

# Widened KRIMP: Better Performance through Diverse Parallelism

Oliver Sampson and Michael R. Berthold

Chair for Bioinformatics and Information Mining  
Department of Computer and Information Science  
University of Konstanz, Germany

**Abstract.** We demonstrate that the previously introduced Widening framework is applicable to state-of-the-art Machine Learning algorithms. Using KRIMP, an itemset mining algorithm, we show that parallelizing the search finds better solutions in nearly the same time as the original, sequential/greedy algorithm. We also introduce Reverse Standard Candidate Order (RSCO) as a candidate ordering heuristic for KRIMP.

## 1 Introduction

Research into parallelism in Machine Learning has primarily focused on reducing the execution time of existing algorithms, e.g., parallelized K-MEANS [23,17,14,26] and DBSCAN [11,4,7]. There have been some exceptions, such as *metalearning* and *ensemble methods* [9], which have employed heterogeneous algorithms in parallel, and [3], which describes the application to simple examples. Recent work [2,15] describes *Widening*, a framework for employing parallel resources to increase accuracy. With Widening, measures of diversity are used to guarantee the parallel search paths' exploration of disparate regions within a solution space, thereby stepping around the common greedy algorithmic tendency to find local optima. Thus far, work has concentrated on a proof-of-concept and demonstrative application to algorithms for solving the SET COVER PROBLEM and the creation of Decision Trees. This document describes the same approach, but with a state-of-the-art algorithm, KRIMP [24].

KRIMP finds "interesting" itemsets from a transactional database via the Minimum Description Length (MDL) principle [21]. The authors summarize the method as "the best set of patterns [being] the set of patterns that describes the data best," where the best set of itemsets is the set that provides the highest compression using MDL. The algorithm not only provides a solution to the problem of pattern explosion, thereby greatly reducing the set of itemsets used to generate association rules, but provides exceptional performance in other applications such as classification [24].

This paper demonstrates that it is possible to apply Widening to find even more interesting sets of itemsets than those found by the standard KRIMP algorithm.

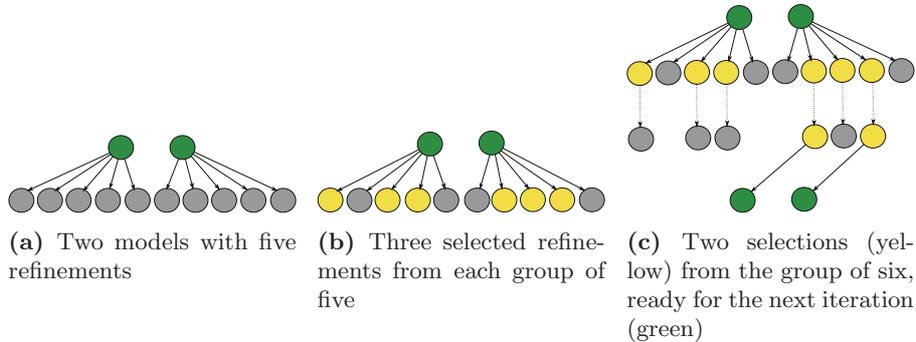


Fig. 1. Refine and Select with  $l = 3$  and  $k = 2$

## 2 Widening

Given the set of all models,  $\mathcal{M}$ , that describe the solution space for a typical greedy machine learning algorithm,  $m(\cdot) \in \mathcal{M}$  is a model which describes a portion of the solution space. It is iteratively refined by a *refinement operator*,  $r(\cdot)$ , based on a subset,  $x$ , from a training dataset  $\mathcal{T}$ , i.e.,  $m'(\cdot) = r(m(\cdot), x)$ ,  $x \subseteq \mathcal{T}$ . The derivation of a Decision Tree is one example of this process [15].

In contrast to that above, in the Widening framework a set of models,  $M \subseteq \mathcal{M}$ , is the result of a refinement operator based on data from  $\mathcal{T}$  and a *diversity metric*,  $\Delta$ , which describes some minimum difference between the resulting models,  $\{m'_1, \dots, m'_l\}$ , i.e.,  $M = r_\Delta(m) = \{m'_1, \dots, m'_l\}$ . For clarity, the data elements from the training data are eliminated from the notation.

A *selection operator*,  $s_{top-k}(\cdot)$  is employed to select the best  $k$  models at each step [15].  $M_{i+1} = s_{top-k}(r_\Delta(M_i))$ . The results of the selection operation are further refined until some stop condition is met. This iterative refine-and-select process, as depicted in Figure 1, is conceptually similar to a *beam search* [19].

## 3 The KRIMP Algorithm

In the area of Itemset Mining, KRIMP finds the set of “most interesting” itemsets in a transaction database based on MDL, i.e., KRIMP defines the best model,  $m$ , as the model that maximizes the compression of a transaction database,  $\mathcal{D}$ , encoded with that model and the compression of the model itself [21,24].

Given a database  $\mathcal{D}$  composed of transactions  $t \in \mathcal{D}$ , KRIMP finds the subset of itemsets,  $X$ , from the set of all itemsets,  $\mathcal{X}$ , that maximally compresses  $\mathcal{D}$ . KRIMP calculates the size of the encoded database using the codelengths of *prefix-free codes*, which are related to the frequency of the appearance of an itemset,  $x \in X$ , in the database and to the Shannon entropy:  $L(x) = -\log_2 P(x)$ , where  $L(x)$  is the codelength measured in bits of an item or itemset in the database, and  $P(x)$  is the relative frequency of the item’s or itemset’s appearance in the database [24].

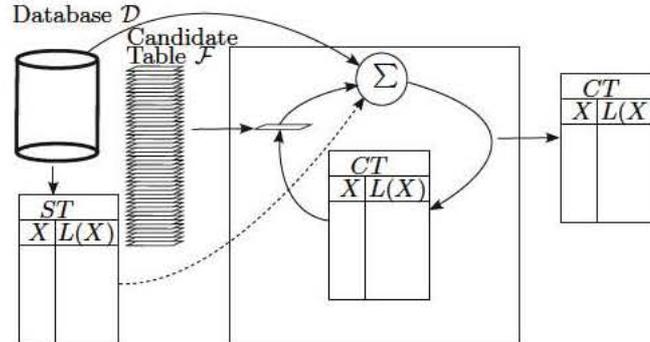


Fig. 2. The KRIMP algorithm

KRIMP begins with the generation of the Standard Code Table,  $ST$ , which is merely a code table comprised of only the individual items from the set of all single items,  $\mathcal{I}$ . The codelength of a given transaction,  $L(t)$  is the sum of its compositional codelengths.

KRIMP then iterates through a list of candidate itemsets,  $\mathcal{F}$ , generated by an algorithm external to KRIMP, such as AFOPT [18] or APRIORI [1]. Each of the candidate itemsets from  $\mathcal{F}$  is temporarily inserted into the code table  $CT$ , where all relative frequencies are determined and the compression evaluated. If it provides better compression, it is kept as part of  $CT$  and if not, it is discarded [24]. A general flow diagram is depicted in Figure 2.

The size of the encoded database,  $L(\mathcal{D}|CT)$ , is the sum of the encoded lengths of all transactions. The size of the encoded Code Table,  $L(CT|\mathcal{D})$ , is the size of each code plus the lengths of the encoded itemsets, for which the single items from  $ST$  are used. The compressed MDL size of the database is the size of the encoded database plus the size of the encoded code table.  $L(\mathcal{D}, CT) = L(\mathcal{D}|CT) + L(CT|\mathcal{D})$  [24]

Both  $\mathcal{F}$  and  $CT$  are ordered heuristically to maximize compression.  $\mathcal{F}$  is ordered according to the *Standard Candidate Order*, which orders primarily by the itemsets' *support* in descending order, secondarily by cardinality in descending order, and tertiarily by lexicographical order, as a tie-breaker. The rationale is that itemsets with larger support are likelier to cover more transactions and are evaluated first. Itemsets with the same support are sorted secondarily by cardinality, because larger itemsets cover more items in each transaction, reducing the number of itemsets or items required to cover a transaction [24].

$CT$  is ordered using the *Standard Cover Order*, which orders primarily by descending cardinality, secondarily by descending support, and tertiarily lexicographically, again as a tie-breaker. The rationale is that larger itemsets are preferred for their ability to cover more of each transaction. Of those, the ones with a larger support are more likely to cover more transactions in the database, thereby providing shorter codes [24].

KRIMP also includes a post-processing step for each iteration called PRUNING. If the relative frequency of any itemset in  $CT$  decreases as a result of adding a

new itemset,  $CT$  is re-evaluated with each itemset in  $CT$  singly removed. If any of the itemsets' temporary absence from  $CT$  enables a better compression, it is discarded. Results in [24] indicate that PRUNING improves overall compression performance marginally, but can dramatically reduce the number of itemsets providing that level of compression.

## 4 Widenend Krimp

A given path through a KRIMP solution space is based on two things: 1) the order with which the candidate itemsets from  $\mathcal{F}$  are evaluated, because the acceptance of a particular itemset into  $CT$  influences which itemsets are accepted in later iterations, and 2) the order in which the itemsets in  $CT$  are used to cover the database. Varying either of these two heuristics' orderings varies the solution path through the solution space and introduces diversity from the other paths taken.

For use as  $\Delta$  in the refining function,  $r_{\Delta}(\cdot)$ , two *explicit* measures and one implicit measure of diversity are investigated here. Explicit measures *p-dispersion-min-sum* and *p-dispersion-sum* select maximally diverse subsets of candidates from the candidate table. Implicit method, *Directed Placement*, is investigated with respect to the ordering of the itemsets evaluated for covering the transactions in  $CT$ .

*p-dispersion-min-sum* maximizes the sum of minimum distances between pairs of members of the selected subset [20].

**Definition 1** *p-dispersion-min-sum*.<sup>1</sup> Given a set  $\mathcal{F} = \{F_1, \dots, F_n\}$  of  $n$  itemsets and  $l$ , where  $l \in \mathbb{N}$  and  $l \leq n$ , and a distance measure  $Jaccard(F_i, F_j) : F_i, F_j \in \mathcal{F}$  between items  $F_i$  and  $F_j$ , the  $l$ -diversity problem is to select the set  $F : F \subseteq \mathcal{F}$ , such that

$$F^* = \max_{\substack{F \subseteq \mathcal{F} \\ |F|=l}} f(F), \text{ where } f(F) = \sum_{i=1}^l \min_{1 \leq j \leq l, i \neq j} Jaccard(F_i, F_j), F_i, F_j \in F [20][16] \quad (1)$$

*p-dispersion-sum* maximizes the distance between all members of the selected subset.

**Definition 2** *p-dispersion-sum*. Given a set  $\mathcal{F} = \{F_1, \dots, F_n\}$  of  $n$  itemsets and  $l$ , where  $l \in \mathbb{N}$  and  $l \leq n$ , and a distance measure  $Jaccard(F_i, F_j) : F_i, F_j \in \mathcal{F}$  between itemsets  $F_i$  and  $F_j$ , the  $l$ -diversity problem is to select the set  $F : F \subseteq \mathcal{F}$ , such that

$$F^* = \max_{\substack{F \subseteq \mathcal{F} \\ |F|=l}} f(S), \text{ where } f(F) = \frac{1}{l(l-1)} \sum_{i=1}^l \sum_{j>1}^l Jaccard(F_i, F_j) [20, 12] \quad (2)$$

<sup>1</sup> The canonical names from the literature, *p-dispersion-min-sum* and *p-dispersion-sum*, are maintained here, even though in this context, they should be called “*l-dispersion-min-sum*” and “*l-dispersion-sum*.”

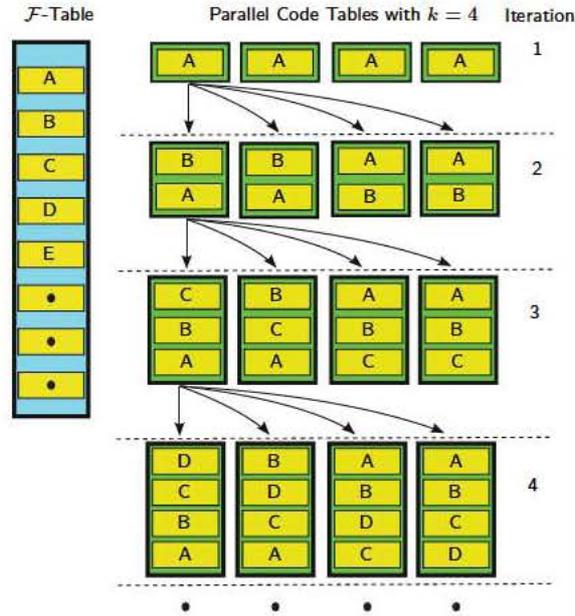


Fig. 3. At each iteration, the next  $F \in \mathcal{F}$  is inserted into each of the parallel code tables at a depth  $\frac{i}{l}$ . Refinement is shown only for one table in each iteration.

$p$ -dispersion-sum has the side-effect of pushing the selected members to the boundaries of the original set. This results in selected sets that are less diverse and representative of the dataset than those that are selected by  $p$ -dispersion-min-sum [20].

The Directed Placement diversity heuristic functions by inserting the next candidate itemset,  $F \in \mathcal{F}$  at a position with different fractional depths into  $l$  parallel instances of  $CT$ . The depth inserted into  $CT$  is a function of  $l$ , where the depth is  $\frac{i}{l}|CT| : i = 1, \dots, l$ . Because the role of each itemset in the covering algorithm is dependent on its position in  $CT$ , positioning  $F$  at different depths explores diverse solution paths. This method of diversity is implicit, because the diversity between different  $CT$  tables is not measured directly. See Figure 3.

An additional heuristic ordering of  $\mathcal{F}$  called *Reverse Standard Candidate Order* (RSCO) is introduced here. It orders the candidate itemsets primarily by cardinality in ascending order, secondarily by support in descending order, and tertiarily by lexicographical order as a tie-breaker. In combination with the Standard Cover Order heuristic for covering transactions, RSCO attempts to mimic the PRUNING subalgorithm; candidate itemsets with larger cardinality are examined later but are inserted before the smaller itemsets already in  $CT$ . With Standard Cover Order, small itemsets whose potential ability to efficiently cover transactions are “shadowed” by larger itemsets and have a lower relative frequency used for the compression calculation. In contrast, using RSCO, smaller itemsets that may have a beneficial effect, yet show up too late in the list to be considered with Standard Candidate Order, can still be evaluated.

## 5 Experimental Results

Compression of  $\mathcal{D}$  with  $CT$ ,  $L(\mathcal{D}, CT)$ , is compared to a baseline compression of  $\mathcal{D}$  encoded using only  $ST$ ,  $L(\mathcal{D}, ST)$ . This paper follows the convention of [24] in using the compression ratio measured in percentages, where lower is better.  $L\% = L(\mathcal{D}, CT)/L(\mathcal{D}, ST) * 100$ . [24]

The notation  $|CT \setminus \mathcal{I}|$  indicates the number of non-singleton itemsets used. For a given compression level, a smaller number of itemsets is considered more interesting. KRIMP optimizes for both  $L\%$  and  $|CT \setminus \mathcal{I}|$  by evaluating  $L\%$  first, and then bettering  $|CT \setminus \mathcal{I}|$  with the PRUNING subalgorithm.

All experiments were conducted in KNIME [6] and used APRIORI [1,8] with a *minsupport* of 1 to generate the set of *closed itemsets*. The datasets used were the LUCS-KDD-DN [10] discretized versions of the Breast Cancer Wisconsin (Original) [25] (*Breast*) and Pima Indians Diabetes Data Set [22] (*Pima*) datasets available from the UCI-ML Data Repository [5].

Evaluations in Sections 5.1 and 5.2 compare three methods, KRIMP<sub>Greedy</sub>, KRIMP<sub>RSCO</sub> and KRIMP<sub>Diverse</sub>. KRIMP<sub>Greedy</sub> refers to the baseline “standard” KRIMP implementation used within KNIME; KRIMP<sub>RSCO</sub> refers to the implementation in KNIME, using RSCO for ordering  $\mathcal{F}$  rather than Standard Candidate Order, because the results with RSCO for the *Breast* and *Pima* datasets were actually better than Standard Candidate Order;<sup>2</sup> and KRIMP<sub>Diverse</sub> refers to KRIMP with a method of diversity being evaluated. KRIMP has two performance metrics for a model; solution pairs are shown in the form of  $\langle L\%, |CT \setminus \mathcal{I}| \rangle$ . Results are shown with and without PRUNING for KRIMP<sub>Greedy</sub> and KRIMP<sub>RSCO</sub>. All experiments with KRIMP<sub>Diverse</sub> were performed without PRUNING because we felt it would introduce another variable of diversity for which we were not controlling.

Experimental solution pairs are also shown at the position found  $\langle l, k \rangle$ , where  $l$  is the number of refined models and  $k$  is the number of models selected according to compression performance.

### 5.1 Diverse Candidate Selection

Both of these methods of subset selection are performed with replacement, because early experiments without replacement on  $\mathcal{F}$  candidate tables generated from closed itemsets demonstrated that there were simply not enough candidate itemsets for evaluation to generate solutions sets of reasonable performance for larger values of  $k$ . With replacement, in order for the algorithm to come to completion, the first element in  $\mathcal{F}$  is removed after each iteration, ensuring the algorithm’s completion after  $|\mathcal{F}|$  iterations. This method naturally entails a dependency on the initial ordering of  $\mathcal{F}$ .

A summary of the results for the *Breast* dataset using the  $p$ -dispersion-min-sum diversity metric for  $\mathcal{F}$  candidate selection can be found in Table 1. The experiments

<sup>2</sup> This is just an artifact for these two datasets. Preliminary results not shown here demonstrate that RSCO does indeed perform better than SCO for some datasets, albeit not consistently across all datasets tested.

**Table 1.** *Breast* Dataset Results Summary

	Heuristic		Pruning	L%	$ CT \setminus \mathcal{I} $
	$\mathcal{F}$ (Candidate Table)	$CT$ (Code Table)			
KRIMP <sub>Greedy</sub>	Standard Candidate Order	Standard Cover Order	no	18.11	29
			yes	17.61	28
KRIMP <sub>RSCO</sub>	RSCO	Standard Cover Order	no	17.86	28
			yes	17.82	26
KRIMP <sub>Diverse</sub>	$p$ -dispersion-min-sum + RSCO	Std. Cover Order	no	17.97	28
	$p$ -dispersion-sum + RSCO	Std. Cover Order		19.42	34
	RSCO	Directed Placement		<b>17.39</b>	<b>26</b>

**Table 2.** *Pima* Dataset Results Summary

	Heuristic		Pruning	L%	$ CT \setminus \mathcal{I} $
	$\mathcal{F}$ (Candidate Table)	$CT$ (Code Table)			
KRIMP <sub>Greedy</sub>	Standard Candidate Order	Standard Cover Order	no	35.6	66
			yes	34.4	53
KRIMP <sub>RSCO</sub>	RSCO	Standard Cover Order	no	34.3	63
			yes	33.7	<b>49</b>
KRIMP <sub>Diverse</sub>	RSCO	Directed Placement	no	<b>32.9</b>	56

were run with all combinations of  $l \in \{5, 10, 20, 30, 40, 50\}$  and  $k \in \{1, 5, 10, 15\}$ . The best solution for KRIMP<sub>Diverse</sub> with  $p$ -dispersion-min-sum, (17.97%, 28) was found at  $\langle l, k \rangle = \langle 50, 10 \rangle$ , which was better than KRIMP<sub>Greedy</sub> without PRUNING (18.11%, 29), but not better than KRIMP<sub>Greedy</sub> with PRUNING (17.61%, 28). KRIMP<sub>RSCO</sub> with PRUNING performed even better at (17.82%, 26).

KRIMP<sub>Diverse</sub> with  $p$ -dispersion-sum was run with all combinations of  $l \in \{5, 10, 20, 30, 40, 50\}$  and  $k \in \{1, 5, 10, 20, 50\}$  and as expected, the best solution pair (19.42%, 34) was not nearly as good as that with  $p$ -dispersion-min-sum, and was found in an even larger search space of  $\langle l, k \rangle = \langle 50, 30 \rangle$ .

The experiments with KRIMP<sub>Diverse</sub> with  $p$ -dispersion-min-sum were run over a smaller search space when compared to KRIMP<sub>Diverse</sub> with  $p$ -dispersion-sum after recognizing that the results had already reached the goal of beating one of the KRIMP<sub>Greedy</sub> scores.

Due to run-time constraints (See Section 6.) experiments were not performed with diversity-based candidate selection on the *Pima* dataset.

## 5.2 Diverse Cover Order

The results of the Directed Placement heuristic with RSCO as the Candidate Selection heuristic are also summarized in Tables 1 and 2. Experiments were performed on the *Breast* dataset with all combinations of  $l \in \{5, 10, 20, 30, 40, 50\}$  and  $k \in \{1, 5, 10, 15\}$ . The heuristic found a solution (17.39%, 26) outperforming the best KRIMP variant. Additionally, the solution was found in a much smaller search space, when compared to Diverse Candidate Selection, with the best solution found first at  $\langle l, k \rangle = \langle 10, 10 \rangle$ .

The *Pima* dataset showed an interesting property in that paths to better solutions were sensitive to increasing  $k$ , the number of selected models to be refined in the next step, but not in  $l$ , the number of refinements in each step. Better results,  $\langle 33.2\%, 61 \rangle$ , than  $\text{KRIMP}_{\text{Greedy}}$  and  $\text{KRIMP}_{\text{RSCO}}$  both without PRUNING were found immediately at  $\langle l, k \rangle = \langle 5, 5 \rangle$ . In fact, all solutions found with the Directed Placement heuristic with RSCO showed better  $L\%$  than either  $\text{KRIMP}_{\text{Greedy}}$  or  $\text{KRIMP}_{\text{RSCO}}$  with or without PRUNING. The best solution found for the *Pima* dataset was  $\langle 32.9\%, 56 \rangle$  at  $\langle l, k \rangle = \langle 5, 50 \rangle$ . This result has significantly better compression and yields nearly the number of itemsets as  $\text{KRIMP}_{\text{Greedy}}$ , but not nearly as good as  $\text{KRIMP}_{\text{RSCO}}$ .

## 6 Discussion and Future Work

In general, absolute timing values are not necessary for timing comparisons. To a first order of approximation, KRIMP runs in  $O(|\mathcal{F}| \times |\mathcal{D}| \times \theta)$  where  $\theta$  is a factor describing the average length of  $CT$  during the entire execution of the algorithm. (It should be noted that the authors of [24] saw a performance improvement in execution speed after implementing the PRUNING subalgorithm, because of a smaller value of  $\theta$ .) Accounting for application of a diversity measure and the use of a performance measurement for selection, KRIMP runs in  $O((|\mathcal{F}| + \Delta + \Psi) \times |\mathcal{D}| \times \theta)$ , where  $O(\Delta)$  is the measure of the complexity of the diversity heuristic, and  $O(\Psi)$  is a measure of the complexity of the performance measurement.

Although  $p$ -dispersion-min-sum was able to find comparable results to the standard KRIMP implementation (better than  $\text{KRIMP}_{\text{Greedy}}$  without PRUNING), the computational cost is significant. Selecting a subset of  $p$  diverse elements from a larger set is a variation of the  $p$ -dispersion problem and is  $\mathcal{NP}$ -hard [13]. Moreover, a comparison of this metric to  $p$ -dispersion-sum demonstrates what could be a pitfall for applying  $p$ -dispersion-sum: a much wider solution space had to be searched,  $\langle l, k \rangle = \langle 50, 10 \rangle$  versus  $\langle l, k \rangle = \langle 50, 30 \rangle$ . Although at least one of these diversity measures fulfills the desire to show that widened data mining can find better solutions than the traditional greedy algorithm, it is insufficient for a requirement of finding better solutions in the same or less time than the traditional greedy algorithm, which is the ultimate goal of Widening.

Directed Placement, however, was able to significantly improve on the solution found by standard KRIMP in [24]. For the *Breast* dataset, the results were even better than the results found with  $\text{KRIMP}_{\text{RSCO}}$ . Directed Placement also showed a partially better solution with the *Pima* dataset. In comparison to the other diversity metrics presented here, Directed Placement has a much smaller overhead for generating diverse solution paths. It must be noted, however, that the claim of “better solutions in the same or faster time” in this case is not strictly accurate. For large values of  $|\mathcal{F}|$  and  $|\mathcal{D}|$ , the influence of  $O(\Delta)$  for the Directed Placement diversity heuristic becomes negligible. The evaluation of the models for selection,  $O(\Psi)$  is also negligible for KRIMP, because it is merely a comparison of the best  $L\%$ . Additionally, Directed Placement provided the best Widened result in a significantly smaller search region than the other diversity heuristics.

Ideally, a Widened algorithm is able to find a better or the best solution in the same time as the traditional greedy algorithm. The pitfall of the methods described here is that both a performance evaluation ( $\Psi$ ) and a synchronized comparison of results from the parallel workers are required. This would be avoided with a *communication-less* [15] approach where the parallel workers would be able to refine and select without requiring a synchronized comparison step. Additionally, although the better solutions found by Widened KRIMP meet the definition of “better,” further research into how well the smaller sets perform as classifiers or in other KRIMP applications is necessary. The effects of including the PRUNING subalgorithm on the dataset compression, and the corresponding solution space paths also require further investigation, as does the magnitude and interplay between  $l$  and  $k$  for different datasets.

## 7 Conclusion

In this paper we have validated Widening for the first time using a state-of-the-art algorithm for itemset mining, KRIMP, and shown that it is possible to use the novel approach of Widening to find significantly better solutions than that of the traditional greedy algorithm by searching diverse regions of a solution space in parallel. We have also introduced RSCO, a new Candidate Table ordering heuristic for KRIMP that can provide even better results for some datasets.

## References

1. Agrawal, R., Srikant, R.: Fast algorithms for mining association rules. In: Proceedings of the 20th International Conference on Very Large Data Bases, vol. 1215, pp. 487–499 (1994)
2. Akbar, Z., Ivanova, V.N., Berthold, M.R.: Parallel data mining revisited. Better, not faster. In: Hollmén, J., Klawonn, F., Tucker, A. (eds.) IDA 2012. LNCS, vol. 7619, pp. 23–34. Springer, Heidelberg (2012)
3. Akl, S.G.: Parallel real-time computation: Sometimes quantity means quality. In: Proceedings of the International Symposium on Parallel Architectures, Algorithms and Networks, I-SPAN 2000, pp. 2–11. IEEE (2000)
4. Arlia, D., Coppola, M.: Experiments in parallel clustering with DBSCAN. In: Sakellariou, R., Keane, J.A., Gurd, J.R., Freeman, L. (eds.) Euro-Par 2001. LNCS, vol. 2150, pp. 326–331. Springer, Heidelberg (2001)
5. Bache, K., Lichman, M.: UCI Machine Learning Repository (2013)
6. Berthold, M.R., Cebron, N., Dill, F., Gabriel, T.R., Kötter, T., Meinl, T., Ohl, P., Sieb, C., Thiel, K., Wiswedel, B.: KNIME: The Konstanz Information Miner. In: Preisach, C., Burkhardt, H., Schmidt-Thieme, L., Decker, R. (eds.) Data Analysis, Machine Learning and Applications - Proceedings of the 31st Annual Conference of the Gesellschaft für Klassifikation e.V (GfKL 2007), Berlin, Germany. Studies in Classification, Data Analysis, and Knowledge Organization, pp. 319–326 (2007)
7. Böhm, C., Noll, R., Plant, C., Wackersreuther, B., Zherdin, A.: Data mining using graphics processing units. In: Hameurlain, A., Küng, J., Wagner, R. (eds.) Transactions on Large-Scale Data- and Knowledge-Centered Systems I. LNCS, vol. 5740, pp. 63–90. Springer, Heidelberg (2009)

8. Borgelt, C., Kruse, R.: Induction of association rules: Apriori implementation. In: *Compstat*, pp. 395–400. Springer (2002)
9. Chan, P., Stolfo, S.J.: Experiments on multistrategy learning by meta-learning. In: *Proceedings of the Second International Conference on Information and Knowledge Management*, pp. 314–323 (1993)
10. Coenen, F.: LUCS-KDD DN software (2003)
11. Dhillon, I.S., Modha, D.S.: A data-clustering algorithm on distributed memory multiprocessors. In: Zaki, M.J., Ho, C.-T. (eds.) *KDD 1999. LNCS (LNAI)*, vol. 1759, pp. 245–260. Springer, Heidelberg (2000)
12. Drosou, M., Pitoura, E.: Comparing diversity heuristics. Technical report, Technical Report 2009-05. Computer Science Department, University of Ioannina (2009)
13. Erkut, E.: The discrete p-dispersion problem. *European Journal of Operational Research* 46(1), 48–60 (1990)
14. Farivar, R., Rebolledo, D., Chan, E., Campbell, R.: A parallel implementation of k-means clustering on GPUs. In: *Proceedings of International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA)*, pp. 340–345 (2008)
15. Ivanova, V.N., Berthold, M.R.: Diversity-driven widening. In: *Proceedings of the 12th International Symposium on Intelligent Data Analysis (IDA 2013)* (2013)
16. Jaccard, P.: Étude comparative de la distribution florale dans une portion des Alpes et des Jura. *Bulletin del la Société Vaudoise des Sciences Naturelles* (1901)
17. Kantabutra, S., Couch, A.L.: Parallel k-means clustering algorithm on nows. *NECTEC Technical Journal* 1(6), 243–247 (2000)
18. Liu, G., Lu, H., Yu, J.X., Wei, W., Xiao, X.: AFOPT: An efficient implementation of pattern growth approach. In: *Proceedings of the ICDM Workshop on Frequent Itemset Mining Implementations* (2003)
19. Lowerre, B.T.: The HARP speech recognition system. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, USA (1976)
20. Meinl, T.: Maximum-Score Diversity Selection. PhD thesis, University of Konstanz (July 2010)
21. Rissanen, J.: Modeling by shortest data description. *Automatica* 14(5), 465–471 (1978)
22. Smith, J.W., Everhart, J.E., Dickson, W.C., Knowler, W.C., Johannes, R.S.: Using the adap learning algorithm to forecast the onset of diabetes mellitus. In: *Proceedings of the Symposium on Computer Applications and Medical Care*, vol. 261, p. 265 (1988)
23. Stoffel, K., Belkoniene, A.: Parallel k/h-means clustering for large data sets. In: Amestoy, P.R., Berger, P., Daydé, M., Duff, I.S., Frayssé, V., Giraud, L., Ruiz, D. (eds.) *Euro-Par 1999. LNCS*, vol. 1685, pp. 1451–1454. Springer, Heidelberg (1999)
24. Vreeken, J., van Leeuwen, M., Siebes, A.: Krimp: Mining itemsets that compress. *Data Mining and Knowledge Discovery* 23(1), 169–214 (2011)
25. Wolberg, W.H., Mangasarian, O.L.: Multisurface method of pattern separation for medical diagnosis applied to breast cytology. *Proceedings of the National Academy of Sciences* 87(23), 9193–9196 (1990)
26. Zhao, W., Ma, H., He, Q.: Parallel k-Means Clustering Based on MapReduce. In: Jaatun, M.G., Zhao, G., Rong, C. (eds.) *Cloud Computing. LNCS*, vol. 5931, pp. 674–679. Springer, Heidelberg (2009)