



Independent Influence of Exploration and Exploitation for Metaheuristic-based Recommendations

Alexandre Bettinger, Armelle Brun, Anne Boyer

► To cite this version:

Alexandre Bettinger, Armelle Brun, Anne Boyer. Independent Influence of Exploration and Exploitation for Metaheuristic-based Recommendations. The Genetic and Evolutionary Computation Conference (GECCO), Jul 2022, Boston, United States. 10.1145/3520304.3528972 . hal-03655953

HAL Id: hal-03655953

<https://inria.hal.science/hal-03655953>

Submitted on 30 Apr 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Independent Influence of Exploration and Exploitation for Metaheuristic-based Recommendations

Alexandre Bettinger, Armelle Brun, Anne Boyer
UL CNRS LORIA
Nancy, Lorraine, FRANCE

ABSTRACT

Exploration and exploitation (E&E) of a search space are two fundamental processes in many fields of artificial intelligence. Indeed, when the search space is vast, it is important to ensure that many of its regions are examined, so as not to get trapped in a local optimum, but also that promising regions are examined more in depth in order to find good local optima. Influencing both processes is thus necessary. The literature has rarely proposed to approach the recommendation task as a vast search space problem. This paper introduces a metaheuristic-based recommendation approach with two contributions: an E&E influence process which is independent from the influenced algorithm and new indicators to represent and explain E&E. Performed on a genetic algorithm (GA) and on a reinforcement algorithm (RA), our experiments confirm that (1) the proposed influence process has a positive impact on evaluation criteria and on E&E which brings better recommendations, (2) proposed indicators contribute to represent E&E from new angles.

CCS CONCEPTS

• **Computing methodologies** → **Heuristic function construction**.

KEYWORDS

Recommendation, metaheuristic, evolutionary algorithm, exploration, exploitation, influence

ACM Reference Format:

Alexandre Bettinger, Armelle Brun, Anne Boyer. 2022. Independent Influence of Exploration and Exploitation for Metaheuristic-based Recommendations. In *Genetic and Evolutionary Computation Conference Companion (GECCO '22 Companion)*, July 9–13, 2022, Boston, MA, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3520304.3528972>

1 INTRODUCTION

Recommending is the act of filtering information in order to target elements likely to interest one or more users. This task can be viewed as the processing of a vast search space which represents all possible recommendations. Depending on the recommendation context, a recommendation can take various forms such as single items, itemsets or item sequences. In what follows, we are interested in the itemset recommendation framework.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GECCO '22 Companion, July 9–13, 2022, Boston, MA, USA

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9268-6/22/07.

<https://doi.org/10.1145/3520304.3528972>

The search space can be crawled through an E&E process to find good recommendations [2]. Exploring means making sure that many regions of the search space are examined [3], so as not to get trapped in a local optimum. Exploiting means that if a promising region of the search space is found, it is examined more carefully [3] to see if it contains good local optima.

The literature highlights that the area of metaheuristics (MH) addresses the issue of proposing recommendations in huge search spaces [3]. Our work aims to contribute to E&E influence and explainability in the MH recommendation frame.

Studies have been conducted to influence the dynamics of E&E in order to ensure better search behaviors and thus find better recommendations regarding predefined evaluation criteria [3]. Influencing E&E implies to know when, why and how. Indeed, E&E have to be controlled regularly through a measure, and operations done to perform the influence must be justified. These elements contribute to explaining E&E behaviors.

The problem we aim to tackle in this work is: How to control, influence and explain E&E in the MH recommendation frame?

Our contribution is threefold. First, we propose an influence process independent of the influenced algorithm in order to promote the full expression of the influence and in order to not undergo structural and/or parametric modifications inside the influenced algorithm. Second, we propose a systematic influence of E&E, at each iteration of the influenced algorithm and as long as the itemsets do not comply with an E&E indicator. Finally, we introduce indicators to measure and represent E&E from new angles.

2 RELATED WORKS

In vast search spaces, filtering information is tedious and performing an exhaustive search is unconceivable. Targeting elements to recommend requires to partially crawl the search space using E&E algorithmic mechanisms [2].

We propose a synthesis of three reviews [3][13][11] which address issues related to the E&E influence.

2.1 MHs for the Recommendation Task

Heuristic (H) methods refer to techniques for problem solving, learning and discovery where exhaustive research is not possible. MH methods are Hs integrated into iterative processes in order to increase their E&E capacities.

Our contribution falls within the domain of MHs because they are transparent and explainable, they perform well in hazardous environments and they are excellent for optimizing decisions.

MHs can be used to address the issue of proposing recommendations in huge search spaces. We found two interesting global surveys regarding MHs [4] [5] and some specific contributions referring the use of MHs for recommendation.

Algorithm 1 Traditional E&E Influence

-
- 1: Step 1: Initial Population
 - 2: Step 2: Itemsets Evaluations
 - 3: Step 3: Itemsets Variations (Selections, Crossovers, Mutations)
 - 4: Step 4: E&E Indicator Evaluation (ex : using diversity)
 - 5: Step 5: Crossover/Mutation Probability Adjustment. Depends on indicator evaluation. Aim to Force Exploration or Exploitation
 - 6: Step 6: Iteration End
 - 7: Step 7: Back To Step 2
-

Algorithm 2 EA Calling The Influence Function

-
- 1: FUNCTION ea() # Traditional EA
 - 2: initialisation()
 - 3: WHILE(condition not reached)
 - 4: ea_variations()
 - 5: influence()
 - 6: iteration_end()
 - 7: finalisation()
-

2.2 E&E Influence and MHs

E&E can be characterized as global and local search [1]. MHs often find themselves trapped in a local optimum, the main reason is the difficulty in properly arbitrating between E&E.

Controlling E&E implies to know when, what and how to control? The measurement of diversity, genotypic or phenotypic, has become widely accepted as control approaches [12].

E&E are often misunderstood by practitioners and researchers because many factors impact E&E.

In the literature, the E&E influence is referred to as compromise or guidance or balance. Influence is crucial in optimization because it improves search efficiency. The literature has proposed many approaches to influence E&E such as adaptive [14] and hybrid [9]. Most of the approaches are based on diversity through maintenance, control or learning. Moreover, E&E phases are always nested, which makes their identification more complex.

There are many open issues such as formally define E&E phases and equilibrium equation, propose (direct) metrics to measure E&E, define when and how to control the balance between E&E.

2.3 E&E Influence and Psycho-sociology

The E&E trade-off [11] depends on the conceptualization of E&E, on environmental, social and individual factors, on the scale at which E&E are considered, on the relationship and types of transitions between them. Indeed, E&E are best conceptualized as points on a continuum and are ubiquitous at many levels of abstraction.

Two main elements are highlighted by [11] for a unification of research in E&E: an E&E continuum and different types of transitions between E&E. Many openings follow: (1) propose a global theory to cover the conceptual dimensions of E&E, (2) capture the interactions between the different levels of abstraction on which E&E is located, (3) propose a hierarchical use of compromise mechanisms at different scales and levels of abstraction, (4) propose a theory on E&E transitions, and (5) explore the extent to which compromise mechanisms are sensitive to various factors.

2.4 Conclusion

Many works have been interested in the trade-off between E&E, without leading to a complete or general answer. The recommendation frame is highly concerned by these problems. Contributions aiming to explain, measure, control and influence E&E are expected.

3 CONTRIBUTIONS

We propose a process to control and influence E&E of a search algorithm (called influenced algorithm). From our view, the influence processes proposed in the literature do not guarantee a highly impactful expression of the influence. One contribution of this work thus aims to give more power to the influence process by designing an independent process from the influenced algorithm. This independence makes that this influence process can be used on most of MHs, without major adaptation.

We choose to present an evolutionary algorithms (EA) among all available MHs because GA are popular. EA generally exploits arrays of digits, traditionally representing recommendable itemsets [10]. Each index in an array corresponds to a recommendable item. Digits can take the value 0 (non-recommendation of the item) or 1 (recommendation of the item). The algorithm iteratively performs variations: generates, crosses and transforms arrays, while trying to optimize and keep the best of them with regard to predefined constraints and evaluation criteria.

In addition, we propose new indicators with the objective to have an understandable representation of E&E behavior.

3.1 E&E Influence Process

In the literature, the influence process has characteristics that limit its impact on E&E. It (1) strongly depends on the influenced algorithm, like in the parameter optimization case (Algorithm 1), (2) does not influence enough, for example when it does not have its own variation operators, (3) does not guarantee any form of compliance with the E&E indicator at each iteration of the influenced algorithm. In this work, we propose an influence process that (1) is dissociated from the influenced algorithm, (2) has its own variation operators, (3) seeks to guarantee the compliance of the solutions with the E&E indicator at each iteration of the influenced algorithm. Thus, only the itemsets are possibly modified and returned to the influenced algorithm, the latter showing no structural and/or parametric modifications. It may bring different decisions regarding E&E influence, which deserves to be studied.

The indicator is the measure of E&E used to decide whether to force exploitation, exploration or both with additional variations. An example of E&E indicator represents the diversity of the solutions found so far by the influenced algorithm. The approach we propose is designed to use any indicator.

The operating principle of the influence process is based on an independent function called during the run of the influenced algorithm (Algorithm 2). In the literature the influence is operated only once per iteration but, in our view, right after an influence decision, the elements used for taking this decision could make it possible to decide to influence again. We thus propose an E&E control, carried out after each influence variation process which is operated as long as the itemsets are not compliant with the limits

of the E&E indicator. The influence process thus guarantees the itemsets compliance with the E&E indicator within each iteration. We consider that doing so allows to dose out the influence impact, the possible issue being to perform a too impacting influence that disfigures what the influenced algorithm does regarding E&E. We must also note that the influence process will inevitably increase the computational cost as it brings more treatments.

Concretely, all solutions of the influenced algorithm are passed to the influence function (Line 5), which will transform them iteratively according to the evolution of the E&E indicator. When the indicator position itself in a specific range of values, the influence function stops the transformations and returns all the solutions to the influenced algorithm, which continues its execution until the next call of the influence function. Note that the influence variations are made as long as the indicator is outside the range. It is therefore necessary to replace the values of the limits in line with the evolution of the E&E indicator and thus to converge towards stabilization around the observed values of the E&E indicator. To do this, after each iteration of the influence function, the minimum and maximum limits of the E&E indicator are adjusted using a parameter added or removed from the average value of the E&E indicator observed so far.

3.2 Indicators to Influence and Explain E&E

As suggested in the literature, E&E would deserve to be represented from new angles [3][13][11]. In response, we propose three types of indicators that measure and represent E&E behaviors and that can be used within the influence function previously introduced.

3.2.1 Constraints Completion (CC). The CC indicator assesses the constraints completion within itemsets by treating constraints as volumes to be filled. To evaluate this indicator, for each itemset the difference between its constraint score (cs) and constraint maximum limit (cml) is evaluated, as well as the average of all differences obtained (Equation 1). The larger this gap, the more space to fill among all itemsets as it is potentially possible to add more items in some of them, while meeting constraints. Indeed, the E&E process did not allow for inserting these items so far.

$$CC = \overline{\{cml - cs\}} \quad (1)$$

3.2.2 E&E Power. The EEP indicator represents the E&E power as the amount of exploration towards exploitation. We consider that each time a digit is passed to 1 for the first time, this constitutes an exploration. An iteration discovering at least one new digit is an exploration iteration and a vector containing at least one new digit is an exploration vector.

We evaluate the number of exploration iterations divided by the number of exploitation iterations (%eit), the number of exploration vectors divided by the number of exploitation vectors (%evec) and the number of digits newly discovered divided by the number of digits already discovered (%dnd). EEP is the average of these three elements (see Equation 2).

$$EEP = \overline{(\%eit, \%evec, \%dnd)} \quad (2)$$

Eventually, a register of undiscovered digits can be used to formally define an E&E phase. Let p be an E&E phase that can take the

values {exploration, exploitation}. Let ndd be the register of digits not yet discovered. This register is updated through iterations by removing digits from it once they are used (passed to value 1) for the first time inside vectors. Let ie be an end of iteration marker of the influenced algorithm. ndd(ie) therefore corresponds to ndd at the end of iteration marker ie. So we can define p as in Equation 3.

$$p(ie + 1) = \begin{cases} \text{exploration,} & \text{if } \overline{ndd(ie)} > \overline{ndd(ie + 1)} \\ \text{exploitation,} & \text{if } \overline{ndd(ie)} = \overline{ndd(ie + 1)} \end{cases} \quad (3)$$

3.2.3 Temporality. We propose to represent the alternation of E&E iterations as gaps between time values. The temporal differences (gaps) between starting time of two iterations are evaluated and bring out the minimum (ming) and maximum (maxg) gaps, the difference between the minimum and the maximum gap (DMMG) (Equation 4), the average gap (AG) (Equation 5). DMMG allows to evaluate most extreme time gaps between exploration iterations. AG implicitly represents the quantity of exploration iterations which implicitly represents their regularity. Depending on the context, a high or low regularity is expected.

$$DMMG = maxg - ming \quad (4)$$

$$AG = \overline{gaps} \quad (5)$$

4 EXPERIMENTS

4.1 Dataset and Evaluation Measures

Our experiments are carried out on the “Movielens 25M” [6] dataset.

Influence functions are applied on a traditional GA that serves as a baseline for comparisons. The objective is to show how it performs when influenced using different E&E indicators.

4.2 Evaluation Criteria

We decide to evaluate our influence process by considering several criteria related to performance, quality and time, that are either those introduced section 3.2 or proposed in the literature [8][7].

4.2.1 Performance Criteria. An item score is defined as the weighted sum of its normalized attribute values and an itemset score is defined as the sum of the scores of its items. *Convergence Iterations* (CI): Average iterations required for reaching convergence points. *Max Score Iterations* (MSI): Average iterations required to reach maximum scores. *Convergence Score* (CS): Average maximum score reached on convergence points. *Max Score* (MS): Average maximum score reached. *Execution Time* (ET): Average execution time.

4.2.2 Quality Criteria. *All Vectors Coverage* (AVC): Coverage considering all generated vectors. It represents the percentage of digits that are at least once at 1 in the vectors considered. The higher the coverage, the better because it means that the search process has found many items. *Constraints Completion* (CC): Average constraints completion percentage. It represents how well recommended items fill the constraints. *E&E Power* (EEP): Part of exploration towards exploitation. See section 3 for more details.

4.2.3 Temporal Criteria. Min-Max Gap (DMMG): Difference between the minimum and the maximum temporal gap between two exploration iterations. **Mean Gap (AG):** Average temporal gap between two exploration iterations.

4.3 Experiments

Aggregated percentages are shown in Table 1. The % ↗, % → and % ↘ lines refer respectively to the percentages of increased, unchanged and decreased evaluation criteria results towards baseline algorithm, considering all execution cases. To complement this information, we perform ranksum tests represented by the %Pass line showing the percentage of passed ranksum tests considering all execution cases. Let us precise that the null hypothesis of a ranksum test is that evaluation criteria values, of the non influenced algorithm and of one of its influenced version, come from the same distribution. If the null hypothesis is rejected, it is a Pass (p-value < 0.05) and otherwise it is a Fail (p-value ≥ 0.05).

Table 1: Results - E&E Influence Impact On GA

	Performance					Quality			Temporal	
	CI	MSI	CS	MS	ET	AVC	CC	EEP	DMMG	AG
% ↗	62	65	40	63	100	50	44	2	82	77
% →	2	0	56	36	0	50	48	13	0	17
% ↘	37	35	4	1	0	0	8	85	18	5
%Pass	17	20	56	68	100	57	52	89	78	96

Significant ratio values for most of the criteria show a positive impact of the influence process. Regarding performance criteria, the influence process makes the algorithm reach better scores at the expense of the number of iterations and execution time. Regarding quality criteria, the influence process makes the baseline algorithm evaluate more solutions. As users preferences are encoded in constraints and as CC remains in almost all cases at least equal to the baseline, we can conclude that the proposed influence process allows to find better recommendations. Besides, the influence process has a significant impact on EEP by favouring exploitation and a noticeable impact on E&E regularity.

To summarize, the influence process has a positive impact on the search behavior. We note that a detailed results analysis shows that some proposed influence cases perform better than the others. We also note that the proposed influence does not have a positive impact in all contexts, this was expected as it depends on the search space structure and on set of constraints. We also stress that results confirm that the influence process does not induce random behaviors as scores remain at least equals to the baseline.

Moreover, we confirm that proposed influence indicators provide explanations regarding E&E. Indeed, they represent E&E behaviors from new angles as shown with CC, EEP, DMMG and AG.

To show the adaptability of the influence process, we apply it to a RA. Results show a positive impact of the influence. Similar conclusions to the GA case can be drawn.

5 DISCUSSION AND FUTURE WORK

We highlight here the genericity of the proposed influence process.

Firstly, the independent influence proposed is adaptable to genetic programming, to MHs and more generally to any algorithm based on the E&E of the search space.

Besides, the moment where the call of the influence function occurs impacts E&E and depends on the influenced algorithm as well as on the way the influence function is designed. We have chosen to position it before the end of each iteration of the influenced algorithm, this choice can be debated.

Moreover, influence variations can be performed in many ways, depending on the implementation context.

To continue, the proposed influence process could include too much randomness and thus alter the influenced algorithm enrichment. Our results show that if used properly, this is not an issue.

Finally, our contributions could be analyzed against other evaluation criteria.

Our future work will be marked by developments around new E&E influence processes and indicators. We will also perform experiments using other MHs to support that conclusions do not only stand for genetic programming and reinforcement. In addition, we intend to carry out more in-depth and diverse experiments in the recommendation frame to thorough study the characteristics of the recommendations provided. Finally, we will keep up with our work regarding the explainability of E&E behaviors.

ACKNOWLEDGMENTS

This work is carried out within the framework of the PEACE agreement with the French Rectorate of the Nancy-Metz Academy as well as with the French Ministry of National Education and Youth.

REFERENCES

- [1] Anne Auger, Marc Schoenauer, and Olivier Teytaud. 2005. Local and global order 3/2 convergence of a surrogate evolutionary algorithm. In *Proceedings of the 7th annual conference on Genetic and evolutionary computation*. 857–864.
- [2] Minmin Chen. 2021. Exploration in Recommender Systems. In *Fifteenth ACM Conference on Recommender Systems*. ACM, Amsterdam Netherlands, 551–553.
- [3] Matej Crepinsek, Shih-hsi Liu, and Marjan Mernik. 2013. Exploration and Exploitation in Evolutionary Algorithms: A Survey. *Comput. Surveys* 45, 35 (06 2013), 1–33. Issue 3. <https://doi.org/10.1145/2480741.2480752>
- [4] Sachin Desale, Akhtar Rasool, Sushil Andhale, and Priti Rane. 2015. Heuristic and meta-heuristic algorithms and their relevance to the real world: a survey. *Int. J. Comput. Eng. Res. Trends* 351, 5 (2015), 2349–7084.
- [5] Tansel Dokeroğlu, Ender Sevinc, Tayfun Kucukylmaz, and Ahmet Cosar. 2019. A survey on new generation metaheuristic algorithms. *C&IE* (2019).
- [6] F Maxwell Harper and Joseph A Konstan. 2015. The movielens datasets: History and context. *ACM tiis* 5, 4 (2015), 1–19.
- [7] Folasade Olubusola Isinkaye, Yetunde O Folajimi, and Bolande Adefowoke Ojokoh. 2015. Recommendation systems: Principles, methods and evaluation. *Egyptian informatics journal* 16, 3 (2015), 261–273.
- [8] Marius Kaminskis and Derek Bridge. 2017. Diversity, serendipity, novelty, and coverage: A survey and empirical analysis of beyond-accuracy objectives in recommender systems. *ACM TiIS* 7, 1 (2017), 2.
- [9] Manuel Lozano and Carlos Garcia-Martinez. 2010. Hybrid metaheuristics with evolutionary algorithms specializing in intensification and diversification: Overview and progress report. *C&OR* 37, 3 (2010), 481–497.
- [10] Silvano Martello and Paolo Toth. 1990. *Knapsack problems: algorithms and computer implementations*. John Wiley & Sons, Inc., Hoboken, New Jersey, US.
- [11] K. Mehlhorn, B. R. Newell, P. M. Todd, M. D. Lee, K. Morgan, V. A. Braithwaite, D. Hausmann, K. Fiedler, and C. Gonzalez. 2015. Unpacking the Exploration-Exploitation Tradeoff: A Synthesis of Human and Animal Literatures. *Decision* 2, 3 (2015), 191–215. <https://doi.org/10.1037/dec0000033>
- [12] Zbigniew Michalewicz and Zbigniew Michalewicz. 1996. *Genetic algorithms+ data structures= evolution programs*. Springer Science & Business Media.
- [13] Junqin Xu and Jihui Zhang. 2014. Exploration-exploitation tradeoffs in meta-heuristics: Survey and analysis. In *Proceedings of the 33rd Chinese control conference*. IEEE, IEEE, Nanjing, China, 8633–8638.
- [14] Jun Zhang, Wei-Neng Chen, Zhi-Hui Zhan, Wei-Jie Yu, Yuan-Long Li, Ni Chen, and Qi Zhou. 2012. A survey on algorithm adaptation in evolutionary computation. *Frontiers of Electrical and Electronic Engineering* 7, 1 (2012), 16–31.