Risk Scoring using Fuzzy Bandits with Knapsacks

Benjamin Otieno
Department of Information Technology
Kibabii University
Bungoma, Kenya

George Musumba
Department of Computer Science
Dedan Kimathi University of Technology
Nyeri, Kenya

Franklin Wabwoba
Department of Information Technology
Kibabii University
Bungoma, Kenya

Abstract:- Bandits with Knapsacks (BwK) is a Multi-Armed Bandit (MAB) problem under supply/budget constraints. Risk scoring is a typical limited-resource problem and as such can be modeled as a BwK problem. This paper tries to solve the triple problem in risk scoring – accuracy, fairness, and auditability by proposing FuzzyBwK application in risk scoring. Theoretical assessment of FuzzyBwK is made to establish whether the regret function would perform better than StochasticBwK and AdversarialBwK functions. An empirical experiment is then set with secondary data (Australian and German credit data) and primary data (Kakamega insurance data) to determine whether the algorithm proposed would be fit for the proposed problem. The results show that the proposed algorithm has optimal regret function and from the empirical test, the algorithm satisfies the test of accuracy, fairness and auditability.

Keywords:- *Multi-Armed Bandits, Bandits With Knapsacks, Fuzzy Logic, FURIA, Risk Scoring.*

I. INTRODUCTION

Multi-armed bandits (MAB) is a model used to balance the acquisition and usage of information. In the model, the tradeoff between exploration and exploitation is established by a decision-making agent. Studies on the likelihood of one unknown probability having a better outcome over another unknown probability, which is the basis of MAB, have been done over time with some of the early works being dated as back as the 1930s [1]. The MAB algorithm chooses in each round a set of alternatives (called arms) and receives a reward for the alternative chosen. The algorithm does not have enough information to answer all "counterfactual" questions about what would have happened if a different action was chosen in each round [2] and as such, choices of alternatives (exploitation) are alternated with learning expeditions (exploration) for optimal rewards.

Typical MAB problems have an assumption that the resources are unlimited and the bandit can endlessly make choices while learning. However, in many real-life situations, there are cases where the resources are limited. In other words, there are constraints in supply or budgets. The financing problem which is modeled in this research is a typical case where there is a limitation in budget. The banks usually have limited lending headroom and insurance companies have maximum risk capacity. The assumption made is that there is a knapsack with limited capacity and the problem is therefore called Bandits with Knapsack (BwK) [3].

BwK problem can be modeled as a linear problem where the context vectors are linear and are generated in identically independent distribution [4]. The problem can also be modeled as a Stochastic BwK in which the probability distribution is fixed but cannot be predicted precisely [5]. For the risk scoring problem, the option for an adversarial distribution was chosen over the stochastic option. Adversarial because of the assumption made that there are unlimited loan-seekers and limited borrowing headroom with the lender but the characteristics of the lenders keep changing over time. The distribution is no longer fixed. In Adversarial BwK, both the reward and resource can change. You do not expect the lending headroom of a bank for example to be static over time. And the risk profiles of loan seekers also change over time depending on many factors including micro and macro factors. Adversarial BwK is a much harder problem compared to Stochastic BwK. The new challenge is that the algorithm needs to decide how much headroom to save for future lending, without being able to predict that there will be corresponding loan-seekers [6].

However, in this work, the Adversarial BwK is extended to use Fuzzy logic in the explore-exploit subroutine. This serves two main purposes. The first reason is that risk scoring is an imprecise process and the use of algorithms giving crisp outcomes can introduce some bias to the outcomes [7]. Other than the challenge of bias (or lack of fairness) introduced by machine learning, banking supervisors worry about the unexplainability of the algorithms used [8]. Many algorithms, including the Adversarial BwK are Blackbox algorithms making auditing the algorithms a difficult task hence the extreme position taken by some authorities such as GDPR not to use automated decision-making mechanisms exclusively [9] for decision making. Fuzzy Unordered Rule Induction Algorithm (FURIA) proposed here as the explore-exploit subroutine is easily auditable since it comprises simple rules and can be regarded as a Whitebox.

II. MAIN CONTRIBUTION

This work extends the bandits with limited resources otherwise called Stochastic BwK [6] to use fuzzy logic in the exploration and exploitation so as to make the stochastic game as close to human imprecise reasoning as possible. Whereas we can gain regret of the order of O(log T) in Stochastic BwK, Fuzzy BwK gives lower regret of the order of p(log(p/t)-log(P/T)) hence making Fuzzy BwK give better rewards over the distribution over arms.

This paper also presents a white-box subroutine as opposed to the Stochastic BwK which is a blackbox subroutine. Whitebox subroutine is easily auditable and hence can be easily accepted by financial services regulators, as opposed to Blackbox which has faced opposition from regulators, with some regulators banning the use of automated risk scoring mechanisms on the account of the inauditability of the algorithms.

III. RELATED WORKS

Multi-armed bandits is a typical solution for a problem where a balance between acquisition of information and the usage of the information is required. Many researchers have worked on the MAB problem with early works dating back to 1930s [1].

The MAB problem considered under limited resources was Bandits with knapsacks as first conceived by [10]. Many different versions have since been investigated, including Linear Contextual BwK [11], and Combinatorial Semi BwK [12].

The regret of the stochastic BwK is in the order O(T) with the regret evaluating to $O\sqrt{T}$ for linear BwK, $O\left(\left(\frac{OPT}{B}+1\right)m\sqrt{T}\right)$ for linear contextual BwK, $O(\sqrt{n})\left(\frac{OPT}{\sqrt{B}}+\sqrt{T+OPT}\right)$ for combinatorial semi BwK. All these are linearly distributed.

Adversarial BwK, which is a harder problem compared to stochastic problem because a condition on limitation of resources/budget, is introduced by [6]. The regret for Adversarial BwK is better than the O(T) achieved by the stochastic BwK. It evaluates to something close to OLog(T). The high-probability algorithm for Adversarial BwK evaluates to O(d log T), the modified uniform exploration under the Adversarial BwK evaluates to $O(T^{3/4})\sqrt{Klog\frac{T}{\delta}}$ while the adaptive Adversary evaluates to $O\left(\sqrt{nT}\log(1/\delta)\right)$ [6].

So far, there is no research that has ventured in use of fuzzy logic in consumer risk scoring. It has however been experimented on work accidents and occupational risks [13], enterprise risk [14], operational and emerging risks [15]. Traditional models based on probability and/or set theory are premised on the fact that a client is either in a set or not. This assumption provides accurate information if there is complete information provided by the entity being assessed. This paper explores the use of Fuzzy as the subroutine in the BwK. The paper will also evaluate the regret function of the Fuzzy BwK and compare with both the Stochastic BwK and Adversarial BwK regrets.

For practicality, the use of Fuzzy Unordered Rule Induction Algorithm (FURIA) was considered. FURIA is a fuzzy eager learner that does not start with default rules [16]. Rather, it learns and updates the rules on an ongoing basis. Essentially, it gets into information acquisition and usage cycles, a typical feature of MAB.

IV. EXPERIMENT

The experiment is divided into two parts. The first part was a theoretical assessment of the Fuzzy BwK to confirm the regret value against Stochastic and Adversarial BwK regrets. The second part is experimentation of FURIA against other algorithms on three different datasets to evaluate for accuracy and fairness. One of the datasets is german credit data[i] , the second is Australian credit data[ii] and the third is Kakamega agricultural insurance data[iii]. The algorithms against which FURIA is evaluated are algorithms commonly used for risk scoring and they include Logistic regression, Naïve Bayes, K-Nearest Neighbors, Bayesian-Network, Additive logistic regression boosting, Adaptive Boosting and Locally weighted learning.

V. FUZZY BWK

In BwK, we take it that there are $d \geq 1$ limited resources to be consumed and that in each round of choice, the algorithm chooses an alternative (arm) from a set of K actions. There are $d \geq 1$ loan-seekers and for each choice, a loan-seeker (arm) $a_t \in [K]$ is chosen for loan advancement from a set of K applicants in each round $t \in [T]$. We also take it that r_t is a reward and $c_{t,i}$ is the consumption of each resource $i \in [d]$. Each resource i is endowed with budget $B_i \leq T$. The game stops early, at some round $\tau_{alg} < T$, when the total consumption of any resource exceeds its budget. The algorithm's objective is to maximize its total reward. In other words, to lend to the optimum lending headroom, to the set of borrowers with the least default probability.

In adversarial learning, given a set of actions A, and a payoff range denoted by $f_t \in [b_{min}, b_{max}]^K$, for some distribution f(t) chosen as a deterministic function in history, the risk-scoring algorithm with known upper bound on regret is given by

$$R = \left[max_{a\in A}\sum t \in [T]\,(f)_t(a)\right] - \left[\sum t \in [T]\,(f_t)(a_t)\right]$$

We consider regret minimization in games with a zero-sum game (A_1, A_2, G) being a game between two players $i \in \{1, 2\}$ with action sets A_1 and A_2 and payoff matrix $G \in R^{A1\times A2}$.Repeated over the Lagrangian game, this evaluates to

$$REW(LagrangeBwK) \leq O\left(\frac{T}{B}\sqrt{T\,K\log\frac{dT}{\delta}}\right) + OPT_{dp}$$

The Adversarial regret as evaluated by [6] is
$O(REG)(Adversarial\ BwK) \leq \left(1 + OPT_{dp}(\pi)\right)$

where π is the Lagrange function.

By greedily growing the rules in FURIA, the upper bound would evaluate to

$$O(REG)(Fuzzy\ BwK) \leq \left(1 + OPT_{dp}(\mathrm{p}(\log(\mathrm{p/t}) - \log(\mathrm{P/T})))\right)$$

In the Adversarial BwK, the algorithm will be as slow as π will allow. The logarithm smoothening introduced in the FuzzyBwK is meant to compensate for this. The FuzzyBwK is of order Olog(T) making it better than the stochastic O(T). As such, FuzzyBwK promises better performance in the explore-exploit MAB problem compared to StochasticBwK and Adversarial BwK.

VI. EMPIRICAL RESULTS

The following section compares regret achieved by various algorithms and this is compared with what is achieved by FuzzyBwK. The experiment is made by testing various algorithms including FURIA using three datasets. The three datasets are the German, Australian and Kakamega datasets and the evaluation is against accuracy and fairness. Fairness is measured by determining the incorrectly classified entities. As for fairness, two common error estimators were used to determine how different the automated assessment scores were from human assessments. These estimators are Mean Absolute Error (MAE) and Relative Absolute Error (RAE). The assumption made is that human assessment is devoid of bias and as such the algorithm that returns the lowest MAE and RAE should be having the lowest bias (highest fairness), or would be regarded as counterfactually fair.

MAE is used to determine the quantum of error between the measured value and the 'true' value.

For n number of items to be classified, MAE is calculated by

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|y_i - x_i|$$

for every pairs of predicted value y_i and true value x_i.

RAE as a measure takes into consideration both bias and accuracy. It is the ratio of the absolute error of a measurement to the measurement being taken. As opposed to MAE, RAE is used to put the error into perspective.

RAE is given by getting the difference between the approximated value v_A from the real value v_E and then dividing the absolute difference by the real value. This is then expressed as a percentage.

$$RAE = \left|\frac{v_A - v_E}{v_E}\right| x100\%$$

From Table 1, it is clear that FURIA returns the lowest regret, evidenced by the lowest incorrectly classified items and the lowest MAE and RAE. FuzzyBwK promises better results compared to other algorithms commonly used in risk scoring.

		Accuracy				Fairness							
		Incorrectly classified (Ger)	Incorrectly classified (Aus)	Incorrectly classified (Kak)	Incorrectly Classified (Aver)	MAE (Ger)	MAE (Aus)	MAE (Kak)	MAE (Aver)	RAE (Ger)	RAE (Aus)	RAE (Kak)	RAE (Aver)
1	Logistic regression	22.65%	16.60%	7.94%	15.73%	0.2986	0.1994	0.021	0.1730	72.02%	40.26%	28.10%	47%
2	Naïve Bayes	22.65%	24.68%	12.94%	20.09%	0.2811	0.2560	0.025	0.1874	67.77%	51.68%	33.50%	51%
3	K-NN	30.00%	21.70%	10.29%	20.66%	0.3006	0.2183	0.020	0.1796	72.49%	44.06%	26.83%	48%
4	B-Net	25.00%	14.04%	10.88%	16.64%	0.3021	0.3540	0.025	0.2270	72.85%	34.90%	33.16%	47%
5	Additive logistic reg. boosting	25.29%	14.47%	7.94%	15.90%	0.3228	0.2106	0.021	0.1848	77.83%	42.52%	28.10%	49%
6	Adapt. Boosting	25.29%	15.32%	14.41%	18.34%	0.3274	0.3367	0.040	0.2347	78.94%	43.52%	53.30%	59%
7	Locally weighted learning	28.82%	14.47%	14.41%	19.23%	0.3598	0.2259	0.037	0.2076	86.76%	45.61%	48.51%	60%
8	FURIA	23.23%	11.06%	10.29%	14.86%	0.2535	0.1264	**0.0160**	0.1320	64.14%	27.73%	**21.29%**	38%

Table 1. German, Australia and Kakamega Accuracy and Fairness

The rule-set in the Whitebox that is FURIA is as presented in Fig. 1, 2 and 3. They are easily auditable and as such can easily find acceptance by banking supervisors/regulators.

i.(checking_status = no checking) => class=good (CF = 0.88)
i.(duration in [-inf, -inf, 20, 21]) and (credit_amount in [1274, 1288, inf, inf]) and (housing = own) => class=good (CF = 0.82)
i.(duration in [-inf, -inf, 11, 12]) => class=good (CF = 0.85)
v.(installment_commitment in [-inf, -inf, 2, 3]) and (credit_amount in [-inf, -inf, 8072, 8086]) => class=good (CF = 0.78)
v.(savings_status = no known savings) and (credit_amount in [1977, 2181, inf, inf]) => class=good (CF = 0.86)
i.(checking_status = <0) and (duration in [15, 16, inf, inf]) and (credit_amount in [-inf, -inf, 2145, 2348]) and (employment = 1<=X<4) => class=bad (CF = 0.91)
i.(checking_status = <0) and (duration in [21, 24, inf, inf]) and (installment_commitment in [2, 3, inf, inf]) and (own_telephone = no) => class=bad (CF = 0.77)
i.(credit_amount in [3990, 4057, inf, inf]) and (age in [-inf, -inf, 29, 30]) and (duration in [30, 36, inf, inf]) => class=bad (CF = 0.75)
x.(checking_status = 0<=X<200) and (duration in [20, 21, inf, inf]) and (savings_status = <100) => class=bad (CF = 0.66)
x.(checking_status = <0) and (age in [-inf, -inf, 35, 36]) and (credit_amount in [-inf, -inf, 1442, 1498]) => class=bad (CF = 0.67)

Fig 1. FURIA Rule Set for German Dataset

(A9 = t) and (A15 in [225, 234, inf, inf]) => A16=+ (CF = 0.95)
(A9 = t) and (A10 = t) => A16=+ (CF = 0.9)
(A9 = t) and (A14 in [-inf, -inf, 110, 120]) and (A4 = u) => A16=+ (CF = 0.93)
(A9 = t) and (A7 = h) => A16=+ (CF = 0.86)
(A9 = f) and (A3 in [0.375, 0.415, inf, inf]) => A16=- (CF = 0.96)
(A9 = f) => A16=- (CF = 0.93)
(A10 = f) and (A14 in [100, 108, inf, inf]) and (A15 in [-inf, -inf, 1, 40]) => A16=- (CF = 0.83)

Fig 2. FURIA Rule Set for Australian Dataset

i. (SOURCE = BROKERS) and (SI BAND = 1M<=2M) => PRICING RATE=2% (CF = 0.78)
ii. (GROSS PREMIUM = 24,500) => PRICING RATE=2% (CF = 0.75)
iii. (SUM INSURED = 18,000) => PRICING RATE=11% (CF = 0.5)
iv. (GROSS PREMIUM = 11,000) and (SOURCE = BROKERS) => PRICING RATE=11% (CF = 0.5)
v. (SUM INSURED = 40,000) and (GROSS PREMIUM = 2,000) => PRICING RATE=5% (CF = 0.88)
vi. (SOURCE = AGENTS) => PRICING RATE=4% (CF = 0.89)
vii. (SOURCE = DIRECT) => PRICING RATE=4% (CF = 0.61)
viii. (SUM INSURED = 50,000) and (GROSS PREMIUM = 2,000) => PRICING RATE=4% (CF = 0.93)
ix. (GROSS PREMIUM = 6,000) => PRICING RATE=4% (CF = 0.89)
x. (SUM INSURED = 55,000) => PRICING RATE=4% (CF = 0.59)
xi. (GROSS PREMIUM = 10,000) => PRICING RATE=4% (CF = 0.56)
xii. (SUM INSURED = 45,000) and (GROSS PREMIUM = 2,000) => PRICING RATE=4% (CF = 0.67)
xiii. (SUM INSURED = 370,000) => PRICING RATE=1% (CF = 0.5)
xiv. (SUB CLASS = CROP INSURANCE-SW) and (SI BAND = 1M<=2M) => PRICING RATE=7% (CF = 0.72)
xv. (SUM INSURED = 30,000) and (GROSS PREMIUM = 2,000) => PRICING RATE=7% (CF = 0.61)
xvi. (SOURCE = DIRECT) and (SUM INSURED = 240,000) => PRICING RATE=6% (CF = 0.61)
xvii. (SOURCE = DIRECT) and (SUM INSURED = 32,000) => PRICING RATE=6% (CF = 0.51)
xviii. (SOURCE = DIRECT) and (SUM INSURED = -240,000) => PRICING RATE=6% (CF = 0.51)
xix. (GROSS PREMIUM = 8,000) and (SUM INSURED = 100,000) => PRICING RATE=8% (CF = 0.5)
xx. (SOURCE = BANCASSUARANCE) and (GROSS PREMIUM = 4,800) => PRICING RATE=3% (CF = 1.0)
xxi. (SOURCE = BANCASSUARANCE) and (SUB CLASS = LIVESTOCK INSURANCE-SW) => PRICING RATE=3% (CF = 0.93)
xxii. (SUM INSURED = 15,000) => PRICING RATE=13% (CF = 0.6)

Fig 3. FURIA Rule Set for Kakamega Dataset

VII. CONCLUSION

Risk scoring requires accuracy, fairness and auditability. Fuzzy logic, applied in the Fuzzy BwK using FURIA as the subroutine in BwK results in higher accuracy scores compared to other algorithms commonly used in risk scoring. Accuracy is important since it inspires confidence among the financial institutions using the algorithm. No financial institution wants to grow a bad loan book. Fairness is also important since loan seekers want to have confidence that they will be treated fairly by the algorithms without being discriminated against. FURIA promises counterfactual fairness, evidenced by the low MAE and RAE in the results. The simple rules used by FURIA are easy to audit unlike the blackbox algorithms that many of the machine learning algorithms are. This means that it can be easily accepted by banking supervisors who for a long time have had issue with algorithms that cannot be explained or audited.

REFERENCES

[1]. W. R. Thompson, "On the likelihood that one unknown probability exceeds another in view of the evidence of two samples," Biometrika 25, no. 3-4 , pp. 285-294, 1933.
[2]. K. A. Sankararaman and A. Slivkins, "Combinatorial semi-bandits with knapsacks," in International Conference on Artificial Intelligence and Statistics, PMLR, 2018.
[3]. A. Badanidiyuru, R. Kleinberg and A. Slivkins, "Bandits with knapsacks," In 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, pp. 207-216, 2013.
[4]. S. a. N. D. Agrawal, "Linear contextual bandits with knapsacks," Advances in Neural Information Processing Systems 29, 2016.
[5]. K. A. Sankararaman and A. Slivkins, "Bandits with knapsacks beyond the worst case," Advances in Neural Information Processing Systems 34, pp. 23191-23204, 2021.
[6]. N. Immorlica, K. A. Sankararaman, R. Schapire and A. Slivkins, "Adversarial Bandits with Knapsacks," in 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), Baltimore, MD, USA, 2019.
[7]. B. Otieno, F. Wabwoba and G. Musumba, "Bandits using Fuzzy to fill Knapsacks: Smallholder Farmers Credit Scoring," In 2021 IST-Africa Conference (IST-Africa), IEEE, pp. 1-10, 2021.
[8]. J. Prenio and J. Yong, "Humans keeping AI in check-emerging regulatory expectations in the financial sector," Bank for International Settlements, Financial Stability Institute, 2021.
[9]. GDPR, "General Data Protection Regulation," 2018.
[10]. A. Badanidiyuru, R. Kleinberg and Slivkins, Bandits with knapsacks, FOCS, 2013.
[11]. S. Agrawal and N. Devanur, "Linear contextual bandits with knapsacks," Advances in Neural Information Processing Systems 29 , 2016.

[12]. K. A. Sankararaman and A. Slivkins, "Combinatorial semi-bandits with knapsacks," in International Conference on Artificial Intelligence and Statistics, 2018.
[13]. I. L. Nunes and M. Simões-Marques, "Applications of Fuzzy Logic in Risk Assessment – The RA_X Case," in Fuzzy Inference System - Theory and Applications, 2012.
[14]. A. F. Shapiro and M.-C. Koissi, "Risk assessment applications of fuzzy logic," Society of Actuaries, 2015.
[15]. K. Shang and Z. Hossen, "Applying Fuzzy Logic to Risk Assessment and Decision-Making," 2013.
[16]. J. Hühn and E. Hüllermeier, "FURIA: an algorithm for unordered fuzzy rule induction," Data Mining and Knowledge Discovery 19, no. 3 , pp. 293-319, 2009.

[i] ftp.ics.uci.edu/pub/machine-learning-databases/statlog/
[ii] http://archive.ics.uci.edu/ml/datasets/statlog+(australian+credit+approval)
[iii] https://drive.google.com/file/d/1zp17LHjK_KMNl2oNqjaSug-_zOULM1th/view?usp=sharing