Evolved Bayesian Network Models of Rig Operations in the Gulf of Mexico

François A. Fournier, John McCall, Andrei Petrovski, Peter J. Barclay

Abstract— **The operation of drilling rigs is highly expensive. It is therefore important to be able to identify and analyse factors affecting rig operations. We investigate the use of two Genetic Algorithms, K2GA and ChainGA, to induce a Bayesian Network model for the real world problem of Rig Operations Management. We sample from a unique dataset derived from the commercial market intelligence databases assembled by ODS-Petrodata Ltd. We observe a trade-off between K2GA, which finds significantly better scoring networks on our dataset, and ChainGA, which uses only one quarter of the computation time. We analyse the best structures produced from an industry standpoint and conclude by outlining a few potential applications of the models to support rig operations.**

I. Introduction

THE oil and gas sector is an active industry constantly seeking to research and apply new technologies. However, the competitiveness of the sector often leads to high levels of secrecy concerning technology, operations, production and support methods and data. Drilling rigs are operated by contractors who hire out their services to oil companies for both exploration and exploitation. The operation of drilling rigs is highly expensive. Typically a rig operating offshore in the Gulf of Mexico can cost from \$400K to \$600K per day.[1] With rig operations lasting weeks or even months at a time, variations in the efficiency with which rigs are operated can affect profitability by millions of dollars. It is therefore important to be able to identify and analyse factors affecting efficiency.

There are many ways of defining efficiency. Drilling Rig efficiency is usually assessed by industry experts on the basis of practical experience [2] but there is currently no industry-standard approach for the objective measurement and prediction of efficiency. Efficiency on its own cannot be directly compared between rigs without considering many influencing factors such as, weather, the specific nature of the geological layers being drilled through, and other environmental or managerial factors. The selection of which factors are relevant and how they are related is largely left to the judgment of managers and other experts in the field. Their approach is mainly based on empirical observations and experience. In some cases, the rig selected for a job will be over-specified or under-specified leading either to unnecessary expense or poor outcomes such as significant delay. It is this uncertainty surrounding the rig selection process which motivates rig operations management as an application area for data modelling.

Manuscript received May 02, 2010. This work was supported by ODS-Petrodata Ltd. (www.ods-petrodata.com) and the Technology Strategy Board under the KTP scheme (Award: KTP006922).

F.A. Fournier, J. McCall and A. Petrovski are with the IDEAS Research Institute at Robert Gordon University, Aberdeen, Scotland. (email: {f.a.j.fournier,j.mccall,a.petrovski}@rgu.ac.uk)

F.A. Fournier and P.J. Barclay are with ODS-Petrodata Ltd., Aberdeen, Scotland. (e-mail: ffournier@ods-petrodata.com).

In this paper, we are interested in the use of evolutionary algorithms to evolve Bayesian Networks relating a range of factors selected from an extensive oil and gas industry dataset. We aim to explore the utility of the models produced. Also we provide a comparison of two published algorithms on a new real world problem.

In the next section, we provide a more detailed account of drilling rig operations and the rig tendering process. We also describe the dataset used and our approach to factor selection. In section 3, we summarise the Bayesian Network modelling approach and describe the two search and score evolutionary algorithms used to build these networks from the data. In Section 4, we describe our experiments with Rig and Wells data. The experimental results are analysed in Section 5 and a discussion of possible applications appears in section 6. The final section contains conclusions and a brief outline of planned future work.

II. Rig Operations and The Gulf of Mexico

A. Offshore Drilling

The offshore drilling process is mainly split into two main steps: exploration and exploitation. Various offshore drilling platform types exist within those two categories. Table 1, derived from Nergaard [3] summarises the different types of offshore drilling platforms available.

Table 1: Offshore Drilling platform types

Exploration	Floaters	Semi-Submersible
		Ships
	Bottom Support	Jack-ups
Production / Exploitation	Surface platforms	Permanent
		Tenders
	Subsea	Semi-Submersible
		Ships
		Jack-ups

Rig owners contract rigs to drilling companies for specific pre-established needs in both exploration and production. The offshore drilling market is dynamic, highly competitive, and regionally-specific. Key differences across regions are legislative and geological variations, however, cultural differences and practices across regions and across companies often also impact on rig results.

Freudenrich [4] explains a simplified path to oil and gas production. Oil is located using various survey methods and tools including geological analysis, gravity meters, magnetometers and seismology technologies. Once a site is selected, it is surveyed to find its boundaries. Then a drilling rig is brought on site and starts drilling. As drilling progresses, mud circulates through the pipe and out of the drill bit to float the rock cuttings out of the hole. When a pre-set depth is reached, the drill bits are removed from the hole and a steel-and-cement casing is installed. When reaching the final depth, various logs and tests are performed and samples are taken for analysis. The well is then secured and installed in order to let the oil flow in a controlled manner. Once the oil is flowing, the oil rig is removed from the site and production equipment is set up to extract the oil from the well. [4]

Regarding performance, Harris, in a 1989 publication [2] explains that no two wells perform the same but that consistently good results are a good indication of a rigs capability. He highlights three main criteria, currently used to select rigs: technical suitability, price, and availability. Osmundsen *et al*. [5], in a more recent publication, highlight more evaluation criteria for selection. In no particular order, he states that typical evaluation criteria can be: expertise, financial strength, day rates, ability to complete on time, compliance with regulations, operational efficiency and achievements, Health and Safety Executive system and culture, high pressure, high temperature (HPHT) expertise and experience. [5]

B. *The Rig Tendering Process*

Rig tendering is the process by which a company contracts a rig for a given operation. According to Harris [2], a successful operation depends on many factors which are difficult to measure. The tendering process for selecting a rig has remained largely unchanged since his publication in 1989.

When selecting a rig for a drilling programme an operator typically has three main criteria: technical suitability, price, and availability. Some technical parameters are absolute and determine the type of rig and equipment. Examples are water depth, pressure and temperature ratings, etc. However, alternatives can sometimes be suitable: semi-submersibles have been seen to operate in jack-up water depth [2]. Many of the other technical requirements included in an invitation to tender are often preferences rather than necessities. It is commonly recognised that, if the well is drilled efficiently, a higher priced bid can lead to a lower overall cost. Also, a low priced bid can become expensive if accidents extend the drilling time. [2] Considering availability, requirements will tend to be stricter in a low-demand rig market compared to when rigs are in short supply. However, the market maintains a system of “extension options” which is one of the main sources of uncertainty on rig availability [2]. Another potential measure is the safety ratings as there is a correlation between a good operation and a good safety record.

The usual process starts by a company in search of a contractor sending out an invitation to tender. The contractor will then respond to the invitation, presenting various options available, depending on the nature of a potential non-compliance. Within all the responses will appear some variation in potential and decisional tradeoffs open to judgement.[2] In recent years, a move toward the search of quality has been made and bidders in Europe are often asked to provide percentage downtime and indicators of drilling efficiency for the past six wells including water depth, mooring time, loss of time, repair time. [5] However this is not often available information in most regions across the globe.

C. *The Gulf of Mexico Rigs and Wells Dataset*

This dataset is covering Rigs and Wells data sourced by ODS-Petrodata Ltd [6] within its market intelligence commercial databases. ODS-Petrodata's RigPoint[1] database covers worldwide offshore drilling rig contract and activity. Currently, it covers over 25 years of historical rig activity. One recent addition to their databases coverage is Wells data. This extension covers both historical and current drilling activities within the offshore industry. We selected data representing Rigs operations and Wells data in Gulf of Mexico. Historical and current data are collected in several tables of ODS-Petrodata’s RigPoint [6] and Wells databases. We compiled those data in a single flat table of related data fields. There are potentially between 700 and 1000 data-fields (or factors) per record in this table ranging from operational data (water depth, footage drilled, operation dates and durations, etc) to technical data (cantilever capacity, water depth rating, age, etc). Not all fields have been included in the experiment because of availability, completeness and accuracy of the data. Overall we identified 37 available factors of particular interest and then reduced it to 17 key factors to reduce the computation load, by removing the fields with insufficient data coverage and the ones with a large number of distinct values. When available, the fields selected covered the information considered in [5] as necessary to estimate rig efficiency. To give an idea of the scope, one of the variables has 72 distinct nominal values while another one had 350. Factors are either taken directly from particular data fields or derived from them when not usable in a meaningful way directly (for example, start and end dates are transformed into durations). All the numerical fields (for example water depth or footage drilled) are discretised in industry-meaningful categories, established using industry expertise, such as rig operating categories or usual operating ranges of particular equipment. The other fields have been left unprocessed and directly copied over to our dataset. The outcome is a dataset consisting of 6670 rows containing related values of 17 factors.

Table 2 shows the variables selected and the number of values they can take.

Table 2: Selected Fields, index and variable count

Field Name	Number of values
Well Phase	6
Well Deviated	4
Well Type	6
Well Status	7
Well Result	17
Days On Location Discretised	11
Number of days to total Depth Discretised	10
Total Vertical Depth Discretised	18
Total Footage Drilled Discretised	18
Average Feet drilled Per Day Discretised	16
Shore Base	54
Region	59
Water Depth Discretised	10
Rig Type	6
Harsh Environment Capability	2
Rig Owner	72
Rig Contractor	70

III. BAYESIAN NETWORKS

Various attempts to use Bayesian Networks applied to the Oil and Gas industry have been made over the years. Some examples are petrophysical decision support [7] [8], safety instrumentation and risk reduction [9].

A. *Bayesian Network theories*

Bayesian Networks are probabilistic models based on Bayesian Inference [10]. They are useful for representing knowledge under uncertainty. They can be represented using a Directed Acyclic Graph associated with a joint probability distribution [11]. Each node of the graph represents a random variable X_i related to a problem domain. Conditional dependencies between variables are represented by edges in the graph and the joint probability distribution can be factorised according to these conditional dependencies. Formally, the joint probability distribution P(X) over the set of random variables $X_1,\ldots,X_n$, given $Pa(X_i)$ as the set of parent nodes for node X_i, is represented by

$$P(X_1, X_2, \ldots, X_n) = \prod_{i=1}^{n} P(X_i \mid Pa(X_i)) \quad (1)$$

To make use of the power of Bayesian Networks in knowledge representation and inference, the network has to be constructed for the given problem. The underlying Directed Acyclic Graph structure representing the network has to be learned and then the conditional probabilities calculated. Learning the underlying structure is a hard problem [12] because the number of possible structures grows super-exponentially with the number of variables [13]. One widely used approach to this problem is search and score. A metaheuristic is used to search a space representing possible networks. Each solution is scored according to how well it represents the observed distribution of the data. Various authors have presented a variety of metaheuristic approaches to this task including Genetic Programming [14] and Genetic Algorithms [15][16][17][18]. Other approaches include hill-climbing methods [19] and Simulated Annealing [20].

B. *The K2 algorithm*

The K2 algorithm was proposed by Cooper and Herskovitz [21]. K2 assumes that a priori all structures are equally likely and that cases in the data occur independently and are complete. Moreover, it assumes the presence of a node ordering and imposes a maximum number of parents for each node (inbound edges). When these conditions are satisfied, K2 starts with an empty ancestor set for each node and incrementally adds links that maximize the score of the resulting structure. The algorithm stops when no more ancestor node additions improve the score. As in [11], we observe that although simple to implement and widely used, K2 is prone to local optima and may not find the globally best structure. Moreover, it relies on prior knowledge of the node ordering and so may return non-equivalent structures given different orderings.

C. *K2GA and ChainGA*

One popular search and score approach is to search the smaller space of variable node orderings using a metaheuristic and use a greedy algorithm to build solutions from each ordering. These solutions are then scored and the result passed back to the metaheuristic. This is more efficient than searching through the space of Bayesian Network structures and it has the additional advantage of eliminating all cyclic structures and structures incompatible with the given ordering. It remains to say however that an exhaustive search through all orderings for large problems remains intractable (($O(n!)$ for a problem of size n). [11]

In [16], Larrañaga *et al*. propose a genetic algorithm to search the space of node orderings rather than the full space of structures. The initial individuals in the population are randomly created node orderings which are then evolved until a good ordering is found. In each generation, a pair of individuals is selected for crossover and mutation given the rank of their fitness values in the population. Only one individual offspring is created at a time and, if better, it replaces the worst individual in the current population. The fitness of each ordering is calculated by running the greedy search algorithm K2 on that ordering and returning the score of the network structure found. For the purpose of this paper, we denote Larrañaga's algorithm by K2GA. Figure 1a provides a schematic representation of its operation.

Kabli *et al*. [11] propose an alternative way of reducing the computational cost related to this by using chain structures to evaluate orderings, replacing the K2 expensive evaluation in K2GA.

ChainGA follows a similar approach to K2GA: it searches the space of node orderings and assigns a value to each ordering based on the K2-CH score [21]. However rather than using K2 to construct a network on each ordering, ChainGA evaluates a fixed chain structure. This low resolution evaluation phase terminates in a set of orderings that have the highest evaluated K2-CH scores found with this structure. The K2-CH score captures the probability of a candidate network structure B_s given a set of data D. Formally the discrete probability $P(B_s,D)$ is given by:

$$P(B_s, D) = P(B_s)\prod_{i=1}^{n}\prod_{j=1}^{qi}\frac{(r_i-1)!}{(N_{ij}+r_i-1)!}\prod_{k=1}^{ri}N_{ijk}! \quad (2)$$

Here q_i denotes the number of possible different instances the parent of variable X_i can take. r_i is the number of values X_i has, N_{ijk} denotes the number of cases in the dataset D in which X_i takes value k of its x_i instance when its parent Pa_i has its j^{th} value. N_{ij} is the sum of all N_{ijk} for all values x_i can take. For the Gulf of Mexico dataset, several variables have large value sets, leading to significant computational cost using this approach. ChainGA then enters a second phase where K2 is run on a percentage of the best orderings found to search for a good structure. Overall, ChainGA results in a reduced computation time since the number of links to evaluate is fixed and in general much smaller than that required by K2GA. In [11], Kabli *et al*. compared K2GA and ChainGA on a set of benchmark problems with known networks and trade-offs were observed between computation cost and the quality of the structure found. In the following section we describe experiments with these algorithms run on our rig operations dataset. Figure 1b illustrates schematically its operation.

IV. Data Preparation And Experiments

Various standard data sets are available for experimenting on Bayesian Networks from such domains as medical diagnosis [11], car diagnosis [12], intensive care patient alarm monitoring [11], interplanetary probe raw data interpretation [22], search heuristic for problem solving [22], virtual office assistants [22] or automatic context detection [22]. Some of these datasets have been used in previous work on K2GA and ChainGA by Kabli *et al*. [12].

A. Oil and Gas Market Intelligence Data Set

In our work, we obtained a new and unique dataset, exposing a brand new problem to this research field. This dataset covers Rigs and Wells data sourced by ODS-Petrodata Ltd. It is an extensive dataset offering a promising and challenging problem for our Bayesian Network learning algorithms In addition to the number of factors, two other elements have a direct influence on the run length: the number of values in each variable and the size of the dataset. For this experiment, we used a subset of 2500 cases randomly selected from the dataset. One reason is that more cases were making this experiment impossible to compute in our available time, given the processing facilities available to us. Table 4 illustrates run times for 2500 cases. For 100 and 2500 cases, using K2GA, the run times were about 20 minutes and up to about 42 hours respectively, while for ChainGA, run times were about 3 minutes and up to about 13 hours respectively. No preliminary run could be completed at present using all the 6670 cases available in the dataset.

B. Experiment setting

Following the steps of the K2GA and ChainGA algorithms described above, we built our Bayesian network model that represents the data we have selected

The K2GA and ChainGA algorithm implementations were run 45 times each with 200 generations with a population size of 30 node orderings. Displacement Mutation and Cycle Crossover rates were 0.05 and 0.9 respectively. The selection used was a tournament selection of size 4. Those values were optimised empirically using test runs with 100 cases randomly selected from our dataset. The best scored resulting network was then chosen as the optimal model for the problem at hand. We ran each algorithm 45 times over the Rig-Well dataset and compared the results using a 2-tailed T-test in SPSS [23] to validate their significance.

V. Experimental Results

A. K2GA vs ChainGA on the Wells-Rigs dataset

Figure 2 illustrates the resulting Bayesian network models as displayed by BNJ [24]. In this figure, we can clearly see some intuitive relationships formed in the models created by both K2GA (Figure 2a) and ChainGA (Figure 2b). These relationships are analysed in more detail below.

As illustrated by the t-test statistical results in Table 3, K2GA is producing significantly better scoring structures than ChainGA on our dataset. The best-ever individual for K2GA scored *-55534* compared to *-60203* for ChainGA. This is consistent with observations made on the car diagnosis dataset in [11]. The results from the Car Diagnosis Problem also showed the score of ChainGA is significantly poorer. The Car Diagnosis Problem and the Rig-Wells Problem have a similar number of cases in their dataset. It is to be noted from [11] that ChainGA produces significantly better scoring structure on using the ALARM dataset which contains 37 nodes and 46 edges. The performance of ChainGA relating to K2GA appears therefore to be highly problem-dependent. The influence of problem-specific features however will require further research.

Table 3: 2-tailed t-test of Best Individuals K2 Score across all runs

	N	Mean Score	Standard Deviation	P
K2GA	45	-56197.44	205.2	< 0.0005
ChainGA	45	-66434.34	1237.7	< 0.0005

Table 4, shows that the ChainGA runtime is about a quarter of the K2GA runtime required to determine the network structure, leading to a tradeoff situation. K2GA took an average of 42 hours and 28 minutes to complete a single run whereas ChainGA only tool about 11 hours. The overall long computation times required on this problem are in a large part due to the number of distinct values taken by many of the variables. Considering the vast amount of data available to us, K2GA might not be feasible in some cases for building larger models whereas ChainGA will allow us to build a model.

Table 4: Time Statistics per run over all runs

	Average	Standard Deviation
K2GA	42h 28min	5h 9min
ChainGA	11h 1min	1h 11min

B. Expert evaluation of the Model

The networks produced by both K2GA and ChainGA have been presented to Rig and Wells data experts. Both algorithms discovered interactions between *Rig Capabilities*, *Rig Types* and *Water Depth* nodes. Experts highlighted that those are linked because specific rig types typically operate at a specific range of water depth. Those rig types have specific capabilities for operating at those depths. Another group of interactions is identifiable between *Well Result*, *Well Status* and *Well Type*. In turn the *Total Footage Drilled* node also interacts with the node representing the *Drilling Phase* and the one representing the *Footage Drilled per Day*. Ignoring directions, we can see edges common to both networks. This includes an interaction between the *Water Depth* and the *Rig Type*. Those will be logically related because of the technical abilities of specific rigs to allow them to work at specified depth. The relationships between the *Rig Type*, the *Rig Owner* and the *Rig Contractor* are justified by the propensity of rig owner and contractors to work together repetitively and to be specialized in specific type of rigs built on the same plans. Our networks also identify a relationship between the *Shore Base* and the *region* where the drilling rig is operating. This is another logical geographical association showing the abilities of both networks to learn valid information from data.

The partial separation between Well-related and Rig-related variables (With the exception of geographical and water depth variables) suggests a potential difficulty in using the model as a prediction across domains, but adding some additional key variables might solve that problem. *Water depth*, originating from the well database has emerged as a key variable that correlates with the rig capabilities and hence is likely to be significant to the choice of rig. In the Gulf of Mexico, which typically has a uniform geological profile, this may be a reasonable assumption. Alternatively there may be other factors in Wells and Rigs, not selected for this experiment, that do correlate more closely. Also we would expect geological and other factors to be relevant in more heterogeneous regions. Given the computation times on the 17 factors chosen there are significant technical challenges to be overcome in selecting useful subsets of factors for modelling.

VI. Application and uses of the Bayesian Network Model

Discussions with our collaborators within the oil and gas industry have highlighted a range of potential applications for validated models of rig operations data.

A. Drilling Rig Selection

The model can constitute the basis for a robust and flexible tool for assisting businesses in finding the best rig for a particular job. Consulting the model for optimal match, it enables the identification of rigs suitable for a specific operational demand. Going further, it would be possible to build a recommender system around this model. A recommender system is a system performing information filtering to bring information items to a user; this information is filtered in a way that it is likely to interest the user [25] Adding variables related to the intended drilling task and user preferences in the model allows filtering relevant rig recommendations, using the Bayesian Network-based Model to provide an expectation of performance on the given task..

B. Rig Performance Forecasting

Rig Performance Forecasting is a particularly interesting application of such a model. This would support businesses in their decision to hire a specific rig for a specific job by using performance expectation. Our discovered structures have shown that rig type is conditionally related to features from the Wells dataset. This supports our intuition that combining Wells and Rigs data will inform the rig selection process.

C. Rig Scheduling

To assist a user or software in scheduling rigs, the application would need to estimate a range of expected completion times for a coordinated set of rig operations parameters. Using the Bayesian Network to provide the expectations of the various task times would enable such an application.

VII. Conclusion and Future Work

In this paper we explored methods for the discovery of Bayesian Network from oil and gas data. We have built a Bayesian network model to represent Rigs and Wells data generated from ODS-Petrodata's databses. With this, we explored the use of both K2GA and ChainGA, Genetic Algorithms based on node orderings. Both algorithms found credible network structures as evaluated by industry experts. Although most of the relationships discovered are obvious at this stage, quantification of the conditional dependencies may be of commercial interest. K2GA found significantly better structures than ChainGA on this dataset; however, the computational effort required for ChainGA is about a quarter, on average. Due to the size and complexity of the datasets being considered, irrespectively of which approach is used, further work will be required to improve the computation time and to meet the challenge of factor selection. One possible improvement would be to use different scoring metrics such as Minimum Description Length (MDL) [26] [14] [18] [14] as suggested by Kabli *et al*. [12]

This research is a step toward a model that could be used to answer various queries relating to applications such as Drilling Rig Selection, Rig Performance forecasting and Rig Operation Scheduling. We obtained the evaluation of the current model by domain experts and saw promising results in spite of the limited number of variables used. The potential of Bayesian networks in this case is to support decision making in a more intuitive and objective way than current human processing methods. Future work on Bayesian Network modelling, and especially on shortening its learning time, is likely to benefit the Bayesian Network community; however, where possible concurrent methods should be evaluated. That is planned in future work. We plan to explore the ability of Bayesian Networks by including more data in the model and to do a larger scale comparison across more variables. Covering larger geographical regions and ultimately worldwide data will in the future allow us to develop the model into a global application. Our planned experimental setup will also allow us to measure the ability of the model to generalise new data. Given the current computation times, there are significant technical challenges to be overcome in selecting useful subsets of factors for modelling. This will be a focus for future work. The positive feedback received from field experts for this initial work encourages us to pursue further model and application building.

ACKNOWLEDGEMENT

The authors would like to thank Robert Steven for his help in providing expert evaluation of the networks.

REFERENCES

[1] ODS-Petrodata Ltd, "http://rigpoint2.ods-petrodata.com/,"
[2] J. Harris. (1989, Selecting an offshore drilling rig - the competitive tendering process. *Offshore Europe*
[3] A. Nergaard Prof. Dr., "Offshore drilling technology," vol. 2010, 14 December 2005. 2005.
[4] C. Freudenrich Ph.D., "How Oil Drilling Works: Extracting the Oil," vol. 2010, 12 April 2001. 2001.
[5] P. Osmundsen, T. Sørenes and A. Toft. (2008, Drilling contracts and incentives. *Energy Policy 36(8),* pp. 3128-3134.
[6] ODS-Petrodata Ltd., "http://www.ods-petrodata.com/," 2010.
[7] W. Wiegerinck, "petrophysical decision support " 2010.
[8] W. Wiegerinck, B. Kappen and W. Burgers. Bayesian networks for expert systems, theory and practical applications. *Networks 20(14),* pp. 7.
[9] P. R. Kannan. (2007, Bayesian networks: Application in safety instrumentation and risk reduction. *ISA Trans. 46(2),* pp. 255-259.
[10] D. Niedermayer. (1998, An introduction to bayesian networks and their contemporary applications. *Retrieved may 6*pp. 2004.
[11] R. Kabli, F. Herrmann and J. McCall, "A chain-model genetic algorithm for bayesian network structure learning," in *GECCO'07,* 2007, pp. 1264-1264-1271.
[12] R. Kabli, F. Herrmann and J. McCall, "Evolved bayesian networks as a versatile alternative to partin tables for prostate cancer management," in *GECCO'08,* 2008, pp. 1547-1547-1554.
[13] Robinson R., "Counting labeled acyclic digraphs," in *New Directions in the Theory of Graphs* Frank Harary, Ed. New York: Academic Press, 1973, pp. 239-239–273.
[14] M. L. Wong, S. Y. Lee and K. S. Leung. A hybrid data mining approach to discover bayesian networks using evolutionary programming. Presented at Proceedings of the Genetic and Evolutionary Computation Conference.
[15] J. Habrant. Structure learning of bayesian networks from databases by genetic algorithms-application to time series prediction in finance. Presented at Proceedings of the 1st International Conference on Enterprise Information Systems.
[16] P. Larranaga, C. Kuijpers, R. Murga and Y. Yurramendi. (1996, Learning bayesian network structures by searching for the bestordering with genetic algorithms. *IEEE Transactions on Systems, Man and Cybernetics, Part A 26(4),* pp. 487-493.
[17] A. J. Novobilski. The random selection and manipulation of legally encoded bayesian networks in genetic algorithms. Presented at The 2003 International Conference on Artificial Intelligence (ICAI).
[18] S. van Dijk, D. Thierens and L. C. van der Gaag. (2003, Building a GA from design principles for learning bayesian networks. *Lecture Notes in Computer Science* pp. 886-897.
[19] W. Buntine. (1994, Operations for learning with graphical models. *Arxiv Preprint cs/9412102*
[20] R. R. Bouckaert. (1995, Bayesian belief networks: From construction to inference. *PhDthesis, University Utrecht*
[21] G. F. Cooper and E. Herskovits. (1992, A bayesian method for the induction of probabilistic networks from data. *Mach. Learning 9(4),* pp. 309-347.
[22] D. Niedermayer, "An introduction to bayesian networks and their contemporary applications," in Anonymous 1998,
[23] SPSS 15.0 for Windows, "Rel. 15.0.1.1," 2006.
[24] Bayesian Network tools in Java, "http://bnj.sourceforge.net/," vol. 2010,
[25] P. Resnick and H. R. Varian, "Recommender systems," in *Recommender Systems* ,CACM ed., vol. 40, Anonymous ACM, 1997,
[26] R. R. Bouckaert. (1993, Probabilistic network construction using the minimum description length principle. *Lecture Notes in Computer Science* pp. 41.

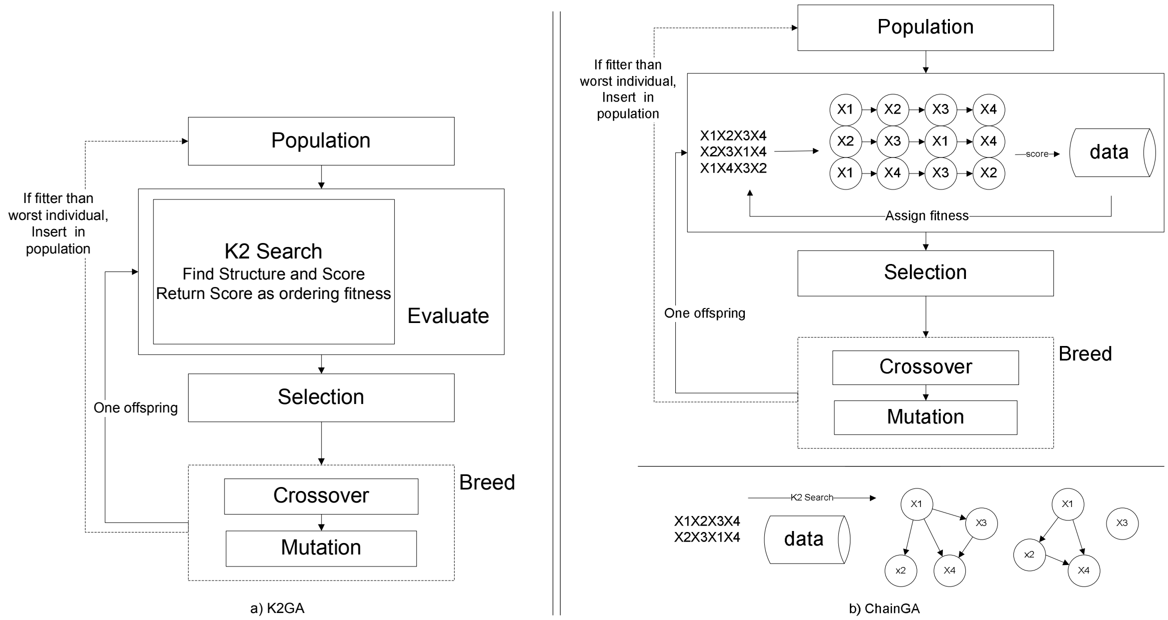

Figure 1: K2GA and ChainGA

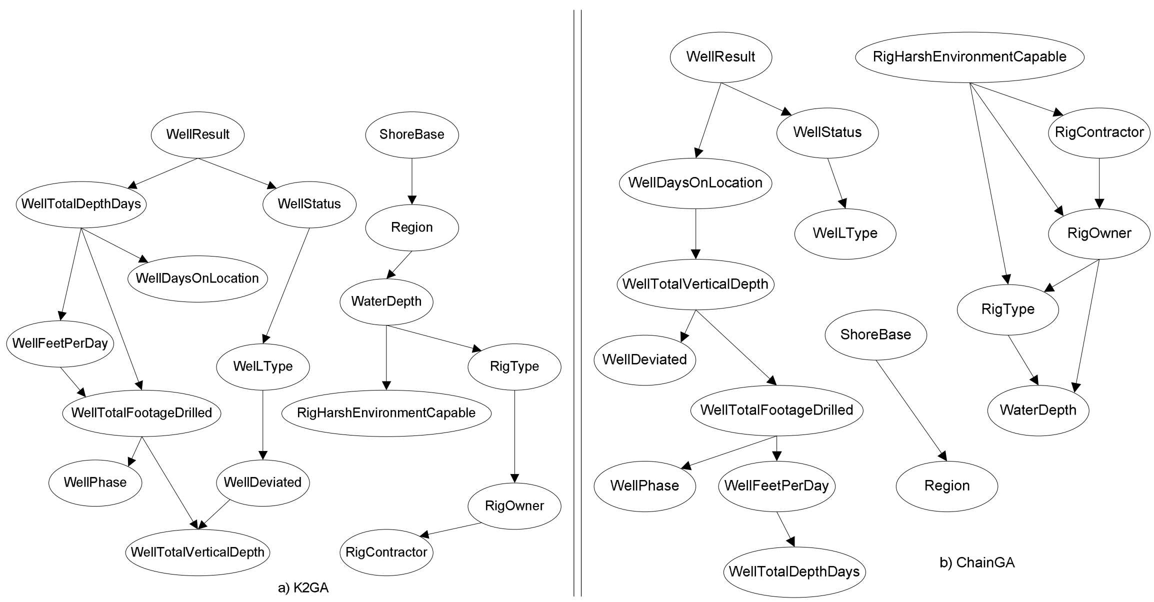

Figure 2: Network Representations for K2GA and ChainGA