

Apache Mahout

Making data analysis easy





Isabel Drost

Nighttime:

Co-Founder, committer Apache Mahout.
Organiser of Berlin Hadoop Get Together.

Daytime:

Software developer.
Guest lecturer at TU Berlin.
Co-Organiser Berlin Buzzwords 2010.



- “Mastering Data-Intensive Collaboration and Decision Making”
- EU funded research project
 - Number of partners: 8
 - Coordinator: Research Academic Computer Technology Institute (CTI), Greece

neofonie*

Hello Devovxx!



Hello Devovx!



Hello Devvxx!



Hello Devovx!



Hello Devvxx!

Machine learning
background?

Hello Devovx!



Hello Devovxx!

Agenda

- Data Mining/ Machine Learning?
- Why is scaling hard?
- Going beyond simple statistics.

Data Mining Applications

- Marketing.
- Surveillance.
- Fraud Detection.
- Scientific Discovery.
- Discover items usually purchased together.

= Extracting patterns from data.

Machine Learning Applications

- E-Mail spam classification.
- News-topic discovery.
- Building recommender systems.

= Extracting prediction models from data.

Machine learning – what's that?






Image by John Leech, from: *The Comic History of Rome* by Gilbert Abbott A Beckett.
Bradbury, Evans & Co, London, 1850s
Archimedes taking a Warm Bath


Archimedes model of nature

$$\frac{\text{Density of Object}}{\text{Density of Fluid}} = .$$

$$\frac{\text{Weight}}{\text{Weight} - \text{Apparent immersed weight}}$$



An SVM's model of nature



The challenge

Mission

Provide scalable data mining algorithms.






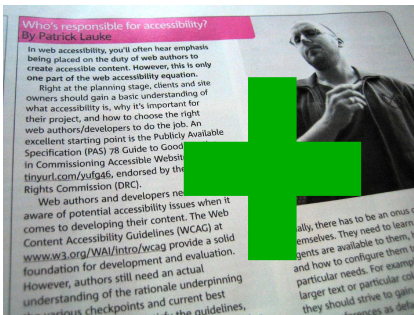


HowTo: From data to information.



January 3, 2006 by Matt Callow
<http://www.flickr.com/photos/blackcustard/81680010>

http://www.flickr.com/photos/29143375@N05/3344809375/in/photostream/

http://www.flickr.com/photos/redux/409356158/





The **HDFS filesystem** is not restricted to **MapReduce jobs**. It can be used for other applications, many of which are under way at Apache. The list includes the **HBase database**, the **Apache Mahout machine learning system**, and **matrix operations**.

DATENSCHUTZ

Tausende demonstrieren für Bürgerrechte im Netz

Für einen besseren Arbeitnehmerdatenschutz und gegen die Gesundheitskarte: 130 Organisationen hatten zur Demonstration aufgerufen. Sie fürchten den Überwachungsstaat.

© Tim Brakemeier/dpa



In Berlin demonstrieren tausende Demonstranten für mehr Datenschutz

Rund 7500 Demonstranten nahmen an dem Protestzug unter dem Motto "Freiheit statt Angst – Stoppt den Überwachungswahn" in Berlin teil.


Die Demonstranten wandten sich unter anderem gegen die Vollerhebung

- DATUM
- QUELL
- KOMM
- EMPF
- Buzz
- ARTIK
- SCHLA
- Datens

- NEU IM R
1. DATENS
 2. FUTURE
 3. URHEB
 4. SPAM A
 5. IPHONE

- NEU AUF
1. NACHR
 2. CDU U
 3. GEDEN

1. BUNDESLIGA Der BVB und die Freiburger siegen
2. DAAD Das Ende der Ära Bode



The National Aquatics Center lies on the... among the most imaginative architec-
tural feats in recent memory.
Critics have incessantly described
the profile projects as bullish
the nation's budding
buildings are
tower.

Europe. Both are part of
bible society that exte
Grand Central Termi
the great train halls
Like Tempelhof,
nal boasts a sweet
evokes the glamo
enclosing a surp
water p



Who's responsible for accessibility?
By Patrick Lauke

In web accessibility, you'll often hear emphasis being placed on the duty of web authors to create accessible content. However, this is only one part of the web accessibility equation. Right at the planning stage, clients and site owners should gain a basic understanding of what accessibility is, why it's important for their project, and how to choose the right web authors/developers to do the job. An excellent starting point is the Publicly Available Specification (PAS) 78 Guide to Good Practice in Commissioning Accessible Websites, at tinyurl.com/yulfg6, endorsed by the Disability Rights Commission (DRC).

Web authors and developers need to be aware of potential accessibility issues when it comes to developing their content. The Web Content Accessibility Guidelines (WCAG) at www.w3.org/WAI/intro/wcag provide a solid foundation for development and evaluation. However, authors still need an actual understanding of the rationale underpinning various checkpoints and current best practices in the guidelines.

Finally, there has to be an onus on themselves. They need to learn what agents are available to them, and how to configure them to particular needs. For example, larger text or particular color contrast should be used where they should strive to gain a better understanding of the guidelines.

Who's responsible for accessibility?
By Patrick Lauke

In web accessibility, you'll often hear emphasis being placed on the duty of web authors to create accessible content. However, this is only one part of the web accessibility equation. Right at the planning stage, clients and site owners should gain a basic understanding of what accessibility is, why it's important for their project, and how to choose the right web authors/developers to do the job. An excellent starting point is the Publicly Available Specification (PAS) 78 Guide to Good Practice in Commissioning Accessible Websites, at tinyurl.com/yulfg6, endorsed by the Disability Rights Commission (DRC).

Web authors and developers need to be aware of potential accessibility issues when it comes to developing their content. The Web Content Accessibility Guidelines (WCAG) at www.w3.org/WAI/intro/wcag provide a solid foundation for development and evaluation. However, authors still need an actual understanding of the rationale underpinning various checkpoints and current best practices in the guidelines.

Finally, there has to be an onus on themselves. They need to learn what agents are available to them, and how to configure them to particular needs. For example, larger text or particular color contrast should be used where they should strive to gain a better understanding of the guidelines.

Who's responsible for accessibility?
By Patrick Lauke

In web accessibility, you'll often hear emphasis being placed on the duty of web authors to create accessible content. However, this is only one part of the web accessibility equation. Right at the planning stage, clients and site owners should gain a basic understanding of what accessibility is, why it's important for their project, and how to choose the right web authors/developers to do the job. An excellent starting point is the Publicly Available Specification (PAS) 78 Guide to Good Practice in Commissioning Accessible Websites, at tinyurl.com/yulfg6, endorsed by the Disability Rights Commission (DRC).

Web authors and developers need to be aware of potential accessibility issues when it comes to developing their content. The Web Content Accessibility Guidelines (WCAG) at www.w3.org/WAI/intro/wcag provide a solid foundation for development and evaluation. However, authors still need an actual understanding of the rationale underpinning various checkpoints and current best practices in the guidelines.

Finally, there has to be an onus on themselves. They need to learn what agents are available to them, and how to configure them to particular needs. For example, larger text or particular color contrast should be used where they should strive to gain a better understanding of the guidelines.

ZEIT ONLINE DATENSCHUTZ

STATISTIK POLITIK WIRTSCHAFT MEINUNG GESELLSCHAFT KULTUR WISSEN DIGITAL STUDIUM KARRIERE LEBENSART NESS AUTO

SPORT INTERNET DATENSCHUTZ MAIL GAMES

DATENSCHUTZ

Tausende demonstrieren für Bürgerrechte im Netz

Für einen besseren Arbeitnehmerdatenschutz und gegen die Gesundheitskarte: 130 Organisatoren hatten zur Demonstration aufgerufen. Sie führten den Überwachungsstaat.

11. Berlin demonstrieren Tausende Demonstration für mehr Datenschutz

Rund 1700 Demonstranten nahmen an dem Protest unter dem Motto "Freiheit statt Angst - Bürger des Überwachungsstaats" in Berlin teil.

Wen Protestveranstaltungen stattfinden, nicht werden aufgerufen, sondern die TeilnehmerInnen

experteer



ZEIT ONLINE DATENSCHUTZ

STATISTIK POLITIK WIRTSCHAFT MEINUNG GESELLSCHAFT KULTUR WISSEN DIGITAL STUDIUM KARRIERE LEBENSART NESS AUTO

SPORT INTERNET DATENSCHUTZ MAIL GAMES

DATENSCHUTZ

Tausende demonstrieren für Bürgerrechte im Netz

Für einen besseren Arbeitnehmerdatenschutz und gegen die Gesundheitskarte: 130 Organisatoren hatten zur Demonstration aufgerufen. Sie führten den Überwachungsstaat.

11. Berlin demonstrieren Tausende Demonstration für mehr Datenschutz

Rund 1700 Demonstranten nahmen an dem Protest unter dem Motto "Freiheit statt Angst - Bürger des Überwachungsstaats" in Berlin teil.

Wen Protestveranstaltungen stattfinden, nicht werden aufgerufen, sondern die TeilnehmerInnen

experteer

Who's responsible for accessibility?
By Patrick Lauke

In web accessibility, you'll often hear emphasis being placed on the duty of web authors to create accessible content. However, this is only one part of the web accessibility equation. Right at the planning stage, clients and site owners should gain a basic understanding of what accessibility is, why it's important for their project, and how to choose the right web authors/developers to do the job. An excellent starting point is the Publicly Available Specification (PAS) 78 Guide to Good Practice in Commissioning Accessible Websites, at tinyurl.com/yulfg6, endorsed by the Disability Rights Commission (DRC).

Web authors and developers need to be aware of potential accessibility issues when it comes to developing their content. The Web Content Accessibility Guidelines (WCAG) at www.w3.org/WAI/intro/wcag provide a solid foundation for development and evaluation. However, authors still need an actual understanding of the rationale underpinning various checkpoints and current best practices in the guidelines.

Finally, there has to be an onus on themselves. They need to learn what agents are available to them, and how to configure them to particular needs. For example, larger text or particular color contrast should be used where they should strive to gain a better understanding of the guidelines.

ZEIT ONLINE DATENSCHUTZ

STATISTIK POLITIK WIRTSCHAFT MEINUNG GESELLSCHAFT KULTUR WISSEN DIGITAL STUDIUM KARRIERE LEBENSART NESS AUTO

SPORT INTERNET DATENSCHUTZ MAIL GAMES

DATENSCHUTZ

Tausende demonstrieren für Bürgerrechte im Netz

Für einen besseren Arbeitnehmerdatenschutz und gegen die Gesundheitskarte: 130 Organisatoren hatten zur Demonstration aufgerufen. Sie führten den Überwachungsstaat.

11. Berlin demonstrieren Tausende Demonstration für mehr Datenschutz

Rund 1700 Demonstranten nahmen an dem Protest unter dem Motto "Freiheit statt Angst - Bürger des Überwachungsstaats" in Berlin teil.

Wen Protestveranstaltungen stattfinden, nicht werden aufgerufen, sondern die TeilnehmerInnen

experteer

From data to information.

✓ • Collect data and define your learning problem.

- Data preparation.
- Training a prediction model.
- Checking the performance of your model.

DATENSCHUTZ

Tausende demonstrieren für Bürgerrechte im Netz

Für einen besseren Arbeitnehmerdatenschutz und gegen die Gesundheitskarte: 130 Organisationen hatten zur Demonstration aufgerufen. Sie fürchten den Überwachungsstaat.



© Tim Brakemeier/dpa

In Berlin demonstrieren tausende Demonstranten für mehr Datenschutz.

Rund 7500 Demonstranten nahmen an dem Protestzug unter dem Motto "Freiheit statt Angst – Stoppt den Überwachungswahn" in Berlin teil.

Die Demonstranten wurden dabei unter anderem gegen die Vollerhebung

DATUM 11.9.2010 - 17:20 Uhr

QUELLE ZEIT ONLINE, AFP, dpa

KOMMENTARE 4

EMPFEHLEN E-Mail verschicken | Facebook, Twitter, Buzz ...

ARTIKEL DRUCKEN Druckversion | PDF

SCHLAGWORTER Datenschutz | Demonstration | Datensicherheit | Medienpolitik

NEU IM RESSORT

1. **DATENSCHUTZ** Tausende demonstrieren für Bürgerrechte im Netz
2. **FUTUREZONE** Kurier darf ORF-Portal Futurezone kaufen
3. **URHEBERRECHTE** "Wir haben als Kind gelernt, Teilen ist gut"
4. **SPAM AUF FACEBOOK** "Ein iPad umsonst, ich halte es in den Händen"
5. **IPHONE & IPAD APPS** Mehr Freiheit im App-Store

NEU AUF ZEIT ONLINE

1. **NACHRUUF** Die Freie - Bärbel Bohley ist tot
2. **CDU Union** streitet über konservatives Profil
3. **GEDENKTAG 9/11** Obama warnt vor religiösen Ressentiments
4. **BUNDESLIGA** Der BVB und die Freiburger siegen
5. **DAAD** Das Ende der Ära Bode

ANZEIGE

experteer.de



ZEIT ONLINE

DATENSCHUTZ

Partnersuche

Immobilien

Automarkt

Jobs

Reiseangebote

STARTSEITE POLITIK WIRTSCHAFT MEINUNG GESELLSCHAFT KULTUR WISSEN DIGITAL STUDIUM KARRIERE LEBENSART REISEN AUTO

SPORT

Internet | Datenschutz | Mobil | Games

ANMELDEN | REGISTRIEREN

DATUM 11.9.2010 - 17:20 Uhr

QUELLE ZEIT ONLINE, AFP, dpa

KOMMENTARE 4

EMPFEHLEN E-Mail verschicken | Facebook, Twitter, Buzz ...

ARTIKEL DRUCKEN Druckversion | PDF

SCHLAGWORTE Datenschutz | Demonstration | Datensicherheit | Medienpolitik

NEU IM RESSORT

DATENSCHUTZ Tausende demonstrieren für Bürgerrechte im Netz

FUTUREZONE Kurier darf ORF-Portal Futurezone kaufen

URHEBERRECHTE "Wir haben als Kind gelernt, Teilen ist gut"

SPAM AUF FACEBOOK "Ein iPad umsonst, ich halte es in den Händen"

IPHONE & IPAD APPS Mehr Freiheit im App-Store

NEU AUF ZEIT ONLINE

NACHRUUF Die Freie - Bärbel Bohley ist tot

CDU Union streitet über konservatives Profil

GEDENKTAG 9/11 Obama warnt vor religiösen Ressentiments

BUNDESLIGA Der BVB und die Freiburger siegen

DAAD Das Ende der Ära Bode

ANZEIGE

experteer.de

Tausende demonstrieren für Bürgerrechte im Netz

Für einen besseren Arbeitnehmerdatenschutz und gegen die Gesundheitskarte: 130 Organisationen hatten zur Demonstration aufgerufen. Sie fürchten den Überwachungsstaat.

© Tim Brakemeier/dpa

In Berlin demonstrieren tausende Demonstranten für mehr Datenschutz

Rund 7500 Demonstranten nahmen an dem Protestzug unter dem Motto "Freiheit statt Angst - Stoppt den Überwachungswahn" in Berlin teil.

Die Demonstranten wurden dabei unter anderem gegen die Vollerhebung

- Remove noise.

ZEITUNG ONLINE
DATENSCHUTZ

[Partnersuche](#)
[Immobilien](#)
[Automarkt](#)
[Jobs](#)
[Reiseangebote](#)

[STARTSEITE](#)
[POLITIK](#)
[WIRTSCHAFT](#)
[MEINUNG](#)
[GESELLSCHAFT](#)
[KULTUR](#)
[WISSEN](#)
DIGITAL
[STUDIUM](#)
[KARRIERE](#)
[LEBENSART](#)
[REISEN](#)
[AUTO](#)

SPORT

[Internet](#)
[Datenschutz](#)
[Mobil](#)
[Games](#)

[Anmelden](#) | [Registrieren](#)

DATENSCHUTZ

Tausende demonstrieren für Bürgerrechte im Netz

Für einen besseren Arbeitnehmerdatenschutz und gegen die Gesundheitskarte: 130 Organisationen hatten zur Demonstration aufgerufen. Sie fürchten den Überwachungsstaat.



In Berlin demonstrieren tausende Demonstranten für mehr Datenschutz.

Rund 7500 Demonstranten nahmen an dem Protestzug unter dem Motto "Freiheit statt Angst – Stoppt den Überwachungswahn" in Berlin teil.

Die Demonstranten wurden dabei unter anderem gegen die Vollerhebung

DATUM 11.9.2010 - 17:20 Uhr
 QUELLE ZEIT ONLINE, AFP, dpa
 KOMMENTARE 4
 EMPFEHLEN E-Mail verschicken | Facebook, Twitter, Buzz ...
 ARTIKEL DRUCKEN Druckversion | PDF
 SCHLAGWORTE Datenschutz | Demonstration | Datensicherheit | Medienpolitik

NEU IM RESSORT
 1. **DATENSCHUTZ** Tausende demonstrieren für Bürgerrechte im Netz
 2. **FUTUREZONE** Kurier darf ORF-Portal Futurezone kaufen
 3. **URHEBERRECHTE** "Wir haben als Kind gelernt, Teilen ist gut"
 4. **SPAM AUF FACEBOOK** "Ein iPad umsonst, ich halte es in den Händen"
 5. **IPHONE & IPAD APPS** Mehr Freiheit im App-Store

NEU AUF ZEIT ONLINE
 1. **NACHRUUF** Die Freie - Bärbel Bohley ist tot
 2. **CDU Union** streitet über konservatives Profil
 3. **GEDENKTAG 9/11** Obama warnt vor religiösen Ressentiments
 4. **BUNDESLIGA** Der BVB und die Freiburger siegen
 5. **DAAD** Das Ende der Ära Bode

ANZEIGE


experteer.de
 

- Remove noise.
- Convert text to vectors.

From texts to vectors


If we looked at two words only:

Sunny weather



High performance computing

Aaron



Binary bag of words

- Imagine a n-dimensional space.
- Each dimension = one possible word in texts.
- Entry in vector is one, if word occurs in text.

$$b_{i,j} = \begin{cases} 1 & \forall x_i \in d_j \\ 0 & \text{else} \end{cases}$$

- Problem:
 - Number of word occurrences not accounted for.

Term Frequency

- Imagine a n-dimensional space.
- Each dimension = one possible word in texts.
- Entry in vector equal to the words frequency.

$$b_{i,j} = n_{i,j}$$

- Problem:
 - Common words dominate vectors.

TF with stop wording

- Imagine a n-dimensional space.
- Each dimension = one possible word in texts.
- Filter stopwords.
- Entry in vector equal to the words frequency.

$$b_{i,j} = n_{i,j}$$

- Problem:
 - Common and uncommon words with same weight.

TF- IDF

- Imagine a n-dimensional space.
- Each dimension = one possible word in texts.
- Filter stopwords.
- Entry in vector equal to **the weighted** frequency.

- Problem:
$$b_{i,j} = n_{i,j} \times \log\left(\frac{|D|}{|\{d : t_i \in d\}|}\right)$$
 - Long texts get larger values.

Normalized TF- IDF

- Imagine a n-dimensional space.
- Each dimension = one possible word in texts.
- Filter stopwords.
- Entry in vector equal to the weighted frequency.
- **Normalize vectors.**

$$b_{i,j} = \frac{n_{i,j}}{\sum_k n_{k,j}} \times \log\left(\frac{|D|}{|\{d : t_i \in d\}|}\right)$$

- Problem:
 - Additional domain knowledge ignored.

Reality

- There are a few more words in news.
- Use all relevant features/ signals available.
 - Words.
 - Header fields.
 - Characteristics of publishing url.
 - ...
- Usually pipeline of feature extractors.

From data to information.



- Collect data and define your learning problem.




- Data preparation.


- Training a prediction model.
- Checking the performance of your model.

Step 2: Similarity


Euclidian




Euclidian




Euclidian








Cosine


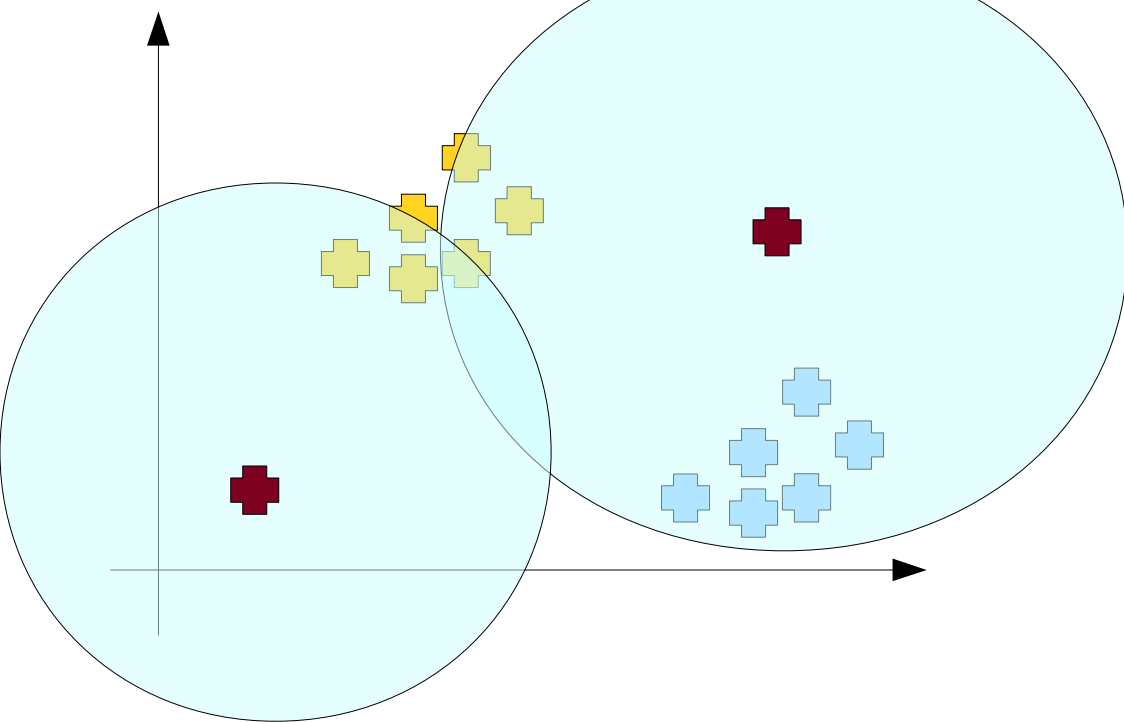




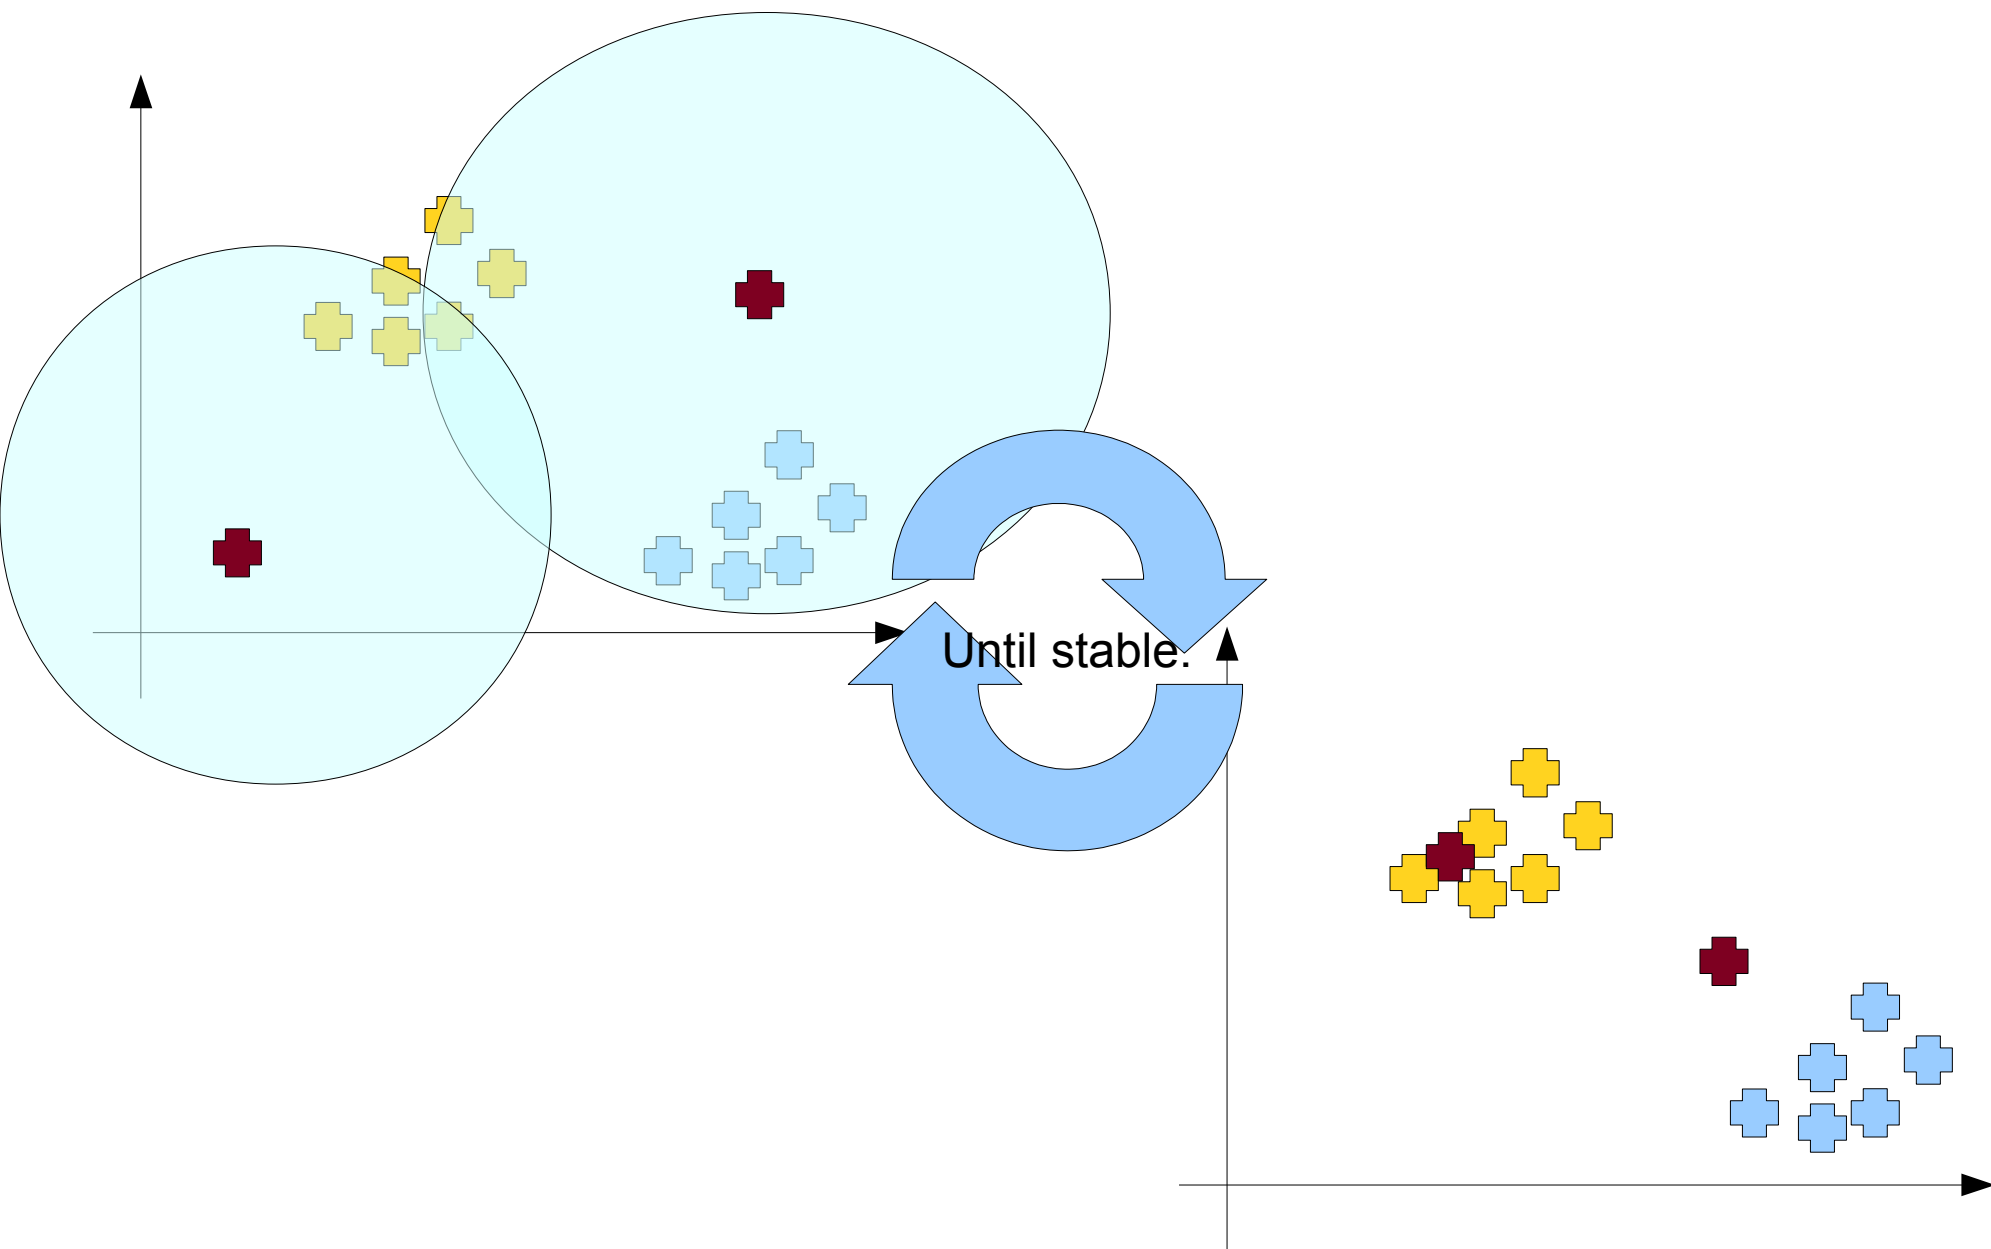
Step 3: Clustering











Reality

- Seed selection.
- Choice of initial k .
- Continuous updates.
- Regular addition of clusters.

From data to information.


- ✓ • Collect data and define your learning problem.
- ✓ • Data preparation.
- ✓ • Training a prediction model.
 - Checking the performance of your model.

Evaluation

- Compare against gold standard.
- Use quality measures.
- Manual inspection.

From data to information.

- ✓ • Collect data and define your learning problem.
- ✓ • Data preparation.
- ✓ • Training a prediction model.
- ✓ • Checking the performance of your model.





What else does Mahout have to offer.

Identify dominant topics


- Given a dataset of texts, identify main topics.

Algorithms: Parallel LDA

- Examples:
 - Dominant topics in set of mails.
 - Identify news message categories.

Assign items to defined categories.

- Given pre-defined categories, assign items to it.





By freezelight, <http://www.flickr.com/photos/63056612@N00/155554663/>

[Advanced Image Search](#)

 SafeSearch: [Moderate](#) ▼

 Images ☐ [Hide options](#)

Results


> Any size

[Medium](#)
[Large](#)
[Icon](#)
[Larger than...](#)
[Exactly...](#)

> Any type

[Face](#)
[Photo](#)
[Clip art](#)
[Line drawing](#)

> Any color

[Full color](#)
[Black and white](#)
[Specific color](#)

 Related searches: [oakland raiders](#)

Oakland Airport

625 x 471 - 103k - jpg

[visitingdc.com](#)

Oakland Ranks Fifth

538 x 359 - 44k - jpg

[bayareahomegirl.com](#)

Oakland

900 x 600 - 171k - jpg

[globalsecurity.org](#)

OAKLAND First

400 x 400 - 27k

[bayassociation.org](#)

Oakland Gaudy Lexus 2

500 x 340 - 73k - jpg

[lexusenthusiast.com](#)

The Oakland Baseball

450 x 305 - 39k - jpg

[sportsbusinesssims.com](#)

Oakland

600 x 320 - 48k - jpg

[webpages.scu.edu](#)


Learn more about

550 x 366 - 56k - jpg

[tripadvisor.com](#)

oakland

Search images

[Advanced Image Search](#)

SafeSearch: [Moderate](#) ▼

[Images](#) > [Face](#) ☐ [Hide options](#)

Results 1 - 20 of about 1,400,000 (0.1)

> Any size

[Medium](#)

[Large](#)

[Icon](#)

[Larger than...](#)

[Exactly...](#)

[Any type](#)

[Face](#)

[Photo](#)

[Clip art](#)


[Line drawing](#)

> Any color

[Full color](#)

[Black and white](#)

[Specific color](#)



[Reset options](#)

Related searches: [oakland raiders](#)



Oakland, CA 94621
513 x 545 - 13k - gif
[nflfootballstadiums.com](#)



All Graphics »
262 x 278 - 43k - gif
[coolchaser.com](#)



Oakland Sideshow and
720 x 480 - 44k
[channels.com](#)



Detroit Tigers v
594 x 396 - 51k
[zimbio.com](#)



Oakland. by Davey D
359 x 512 - 26k - jpg
[sfbayview.com](#)



Detroit Tigers v
443 x 594 - 74k
[zimbio.com](#)



Dallas Cowboys v
594 x 404 - 60k
[zimbio.com](#)



Distributed by Tubemogul. The
720 x 480 - 18k
[channels.com](#)



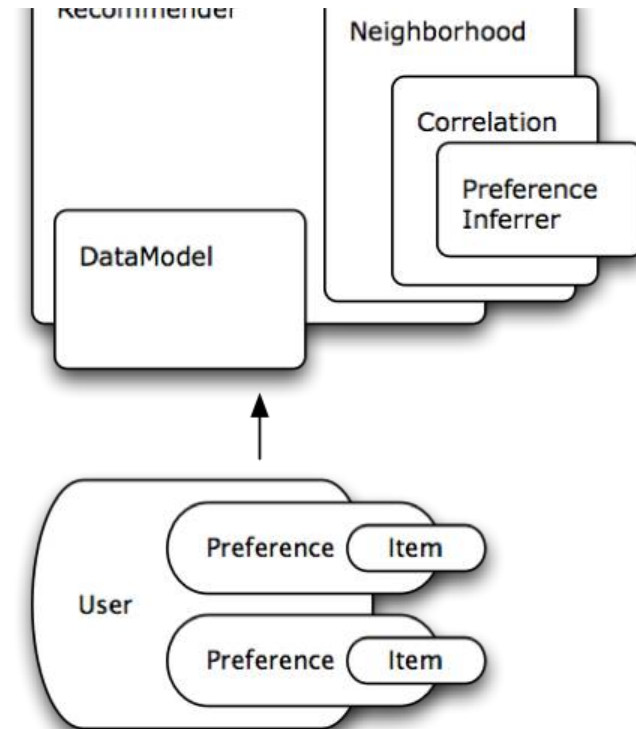
oakland@coe.ufl.edu
379 x 471 - 110k - jpg
[coe.ufl.edu](#)



#20 of the Oakland
467 x 594 - 86k
[zimbio.com](#)

Recommendation mining.

- Collaborative filtering.



Show most relevant ads

clearasil®

erwandte Suchbegriffe: clearasil pickelstift, world of warcraft.

este Ergebnisse

« Zurück | Seit

1:



Clearasil 44161 Tiefenreinigung Antibakterielle Reinigungspads, 60er

Neu kaufen: **EUR 5,99** 2 Angebote ab EUR 5,15

Lieferung bis **Samstag, 22. März**: Bestellen Sie innerhalb der nächsten **23 Stunden**
Kostenlose Lieferung möglich.

★★★★★ (1)

Drogerie & Bad: Alle 13 Artikel ansehen

2:




Clearasil Ultra Anti-Pickel Reinigungspads, 65 Stück von Clearasil (Ba

Neu kaufen: **EUR 7,99**

Gewöhnlich versandfertig in 1 bis 3 Wochen.
Kostenlose Lieferung möglich.

Drogerie & Bad: Alle 13 Artikel ansehen

3:




World of WarCraft: Wrath of the Lich King (Add-on) von Vivendi Unive
Vista / XP)

Neu kaufen: **EUR 39,99**

Vorbestellbar
Kostenlose Lieferung möglich.

Games: Alle Artikel ansehen

Show most relevant ads



[See larger image](#)

[Share your own customer images](#)

[Publisher: learn how customers can search inside this book.](#)

~ [Otis Gospodnetic](#) (Author), [Erik Hatcher](#) (Author)
★★★★☆ (20 customer reviews)

List Price: ~~\$44.95~~

Price: **\$29.67** & this item ships for **FREE with Super Saver Shipping.** [Details](#)
You Save: **\$15.28 (34%)**

In Stock.

Ships from and sold by **Amazon.com**. Gift-wrap available.

[23 new](#) from \$19.99 [15 used](#) from \$17.11


	List Price:	Our Price:	Other Offers:
Paperback (2)	\$44.99	\$29.69	

Frequently Bought Together

Customers buy this book with [Building Search Applications: Lucene, LingPipe, and Gate](#) by Manu Konchady



+



Price For Both: \$70.12



Add both to Cart


Add both to Wish List

[Show availability and shipping details](#)

Customers Who Bought This Item Also Bought



[Solr 1.4 Enterprise Search Server](#) by David




[Hadoop: The Definitive Guide](#) by Tom White



[Hibernate Search in Action](#) by Emmanuel




[Introduction to Information Retrieval](#) by



[Collective Intelligence in Action](#) by Satnam Alag

Recommending places



Thanks to Falko Menge for the pictures of Brussels.

Recommending people

People you may know

Shane Curcuru
Conference Lead at Apache
Softw

Ugo
Com
Profe

Arjé
CTO
lead,
Mem
Foun

Invite | x

Invite | x

Invite | x



Recommendation mining.

- Online collaborative filtering on single machine.
- Offline Map/Reduce based version.
- Content similarity can be integrated.
- Based on former Taste project.

Frequent pattern mining

- Given groups of items, find commonly co-occurring items.
- Examples:
 - In shopping carts find items bought together.
 - In query logs find queries issued in one session.



By quinnanya, <http://www.flickr.com/photos/quinnanya/2806883231/>



By crypto, <http://www.flickr.com/photos/crypto/3201254932/sizes//>

By libraryman, <http://www.flickr.com/photos/libraryman/78337046/sizes//>



A close-up photograph of a pile of wooden alphabet blocks. The blocks are made of light-colored wood and are painted with various colors including red, green, blue, and yellow. Some blocks have letters, while others have geometric patterns or abstract designs. The blocks are scattered and overlapping, creating a textured and colorful background.

Requirements to get started

March 14, 2009 by Artful Magpie
<http://www.flickr.com/photos/kmtucker/3355551036/>



[AWS](#)[Products](#)[Developers](#)[Community](#)[Support](#)[Account](#)[Products & Services](#)

Amazon Elastic Compute Cloud (Amazon EC2)

Amazon EC2 Details

- **EC2 Overview**
- [FAQs](#)
- [Amazon EC2 SLA](#)
- [EC2 Instance Types](#)

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers.

Amazon EC2's simple web service interface allows you to obtain and configure capacity with minimal friction. It provides you with complete control of your computing resources and lets you run on Amazon's proven computing environment. Amazon EC2 reduces the time required to obtain

[Sign Up For Amazon EC2](#)

Amazon Elastic MapReduce

Amazon Elastic MapReduce is a web service that enables businesses, researchers, data analysts, and developers to easily and cost-effectively process vast amounts of data. It utilizes a hosted Hadoop framework running on the web-scale infrastructure of Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Simple Storage Service (Amazon S3).

Using Amazon Elastic MapReduce, you can instantly provision as much or as little capacity as you like to perform data-intensive tasks for applications such as web indexing, data mining, log file analysis, machine learning, financial



(Thanks to Thilo for helping set up the cluster, Thanks to packet and masq for two of the three machines.)

Requirements to get started

March 14, 2009 by Artful Magpie
<http://www.flickr.com/photos/kmtuck/3355551036/>

MacBook

Why go for Apache Mahout?

Jumpstart your project with proven code.

January 8, 2008 by dreizehn28
<http://www.flickr.com/photos/1328/2176949559>

A photograph of two men in a dimly lit room. The man on the left is sitting on a couch, wearing a dark hoodie, with his hands near his face in a thoughtful or listening pose. The man on the right is sitting in a chair, wearing a dark t-shirt and glasses, with his hand on his chin, looking towards the first man. The background is a plain, light-colored wall. The overall mood is contemplative and focused.

Discuss ideas and problems online.

November 16, 2005 [phil h]
<http://www.flickr.com/photos/hi-phi/64055296>

WARNING





Become a committer.





Sebastian Schelter
Jake Mannix
Benson Margulies
Robin Anil
David Hall
AbdelHakim Deneche
Karl Wettin
Sean Owen
Grant Ingersoll
Otis Gospodnetic
Drew Farris
Jeff Eastman
Ted Dunning
Isabel Drost



Become a committer: Of Apache Mahout



Emeritus:
Niranjan Balasubramanian
Erik Hatcher
Ozgur Yilmazel
Dawid Weiss

*-user@mahout.apache.org

*-dev@mahout.apache.org



Image by: Patrick McEvoy

Interest in solving hard problems.

Being part of lively community.

Engineering best practices.

Bug reports, patches, features.

Documentation, code, examples.



FOSDEM

FREE AND OPEN SOURCE SOFTWARE DEVELOPERS' EUROPEAN MEETING

FREE AND OPEN SOURCE SOFTWARE DEVELOPERS' EUROPEAN MEETING

FOSDEM is the biggest free and non-commercial event organized by and for the community. Its goal is to provide Free and Open Source developers a place to meet. No registration necessary.

Front page

About

FOSDEM

Sponsors

Contact

[Home](#)

Accepted devrooms

We were very pleased with the devroom proposals we received, as many were collaborative proposals with ample opportunity for cross-pollination between projects. As a side effect, this did make the decision process even harder.

- **Data Analytics**

Open Source tools for Data Analytics. This includes Machine Learning, Text Mining and Natural Language Processing, Computer Vision, Recommender Engines, Visualisation, Data Mining, Large Scale and distributed data processing, etc.



Thanks to Tim Lossen et. al for taking amazing pictures of the conf.

Berlin Buzzwords 2011

Search/ Store/ Scale

May/ June 2011

Thanks to Tim Lossen et. al for taking amazing pictures of the conf.

Berlin Buzzwords 2011

Search/ Store/ Scale

May/ June 2011

Thanks to Tim Lossen et. al for taking amazing pictures of the conf.

[illegible][illegible]

*-user@mahout.apache.org

*-dev@mahout.apache.org



Image by: Patrick McEvoy

Interest in solving hard problems.

Being part of lively community.

Engineering best practices.

Bug reports, patches, features.

Documentation, code, examples.

