Parameter Tuning of Reranking-based Diversification Algorithms using Total Curvature Analysis

Shameem A Puthiya Parambath Siwei Liu Christos Anagnostopoulos Roderick Murray-Smith Iadh Ounis

School of Computing Science
Essence: Data Science & Distributed Computing Lab

Introduction

Goals

- ✔ Analyse popular diversification objective functions to find common pattern
- ✔ Find a simple way to estimate/validate the hyperparameters for balanced relevance-diversity trade-off

Contributions

- ✔ Popular diversification objectives can be unified under the scheme of maximizing submodular or modular functions from the class of parameterized concave or linear over modular functions.
- ✔ The total curvature of the objective functions provides insights about the relevance-diversity trade off
- ✔ The total curvature value serves as a 'vehicle of validation' to seek hyperparameters that balances the relevance-diversity.

Background

Definitions

Definition 1 *A function F defined on the subsets of a ground set $\mathcal{Z}$ is called submodular, if for all subsets $\mathcal{A}, \mathcal{B} \subseteq \mathcal{Z}$,*

$$F(\mathcal{A}) + F(\mathcal{B}) \geq F(\mathcal{A} \cup \mathcal{B}) + F(\mathcal{A} \cap \mathcal{B}).$$

F is modular if strict equality holds, while F is monotone if for every $\mathcal{A} \subseteq \mathcal{B}$, $F(\mathcal{A}) \leq F(\mathcal{B})$.

Definition 2 *A real-valued function g on a convex set $\mathcal{C}$ is said to be concave if, for any x and $y \in \mathcal{C}$ and for any $a \in [0,1]$*

$$g((1-a)x + ay) \geq (1-a)g(x) + ag(y)$$

Diversification Objective Analysis

Generic Form

- Re-ranking based diversification objective functions can be expressed in the generic form:

$$F(\mathcal{S}) = f(\mathcal{S}) + \beta g(h(\mathcal{S})),$$

- ✔ $f(\cdot)$ represents relevance of $\mathcal{S}$. It is a **modular** function.
- ✔ $g(h(\mathcal{S}))$ represents the diversity of $\mathcal{S}$. It is the composition of a **linear or concave** function, $g(\cdot)$, and a **modular** function h.
- ✔ $h(\cdot)$ is defined in terms of different diversification concepts such as item coverage, popularity bias, item novelty and long-tail recommendation.
- ✔ The β is the hyperparameter to be tuned for the relevance-diversity trade-off.

Functional Form of $g(x)$

Algorithm/Method	$g(x)$
Onuma et al. [5], Su et al. [7]	$-\frac{1}{x}$, $x > 0$,
Oh et al. [4]	$\log(x)$, $x > 0$
Puthiya Parambath et al. [6], Vargas et al. [8]	x^{β}, $x \geq 0, \beta \in (0,1)$
Wu et al. [10]	$\frac{x}{1+x}$, $x \geq 0$
Wasilewski and Hurley [9]	βx, $\beta > 0$

Monotonicity of the Objective

Algorithm/Method	Monotonicity
Carbonell and Goldstein [2], Su et al. [7]	No
Abdollahpouri et al. [1], Oh et al. [4], Onuma et al. [5]	Yes
Puthiya Parambath et al. [6], Vargas et al. [8]	Yes
Wasilewski and Hurley [9], Wu et al. [10]	Yes

Total Curvature

Definition 3 *The total curvature α of a non-decreasing submodular set function with respect to a set $\mathcal{S}$ is defined as:*

$$\alpha = \max_{j \in \mathcal{S}} \frac{F(\mathcal{S} \setminus \{j\}) + F(\{j\}) - F(\mathcal{S})}{F(\{j\})}.$$

- ✔ *total curvature* measures how far $F(\cdot)$ is from being modular i.e. it represents the distance of a monotone submodular function to the modularity.
- ✔ *total curvature* takes value between 0 and 1.
- ✔ *total curvature* is zero for modular functions and one for matroid rank function.
- ✔ *total curvature* relys only on the marginal gains obtained by adding an item and not on specific functional form of $F(\cdot)$.

Optimality bound in terms of *total curvature*

The bounds that can be obtained using the greedy algorithm for submodular maximization can be tightened using *total curvature*. According to Conforti and Cornuéjols [3],

$$F(\mathcal{S}^*) \geq \frac{(1 - e^{-\alpha})F(\mathcal{S}^{opt})}{\alpha}$$

α	Optimality Gap
0.1	$F(S^*) \geq 0.95F(S^{opt})$
0.5	$F(S^*) \geq 0.79F(S^{opt})$
0.9	$F(S^*) \geq 0.66F(S^{opt})$

- ✔ By changing the value of the hyperparameter β in the re-ranking objective function, one can effectively change the *total curvature* of the objective.

Balanced Hyperparameter Tuning

- ✘ In practice, the balanced value of β is found by the grid search.
- ✔ The balanced β value depends on the objective function F, whereas α is independent of F.
- ✔ An α value closer to 0.5 gives the balanced trade-off between being completely relevant and diverse.
- ✔ Practitioner can choose β such that the corresponding α is closer to 0.5.

Experiments

We used the benchmark MovieLens 20M dataset. We tested with two algorithms *(i)* the coverage maximization algorithm in [6] and *Binary xQuAD* algorithm in [1].

Results

Relevance-Diversity trade-off for different values of α and β. Top and bottom rows represent the relevance-diversity trade-off for the algorithms in [6] and [1], respectively.

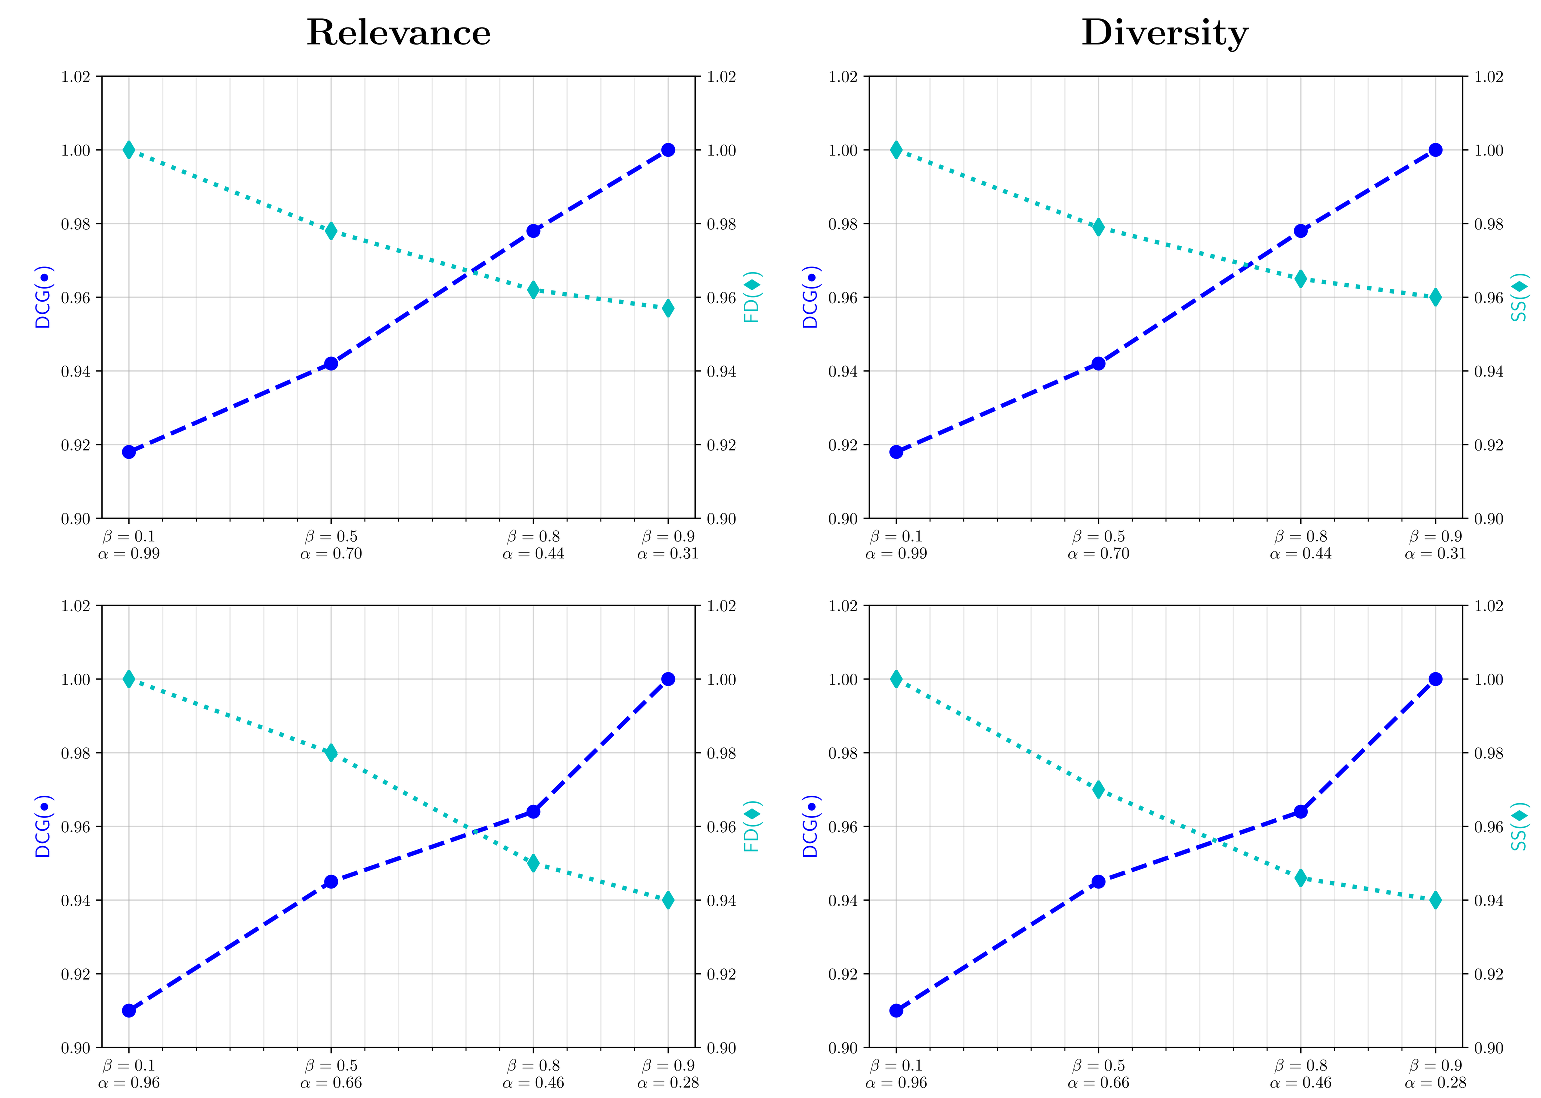

References

[1] Himan Abdollahpouri, Robin Burke, and Bamshad Mobasher. Managing popularity bias in recommender systems with personalized re-ranking. In *AAAI Conference on Artificial Intelligence*, 2019.

[2] Jaime Carbonell and Jade Goldstein. The use of mmr, diversity-based reranking for reordering documents and producing summaries. In *SIGIR Conference on Research and development in information retrieval*, pages 335–336, 1998.

[3] Michele Conforti and Gérard Cornuéjols. Submodular set functions, matroids and the greedy algorithm: tight worst-case bounds and some generalizations of the rado-edmonds theorem. *Discrete applied mathematics*, 7 (3):251–274, 1984.

[4] Jinoh Oh, Sun Park, Hwanjo Yu, Min Song, and Seung-T Park. Novel recommendation based on personal popularity tendency. In *International Conference on Data Mining*, pages 507–516, 2011.

[5] Kensuke Onuma, Hanghang Tong, and Christos Faloutsos. Tangent: a novel,'surprise me', recommendation algorithm. In *SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 657–666, 2009.

[6] Shameem A Puthiya Parambath, Nicolas Usunier, and Yves Grandvalet. A coverage-based approach to recommendation diversity on similarity graph. In *ACM Conference on Recommender Systems*, pages 15–22, 2016.

[7] Ruilong Su, Li'Ang Yin, Kailong Chen, and Yong Yu. Set-oriented personalized ranking for diversified top-n recommendation. In *ACM conference on Recommender systems*, pages 415–418, 2013.

[8] Saúl Vargas, Linas Baltrunas, Alexandros Karatzoglou, and Pablo Castells. Coverage, redundancy and size-awareness in genre diversity for recommender systems. In *ACM conference on Recommender systems*, pages 209–216, 2014.

[9] Jacek Wasilewski and Neil Hurley. Intent-aware item-based collaborative filtering for personalised diversification. In *Conference on User Modeling, Adaptation and Personalization*, pages 81–89, 2018.

[10] Le Wu, Qi Liu, Enhong Chen, Nicholas Jing Yuan, Guangming Guo, and Xing Xie. Relevance meets coverage: a unified framework to generate diversified recommendations. *ACM Trans. Intell. Syst. Technol.*, 7(3), 2016.