Can Who-Edits-What Predict Edit Survival?

Ali Batuhan Yardim[1]; Victor Kristof[2]; Lucas Maystre[2]; Matthias Grossglauser[2]
[1]Bilkent University, [2]Ecole Polytechnique Fédérale de Lausanne

Motivation

Online collaborative projects: Wikipedia, Linux, OpenStreetMap, …

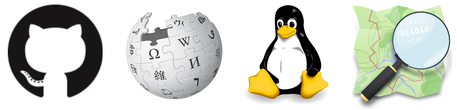

User bases of tens of thousands…

Goal:

- Predict quality of contributions
- Match contributors to suitable tasks

Existing Solutions

Specialized classifiers: Powerful yet brittle

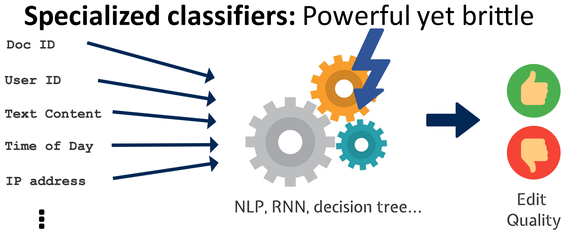

Features limited / not transferable

Tools not always available (e.g. NLP for Turkish?)

User reputation: Information lost

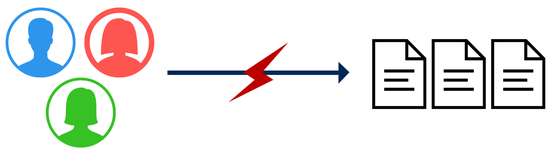

No consideration of item differentiation

Bridging the gap: INTERANK

Uses only who-edits-what data, based on discrete choice models

- Simple
- Powerful
- Easy-to-interpret
- Generalizable

$$p_{ui} = \frac{1}{1 + \exp[-(s_u - d_i + \boldsymbol{x}_u^\top \boldsymbol{y}_i + b)]}$$

User skill

Item difficulty

User & item embedding (in full variant)

Wikipedia

Performance: comparable to state-of-the-art specialized classifier.

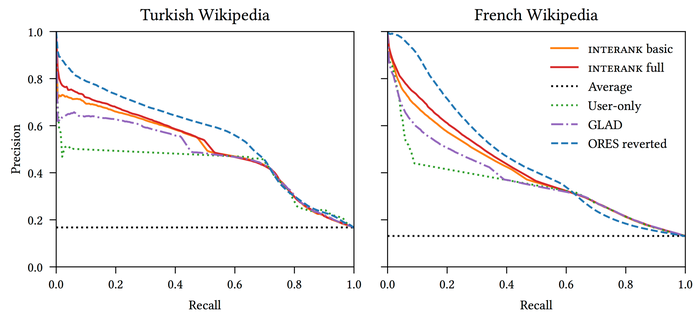

Interpretation: The most controversial articles on French Wikipedia [Yasseri et al. 2014] and the percentile difficulty for our model.

Rank	Article Title	Percentile of d_i
1	Ségolène Royal	99.840%
2	Unidentified flying object	99.229%
3	Jehovah's Witnesses	99.709%
4	Jesus	99.953%
5	Sigmund Freud	99.841%

Interpreting the latent terms: t-SNE plots of latent features from 80 French articles.

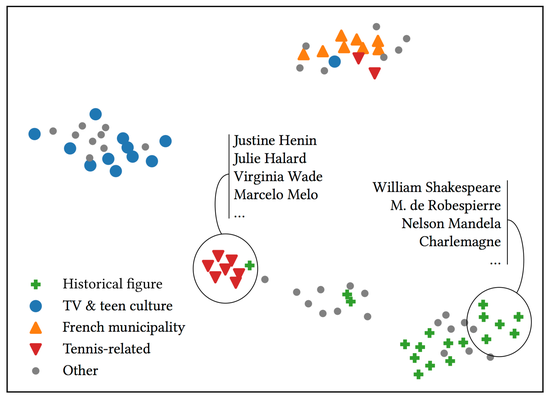

High culture vs. popular culture: Highest & lowest articles along the first PCA axis of latent vectors on the Turkish Wikipedia.

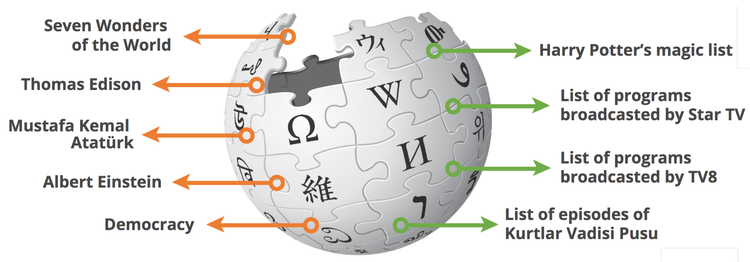

Linux

Performance: Better performance than state-of-the-art specialized classifiers.

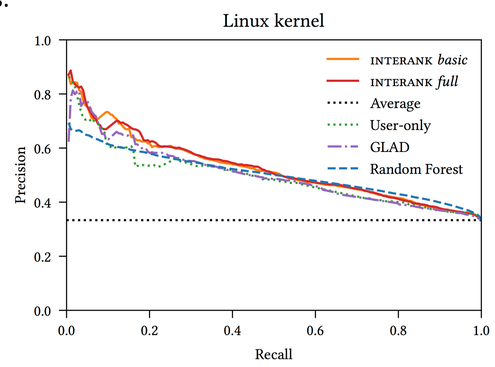

Interpreting the difficulty: The five most easy and difficult subsystems to contribute to, with their acceptance rate and number of developers.

Difficulty	Subsystem	%Acc.	#Dev.
+2.664	usr	1.88%	70
+1.327	include	7.79%	101
+1.038	lib	15.99%	707
+1.013	drivers/clk	34.34%	81
+0.865	include/trace	17.73%	81
-1.194	drivers/addi-data	78.31%	8
-1.080	net/tipc	43.11%	44
-0.993	drivers/ps3	44.26%	9
-0.936	net/nfc	73.04%	26
-0.796	arch/mn10300	45.40%	63

INTERANK at a glance

- **Goal:** predict edit quality in online collaborative projects
- Simple model, yet matches state-of-the-art predictive performance
- Evaluated on Wikipedia and Linux datasets
- **Key feature - Interpretability:** Easily discover insights into the project from model parameters

Code available at github.com/lca4/interank