UNIVERSITY OF AMSTERDAM

Multileave Gradient Descent for Fast Online Learning to Rank

Anne Schuth, Harrie Oosterhuis, Shimon Whiteson, Maarten de Rijke

Search Engines that Learn

Search Engines that Learn

- Search engines are complex machines
  - Combining hundreds of ranking signals

Search Engines that Learn

- Search engines are complex machines
  - Combining hundreds of ranking signals
- Learning to rank
  - Learning how to combine features

Search Engines that Learn

- Search engines are complex machines
  - Combining hundreds of ranking signals
- Learning to rank
  - Learning how to combine features
- Offline
  - Using (labeled) static datasets

Search Engines that Learn

- Search engines are complex machines
  - Combining hundreds of ranking signals
- Learning to rank
  - Learning how to combine features
- Offline
  - Using (labeled) static datasets
- Online
  - Directly from users

Dueling Bandit Gradient Descent (DBGD)

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

Dueling Bandit Gradient Descent (DBGD)

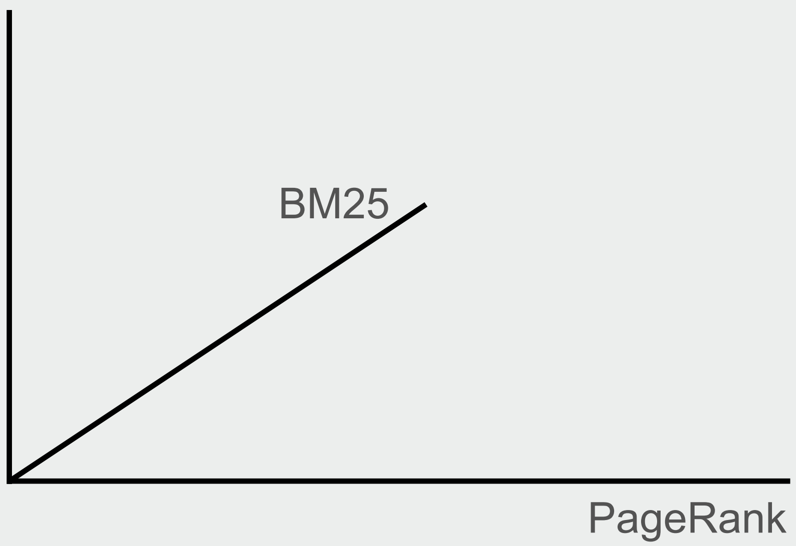

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Dueling Bandit Gradient Descent (DBGD)

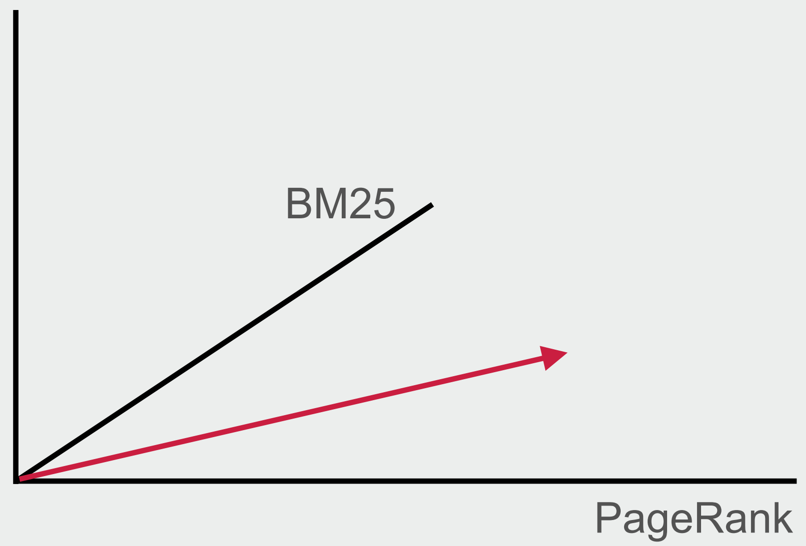

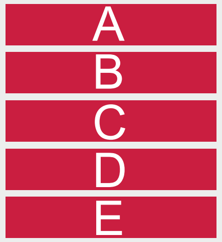

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Dueling Bandit Gradient Descent (DBGD)

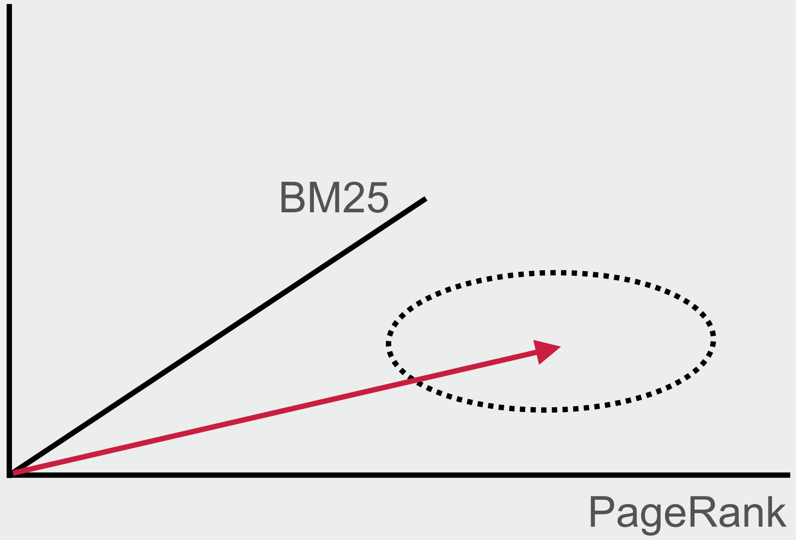

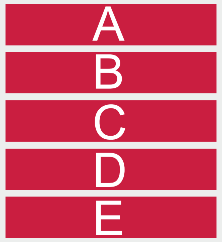

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Dueling Bandit Gradient Descent (DBGD)

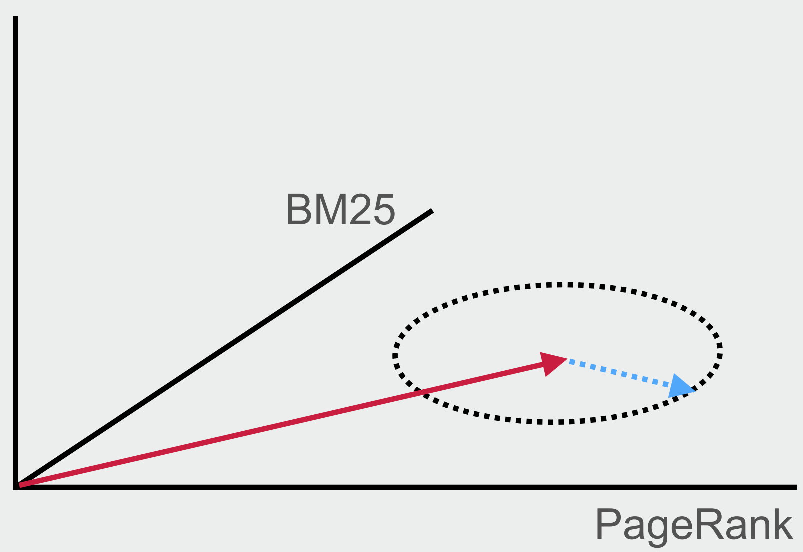

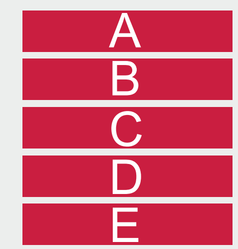

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Dueling Bandit Gradient Descent (DBGD)

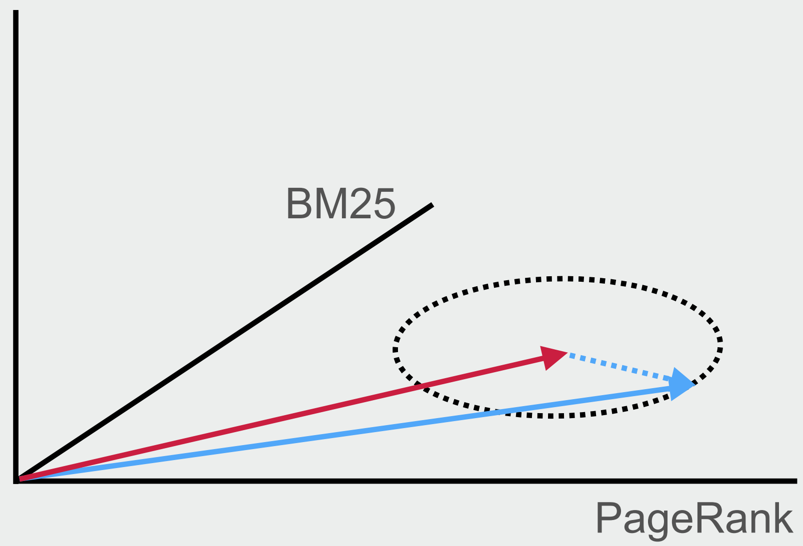

A	F
B	A
C	E
D	B
E	C

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

Dueling Bandit Gradient Descent (DBGD)

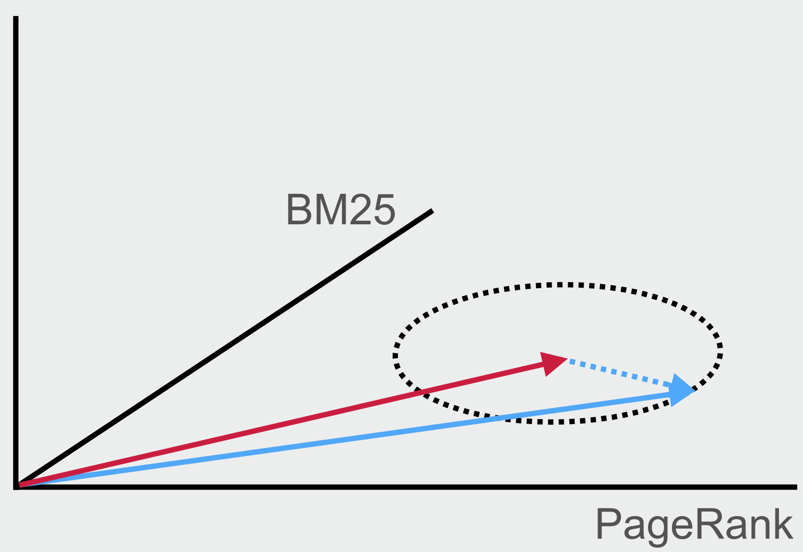

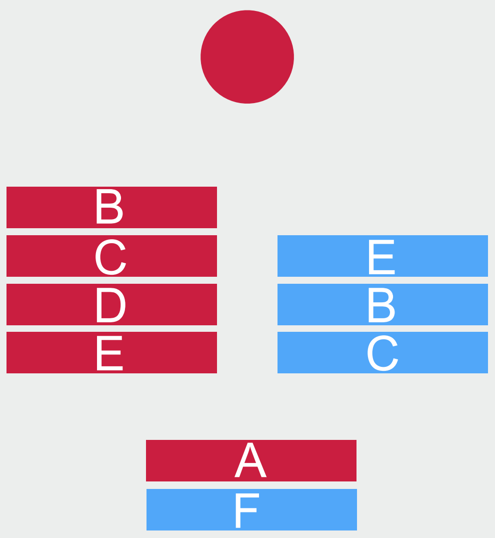

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Dueling Bandit Gradient Descent (DBGD)

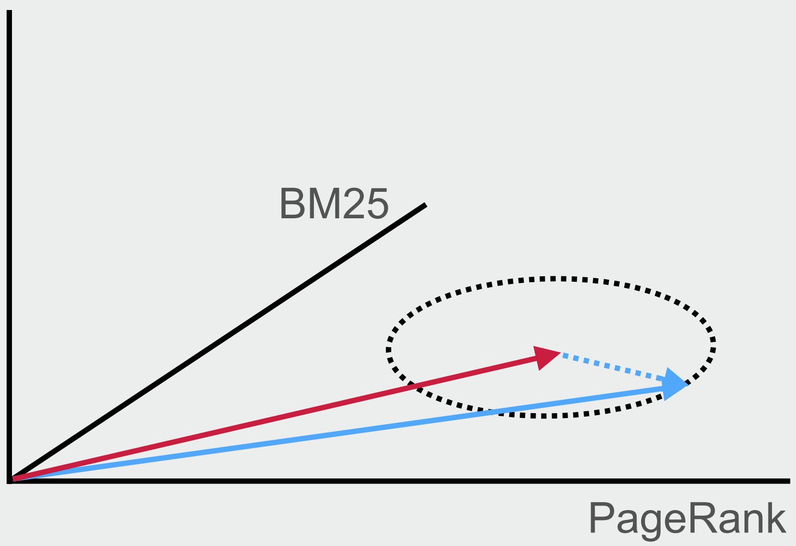

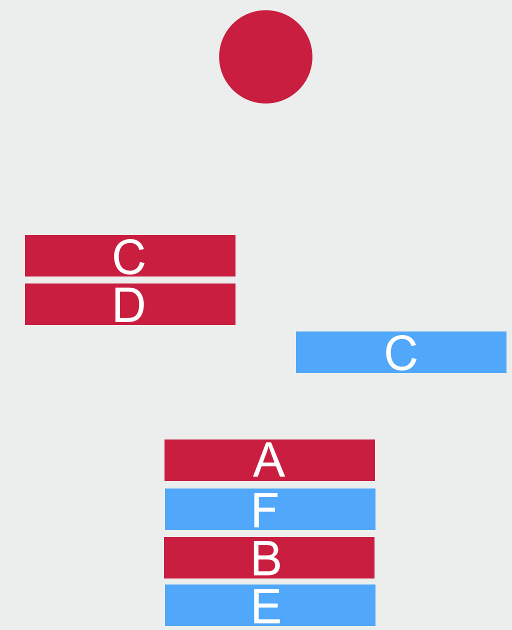

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Dueling Bandit Gradient Descent (DBGD)

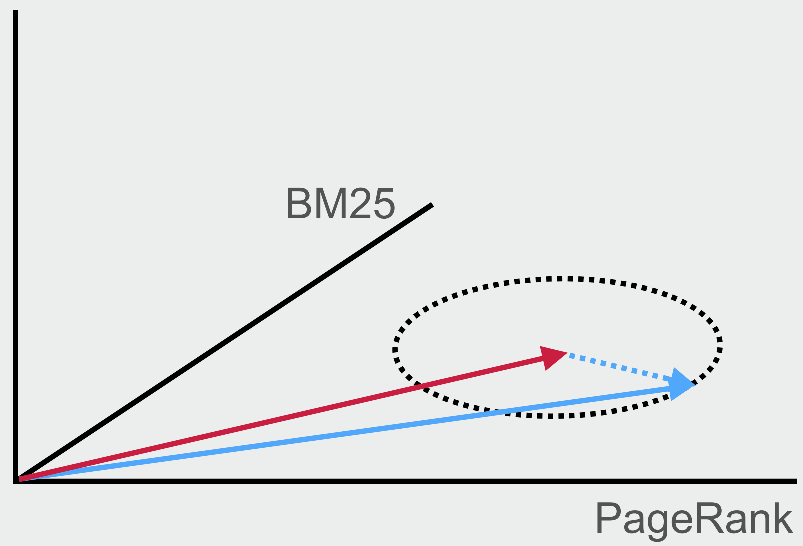

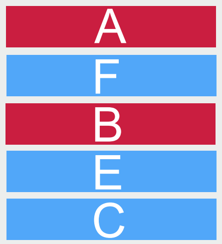

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Dueling Bandit Gradient Descent (DBGD)

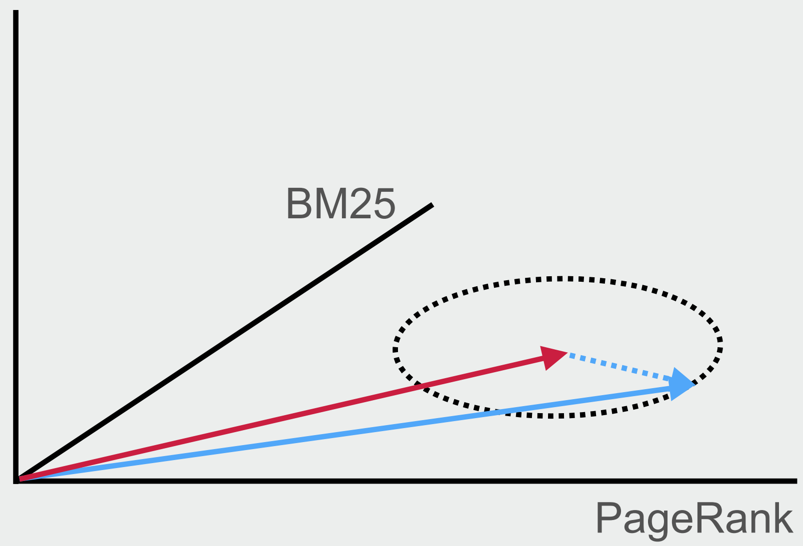

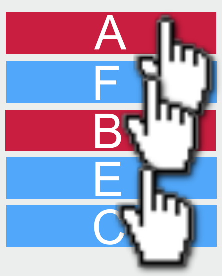

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Dueling Bandit Gradient Descent (DBGD)

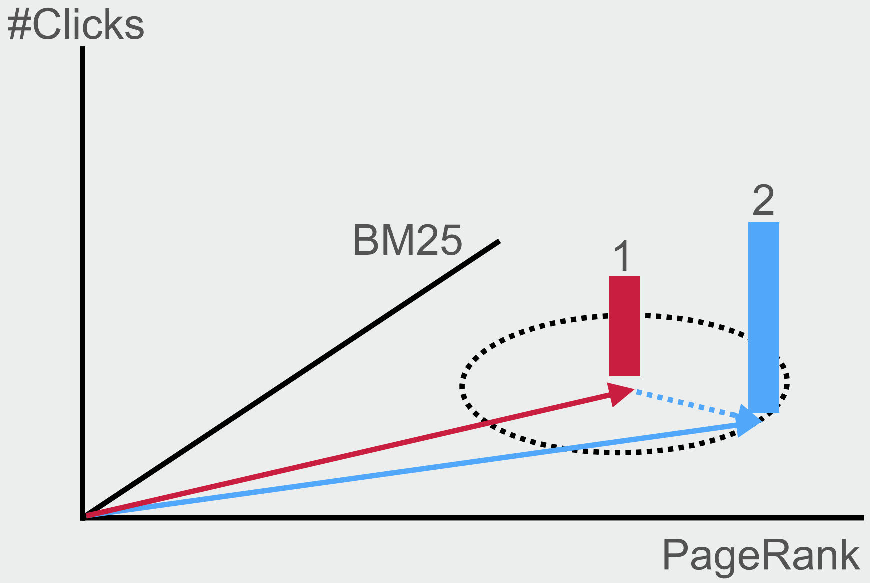

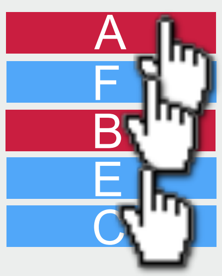

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Dueling Bandit Gradient Descent (DBGD)

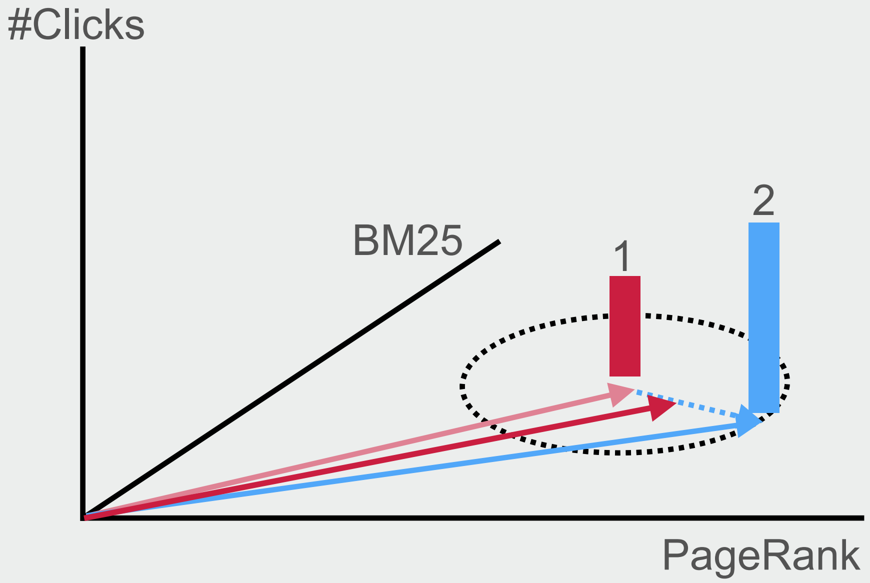

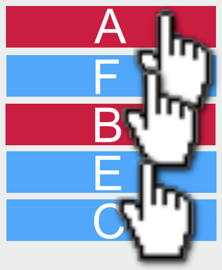

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Dueling Bandit Gradient Descent (DBGD)

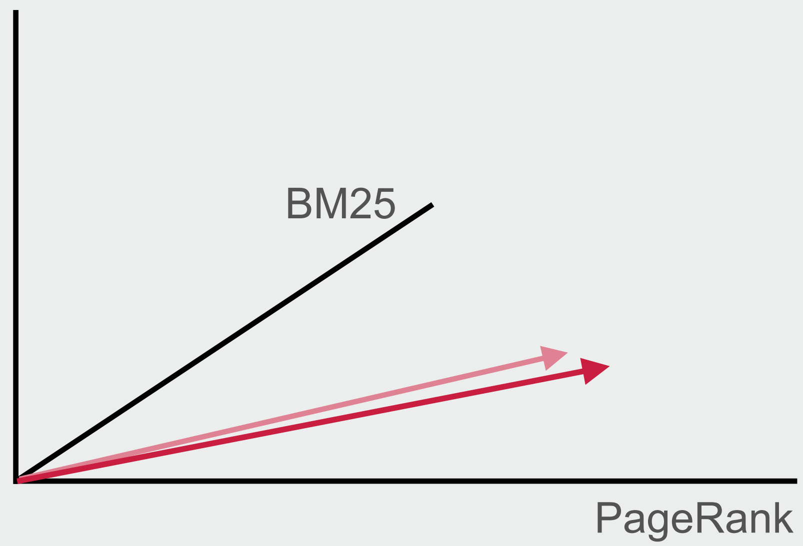

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008]

Dueling Bandit Gradient Descent (DBGD)

- Updates after exploring a **single** direction
- Exploring multiple directions before updating would be beneficial
  - Fewer updates would lead to a better ranker
- But would be **expensive** when interleaving was used
  - All directions require pairwise comparisons
- Multileaved comparisons come to the rescue

Multileave Gradient Descent (MGD)

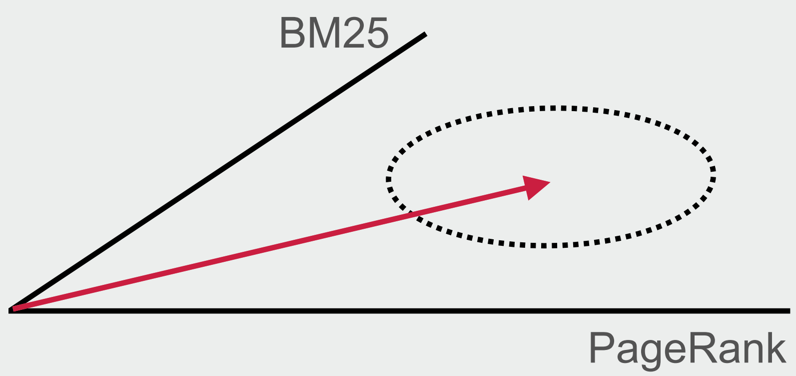

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Multileave Gradient Descent (MGD)

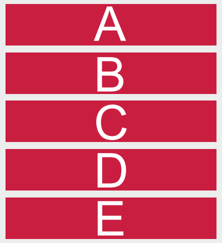

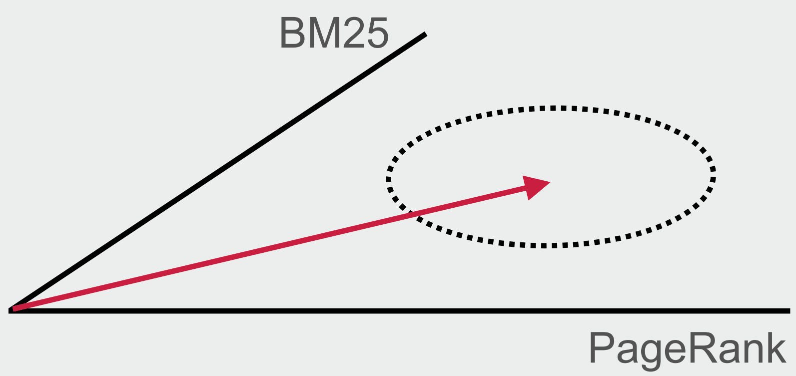

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Multileave Gradient Descent (MGD)

A
B
C
D
E

C
A
B
E
G

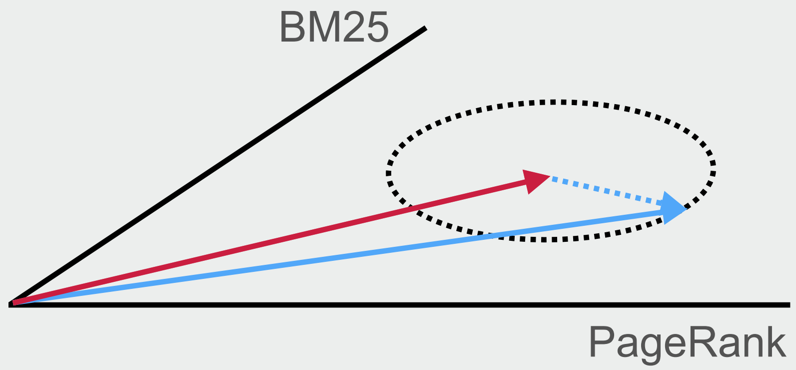

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Multileave Gradient Descent (MGD)

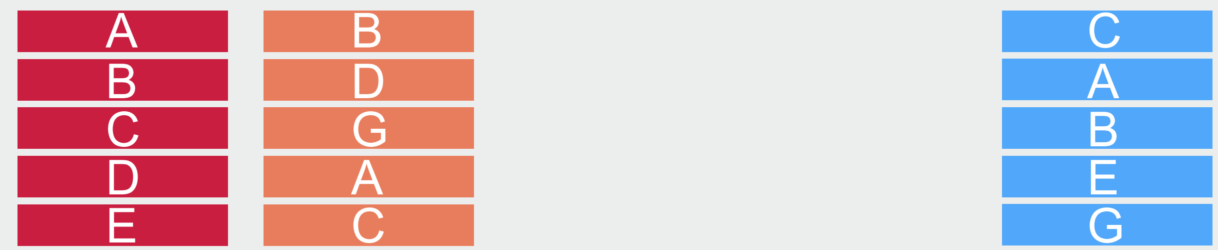

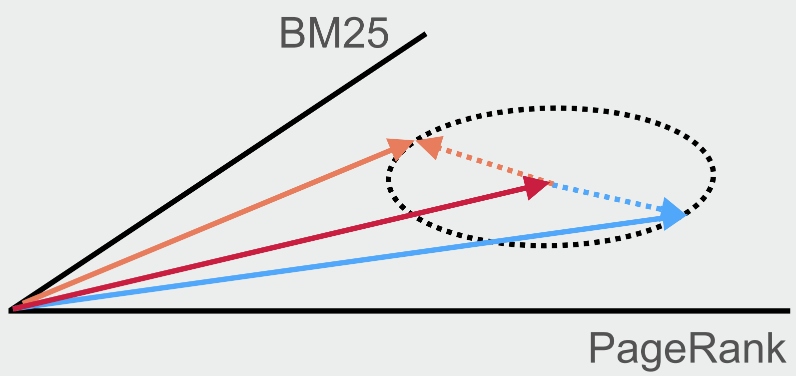

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Multileave Gradient Descent (MGD)

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

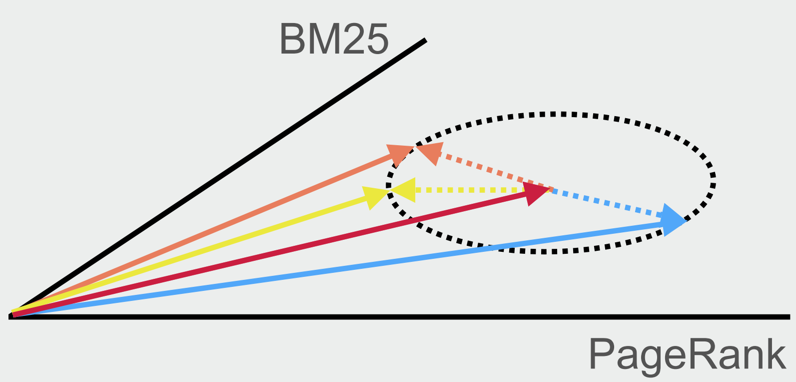

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

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

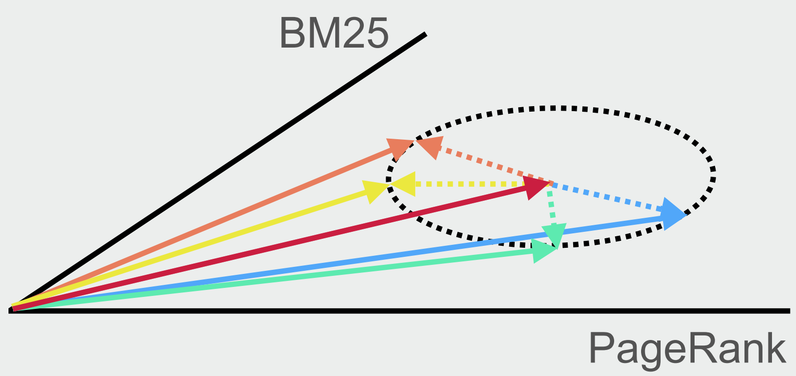

[Yue et al., 2009; Hofmann et al., 2011;
Radlinski et al., 2008; Schuth et al., 2014]

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

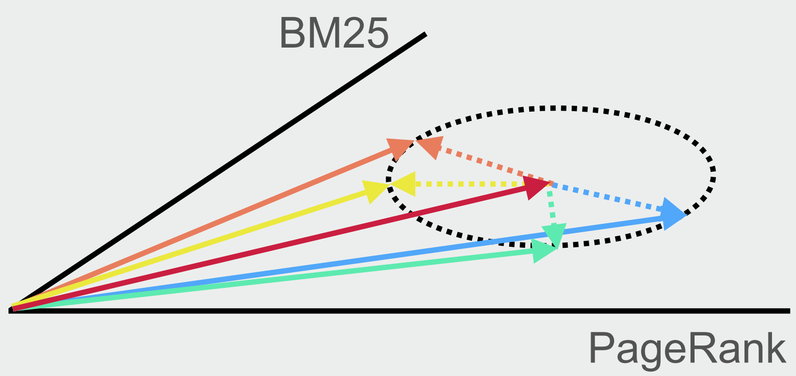

[Yue et al., 2009; Hofmann et al., 2011;
Radlinski et al., 2008; Schuth et al., 2014]

Multileave Gradient Descent (MGD)

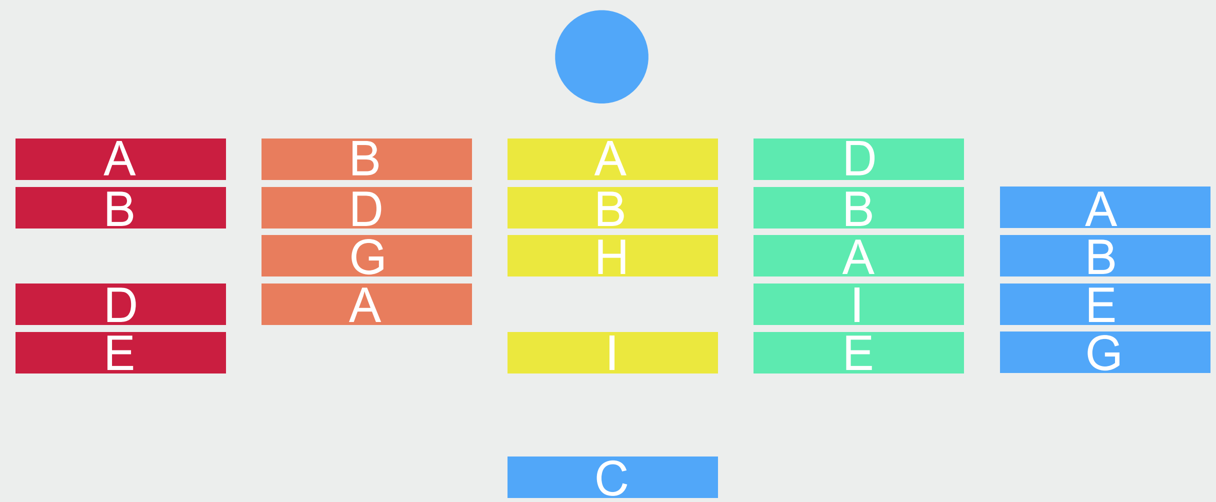

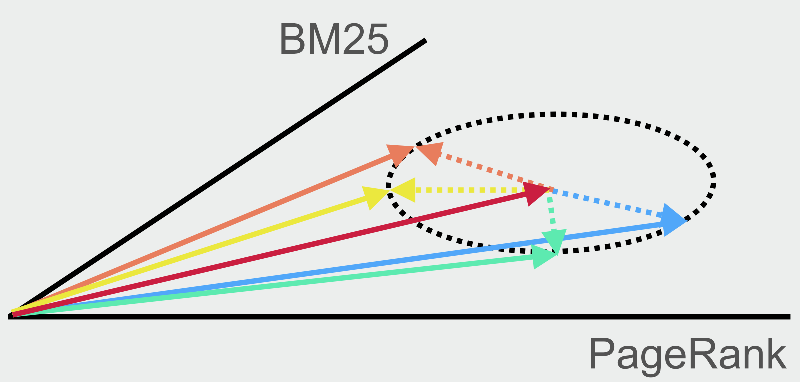

[Yue et al., 2009; Hofmann et al., 2011;
Radlinski et al., 2008; Schuth et al., 2014]

Multileave Gradient Descent (MGD)

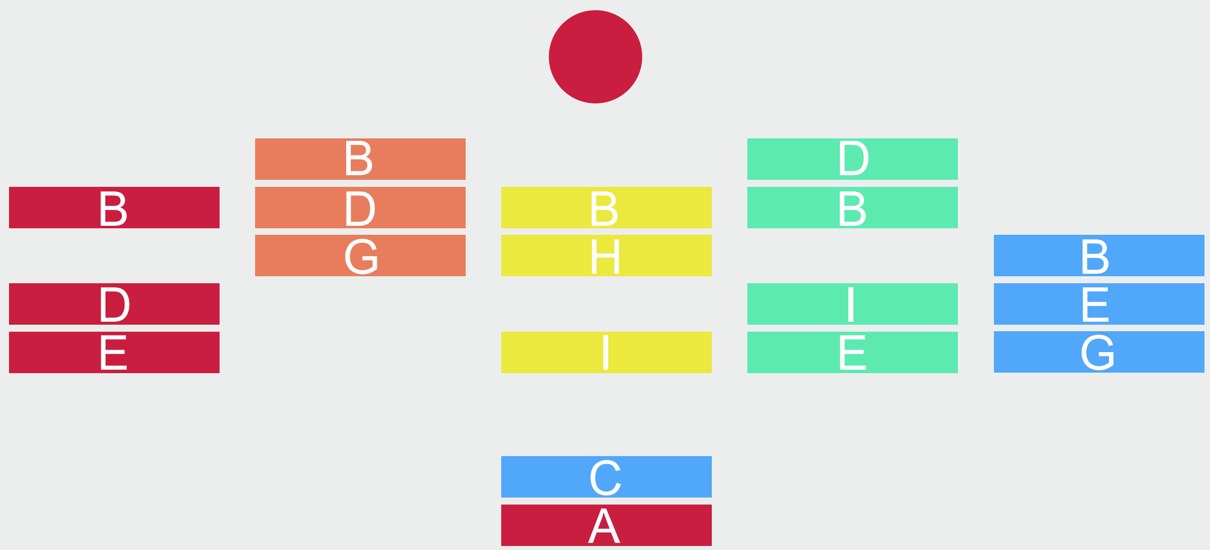

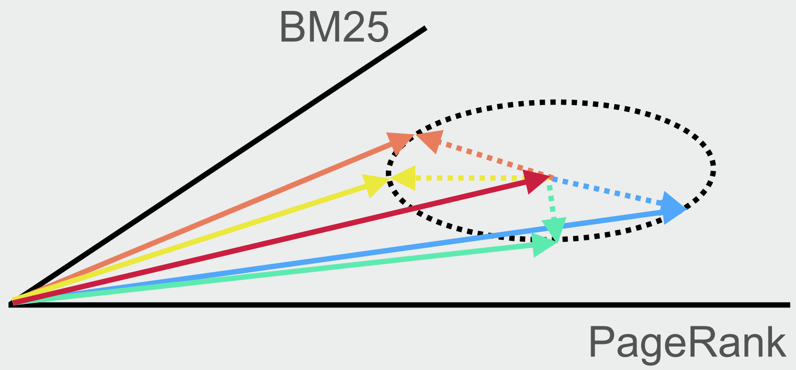

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Multileave Gradient Descent (MGD)

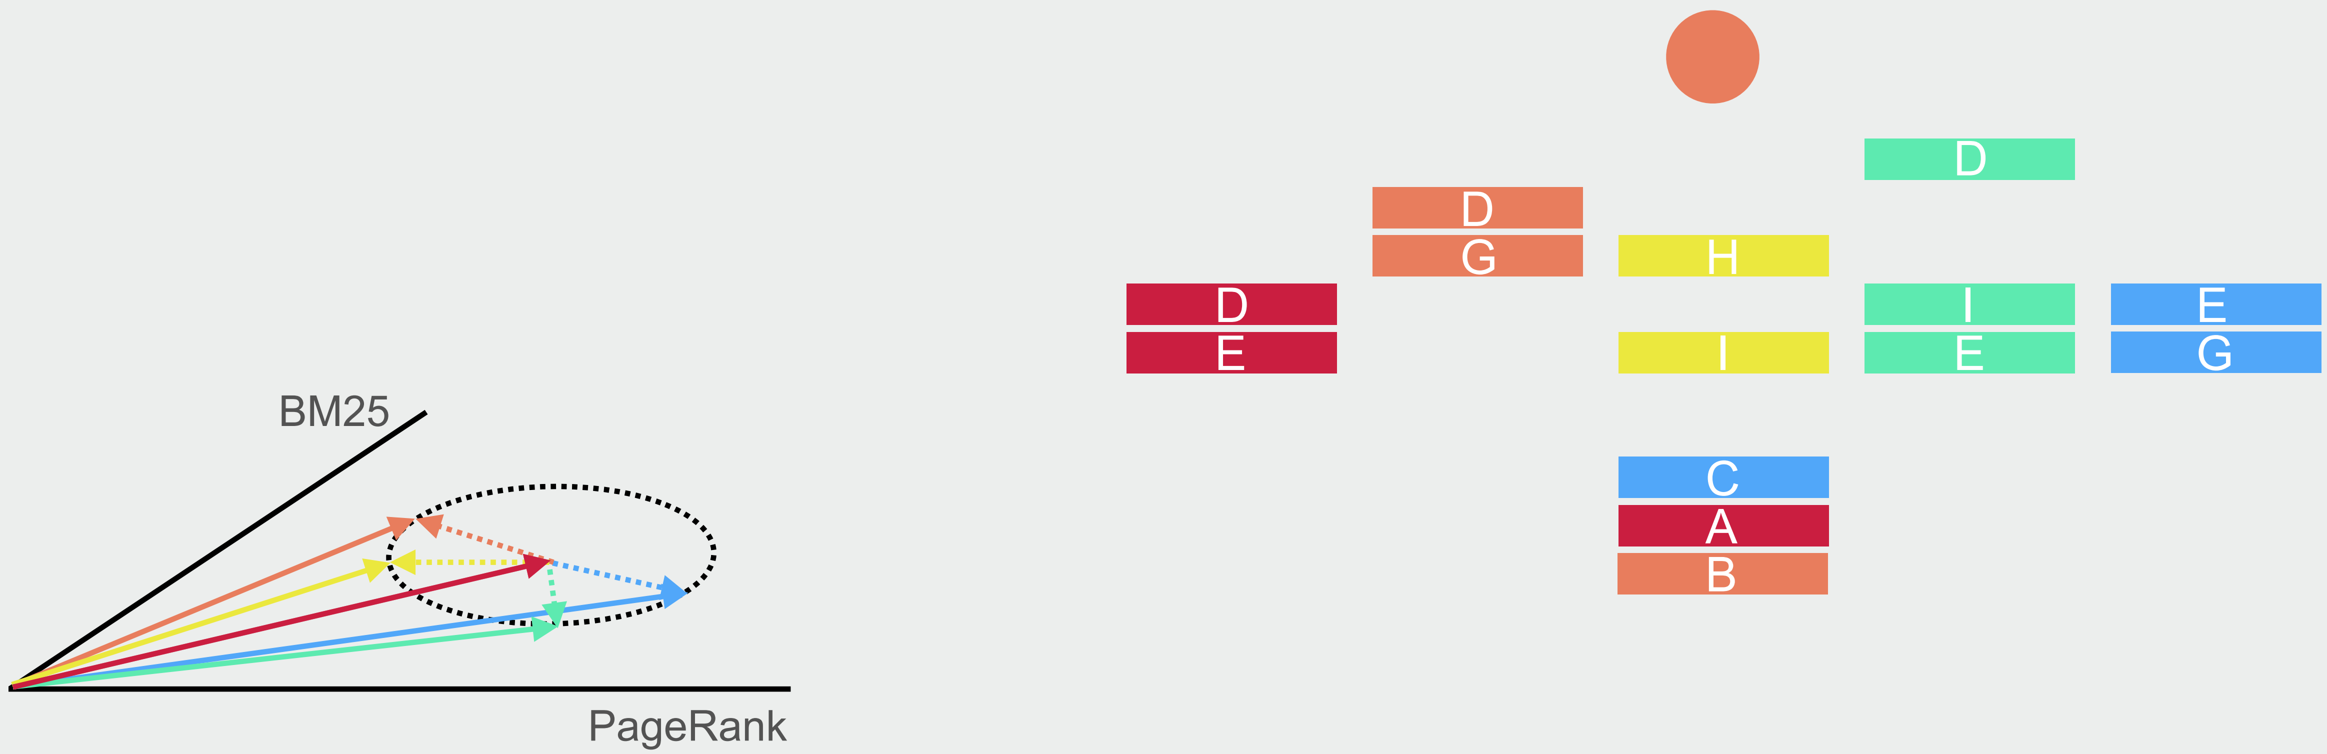

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Multileave Gradient Descent (MGD)

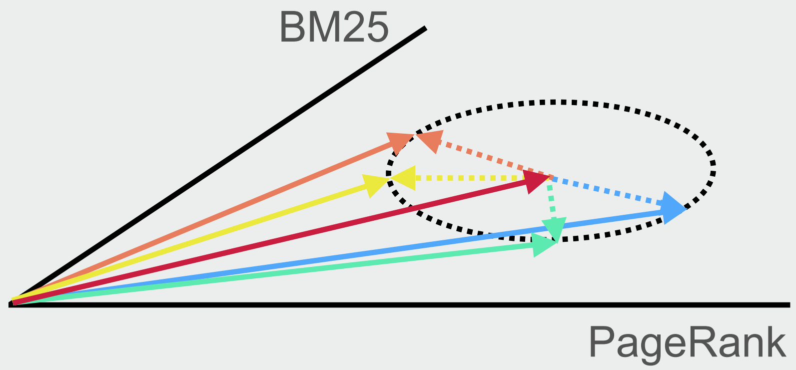

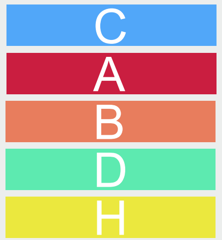

[Yue et al., 2009; Hofmann et al., 2011; Radlinski et al., 2008; Schuth et al., 2014]

Multileave Gradient Descent (MGD)

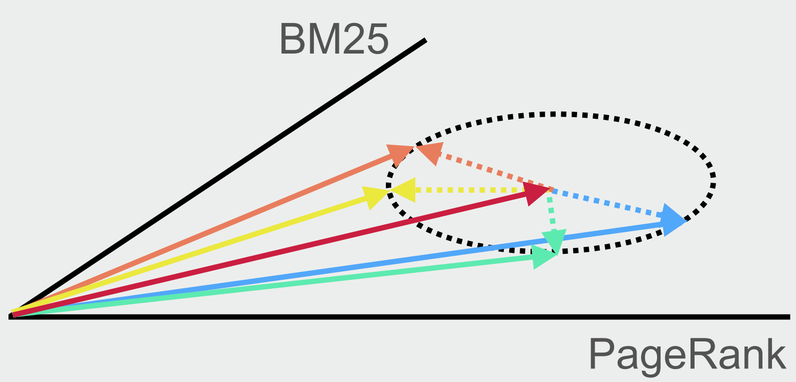

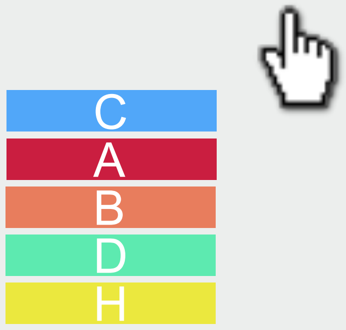

[Yue et al., 2009; Hofmann et al., 2011;
Radlinski et al., 2008; Schuth et al., 2014]

Multileave Gradient Descent (MGD)

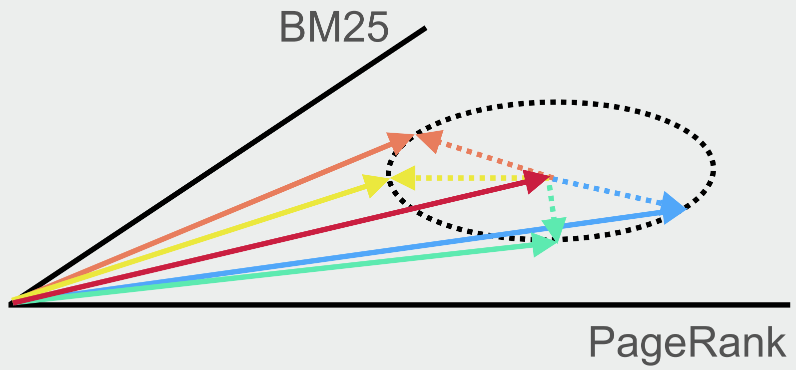

[Yue et al., 2009; Hofmann et al., 2011;
Radlinski et al., 2008; Schuth et al., 2014]

Multileave Gradient Descent (MGD)

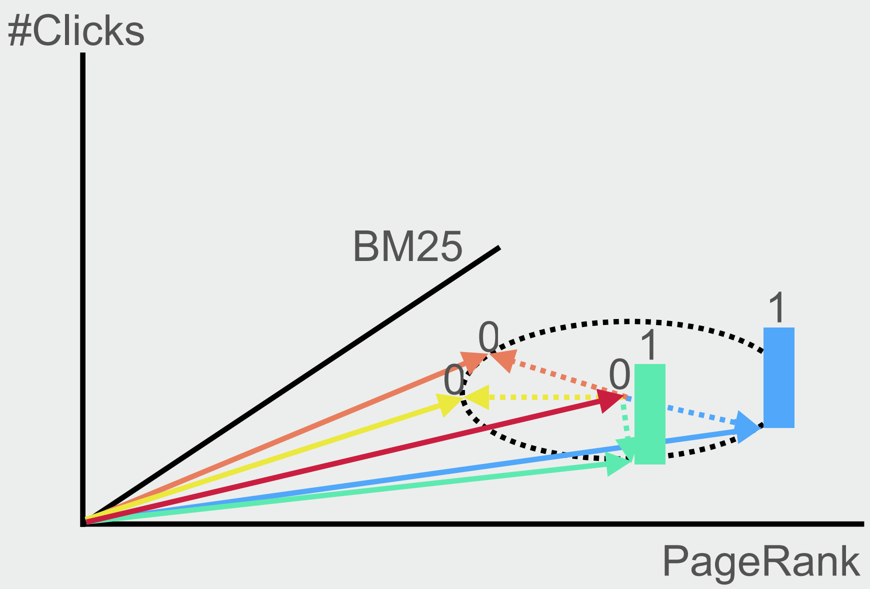

[Yue et al., 2009; Hofmann et al., 2011;
Radlinski et al., 2008; Schuth et al., 2014]

Multileave Gradient Descent (MGD)

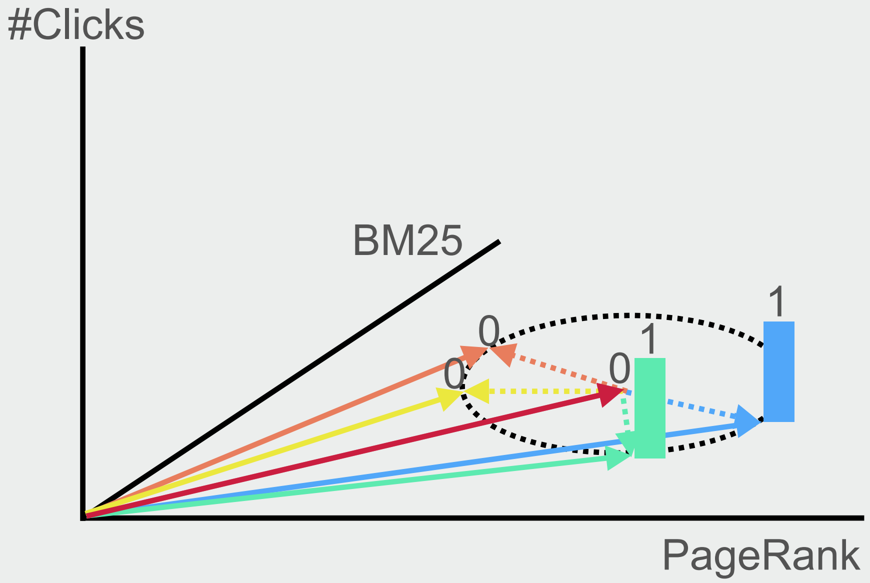

Multileave Gradient Descent (MGD)

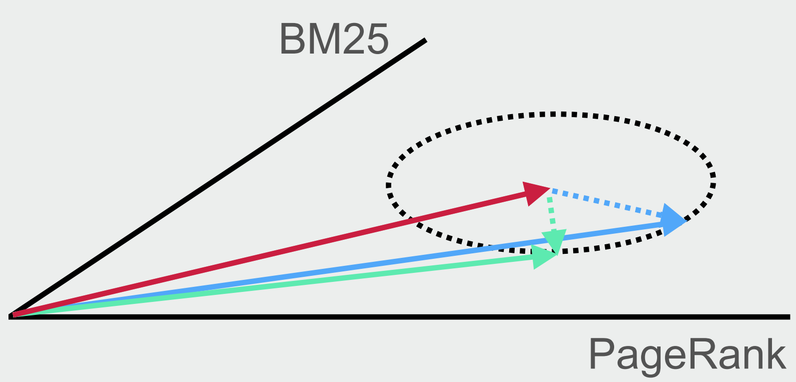

Multileave Gradient Descent (MGD)

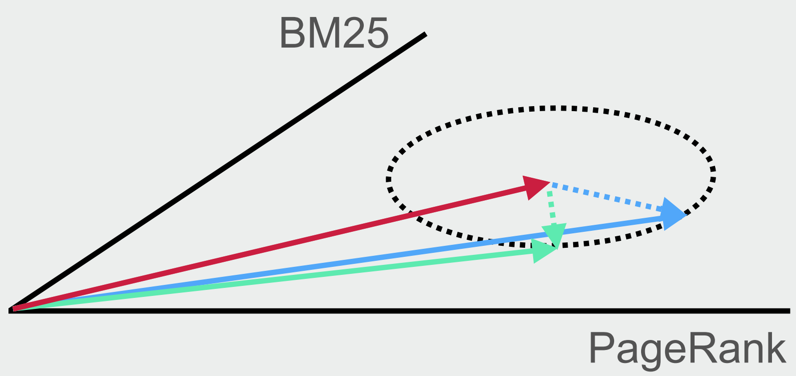

Multileave Gradient Descent (MGD)

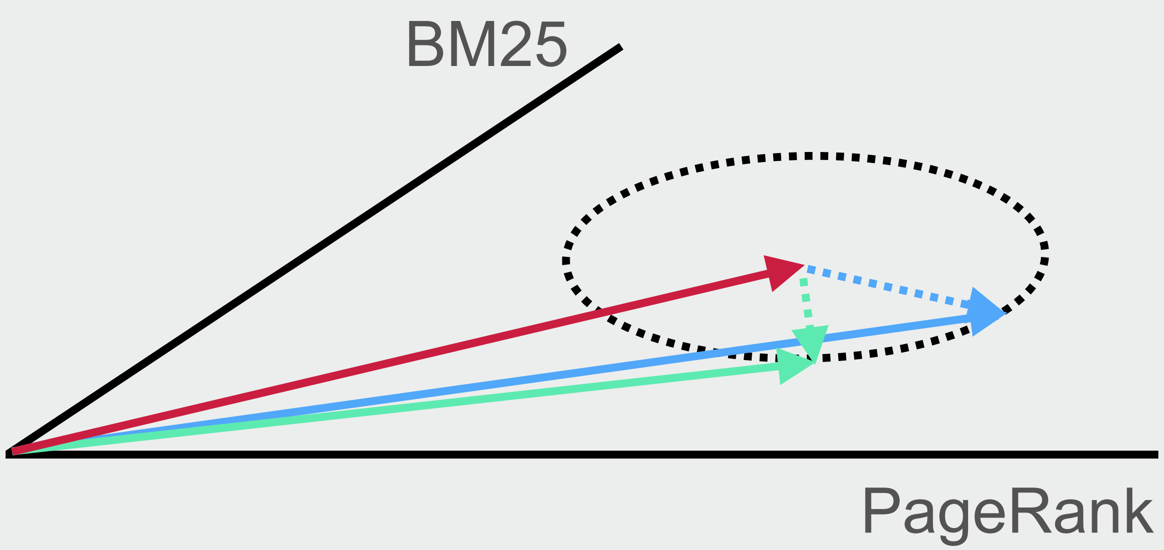

Multileave Gradient Descent (MGD)

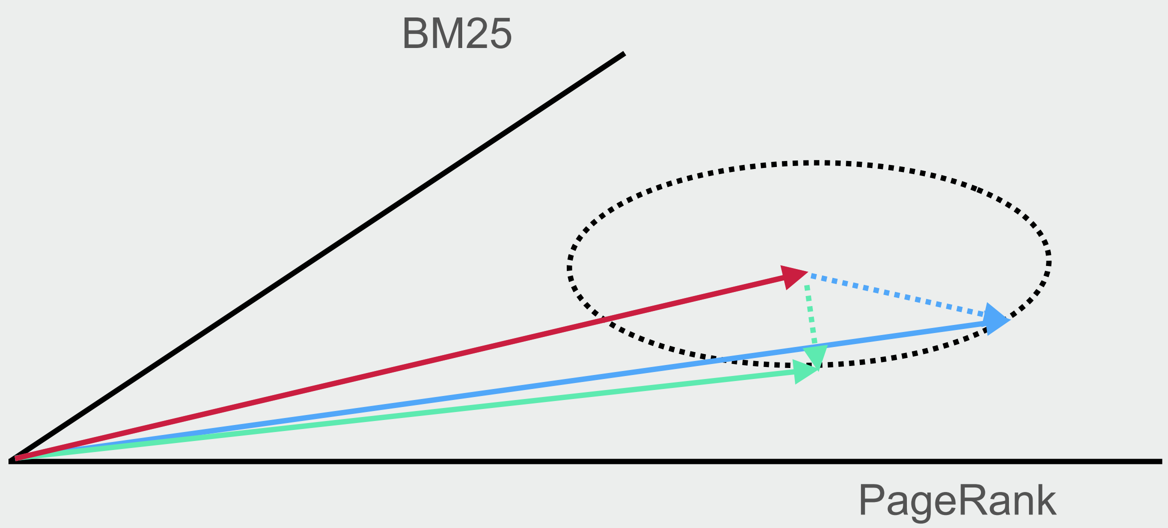

Multileave Gradient Descent (MGD)

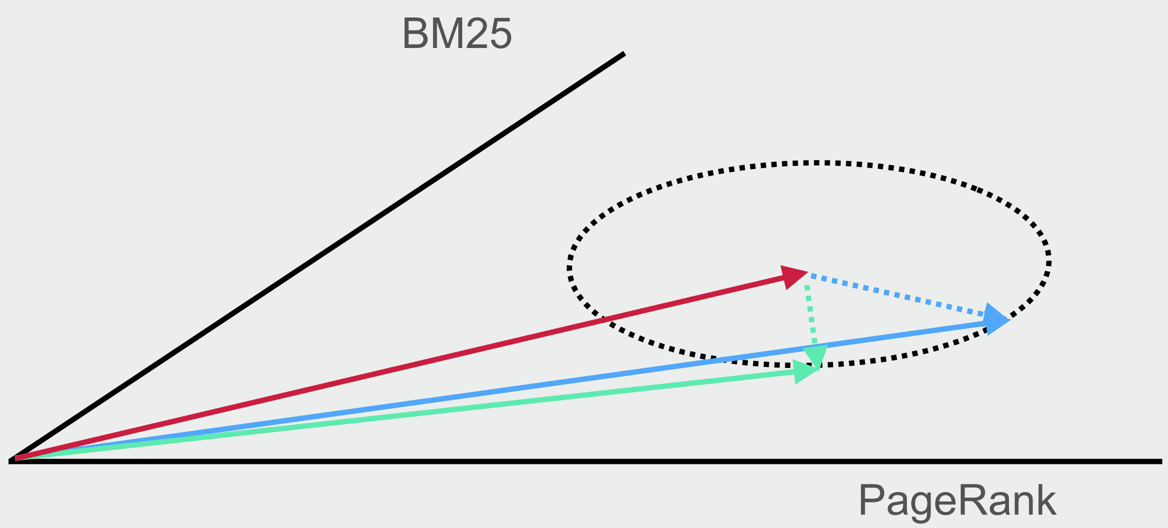

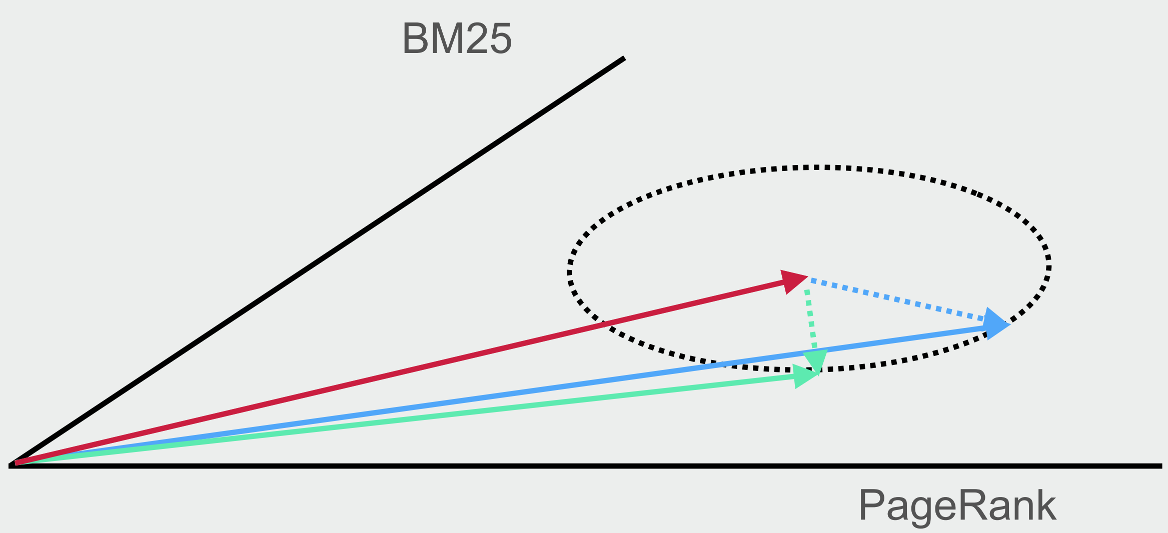

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

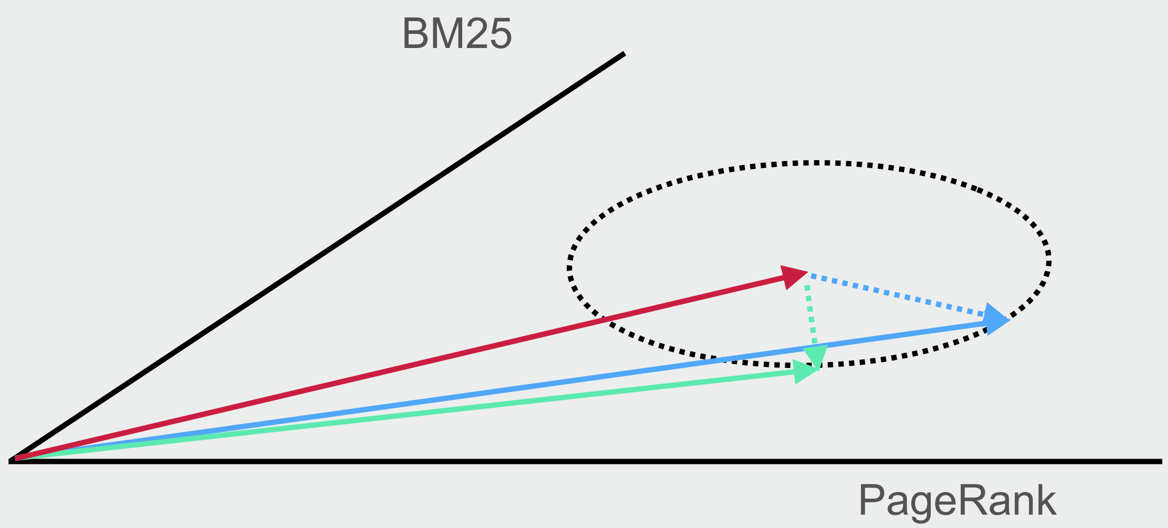

Mean winner (MGD-M)

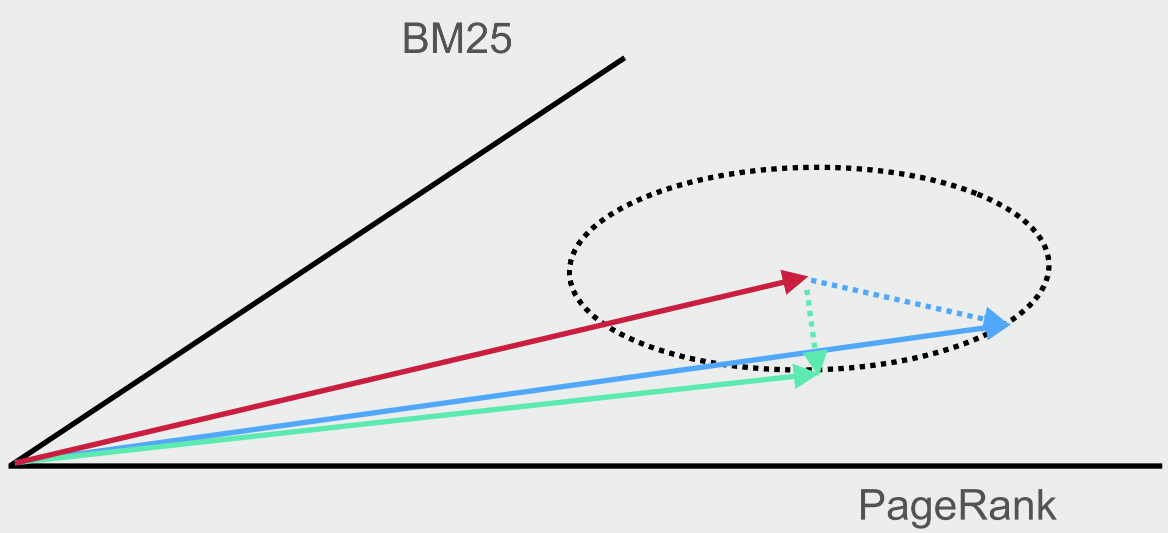

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

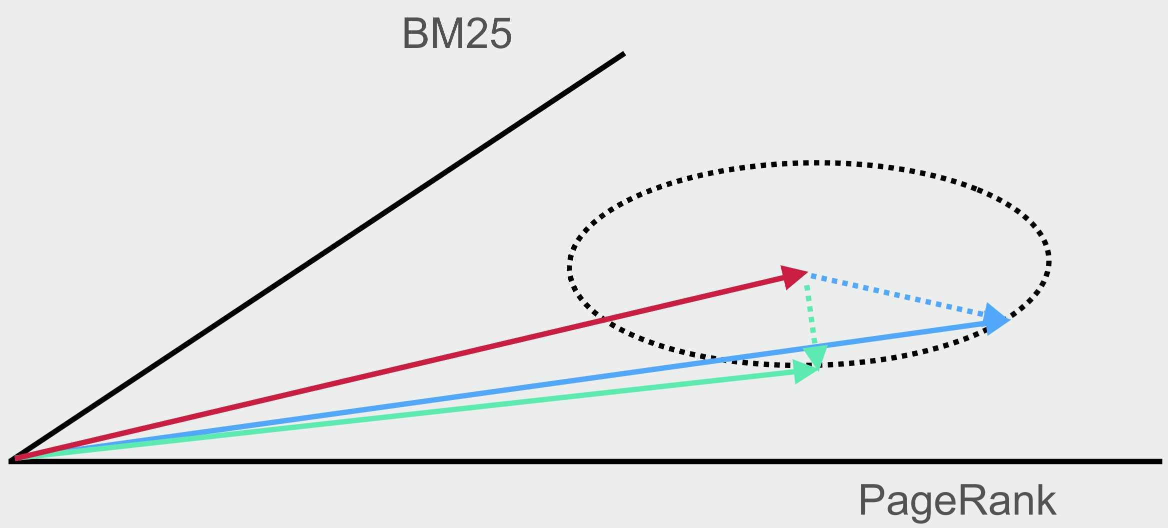

Mean winner (MGD-M)

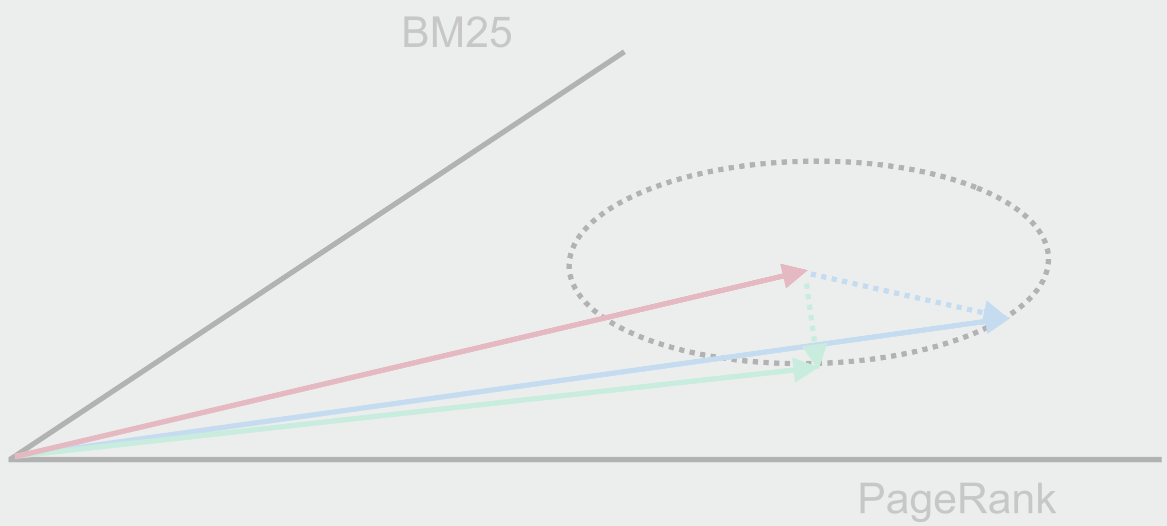

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

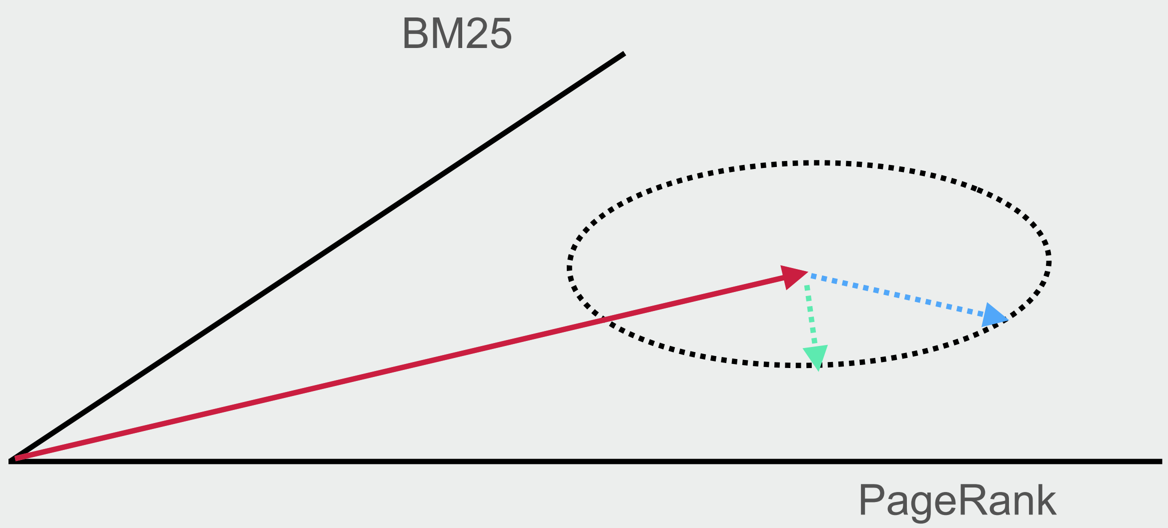

Mean winner (MGD-M)

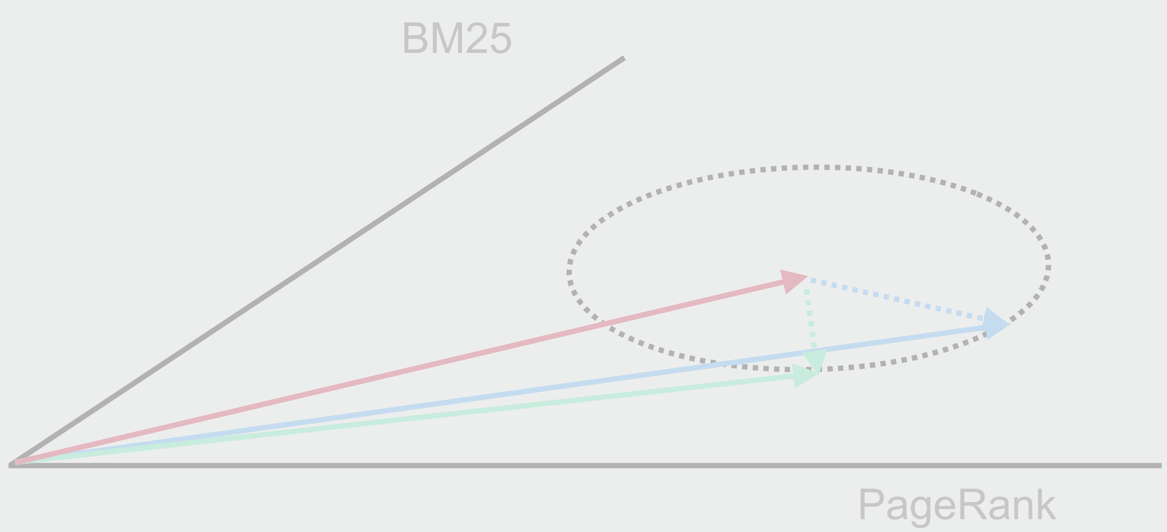

UNIVERSITY OF AMSTERDAM

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

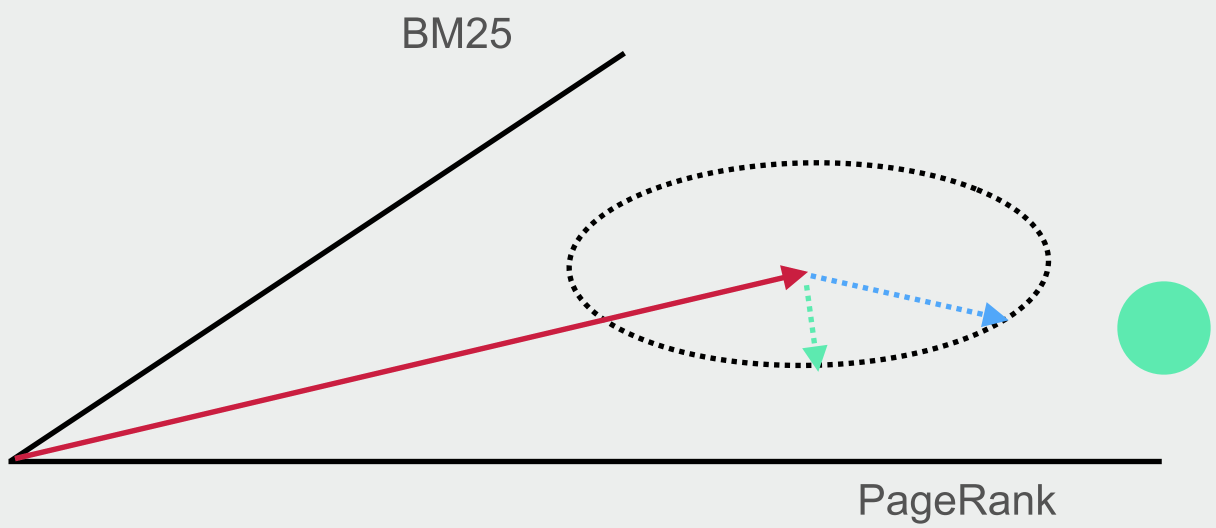

Mean winner (MGD-M)

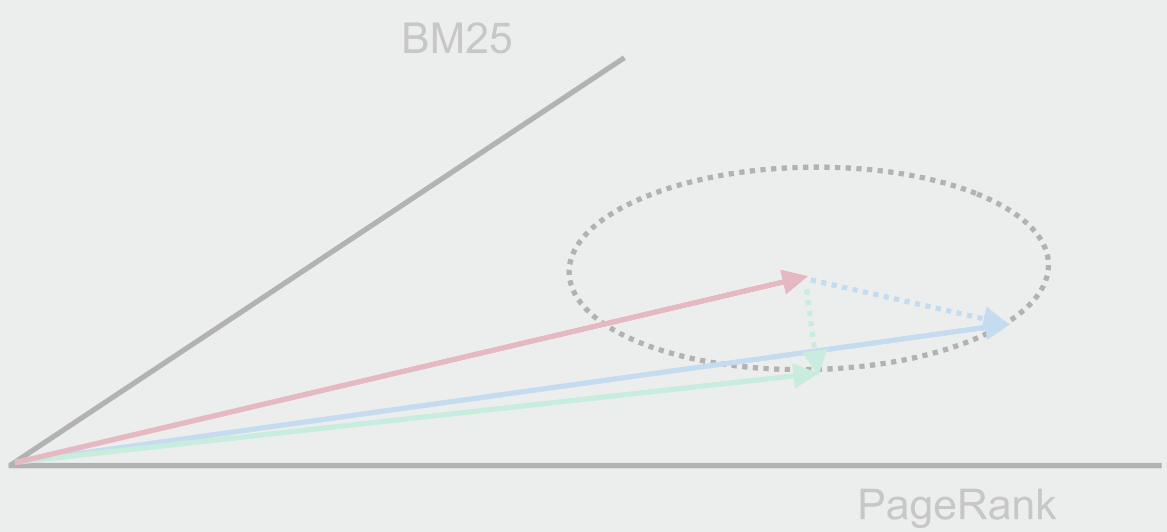

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

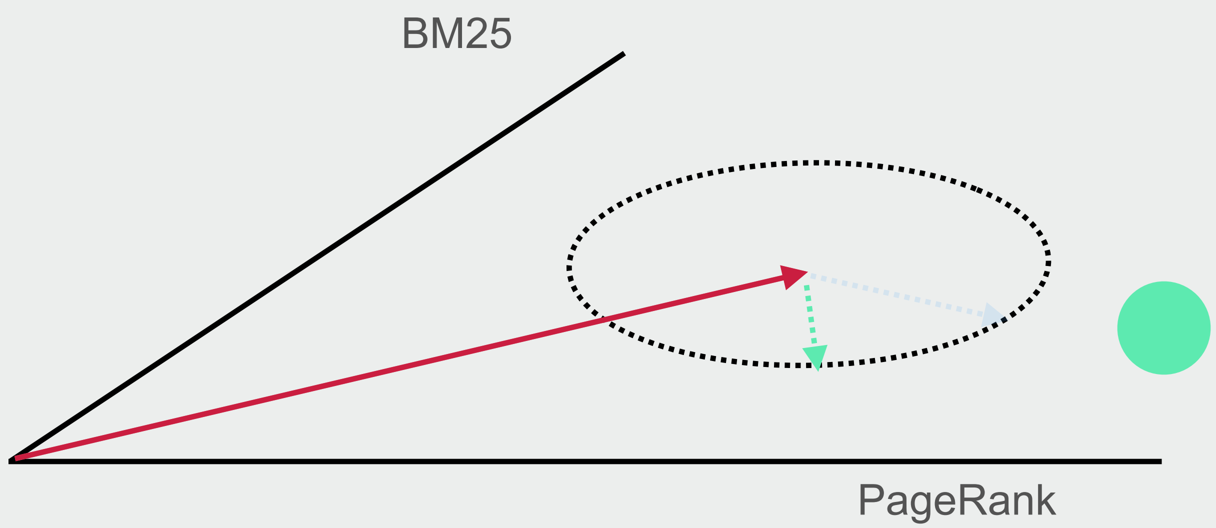

Mean winner (MGD-M)

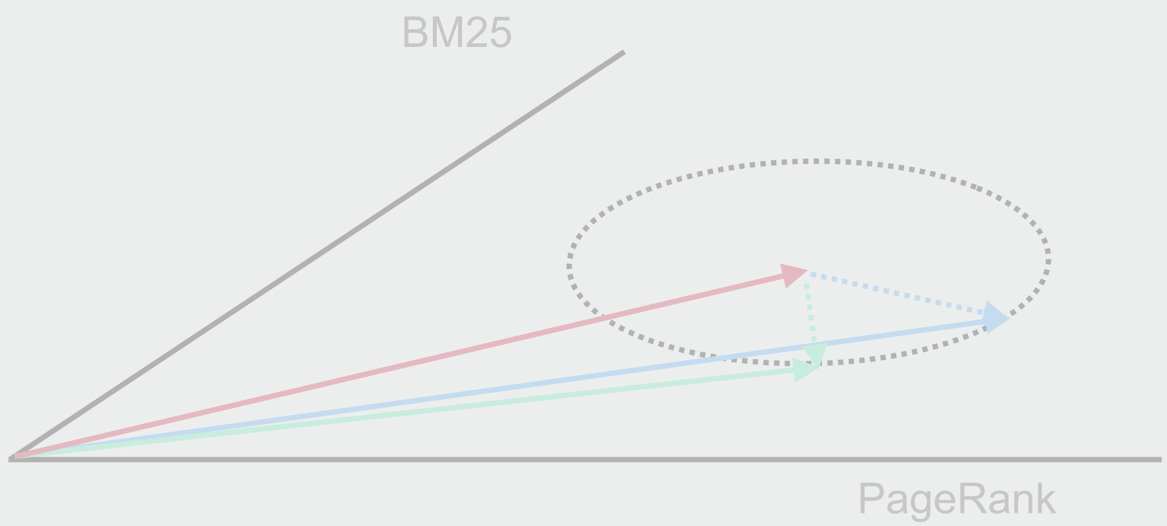

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

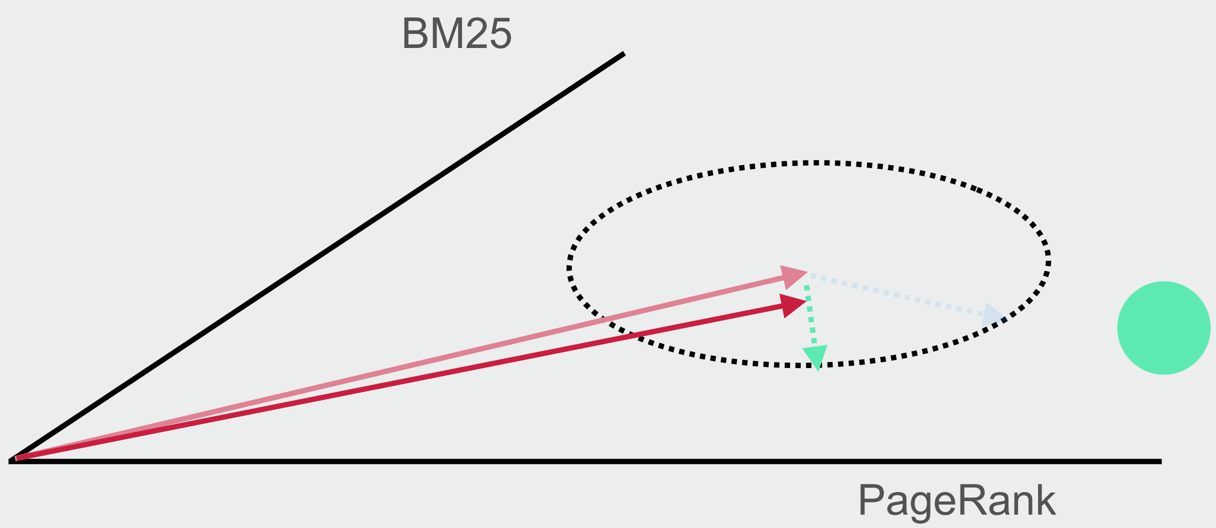

Mean winner (MGD-M)

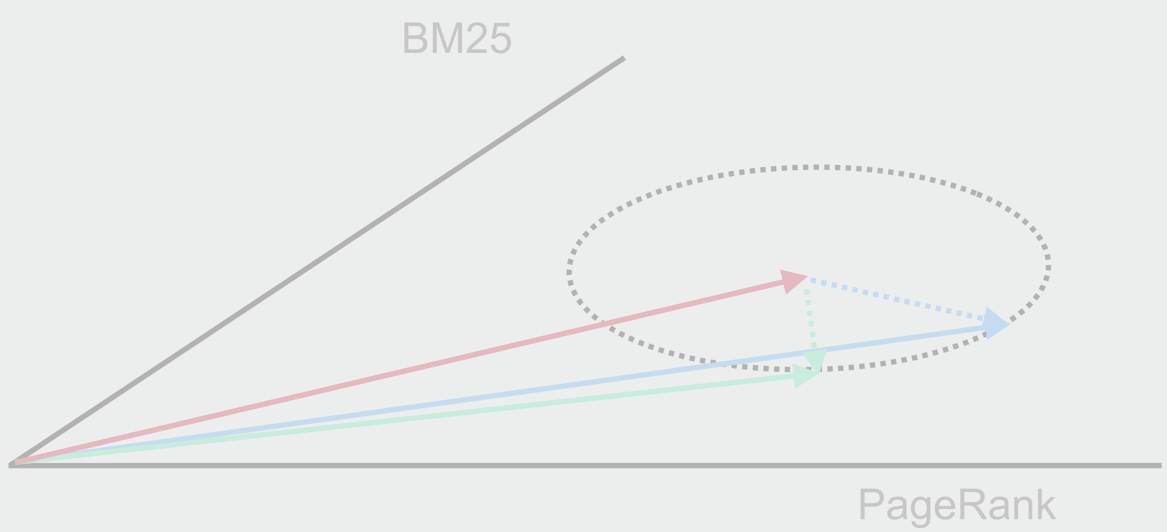

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

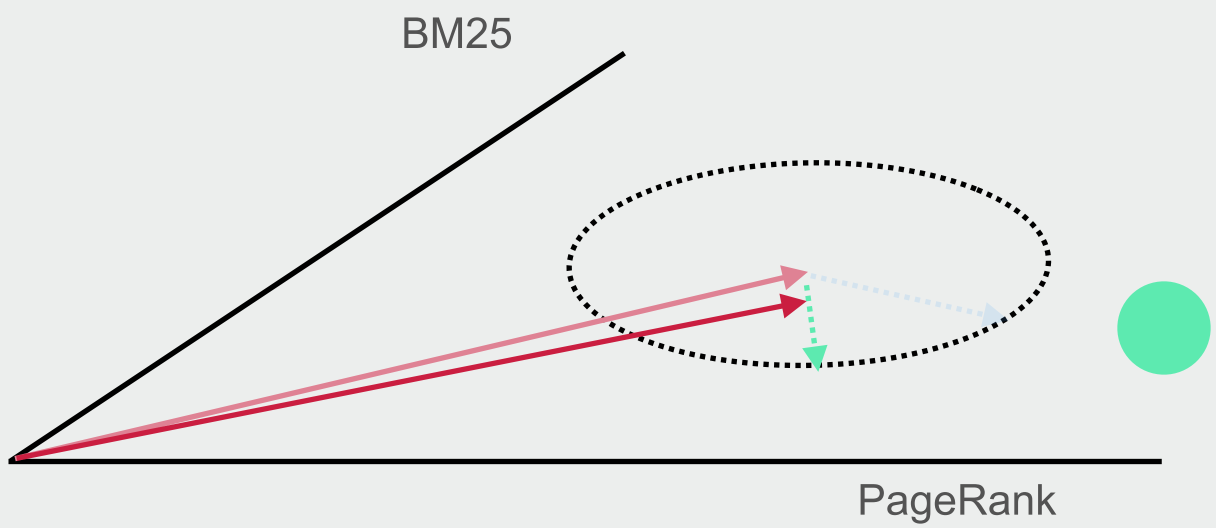

Mean winner (MGD-M)

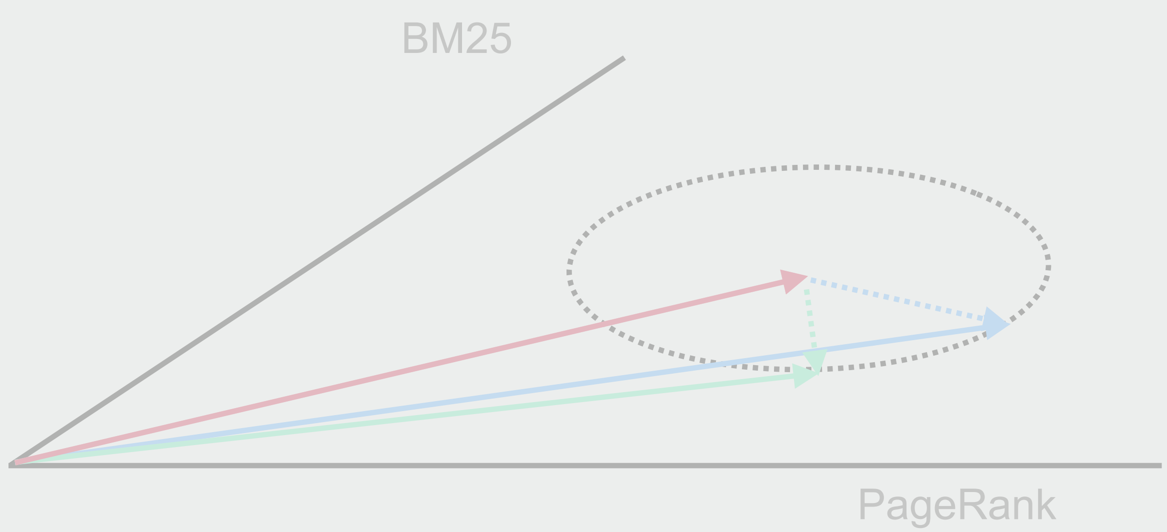

- Pick one of the winners
- Update with an alpha step

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

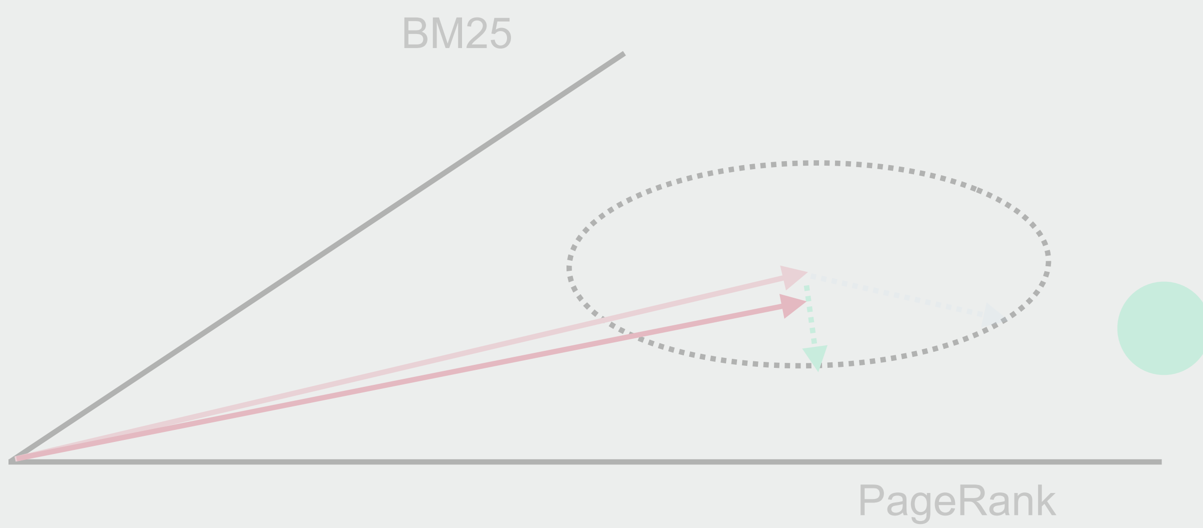

- Pick one of the winners
- Update with an alpha step

Mean winner (MGD-M)

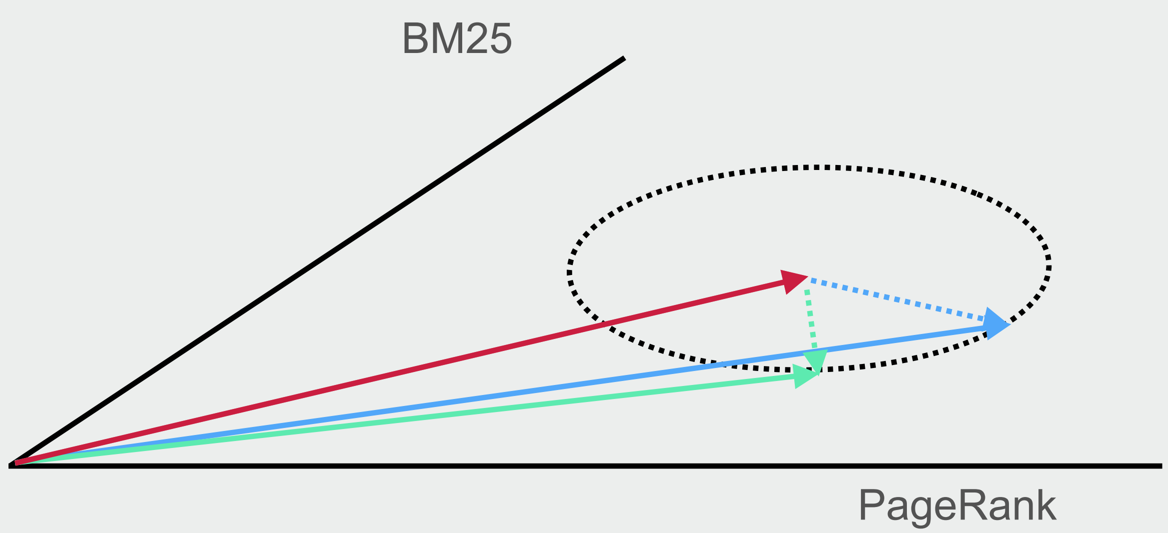

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

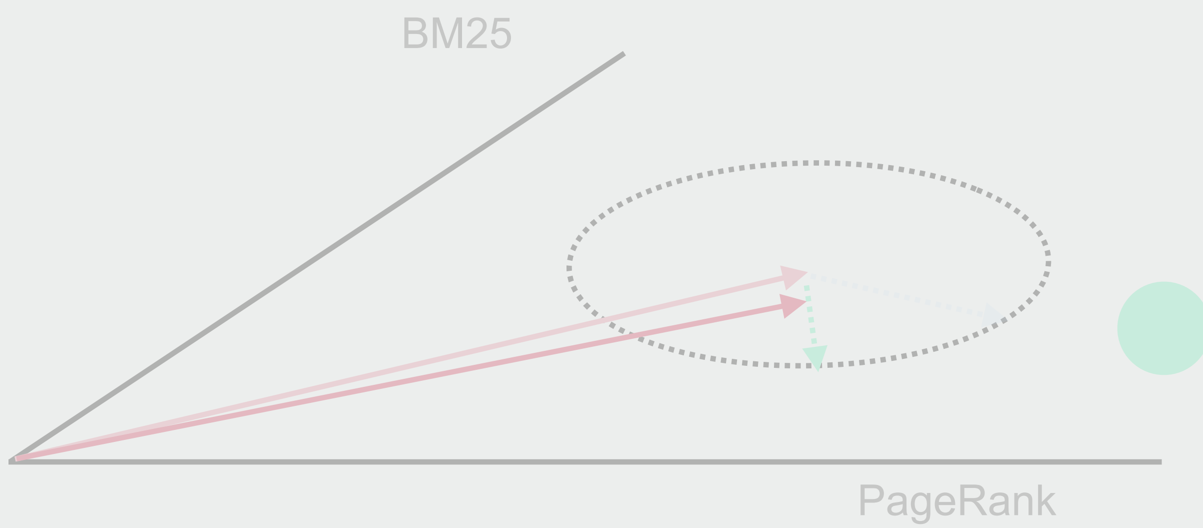

Mean winner (MGD-M)

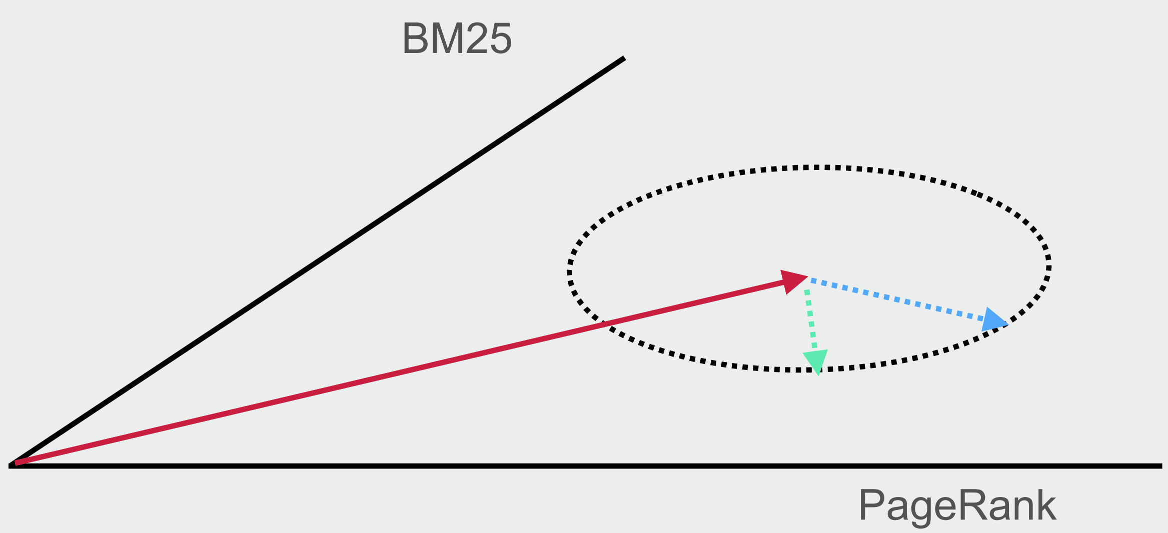

- Pick one of the winners
- Update with an alpha step

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

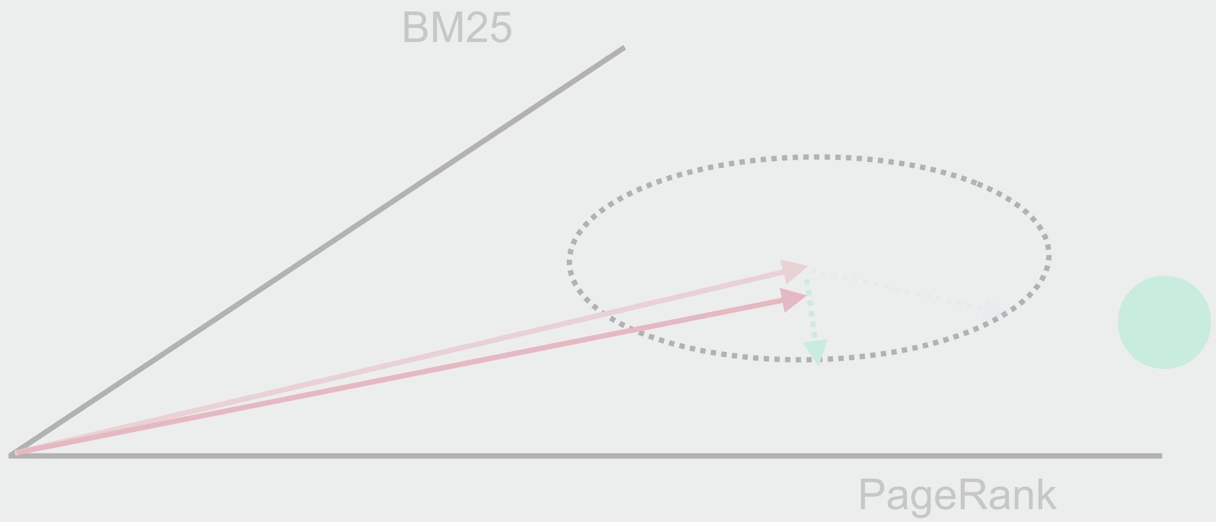

Mean winner (MGD-M)

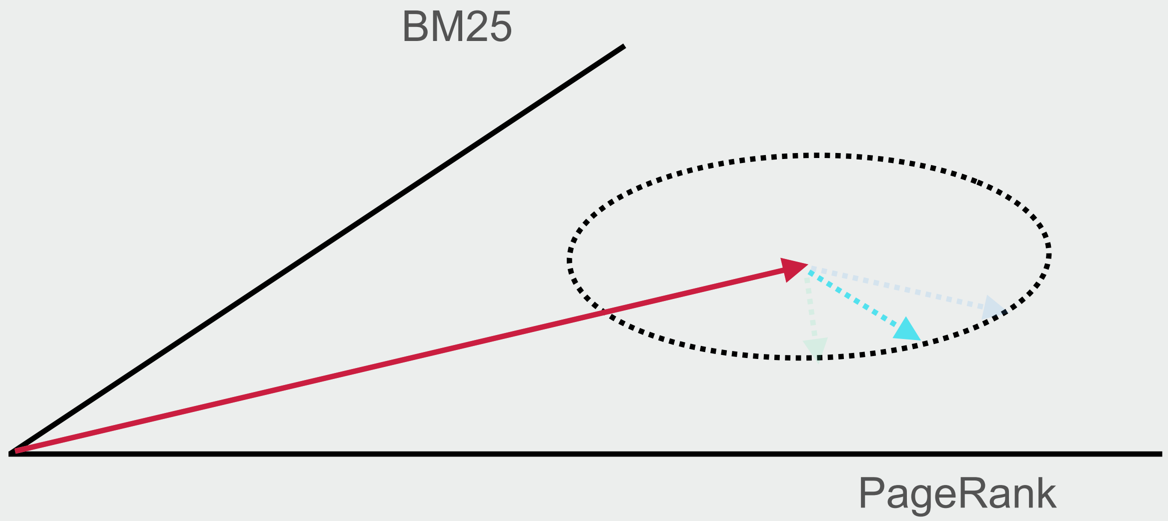

- Pick one of the winners
- Update with an alpha step

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

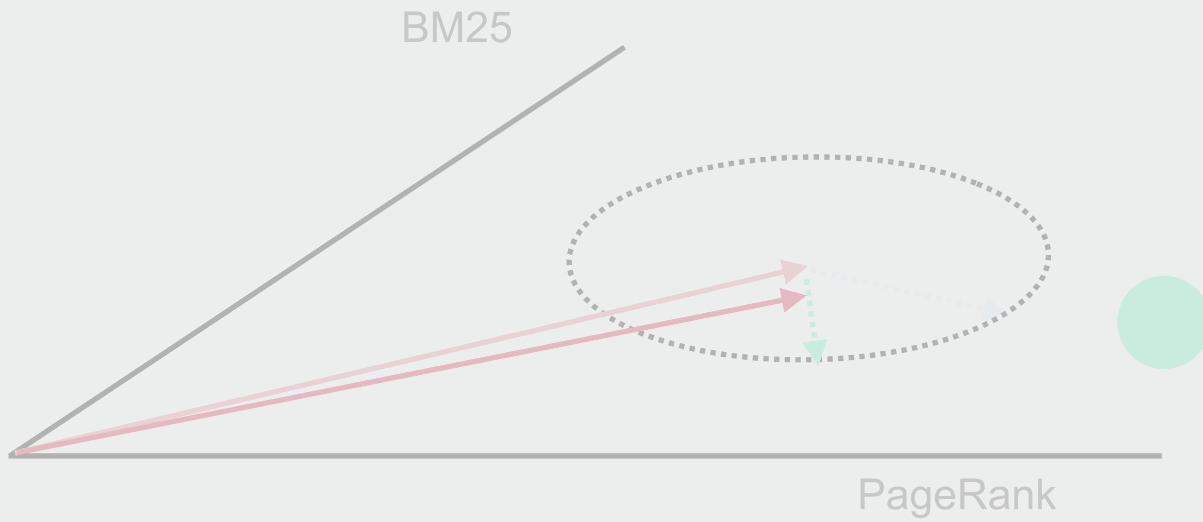

Mean winner (MGD-M)

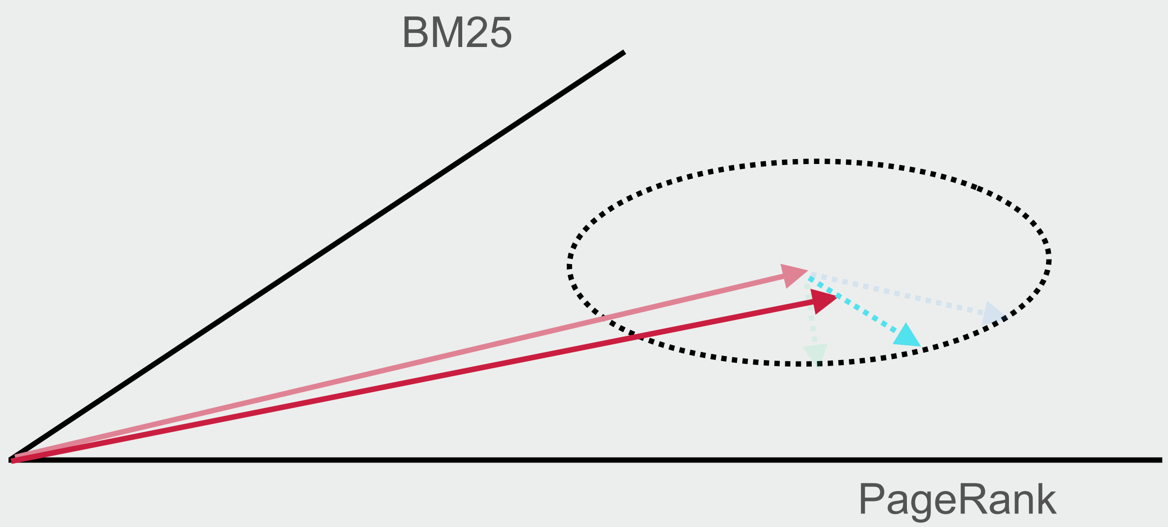

- Pick one of the winners
- Update with an alpha step

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

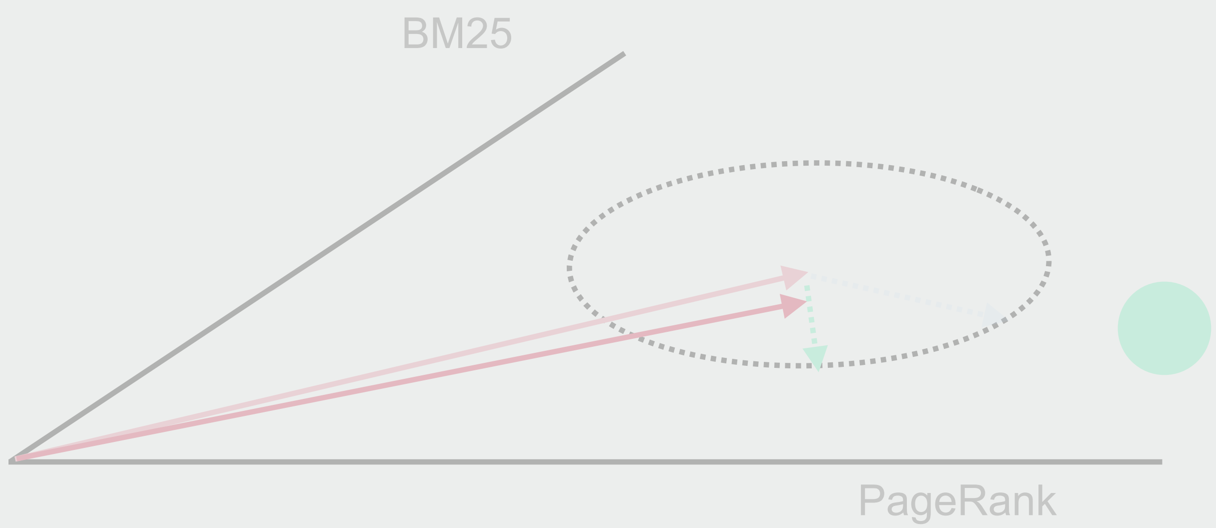

- Pick one of the winners
- Update with an alpha step

Mean winner (MGD-M)

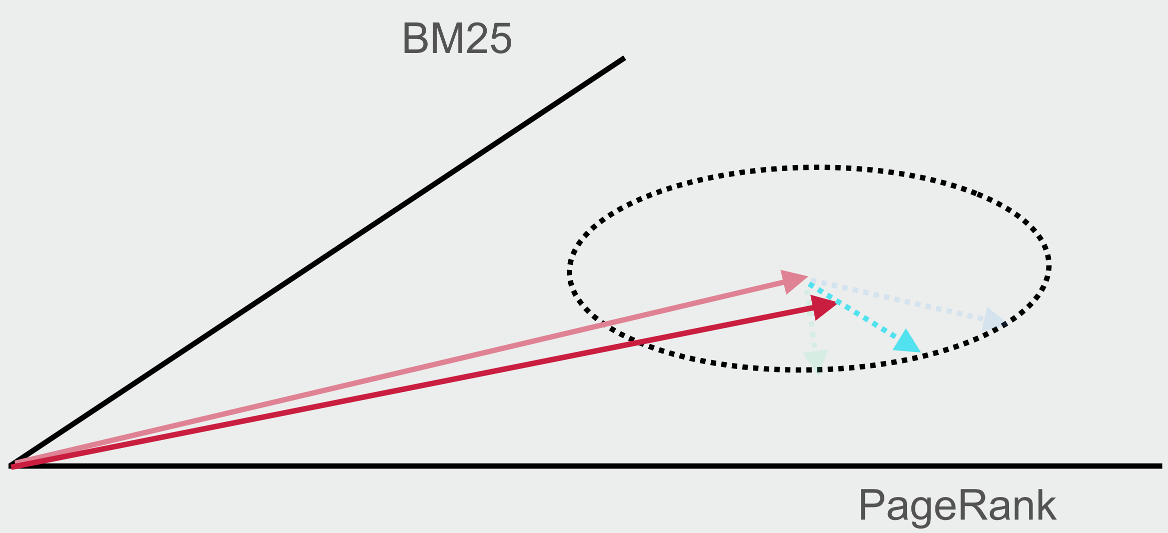

- Compute the mean of the winners
- Update with an alpha step

Multileave Gradient Descent (MGD)

Winner takes all (MGD-W)

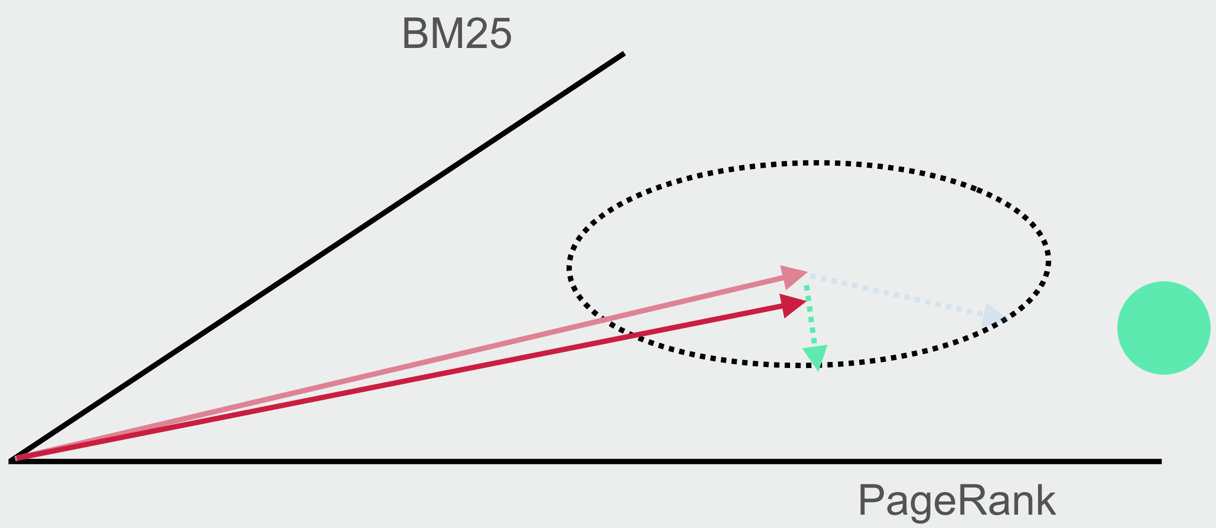

- Pick one of the winners
- Update with an alpha step

Mean winner (MGD-M)

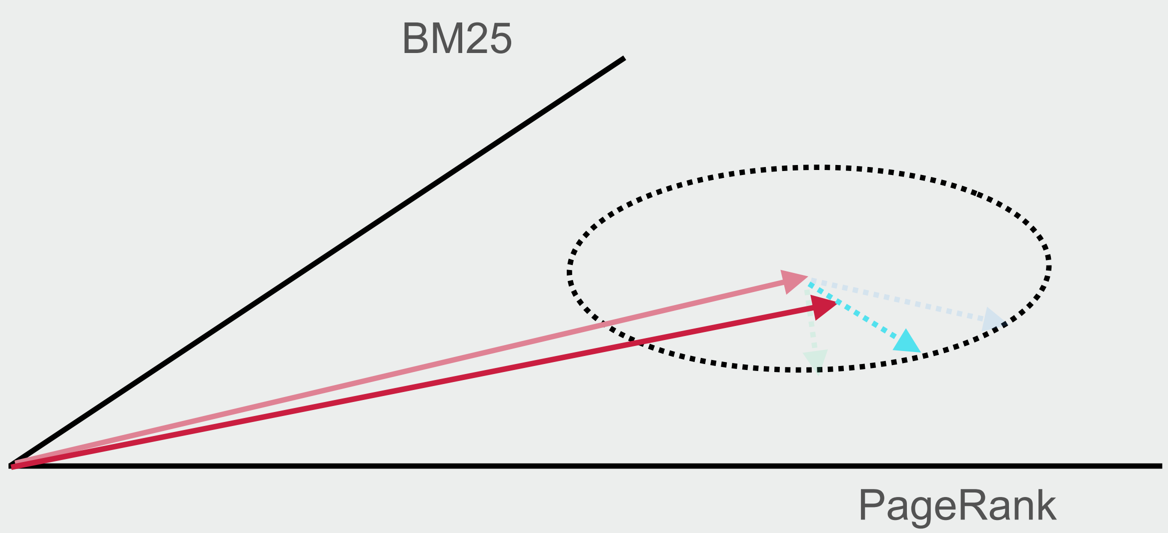

- Compute the mean of the winners
- Update with an alpha step

Experiments

Experiments

- Experimental run

Experiments

- Experimental run
    - queries sampled from L2R dataset

Experiments

- Experimental run
  - queries sampled from L2R dataset
  - clicks generated by cascade click model conditioned on relevance assessments

Experiments

- Experimental run
  - queries sampled from L2R dataset
  - clicks generated by cascade click model conditioned on relevance assessments
  - 25 repetitions * 5 folds

Experiments

- Experimental run
  - queries sampled from L2R dataset
  - clicks generated by cascade click model conditioned on relevance assessments
  - 25 repetitions * 5 folds
- 9 Datasets

Experiments

- Experimental run
  - queries sampled from L2R dataset
  - clicks generated by cascade click model conditioned on relevance assessments
  - 25 repetitions * 5 folds
- 9 Datasets
- NDCG

Results - **MGD-M vs MGD-W**

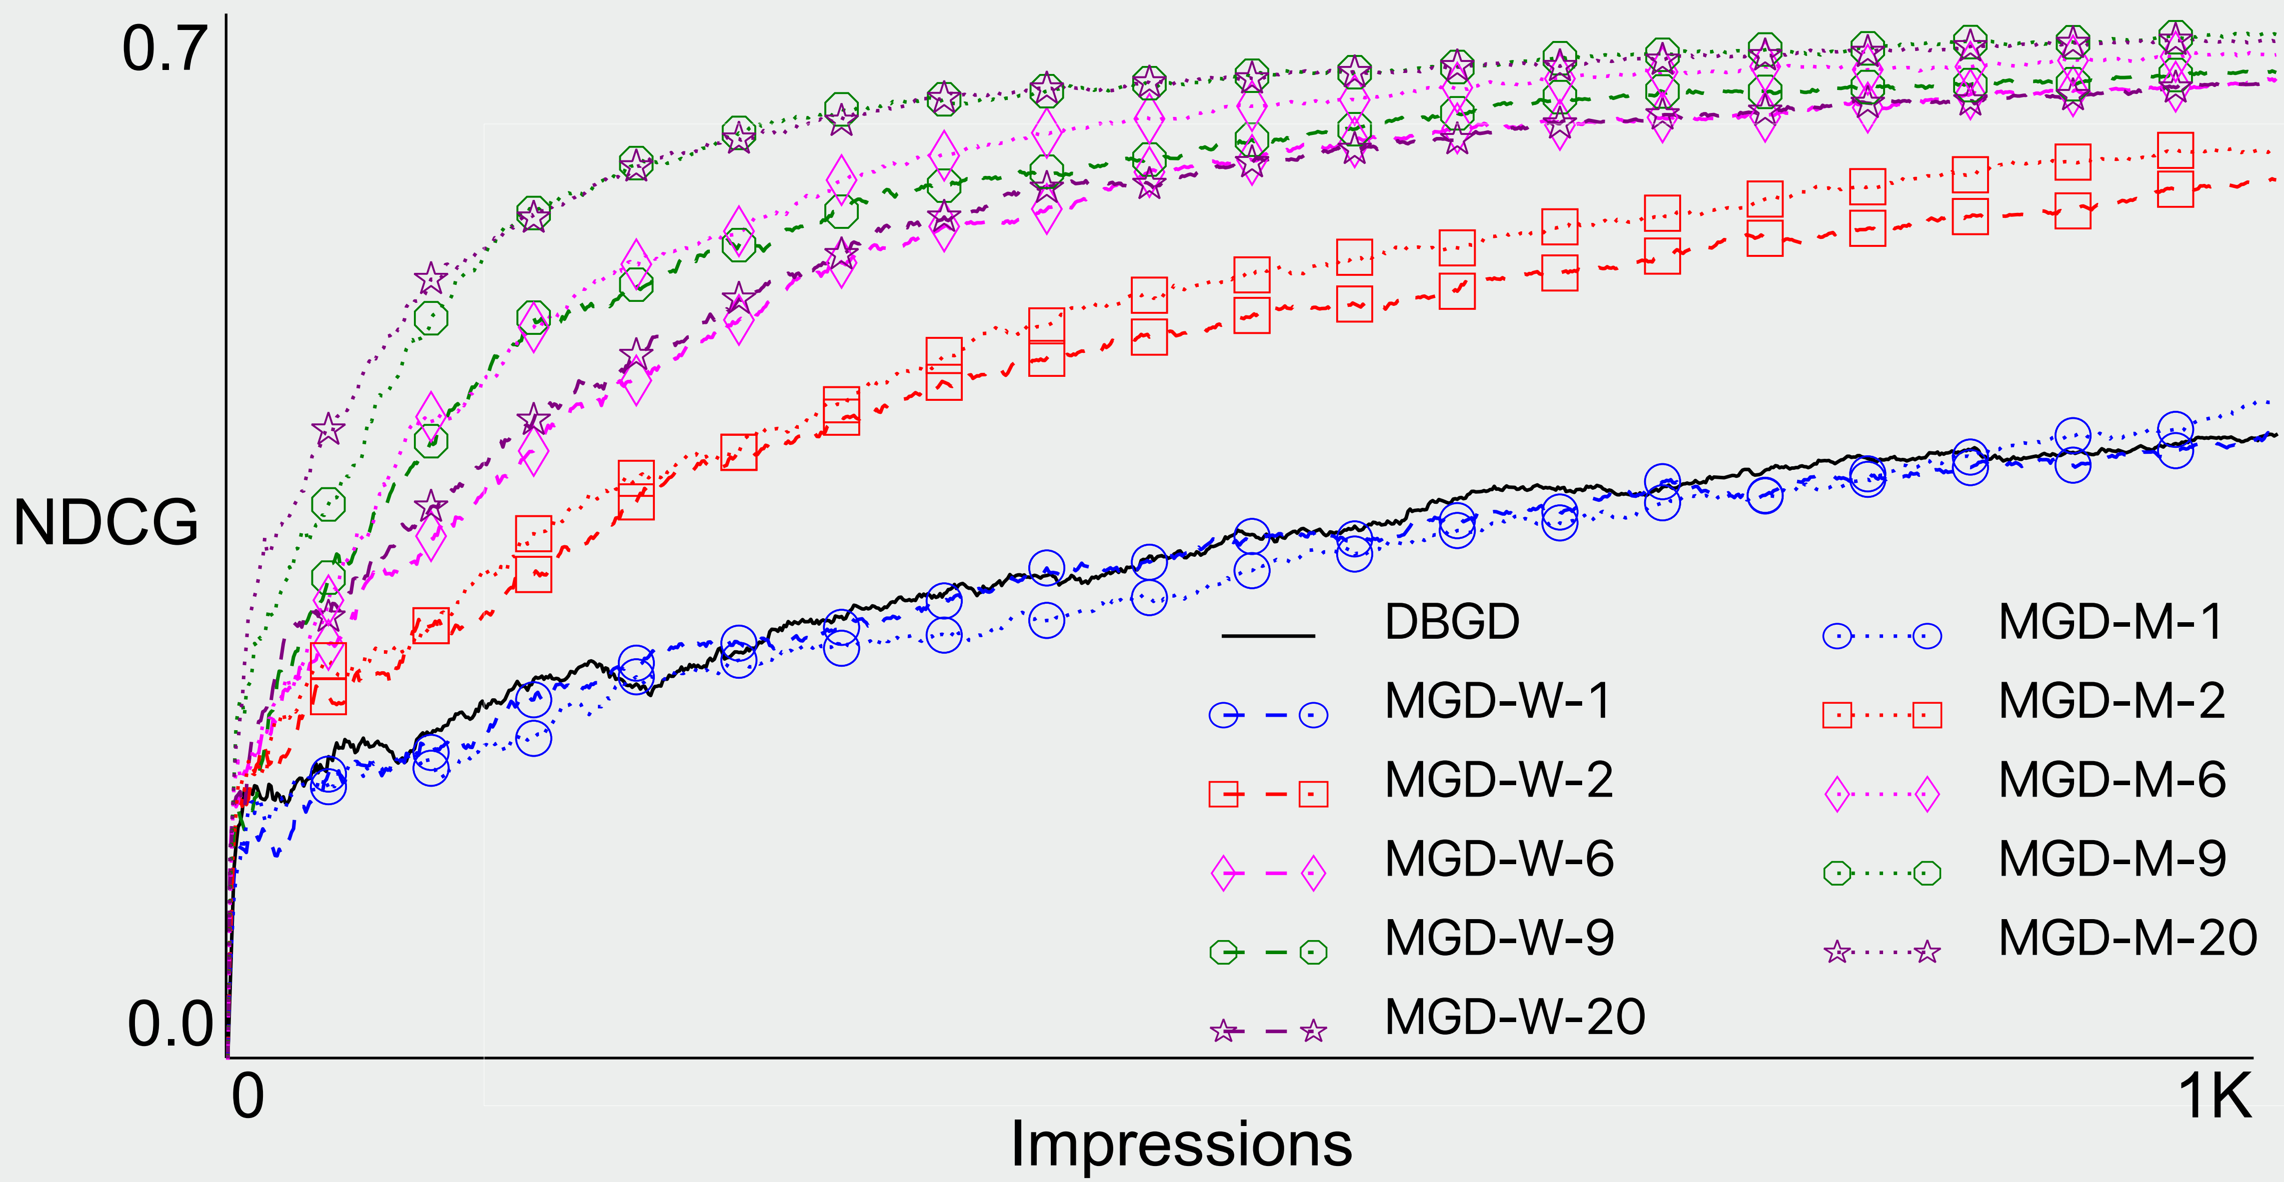

Results - **Long Run**

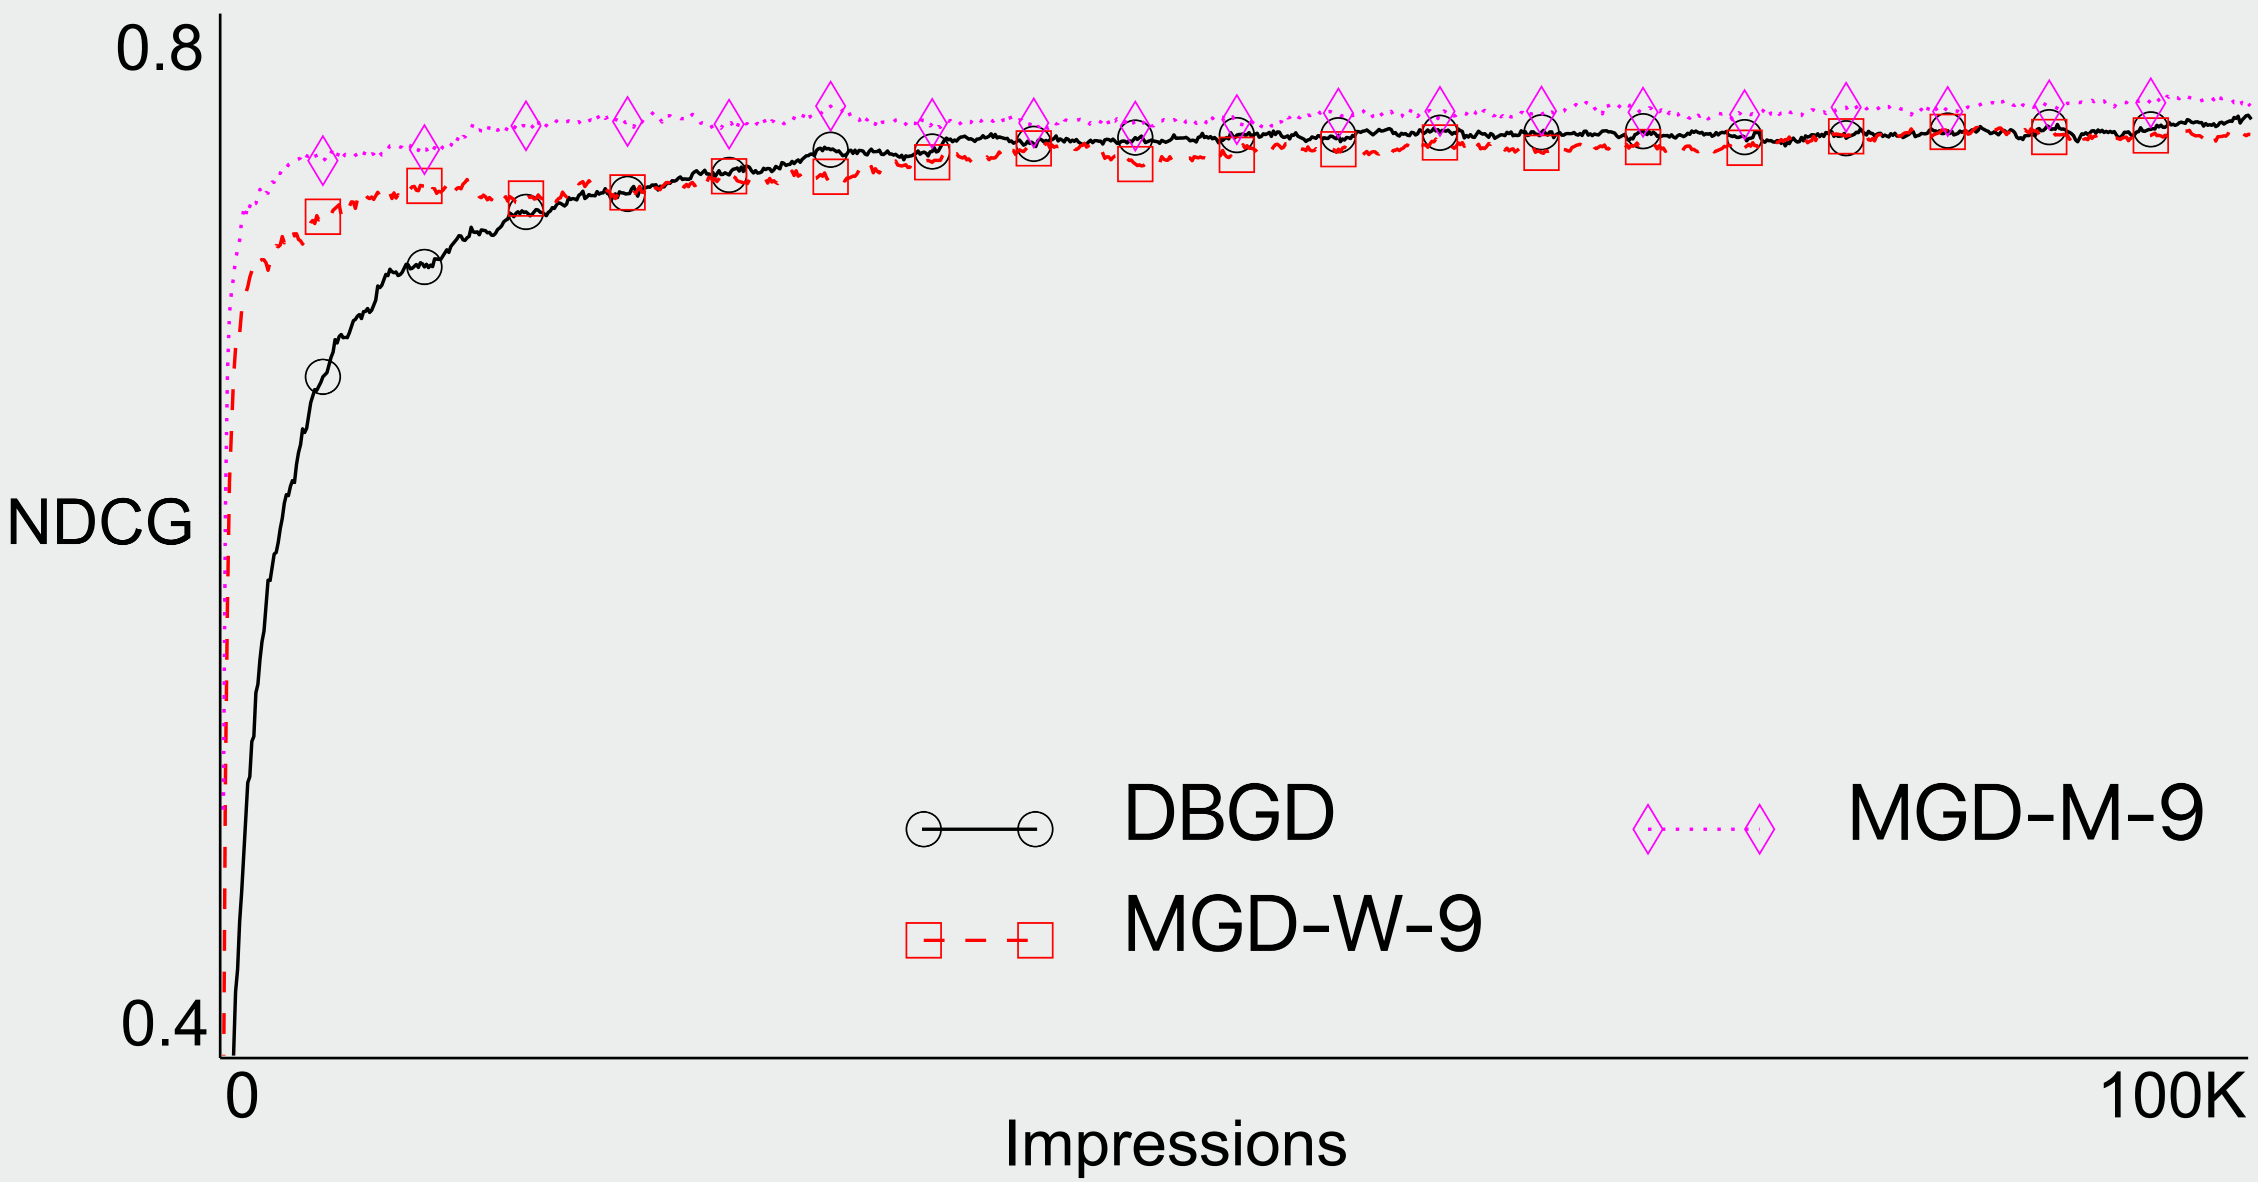

Conclusions

Conclusions

- MGD is an extension of DBGD

Conclusions

- MGD is an extension of DBGD
- Multileaving instead of Interleaving

Conclusions

- MGD is an extension of DBGD
- Multileaving instead of Interleaving
- Two update methods MGD-M and MGD-W

Conclusions

- MGD is 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

- MGD is 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
- Implication
  - Orders of magnitude less interaction data required with MGD
  - Search engines can adapt much faster

Thank you