

Optimal ranking of online search requests for long-term revenue maximization

Pierre L'Ecuyer



Patrick Maillé, Nicolás Stier-Moses, Bruno Tuffin



Technion, Israel, June 2016

Search engines

- ▶ Major role in the Internet economy
- ▶ most popular way to reach web pages
- ▶ 20 billion requests per month from US home and work computers only

Search engines

- ▶ Major role in the Internet economy
- ▶ most popular way to reach web pages
- ▶ 20 billion requests per month from US home and work computers only

For a given (set of) keyword(s), a search engine returns a ranked list of links: the **organic results**.

Organic results are supposed to be based on **relevance only**

Is this true?

Each engine has its own formula to measure (or estimate) relevance. May depend on user (IP address), location, etc.

Refine your search for specialized s work venge 52

Categories

All

Sporting Goods

Bicycles

Condition

see all

Used (1)

Price

\$ to \$ >>

Format

see all

All Listings (1)

Auction (0)

Buy It Now (1)

Delivery Options

see all

Free international shipping

Show only

see all

Returns accepted

Completed listings

Sold listings

More refinements...

Recently viewed items

All Listings

Auction

Buy It Now

Sort:

Best Match

View:

Grid

1 result for specialized s work ... [Follow this search](#)



SPECIALIZED S WORKS McLaren VENGE, DI2, SIZE 52cm, MARK CAVENDISH'S HTC BIKE

\$20,000.00

or Best Offer

From Ireland

More items related to specialized s work venge 52



2012 Specialized Venge Pro SRAM Red 52cm 10 Spd Road Race Bike-Carbon-Nice

\$2,800.00

or Best Offer

From United States

20 watching



2015 S-Works Venge 52cm Dura Ace Enve w/ Quarq

\$6,400.00

or Best Offer

From United States

Shop by
Department ▾

pierre's Store Deals Store Gift Cards Sell Help en français

Hello, pierre
Your Account ▾Try
Prime ▾

Books Advanced Search Browse Subjects New Releases Best Sellers Globe and Mail Best Sellers Best Books of the Month Children's Books Textbooks Today's Deals

1-16 of 83 results for Books : "Pierre l'ecuyer"

Sort by Relevance ▾

Show results for

< Any Category

Books

- Histoire et Sciences humaines (1)
- Sociology (1)
- Networking & Cloud Computing (10)
- Programming (30)
- Computers & Technology (55)
- Hardware (9)
- Algorithms Textbooks (6)
- Algorithms (14)
- Textbooks (56)

• See more

Refine by

Amazon Prime

 Prime

Shipping Option (What's this?)

Free Super Saver Shipping

Author

- Mohamed Cherkaoui
- Raymond Boudon
- Thierry Martin

New Releases

- Last 30 days
- Last 90 days (1)
- Coming Soon

Discount

- 10% Off or More (70)
- 25% Off or More (60)
- 50% Off or More (37)
- 70% Off or More (27)

[(Monte Carlo and Quasi-Monte Carlo Methods 2008)] [Author: Pierre L' Ecuyer] [Dec-2009] Dec 1
by Pierre L' Ecuyer

Hardcover

CDN\$ 354.92 used & new (1 offer)



Monte Carlo and Quasi-Monte Carlo Methods 2008 Nov 25 2009

by Pierre L' Ecuyer and Art G. Owen

Hardcover

CDN\$ 314.33 ~~CDN\$ 346.78~~ Prime

Only 1 left in stock - order soon.

More buying choices

CDN\$ 122.44 used & new (20 offers)

Eligible for FREE Super Saver Shipping.

Paperback

CDN\$ 312.95 ~~CDN\$ 346.78~~ Prime

Not in stock; order now and we'll deliver when available

More buying choices

CDN\$ 299.62 used & new (9 offers)



Modeling Uncertainty: An Examination of Stochastic Theory, Methods, and Applications Jan 31 2002

by Moshe Dror and Pierre L'Ecuyer

Hardcover

CDN\$ 538.65 ~~CDN\$ 542.13~~ Prime

Not in stock; order now and we'll deliver when available

More buying choices

CDN\$ 403.28 used & new (18 offers)

Eligible for FREE Super Saver Shipping.

Perfect Paperback

CDN\$ 524.64 used & new (2 offers)



Computational Science and Its Applications - ICCSA 2003: International Conference, Montreal, Canada, May 18-21... Jul 15 2003

by Vipin Kumar and Marina L. Gavrilova

Paperback

CDN\$ 248.91 ~~CDN\$ 250.52~~ Prime

Not in stock; order now and we'll deliver when available

Eligible for FREE Super Saver Shipping.

How are items ranked? Relevance vs expected revenue?

Web News Videos Images Shopping More Search tools



About 993,000 results (0.31 seconds)

Justin Trudeau vs. Patrick Brazeau Boxing Match p1 ...



www.youtube.com/watch?v=tb-sqWJyWqA

Mar 31, 2012 - Uploaded by CIRaport ca
Skinny Indian? What part of Brazeau is skinny? Secondly please take a good look at Patrick Brazeau, if he is ...

Justin Trudeau -- Patrick Brazeau Charity Boxing Match ...



www.youtube.com/watch?v=XuSpZ3_5pTc

Apr 1, 2012 - Uploaded by FactPointVideo
Justin Trudeau -- Patrick Brazeau Charity Boxing Match -- English ... they went under, I love watching the ...

Trudeau vs Brazeau BOXING FOR CANCER (Full fight ...



<https://www.youtube.com/watch?v=fYfWZMhaLE>

Apr 1, 2012 - Uploaded by Reasoning Conspiracy
... link <http://www.sunnewsnetwork.ca/video/1540728713001>
Watch the ... Patrick Brazeau and Justin ...

Meet Justin Trudeau: The Tattooed, Boxing New Canadian ...



www.nbcnews.com/.../video/meet-justin-trudeau--the-t...

16 hours ago
Meet Justin Trudeau: The Tattooed, Boxing New Canadian Prime Minister ... Stunning Video Shows Qantas ...



Justin Trudeau vs Patrick Brazeau Le combat de boxe ...



www.youtube.com/watch?v=go7nf6xwfyA

Mar 31, 2012 - Uploaded by LP - Mma Nouvelles
Justin Trudeau vs Patrick Brazeau Le combat de boxe Justin Trudeau -- Patrick Brazeau Charity Boxing ...

Justin Trudeau Boxing Match Weigh In With Senator Patrick ...



www.youtube.com/watch?v=M0e9D2rBN9w

Mar 28, 2012 - Uploaded by althiaraj
Liberal MP Justin Trudeau meets his competitor Conservative Senator Patrick Brazeau for a weigh in before ...


Justin Trudeau Boxing - YouTube



www.youtube.com/watch?v=JLQWA35er3g

Auo 6, 2015 - Uploaded by Alex Guibord


Durée ▾ Date ▾ Résolution ▾ Source ▾ Prix ▾ Filtre adulte: **Modéré** ▾



Justin Trudeau Beats Patrick Brazeau In Charity Boxing Match

▶ 01:29 - HD - 2012-03-31 - Par eCanadaNow Breaking News

▶ YouTube - 6 000+ vues



Justin Trudeau Boxing Weigh-In With Senator Patrick ...

▶ 04:20 - 2012-03-28 - Par | By


The Huffington Post



Justin Trudeau vs. Patrick Brazeau Boxing Match p1

▶ 12:03 - HD - 2012-03-31 - Par CIReport.ca

▶ YouTube - 22 000+ vues



Justin Trudeau -- Patrick Brazeau Charity Boxing Match ...

▶ 16:14 - HD - 2012-04-01 - Par FactPointVideo


▶ YouTube - 101 000+ vues



Trudeau Admits Wife Thought He Might Lose Boxing Match

▶ 00:26 - 2012-10-16 - Par HuffPost Canada Studios

Aol.On - 10 000+ vues



Trudeau vs Brazeau BOXING FOR CANCER (Full fight)

▶ 11:26 - 2012-04-01 - Par Reasoning Conspiracy

▶ YouTube - 25 000+ vues



Justin Trudeau gets some boxing practice in ahead of first ...

▶ 01:35 - 2015-08-06 - Par Global News


YAHOO! NEWS



Justin Trudeau and Patrick Brazeau - bet after boxing ...

▶ 09:50 - 2015-07-21 - Par Marksminstrel

Dailymotion - vues 16 fois



Justin Trudeau Works Out in Boxing Ring Ahead of Debate

▶ 01:12 - 2015-08-06 - Par Canadian Press

Aol.On

 Commentaires



Web News **Videos** Images Maps More ▾ Search tools

About 11,400 results (0.18 seconds)

Video: Evo Morales caught kneeling player in groin during ...



www.telegraph.co.uk > ... > South America > Bolivia ▾

Oct 5, 2010

Evo Morales, the Bolivian president, has been caught on camera kneeling an opposing player in the groin ...



Bolivian president Morales plays football for UN - YouTube



www.youtube.com/watch?v=uHIT3T9o6GI

Sep 27, 2013 - Uploaded by AFP news agency

President of Bolivia Evo Morales during a soccer match at a park near the UN on the sidelines of the 68th ...

Bolivian President kicks man in bollocks (in football match ...



<https://www.youtube.com/watch?v=5XkbNUix5Bo>

Oct 6, 2010 - Uploaded by TrashTVChannel

Bolivian president Evo Morales caught on camera knee-ing opponent (Daniel Gustavo Cartagenam) in the ...

Ahmadinejad plays football with Bolivian President Evo ...



www.youtube.com/watch?v=ragkI2EPqww ▾

Oct 27, 2010 - Uploaded by sherlock72

During a visit of Bolivian President Evo Morales to Iran , he played a friendly football game with Iranian ...

Bolivian president Morales plays football for UN - Video ...



www.dailymotion.com/video/x158t77_bolivian-presid...

Sep 27, 2013

President of Bolivia Evo Morales during a soccer match at a park near the UN on the sidelines of the 68th ...



Bolivia: Watch Evo Morales play against ex-Real Madrid ...



<https://www.youtube.com/watch?v=v9oz0vvD2U4>

Oct 8, 2014 - Uploaded by Ruptly TV

Video ID: 20141008-057 M/S Evo Morales and former Real Madrid ... Bolivian President Evo Morales ...



Bolivian President Morales plays in friendly football match ...





www.youtube.com/watch?v=PUxWjVVGLOI

Jul 31, 2015 - Uploaded by AP Archive

Bolivian President Evo Morales walking onto pitch in Cadiz stadium 2 ... Bolivian President Morales plays ...





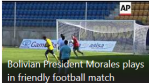








 

Sign in  

Web Images **Videos** Maps News Explore

Evo Morales **Debates De Evo Morales** **Noe Morales Mix** **Evo Morales Documentary** **Evo Morales Real News** **Evo Morales History** **Evo Morales Presidente De Bolivia** **Evo Morales Biography**

Length ▾ Date ▾ Resolution ▾ Source ▾ Price ▾ SafeSearch: **Moderate** ▾

 <p>Bolivian president Morales plays football for UN</p> <p>▶ 00:40 · 2013-09-27 · By AFP news agency</p> <p> YouTube · 1,000+ views</p>	 <p>Bolivian President Evo Morales signed by football club</p> <p>▶ 01:39</p> <p> BBC</p>	 <p>Bolivian President Morales plays in friendly football match</p> <p>▶ 01:39</p> <p> WN</p>	 <p>Bolivian president seeks Carter's advice</p> <p>▶ 02:35</p> <p> CNN</p>
 <p>Bolivian president Morales plays football for UN</p> <p>▶ 00:40 · 2013-09-27 · By Agence France-Presse (AFP)</p> <p> DailyMotion · 26 views</p>	 <p>Evo Morales caught kneeling player in groin during football ...</p> <p>▶ 00:27 · 2010-10-05 · By Our Foreign Staff</p> <p>The Telegraph</p>	 <p>Bolivian football club sign new midfielder: President Evo ...</p> <p>▶ 02:27</p> <p>Mail Online</p>	 <p>Morales the McQueen</p> <p>Being good 80% is easier. It's only being a bad person. We'd play for 20-30 minutes</p> <p>Football Team Signs Bolivian President</p> <p>▶ 01:18 · 2014-05-20 · By BBC World News</p> <p>On Aol · 424 views</p>

Google+ Search Images Maps Play YouTube News Gmail More - Sign in 

Google

Web Images Videos News Maps Books

About 11,000,000 results

Any country

Country: Canada

Any time

Past hour
Past 24 hours
Past week
Past month
Past year

All results

Verbatim



Barack Obama playing Basketball Game. AMAZING FOOTAGE

<https://www.youtube.com/watch?v=0OIdGQQ0L8>

Images for barack obama basketball video



Barack Obama's basketball fail - YouTube

<https://www.youtube.com/watch?v=gmTfKPx1Cug>
1 Apr 2013 - 1 min - Uploaded by The Telegraph



putin judo video

Web **Videos** Images News Maps More ▾ Search tools

About 54,600 results (0.24 seconds)

All YouTube:

'Judo Knight' Putin shows off martial arts skills in wrestling ...



<https://www.youtube.com/watch?v=0lxMglj8LuM>

Dec 23, 2010 - Uploaded by RT

Prime Minister Vladimir Putin has taken another step to boost his charisma by ... Bodyguard face-off video ...

Putin's Judo - YouTube



https://www.youtube.com/watch?v=dYIDQ_t_HQ

Apr 5, 2012 - Uploaded by strawisinthetack

Vladimir Putin has a black belt in the martial art and has competed at a high ... One day...a dream face off ...

Let's Learn Judo with Vladimir Putin (2008) - YouTube



<https://www.youtube.com/watch?v=162myM2fJE>

Nov 7, 2012 - Uploaded by Putinery

... Russian President Vladimir Putin shares his favourite judo moves. ... Let's Learn Judo with Vladimir ...

Putin may join Russian Olympic judo team - YouTube



<https://www.youtube.com/watch?v=2c3wzbyikw>

Dec 20, 2009 - Uploaded by RT

The Russian Judo Olympic team may have a new member. ... Russia's Got Talent: Extra video of Putin ...

Putin practices judo - No comment - YouTube



<https://www.youtube.com/watch?v=V5z5UAtCvgE>

Dec 21, 2009 - Uploaded by No Comment TV

<http://www.euronews.com/nocomment/> Vladimir Putin announcing he could help the Russian Judo Team if ...

Putin Judo - YouTube



https://www.youtube.com/watch?v=1dXwu2MXR_s

Jun 14, 2006 - Uploaded by timourkh

Russian President Vladimir Putin has a black belt in Judo.

Putin & Co to teach you judo - YouTube



<https://www.youtube.com/watch?v=C7yjCaCQP08>

Oct 7, 2008 - Uploaded by RT

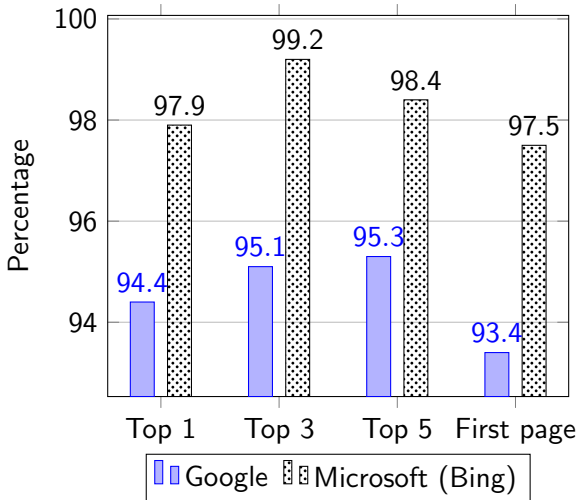
Do search engines return biased results?

Comparison between Google, Bing, and Blekko (Wright, 2012):

- ▶ Microsoft content is 26 times more likely to be displayed on the first page of Bing than on any of the two other search engines
- ▶ Google content appears 17 times more often on the first page of a Google search than on the other search engines

Search engines do favor their own content

Do search engines return biased results?



Percentage of Google or Bing search results with **own content not ranked similarly** by any rival search engine (Wright, 2012).

Search Neutrality (relevance only)

Some say search engines should be considered as a **public utility**.

Idea of **search neutrality**: All content having equivalent relevance should have the same chance of being displayed. Content of higher relevance should never be displayed in worst position.

More fair, better for users and for economy, encourages quality, etc.

What is the precise definition of “**relevance**”? Not addressed here ...

Debate: Should neutrality be imposed by law? Pros and cons.

Regulatory intervention: The European Commission, is progressing toward an antitrust settlement deal with Google.

“Google must be even-handed. It must hold all services, including its own, to exactly the same standards, using exactly the same crawling, indexing, ranking, display, and penalty algorithms.”

In general: trade-off in the rankings

From the viewpoint of the SE: Tradeoff between

- ▶ **relevance** (long term profit)

versus

- ▶ **expected revenue** (short term profit)

Better relevance brings more customers in the long term because it builds reputation.

What if the provider wants to **optimize its long-term profit**?

Simple model of search requests

Request: random vector $Y = (M, R_1, G_1, \dots, R_M, G_M)$ where
 M = number of pages (or items) that match the request, $M \leq m_0$;
 $R_i \in [0, 1]$: measure of **relevance** of item i ;
 $G_i \in [0, K]$: **expected revenue** (direct or indirect) from item i .
 has a prob. distribution over $\Omega \subseteq \mathbb{N} \times ([0, 1] \times [0, K])^{m_0}$.
 Can be discrete or continuous.

$y = (m, r_1, g_1, \dots, r_m, g_m)$ denotes a realization of Y .

$c_{i,j}(y) = \mathbb{P}[\text{click page } i \text{ if in position } j] = \text{click-through rate (CTR)}$.

Assumed \nearrow in r_i and \searrow in j . Example: $c_{i,j}(y) = \theta_j \psi(r_i)$

Simple model of search requests

Request: random vector $Y = (M, R_1, G_1, \dots, R_M, G_M)$ where
 $M =$ number of pages (or items) that match the request, $M \leq m_0$;
 $R_i \in [0, 1]$: measure of **relevance** of item i ;
 $G_i \in [0, K]$: **expected revenue** (direct or indirect) from item i .
 has a prob. distribution over $\Omega \subseteq \mathbb{N} \times ([0, 1] \times [0, K])^{m_0}$.
 Can be discrete or continuous.

$y = (m, r_1, g_1, \dots, r_m, g_m)$ denotes a realization of Y .

$c_{i,j}(y) = \mathbb{P}[\text{click page } i \text{ if in position } j] = \text{click-through rate (CTR)}$.

Assumed \nearrow in r_i and \searrow in j . Example: $c_{i,j}(y) = \theta_j \psi(r_i)$

Decision (ranking) for any request y : Permutation $\pi = (\pi(1), \dots, \pi(m))$ of the m matching pages. $j = \pi(i) =$ position of i .

Local relevance and **local revenue** for y and π :

$$r(\pi, y) = \sum_{i=1}^m c_{i,\pi(i)}(y) r_i, \quad g(\pi, y) = \sum_{i=1}^m c_{i,\pi(i)}(y) g_i.$$

Deterministic stationary ranking policy μ

It assigns a permutation $\pi = \mu(y) \in \Pi_m$ to each $y \in \Omega$.

Long-term **expected relevance** per request (reputation of the provider) and **expected revenue** per request (from the organic links), for given μ :

$$\begin{aligned}r &= r(\mu) = \mathbb{E}_Y[r(\mu(Y), Y)], \\g &= g(\mu) = \mathbb{E}_Y[g(\mu(Y), Y)].\end{aligned}$$

Deterministic stationary ranking policy μ

It assigns a permutation $\pi = \mu(y) \in \Pi_m$ to each $y \in \Omega$.

Long-term **expected relevance** per request (reputation of the provider) and **expected revenue** per request (from the organic links), for given μ :

$$\begin{aligned} r &= r(\mu) = \mathbb{E}_Y[r(\mu(Y), Y)], \\ g &= g(\mu) = \mathbb{E}_Y[g(\mu(Y), Y)]. \end{aligned}$$

Objective: Maximize long-term utility function $\varphi(r, g)$.

Assumption: φ is **strictly increasing** in both r and g .

Example: expected revenue per unit of time

$$\varphi(r, g) = \lambda(r)(\beta + g), \quad \text{where}$$

$\lambda(r)$ = **arrival rate** of requests, strictly increasing in r ;

β = \mathbb{E} [revenue per request] from **non-organic** links (ads on root page);

g = \mathbb{E} [revenue per request] from **organic** links.

Deterministic stationary ranking policy μ

It assigns a permutation $\pi = \mu(y) \in \Pi_m$ to each $y \in \Omega$.

Q: Is this the most general type of policy?

Long-term **expected relevance** per request (reputation of the provider) and **expected revenue** per request (from the organic links), for given μ :

$$\begin{aligned} r &= r(\mu) = \mathbb{E}_Y[r(\mu(Y), Y)], \\ g &= g(\mu) = \mathbb{E}_Y[g(\mu(Y), Y)]. \end{aligned}$$

Objective: Maximize long-term utility function $\varphi(r, g)$.

Assumption: φ is **strictly increasing** in both r and g .

Example: expected revenue per unit of time

$$\varphi(r, g) = \lambda(r)(\beta + g), \quad \text{where}$$

$\lambda(r)$ = **arrival rate** of requests, strictly increasing in r ;

β = \mathbb{E} [revenue per request] from **non-organic** links (ads on root page);

g = \mathbb{E} [revenue per request] from **organic** links.

Randomized stationary ranking policy $\tilde{\mu}$

$$\tilde{\mu}(y) = \{q(\pi, y) : \pi \in \Pi_m\}$$

is a probability distribution, for each $y = (m, r_1, g_1, \dots, r_m, g_m) \in \Omega$.

Expected relevance

$$r = r(\tilde{\mu}) = \mathbb{E}_Y \left[\sum_{\pi} q(\pi, Y) \sum_{i=1}^M c_{i, \pi(i)}(Y) R_i \right]$$

Expected revenue

$$g = g(\tilde{\mu}) = \mathbb{E}_Y \left[\sum_{\pi} q(\pi, Y) \sum_{i=1}^M c_{i, \pi(i)}(Y) G_i \right]$$

Randomized stationary ranking policy $\tilde{\mu}$

$$\tilde{\mu}(y) = \{q(\pi, y) : \pi \in \Pi_m\}$$

is a probability distribution, for each $y = (m, r_1, g_1, \dots, r_m, g_m) \in \Omega$.

Let $z_{i,j}(y) = \mathbb{P}[\pi(i) = j]$ under $\tilde{\mu}$.

Expected relevance

$$r = r(\tilde{\mu}) = \mathbb{E}_Y \left[\sum_{\pi} q(\pi, Y) \sum_{i=1}^M c_{i,\pi(i)}(Y) R_i \right] = \mathbb{E}_Y \left[\sum_{i=1}^M \sum_{j=1}^M z_{i,j}(Y) c_{i,j}(Y) R_i \right]$$

Expected revenue

$$g = g(\tilde{\mu}) = \mathbb{E}_Y \left[\sum_{\pi} q(\pi, Y) \sum_{i=1}^M c_{i,\pi(i)}(Y) G_i \right] = \mathbb{E}_Y \left[\sum_{i=1}^M \sum_{j=1}^M z_{i,j}(Y) c_{i,j}(Y) G_i \right].$$

In terms of (r, g) , we can redefine (simpler)

$\tilde{\mu}(y) = \mathbf{Z}(y) = \{z_{i,j}(y) \geq 0 : 1 \leq i, j \leq m\}$ (doubly stochastic matrix).

Q: Here we have a **stochastic dynamic programming** problem, but the rewards are not additive! Usual DP techniques do not apply.

How can we compute an optimal policy?

Seems very hard in general!

Optimization problem

$$\max_{\tilde{\mu} \in \tilde{\mathcal{U}}} \varphi(r, g) = \lambda(r)(\beta + g)$$

subject to

$$r = \mathbb{E}_Y \left[\sum_{i=1}^M \sum_{j=1}^M z_{i,j}(Y) c_{i,j}(Y) R_i \right]$$

$$g = \mathbb{E}_Y \left[\sum_{i=1}^M \sum_{j=1}^M z_{i,j}(Y) c_{i,j}(Y) G_i \right]$$

$$\tilde{\mu}(y) = \mathbf{Z}(y) = \{z_{i,j}(y) : 1 \leq i, j \leq m\} \quad \text{for all } y \in \Omega.$$

Optimization problem

$$\max_{\tilde{\mu} \in \tilde{\mathcal{U}}} \varphi(r, g) = \lambda(r)(\beta + g)$$

subject to

$$r = \mathbb{E}_Y \left[\sum_{i=1}^M \sum_{j=1}^M z_{i,j}(Y) c_{i,j}(Y) R_i \right]$$

$$g = \mathbb{E}_Y \left[\sum_{i=1}^M \sum_{j=1}^M z_{i,j}(Y) c_{i,j}(Y) G_i \right]$$


$$\tilde{\mu}(y) = \mathbf{Z}(y) = \{z_{i,j}(y) : 1 \leq i, j \leq m\} \quad \text{for all } y \in \Omega.$$


To each $\tilde{\mu}$ corresponds $(r, g) = (r(\tilde{\mu}), g(\tilde{\mu}))$.

Proposition: The set $\mathcal{C} = \{(r(\tilde{\mu}), g(\tilde{\mu})) : \tilde{\mu} \in \tilde{\mathcal{U}}\}$ is **convex**.

Optimal value: $\varphi^* = \max_{(r,g) \in \mathcal{C}} \varphi(r, g) = \varphi(r^*, g^*)$ (optimal pair).

Idea: find (r^*, g^*) and recover an optimal policy from it.





Optimization

$$\nabla\varphi(r^*, g^*)'(r - r^*, g - g^*) = \varphi_r(r^*, g^*)(r - r^*) + \varphi_g(r^*, g^*)(g - g^*) = 0.$$

Let $\rho^* = \varphi_g(r^*, g^*)/\varphi_r(r^*, g^*) =$ slope of gradient.

Optimal value = $\max_{(r,g) \in \mathcal{C}} \varphi(r, g)$.

Optimization

$$\nabla\varphi(r^*, g^*)'(r - r^*, g - g^*) = \varphi_r(r^*, g^*)(r - r^*) + \varphi_g(r^*, g^*)(g - g^*) = 0.$$


Let $\rho^* = \varphi_g(r^*, g^*)/\varphi_r(r^*, g^*) = \text{slope of gradient}$.

Optimal value = $\max_{(r,g) \in \mathcal{C}} \varphi(r, g)$.

Optimal "solution" satisfies $(r^*, g^*) = \arg \max_{(r,g) \in \mathcal{C}} (r + \rho^* g)$.

The **optimal (r^*, g^*) is unique** if the contour lines of φ are strictly convex.
True for example if $\varphi(r, g) = r^\alpha(\beta + g)$ where $\alpha > 0$.

The **arg max for the linear function is unique** if and only if green line $\varphi_r(r^*, g^*)(r - r^*) + \varphi_g(r^*, g^*)(g - g^*) = 0$ touches \mathcal{C} at a single point.



One more assumption

Standard **assumption**: click-through rate has **separable** form:

$$c_{i,j}(y) = \theta_j \psi(r_i),$$

where $1 \geq \theta_1 \geq \theta_2 \geq \dots \geq \theta_{m_0} > 0$ (ranking effect)
and $\psi : [0, 1] \rightarrow [0, 1]$ increasing.

Let $\tilde{R}_i := \psi(R_i)R_i$, $\tilde{G}_i := \psi(R_i)G_i$, and similarly for \tilde{r}_i and \tilde{g}_i . Then

$$r = \mathbb{E}_Y \left[\sum_{i=1}^M \sum_{j=1}^M z_{i,j}(Y) \theta_j \tilde{R}_i \right] \quad \text{and}$$

$$g = \mathbb{E}_Y \left[\sum_{i=1}^M \sum_{j=1}^M z_{i,j}(Y) \theta_j \tilde{G}_i \right].$$

Optimality conditions, discrete case

Definition. A linear ordering policy with ratio ρ (LO- ρ policy) is a (randomized) policy that ranks the pages i by decreasing order of their score $\tilde{r}_i + \rho\tilde{g}_i$ with probability 1, for some $\rho > 0$, except perhaps when $\theta_{j'} = \theta_j$ where the order does not matter.

Theorem. Suppose Y has a discrete distribution, with $p(y) = \mathbb{P}[Y = y]$. Then any optimal randomized policy must be an LO- ρ^* policy.

Idea of proof: by an interchange argument. If for some y with $p(y) > 0$, page i at position j has lower score $\tilde{r}_i + \rho^*\tilde{g}_i$ than the page at position $j' > j$ with probability $\delta > 0$, we can gain by exchanging those pages, so this is cannot be optimal.

Optimality conditions, discrete case

Definition. A linear ordering policy with ratio ρ (LO- ρ policy) is a (randomized) policy that ranks the pages i by decreasing order of their score $\tilde{r}_i + \rho\tilde{g}_i$ with probability 1, for some $\rho > 0$, except perhaps when $\theta_{j'} = \theta_j$ where the order does not matter.

Theorem. Suppose Y has a discrete distribution, with $p(y) = \mathbb{P}[Y = y]$. Then any optimal randomized policy must be an LO- ρ^* policy.

Idea of proof: by an interchange argument. If for some y with $p(y) > 0$, page i at position j has lower score $\tilde{r}_i + \rho^*\tilde{g}_i$ than the page at position $j' > j$ with probability $\delta > 0$, we can gain by exchanging those pages, so this is cannot be optimal.

One can find ρ^* via a linear search on ρ (various methods for that). For each ρ , one may evaluate the LO- ρ policy either exactly or by simulation.

Just finding ρ^* appears sufficient to determine an optimal policy

Nice!

Optimality conditions, discrete case

Definition. A linear ordering policy with ratio ρ (LO- ρ policy) is a (randomized) policy that ranks the pages i by decreasing order of their score $\tilde{r}_i + \rho \tilde{g}_i$ with probability 1, for some $\rho > 0$, except perhaps when $\theta_{j'} = \theta_j$ where the order does not matter.

Theorem. Suppose Y has a discrete distribution, with $p(y) = \mathbb{P}[Y = y]$. Then any optimal randomized policy must be an LO- ρ^* policy.

Idea of proof: by an interchange argument. If for some y with $p(y) > 0$, page i at position j has lower score $\tilde{r}_i + \rho^* \tilde{g}_i$ than the page at position $j' > j$ with probability $\delta > 0$, we can gain by exchanging those pages, so this is cannot be optimal.

One can find ρ^* via a linear search on ρ (various methods for that). For each ρ , one may evaluate the LO- ρ policy either exactly or by simulation.

Just finding ρ^* appears sufficient to determine an optimal policy

Nice! Right?

Beware of equalities!

What if two or more pages have the same score $\tilde{R}_i + \rho^* \tilde{G}_i$?
Can we rank them arbitrarily?

Beware of equalities!

What if two or more pages have the same score $\tilde{R}_i + \rho^* \tilde{G}_i$?

Can we rank them arbitrarily?

The answer is **NO**.


Beware of equalities!

What if two or more pages have the same score $\tilde{R}_i + \rho^* \tilde{G}_i$?

Can we rank them arbitrarily?

The answer is **NO**.

Specifying ρ^* is not enough to uniquely characterize an optimal policy when equality can occur with positive probability.



Counter-example.

Single request $Y = y = (m, r_1, g_1, r_2, g_2) = (2, 1, 0, 1/5, 2)$.

$\psi(r_i) = 1$, $(\theta_1, \theta_2) = (1, 1/2)$, $\varphi(r, g) = r(1 + g)$.


For each request, $\mathbb{P}[\text{ranking } (1, 2)] = p = 1 - \mathbb{P}[\text{ranking } (2, 1)]$.

One finds that $\varphi(r, g) = (7 + 4p)(3 - p)/10$, maximized at $p^* = 5/8$.

This gives $r^* = 19/20$, $g^* = 11/8$, $\varphi(r^*, g^*) = 361/160$.

$p = 0$ gives $\varphi(r, g) = 336/160$ and $p = 1$ gives $\varphi(r, g) = \varphi^* = 352/160$.

No optimal deterministic policy here!



Continuous distribution for Y

Definition. A randomized policy $\tilde{\mu}$ is called an **LO- ρ policy** if for almost all Y , $\tilde{\mu}$ sorts the pages by decreasing order of $\tilde{R}_i + \rho\tilde{G}_i$, except perhaps at positions j and j' where $\theta_j = \theta_{j'}$, at which the order can be arbitrary.

Theorem (necessary conditions). Any optimal policy must be an LO- ρ policy with $\rho = \rho^*$.

Continuous distribution for Y

Assumption A. For any $\rho \geq 0$ and $j > i > 0$,

$$\mathbb{P}[M \geq j \text{ and } \tilde{R}_i + \rho \tilde{G}_i = \tilde{R}_j + \rho \tilde{G}_j] = 0.$$

Theorem (sufficient condition). Under Assumption A, for any $\rho \geq 0$, a deterministic LO- ρ policy sorts the pages for Y uniquely with probability 1. For $\rho = \rho^*$, this policy is optimal.

Idea: With probability 1, there is no equality.

Continuous distribution for Y

Assumption A. For any $\rho \geq 0$ and $j > i > 0$,

$$\mathbb{P}[M \geq j \text{ and } \tilde{R}_i + \rho \tilde{G}_i = \tilde{R}_j + \rho \tilde{G}_j] = 0.$$

Theorem (sufficient condition). Under Assumption A, for any $\rho \geq 0$, a deterministic LO- ρ policy sorts the pages for Y uniquely with probability 1. For $\rho = \rho^*$, this policy is optimal.

Idea: With probability 1, there is no equality.

In this case, it suffices to find ρ^* , which is a root of

$$\rho = \tilde{h}(\rho) := h(r, g) := \frac{\varphi_g(r, g)}{\varphi_r(r, g)} = \frac{\lambda(r)}{(\beta + g)\lambda'(r)}.$$

Can be computed by a root-finding technique.

Proposition. (i) If $h(r, g)$ is bounded over $[0, 1] \times [0, K]$, then the fixed-point equation $\tilde{h}(\rho) = \rho$ has at least one solution in $[0, \infty)$.
 (ii) If the derivative $\tilde{h}'(\rho) < 1$ for all $\rho > 0$, then the solution is unique.

Proposition.

Suppose $\varphi(r, g) = \lambda(r)(\beta + g)$.

Then $h(r, g) = \lambda(r)/(\beta + g)\lambda'(r)$ and

- (i) If $\lambda(r)/\lambda'(r)$ is bounded for $r \in [0, 1]$ and $g(\rho(0)) > 0$, then the fixed point equation has at least one solution in $[0, \infty)$.
- (ii) If $\lambda(r)/\lambda'(r)$ is also non-decreasing in r , then the solution is unique.

Often, $\rho \mapsto \tilde{h}(\rho)$ is a contraction mapping. It is then rather simple and efficient to compute a fixed point iteratively.

Proposition.

Suppose $\varphi(r, g) = \lambda(r)(\beta + g)$.

Then $h(r, g) = \lambda(r)/(\beta + g)\lambda'(r)$ and

- (i) If $\lambda(r)/\lambda'(r)$ is bounded for $r \in [0, 1]$ and $g(\rho(0)) > 0$, then the fixed point equation has at least one solution in $[0, \infty)$.
- (ii) If $\lambda(r)/\lambda'(r)$ is also non-decreasing in r , then the solution is unique.

Often, $\rho \mapsto \tilde{h}(\rho)$ is a contraction mapping. It is then rather simple and efficient to compute a fixed point iteratively.

In this continuous case, computing an optimal deterministic policy is relatively easy. It suffices to find the root ρ^* and use the LO- ρ^* policy, which combines optimally relevance and profit.

What to do for the discrete case?

Often, only a randomized policy can be optimal. But such a policy is very hard to compute and use in general! Not good!

Much simpler and better solution:


Select a small $\epsilon > 0$ (e.g., $\epsilon = 10^{-10}$) and whenever some items i have equal scores, add a random perturbation, uniform over $(-\epsilon, \epsilon)$, to each G_i .

Under this perturbed distribution of Y , the probability of equal scores becomes 0 and one can compute the optimal $\rho^*(\epsilon)$, which gives the optimal policy w.p.1.

Proposition. Let φ^* be the optimal value (average gain per unit of time) and φ^{**} the value when applying the optimal policy for the perturbed model (with an artificial perturbation), both for the original model. Then

$$0 \leq \varphi^* - \varphi^{**}(\epsilon) \leq \lambda(r^*(\epsilon))(\theta_1 + \dots + \theta_{m_0})\epsilon.$$

Application to previous example



ϵ	$p^*(\epsilon)$	$\rho^*(\epsilon)$	$r^*(\epsilon)$	$g^*(\epsilon)$	$\varphi^*(\epsilon)$	$\varphi^{**}(\epsilon)$
0.0	0.625	0.4	0.95	1.375	2.25625	2.25625
0.001	0.62491	0.39995	0.94996	1.37521	2.25636	2.25615
0.01	0.62411	0.39950	0.94964	1.37006	2.25736	2.25539
0.1	0.61705	0.39537	0.94682	1.39476	2.26741	2.24869
0.5	0.59771	0.38137	0.93908	1.46240	2.31240	2.23031

$\varphi^* = \varphi^*(0) =$ optimal value.

$\varphi^*(\epsilon) =$ optimal value of perturbed problem.

$\varphi^{**}(\epsilon) =$ value using optimal policy of perturbed problem.

Conclusion

Even if our original model can be very large and an optimal policy can be complicated (and randomized), we found that for the continuous model, the optimal policy has a simple structure, with a single parameter ρ that can be optimized via simulation.

In real-life situations in which our model assumptions may not be satisfied completely, it makes sense to adopt the same form of policy and optimize ρ by simulation. This is a viable strategy that should be often close to optimal.

Other possible approach (future work): a model that uses discounting for both relevance and gains.

Model can be refined in several possible directions.

More details

P. L'Ecuyer, P. Maille, N. Stier-Moses, and B. Tuffin.
Revenue-Maximizing Rankings for Online Platforms with Quality-Sensitive Consumers.

GERAD Report

On my web page.

Also at <https://hal.inria.fr/hal-00953790>.