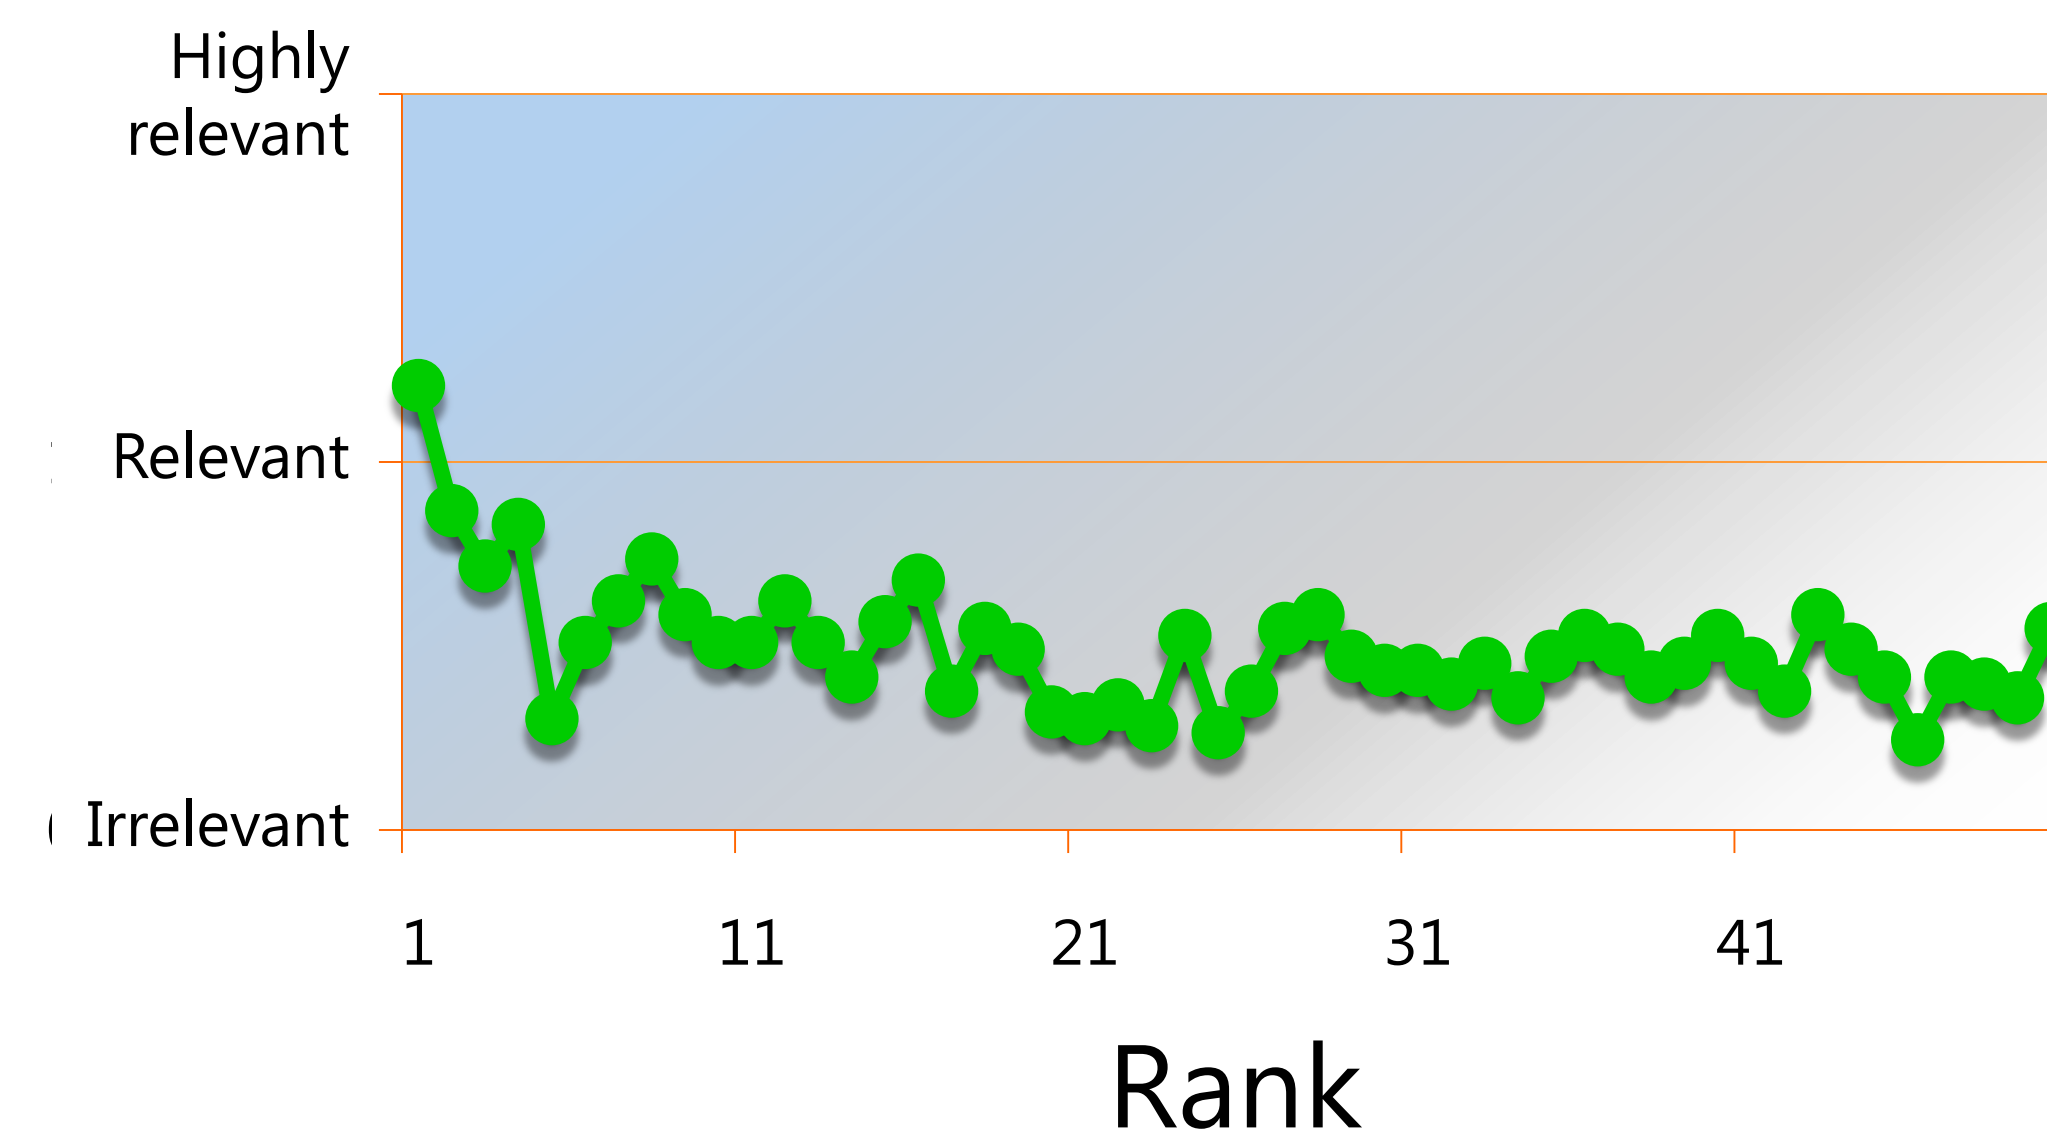


Characterizing the Value of Personalization

Jaime Teevan, Susan Dumais, and Eric Horvitz
{teevan,sdumais,horvitz}@microsoft.com

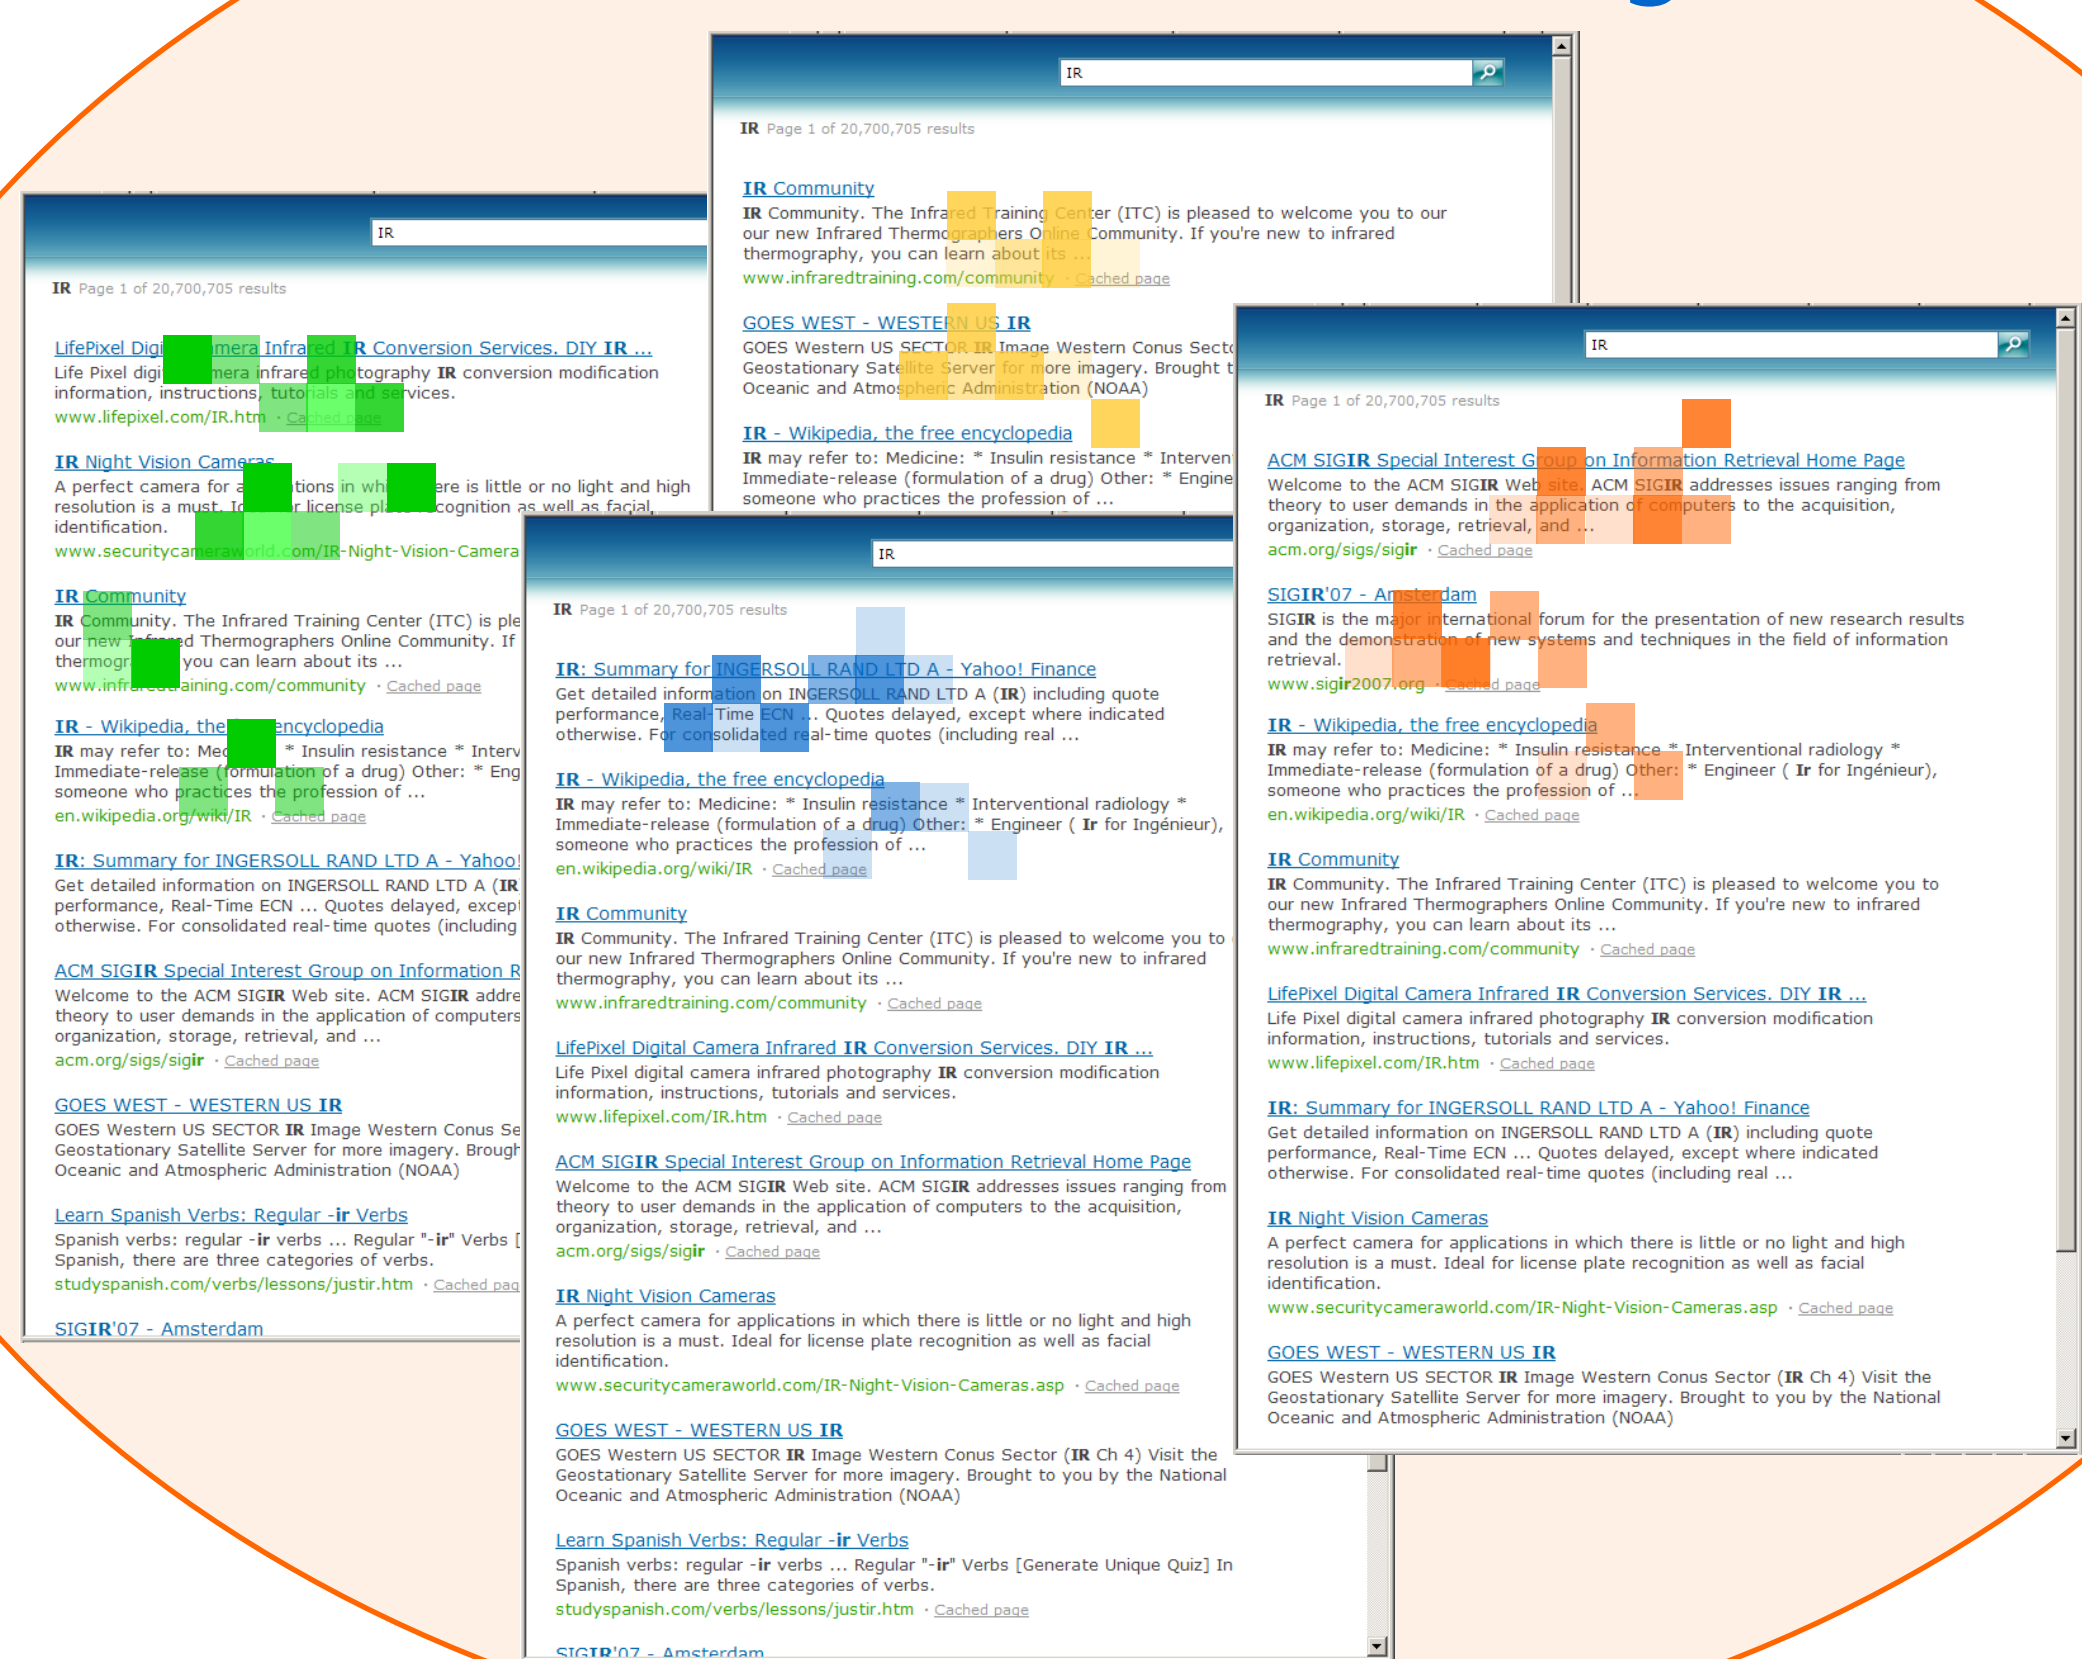
Relevant Results Ranked Low

Relevant results often ranked low (see graph)
Global ranking attempts to address this
But... Different users rate different results relevant
Study: People asked to rate relevance to *them* v. relevance to query



Understanding the Best Possible Ranking

Best Individual Rankings



The **best individual ranking** places
Highly relevant documents first
Relevant documents next
Irrelevant documents last

Can a global ranking satisfy everyone?

Result quality can be measured using DCG

$$DCG(i) = \begin{cases} G(1) & \text{if } i=1, \\ DCG(i-1) + G(i)/\log(i) & \text{otherwise.} \end{cases}$$

The best result list maximizes DCG
The **best group ranking** puts results with highest collective gain first

Best Group Ranking




Kendall-Tau distance measures list differences
KT distance of 0: Lists are the same
KT distance of 1: Lists in reverse order

$$KT(\text{web}, \text{best individual}) = 0.47$$

$$KT(\text{web}, \text{best group}) = 0.44$$

The best group ranking is significantly more like the Web ranking


The Value of Personalization



Value of personalizing

There is also a significant difference in list quality between
The **best individual ranking** and
The **best group ranking**

As group size grows, DCG for each individual drops (see graph)
But... No ranking seen by all can do better than group ranking
Gap is *value of personalizing* result list returned to the individual



Personalization unlikely to be valuable

Personalization likely to be very valuable

Future work: Characterize queries that personalize best
Value measured using:
Explicit judgments
Content match
Click behavior

Search engines do a good job of ranking results to maximize global happiness, but do not do a very good job for the specific individual.