

CATALANO, G.A.P.I., BROWNLEE, A.E.I., CAIRNS, D., MCCALL, J. and AINSLIE, R. 2024. Mining potentially explanatory patterns via partial solutions. In *GECCO'24 companion: proceedings of the 2024 Genetic and evolutionary computation conference companion 2024 (GECCO'24 companion)*, 14-18 July 2024, Melbourne, Australia. New York: ACM [online], pages 567-570. Available from: <https://doi.org/10.1145/3638530.3654318>

Mining potentially explanatory patterns via partial solutions.

CATALANO, G.A.P.I., BROWNLEE, A.E.I., CAIRNS, D., MCCALL, J. and AINSLIE, R.

2024

© 2024 Author(s). This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *GECCO'24 Companion*, <https://doi.org/10.1145/3638530.3654318>

Mining Potentially Explanatory Patterns via Partial Solutions

GianCarlo A.P.I. Catalano

Alexander E.I. Brownlee

David Cairns

g.a.catalano@stir.ac.uk

alexander.brownlee@stir.ac.uk

dec@cs.stir.ac.uk

University of Stirling

Computing Science and Mathematics

Stirling, UK

John McCall

j.mccall@rgu.ac.uk

Robert Gordon University

National Subsea Centre

Aberdeen, UK

Russell Ainslie

russell.ainslie@bt.com

BT Technology

Applied Research

Ipswich, UK

ABSTRACT

We introduce Partial Solutions to improve the explainability of genetic algorithms for combinatorial optimization. Partial Solutions represent beneficial traits found by analyzing a population, and are presented to the user for explainability, but also provide an explicit model from which new solutions can be generated. We present an algorithm that assembles a collection of explanatory Partial Solutions chosen to strike a balance between simplicity, high fitness and atomicity, that are shown to be able to solve standard optimization benchmarks.

CCS CONCEPTS

• **Computing methodologies** → **Genetic algorithms**; • **Theory of computation** → *Models of computation*.

KEYWORDS

Genetic Algorithms, Explainable AI (XAI), Combinatorial Optimization Problems

ACM Reference Format:

GianCarlo A.P.I. Catalano, Alexander E.I. Brownlee, David Cairns, John McCall, and Russell Ainslie. 2024. Mining Potentially Explanatory Patterns via Partial Solutions. In *Genetic and Evolutionary Computation Conference (GECCO '24 Companion)*, July 14–18, 2024, Melbourne, VIC, Australia. ACM, New York, NY, USA.

INTRODUCTION

Genetic Algorithms (GAs) have proven their ability to solve many optimization problems, but have made relatively little progress in the aspect of **explainability** and still struggle to gain the user's trust. As GAs are adopted in critical applications [9, 16], the proposed solutions cannot be accepted at face value due to the many issues that can negatively affect the search process [16].

Combinatorial optimization problems are often tackled using GAs, since they are known to be computationally hard (such as Graph Coloring [6], Knapsack [15] and resource allocation [12]).

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 '24 Companion, July 14–18, 2024, Melbourne, VIC, Australia

© 2024 Copyright held by the owner/author(s).

These kinds of problems are the topic of this paper, as they present open questions relating to explainability:

- **What characterizes good solutions?**
- **Why is a solution good, or better than another?**

This work offers one way to address these questions and introduces the concept of a *Partial Solution* (PS), a positive trait commonly found in high-fitness solutions. More specifically, in this work they are patterns where certain variables are fixed and others can vary (indicated using $*$). For example, if the solutions 0001, 0011, 0101 and 0111 have high fitness, these would be described using $0**1$. A population of solutions is analyzed to find a set of recurring positive traits, the **Partial Solution catalog**. This acts as a model that is then sampled to generate new solutions. In doing so, the algorithm is inherently explainable, offering the ability for the end user to identify key components of the recommended solutions. Partial Solutions are an explicit decomposition of the problem, which offers the benefits suggested in [4]. Partial Solutions are **interpretable**:

- **Global explanations**: describing the fitness landscape by finding simple and atomic Partial Solutions which are associated with high-fitness.
- **Local explanations**: describing a solution by pointing out the Partial Solutions it contains.

Additionally, Partial Solutions **assist the search process**, because they act as a model that describes high-fitness regions of the solution space (similar to Probabilistic Graphical Models) from which solutions may be constructed. The system we propose consists of:

- **PS Miner**: an algorithm which obtains the PS catalog from a reference population (Section 4),
- **Pick & Merge**: the algorithm which forms full solutions by combining elements from the PS catalog (Section 5).

They are composed in the following (one pass) sequence:

- (1) Generate a reference population P_{Ref} and evaluate it.
- (2) Apply the PS Miner on P_{Ref} to generate the PS catalog.
- (3) Apply Pick & Merge on the PS catalog to obtain full solutions.

These Partial Solutions offer a novel way of providing both global and local explanations, but at the same time improve the search efficiency when solving the problem. Using benchmark functions which have known underlying structures, in Section 6 we test whether we're able to consistently find these Partial Solutions, and how their ability to solve the optimization problem by finding the optima compares against traditional methods.

2 RELATED WORK

Schemata and **backbones** are similarly structured to Partial Solutions, but they are used very differently. **Schemata** are abstractions that were used for theoretical work on GAs, although there are works which do explicitly construct them in order to improve the search efficiency [5, 8], but they are not used for explainability.

Backbones relate to SAT problems where multiple satisfactory solutions share a common sub-configuration. For most problems there is only one backbone, and it is preferred for it to have many fixed variables in order to shrink the search space [13], which is in contrast with the objectives used here.

Since in this work the PS catalog is treated as a model, EDAs are also an adjacent topic. In this sense, there are some related works: **Linkage learning** [1, 10] and **Family-of-Subsets** [14] aim to find which variables are interacting in the fitness function, and are then assembled into structures such as sets or cliques. These structures only determine which variables are interacting, but their optimal value assignments are not specified, whereas PSs contain both. In other words, PSs are instances of the cliques / linked sets found by the aforementioned approaches.

3 PARTIAL SOLUTIONS

3.1 Formal definition

Table 1: Notation

Symbol	Meaning
F	Fitness function to be maximized
F^ψ	Fitness function for partial solutions
n	The number of parameters in the solutions
a, b, \dots, x	Variables used to denote a <i>full</i> solution
x_1, x_2, \dots, x_n	The parameters of x , a full solution
$a^\psi, b^\psi, \dots, x^\psi$	Variables used to denote a <i>Partial</i> Solution
$x_1^\psi, x_2^\psi, \dots, x_n^\psi$	The parameters of x^ψ , a <i>Partial</i> Solution
P_{Ref}	A collection of evaluated full solutions

Table 1 summarizes our notation. A Partial Solution x^ψ can be described as “a sub-configuration of parameter values which is associated with high-fitness solutions within a reference population”, where that reference population will be denoted as P_{Ref} .




Figure 1: The positive traits in a collection of full solutions P_{Ref} can be described the PS catalog X^ψ

In the same way that a full solution is a tuple of values $x = (x_1, x_2, \dots, x_n)$, a Partial Solution is $x^\psi = (x_1^\psi, x_2^\psi, \dots, x_n^\psi)$ where each x_i^ψ can take any value that x_i can, as well as “*”, the “any” value. A solution x is said to “contain” the Partial Solution x^ψ when they agree on the fixed values found in x^ψ (Equation (1)). Conversely, it is possible to find all the solutions in P_{Ref} which contain the Partial Solution x^ψ , the **observations** (Equation (2)).

$$\text{contains}(x, x^\psi) \Leftrightarrow x_i = x_i^\psi \forall i \in [1, n], x_i^\psi \neq * \quad (1)$$

$$\text{obs}_{P_{\text{Ref}}}(x^\psi) = \{x \in P_{\text{Ref}} \mid \text{contains}(x, x^\psi)\} \quad (2)$$

Within all of possible PSs for an optimization problem, there is a small subset which is preferable for explainability: the **catalog**. These are the PSs which are optimal with respect to three metrics described in Section 3.3: **simplicity**, **mean fitness** and **atomicity**.

3.2 Explainability of Partial Solutions

A collection of PSs, found by analyzing a population, provides a succinct and expressive representation of positive traits (as seen in Figure 1), also described as “**innovization**” in the literature [4]. In a real problem, PSs can offer useful insight into what variables are linked, and what values they should take.

In an employee-allocation problem, for instance, these subconfigurations might indicate that it is beneficial to assign certain groups of employees to specific tasks. The solution 111111 might have its high fitness explained by the presence of 111*** and 1***1, although 111001 is even better because of the presence of ***00*.

Additionally, PSs might show the presence of positive traits that cannot coexist (e.g., **00 and *11*), which might be otherwise unclear. The ability to find such disagreeing PSs can help in explaining situations where generally positive traits are not necessarily found in the global optima.

3.3 Partial Solution metrics

This section will define three metrics, which are used as objectives in the search process to find PSs. **simplicity** (Equation (3)) represents the idea that having more parameters with value * (as opposed to fixed values) improves interpretability. **meanFitness** (Equation (4)) is the mean fitness of the observations of a PS, measuring to the effect the PSs generally has on fitness. **atomicity** (Equation (5)) is the more novel of the metrics, and represents the idea that a PSs should be “irreducible” and not “composite”. This is motivated by explainability, since we’re aiming to break down the problem into small, atomic parts. This is detected by measuring the contribution of each variable to the rest of the PS (Equation (6)).

$$\text{simplicity}(x^\psi) = |\{i \text{ s.t. } x_i^\psi = *\}| \quad (3)$$

$$\text{meanFitness}(x^\psi) = \text{average}(F(\text{obs}_{P_{\text{Ref}}}(x^\psi))) \quad (4)$$

$$\text{atomicity}(x^\psi) = \min \left\{ \text{contribution}(x^\psi, k) \mid k \in [1, n] \atop x_k^\psi \neq * \right\} \quad (5)$$

$$\text{contribution}(x^\psi, k) = F^\Sigma(x^\psi) \cdot \log \left(\frac{F^\Sigma(x^\psi)}{F^\Sigma(k^\psi) \cdot F^\Sigma(\bar{k}^\psi)} \right) \quad (6)$$

where $k_i^\psi = (x_i^\psi \text{ if } i = k, * \text{ otherwise})$

$\bar{k}_i^\psi = (* \text{ if } i = k, x_i^\psi \text{ otherwise})$

$m = \min_{x \in P_{\text{Ref}}} F(x)$

$S = \text{sum}\{F(x) - m \mid x \in P_{\text{Ref}}\}$

$F^\Sigma(x^\psi) = \text{sum}\{(F(x) - m) / S \mid x \in \text{obs}_{P_{\text{Ref}}}(x^\psi)\}$

4 FINDING PARTIAL SOLUTION CATALOGS

The algorithm to find the PS catalog (Algorithm 1) is implemented using the metrics in Section 3.1, but the design of the search method involves some non-trivial design choices:

- The **Reference Population** remains fixed throughout the process, making the system usable within a GA
- The three metrics are treated as a **single objective** by normalizing all their values in the range $[0, 1]$ between individuals, and averaging them within individuals.
- An **archive** is used in the search, storing all of the individuals which have been selected and that are not allowed in any future population. The final outputs come from the archive.
- The search could be implemented using a standard GA, or by starting from empty PSs which get progressively "**specialized**", or even from full solutions which get progressively "**simplified**" as PSs. All are tested via different implementations of `get_local` and `get_init` in Algorithm 1.

5 PICK & MERGE ALGORITHM

The PS catalog X^ψ obtained via Algorithm 1 can be used as a model to construct high-quality full solutions. **Pick & Merge** (Algorithm 2) simply picks items from the PS catalog and merges them when possible. At the end, any remaining unfixed parameters are filled with random values, so that a full solution can be returned.

6 EXPERIMENTS

We designed tests to answer the following research questions:

- **RQ1:** Which algorithm parameters are best for finding the PS catalog?
- **RQ2:** How is the ability to extract PSs from a reference population affected by evolution
- **RQ3:** Can the PS catalog be used to construct good solutions?

These were answered with testing rounds T1, T2 and T3 respectively. T1 and T2 consist of hyperparameter tuning, and T3 checks whether the original optimization problem is still being solved, all using the benchmark problems discussed in the next section.

6.1 Benchmark Problems

Three problems were chosen as benchmarks: Royal Road (RR, defined as [11], Trap-k (defined as [7]) and Royal Road with Overlaps (RRO). Each problem instance was constructed to have 5 "target" PSs: the cliques that we know the algorithm should find (e.g. 111**..** for Royal Road). Royal Road with Overlaps is a novel

Algorithm 1: Archive-based PS Miner

```

1 Def normalize(values):
2   return  $\frac{\text{values} - \min(\text{values})}{\max(\text{values}) - \min(\text{values})}$ 
3 Def  $F^\psi(X^\psi)$ :
4    $M \leftarrow \text{normalize}(\text{meanFitness}(X^\psi))$ 
5    $S \leftarrow \text{normalize}(\text{simplicity}(X^\psi))$ 
6    $A \leftarrow \text{normalize}(\text{atomicity}(X^\psi))$ 
7   return  $\{\text{avg}(m, s, a) \text{ for } m, s, a \in \text{zip}(M, S, A)\}$ 
8 Def top( $X^\psi$ , quantity):
9   sorted  $\leftarrow \text{sort}(X^\psi, \text{key} = F^\psi, \text{order} = \text{descending})$ 
10  return sorted[:quantity]

```

`get_init` and `get_local` are discussed in Section 4

```

11
12 Def mine_ps(get_init, get_local, pop_size, qty_ret):
13    $X^\psi \leftarrow \text{get\_init}()$ 
14   archive  $\leftarrow \{\}$ 
15   while termination_criteria_not_met do
16     selected  $\leftarrow \text{TournamentSelection}(X^\psi)$ 
17     localities  $\leftarrow \bigcup_{x^\psi \in \text{selected}} \text{get\_local}(x^\psi)$ 
18     archive  $\leftarrow \text{archive} \cup \text{selected}$ 
19      $X^\psi \leftarrow (X^\psi \cup \text{localities}) \setminus \text{archive}$ 
20      $X^\psi \leftarrow \text{top}(X^\psi, \text{pop\_size})$ 
21   return top(archive, qty_ret)

```

Algorithm 2: Pick & Merge algorithm

```

1 Def merge( $x^\psi, y^\psi$ ):
2   return  $z^\psi \text{ s.t. } z_i^\psi = x_i^\psi \text{ if } x_i^\psi \neq * \text{ else } y_i^\psi$ 
3 Def merge_from( $X^\psi$ , limit =  $\lceil \sqrt{n} \rceil$ ):
4    $Y^\psi \leftarrow \text{copy}(X^\psi)$ 
5    $x^\psi \leftarrow **, **, **$ 
6   added  $\leftarrow 0$ 
7   while  $Y^\psi \neq \emptyset$  & added < limit & has_*( $x^\psi$ ) do
8      $y^\psi \leftarrow \text{weighted\_random\_choice}(Y^\psi)$ 
9      $Y^\psi \leftarrow Y^\psi \setminus \{y^\psi\}$ 
10    if mergeable( $x^\psi, y^\psi$ ) then
11       $x^\psi \leftarrow \text{merge}(x^\psi, y^\psi)$ 
12      added  $\leftarrow \text{added} + 1$ 
13  return  $x^\psi$ 
14 Def fill_gaps( $x^\psi$ ):
15   for  $x_i^\psi \in x^\psi$  do
16     if  $x_i^\psi = *$  then
17        $x_i^\psi \leftarrow \text{random.randrange}(\text{cardinalities}[i])$ 
18  return as_solution( $x^\psi$ )
19 Def generate_via_pick_and_merge( $X^\psi$ , merge_limit):
20  return fill_gaps(merge_from( $X^\psi$ , merge_limit))

```

problem introduced here to test whether the algorithms struggle with overlapping building blocks. The fitness function of the problem is similar to Royal Road, but the groups of interest are allowed to overlap and to consists of all 0's as well. The cliques of RR and RRO consist of 4 bits, whereas Trap-k will have cliques of 5 bits.

6.2 Setup

Each algorithm configuration is executed 100 times, with a freshly generated P_{Ref} , and restricted to return at most 50 PSs.

T1 consists of Hyper-parameter tuning the PS Miner in PS population size, use of archive and other aspects, with fixed $|P_{\text{Ref}}| = 10^4$ (uniformly randomly generated) and F^ψ evaluation budget of 10^5 . A run is considered successful if all the target PSs of the problem are returned.

T2 investigates the effect of P_{Ref} on the miner, based on its size and the generations it has been evolved for.

Finally, in **T3** the entire system is used: a run consists of finding the PS Catalog using the PS miner (configured using T1 and T2) and passing it to Pick & Merge to generate 100 full solutions. The total evaluation budget varies, with differing percentages being dedicated to F^ψ evaluations. A run is considered successful when the global optima is found.

7 CONCLUSION

The results of the experiments are the following:

- **RQ1:** The best approach to finding PSs was **specialization**: to start from empty items and gradually fix the variables.




Figure 2: Results from T2, plotting the success rate (finding PS Catalogs) with a fixed P_{Ref} size of 10^4




Figure 3: Plot for T3, percentage of runs where the global optima was found, with F^ψ consuming 50% of the budget.

- **RQ2:** The proposed PS mining algorithm performs better on a large unevaluated population
- **RQ3** Applying Pick & Merge on the PS Catalog allows for the global optima to be found consistently, especially when F budget is close to the F^ψ budget.

The system proposed in this paper is capable of finding the intended PSs for the benchmark problems, and to solve the original optimization problem, at the cost of extra F^ψ evaluations. Source code for the system can be found at [2], and an extended version of this document can be found at [3].

8 FUTURE WORK

Overall, the system could be readapted to be a full EDA, for example by feeding the outputs of pick_and_merge into P_{Ref} .

The problems discussed in this document are admittedly simpler than **real-world problems**, and therefore it would be useful to see how well the PSs perform with some **inherently inelegant problems**, such as rostering and scheduling tasks.

Acknowledgment: PhD project supported by The Data Lab and BT Group plc. Many thanks go to the patience of my supervisors.

REFERENCES

- [1] Shumeet Baluja and Scott Davies. 1997. Using Optimal Dependency-Trees for Combinational Optimization. In *Proceedings of the Fourteenth ICML*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 30–38.
- [2] Giancarlo Catalano. 2024. *PS Assisted Explainability*. https://github.com/Giancarlo-Catalano/PS_Minimal_Showcase
- [3] Giancarlo A.P.I Catalano, Alexander E. I. Brownlee, David Cairns, John McCall, and Russell Ainslie. 2024. Mining Potentially Explanatory Patterns via Partial Solutions. arXiv:2404.04388
- [4] Kalyanmoy Deb and Aravind Srinivasan. 2008. Innovization: Discovery of innovative design principles through multiobjective evolutionary optimization. *Multiobjective Problem Solving from Nature* (2008), 243–262.
- [5] Jack McKay Fletcher and Thomas Wennekers. 2017. A natural approach to studying schema processing. arXiv:1705.04536 [cs.NE]
- [6] Michael R. Garey and David S. Johnson. 2006. Computers and Intractability: A Guide to the Theory of NP-Completeness. 24 (7 2006), 90–91. Issue 1. <https://doi.org/10.1137/1024022>
- [7] David E Goldberg. 1989. Genetic algorithms and Walsh functions: Part 2, Deception and its analysis. *Complex systems* 3 (1989), 153–171.
- [8] David E. Goldberg, Kumara Sastry, and Yukio Ohsawa. 2003. *Discovering Deep Building Blocks for Competent Genetic Algorithms Using Chance Discovery via KeyGraphs*. Springer, Berlin, Heidelberg, Tokyo, Japan, 276–301. https://doi.org/10.1007/978-3-662-06230-2_19
- [9] Katja Grace, John Salvatier, Allan Dafoe, Baobao Zhang, and Owain Evans. 2018. Viewpoint: When Will AI Exceed Human Performance? *Journal of Artificial Intelligence Research* 62 (7 2018), 729–754. <https://doi.org/10.1613/JAIR.1.11222>
- [10] Shih-Huan Hsu and Tian-Li Yu. 2015. Optimization by pairwise linkage detection, incremental linkage set, and restricted/back mixing: DSMGA-II. In *Proceedings of GECCO 2015*. ACM, New York, NY, USA, 519–526.
- [11] Melanie Mitchell, Stephanie Forrest, and John Holland. 1992. The Royal Road for Genetic Algorithms: Fitness Landscapes and GA Performance. *European Conference on Artificial Life* 1 (11 1992), 23–33.
- [12] L. Rabiner. 1984. Combinatorial optimization: Algorithms and complexity. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 32 (12 1984), 1258–1259. Issue 6. <https://doi.org/10.1109/TASSP.1984.1164450>
- [13] John Slaney and Toby Walsh. 2001. Backbones in Optimization and Approximation. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence - Volume 1* (Seattle, WA, USA) (IJCAI'01). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 254–259. <https://doi.org/10.5555/1642090.1642125>
- [14] Dirk Thierens and Peter Bosman. 2011. Optimal mixing evolutionary algorithms. In *Proceedings of the 13th GECCO conference*. ACM, New York, NY, USA, 617–624.
- [15] Pamela H. Vance. 2006. Knapsack Problems: Algorithms and Computer Implementations (S. Martello and P. Toth). 35 (7 2006), 684–685. Issue 4. <https://doi.org/10.1137/1035174>
- [16] Feiyu Xu, Hans Uszkoreit, Yangzhou Du, Wei Fan, Dongyan Zhao, and Jun Zhu. 2019. Explainable AI: A Brief Survey on History, Research Areas, Approaches and Challenges. 11839 LNAI (2019), 563–574. https://doi.org/10.1007/978-3-030-32236-6_51