

Ranking article comments using reinforcement learning

Lester Solbakken | October 28th 2019



Google announces major breakthrough

The search company claims that it's achieved quantum supremacy, a major computing milestone.

What this means »







Rumors swirl about potential 'I realize I am taking a impeachment circus

personal and a career risk' U.S. AOL.Com

Anchor's devastating

miscarriages while on the job future looms before Brady

'They'll have no choice but to step back and take notice'

3891

5415



Mourners pay tribute to US Rep. Cummings USA TODAY





Ad MacKeeper Is Your Mac Catalina-ready?

Free up your Mac's disk space so it's ready for macOS Catalina.

challenges' -- here's a look at the legacy and wealth he left behind.



Politics The Independent

Hillary Clinton says she would run again in 2020 if she thought she could win, report says

Rep. Elijah E Cummings' net worth at the time of his passing

Elijah Cummings passed away in October due to complications from 'longstanding health

Hillary Clinton has told people privately that she would consider joining the 2020 Democratic primary, but only if she thought she could win. The private conversations were aired publicly i...

Hillary Clinton's attacks on Tulsi Gabbard are embarrassing The Guardian



Hillary Clinton weighs in on Democrats in 2020 race Yahoo News Video



Encourages meaningful discussion?





now."

verizon media



YAHOO! TechCrunch IHUFFPOSTI RYOT MAKERS tumblr. #BUILTBYGIRLS Engadget © FLURRY AUTODIO





Vespa at

verizon media

Hundreds of Vespa applications (Flickr, Tumblr, TechCrunch, Huffington Post, Aol, Gemini, Engadget, Yahoo News Sports Finance Mail etc.):

- serving over a billion users,
- hundreds of thousands of queries per second,
- billions of content items.







Around 30 developers in Trondheim, Norway





Baseline - existing solution

Comments found on many Yahoo properties such as Yahoo Finance, Yahoo News, and Yahoo Sports

- ~ 1 billion comments stored
- ~ 12.000 queries per second
- 2x that for updates

Some articles have > 100.000 comments!

| manter R | achel England, Engadget 18 hours ago 🗦 🛱 🕇 | 9 8 |
|--------------|--|---------------|
| Read | full article 🔻 | |
| () 1, | 871 reactions | |
| Sign | in to post a message. | |
| Top R | eactions 🗸 📫 2 | 99 viewin |
| | 2 people reacting •••• | |
| Ŷ | $\label{eq:constraint} \begin{array}{l} \textbf{XC} & 12\text{hours ago} \\ \hline \text{One of the biggest impacts quantum computers will have is its impact on cryptograd Current encryption could be broken in seconds with quantum computers.} \\ \hline \ensuremath{\mathfrak{S}_{2}} \text{ Reply } \text{Replies (119) } & \ensuremath{\Delta}_{2} \text{ Osc} \ensuremath{\mathcal{Q}_{2}} \text{ 2s} \end{array}$ | |
| | | |
| M | Mason 1 hour app Google claims to have made a quantum computing breakthrough with a processor performs in minutes work that would take classical computers 10,000 years. This demonstration of so-called 'quantum supremacy' over conventional computers More ~ | that has |
| | Meson Theor app Google calins to have made a quantum computing breakthrough with a processor of performs in minutes work that would take classical computers 10,000 years. This demonstration of so-calied 'quantum supremacy' over conventional computers to the so-calied 'quantum supremacy' over conventional computers to the source of the s | that 6 has |
| | Macon 1 hoar app Google claims to have made a quantum computing breakthrough with a processor performs in minutes work that would take classical computers 10,000 years. This demonstration of so-called 'quantum supremacy' over conventional computers More • Core performs Google clouts quantum computing milestone neverthreva.top convertineva.top convertineva.top | that s has |
| | Nason 1 hour app Google claims to have made a quantum computing breakthrough with a processor inperforms in minutes work that would take classical computers 10,000 years. This demonstration of so-called 'quantum supremacy' over conventional computers More Cogle could quantum supremacy' over conventional computers More Cogle touts quantum computing milestone newsenswa.sp ch Reply Reples (2) 66 Q 4 Magent 11 hours apo This means that they con solve the protein folding prediction problem Ch Reply Replies (2) 2 4 2 3 | that 5 has |
| | Name 1 hour and Stoogle claims to have made a quantum computing breakthrough with a processor performs in minutes work that would take classical computers 10,000 years. This demonstration of so-called 'quantum supremacy' over conventional computers Nore > Store > Image: Store > C_{31} Reply Reples (2) C_{62} C_{32} Reply Reples (2) C_{62} | that |
| | Name Theorem Stoogle claims are made a quantum computing breakthrough with a processor. Stoogle claims | that s has |



Potential features

Community How users interacted with comment

- **Comment** Relevance to topic, moderation
 - Author Reputation
 - **User** Preferences
 - **Other** Time

Wilson score:* probability of comment being overwhelmingly liked by all users



$$s(n_{\uparrow},n_{\downarrow}) = \frac{\hat{p} + \frac{z_{1-\alpha/2}^{2}}{2n} - \sqrt{\frac{z_{1-\alpha/2}^{2}}{n} \left[\hat{p}(1-\hat{p}) + \frac{z_{1-\alpha/2}^{2}}{4n}\right]}}{1 + \frac{z_{1-\alpha/2}^{2}}{n}},$$

Conversation AI (https://conversationai.github.io)



shut up you fucking idiot

Perspective



(*) Zhang et. al. 2011. How to Count Thumb-Ups and Thumb-Downs: User-Rating Based Ranking of Items from an Axiomatic Perspective.

Previous ranking algorithm







Scoring

Ranking

Learning





Scoring

How should features be **combined intelligently**?

Ranking

How can we overcome **position bias**?

Learning

How do we learn directly from **user behavior**?





Scoring

How should features be **combined intelligently**?

Neural network over comment features

Ranking

How can we overcome **position bias**?

Learning

How do we learn directly from **user behavior**?





Scoring

How should features be **combined intelligently**?

Neural network over comment features

Ranking

How can we overcome **position bias**?

Exploration with sampling

Learning

How do we learn directly from **user behavior**?





Scoring

How should features be **combined intelligently**?

Neural network over comment features

Ranking

How can we overcome **position bias**?

Exploration with sampling

Learning

How do we learn directly from **user behavior**?

Reinforcement learning with dwell time rewards



Reinforcement learning in general

RL is a general-purpose framework for artificial intelligence

- RL is for an agent with the capacity to act
- Each action influences the agent's future state
- Success is measured by a scalar reward signal
- Select actions to maximise future reward





Contextual bandits

Multi-arm bandits with context

features **x**

$$\begin{bmatrix} 0.1, 0.7, \dots, 0.2 \end{bmatrix}$$
score **v** = **f**(**x**)
$$a_1 : 0.4 \quad a_2 : 0.2 \quad a_3 : 0.3$$

$$a_1 \qquad a_1 \qquad a_$$

Reward **r** is conditioned on chosen action - *feedback is partial*

Canonical example: ad serving



Source: Microsoft research



Contextual bandits in ranking

Sometimes called contextual semibandits*



Importance weighted sampling to construct unbiased estimates for rewards



Comment





Comment

Features

Community

Comment

| 0 | Richard | 10 hc | urs ago | | | | | | |
|---|--|--------------|---------|-------|--|--|--|--|--|
| | As a physicist myself who teaches QM, Particle Physics at one of the | | | | | | | | |
| | UC's, the Quantum computing is the same farce as the cold fusion | | | | | | | | |
| | turned out to be. I've been to Google and have seen the system. | | | | | | | | |
| | More ~ | | | | | | | | |
| | S Reply | Replies (52) | A 157 | 52 36 | | | | | |

Author

User

Other



Comment





Community



Comment



User

Other





Model θ Comment Features Positive score Community Comment Richard 10 hours ago As a physicist myself who teaches QM, Particle Physics at one of the UC's, the Quantum computing is the same farce as the cold fusion Author turned out to be. I've been to Google and have seen the system. More v A Reply Replies (52) 157 🖓 36 User Other







Comments

















Sampling





Comments





Scores







Charles 11 hours ago It reminds me of the joke where the U.S. Government gathered the best computer scientists together to build the biggest computer ever to answer age old questions. They built it and the first question they asked was "Is there a God?". The computer answered, "There is now."



Sampling

Ranking π





Comments



What no one says about quantum computing is that superposition makes abundant errors part of any quantum computing process. Until we know more about quantum mechanics and superposition specifically, 1 don't see much in the way of true quantum breakthroughs. After we solve the error ridied process however, every code is unbreakable, and the sky is the limit.







vespa.ai











Ranking π



Comments





Scores



Sampling





Ranking π





vespa.ai



Comments



It reminds me of the joke where the U.S. Government gathered the best computer scientists together to build the biggest computer ever to answer age old questions. They built it and the first question they asked was "is there a God?". The computer answered, "There is now."



Sampling







Ranking π





Model θ

Rankings π



| C | shrik 11 What no is makes also we know is geothcall | ne says abou ne says abou note about o p. 1 don't see | f quantum comp part of any quar- santum mechani much in the way | uting is the turn comp (is and sig (if the gu | i specielice dig provisi Un episitice dura | | | | | | |
|---|---|--|---|---|---|---|--|--|--|--|--|
| | every cod | one ware | to, and the stry | riddod pro | ioss licerener, | | | | | | |
| 0 | Gunnes 1 Enventionals of Dest comp to answer a mixed was reper. ¹ | Replect() Investory ne of the paid der scientiste ge old quest 'to there a Gr | where the U.S. I rangether to cust one. They built it eC ² . The compar- | Downvised the Doge and the fit for attained | patrenol the computer ever question they c. "There is | | | | | | |
| • | G, Rejo Richard P Al A Dityl UC'h, Per Saned au Mare v | Suprim (12) Cast reposit v Guershum com No Sel. Free D | C 40° Q 10 to baches QM, rowing is the or een to bacque a | Particle Pt eno farce a la name se | ytics at one of the total days | i. | | | | | |
| J | Contrapio Traver 12 The also be ignore that | Replice (12) hours app addressing in the register | 010 QH | ion that I g | The Fee 206 | -0.5 505 600 800 800 800 800 800 800 800 800 800 | f quantum er pert of any c santum mech much in the o | onputing is partial tents haves and s may of true is not readed of | Pet superpesition sputing process. Un sperpesition publics publics | 41 | |
| | 68/8 111 | Pageta de | 0 4 99 | | | - | Do, and the s | sky 16 The 30 | 4 | | |
| | Whet no a makes all we know a specifical beaktivas every cash | ne søys abox undant errors nore about op p. I don't see ops. After se o is unbrooka | d quantum comp pert of any quan- antum mechani much in the way b salue the encor be, and the sky | sting is the turn comp of true pur robbed pro- is the and | i spepelion iting process. Un epielion nturk loss however, | 2222 | rates the G ragether to C ons. They but sc?". The com | 3. Governmonth outd the bog in it and the router array | ent gatesed the pat computer new first question they and, "There is | | |
| ~ | Guilley I | Papies (I) | 635 Q2 | | | | 647 Q1 | | | | |
| U | Erweinick o best comp to answer a mixed was now." | ne of the poor new accention ge off quest 'to there a Ge | where the U.S. I respective to cure one. They built it re?". The compa- | The Doge and the Tri for answers | all pathened the computer ever computer ever computer they c. There is | er w t con te b | to teaches of riputing is the een to Geogr | 244, Particle e same farce le and hore i | Physica at one of the as the cold fusion seen the system. | | |
| | G, Berly | Replan (18) | 00.64 | | | P 2) | $\Delta = Q$ | 28 | | | |
| | As a physic of the terms of | cist repetit v Guentum con 130 Sei. Frie D | to teaches QM, routing is the so een to bacque a | Particle Pt mo farce a ht name an | ysics at one of the the celld fusion is the celld fusion is the cystem. | | e2020USe 스< 구) | lection But | The Few area o | er computing is that superposition any quantum computing process. Until mechanics and superposition the way of true quantum | |
| - | Suffering the | Replice (42) hours ago | Q = Q = | | | -2 | coartern m | onputing is | hel supepesitor | e sky is the limit | |
| J | IT'S 2000 HI | adarang is th | e 21223 US elect | ion Ball 1 | The Few 22% | - | antum moot | hands and s | cepesitos | | |
| | du hiqiy | Feplies (I) | $Q \in (\hat{\Delta})$ | | | H M | salue the or too, and the r | nor ridded p sky 16 the an | racess however, d | U.S. Government gathered the I build the biggest computer ever | |
| | | | | - | G, Paply Repla | -00 | 635 Q2 | | | provide answered. There is | |
| | | | | ۲ | Energiands are at the set computer according to the set of the set | a pina citata ponti | where the G ragether to C ons. They but will The cost | 3. Governmonth outside the bigs in it and the | ert gateend the put computer ever fest question they and . "Data is | U III | |
| | | | | | ican' | | | | | is the same falce as the cold fasion copie and have seen the system. | |
| | | | | + | RCRAH | 111 | tere ago | M Darbie | ituara e ano e e | - Q.S. | |
| | | | | | UC's, the Quest, turned out to be | THE D | routing is the een to bacogi | e some face is and have | as the cold fusion loost the cythers. | election. But I guess we'll just | |
| | | | | | Mark v Galleniy Realer | 101 | 0-0 | 24 | | | |
| | | | | | | | | where the | - | re convoline is that supervisition | |
| | | | | - | grow that right? | | | | | my quantum computing process. Until mechanics and superposition | |
| | | | | | | | | beat in the | opis. Alter we solve | the way of Irue quarters The error statled process however, of the size is the hold | |
| | | | | | | | | So Reply | Rein (1 63 | 0. | |
| | | | | | | | 0 | Charles I Eremonds in | I have age for all the pilos where | Ite 2.3. Government gattered the | |
| | | | | | | | | to answer a mixed was now." | pe old questions. The there is Get ?". T | er to build the biggest computer ever wy built it and the first question they be computer answered. "There is | |
| | | | | | | | | G, Reply | Replan (14) (2) (2) | Q 10 | |
| | | | | | | | - | AL a pres UC 1, the turned ou More - | Quartur computer to be | anes CM, Particle Physics at one of the g is the same face as the cold fasion Goople and have seen the system. | |
| | | | | | | | ۲ | College Dever 12 The also be | Poples (12) 🔥 1 hout age advising in the 2020 | IT Q'SI IUS RECTOR. But I guess we'l put | |
| | | | | | | | - | ignow that is inspire | Replace 🛆 e | φ) | |
| | | | | | | | | | | | |



Model θ

Rankings π — Reward $R(\pi)$



| | diffiti. 11 hours age | | | | | |
|---|--|---|---|--|--|--|
| C | What no one says about quantum computing it makes abundant errors part of any quantum ce | that superposition reputing process. Un- | | | | |
| | we know more about quantum mechanics and specificate. I don't see much in the way of the | superposition | | | | |
| | beakthroughs. After we salve the enty ridded | process however, | | | | |
| | So Party Restance Architectory and the set of | | | | | |
| | Charles 11 hours age | | | | | |
| • | Eventuals the of the place where the U.S. Covers's beef computer accentists together to curd the bid | nert gatered the cent concuber ever | | | | |
| | to answer age oil quantions. Twy built it and the minut ass. In these a Gest?" The corrector area | finz question they | | | | |
| | now." | | | | | |
| | Galleto Restanci (Chill Chill) | | | | | |
| + | As a physicial myself who teaches QM, Particle | Physics at one of th | | | | |
| | UC 5. the Quantum computing is the same fam turned aut to be. Fre been to bacque and have | e as the cold fusion deet the system. | | | | |
| | Mare v | | | | | |
| - | Contractor and American State | | out quantum to | emputing is that superposition | | |
| J | It's also interfering in the 2020 LtS election. But | Tauj trev armigi | of guardam med | sation random process or satios and superposition | | |
| | ignore that right? | | see much in the in trace subsective or | way of true-quantum for reliabled process bourseer. | | |
| | data There are | | Heatie, and the r | ay is the and | | |
| C | What no one says about quantum computing it | Pel supepestion | 4.8 Q2 | | | |
| | we know more about quantum mechanics and | infection and because out | to where the G | 3. Government gathered the | | |
| | specifically, I don't see much in the way of true beauthroughs. After we salve the error ridded | quantum process however, | ations. They but | in it and the finit question they | | |
| | every code is unbroakable, and the sky is the to | 54 | Ged?". The com | puter answered. 'There is | | |
| - | Carles (Lineran | | - 60 QH | | | |
| U | Eventeds the of the pilos where the U.S. Government of the pilos where the U.S. Government of the pilos of th | nett gattered the | Distance approaches and | W Ratio Resource of some of the | - | |
| | to answer age old questions. They built it and the | Int question they | I computing is the | anno tarce as the cold fusion | | |
| | now." | and there | R 088777 0809 | a constraint president of a system of | | |
| | Gallety Replace() () (C (C (C))) | | 10 0 × 0 | | | |
| + | As a physicial myself who teaches QM, Particle | Physics at one of th | e the 2020 US e | ectors that I guess with just | in computing is that superposition | |
| | UC 5. the Quantum computing is the same fam turned aut to be. Fre been to bacque and have | e as the cold fusion also the cythers. | | | my quantum computing process. Until hechanics and superposition | |
| | Mare v | | 0 < 91 | | the way of true quantum | |
| - | Contractor and American State | | out quantum to | emputing is that superposition | e sky is the limit | |
| J | It's also interfering in the 2020 US election. But | Tory I've assug 1 | of quantum med | anos and superposition | 2 | |
| | A Real Particle (A C (2)) | | see much in the or or we salve the or | kay of true-guantum or ridded process however, | U.S. Government gathered the | |
| | | | -solution, and the s | ay is the and | Louid the biggest computer ever suit it and the first question they | |
| | | Charles II haven | | | propular answared, "There is | |
| | | Erennik ne at t | a joho where the G | 3. Government gathered the | DH- | |
| | | to answer age old | questions. They but | it it and the first question they | the FRA Restore Restore of one of the | |
| | | NOW, | e a Georgia Nelcon | CON REPART. LINES | is the same falce as the cold fusion | |
| | | G, Reply Replan | 08 047 010 | | tops and the second second | |
| | 4 | As a physician re- | to have ago out who happed of | M, Particle Physics at one of th | 2 P.M. | |
| | | UC's, the Quest, SameDaut to be. | re computing is the rise been to bacge | a same farce as the cold fusion and faire spect the system. | election. But I guess we'll put | |
| | | Mare v | | | | |
| | - | Trever Street | 994 (310 Q | · . | | |
| | 4 | The alloca between the | a the 2020 US et | ection. But I guess we'll just | ex computing is that superposition by autobar computing process from | |
| | | A Professional April 1 | 0.01.01 | | nechanics and superposition | |
| | | | | broaklyingts. After we solve | the error sollied process however, | |
| | | | | every lock is unbroketing, or | d the sky is the limit | |
| | | | | Charles II have apr | | |
| | | | | Eremands one of the jobs where | Ine G.B. Government gathered the or to fourth the force of control for any | |
| | | | | to answer age old questions. To | ey built it and the first question they | |
| | | | | now." | A COLORIA SUBMARY, 1280-10 | |
| | | | | Skiller Relation (St. 15 | Q III | |
| | | | - | Robert Science of Auto State | ites OM, Particle Physics at one of the | |
| | | | | UC's, the Quantum computer turned to the fue heat to | is the same falce as the cold fusion Goods and how seen the potent | |
| | | | | Mon - | the spectrum of the spectrum o | |
| | | | - | Collegely Replace (C) (C) | r Qai | |
| | | | | IT's goo starfing to the 2020 | LEX etection. But I pueses with hull | |
| | | | - | | | |
| | | | - | sprove that right? | 0) | |





Model θ



| Rar | ıkin | gs | π |
|---|--------|--|--|
| Part = net al. Part = net al.< | | | |
| | • ÷ | Control of the pair when any control of | The U.S. Obverview garment the the U.S. Obverview garment the the built the bigging compare war by built taked the fait quarticol trap comparing anywhere. "There is The comparing anywhere the The CDC, Parocce Physics at one of the |
| | | UCIL THE Quantum computing thread out to be. The been to the Mark | In the same mode at the cold fusion looping and the sponten ${\rm the sponten}$, ${\rm e}_{\rm CO}$, ${\rm e}_{\rm $ |



Gradient ascent in direction of expected reward





Model θ





Model θ



Can use any reward



Bootstrapping and testing

Cold start: pre-train neural network to emulate previous ranking

• Gradient ascent with Kendall's tau coefficient as reward

Off-policy evaluation: interactions are logged as (*x*, *a*, *r*, *p*), where *p* is the policy's probability of choosing *a* given *x*.

 Inverse-Propensity Scoring* for estimating average reward of a some policy from data collected by another policy



Elements of a solution





Elements of a solution is a physicist myself who teaches QM. UC's, the Quantum computing is the same f What no one says about quantum computing specifically. I don't see much in the way of true quantum souchs. After we solve the error riddled prokable, and the sky is the li Trevor 12 hours ago It's also interfering in the 2020 US election. But I guess we'll just ignore that right? to answer age old questions. They built it and the first question they asked was "is there a God?". The computer answered, "There is now." Comments Scoring model ÷ the 2020 LDS electron. But I guess we'll just Ranking log **(r) Distributed DB Reward instrumentation**



Elements of a solution s a physicist myself who teaches QN UC's, the Quantum computing is the same f What no one says about quantum compo makes abundant errors part of any quantum computing we know more about quantum mechanics and superpos specifically. I don't see much in the way of true quantum uchs. After we solve the error riddled Trevor 12 hours app It's also interfering in the 2020 US election. But I guer ignore that right? build the biggest computer eve to answer age old questions. They built it and the first question they asked was "is there a God?". The computer answered, "There is Comments Scoring model 109 (X, a, p) 2020 US election that I guess with just Ranking log **(r) Distributed DB Reward instrumentation**







Implementation





Vespa

A platform for low latency computations over large, **evolving** data sets:

- Search and filter over structured and unstructured data
- Query time organization and aggregation of matching data
- Real-time writes

- Advanced relevance scoring with tensors as first class citizens*
- Scaleable and fast
- Elastic and fault tolerant
- Pluggable
- Easy to operate



Typical use cases: text search, personalization, recommendation, targeting, real-time data display



Vespa as comment serving system

Scaleable and fast

- About 1 billion comments / ~12.000 queries per second
- Read latency 7ms for 10k comments including model evaluation
- Write latency ~1ms

Direct deployment of ML scoring models

Advanced computation framework for complex features

Custom logic for implementing sampling and logging

Hosted for simpler architecture *





Scalable low latency execution

How to bound latency:

- Parallelization 1)
- 2) Prepared data structures (indexes etc.)
- 3) Move execution to data nodes



Querv



Deploying ML models to Vespa

- 1. Model in application package
- 2. Download model from external source during (re-)deployment
- 3. Feed model weights as tensors





Deployment strategy





Results and ongoing work

~25% increase in time spent

Experimenting with

- more features for a larger neural networks
- personalized comment ranking
- more sophisticated rewards



Generalizing the implementation





Thanks to

Verizon Media Engineering

Sreekanth Ramakrishnan Aaron Nagao Zhi Qu Xue Wu

Verizon Media Science

Akshay Soni Kapil Thadani





Thank you!

https://vespa.ai/cloud