

How (much) to trust the Wikipedia?

Luca de Alfaro

UC Santa Cruz

Joint work with Bo Adler, Ian Pye

With contributions by Jason Benterou,
Krishnendu Chatterjee, Marco Faella, Vishwanath Raman

CITRIS, February 2008





WIKIPEDIA
The Free Encyclopedia

navigation

- [Main Page](#)
- [Community portal](#)
- [Current events](#)
- [Recent changes](#)
- [Random page](#)
- [Help](#)
- [Donations](#)

[article](#)[discussion](#)[view source](#)[history](#)

Italian cuisine

Revision as of 04:20, 30 January 2007 by [69.210.149.199](#) (Talk)
(diff) ←[Older revision](#) | [Current revision](#) (diff) | [Newer revision](#) → (diff)

Italian cuisine is extremely varied: the country of [Italy](#) was only unified in **1861**, and its cuisines reflect the cultural variety of its [regions](#) and its diverse history (with culinary influences from Greek, Roman, Norman and Arab civilizations). Italian cuisine is imitated all over the world. It also is way better then French food, the losers.

To a certain extent, there is really no such thing as

This article is part
of the [Cuisine](#)
series

**Preparation
techniques and
cooking items**

[Techniques - Utensils](#)

Italian cuisine – The UCSC Wikipedia Trust Project

Log in / create account

article discussion view source history

Italian cuisine

Revision as of 04:20, 30 January 2007 by 69.210.149.199 (Talk)
(diff) ←Older revision | Current revision (diff) | Newer revision→ (diff)

Italian cuisine is extremely varied: the country of Italy was only unified in 1861, and its cuisines reflect the cultural variety of its regions and its diverse history (with culinary influences from Greek, Roman, Norman and Arab civilizations). Italian cuisine is imitated all over the world. It also is way better than French food, the losers.

To a certain extent, there is really no such thing as

This article is part of the Cuisine series

Preparation techniques and cooking items

Techniques - Utensils

navigation

- Main Page
- Community portal
- Current events
- Recent changes
- Random page
- Help
- Donations

Anyone can edit the Wikipedia

- This has been the key to its success (get knowledge from all sources).
- But if anyone could have added it, how do we know whether we can trust it?
 - Typos, misleading information, slander, information deletion, ...

Author Reputation + Text Trust

Author Reputation:


- Promote constructive behavior
- Provide information on author reliability

Text Trust:

- Give a guide to text reliability
- Provide alert for attempts to tamper with content

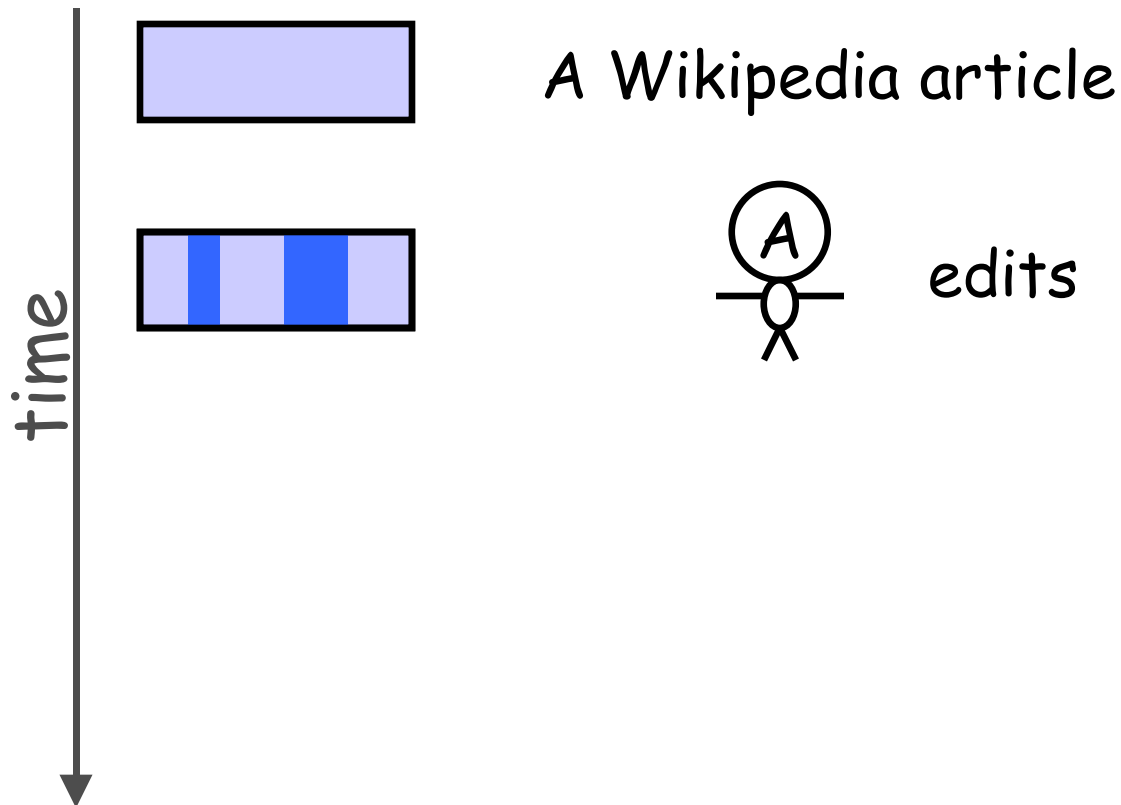
Content-driven reputation

- Authors of long-lived contributions gain reputation
- Authors of reverted contributions lose reputation



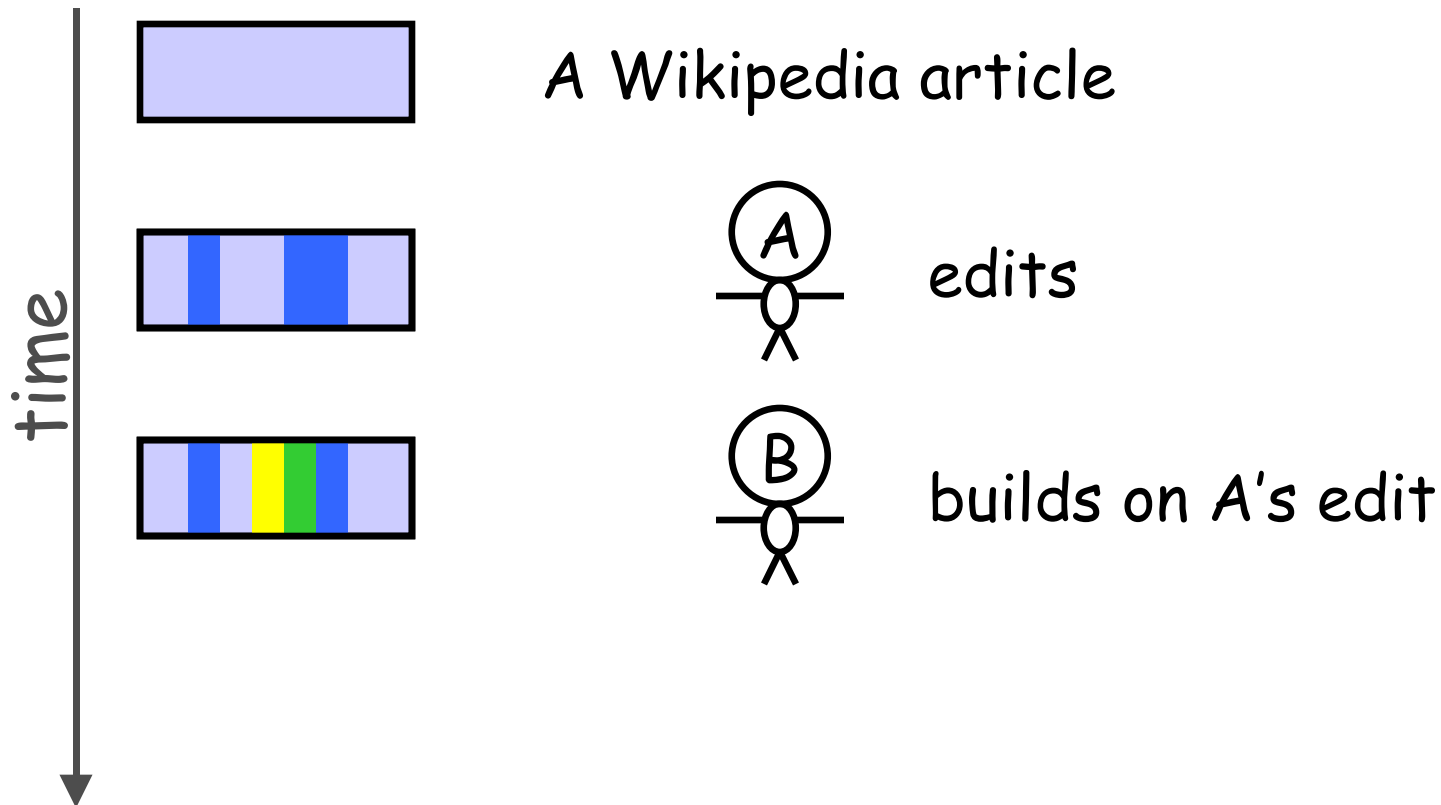
Content-driven reputation

- Authors of long-lived contributions gain reputation
- Authors of reverted contributions lose reputation



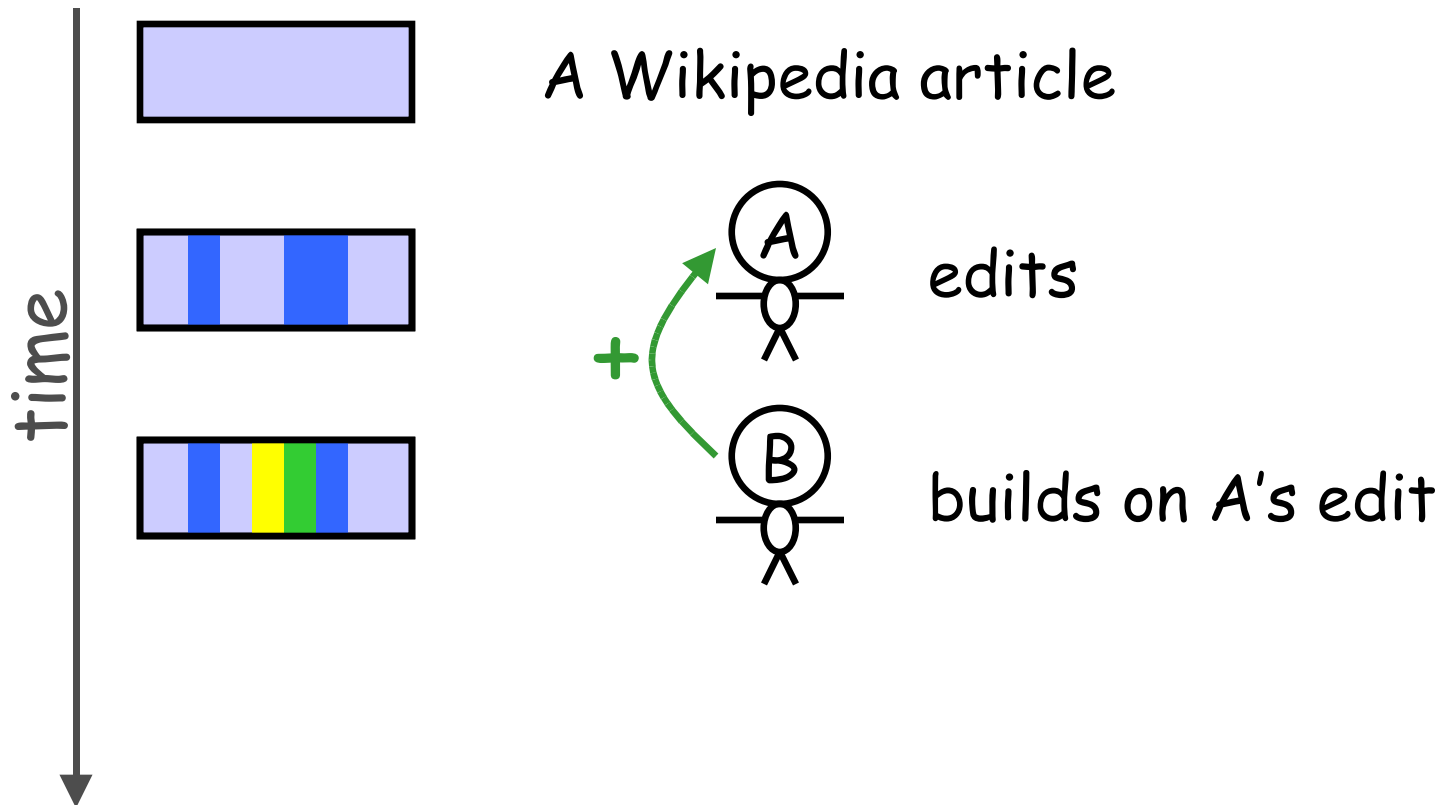
Content-driven reputation

- Authors of long-lived contributions gain reputation
- Authors of reverted contributions lose reputation



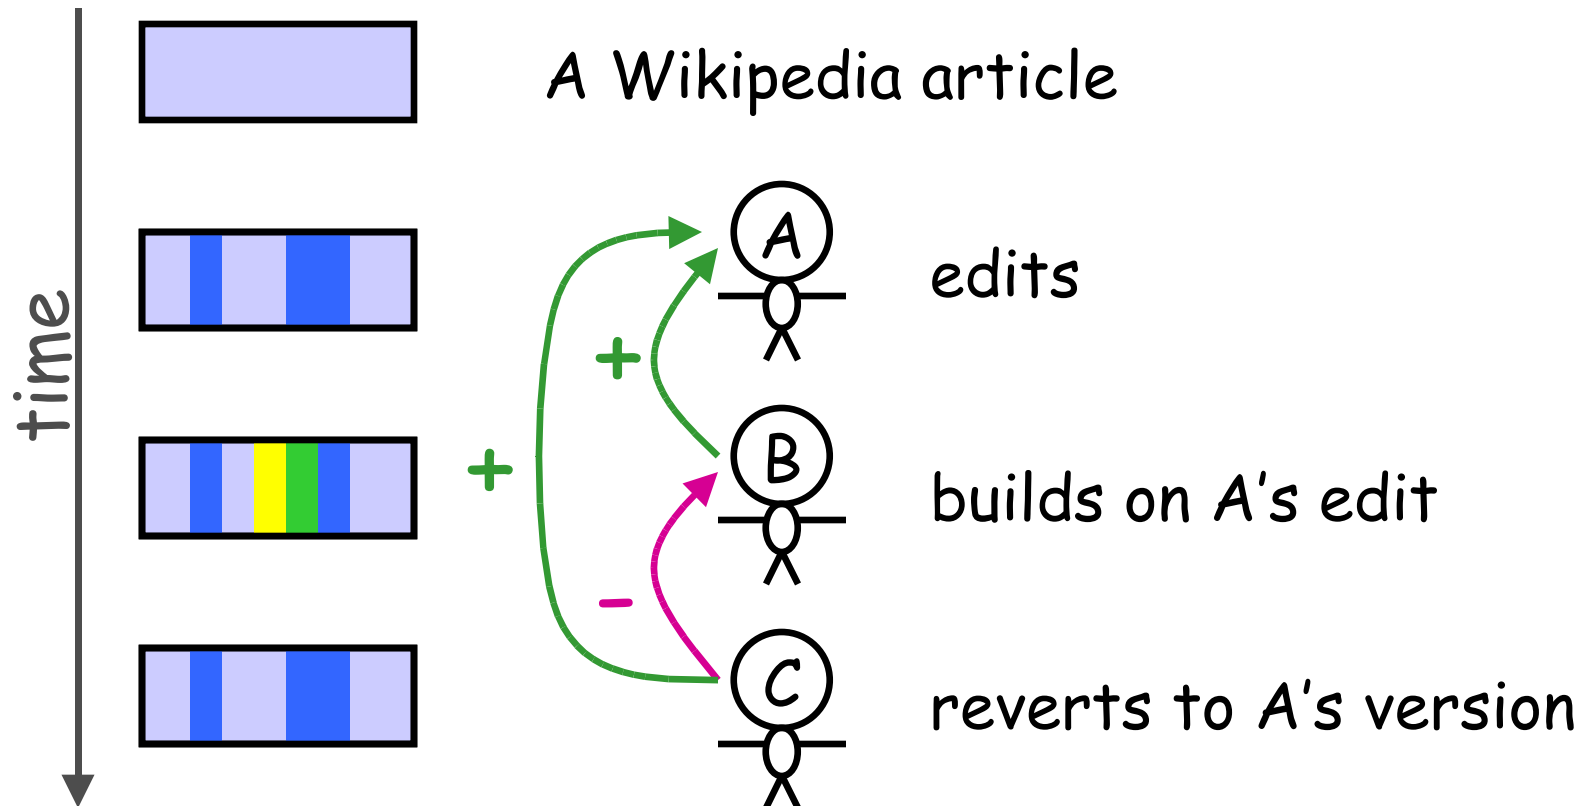
Content-driven reputation

- Authors of long-lived contributions gain reputation
- Authors of reverted contributions lose reputation



Content-driven reputation

- Authors of long-lived contributions gain reputation
- Authors of reverted contributions lose reputation



Why content-driven reputation?


- **No change to Wikipedia user experience**
 - Transparent to casual users
 - No need to explicitly rate others and be rated, less stress
- **Everybody votes (via their edits)**
 - Technically dedicated users do not carry more weight
- **Less prone to reputation wars**
- **We have the data**
 - The whole Wikipedia history, many million article versions.

We would like to avoid displaying reputation values directly; our main use for reputation is as a means of computing text trust.

Goals of reputation in Wikipedia

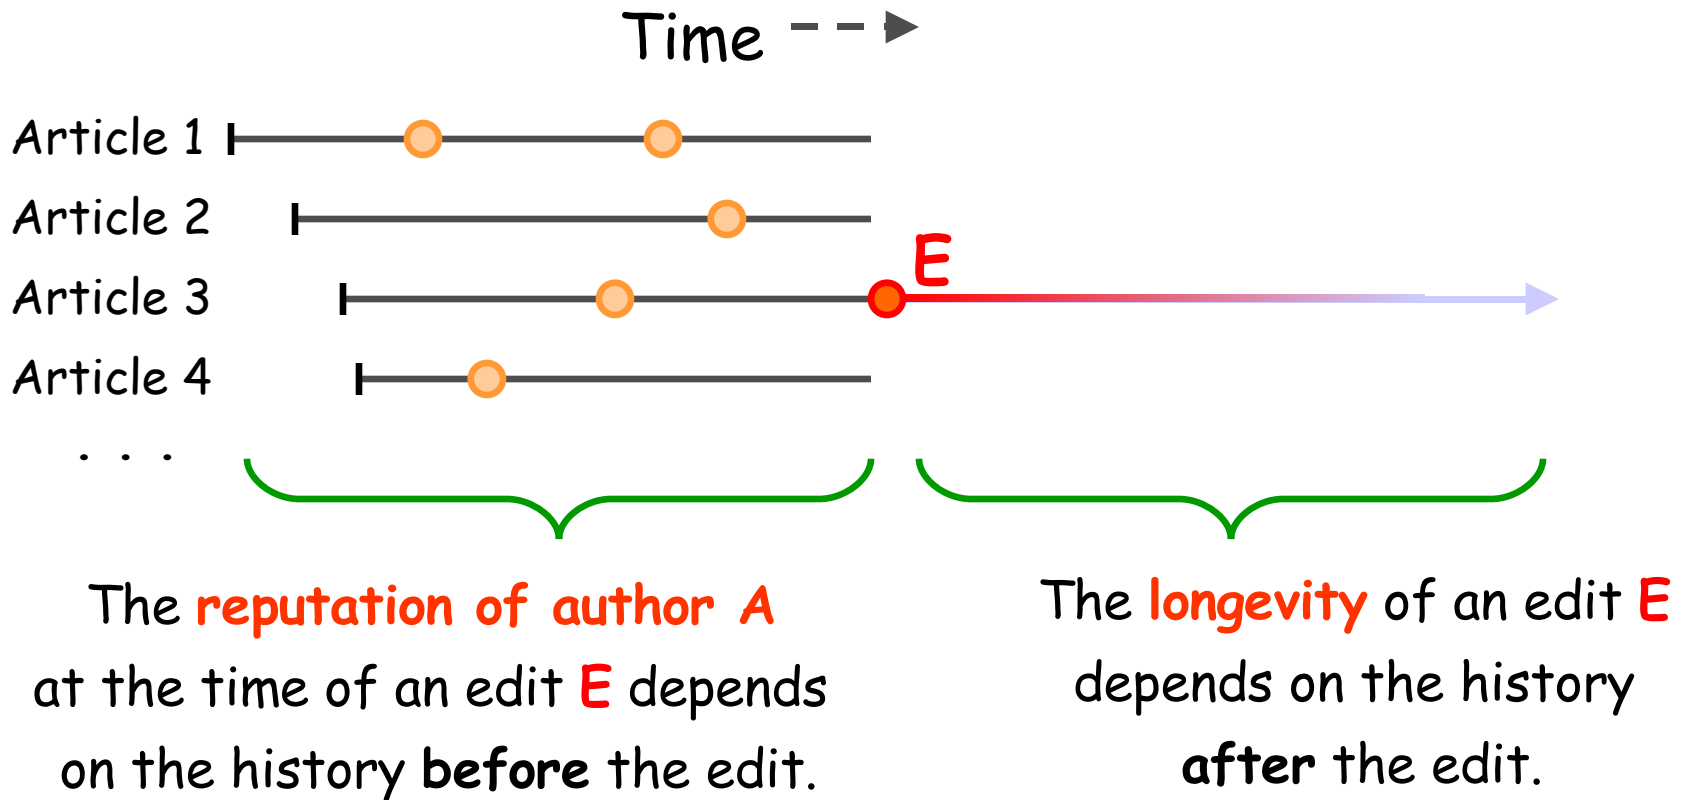
- **Prescriptive:** encourages people to behave in a good way (e.g., Ebay system).
 - We want to encourage lasting contributions.
- **Descriptive:** gives information to users (e.g., Pagerank, Ebay system).
 - Author reputation can be used as a rough guide to the trust in new text/edits.
- **Predictive:** Is reputation a good predictor for future behaviour?
Few systems make this claim!
 - We use this as our evaluation criterion, and we show that our reputation can predict edit quality.

Does our reputation have predictive value?



○ = edits by user **A**

Does our reputation have predictive value?




We will show a correlation between **author reputation** and **edit longevity**

Building a content-driven reputation system

What is a "contribution"?

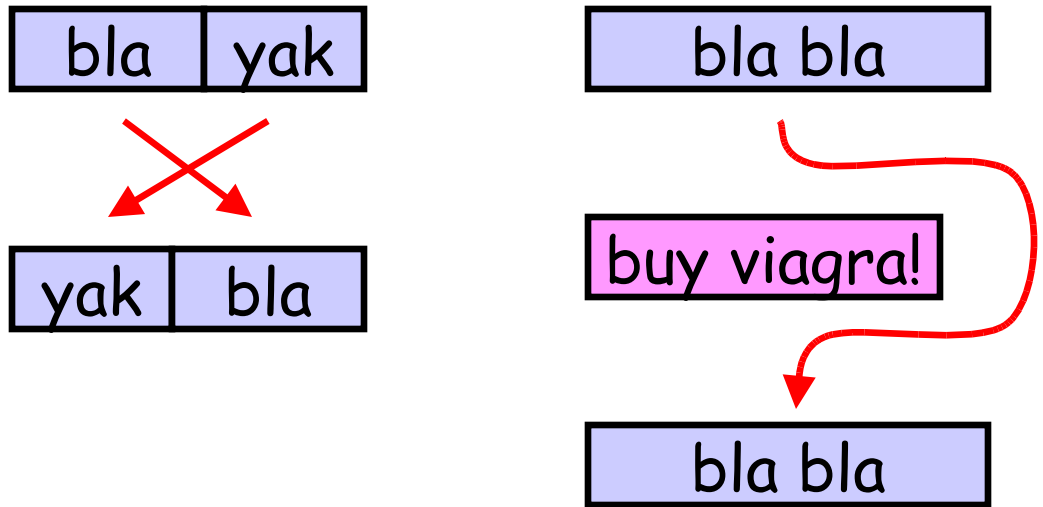
Text



We measure how long the added text survives.

Based on text tracking.

Edit



We measure how long the "edit" (reorganization) survives.

Based on edit distance.

Text

version 9

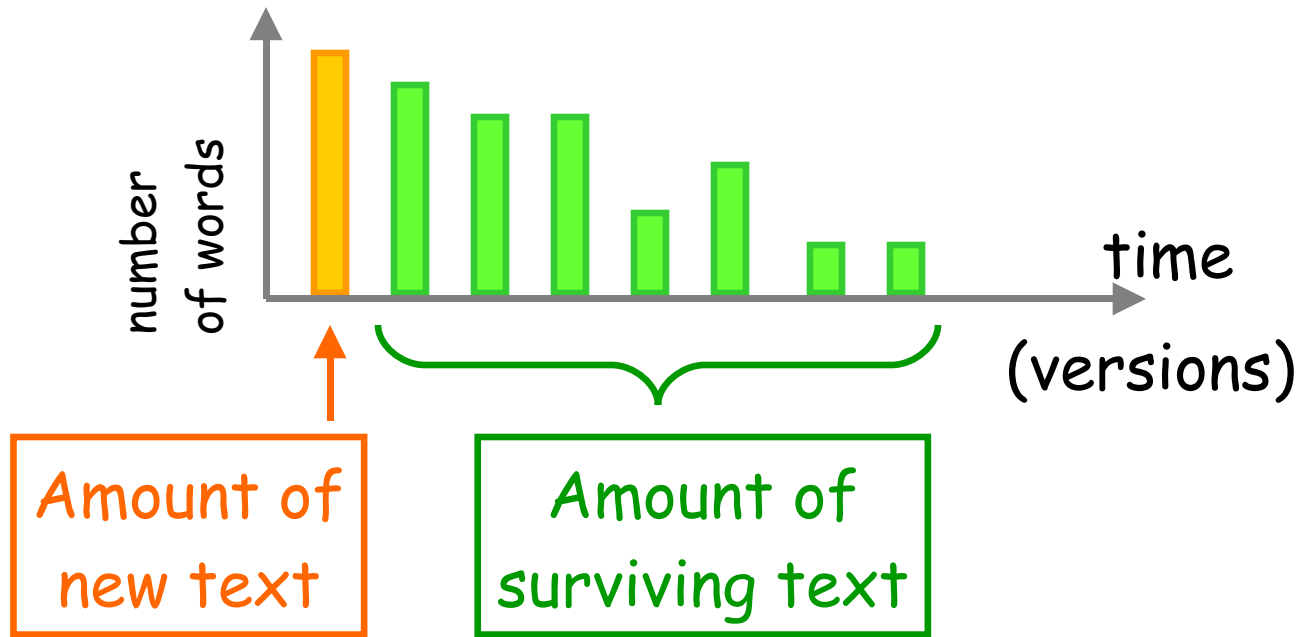
bla	bla	wuga	boink
5	8	9	6

version 10

bla	bla	wuga	wuga	wuga	boink
5	8	10	10	9	6

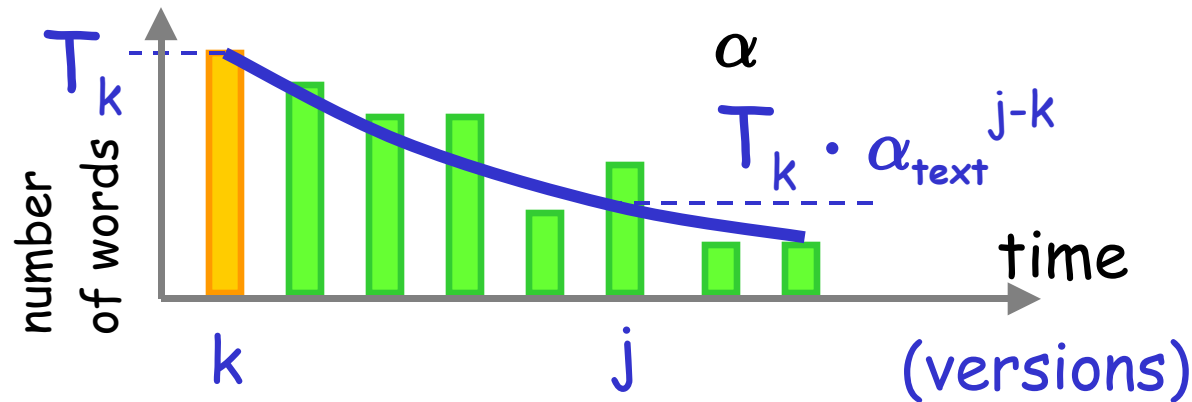
We label each word with the version where it was introduced. This enables us to keep track of how long it lives.

Text



The life of the text introduced at a revision.

Text: Longevity

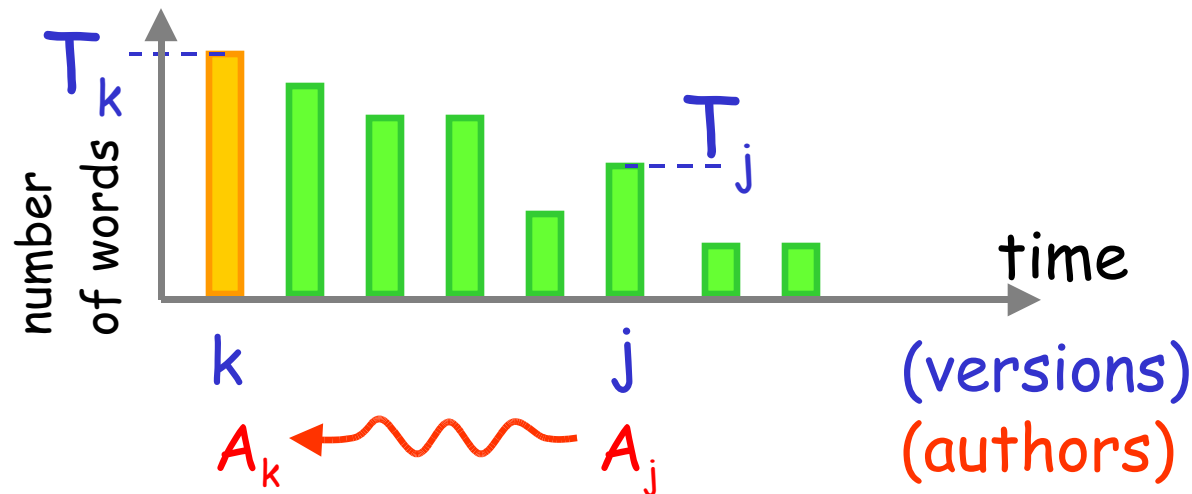


Text Longevity α_{text} :

- We find the $\alpha_{\text{text}} \in [0,1]$ that yields the best geometrical approximation for the amount of residual text.
- We call α_{text} the *text longevity* of edit k .

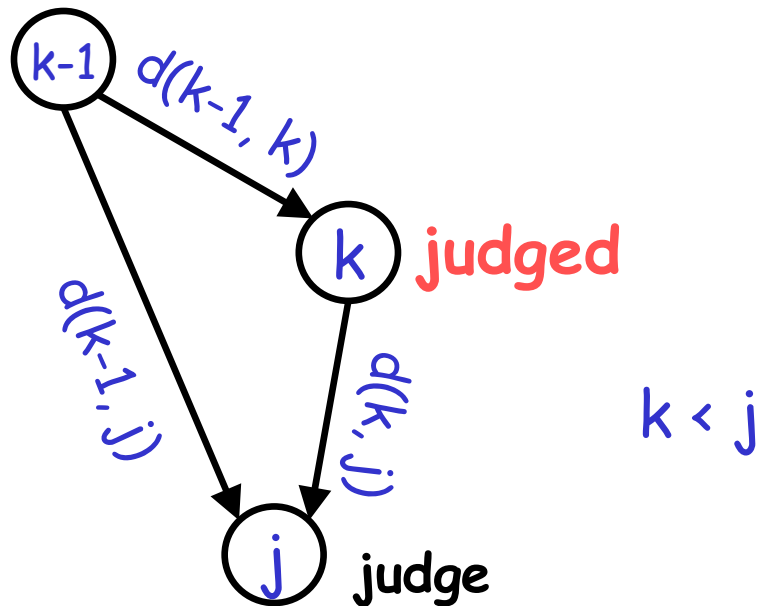
$\alpha_{\text{text}} \simeq 1$: long-lived; $\alpha_{\text{text}} = 0$: removed immediately.

Text: Reputation update




- As a consequence of version j , we increment the reputation of A_k in proportion to T_j , and to the reputation of A_j .

Edit

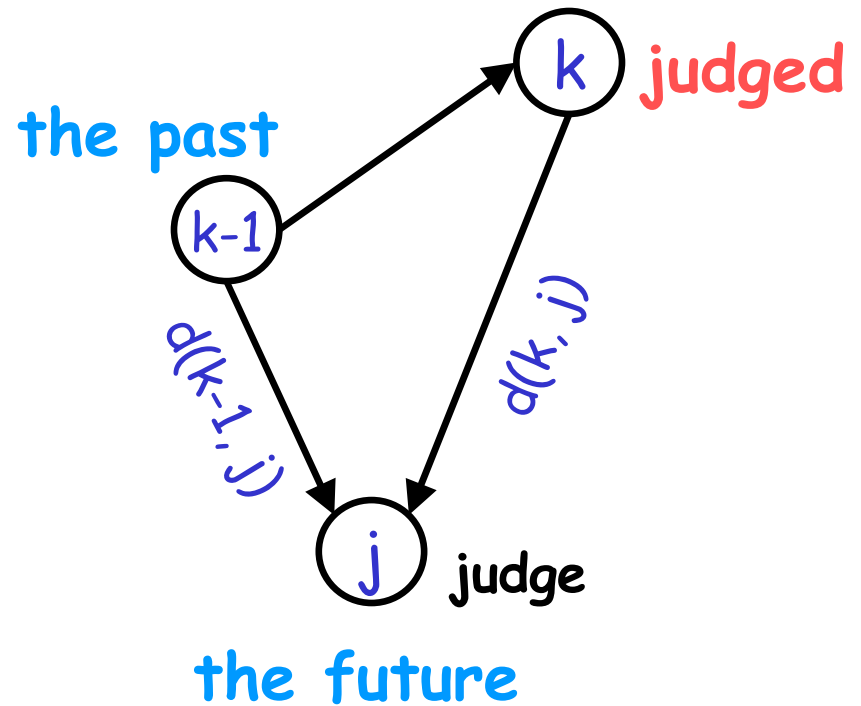


We compute the edit distance between versions $k-1$, k , and j , with $k < j$ (see paper for details on the distance).

Edit: good or bad?

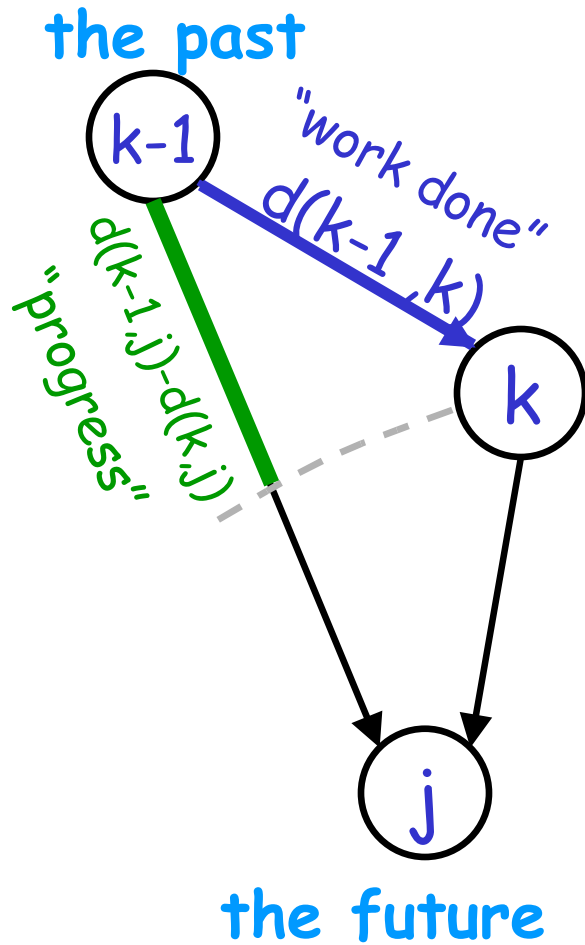


k is **good**: $d(k-1, j) > d(k, j)$
"k went towards the future"



k is **bad**: $d(k-1, j) < d(k, j)$
"k went against the future"

Edit: Longevity



Edit Longevity:

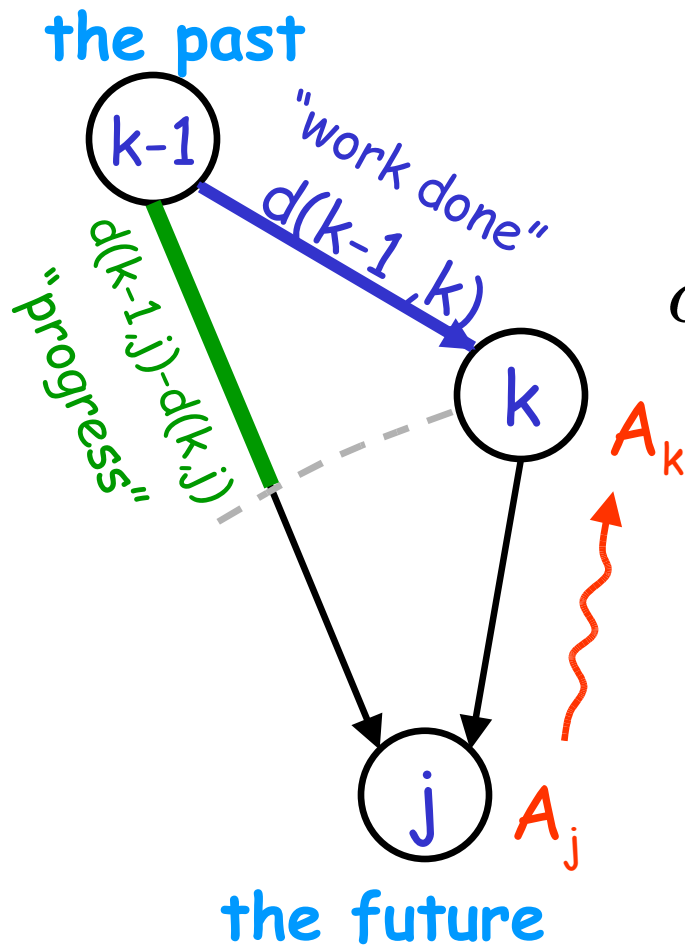
$$\alpha_{edit} = \frac{d(k-1, j) - d(k, j)}{d(k-1, k)}$$

Edit longevity measures the fraction of change that agrees with the future page evolution.

- $\alpha_{edit} \simeq +1$: edit k is good
- $\alpha_{edit} \simeq -1$: edit k is reverted

Corollary: we can detect reversions automatically.

Edit: Updating reputation



$$\alpha_{edit} =$$

Edit Longevity:

$$\frac{d(k-1, j) - d(k, j)}{d(k-1, k)}$$

Reputation update:

The reputation of A_k

- increases if $\alpha_{edit} > 0$,
- decreases if $\alpha_{edit} < 0$.

(see our WWW07 paper for details)

Results: English Wikipedia, in detail

% of edits below a given longevity →

	Bin	%_data	$l < 0.8$	$l < 0.4$	$l < 0.0$	$l < -0.4$	$l < -0.8$
$\log(1 + \text{reputation})$ ↓	0	16.922	93.11	91.65	89.15	83.76	73.53
	1	1.191	77.24	69.83	65.60	61.11	56.00
	2	1.335	69.53	57.08	49.79	45.71	41.25
	3	1.627	38.00	28.61	20.23	16.16	13.62
	4	2.780	32.84	22.31	13.32	9.57	8.04
	5	4.408	41.70	15.76	5.90	3.80	2.57
	6	6.698	29.40	16.74	7.54	4.35	3.12
	7	8.281	32.04	15.16	5.44	2.25	1.40
	8	12.233	34.06	16.64	6.78	3.79	2.73
	9	44.524	32.55	15.51	5.05	1.88	1.14

Results: English Wikipedia, in detail

% of edits below a given longevity

	Bin	% data	$l < 0.8$	$l < 0.4$	$l < 0.0$	$l < -0.4$	$l < -0.8$
low rep	0	16.922	93.11	91.65	89.15	83.76	73.53
	1	1.191	77.24	69.83	65.60	61.11	56.00
log(1 + reputation) ↓	2	1.335	69.53	57.08	49.79	45.71	41.25
	3	1.627	38.00	28.61	20.23	16.16	13.62
	4	2.780	32.84	22.31	13.32	9.57	8.04
	5	4.408	41.70	15.76	5.90	3.80	2.57
	6	6.698	29.40	16.74	7.54	4.35	3.12
	7	8.281	32.04	15.16	5.44	2.25	1.40
	8	12.233	34.06	16.64	6.78	3.79	2.73
	9	44.524	32.55	15.51	5.05	1.88	1.14

Short-Lived

Predictive power of low reputation

- **Recall:** Low-reputation authors (those in the bottom 1/5 of reputation) account for 18.1% of the edits, and for 82.9% of short-lived edits.
- **Precision:** An edit has a 5.7% probability of being short-lived. However, if the edit is done by a low-reputation author, this probability raises to 48.9% .

Recall and precision are high, even though *this is a human behavior prediction problem!*

From author reputation
to text trust

Text Trust: Principles

Compute trust at **word** level to track changes:

- New text starts with a trust value proportional to the author's reputation.
- Text trust can raise when the text is revised by higher-reputation authors.

Trust must alert visitors to information tampering.

Display trust via background colors:

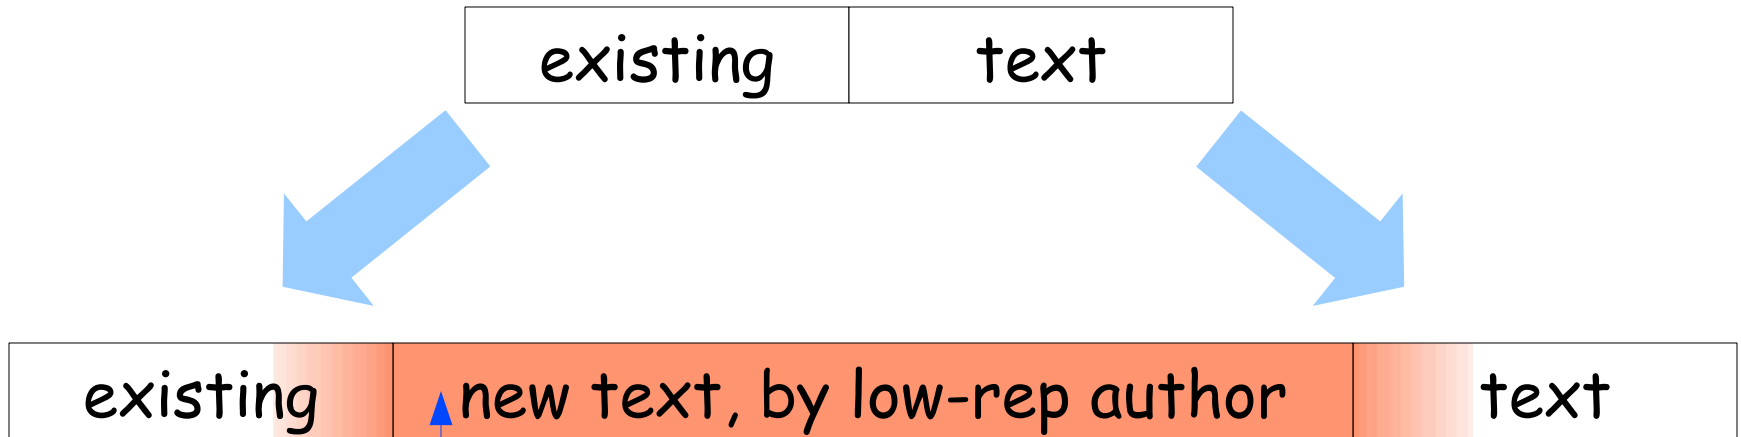
trusted text	decreasing trust	untrusted text
--------------	------------------	----------------

[Related work: Zeng, Alhousaini, Ding, Fikes, McGuinness 06]

Trust: New Text

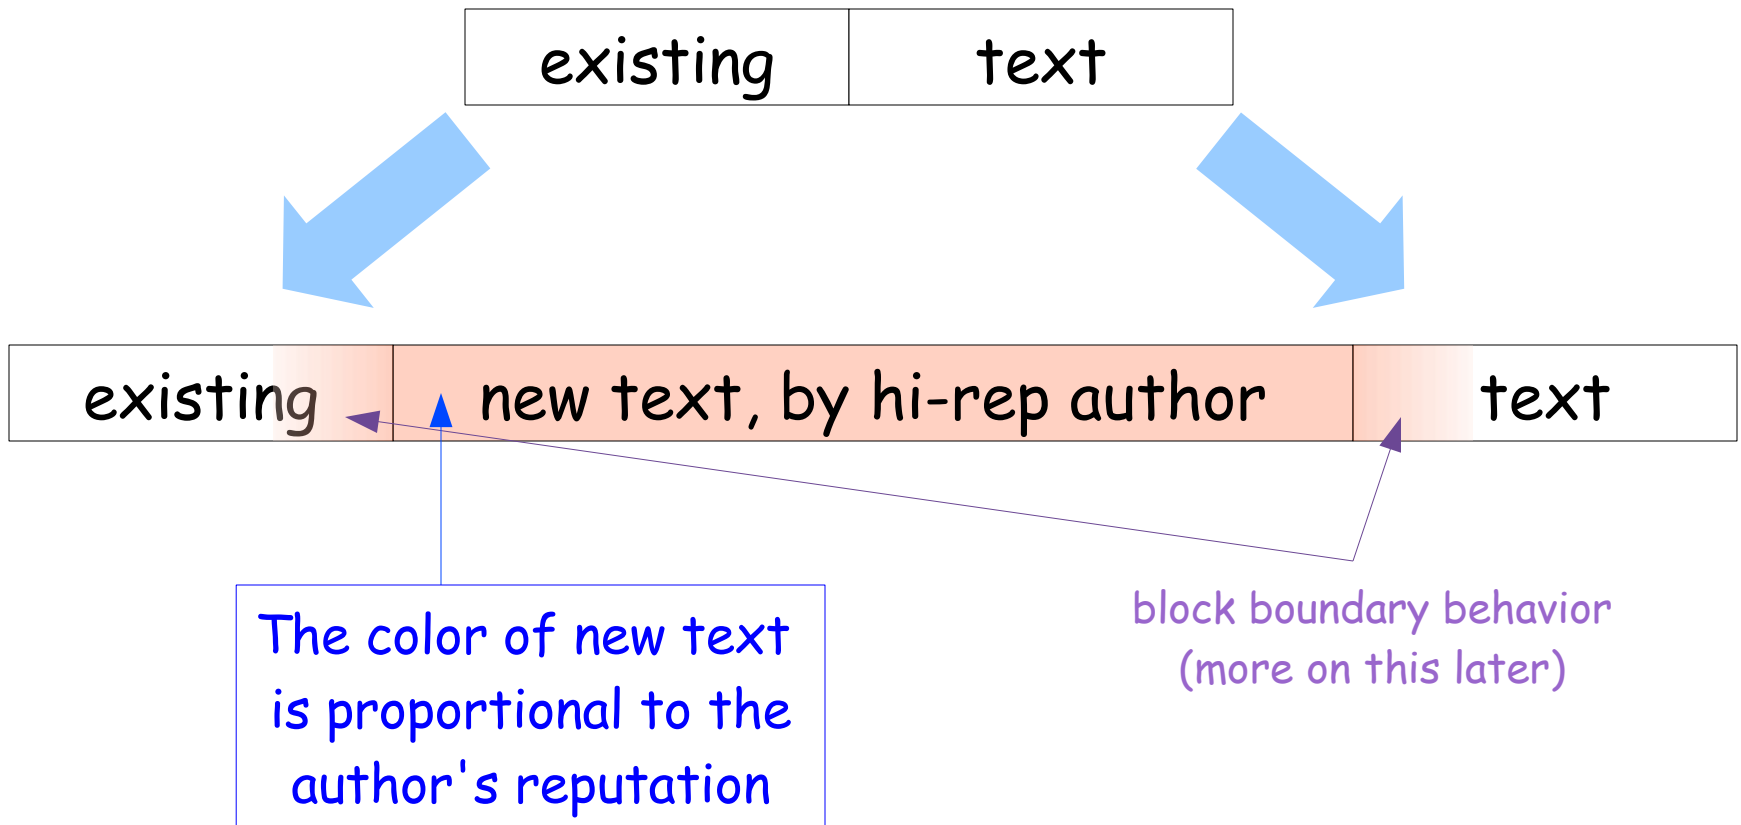
existing	text
----------	------

Trust: New Text




The color of new text is proportional to the author's reputation

Trust: New Text




Even top-reputation authors cannot single-handedly create trusted text: **trust always requires consensus.**

Trust: Rearranging Blocks of Text



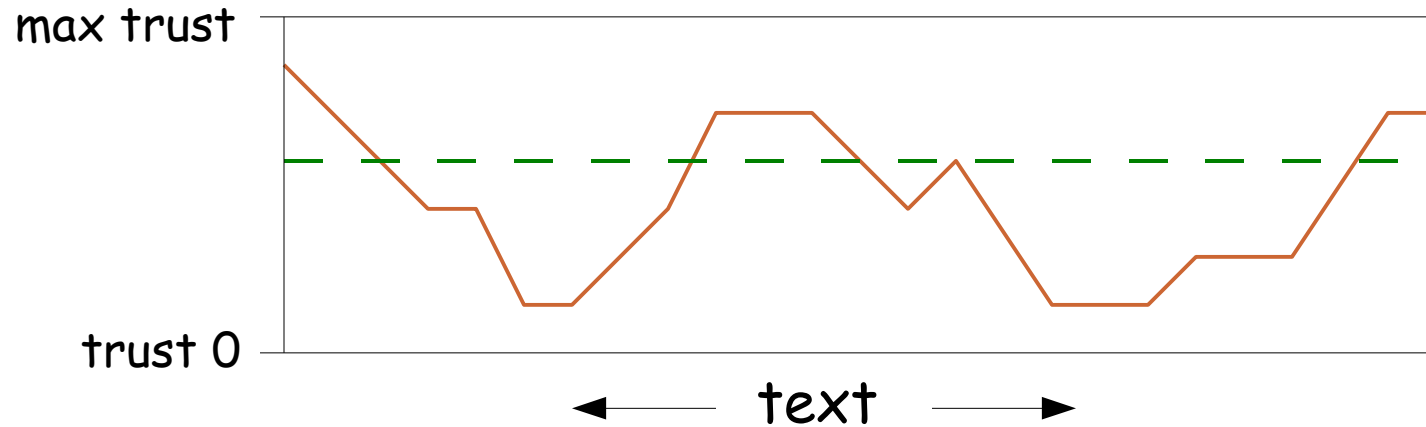
At every border between new neighbours, the text has the same trust as new text. The trust gradually returns to the original value as the distance from the disrupted border increases.

Trust: Text revision effect



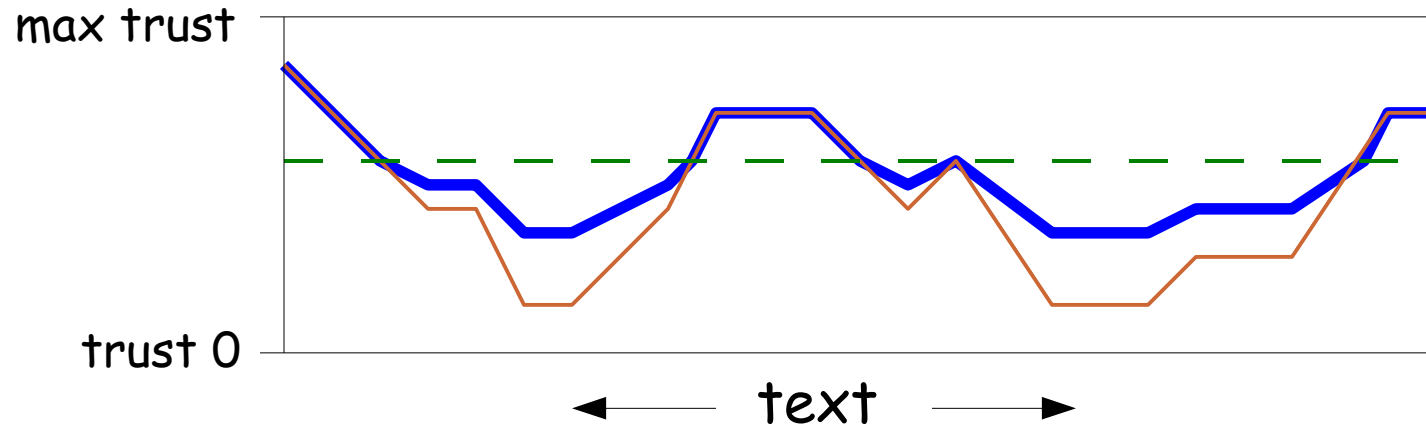
— trust of existing text

Trust: Text revision effect



- trust of existing text
- - - author reputation

Trust: Text revision effect

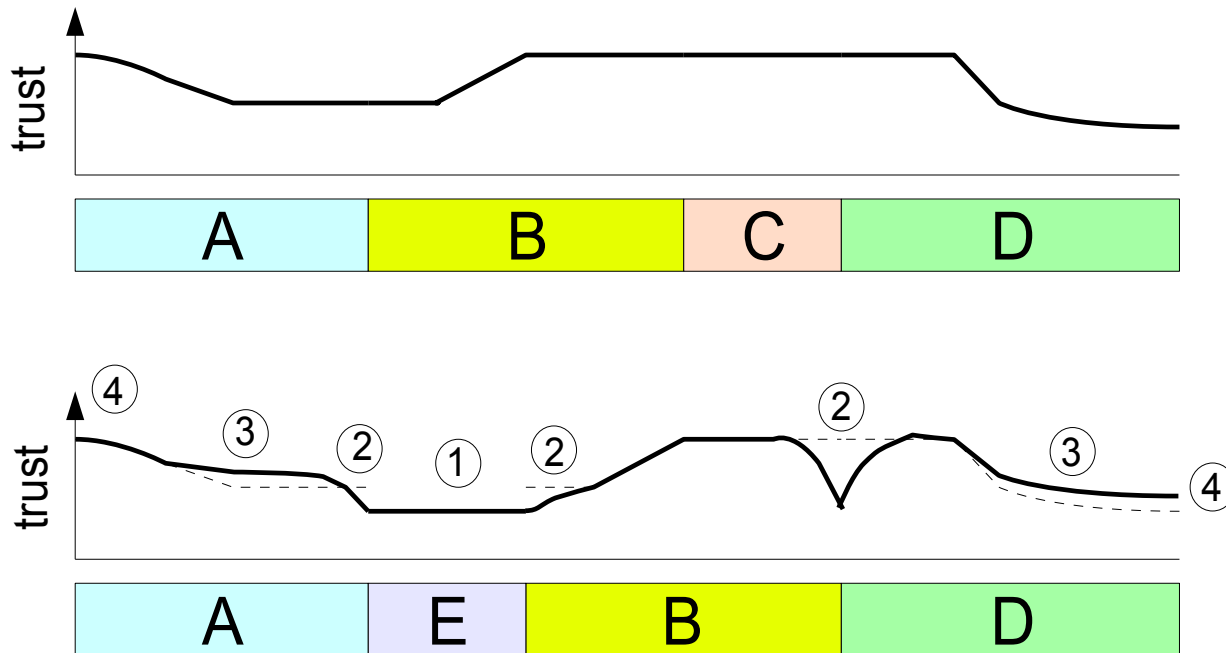


- trust of existing text
- trust of new text
- - - author reputation

If text has trust lower than the author's reputation, we update it as follows:

$$trust_{new} = trust_{old} + \alpha \cdot (reputation - trust_{old})$$

All (trust updates) together now!



① Trust of new text

② New block borders have the same trust as new text

③ The revision effect increases the trust of existing text

④ Note: this is not a new border

Italian Cuisine, revisited

Italian cuisine – The UCSC Wikipedia Trust Project

Log in / create account

[article](#) [discussion](#) [view source](#) [history](#)

Italian cuisine

Revision as of 04:20, 30 January 2007 by 69.210.149.199 (Talk)
(diff) ←Older revision | Current revision (diff) | Newer revision→ (diff)

Italian cuisine is extremely varied: the country of **Italy** was only unified in **1861**, and its cuisines reflect the cultural variety of its **regions** and its diverse history (with culinary influences from Greek, Roman, Norman and Arab civilizations). Italian cuisine is imitated **all over the world**. It also is way better than French food, the losers.

To a certain extent, there is really no such thing as

This article is part of the **Cuisine** series

Preparation techniques and cooking items

[Techniques - Utensils](#)

WIKIPEDIA
The Free Encyclopedia

navigation

- [Main Page](#)
- [Community portal](#)
- [Current events](#)
- [Recent changes](#)
- [Random page](#)
- [Help](#)
- [Donations](#)


Cheerleading and word spotting

History

Spam



The real starter of cheerleader was Marte Klopstad! She was a happy little child who loved to dance and play with dolls. Cheerleading first started at Princeton University in the 1880s with the crowd chant, "Rah rah rah, tiger lion bear, sis sis sis, boom boom boom ahhhhhhh, Princeton Princeton Princeton!!" as a

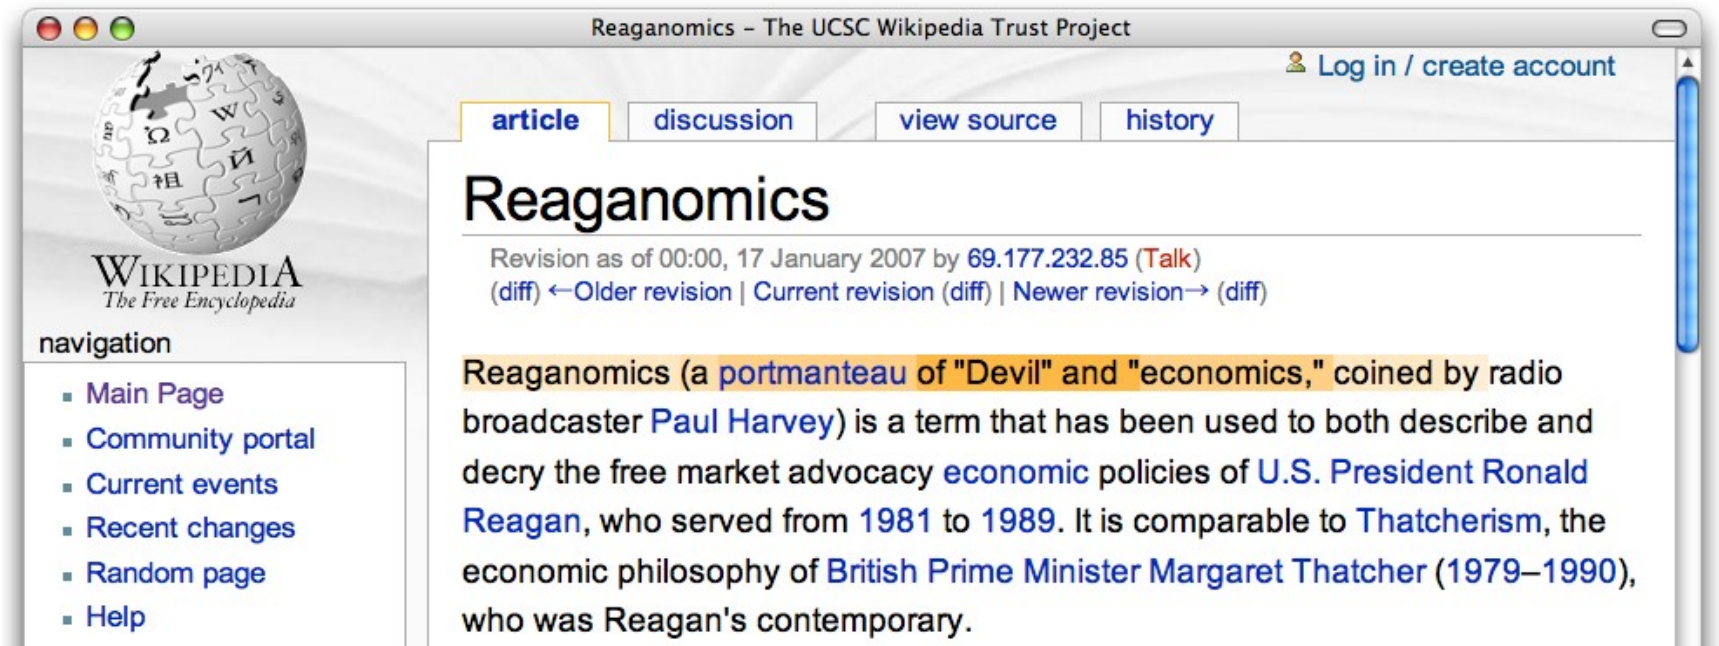


Not spam

Many tell us that we should use semantic analysis as well.
But the present method is simpler, and language independent.

Word or sentence level?

Word level gives more information.



The screenshot shows a web browser window with the title "Reaganomics - The UCSC Wikipedia Trust Project". The browser's address bar is empty. The page content includes a navigation menu with "article", "discussion", "view source", and "history" tabs. The main heading is "Reaganomics". Below the heading, it says "Revision as of 00:00, 17 January 2007 by 69.177.232.85 (Talk)" and "(diff) ←Older revision | Current revision (diff) | Newer revision→ (diff)". The main text of the article is: "Reaganomics (a portmanteau of "Devil" and "economics," coined by radio broadcaster Paul Harvey) is a term that has been used to both describe and decry the free market advocacy economic policies of U.S. President Ronald Reagan, who served from 1981 to 1989. It is comparable to Thatcherism, the economic philosophy of British Prime Minister Margaret Thatcher (1979–1990), who was Reagan's contemporary." The word "portmanteau" is highlighted in yellow. On the left side, there is a sidebar with the Wikipedia logo and a "navigation" section containing links to "Main Page", "Community portal", "Current events", "Recent changes", "Random page", and "Help".

Reaganomics - The UCSC Wikipedia Trust Project

Log in / create account

article discussion view source history

Reaganomics

Revision as of 00:00, 17 January 2007 by 69.177.232.85 (Talk)
(diff) ←Older revision | Current revision (diff) | Newer revision→ (diff)

Reaganomics (a portmanteau of "Devil" and "economics," coined by radio broadcaster Paul Harvey) is a term that has been used to both describe and decry the free market advocacy economic policies of U.S. President Ronald Reagan, who served from 1981 to 1989. It is comparable to Thatcherism, the economic philosophy of British Prime Minister Margaret Thatcher (1979–1990), who was Reagan's contemporary.

WIKIPEDIA
The Free Encyclopedia

navigation

- Main Page
- Community portal
- Current events
- Recent changes
- Random page
- Help

Finding non-obvious tampering

together with the leaders of his coalition partners selects the other Ministers which make up the Government and acts as political heads of the various government departments. Cabinet members are occasionally recruited from outside the [Folketing](#).

the United States foreign policies

Since [27 November 2001](#), the [economist Anders Fogh Rasmussen](#) has been Prime Minister to Denmark.

As known in other parliamentary systems of government, the executive, i.e. the Government, is answerable to the Folketing. Under the [Danish constitution](#), no government may exist with a majority against it, as opposed to the more common rule of government needing a majority for it. It is because of this rule, Denmark often sees minority governments.

The correct spelling is Fogh.

In Danish, "fjog" means "fool" or "goofy"

How to evaluate our trust computation?

Via a human trial?

- Long, boring, expensive
- Not suited to algorithm optimization
- Surprising difficult even for humans to judge whether information is trustworthy.

Idea: Trust should predict text stability [Zeng, Alhoussaini, Fikes, McGuinness 96]

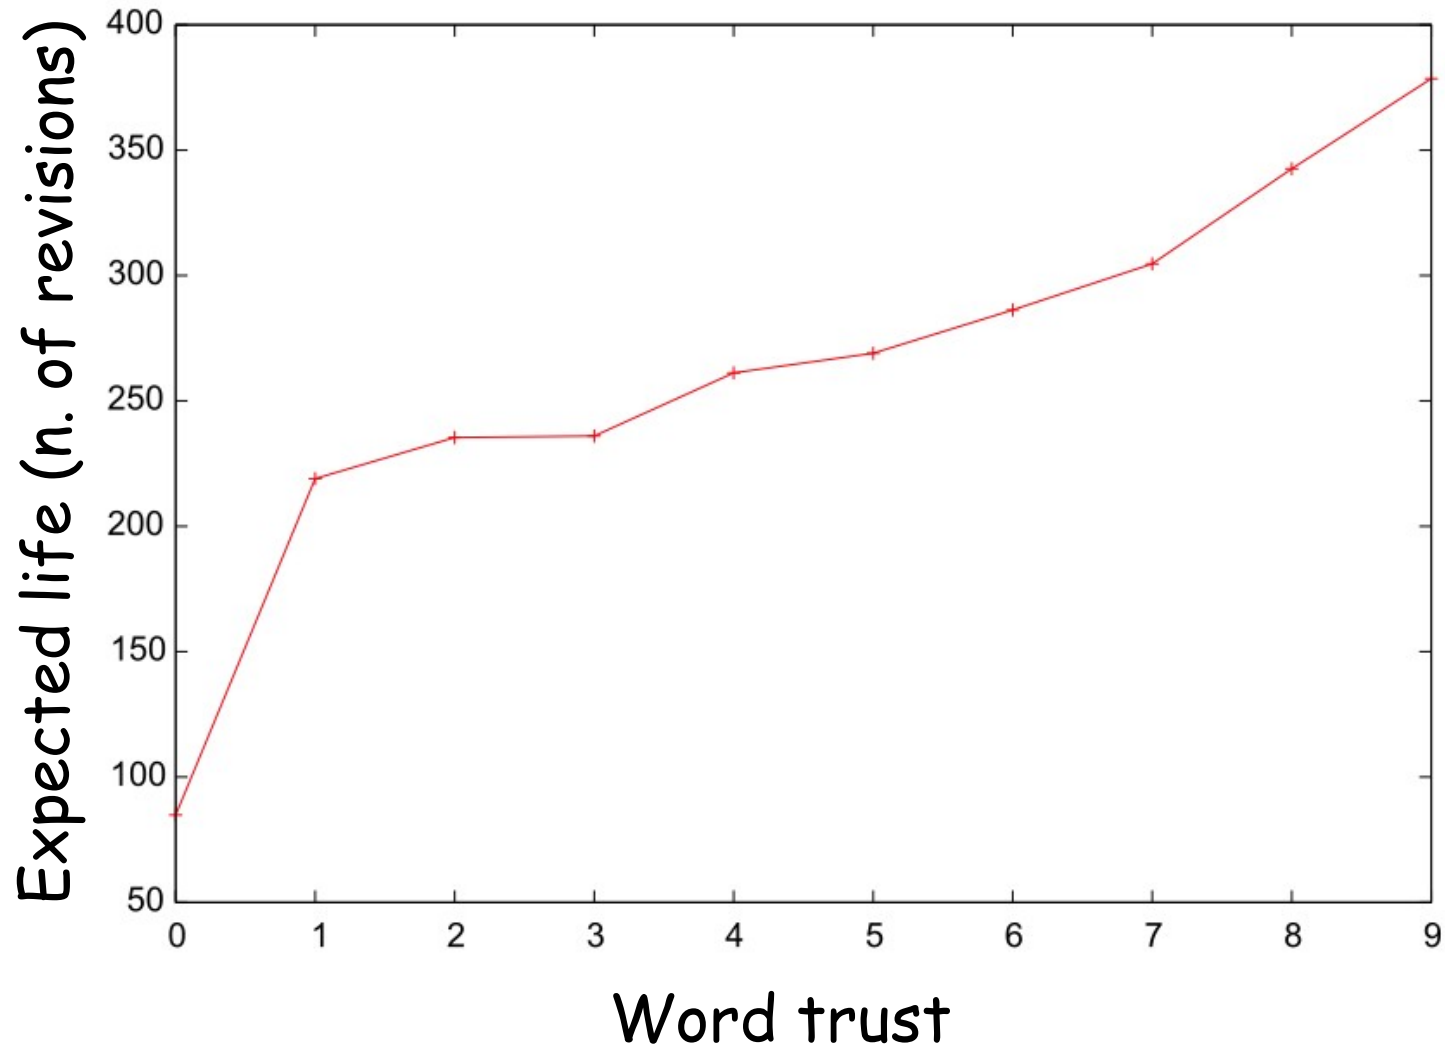
- Text that is trusted should be less likely to change.

Low trust predicts text deletion

- **Recall wrt. deletions:** Text in the bottom half of trust values constitutes 3.4% of the text, yet corresponds to 66% of the text that is deleted in the next revision.
- **Precision wrt. deletions:** Text in the bottom half of trust values has a probability of 33% of being deleted in the very next revision, compared with 1.9% for general text. The probability raises to 62% for text in the bottom fifth of trust values.

Data obtained by analyzing 1,000 articles selected at random among those with at least 200 revisions.

Trust predicts text lifespan



From new text to trusted text - animation

Executive branch

Denmark is a constitutional monarchy with an almost unbroken link of monarchs for more than 1,000 years (except for an interregnum of eight years from 1332 to 1340). The current monarch, Queen Margaret II, has largely ceremonial functions; perhaps her most significant formal power lies in her right to appoint the prime minister (Statsminister) and the cabinet of Denmark, who are responsible for administration of the government. However, she must consult with parliamentary leaders to determine the public's will, since the cabinet may be dismissed by a vote of no confidence in the Folketing, the Danish parliament. Cabinet members are occasionally recruited from outside the Folketing. The vote of no confidence is an essential part of danish parliamentary work, since the constitution states "no government can work with a majority against it", opposed to the more common international rule that says government need a majority for it. The rule can seem to have the same meaning, but in politics nothing is absolute, and as proof of this the socialdemocrats has once used this rule to form goverment although the liberal-conservative block was of greater number.

From new text to trusted text - animation

Executive branch

Denmark is a constitutional monarchy with an almost unbroken link of monarchs for more than 1,000 years (except for an interregnum of eight years from 1332 to 1340). The current monarch, Queen Margaret II, has largely ceremonial functions; perhaps her most significant formal power lies in her right to appoint the prime minister (Statsminister) and the cabinet of Denmark, who are responsible for administration of the government. However, she must consult with parliamentary leaders to determine the public's will, since the cabinet may be dismissed by a vote of no confidence in the Folketing, the Danish parliament. Cabinet members are occasionally recruited from outside the Folketing. The vote of no confidence is an essential part of Danish parliamentary work, since the constitution states "no government can work with a majority against it", opposed to the more common international rule that says a government needs a majority for it. These rules may seem to have the same meaning, but in politics nothing is absolute, and as proof of this the Social Democrats once used this rule to form government although the liberal-conservative block was of greater number.

From new text to trusted text - animation

Executive branch

Denmark is a constitutional monarchy with an almost unbroken link of monarchs for more than 1,000 years (except for an interregnum of eight years from 1332 to 1340). The current monarch, Queen Margaret II, has largely ceremonial functions; perhaps her most significant formal power lies in her right to appoint the prime minister (Statsminister) and the cabinet of Denmark, who are responsible for administration of the government. However, she must consult with parliamentary leaders to determine the public's will, since the cabinet may be dismissed by a vote of no confidence in the Folketing, the Danish parliament. Cabinet members are occasionally recruited from outside the Folketing. The vote of no confidence is an essential part of Danish parliamentary work, since the constitution states "no government can work with a majority against it", opposed to the more common international rule that says a government needs a majority for it. These rules may seem to have the same meaning, but in politics nothing is absolute, and as proof of this the Social Democrats once used this rule to form government although the liberal-conservative block was of greater number.

From new text to trusted text - animation

Executive branch

Denmark is a constitutional monarchy with an almost unbroken link of monarchs for more than 1,000 years (except for an interregnum of eight years from 1332 to 1340). The current monarch, Queen Margaret II, has largely ceremonial functions; perhaps her most significant formal power lies in her right to appoint the prime minister (Statsminister) and the cabinet of Denmark, who are responsible for administration of the government. However, she must consult with parliamentary leaders to determine the public's will, since the cabinet may be dismissed by a vote of no confidence in the Folketing, the Danish parliament. Cabinet members are occasionally recruited from outside the Folketing. The vote of no confidence is an essential part of Danish parliamentary work, since the constitution states "no government can work with a majority against it", opposed to the more common international rule that says a government needs a majority for it. These rules may seem to have the same meaning, but in politics nothing is absolute, and as proof of this the Social Democrats once used this rule to form government although the liberal-conservative block was of greater number.

From new text to trusted text - animation

Executive branch

Denmark is a constitutional monarchy with an almost unbroken link of monarchs for more than 1,000 years (except for an interregnum of eight years from 1332 to 1340). The current monarch, Queen Margaret II, has largely ceremonial functions; perhaps her most significant formal power lies in her right to appoint the prime minister (Statsminister) and the cabinet of Denmark, who are responsible for administration of the government. However, she must consult with parliamentary leaders to determine the public's will, since the cabinet may be dismissed by a vote of no confidence in the Folketing, the Danish parliament. Cabinet members are occasionally recruited from outside the Folketing. The vote of no confidence is an essential part of Danish parliamentary work, since the constitution states "no government can work with a majority against it", opposed to the more common international rule that says a government needs a majority for it. These rules may seem to have the same meaning, but in politics nothing is absolute, and as proof of this the Social Democrats once used this rule to form government although the liberal-conservative block was of greater number.

From new text to trusted text - animation

Executive branch

Denmark is a constitutional monarchy with an almost unbroken link of monarchs for more than 1,000 years (except for an interregnum of eight years from 1332 to 1340). The current monarch, Queen Margaret II, has largely ceremonial functions; perhaps her most significant formal power lies in her right to appoint the prime minister (Statsminister) and the cabinet of Denmark, who are responsible for administration of the government. However, she must consult with parliamentary leaders to determine the public's will, since the cabinet may be dismissed by a vote of no confidence in the Folketing, the Danish parliament. Cabinet members are occasionally recruited from outside the Folketing. The vote of no confidence is an essential part of Danish parliamentary work, since the constitution states "no government can work with a majority against it", opposed to the more common international rule that says a government needs a majority for it. These rules may seem to have the same meaning, but in politics nothing is absolute, and as proof of this the Social Democrats once used this rule to form government although the liberal-conservative block was of greater number.

From new text to trusted text - animation

Executive branch

Denmark is a constitutional monarchy with an almost unbroken link of monarchs for more than 1,000 years (except for an interregnum of eight years from 1332 to 1340). The current monarch, Queen Margaret II, has largely ceremonial functions; perhaps her most significant formal power lies in her right to appoint the prime minister (Statsminister) and the cabinet of Denmark, who are responsible for administration of the government. However, she must consult with parliamentary leaders to determine the public's will, since the cabinet may be dismissed by a vote of no confidence in the Folketing, the Danish parliament. Cabinet members are occasionally recruited from outside the Folketing. The vote of no confidence is an essential part of Danish parliamentary work, since the constitution states "no government can work with a majority against it", opposed to the more common international rule that says a government needs a majority for it. These rules may seem to have the same meaning, but in politics nothing is absolute, and as proof of this the Social Democrats once used this rule to form government although the liberal-conservative block was of greater number.

From new text to trusted text - animation

Executive branch

Denmark is a constitutional monarchy with an almost unbroken link of monarchs for more than 1,000 years (except for an interregnum of eight years from 1332 to 1340). The current monarch, Queen Margaret II, has largely ceremonial functions; perhaps her most significant formal power lies in her right to appoint the prime minister (Statsminister) and the cabinet of Denmark, who are responsible for administration of the government. However, she must consult with parliamentary leaders to determine the public's will, since the cabinet may be dismissed by a vote of no confidence in the Folketing, the Danish parliament. Cabinet members are occasionally recruited from outside the Folketing. The vote of no confidence is an essential part of Danish parliamentary work, since the constitution states "no government can work with a majority against it", opposed to the more common international rule that says a government needs a majority for it. These rules may seem to have the same meaning, but in politics nothing is absolute, and as proof of this the Social Democrats once used this rule to form government although the liberal-conservative block was of greater number.

Social interaction on the Wikipedia

- We can automatically detect the amount of contribution, reversions, edit wars, and other phenomena!
- This enables all sorts of sociological studies of the Wikipedia, and of how people collaborate.
- We are just starting to look into the data...

The demo, and the tool: WikiTrust

- Live demo: <http://trust.cse.ucsc.edu/>
- The tool WikiTrust is now available (BSD license):
<http://trust.cse.ucsc.edu/WikiTrust>
- We are currently developing an on-line version of the tool, that can compute reputation and trust in real-time, as edits are made.
- Goal: deployment on the live Wikipedia.
- We thank CITRIS for the support!!

The End