## BAŞARIM 2020.

Yüksek High Başarımlı Performance Hesaplama Computing Konferansı Conference

HPC in small packages: Viability of Small Machines facing Big Data, a study in leveraging highexpressive data helps improve iterative algorithms.

> Mehmet M. Dalkılıç Dr. Hasan Kurban, Dr. Mark Jenne, Parichit Sharma Indiana University, Bloomington, IN, USA

Ankara, Turkey

NIH NATIONAL CANCER INSTITUTE

#### Outline

#### I. Context

Data Explosion and consequences

#### II. Data-Driven Problem Statement (Clustering)

- Using k-means and EM (iterative)
- Low Expression (LE), High Expression (HE)
  - HE affects the outcome the most

#### III. Solution

- Separate LE, HE, using the HE as much as possible, ignoring LE as much as possible
- Implement DCEM R package (Data Clustering using EM\*)

#### IV. Results

- EM\* work on big data
- Case Studies

#### V. Future Work

Parallelize, Distributed algorithm (Extend DCEM)

Başarim 2020

Zettabytes of Data Generated and Colleged as a Function of Time



Zettabytes of Data Generated and Colleged as a Function of Time









#### A taxonomy of data reduction techniques



#### I. Wikipedia pages...



#### I. Related Publications

- Hussein Mohsen, Hasan Kurban, Mark Jenne, Mehmet Dalkilic: A new set of Random Forests with varying dynamic data reduction and voting techniques, IEEE International Conference on Data Science and Advanced Analytics (DSAA'14), Shanghai, China, 2014
- Mark Jenne, Owen Boberg, Hasan Kurban and Mehmet M. Dalkilic: Studying the Milky Way Galaxy using ParaHeap-k, a parallel heap-based k-means, Computer, vol.47, no.9, pp.26-33, 2014
- Hasan Kurban, Mark Jenne and Mehmet M. Dalkilic: EM\*: An EM algorithm for Big Data, IEEE International Conference on Data Science and Advanced Analytics (DSAA'16), Montreal, Canada, 2016 (received "Honorable Mention Paper Award", best paper awards)
- Hussein Mohsen, Hasan Kurban, Kurt Zimmer, Mark Jenne, Mehmet Dalkilic: Red-RF: Reduced Random Forest for big data using priority voting & dynamic data reduction, International Congress on Big Data, IEEE BigData Congress 2015, New York, USA, 2015



## I. Related Publications/Products

Hasan Kurban, Mark Jenne and Mehmet M. Dalkilic: Using data to build a better EM: EM\* for big data, International Journal of Data Science and Analytics (JDSA), pp.1-15, 2017

 H. Kurban, M.M. Dalkilic: A novel approach to optimization of iterative machine learning algorithms: Over heap structure, 2017 IEEE International Conference on Big Data (Big Data), Boston, MA, 2017, pp. 102-109, doi: 10.1109/BigData.2017.8257917

Hasan Kurban, Can Kockan, Mark Jenne and Mehmet M. Dalkilic: Case study: clustering big stellar data with EM\*, IEEE/ACM International Conference on Big Data Computing. Applications and Technologies (BDCAT 2017)

DCEM: Clustering Big Data using Expectation Maximization Star (EM*) Algorithm			
Implements the Improved Expectation Maximisation EM* and the traditional EM algorithm for clustering big data (gaussian mixture mode structure. The implementation supports both random and K-means++ based initialization. Reference: Hasan Kurban, Mark Jenne, Mehmet			
Version:	2.0.4		
Depends:	R (≥ 3.2.0)		
Imports:	mvtnorm (≥ 1.0.7), matrixcalc (≥ 1.0.3), MASS (≥ 7.3.49), Rcpp (≥ 1.0.2)		
LinkingTo:	<u>Rcpp</u>		
Suggests:	knitr, <u>rmarkdown</u>		

Under Journal Review



## II. Data-Driven Problem Statement (Clustering)

We conducted a series of experiments comparing

- KM\* to KM
- EM\* to EM
- Parallel implementations, ParaHeap-k to k-means++ and k-means||

with both real world and synthetic data sets through different kinds of testing: scale, dimension, and separability.

## II. Data-Driven Problem Statement (Clustering)

- 14
- 1<sup>st</sup> Call from IEEE Computer on Computational Astronomy
  - Question: Cluster the Milky Way?
  - Yes: Difficult because of the data
- Let's do it!
  - Each star is "given birth" from one of the galactic components
  - We can check the quality of the cluster using metallicity





## II. Data-Driven Problem Statement (Clustering)

- What, at that time, would be the best clustering
- Quick change in nomenclature

For a non-empty finite set X, a partition is  $\{X_1, X_2, \ldots, X_n\}$  s.t.

•  $\cup_{i=1}^n X_n = X$ 

16

- for  $i, j = 1, 2, \dots, n; i \neq j \rightarrow X_i \cap X_j = \emptyset$
- for  $i = 1, 2, \ldots, n; X_i \neq \emptyset$
- for  $i = 1, 2, \ldots, n; X_i \subseteq X, X_i$  is called a <u>block</u>
- So the question becomes, can we partition the Milky Way so that each star belongs to its proper block (in this case, galactic component)

- 17
- Can we materially improve k-means from a datacentric perspective to cluster big stellar data?
   YES!
- Can this be generalized for other iterative machine learning algorithms?
  - YES!
    - Separating LE/HE data

- □ k-means algorithm (60 yrs.) [hard]
  - Assign each datum to <u>one</u> block (Centroid ~ "Best Representative" ~ Ave)
  - Use simple Euclidean distance
  - Re-calculate Centroid until convergence
- Gaussian Mixture Model (50 yrs.) [soft]
  - Assign datum to <u>every</u> Gaussian (block) with a probability
  - Re-calculate properties of Gaussians
  - Iterate until convergence though it can be very slow (another ascent method)

## II. Data-Driven Problem Statement (Clustering)

#### 19

- We decided on k-means because of its general success and ubiquitous use and remains one of the most popular algorithms (Lloyd's algorithm)
  - Data Clustering: 50 Years Beyond K-Means [Jain] 2009
  - Initialization (seeding)
    - k-means++ [Arthur, Vassilvitskii] 2006, 2007
    - <u>k-means</u> | (Scalable k-means) [Bahman et al.] 2012
    - kd-trees pre-filtering [Kanungu, et al.] 2000
    - kd-trees [Pelleg, Moore] 1999
  - Triangle Inequality + distance bounds
    - [Elkan] 2003 (Excellent paper for bibliography on speed-ups for k-means)
    - [Hamerly] 2010
    - [Drake] 2012
    - [Hamerly, Drake] 2014 use priority queues to prune points close to assigned centers; use heaps to store differences between lower and upper bounds of points
- Witness where it's used: ELKI, graphlab, Mahout, MATLAB, MLPACK, Octave, OpenCV, R (various), SciPy, Weka, Yael, etc.<sup>Başarim 2020</sup>

## II. Data-Driven Problem Statement (Clustering)

20



**Figure 1** Learning problems: dots correspond to points without any labels. Points with labels are denoted by plus signs, asterisks, and crosses. In (c), the must-link and cannot-link constraints are denoted by solid and dashed lines, respectively (figure taken from [Lange et al., 2005]).

[Jain] 2009 Pattern Recognition Letters

#### Similis simili gaudet

#### 21

repeat

Algorithm 1: K-Means (KM) Algorithm

 $\begin{array}{l} \text{K-Means}(\mathbf{D}, \, k, \, \epsilon): \\ t \leftarrow 0 \\ // \text{ Set of Centroids } \mathbf{C}^t = \{C_1^t, \dots, C_k^t\} \\ // \text{ Assume Centroid is Class } C_i = (\bar{\mathbf{x}}_i, D_i), \, D_i \subseteq D \\ \text{Randomly initialize } C_i.\bar{\mathbf{x}}_1, \dots, C_k.\bar{\mathbf{x}}_k \end{array}$ 

Iteration continues until the set of centroids is stable; in other words, convergence is guaranteed in a finite number of steps by showing that for some non-negative error function, monotonically decreasing during each iteration



## II. Data-Driven Problem Statement (Clustering)

- $\Box$  The run-time of k-means is O(inkd) where
  - i is the iteration
  - $\bullet$  n is data size

22

- k is the number of clusters
- d is the dimension of data, i.e.,  $\mathbf{x} \in \mathcal{R}^d$

#### Başarim 2020

- Each cluster  $C_i = {\mathbf{x}_{i1}, \ldots, \mathbf{x}_{i\ell}}$  contributes to sum of pairwise dissimilarities to sum of total measure
- For cluster  $C_i$ , we can write total measure measure(i), where  $|C_i D_i| = \ell$  as:

$$J(i) = \frac{1}{2} \sum_{j=1}^{\ell} \sum_{l=1}^{\ell} \|\mathbf{x}_{ij} - \mathbf{x}_{il}\|$$
(1)  
inear algebra 
$$= \ell \sum_{j=1}^{\ell} \|\mathbf{x}_{ij} - C_i \cdot \bar{\mathbf{x}}_i\|$$
(2)

24

Cost: J is minimized monotonically

$$J = \sum_{i=1}^{k} \sum_{\mathbf{x} \in C_i} \min ||\mathbf{x} - C_i \cdot \bar{\mathbf{x}}||^2$$



- Can we materially improve k-means from a datacentric perspective to partition big stellar data?
- Can this be generalized for other iterative machine learning algorithms?
- □ YES!
  - Observing that data, in the original algorithm, is visited continually no matter its affect on the algorithm *i*, *n*
  - Data as it is used <u>changes</u> in its "expressiveness"—this notion is that the <u>affect of data is dependent on when</u> <u>it's used</u>.

- Low Expression (LE) and High Expression (HE) data affect the outcome minimally and significantly, respectively
- LE and HE can switch
- There is a tendency as k-means
- converges that most of the data
- becomes LE
- Can we capture this difference?



Başarim 2020

## **II. Failures to Capture LE and HE**

- Transparency—there were many failures to capture this notion—which made us doubt, maybe, this was a sound approach
  - Entropy failed
  - Itemset (level set) failed
  - Triangulation failed—bounds
  - Eureka—heaps! Now we needed to re-examine heaps

## II. Taste of Using LE,HE vs. k-

28

	data set		<i>k</i> = 1000	<i>k</i> = 5000	<i>k</i> = 10000	
010		init time (s)	0	0	0	
9101	GS1	total run time (s)	66	242	354	
ParaHeap-k		cluster error	0.097	0.089	0.084	
		init time (s)	0	0	0	
95N	GS2	total run time (s)	2540	6226	11785	
		cluster error	0.096	0.102	0.098	
	data set		<i>k</i> = 1000	<i>k</i> = 5000	k = 10000	
		init time (s)	46	158	313	
910K	GS1	total run time (s)	69	214	412	
<i>k</i> -means		cluster error	0.084	0.081	0.077	
		init time (s)	6934			
95M	GS2	total run time (s)	8713	failed	failed	
		cluster error	0.092			

#### II. HPC (system architecture)

29



Figure 2. An overview of the system architecture highlighting the heap data structures and parallel implementation for a single ParaHeap-*k* iteration in its entirety. Each iteration includes data association threads (DATs) and centroid update threads (CUTs).



- Like k-means lots of work done (mostly by statisticians) to improve movement
- Unlike k-means (hard assignment), EM is soft assignment—each datum belongs to typically a Gaussian.
- Used everywhere like k-means
- Expectation-maximization algorithm as a clustering algorithm(EM-T): Each block can be characterized as a probability distribution and the goal is to iteratively find the MLE of unknown parameters of blocks.

32

Algorithm 2: Expectation-Maximization (EM) Algorithm



33

EM-T general run-time:  $O(ik(d^3 + nd^2); \text{ best run-time } O(inkd))$ 

**E-step** invert  $\Sigma$ ; compute  $|\Sigma_i|$ ;  $O(d^3) \stackrel{k}{\to} O(kd^3)$ evaluate density;  $O(d^2) \stackrel{k,n}{\to} O(knd^2)$ **M-step** update  $\Sigma \stackrel{k}{\to} O(knd^2)$ Use D = KM-T  $Use D^{HE} EM^*$  $D = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}, \mathbf{x}_i \in \mathcal{R}^d$  $C_i^t = (\mu_i \in \mathcal{R}^d, \Sigma_i \in \mathcal{R}^{d \times d}, Pr(C_i) \in [0, 1], C_i.D_i \subseteq \mathbf{D})$  $C_i \sim \mathsf{N}(C_i.\mu, C_i.\Sigma)$  $f_i(\mathbf{x}|C_i) = (2\pi)^{-\frac{d}{2}} \det(\Sigma)^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{x}-\bar{\mathbf{x}})^T \Sigma^{-1}(\mathbf{x}-\bar{\mathbf{x}})}$ 

Maximum Likelihood Estimate (Calculus and Linear Algebra):  $\hat{\mathbf{x}} = \frac{1}{n} \sum_{i} \mathbf{x}_{i}$  $\hat{\mathbf{\Sigma}} = \Sigma_{i} (\mathbf{x}_{i} - \hat{\boldsymbol{\mu}})^{\mathsf{T}} (\mathbf{x}_{i} - \hat{\boldsymbol{\mu}})$ Başarim 2020

34



Data





Başarim 2020

36

- Heaps have excellent computational properties
- Use these to store and separate LE, HE
- A complete/nearly complete balanced binary tree with the "heap" property
  Binary (min) h



Typebinary tree/heapInvented1964Invented byJ. W. J. WilliamsJ. W. J. WilliamsSpaceO(n)O(n)SearchO(n)O(n)InsertO(1)O(n)Find-minO(1)O(log n)Delete-minO(log n)O(log n)	Binary (min) heap					
Invented1964J. W. J. WilliamsTime complexity in big O notationAlgorithmAverageWorst caseSpaceO(n)O(n)SearchO(n)O(n)InsertO(1)O(log n)Find-minO(1)O(log n)Delete-minO(log n)O(log n)	Туре	binary ti	ree/heap			
Invented byJ. W. J. WilliamsTime complexity in big O notationAlgorithmAverageWorst caseSpaceO(n)O(n)SearchO(n)O(n)InsertO(1)O(log n)Find-minO(1)O(log n)Delete-minO(log n)O(log n)	Invented	1964				
SearchO(n)O(n)O(n)InsertO(1)O(log n)Find-minO(1)O(log n)Delete-minO(log n)O(log n)	Invented by	vented by J. W. J. Williams				
AlgorithmAverageWorst caseSpaceO(n)O(n)SearchO(n)O(n)InsertO(1)O(log n)Find-minO(log n)O(log n)	Time complexity in big O notation					
Space $O(n)$ $O(n)$ Search $O(n)$ $O(n)$ Insert $O(1)$ $O(\log n)$ Find-min $O(1)$ $O(\log n)$ Delete-min $O(\log n)$ $O(\log n)$	Algorithm	Average	Worst case			
SearchO(n)O(n)InsertO(1)O(log n)Find-minO(1)O(1)Delete-minO(log n)O(log n)	Space	O(n)	O(n)			
Insert $O(1)$ $O(\log n)$ Find-min $O(1)$ $O(1)$ Delete-min $O(\log n)$ $O(\log n)$	Search	O(n)	O(n)			
Find-min $O(1)$ $O(1)$ Delete-min $O(\log n)$ $O(\log n)$	Insert	O(1)	O(log n)			
Delete-min O(log n) O(log n)	Find-min	O(1)	O(1)			
	Delete-min	O(log n)	O(log n)			

Bașarim 20

wikipedia

- Intuition was LE resides at top, HE resides at bottom of heap (min for k-means, max for EM) (investigating Laplacian)
- □ Why should LE remain at the top if it's LE?
- We identified a new property of heaps that we call strong and weak
  - Strong means from the root to some level *l* each level can be permuted, and the heap property remains true.
  - Weak means from some level *l* below the root to the leaves, one or more levels cannot be permuted without violating the heap property
  - Strong heaps are "semi"-monotonic—tends to want to be monotonic
    Basarim 2020



39



Başarim 2020

40



 $0 \ | \ 2 \ | \ 1 \ | \ 4 \ | \ 5 \ | \ 3 \ | \ 7 \ | \ 8 \ | \ 9 \ | \ 10 \ | \ 11 \ | \ 6$ 

41



 $0 \ | \ 2 \ | \ 1 \ | \ 4 \ | \ 5 \ | \ 3 \ | \ 7 \ | \ 8 \ | \ 9 \ | \ 10 \ | \ 11 \ | \ 6$ 

42



Observe that  $\mu_s h$  exist from the root to some level  $\ell$  and  $\mu_w h$  exist from  $\ell + 1$  to the root. For *N* distinct elements, the limit of the ratio of  $\mu_s h$  to  $\mu h$  goes to zero:

#### Lemma

$$\lim_{N\to\infty}\frac{f(\mu_s h)}{f(N)} = 0$$

As an example, for seven distinct elements, 48 of the 80 different heaps (60%) are  $\mu_s h$ . Adding only one more distinct element drops the percentage to less than 23% for possible 210 heaps.

44

					KM*
Data set	Never	Once	All But Once	Always	
Ringnorm	17%	18%	8%	26%	_
US Census	19%	14%	9%	25%	
KDD Cup	25%	6%	4%	28%	

Table of what percentage of data vs. how many times showed up in the leat nodes. This table reveals that a significant portion of the data can be clustered without any computation after first iteration.



- Why should LE remain at the top if it's LE?
- We observed what we identified as strong and LE DATA weak heaps **HE** Data LE Data HE DATA LE strong weak HE

#### II. Where is HE mostly?



Proportion of HE data in the  $\mu h$  (%)

Başarim 2020

### II. Where is HE mostly?



## III. Solutions (Model based)

48

There are several variants of the initial expectation-maximization algorithm (EM-T for *traditional* EM) that exist and most of them have been designed by the statistical community. Broadly, we can aggregate the strategies in three categories:

- (1) Employing new models Langari et al. (1997)
- (2) Improving the accuracy Tang et al. (2007)
- (3) Expediting convergence rates and reducing computing times Neal and Hinton (1998), Booth and Hobert (1999)

<u>Ann. Statist.</u> Volume 11, Number 1 (1983), 95-103. On the Convergence Properties of the EM Algorithm <u>C. F. Jeff Wu</u>

*J. R. Statist. Soc.* B (1999) **61**, *Part* 1, *pp*. 265–285

#### Maximizing generalized linear mixed model likelihoods with an automated Monte Carlo EM algorithm

James G. Booth and James P. Hol University of Florida, Gainesville, USA A VIEW OF THE EM ALGORITHM THAT JUSTIFIES INCREMENTAL, SPARSE, AND OTHER VARIANTS

![](_page_48_Figure_0.jpeg)

![](_page_49_Figure_0.jpeg)

![](_page_50_Figure_0.jpeg)

#### IV. Results

#### The CRAN landscape

Pkg.	Vrs.	Clust./Class	Density Est.	MB Clust.	Lang.	Parallel
mclust	5.4.6	Y/Y	Y	Y	R	Ν
mixture	1.5	Y/Y	Ν	Y	R	Ν
Rmixmod	2.1.5	Y/Y	Y	Y	R	Ν
flexmix	2.3-15	Y/N	Ν	Y	R	Ν
EMCluster	0.2-12	Y/Y	Ν	Y	R	Ν
mixtools	1.2.0	Y/N	Y	Ν	R	Ν
ClusterR	1.2.2	Y/Y	Ν	Ν	R,C++	Y
DCEM (EM*)	2.0.2	Y/Y	Ν	Ν	R	Ν

#### IV. Results 700 points, 11 features

#### Wisconsin breast cancer data

The first experiment is done with the Wisconsin breast cancer data Wolberg et al. (1995). This data set is small: 699 records and 11 attributes. It is publicly available from the UCI (University of California, Irvine) machine learning repository Dua and Graff (2017) and was created to identify features that indicate tumor classification as either malignant or benign. Originally, the two groups have 444 and 239 data points, respectively. After removing missing records (a minor, insignificant number), EM\* and EM-T algorithms are running over the data size (683 tuples using 9 features).

The results show that EM-T is marginally better for small values of k but, EM\* takes less time and fewer number of iterations, as the number of clusters are increased beyond k = 2 (Figure 2 and Table 3). Results are similar in terms of clustering error and the accuracy of EM\* is in the range 81.82 – 96.77% and for EM-T, it is 78.74 – 95.75%. However, EM\* has slightly higher accuracy than EM-T for true numbers of clusters (k = 2) (Table 3).

Başarim 2020

#### Wisconsin breast cancer data

![](_page_53_Figure_1.jpeg)

(A): Plots showing averaged results of 5 runs for training time, iterations and accuracy.

![](_page_53_Figure_3.jpeg)

(B): Boxplots for training time, number of iterations and accuracy.

#### IV. Results 50K points, 14 features

#### **US** Census data

The second experiment is done on the US census income data Meek et al. (2002). This data set contains approximately 50K points with 14 features (both continuous and discrete) and is available for download from the UCI (University of California, Irvine) machine learning repository Dua and Graff (2017). After removing missing values, the data had 45K points and 6 dimensions (only continuous attributes). The original classification objective was to separate people who earn more than 50K from those who do not.

The results indicate that EM\* quickly outperforms EM-T (Figure 3) as the number of clusters are increased, particularly beyond k = 2 (Table 4). EM-T fails to converge in 1000

#### US Census data

![](_page_55_Figure_1.jpeg)

(A): Plots showing averaged results of 5 runs for training time, iterations and accuracy.

![](_page_55_Figure_3.jpeg)

(B): Boxplots for training time, number of iterations and accuracy.

#### IV. Results 5K points, 57 features

#### Spambase data

In the third experiment, we used the Spambase data set Hopkins et al. (1999). The original purpose of this data was to create a spam filter by identifying the characteristic patterns (words, sequences *etc.*) that can distinguish between spam and non-spam emails. This data set has 4601 records with 57 features and in spite of its moderate size, presents a sharp increase in number of dimensions, scale and type of features, as compared to the earlier experiments. It includes 2788 spam and 1813 non-spam emails and is publicly available from the UCI (University of California, Irvine) machine learning repository Dua and Graff (2017).

#### Spambase data

![](_page_57_Figure_1.jpeg)

(A): Plots showing averaged results of 5 runs for training time, iterations and accuracy.

![](_page_57_Figure_3.jpeg)

(B): Boxplots for training time, number of iterations and accuracy.

#### IV. Results 7.5K points, 20 features

#### **Ringnorm data**

The first experiment was done on the ringnorm data Breiman (1996). This data set was initially used to study the bias and variance properties in arcing classifiers Breiman (1996) and is openly available from the UCI (University of California, Irvine) machine learning repository Dua and Graff (2017). It contains 7400 data points with 20 features. There are two groups that have 3664 and 3736 instances, respectively and both are drawn from multivariate normal distribution.

adaptively, resample, and combine (arc)

Başarim 2020

![](_page_59_Figure_0.jpeg)

(A): Plots showing averaged results of 5 runs for training time, iterations and accuracy.

![](_page_59_Figure_2.jpeg)

(B): Boxplots for training time, number of iterations and accuracy.

Synthetic

#### Number of clusters experiment

![](_page_60_Figure_2.jpeg)

![](_page_60_Figure_3.jpeg)

![](_page_60_Figure_4.jpeg)

EM-T failed to converge for d = 100 in multiple runs.

61

#### IV. Results

- Comparison to other packages in case we did "something special" to our code
- We used large synthetic data sets so no package would have an advantage
- We demonstrate some of the metrics we used to compare EM-T with EM\*

#### Experiment on number of clusters

![](_page_62_Figure_1.jpeg)

Comparing DCEM with mixtools and EMCluster on data with large number of clusters. Plot for execution time show that DCEM is significantly faster than EMCluster and mixtools. In terms of iterations, DCEM requires fewer iterations than EMCluster, iterations for mixtools are not reported as it could not converge in the threshold of 2 hours.

Başarim 2020

#### Scalability experiment

![](_page_63_Figure_1.jpeg)

Comparing DCEM with mixtools and EMCluster on large data. DCEM always converges (maximum execution time of ~33 minutes for 2 million points) in all cases whereas EMCluster and mixtools failed to converge within the time limit of 2 hours. Iterations are not reported for mixtools and EMCluster because they failed to converge for all values of *n*.

Başarim 2020

IV. HPC in small packages: Viability of Small Machines facing Big Data, a study in leveraging high-expressive data helps improve iterative algorithms.

#### **Computational details**

All experiments were done on a dedicated SICE (School of Informatics, Computing and Engineering) VM instance - Schmuck. Schmuck is a 64-bit system running the bionic-18.04 release of the Ubuntu desktop. It has two intel Xeon cores (2.70GHz) with 16 GB of main memory and 8GB of swap partition. To ensure precision of results, the system was maintained in ideal conditions i.e., no other resource intensive jobs were running during the experiments. DCEM is implemented and tested in R 3.6.3 (2020-02-29).

![](_page_64_Figure_4.jpeg)

## V. Future Work

- Implement parallel/distributed heap-based optimization
- Examine other structures (maybe total order is better than partial order)
- Add additional functionality to EM\* (Specifically kmeans\*, other iterative algorithms
- Maybe (maybe) make the Python upgradable to 3.8 instead of 2.7 (we have a good number of downloads)

## Questions?

![](_page_66_Picture_2.jpeg)