

## A Novel Parallel Implementation of the CARE Algorithm to Demonstrate the Performance Benefits on Big Patient Data

Sridhar Gujjeti<sup>1</sup> & Dr. Suresh Pabboju<sup>2</sup>

<sup>1</sup>Assistant Professor, Department of CSE, Kakatiya Institute of Technology and Science, Warangal

<sup>2</sup>Professor, Department of Information Technology, Chaitanya Bharathi Institute of Technology (CBIT), Hyderabad.

Received: April 12, 2018

Accepted: May 21, 2018

### ABSTRACT

*Today the healthcare industry is experiencing one of the most imperative and testing transitions to date, the move from paper to electronic healthcare records. While the healthcare industry has for the most part been an incrementally propelling field, this change can possibly be revolutionarily. Utilizing the data gathered from these electronic records energizing apparatuses, for example, disease recommendation systems have been made to convey personalized models of a person's health profile. Anyway in spite of their initial achievement, apparatuses, for example, these will soon experience a huge issue. The measure of healthcare experience data gathered is expanding radically, and the computational time for these applications will soon achieve a time when these systems can never again function in a viable timeframe for clinical use. This paper will start by dissecting the performance limitations of the personalized disease prediction engine CARE (Collaborative Assessment and Recommendation Engine). Next it will detail the creation and performance of another single patient implementation of the algorithm. At last this work will demonstrate a novel parallel implementation of the CARE algorithm, and demonstrate the performance benefits on big patient data.*

**Keywords:** CARE, personalized healthcare, Collaborative Assessment and Recommendation Engine, big data, CARE

### I. Introduction

Medical research has occurred for decades. It has given what we as a general public feel are a portion of the best present day accomplishments, from the disclosure of microscopic organisms and viruses to the advancement of antibiotics. Today, as the healthcare industry starts its transition into the advanced age it is anything but difficult to see the occasion as an insignificant transitioning, essentially the transformation of the medical network's paper documentation into electronic form. Be that as it may, it gives a great deal more. This transition has established the framework for another essential headway in the field of healthcare, the evolution from safeguard care into personalized treatment designs.

It has been very much archived that early detection and treatment of numerous diseases is specifically associated with enhanced health results for the patient (Etzioni et al., 2003) (Wilkinson, Donaldson, Hurst, Seemungal, and Wedzicha, 2004) (Lard et al., 2001). Accordingly, consistent supposed "health visit programs" have been actualized by numerous organizations and care suppliers keeping in mind the end goal to advance preemptive testing for specific conditions (Kickbusch and Payne, 2003; Ozminkowski et al., 2002). Notwithstanding, as the identification and treatment of these diseases are performed in a similar way for numerous people construct fundamentally with respect to their flow health state, i.e. age, sexual orientation, race, earlier lab comes about, and so forth this sort of care falls nearer to protection medication than personalized care.

Where as earlier investigations have guided safeguard solution treatment systems by giving verifiable probabilistic models in light of the results of patients who created comparable conditions, new prescient techniques can help make personalized models of a patients future health dangers custom fitted to the person's health profile. Keeping in mind the end goal to make these personalized models, data mining techniques have been connected to population-level health data amassed from electronic healthcare records (EMR).

While traditional data mining procedure, for example, grouping, decision trees and associate investigation created empowering comes about, there was unfortunately an issue (Jensen, Jensen, and Brunak, 2012; Bellazzi and Zupan, 2008). Similarly as with paper records, each additional medical experience by a patient brought about additional data added to their electronic health record, and the amount of data soon surpassed the capacity of standard data preparing techniques. In response, new data preparing techniques and algorithms are being made, for example, Google's MapReduce, Yahoo's Hadoop, and so forth. (Mayer-Schönberger and Cukier, 2013). These techniques use the concepts of errand segmentation and circulated figuring so as to mitigate a portion of the computational load from a solitary machine, and take into account fundamentally enhanced runtimes for parallelizable tasks.

Because of the time basic nature of medical conditions, the utility of any model made is straightforwardly proportional to the time taken to make it. In that capacity we should center around preparing time of a model request to permit personalized healthcare models to be made inside a useful timeframe.

This paper will center around the essential issue of CARE's convenience in a clinical setting. It will start by recognizing the limitations of the present CARE algorithm, and mean to give an arrangement of ideal performance parameters. Next a portion of the precision limitations will be tended to through the creation of a solitary patient version of CARE. At last this work will demonstrate a novel circulated processing implementation of the CARE algorithm. This implementation will address issues with both execution time and disease scope while endeavoring to give close industry level performance.

## II. Current CARE Architecture

The current CARE architecture can be seen in Figure 1 and is fairly straightforward. The basic steps for the algorithm are detailed below.

1. CARE begins with an individual presenting a set of diseases. This set is the accumulation of diseases over their personal medical history.
2. The individual's disease similarity is then compared to all other patients in the provider's existing database and an initial filtering is done.
3. Collaborative filtering is then performed on this filtered dataset.
4. Finally a probabilistic ranking of diseases for the individual is returned.

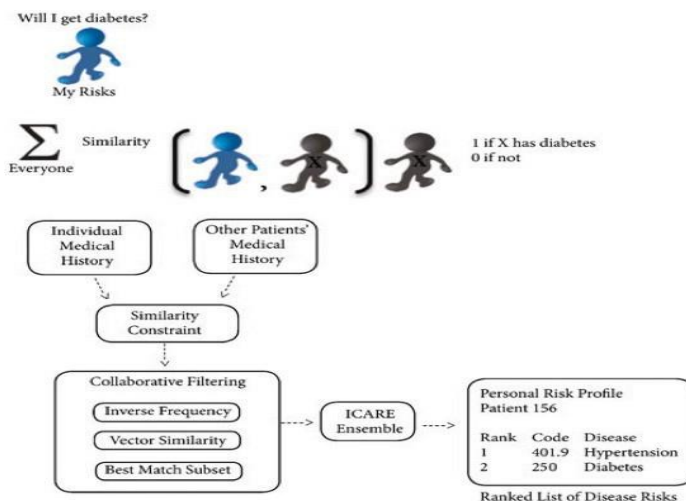


Figure 1: Current CARE Implementation (Chawla & Davis, 2013).

### Data

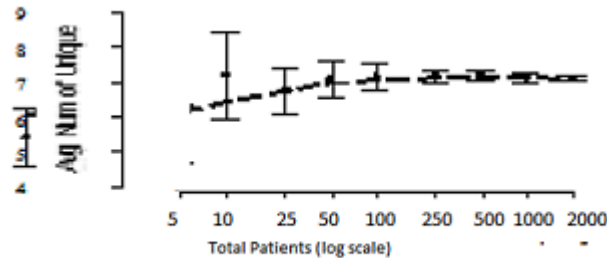
The data used in this investigation is the same dataset used inside the CARE research. The data consists of anonymized Medicare claims records gathered by the Harvard University Medical School. There are around 13 million individual patients, representing a little more than 32 million healing facility visits, and contains a sum of 18 thousand special disease codes. Each record speaks to a solitary doctor's facility visit and is contained a patient ID number and up to 10 singular findings from the visit. The finding codes are characterized by the International Classification of Diseases, Ninth Revision, Clinical Modification (ICD-9-CM), distributed by the World Health Organization (WHO) (Slee, 1978). Through the ICD-9-CM code every disease is given a one of a kind code, which can be up to 5 characters long. These codes may incorporate specifics of the condition, for example, the anatomical location. In any case, these fine-grained subtle elements are not required for the CARE algorithm, and thus the 5 digit finding codes can be fell to a 3-digit generalization of the diseases. For instance codes 461.0 and 461.1 can be fell into the nonexclusive conclusion code 461. The accuracy of this all inclusive statement is archived inside the CARE paper, and in that capacity will be used going ahead in this work also (Davis et al., 2010).

Note that a disease might be analyzed to an individual different circumstances all through their medical history. Be that as it may, as different diseases are not useful when looking at patient's disease

sets, only one of a kind diseases are required for recommendations. Figure 2 demonstrates that the normal number of one of a kind diseases converges to around 7 for each patient over the full dataset. This esteem will be used while distinguishing anomalies from the randomly chose patients, lessening the inclination amongst datasets and execution time.

#### Sequential CARE Implementation

As stated prior the current implementation of CARE is executed in a sequential manner. Our evaluation began with a detailed benchmarking on the existing CARE architecture. The goal of these benchmarks was to identify areas of computational resource restrictions, as well as any areas of algorithmic complexity that could not be solved with improved hardware.



**Figure 2: Average Number of Unique Diseases Per Patient**

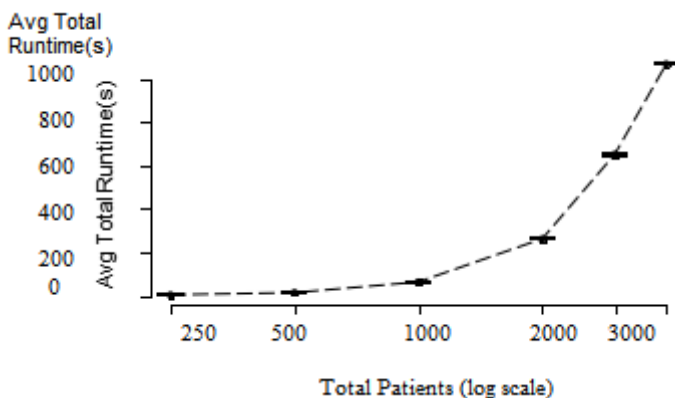
CARE was designed to operate within a clinical healthcare environment, and as this environment operates in a time-critical manner the primary benchmark metric used was execution time. Further a secondary metric for evaluation is the total number of patients used for training, as a proxy for recommendation accuracy. This proxy is a result of the sparse nature of disease networks. A large sampling of patients is required in order to ensure sufficient similarity comparisons for the disease ranking calculations.

#### Evaluation Environment

To perform the evaluation of CARE the Opteron machine at the Notre Dame Center for Re-look Computing (CRC) was used. This machine contains 4, 16-center 2.3 GHz AMD Opteron processors with 128 GB RAM. As this machine is a period shared framework it is probably going to encounter performance fluxuations, because of user process stack. With a specific end goal to represent this all simulations were run 5 times, and the outcomes arrived at the midpoint of to get steady performance gauges. Additionally, preceding every benchmark CARE was run once in an endeavor warm the reserve for the new arrangement of diseases and visits. Further to guarantee that the performance benchmarks were precise and repeatable CARE was agreed utilizing the -OO banner to cripple a particular complier optimizations. This was chosen as the parallel CARE implementation was kept running on the CRC Sun Grid Engine (SGE). Given that CARE has no particular framework necessities, SGE specialists could hypothetically be dispatched to any machine with an extra center. In that capacity, optimizations were evacuated to guarantee uniform execution designs over all specialist machines. At long last all unessential framework calls, expanding and logging were altogether expelled from the CARE source code. This was done in an effort to settle the code from any branch prediction inside processors, to guarantee as consistent a pipeline as conceivable between each consequent benchmark.

#### Execution Mode

In the course of this evaluation it became clear that CARE was constructed with some specific design considerations. The current implementation CARE effectively performs an all-pairs computation for diseases similarity against every patient within the database. This design has since been designated Batch Mode, as this is a process that would likely be executed as an overnight batch job. In contrast this work details a new version of CARE, which performs on-demand recommendations for a single patient. This implementation is similarly designed Individual Mode, and is detailed below. The following sections will detail the performance evaluations of each recommendation mode.



**Figure 3: Average Total Runtime of Unmodified CARE**

Total	I/O Execution	CPU Computation
5	13.27	24.95
10	8.49	61.31
25	1.73	88.55
50	0.73	90.96
100	0.35	92.64
250	0.15	93.00

**Table 1: I/O Vs. CPU Breakdown as a Percent of Total Execution Time**

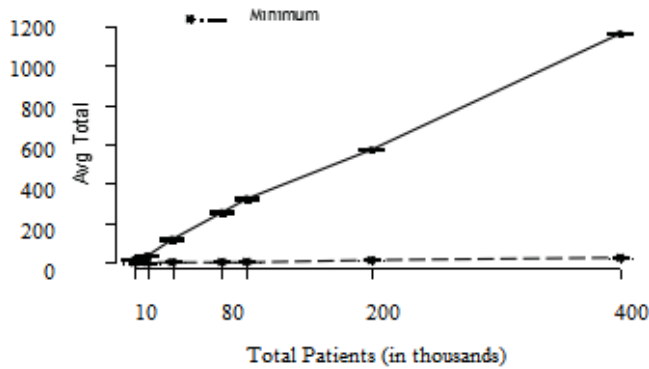
### III. Results and Analysis

The underlying performance examination was performed on an unmodified version of the CARE algorithm, and base runtimes were set up finished a shifting number of patients (Figure 3). Further as the objective of this work was to demonstrate the implementation of a scalable disease prediction algorithm, which could use big data from the biggest of supplier arranges inside an industry practical timeframe, it was vital to first figure out which computational territory would profit most from optimization. Keeping in mind the end goal to accomplish this the major inner CARE functions were partitioned into the two crucial computational classes, I/O and CPU bound, and their percent of aggregate computation time was figured (Table 1). Note the rest of the level of execution time for littler patient datasets was consumed by standard framework calls and data structure initialization. As should be obvious these calls are not a bouncing component as they lessen to under six percent with a dataset of only 250 patients, well underneath even the littlest supplier systems. The essential result of these underlying evaluations was the confirmation that the CARE algorithm is exceptionally CPU bound. Further, the program was CPU bound, as well as more terrible the all-sets computation for disease rankings brought about an exponential runtime as a factor of aggregate patients. Indeed as should be obvious from the comparison between Figure 3 and Table 1, the percent of execution time from I/O ends up negligible before the runtime even starts it's fast increment. This further reinforces the need to center around CPU bound functionality.

The standard CARE algorithm figures the disease rankings for all patients contained in a four thousand person dataset in a little more than 17 minutes. At only four thousand patients this dataset is a reasonable simple for even the littlest of medical practices. As expressed before patient considers can be used an intermediary for CARE precision, and a dataset this little is a huge concern. Additionally these base implementation comes about feature another real issue. Studies have demonstrated that a whole clinical experience, including recording a patient's medical history and vitals endures around only 17 minutes (Mechanic, McAlpine, and Rosenthal, 2001). At a runtime of coordinating this figure, these outcomes viably wipe out the use of Batch Mode as an achievable strategy for intra-day and on-demand disease rankings.

After modifying the CARE algorithm to perform on-demand computation, the logical first step was to compare the execution time to that of Batch Mode (Figure 4). By removing the all-pairs comparison, CARE is able to handle substantially larger datasets. It was able to process just under 400 thousand

patients in the equivalent timespan of the four thousand patient dataset, at 17 minutes. Again, as the number of patients in the provider database can be used as a proxy for ranking the Individual Mode would provide disease rankings at a significantly higher confidence level. It is important to note that the Batch Mode would create disease rankings for 3,999 more patients than would the Individual Mode implementation. However when diagnosing a patient, accuracy for that individual may be more important than would be the diagnosis for multiple potentially unrelated patients at once. In fact, if Individual Mode CARE were initialized at the start of a patients primary care visit, their diseases rankings could be calculated against a dataset of up to 300 thousand patients.



**Figure 4: Average Total Runtime of Unmodified CARE**

In these situations CARE could be instated at the season of the patients admission. CARE would then have the capacity to deliver a preparatory set diseases chance inside the early minutes of treatment, a basic timeframe for high-chance patients. Breaking down the result from Individual Mode yields a fascinating outcome as identified with the general CARE algorithm. The outcomes demonstrate that execution time isn't autonomously connected with the measure of the patient dataset used. Rather, the execution time is additionally specifically associated to the quantity of diseases per patient. This comes because of the preprocessing step taken via CARE, noted prior, where only those patients that offer a common disease are used for the collaborative sifting algorithm. The aftereffects of this implementation decision are appeared by Figure 4 where the patient with the most extreme number of diseases has a fundamentally higher runtime, as does the patient with the base number of diseases when kept running on the same dataset.

**IV. Parallel CARE Implementation**

While crafted by making a solitary user implementation of CARE demonstrates that it is conceivable to enhance the execution time, the computational prerequisites of collaborative sifting limit the maximal performance picks up that can be accomplished by the present CARE engineering. In the wake of assessing both the present implementation and enhanced single patient CARE algorithms, it turned out to be evident that the major CARE design would should be changed to acquire any further performance upgrades. Keeping in mind the end goal to understand how the CARE algorithm could profit by optimizations, for example, parallel execution, it was first essential to understand where inside cared slow down.

This paper has beforehand demonstrated that CARE is CPU bound, however now going further it is imperative to characterize where this happens. Keeping in mind the end goal to answer this the CPU bound components were separated into the individual functions as a percent of aggregate runtime (Table 2). Unmistakably the CPU bouncing is commanded by one function, Best Match.

Total	Best Match	LoadPatient	Load Disease
5	65.28	15.86	18.86
10	87.84	6.69	5.47
25	98.08	1.36	0.55
50	99.20	0.64	0.15
100	99.62	0.33	0.04
250	99.88	0.15	0.01

Table 2: Function Breakdown of CARE Execution

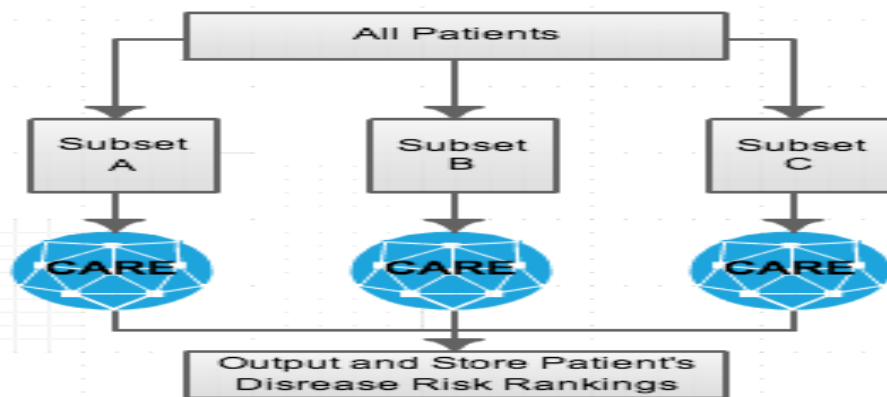
Total	Percent of Time Per Function		
	VectorSimilarity	Merge Visits	System Calls
25	41.78	10.32	47.90
50	41.69	10.05	48.26
100	42.24	8.22	49.54
250	42.13	7.52	50.36
Average	42.51	7.85	49.64
SD	1.30	2.19	1.35

**Table 3: Component Breakdown of Best Match Function**

Taking this further the Best Match function was separated to examine precisely what was causing the bottleneck (Table 3). Note that because of the short execution time and inadequate nature of the disease arrange, investigating datasets containing under 10 patients makes exceedingly factor and non-convergent outcomes. Along these lines all datasets underneath 25 patients were barred from this evaluation. Taking a gander at the table it is fascinating to see that the percent of time spent in every component of the function stays unaltered as a result of number of patients in the dataset. This outcome loans itself well to the potential benefits of parallelization as it demonstrates that despite the fact that the algorithm has an exponential runtime the measure of time spent in each function is steady.

**Distributed Architecture**

The distributed model for the CARE algorithm is an aftereffect of the current CARE system, where every patient's disease positioning is figured freely. While all patient's medical narratives are used as preparing for the model, making the disease organize for a patient is a free occasion. Therefore, the distributed version of CARE partitions the arrangement of all patients into measure up to estimate subsets. These subsets are then distributed to up to 50 singular worker nodes.



**Figure 5: Distributed CARE Architecture**

**Evaluation Metric Extension**

For the investigation of Distributed CARE model this paper will use an additional performance measurements assigned Derived Time. Because of the time-shared nature of tasks on the CRC machines a precise time for extensive scale tasks can't be estimated. Notwithstanding, exact execution times can be estimated for each errand. Determined Time at that point is an extension of standard time and is figured accepting each errand will start as the assignment in front of it closes. Each piece of tasks will execute in parallel and the aggregate Derived Time will be a consequence of the normal I/O overhead added to the normal computation time for an errand increased by the quantity of worker cycles that must be finished for the whole dataset. This strategy is used for all datasets more than 104 patients.

## Analysis and Results

By utilizing the exponential decay aspect of the Batch Mode implementation it is clear that the execution time of CARE can be significantly improved. Additionally this improvement does not suffer from the trade off of reduced numbers of patient disease risk calculations, as does the current Individual Mode improvement.

Task Count	Avg IO Overhead (s)	Avg CPU Time (s)	Num Worker	Derived Execution
10000	0.017	134.358	200	7.465
25000	0.021	43.811	500	6.088
50000	0.098	19.412	1000	5.419
100000	0.101	6.658	2000	3.755

**Table 4: Avg Total Runtime per Task of Distributed CARE**

Task Count	1000	10000	25000	50000
Patients Per Task	50	5	2	1
No Cache	43618.23	43609.78	46276.27	57924.65
Cache First Access	58597.37	50521.59	57008.15	54396.94
Cache Subsequent Access	29.92	32.11	29.89	31.43

**Table 5: Avg IO Overhead for 50000 Patient Dataset in  $\mu$ s**

As you can see Figure 6 demonstrates the significant improvements of Distributed CARE. It is important to note that Distributed CARE still exhibits an exponential execution time as a function of total patients in the dataset. This is expected, as Distributed CARE makes no effort to change the internal framework of CARE. The goal was to process significantly larger patient datasets in a practical timeframe, both of which Distributed CARE achieves. While another exponential program may not appear like a lot of a change the accompanying demonstrates that it isn't only a change, yet in addition essential for boundless adoption of CARE. Utilizing the standard Batch Mode execution times above we could fit the exponential function  $60.3882e0.00073x$  to model the anticipated runtime of CARE. At this estimation the computation for 100000 patients, a number effectively achievable by significant practices, would take  $7.04e25$  years. This figure viably disposes of CARE's utilization on a training wide premise, forcing it to be used only in specific cases. This figure adequately takes out CARE's utilization on a training wide premise, forcing it to be used only in particular cases. As should be obvious through Table 4 expanding the assignment check, and along these lines diminishing the patients per worker errand drastically diminishes the computation time down to just shy of 4 hours. With runtimes at only 1/6 of a day, essentially more patients could profit by the utilization of CARE.

Anyway so as to completely use the Distributed CARE algorithm, you should look further and consider optimizations inside each run. This work chose to use the same number of workers as permitted, for this situation 50, under the assumption that the any medical work on utilizing CARE will have had more than 50 patients go through their framework, taking into account no less than at least one patient for each worker. All things considered, all optimizations will come through variations in the distributed subset estimate. While breaking down a distributed framework for performance optimizations it winds up imperative to look not just at the execution time for an undertaking, yet in addition at the time required for communication overhead between the ace and worker nodes. Once more, as should be obvious from Table 4 the greater part of an errand's execution time is spent in CPU bound calculations. It ought to be noticed that while computation time exponentially diminishes with the undertaking check, I/O overhead increments. This demonstrates in principle for extensive datasets if the errand tally is too high I/O overhead may really expand the aggregate execution time of Distributed CARE. Unfortunately it was unrealistic to achieve that cutoff notwithstanding while using all patients with more than 5 diseases. Nonetheless, it ought to be considered that medical practices with datasets containing millions of patients may accomplish imperfect performance by endeavoring to partition the dataset too firmly.

As an additional note, distributed systems today are ending up essentially more common, and accordingly numerous controller systems are executing apparatuses to enhance performance of these distributed systems. For instance WorkQueue algorithm can store enter records for distribution to worker nodes. As should be obvious in Table 5 empowering reserving adequately takes out all overhead from the ace,

lessening the requirement for dataset particular optimizations. While there might lessen returns on diminished execution time, a training may use the same number of computational assets as accessible and be confident that they are working at ideal performance levels.

## V. Conclusion

To recap, this paper has demonstrated the performance limitations of the present CARE algorithm. While some claim that an overnight batch execution is adequate, as it can process a substantial patient dataset with a high level of precision, this strategy is non-feasible for medical utilization. Big data suppliers, for example, Facebook do use comparable clump occasions to help with data preparing, yet the information produced does not have the wellbeing basic nature of healthcare data. If a disease is inaccurately recorded, a patient may need to hold up to 24 hours to get refreshed disease dangers. This turnaround time might be inadmissible, particularly for time basic units. Keeping in mind the end goal to understand the issue of computation time this paper has plot two unmistakable techniques. Initial a solitary patient version of CARE, which can be used to perform disease chance rankings on demand with a genuinely high level of exactness. This strategy is expected to be used for the situation above where refreshed rankings must be recovered because of blunder, or for another patient who was absent in the database when the last cluster work was run. The second strategy is a distributed computation of the CARE algorithm. This implementation can be used to produce on-demand rankings for a solitary patient with a high level of precision, or executed as a daily clump work on fundamentally bigger patient sets for vast practices or clinics.

## References

- AbuKhoua, E., and Campbell, P. (2012). Prescient data mining to help clinical decisions: A diagram of coronary illness prediction systems. In *Innovations in information innovation (iit), 2012 international conference on* (pp. 267- 272).
- Austin, P. C., Tu, J. V., Ho, J. E., Levy, D., and Lee, D. S. (2013). Utilizing strategies from the data-mining and machine-learning writing for disease classification and prediction: a contextual analysis examining classification of heart disappointment subtypes. *Diary of clinical the study of disease transmission*.
- Bellazzi, R., and Zupan, B. (2008). Prescient data mining in clinical pharmaceutical: current issues and rules. *international diary of medical informatics*, 77 (2), 81- 97.
- Berkovsky, S., Eytani, Y., Kuflik, T., and Ricci, F. (2007). Improving security and protecting precision of a distributed collaborative sifting. In *Proceedings of the 2007 acm conference on recommender systems* (pp. 9- 16).
- Bui, P., Rajan, D., Abdul-Wahid, B., Izaguirre, J., and Thain, D. (2011). Work queue+ python: A structure for scalable logical troupe applications. In *Workshop on python for superior and logical processing at sc11*.
- Chawla, N. V., and Davis, D. A. (2013). Bringing big data to personalized healthcare: A patient-focused system. *Diary of general interior medication*, 28 (3), 660- 665.
- Davis, D. A., Chawla, N. V., Christakis, N. An., and Barab'asi, A.- L. (2010). Time to care: a col-laborative engine for reasonable disease prediction. *Data Mining and Knowledge Discovery*, 20 (3), 388- 415.
- Dignitary, J., and Ghemawat, S. (2008). Mapreduce: disentangled data handling on expansive groups. *Communications of the ACM*, 51 (1), 107- 113.
- Etzioni, R., Urban, N., Ramsey, S., McIntosh, M., Schwartz, S., Reid, B., . . . Hartwell, L. (2003). The case for early detection. *Nature Reviews Cancer*, 3 (4), 243- 252.
- Goil, S., and Choudhary, A. (1997). Superior olap and data mining on parallel comput-ers. *Data Mining and Knowledge Discovery*, 1 (4), 391- 417.
- Jensen, P. B., Jensen, L. J., and Brunak, S. (2012). Mining electronic health records: towards better research applications and clinical care. *Nature Reviews Genetics*, 13 (6), 395- 405.
- Kickbusch, I., and Payne, L. (2003). Twenty-first century health promotion: the general health revolution meets the wellbeing revolution. *Health promotion international*, 18 (4), 275- 278.
- Grease, L. R., Visser, H., Speyer, I., vander Horst-Bruinsma, I. E., Zwinderman, A. H., Breedveld, F. C., and Hazes, J. M. (2001). Early versus postponed treatment in patients with later onset rheumatoid joint pain: comparison of two associates who got diverse treatment procedures. *The American diary of medication*, 111 (6), 446- 451.
- Lathia, N., Hales, S., and Capra, L. (2007). Private distributed collaborative separating utilizing esti-mated concordance measures. In *Proceedings of the 2007 acm conference on recommender systems* (pp. 1- 8).
- Linden, G., Smith, B., and York, J. (2003). Amazon. com recommendations: Item-to-thing collaborative sifting. *Web Computing, IEEE*, 7 (1), 76- 80.
- Mayer-Sch'onberger, V., and Cukier, K. (2013). *Big data: A revolution that will transform how we live, work, and think*. Houghton Mifflin Harcourt.

**McCormick, T. H., Rudin, C., and Madigan, D. (2012). Bayesian progressive govern modeling for foreseeing medical conditions. *The Annals of Applied Statistics*, 6 (2), 652- 668.**