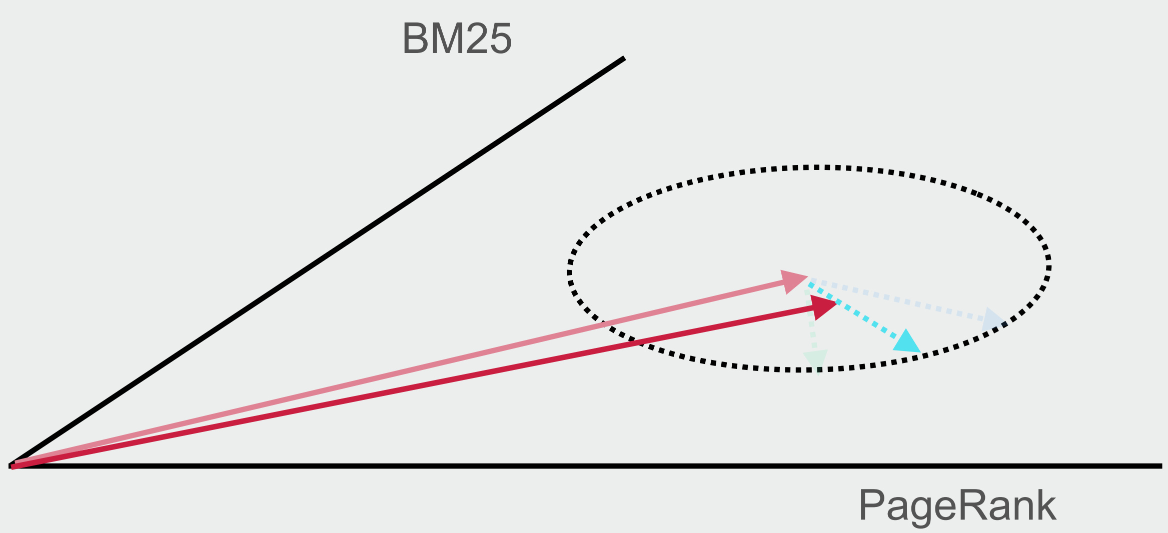

- Compute the mean of the winners
- Update with an alpha step

Our work

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

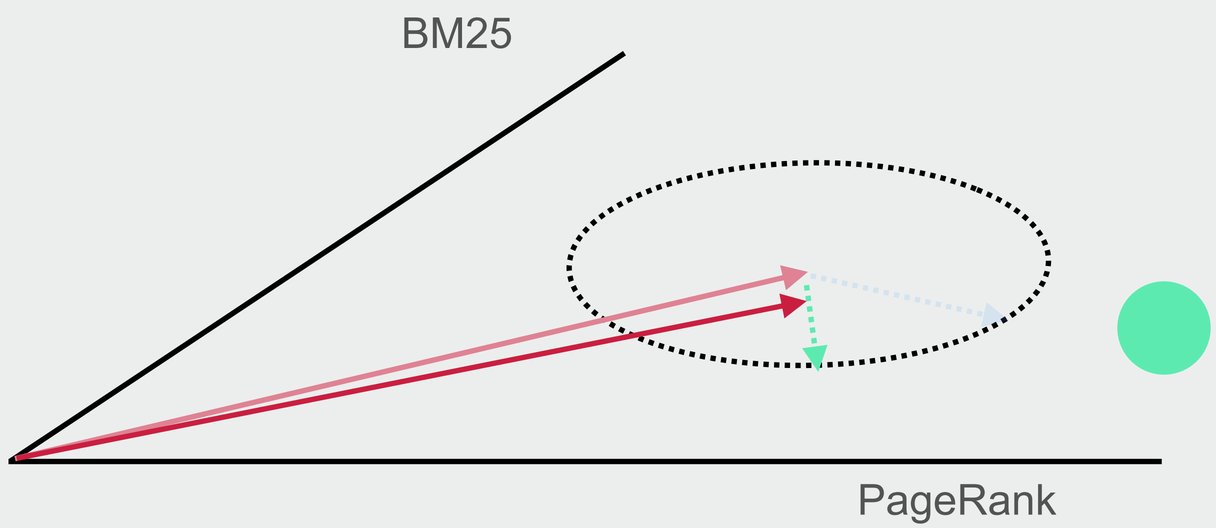

- Pick one of the winners
- Update with an alpha step

Mean winner (MGD-M)

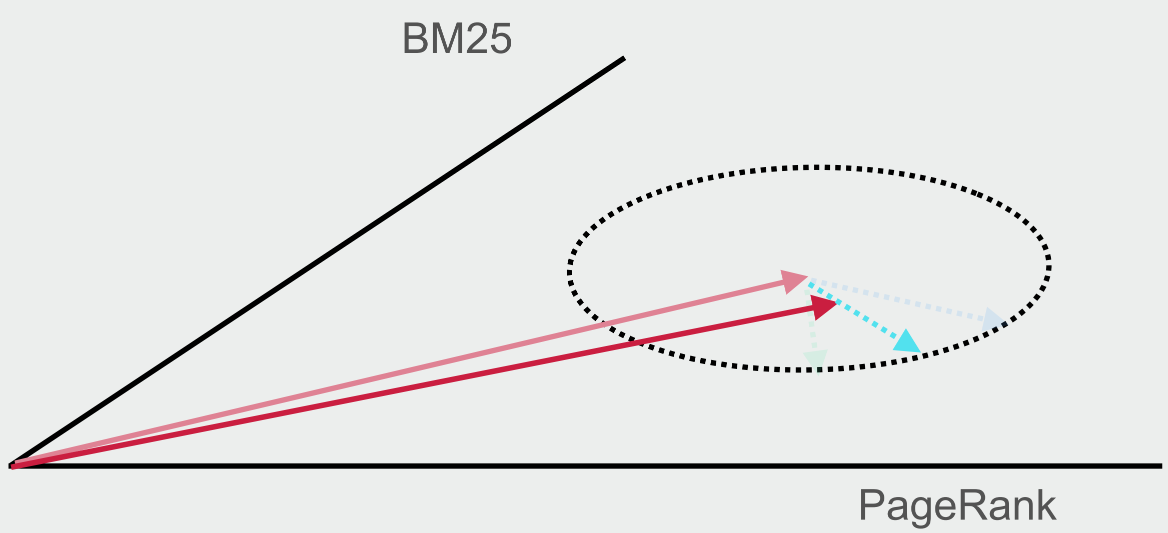

- Compute the mean of the winners
- Update with an alpha step

Experiments

- Queries
  - Sampled from a L2R dataset

Experiments

- Queries
  - Sampled from a L2R dataset
- Clicks
  - Generated by a *cascade click model*

Experiments

- Queries
  - Sampled from a L2R dataset
- Clicks
  - Generated by a *cascade click model*
- NDCG
  - Measured on held out data

Results

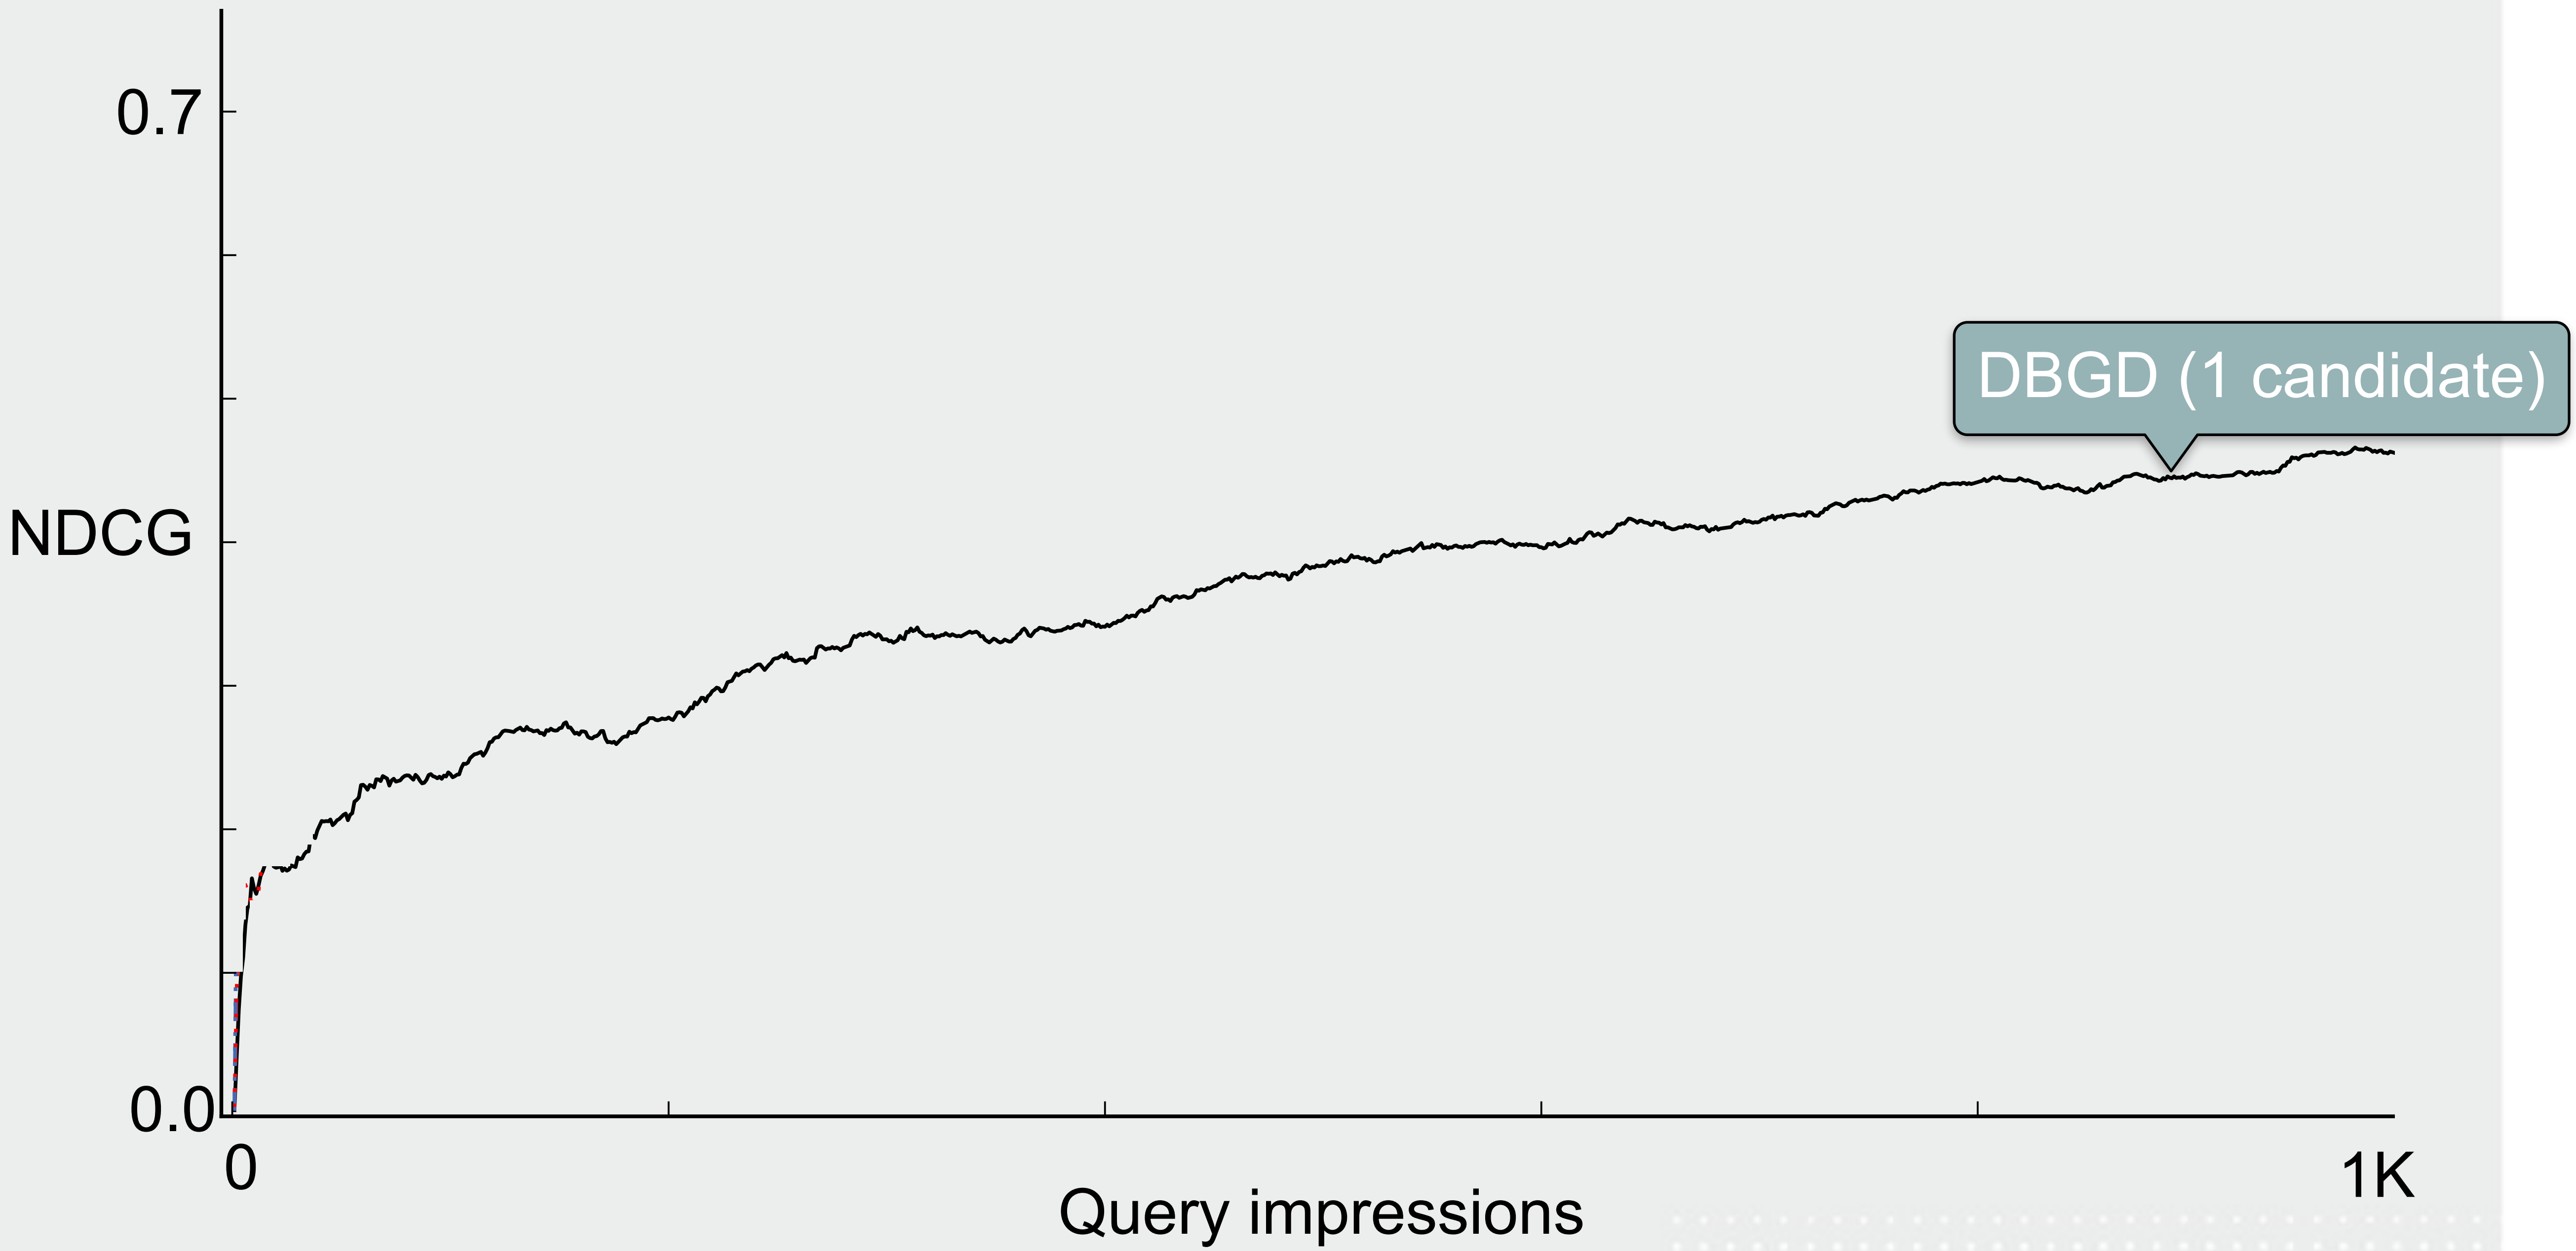

Results

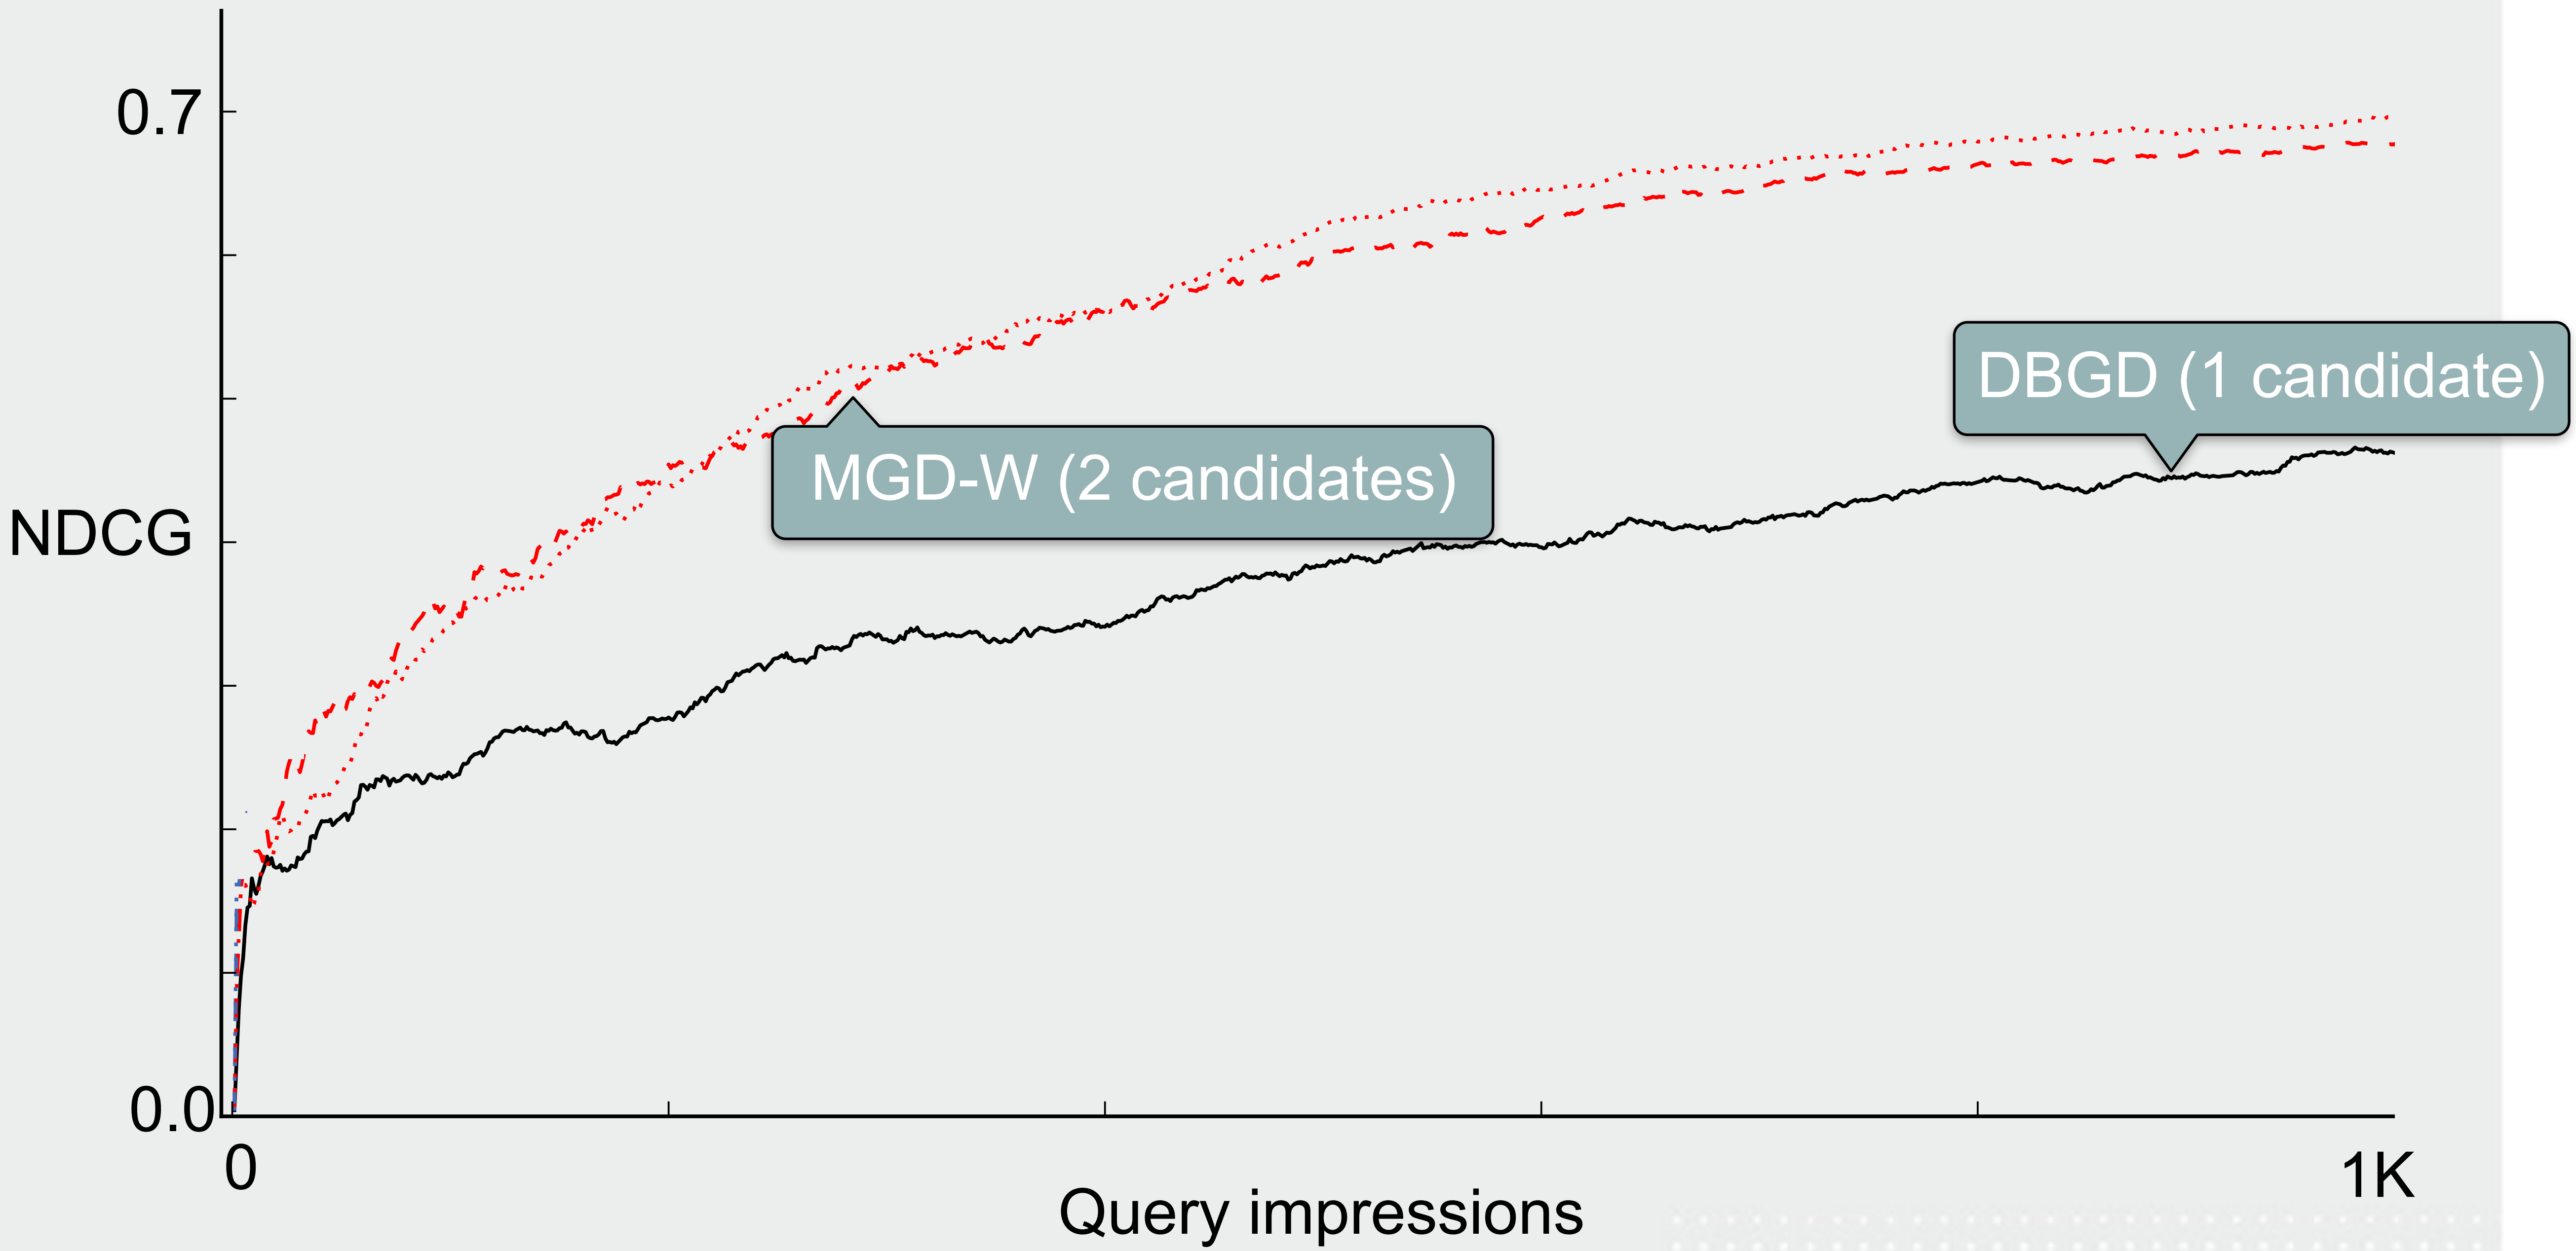
0.7
NDCG
0.0
0
1K
Query impressions
MGD-W (2 candidates)
DBGD (1 candidate)

Results

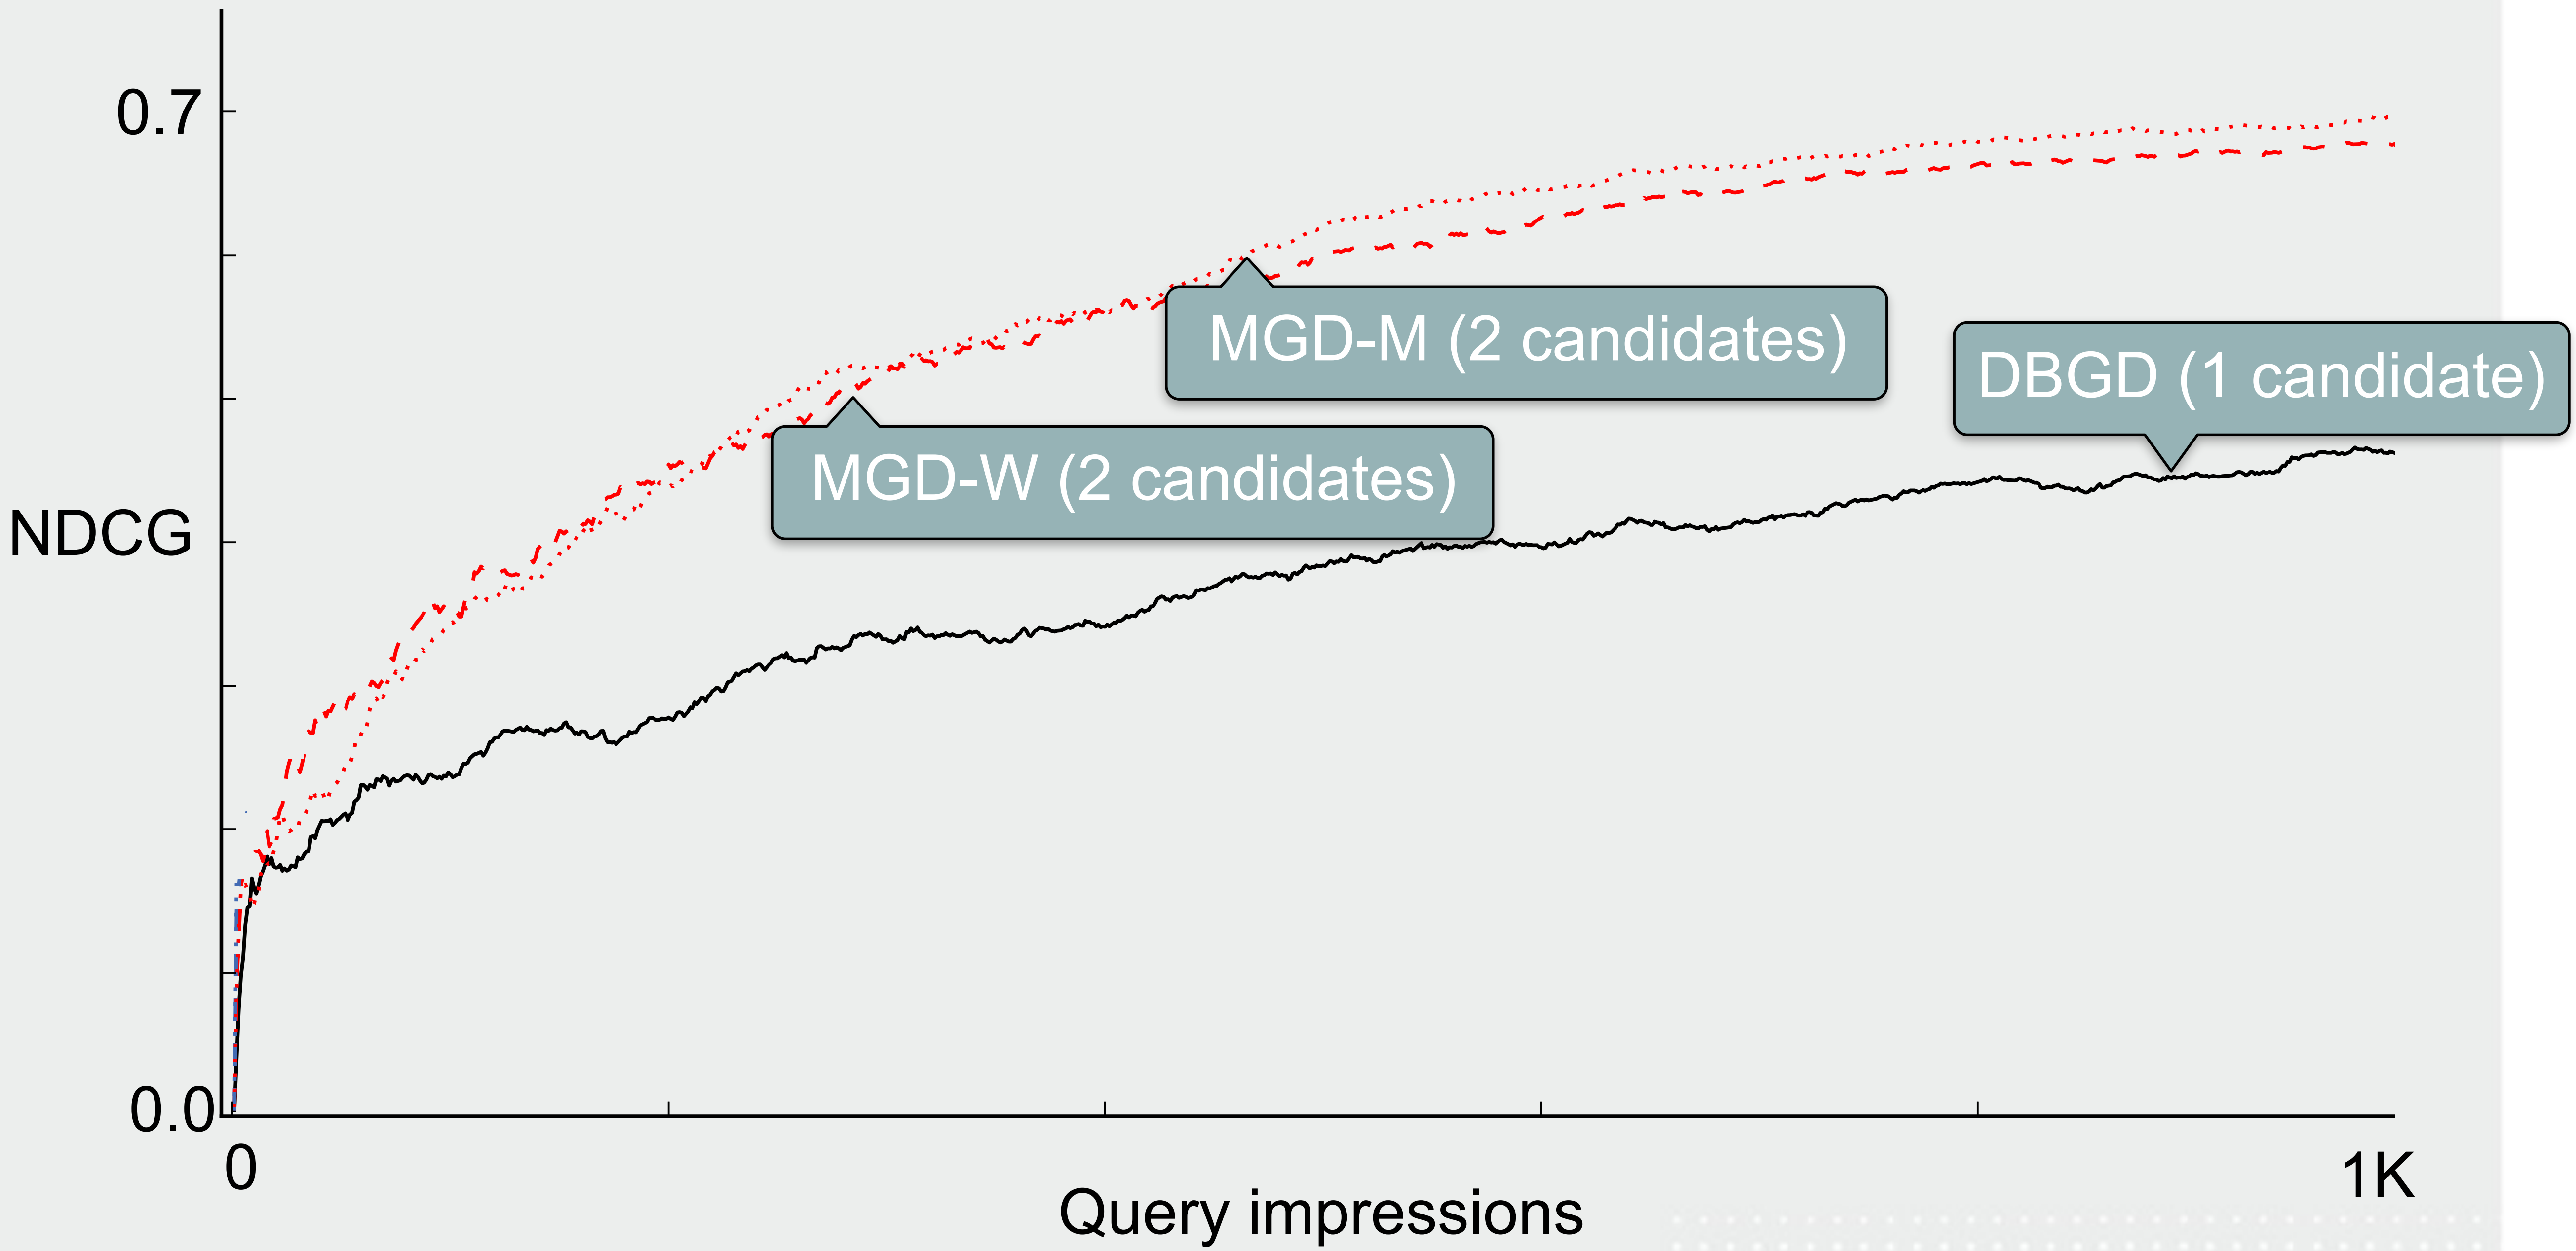

0.7
NDCG
0.0
0
1K
Query impressions
MGD-M (2 candidates)
DBGD (1 candidate)
MGD-W (2 candidates)

Results

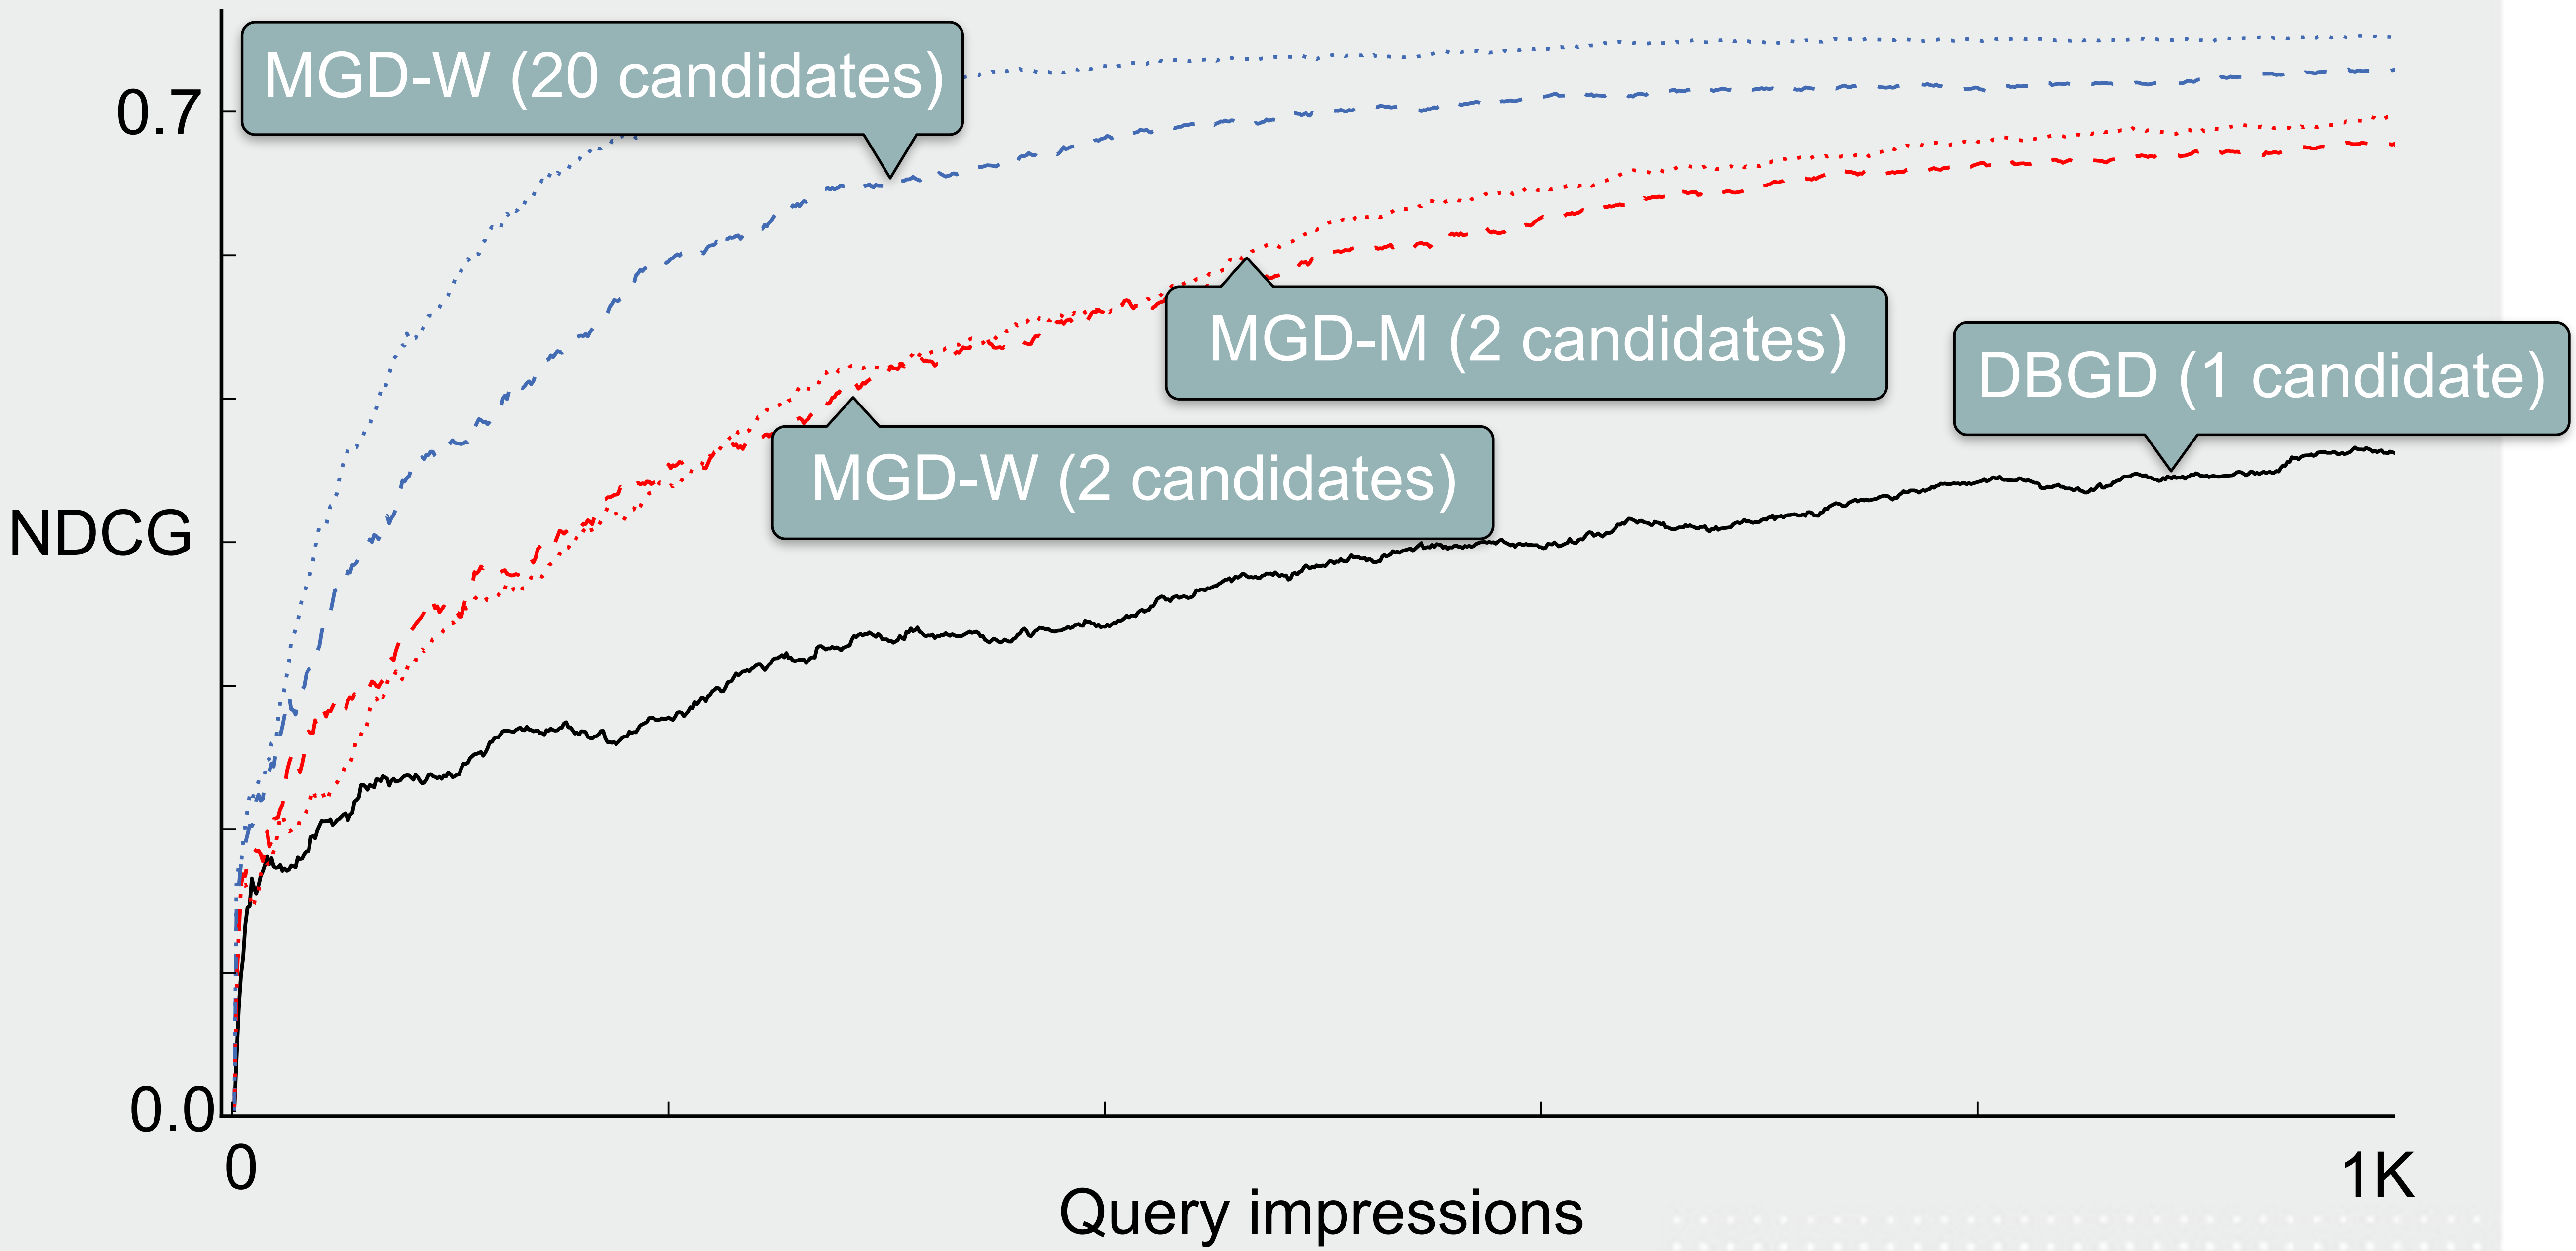
MGD-W (20 candidates)
0.7
NDCG
0.0
MGD-M (2 candidates)
DBGD (1 candidate)
MGD-W (2 candidates)
0
1K
Query impressions

Results

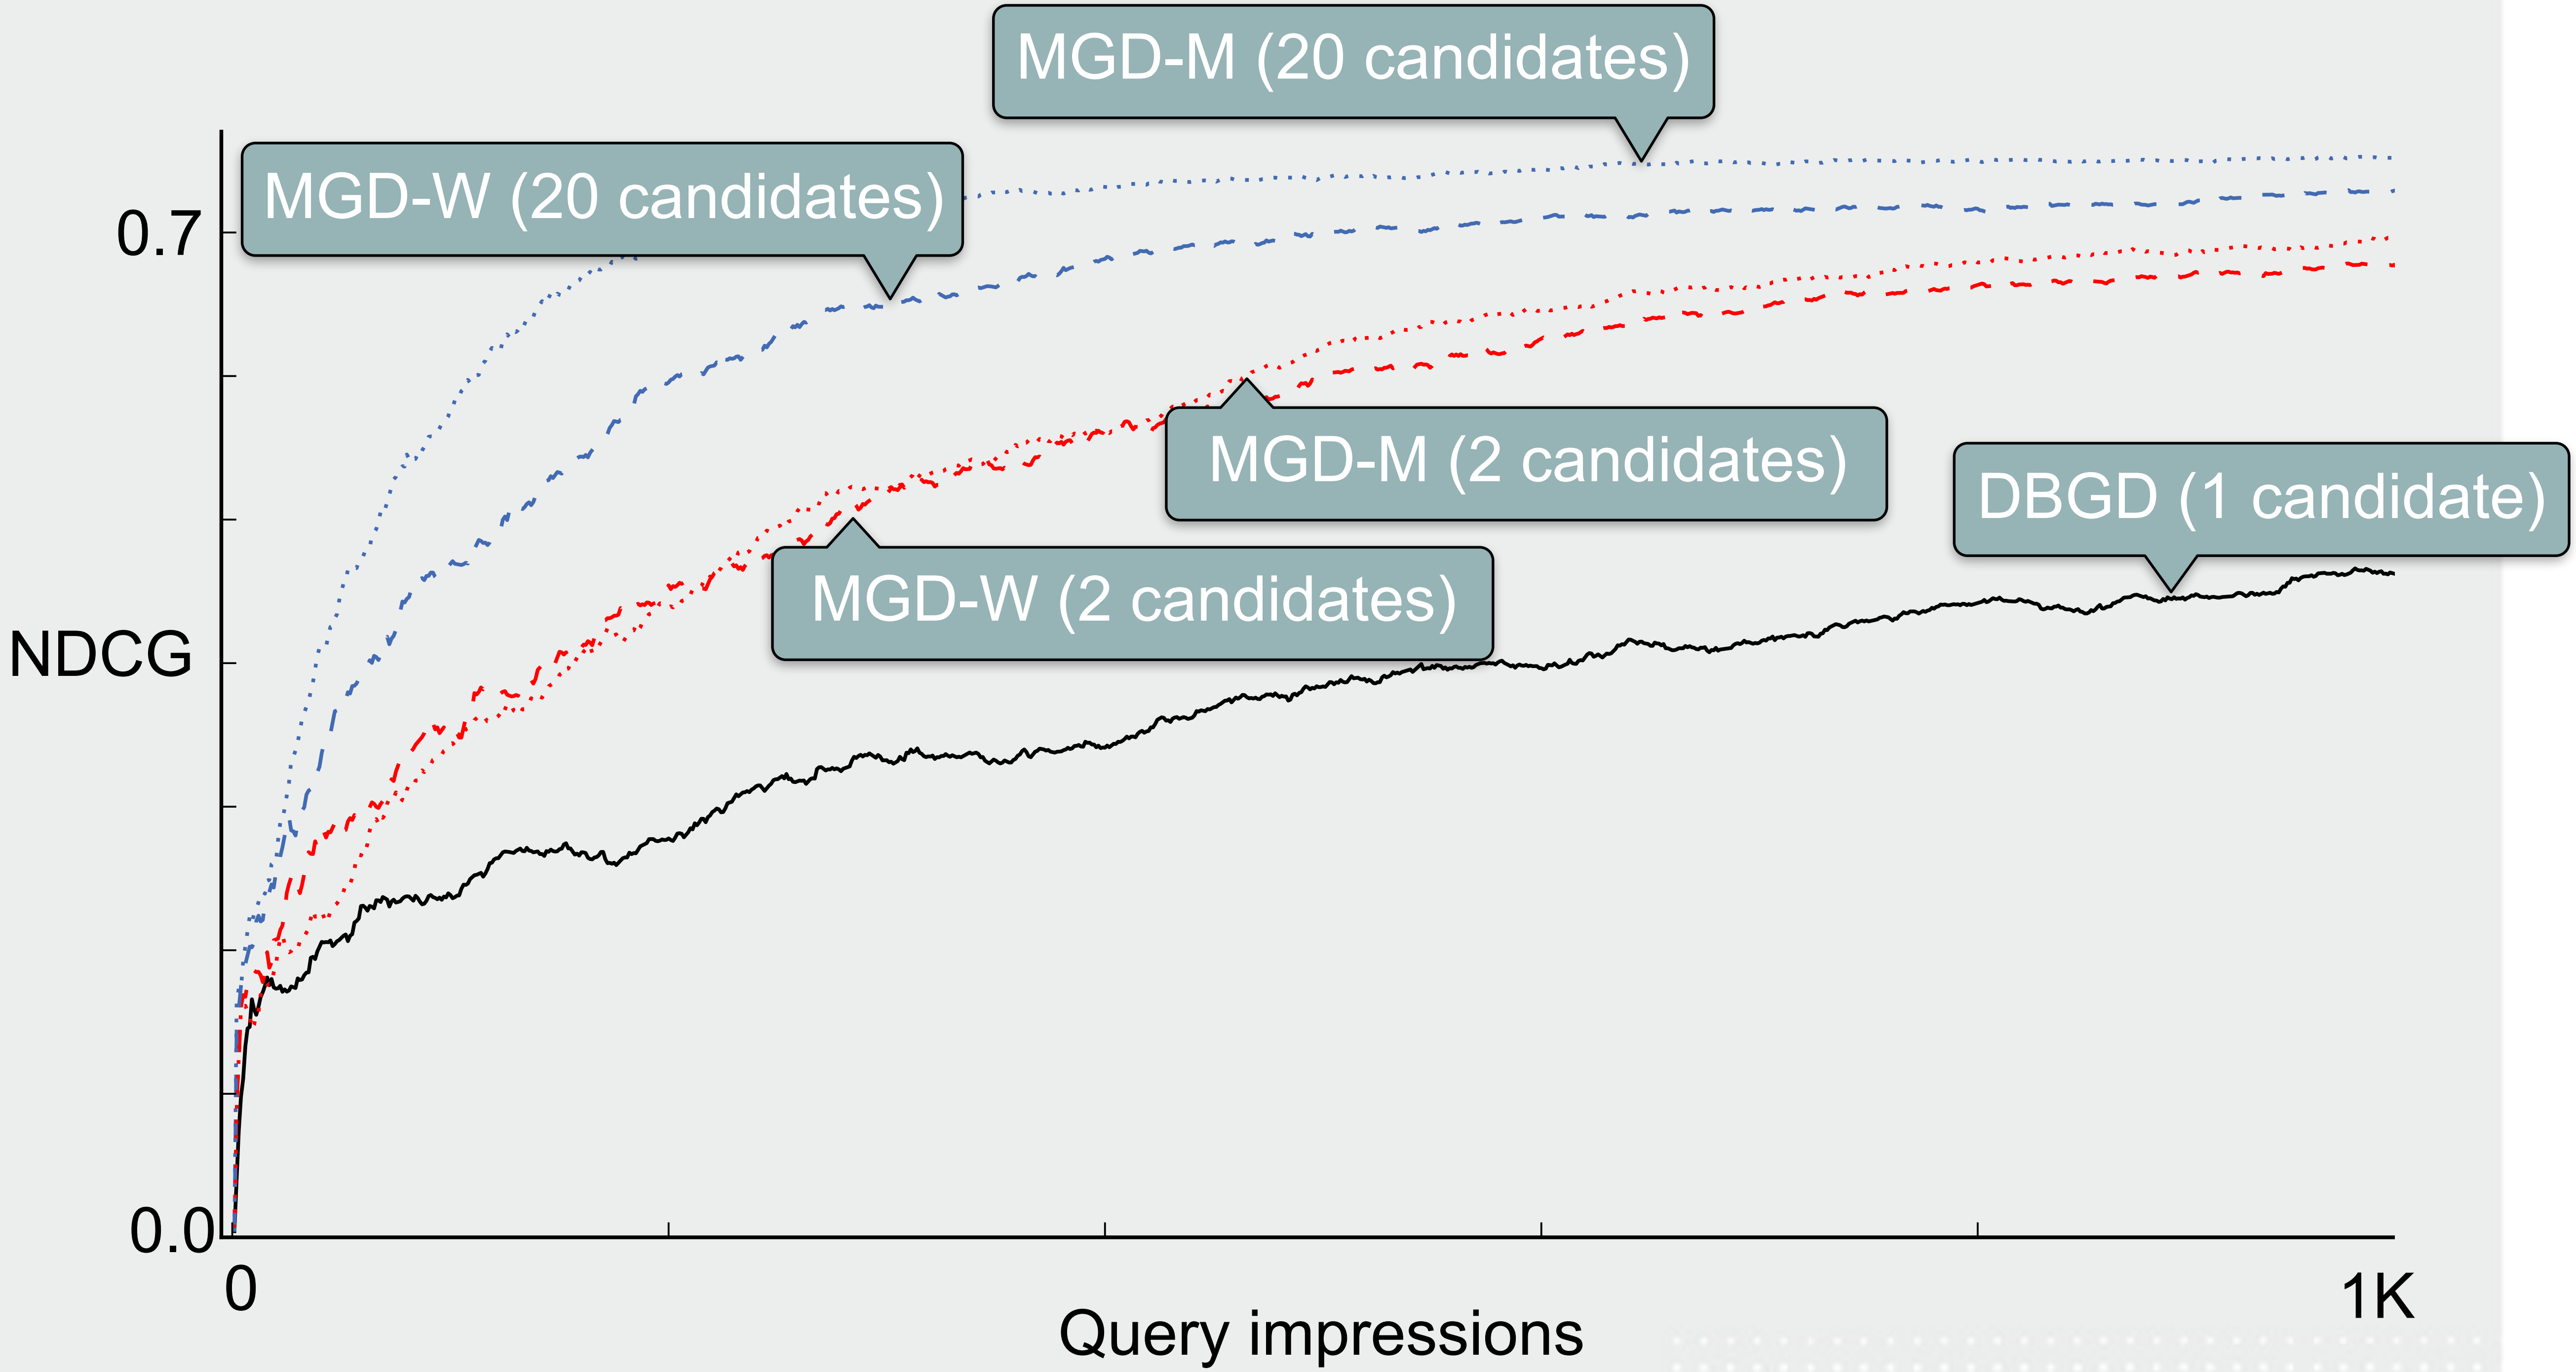
MGD-M (20 candidates)
MGD-W (20 candidates)
0.7
NDCG
MGD-M (2 candidates)
DBGD (1 candidate)
MGD-W (2 candidates)
0.0
0
1K
Query impressions

Conclusions

- We introduce MGD, an extension of DBGD
  - Multileaving instead of Interleaving
  - Two update methods MGD-M and MGD-W

Conclusions

- We introduce MGD, an extension of DBGD
  - Multileaving instead of Interleaving
  - Two update methods MGD-M and MGD-W
- Experimental validation
  - Large improvements over baseline
  - Especially with noise in feedback

Conclusions

- We introduce MGD, an extension of DBGD
  - Multileaving instead of Interleaving
  - Two update methods MGD-M and MGD-W
- Experimental validation
  - Large improvements over baseline
  - Especially with noise in feedback
- Implications
  - Orders of magnitude **less interaction data required** with MGD
  - Search engines can adapt much **faster**

Thank you

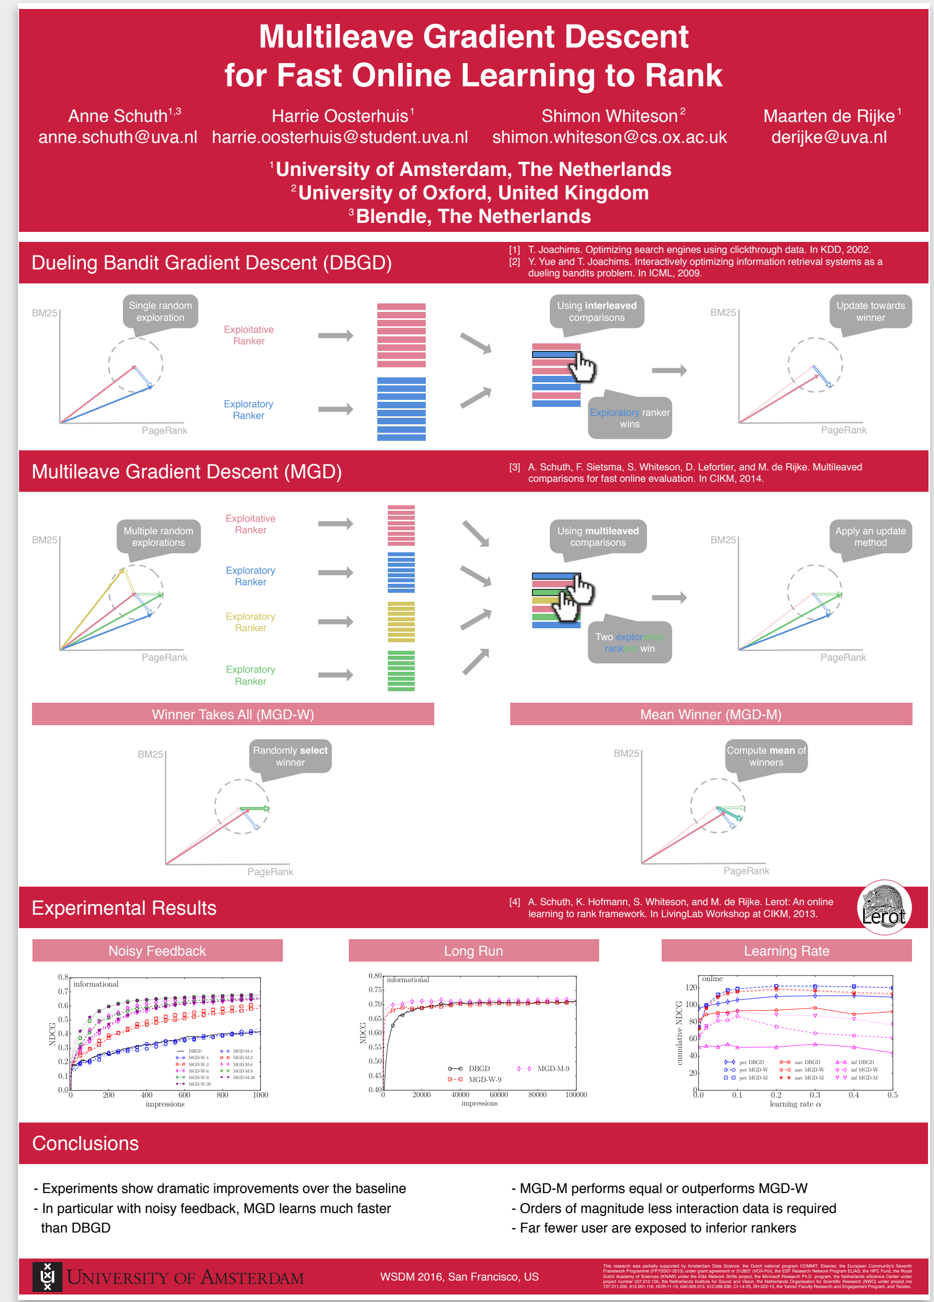
Multileave Gradient Descent
for Fast Online Learning to Rank
Anne Schuth[1,3] anne.schuth@uva.nl
Harrie Oosterhuis[1] harrie.oosterhuis@student.uva.nl
Shimon Whiteson[2] shimon.whiteson@cs.ox.ac.uk
Maarten de Rijke[1] derijke@uva.nl
[1]University of Amsterdam, The Netherlands
[2]University of Oxford, United Kingdom
[3]Blendle, The Netherlands
Dueling Bandit Gradient Descent (DBGD)
Multileave Gradient Descent (MGD)
Winner Takes All (MGD-W)
Mean Winner (MGD-M)
Experimental Results
Noisy Feedback
Long Run
Learning Rate
Conclusions
- Experiments show dramatic improvements over the baseline
- In particular with noisy feedback, MGD learns much faster than DBGD
- MGD-M performs equal or outperforms MGD-W
- Orders of magnitude less interaction data is required
- Far fewer user are exposed to inferior rankers
WSDM 2016, San Francisco, US