

Customs Fraud Detection in the Presence of Concept Drift

Tung-Duong Mai*, Kien Hoang*, Aitolkyn Baigutanova*, Gaukhartas Alina*,
Korea Advanced Institute of Science and Technology
 {maitungduong2605, kienapp286, baiaitolkyn, gaukhar112200}@gmail.com

Sundong Kim
Institute for Basic Science
 sundong@ibs.re.kr

Abstract—Capturing the changing trade pattern is critical in customs fraud detection. As new goods are imported and novel frauds arise, a drift-aware fraud detection system is needed to detect both known frauds and unknown frauds within a limited budget. The current paper proposes ADAPT, an adaptive selection method that controls the balance between exploitation and exploration strategies used for customs fraud detection. ADAPT makes use of the model performance trends and the amount of concept drift to determine the best exploration ratio at every time. Experiments on data from four countries over several years show that each country requires a different amount of exploration for maintaining its fraud detection system. We find the system with ADAPT can gradually adapt to the dataset and find the appropriate amount of exploration ratio with high performance.

Index Terms—Concept Drift, Customs Fraud Detection, Exploration-Exploitation Dilemma, Multi-Armed Bandit

I. INTRODUCTION

With enormous daily trade traffic, effective trade handling becomes the main task of customs administrations, and the urge for AI-based fraud detection mechanisms becomes apparent [1], [2]. However, maintaining a sustainable fraud detection system is challenging due to the changes in customs trades and fraudulent transaction trends with time. A fraud detection model trained on historical data often falls into a confirmation bias and results in performance degradation [3]. The change in data characteristics and distribution over time is referred to as concept drift, which can be gradual, incremental, recurrent, or sudden [4]. In customs, concept drift is caused by alterations of importers, goods types, and business partners [5].

The customs workflow follows the human-in-the-loop inspection format, where physical inspections are carried out with the help of a fraud detection model, and customs officers confirm whether the declared item is fraud or not. If the inspection reveals fraud, the officers can levy extra duties and the results will be used to update the fraud detection model. Usually, the model provides the most suspicious items for inspection, which can be problematic in countries facing concept drift in their trade pattern.

While the fraud detection model should keep on catching the known frauds and secure the revenue, it should also acquire knowledge about new fraudulent behavior. These conflicting objectives can be described as the famous exploration-exploitation dilemma [6], where there must be a balance

*Work done while interning at Institute for Basic Science

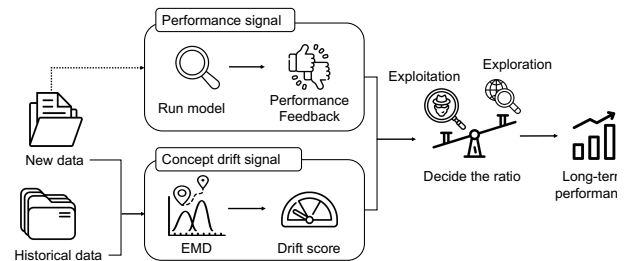


Fig. 1: ADAPT uses two signals for model sustainability.

between selecting illicit items to secure immediate revenue (exploitation) and discovering new items to maintain long-term performance (exploration).

One practical solution is to randomly inspect a small set of declared items. This is called random selection, which can estimate their trade statistics and learn new fraud patterns [7]. Adding some random exploration together with exploitation strategy is shown to be effective in the presence of concept drift [3]. However, there are no studies on how to determine the best exploration amount in customs settings. Empirically setting the exploration ratio is problematic where the amount of concept drift in the underlying data is unpredictable, and the fraud detection performance is susceptible to the amount of exploration.

This research introduces an Adaptive Drift-Aware and Performance Tuning (ADAPT) method to decide the balance between exploration and exploitation in detecting customs frauds. We measure the amount of concept drift to set a baseline of how much exploration should be used. In addition, we incorporate a performance signal to account for the fitness of the current model to the latest data. We utilize a multi-armed bandit framework with each arm corresponding to an exploration ratio. The changes in data distribution are measured and used to select a range of candidate arms for further consideration. Then, the model considers the historical performance of the candidate arms for its final decision. Fig. 1 illustrates the concept of the method.

To evaluate ADAPT, we benchmark its performance against the optimal strategy, which determines the exploration ratio by accessing the whole dataset in advance. This setting is unrealistic but provides a meaningful “gold standard” for the model

assessment. Experiments on multi-year and million-item import declarations from the four countries report that customs fraud detection system with ADAPT operates sustainably for a long period. Further analysis on the import declarations with previously unencountered characteristics demonstrates that ADAPT is effective in detecting new types of frauds. Finally, we demonstrate the importance of concept drift by examining its correlation with the model performance.

For risk management in customs administrations, ADAPT provides two benefits. It can automatically balance the exploration and exploitation rate, without the need for external information, such as other models' performance or extensive hyperparameter tuning. This will save time and guarantee efficiency in the customs workflow. It can also act as a warning trigger for the offices and inform that trade pattern changes.

II. RELATED WORK

A. Customs fraud detection

Previous works in the customs fraud detection domain suggest solutions in two major directions. The first body of work proposes approaches for exploring new cases, from the simple but intuitive random examination [7] to more advanced active learning solutions using uncertainty [8] and diversity [9], [10]. The second group of work focuses on exploiting based on the previously obtained knowledge, which includes adapting heuristic approaches and off-the-shelf machine learning techniques, such as XGBoost [11], SVM [12], or sophisticated deep-learning-based methods like DATE [13]. Inspecting some random items helps to improve the fraud detection performance even it sacrifice some known frauds [3]. Determining the exploration ratio in the human-in-the-loop fraud detection problem is under-explored, which is the goal of this paper.

B. Concept drift detection

Due to the proven effectiveness of concept drift for data analysis and performance of machine learning algorithms, there is extensive literature targeting this problem. Lu et al. classified these approaches [4] in three categories; error-based drift detection [14], data distribution-based drift detection [15], [16], and multiple hypothesis test drift detection [17]. Error-based approaches refer to the performance of the classifier, while the distribution-based methods utilize the concept of dissimilarity between new and historical data distributions [18]. The third category uses statistical tests to decide the final concept drift value. In this work, we conduct distribution-based drift detection analysis and further use this result to decide the exploration ratio of a fraud detection system.

C. Multi-armed bandit problem

The Multi-Armed Bandit Problem (MABP) addresses the dilemma of balancing between exploration and exploitation. The goal of the problem is to decide the best strategy for achieving a higher reward. Specifically, the decision is made

between either relying on the previously discovered information or giving a chance for exploring new choices. A representative approach for dynamically adjusting the exploration-exploitation ratio is the exponential weighted framework. RPI algorithm leverages an online learning mechanism using this framework to dynamically tune the ratio [19]. MABP has two settings, stochastic and adversarial. They differ by the reward model being applied [20]. This work belongs to the adversarial bandits, where the agent makes no assumption regarding the reward generation process. A representative approach for the adversarial bandits in an online setting is the EXP3 algorithm [21]. It works by assigning weights for each action, picking an action based on assigned weight, and then updating its weight based on the observed payoff. EXP3.S [22] is a variant of the EXP3 algorithm, which adds regularization to ease arm switching. We improve upon their method by additionally accounting for the effect of concept drift.

III. PROBLEM DEFINITION

A. Customs selection problem

The human-in-the-loop customs selection problem, introduced in [3], is formally described as follows:

At each timestamp t , customs receive a batch of items \mathcal{B}_t from trade flows \mathcal{B} . Based on a strategy f trained with labeled data X_t , customs officers select a batch of items \mathcal{B}_t^S for manual inspection. After inspection, the newly annotated results are added into the training data for the next iteration X_{t+1} . The goal is to devise a strategy f^* that maximizes the precision and revenue in the long term:

$$f^* = \operatorname{argmax}_f \sum_t m(\mathcal{B}_t^S(f)), \quad (1)$$

where m is the key performance index for the fraud detection system such as precision and revenue.

B. Customs selection strategies

A customs selection strategy f generally falls into one of the two categories:

- Exploitation strategy f_{exploit} : Selecting likely fraudulent and lucrative items. It guarantees customs to collect tariffs, ensuring revenue.
- Exploration strategy f_{explore} : Selecting unseen items. Adequately exploring diverse trade items allows to detect novel frauds and prevent confirmation bias.

A mixture of exploitation and exploration strategies—a hybrid approach is proven to effectively maintain long-term performance [3]. The hybrid approach selects $k|\mathcal{B}_t^S|$ from f_{exploit} and $(1-k)|\mathcal{B}_t^S|$ items from f_{explore} , with k being the constant exploration ratio. The concept of hybrid strategy using a constant amount of exploration is similar to the ϵ -greedy algorithm [20]. However, unlike ϵ -greedy, the hybrid approach can support any exploration strategy f_{explore} other than random exploration.

C. Finding the best exploration ratio

In this paper, we extend the problem to the case where k is no longer a constant. The algorithm aims to find the best exploration ratio k_t for every timestamp t . The problem can be formally defined as follows: Given the exploration strategy f_{explore} and exploitation strategy f_{exploit} , choose the exploration ratio k_t for current timestamp t .

$$k_t^* = \underset{k_t}{\operatorname{argmax}} \sum_t m(\mathcal{B}_t^s(k_t; f_{\text{explore}}, f_{\text{exploit}})). \quad (2)$$

IV. METHOD

The performance of the machine learning model for customs fraud detection is sensitive to the degree of exploration [3]. However, there is little work regarding how to manage this trade-off effectively. To make a robust selection model, we consider an exploration-exploitation trade-off in our proposed method, updating the ratio according to two signals; model performance and concept drift.

Those two signals relate to two concepts: the multi-armed bandit problem, where actions are chosen based on rewards, and concept drift when data shift its distribution over time. When the performance drops or the underlying data distribution changes, the model adjusts the exploration level. From MABP's perspective, the probability of choosing each arm is updated by the historical result. Meanwhile, concept drift narrows down the choice to a smaller set of arms, acting as a filter. Figure 2 illustrates this concept.

A. Performance signal with multi-armed bandit

The multi-armed bandit problem is a classic example of solving the exploration-exploitation dilemma. The performance of the model can be used as the reward to update the bandit and guiding the bandit to the region of higher performance. Due to the arbitrary nature of the incoming trade flow, we cannot make any well-defined statistical assumptions about the generation of rewards, which corresponds to the adversarial setting [21].

EXP3 (stands for exploration, exploitation, exponentiation) is a widely used algorithm for adversarial bandit [23]. This framework assumes a set of actions (arms) $A = \{a_1, a_2, \dots, a_n\}$ which essentially corresponds to different exploration ratios $\{k_1, k_2, \dots, k_n\}$. At the beginning of each timestamp, the learner must choose an action a_t to minimize the cumulative regret. The regret of each timestamp is defined as the reward difference between the optimal action and the actual action taken.

$$\text{Regret} = \sum_{t=0}^T \max_{a \in A} R(a, B_t) - R(a_t, B_t) \quad (3)$$

The framework maintains a guiding distribution p_t over the set of actions A and uses it to sample the action. The algorithm will receive a reward R (preferably for each action), and the guiding distribution is updated by decreasing the weights of 'bad' actions and exponentially raising the weights of 'good'

actions. Since we only get the reward for the taken action, we use an unbiased estimator [24]:

$$\hat{R}_t(a_i) = \frac{R_t(a_t) \mathbb{1}_{(a_i=a_t)}}{p_t(a_i)}. \quad (4)$$

The distribution is updated by

$$p_{t+1}(a_i) = p_t(a_i) e^{\eta \hat{R}_t(a_i)}, \quad (5)$$

where η is the learning rate. The reward R is usually normalized to be bounded by 1. EXP3 is built to find the best arm on the entire run and it attains an asymptotic regret $O(\sqrt{n})$, where n is the number of rounds.

A variant of EXP3, namely EXP3.S [23], adds a regularization term to the update function to ease arm switches and better adapt to a non-stationary environment. It achieves a similar asymptotic regret bound.

To adapt to our setting, each arm will correspond to an exploration ratio. We use a 21-arm framework corresponding to 21 ratios ranging from 0 to 1 with a step size of 0.05. In order to maintain a robust performance, precision is given as the feedback, such that the exponential-weighted framework could timely respond in the event of a performance drop. The reward in our model is the current precision compared to the weighted averaged precision of all rounds. A discount factor of γ is used to put more weight on recent results. We call this strategy Adaptive Performance Tuning (APT), which is formally presented in Algorithm 1.

B. Measuring concept drift by optimal transport

The second way to decide the exploration rate is by measuring the concept drift. Optimal transport theory is used to measure the distribution difference between data snapshots. Optimal transport aims to find the most efficient way to move mass between distributions [25]. The cost of moving a unit of mass between two positions is called the ground cost, and the objective is to minimize the overall cost of moving one mass distribution to another one. Wasserstein has formulated this problem as the following: [25]

$$W_p(\mu, \nu) = \left[\inf_{\gamma \in \Gamma(x, y)} \int \mathcal{D}(x, y)^p d\gamma(x, y) \right]^{\frac{1}{p}} \quad (6)$$

with μ and ν are the marginal probability distributions of X and Y with realized values of x and y . $\Gamma(x, y)$ denotes the set of all possible joint distributions between X and Y . \mathcal{D} is a distance function. p is the power of the distance, which can be an integer.

Earth mover's distance (EMD) is a measure of the distance between two data distributions. It measures the minimum cost of turning one pile of distribution into the other. To bound EMD to have a specific range of measuring concept drift, we consider how much the optimal EMD value has

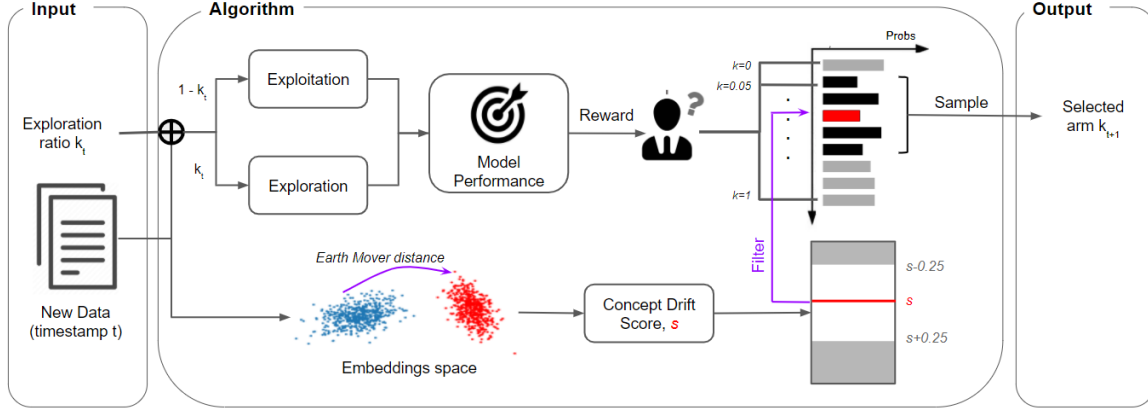


Fig. 2: Detailed workflow of the model. The performance signal helps the multi-armed bandit algorithm to update the probabilities of choosing each arm. Meanwhile, the drift signal suggests the exploration ratio be in a certain range. The arms lying outside the range are eliminated, and the probabilities are normalized for final arm selection.

Algorithm 1 Adaptive Performance Tuning algorithm

Parameters: Learning rate η , randomness $\epsilon \in [0, 1]$ and regularization factor $\alpha > 0$, discount factor γ .

Initialization: $w_0 = [1, 1, \dots, 1]$

for $t = 1, 2, 3, \dots, T$ **do**

1. Set selection probability:

$$p_t(a_i) = \frac{\epsilon}{k} + (1 - \epsilon) \frac{w_t(a_i)}{\sum_i w_t(a_i)}$$

2. Draw an arm a_t according to the probabilities p_t , and get the corresponding precision π_t
3. Calculate weighted average performances of all round, with discount factor γ :

$$\bar{\pi} = \frac{\pi_t + \gamma\pi_{t-1} + \dots + \gamma^{t-1}\pi_1}{1 + \gamma + \dots + \gamma^{t-1}}$$

4. Calculate the reward :

$$R_t = \frac{\pi_t - \bar{\pi}}{\pi_t}$$

5. Update:

$$\hat{R}_t(a_i) = \frac{R_t(a_t) \mathbb{1}_{(a_i=a_t)}}{p_t(a_i)}$$

$$w_{t+1}(a_i) = w_t(a_i) e^{-\eta \hat{R}_t(a_i)} + \frac{e\alpha}{k} \sum_i w_t(i)$$

end

improved from its upper bound, obtained by applying the norm inequality [25]:

$$\begin{aligned} W_1(\mu, \nu) &\leq \inf_{\gamma \in \Gamma(x, y)} \int [\mathcal{D}(x) + \mathcal{D}(y)] d\gamma(x, y) \\ &\leq \int \mathcal{D}(x) d\mu(x) + \int \mathcal{D}(y) d\nu(y). \end{aligned} \quad (7)$$

By dividing the left-hand side by the right-hand side, the concept drift score s is obtained with the range of 0 to 1, allowing us to obtain an interpretable threshold. Especially in an online learning setting, the bounding is necessary since the distribution and values frequently change. The closer the value is to 1, the more statistically different the two data snapshots are.

In the customs setting, we compare the distance between recent historical data and incoming data to measure the concept drift. In our online setting, we use data from the recent four weeks as the validation data. Hence, it becomes handy to use the validation set as the historical data. We encode each import declaration to 16 dimensions by using a transaction encoder [13], and the concept drift is measured by calculating the normalized earth mover's distance between the two embedding sets [26]. Equal probability is assigned to each data point, and the bootstrap technique is used to reduce training time and obtain robust results. The power of the earth mover's distance is chosen as 1, with Euclidean distance as the distance function [27]. The result of normalized earth mover's distance can be thought of as the irreducible cost to turn one distribution to another or the new dataset's novelty from the original dataset and will be used as the ratio for data exploration.

We can use this concept drift score directly as the exploration ratio since if we have more concept drift, there are more patterns to discover and potentially need more exploration. We call this method Adaptive Drift-Aware (ADA) strategy.

C. Adaptive Drift-Aware Performance Tuning algorithm

APT makes its decision from a probability distribution spread over different exploration ratios, while ADA measures the concept drift and puts its bet on a single point in the unit interval. The proposed model, ADAPT (Adaptive Drift-Aware Performance Tuning), makes the most out of two strategies by choosing arms within a more targeted range. More specifically, after getting the concept drift score s , we only consider the

arms corresponding to the exploration ratio inside the interval $[max(0, s - l), min(1, s + l)]$. Only these arms are granted probabilities of being chosen by the APT algorithm, while all other arms outside of the region are given a probability of 0. The probabilities for all arms are then normalized and ready for the next iteration. We can consider this as applying a filter center at s , with $2l$ as the window width. We set l as 0.25 to consider a reasonably wide range of arm candidates.

V. EXPERIMENT

For evaluation, we first review our experiments to answer the following questions:

- Identify the need for exploration (Q1): How much exploration do we need for maintaining a sustainable system?
- Effectiveness of the adaptive strategies (Q2): How well does ADAPT find the ideal exploration ratio?
- Concept drifts analysis (Q3): Does ADAPT capture the concept drift effectively, and does concept drift substantially affect model performance?
- Novel fraud detection (Q4): Is ADAPT effective in detecting frauds among novel trades?
- Ablation analysis (Q5): How important is each component?

A. Settings

1) *Datasets*: For experiments, we employed item-level import declarations from four countries in Africa. The datasets span multiple years allowing us to observe concept drift. The import declarations include the item's free on board (FOB) price, gross weight, quantity, tariff code, importing country, and handler information such as importer, declarant, customs office, and its estimated taxes. The dataset also contains inspected results of whether a transaction is a fraud or not. Due to the data confidentiality policy, we call these countries M, C, N, and T. Table. I shows the statistics of the datasets.

Customs administrations in the four countries conducted nearly 100% of manual inspections of their imported goods. Since the data obtained so far was under a complete inspection, illicitness of the transaction and charged tariffs are accurately labeled at the single-goods level. But this practice is not sustainable, and the customs offices of these countries plan to reduce the inspection rate in the future.

TABLE I: Statistics of the datasets

Datasets	Country M	Country C	Country N	Country T
Periods	2013–2016	2016–2019	2013–2017	2015–2019
# imports	0.42M	1.90M	1.93M	4.17M
# importers	41K	9K	165K	133K
# tariff codes	1.9K	5.5K	6.0K	13.4K
GDP per capita	\$400	\$1,500	\$2,200	\$3,300
Illicit rate	1.6%	1.7%	4.1%	8.2%

2) *Long-term simulation setting*: We conduct experiments to check the ability of our proposed ADAPT method to choose the exploration ratio that maintains the customs selection model in the long run. We used the simulation setting proposed in [3], where a selection model is deployed, updated, and

maintained for multiple years. A limited amount of annotated training data is given to initialize the model. For each timestamp, the method updates the ratio between exploration-exploitation. With the ratio, the model selects a batch of items from the incoming import declarations stream. The custom officers manually inspect these items and their fraud labels are obtained. The model is re-trained with the accumulated training set, and the most recent four weeks of data are used to validate the model. The performance for each timestamp is recorded to evaluate the method. The average performances over the whole period, last two years, last one year, and last six months are additionally reported to summarize the performance.


3) *Evaluation metrics*: If $n\%$ of all import declarations are inspected, the performance of the customs selection model in the online setting is measured by two metrics introduced in [3]: Norm-Precision@ $n\%$ and Norm-Revenue@ $n\%$. They are the precision and revenue evaluated on the selected items, divided by the maximum achievable value. Precision indicates the proportion of fraud cases among the inspected items. Revenue indicates the proportion of revenue secured by examining the set of items, which puts more weight on the lucrative items. The normalization mitigates the change in difficulty of the dataset at different times.

4) *Training Details*: We use XGBoost (GBDT) [11] as our exploitation strategy and random [7] as our exploration strategy. For Exp3.S, we set the learning rate as 3.0, randomness as 0.1, regularization weight as 0.001, and discount factor as 0.9. The inspection rate is 10%, except for country M we use 2%. This is because data M has a low fraud rate and the fraud pattern is relatively easy so that many algorithms reach 100% precision if an inspection rate is kept at 10%. Each experiment is run five times and the averaged results are reported. To smooth out short-term fluctuations and highlight longer-term trends, we show the moving average of 14-weeks in the following figures.

B. Performance Evaluation


1) *Identify the need for exploration (Q1)*: We first identify the need for exploration by simulating the customs selection process in four countries. Assuming that each country faces a different amount of concept drift, we first run hybrid strategies (Sec. III-B) with different amounts of exploration to determine the most appropriate exploration ratio for each country. We record the precision and revenue over time by varying the exploration ratio ranging from 0 to 1. The precision trend is reported in Fig. 3. Results reveal that each country has its own suitable exploration rate. For country C and country T, 20% and 10% of exploration yield the best performance, respectively. Country M and N do not experience significant concept drift, and the full exploitation strategy works the best. Therefore, our main experiment will focus on the dataset with the concept drift, namely country C and T.

In country T, the performance of the full exploitation strategy degrades over time. However, injecting some explorations into the strategy keeps the performance more robust as time



(a) In country T, the performance of the full exploitation strategy drops over time but injecting exploration keeps the performance robust. 10% of exploration yields the best performance.

(b) In country C, 20% of exploration yields the best performance. This setting underperforms full exploitation at the beginning, but the performance improved as time goes.



(c) In country M, fully exploitation yields the best performance.

(d) In country N, fully exploitation yields the best performance.

Fig. 3: Performance of hybrid with different exploration ratios in 4 countries (T, C, M, N). The ratio between exploration and exploitation greatly affects the overall performance in all countries. Adding some exploration helps to improve the performance in data T and C, while fully exploitation strategy works best in data M and N. Results suggest that each country has its own suitable exploration rate.

passes. Model performance is sensitive to the exploration rate. Models with 10–30% exploration achieve the highest precision, while models with 50–70% exploration perform moderately. Fully exploration or fully exploitation strategies performance is worse than any mixtures, with a difference of more than 0.3 at some timestamps.

For country C, full exploitation performed the best at first. However, since 2018, the 20% exploration rate takes the lead in performance, followed by the 10% exploration rate. The strategies with 60% or more exploration perform far worse than others. The best ratio does not stay the same over time but also depends on the period in consideration. The exploration rate of 20% is ultimately the best strategy.

For brevity, we refer to the best performing hybrid in the last six-month period as the *oracle*. Performance statistics of oracle are reported in Table. II.

2) *Effectiveness of the adaptive strategies (Q2)*: Through this experiment, we confirmed how much the risk management system improved when the adaptive exploration rate is determined by the proposed ADAPT method. Precision and revenue are reported in Table. II. Precision trends compared with some baselines are shown in Fig. 4.

In country C, ADAPT follows the similar trend with the fully exploitation model, which was the best performing one until the middle of 2018. From that point, hybrid with 20% exploration (oracle) outperforms all hybrid baselines and fully exploitation strategy. Likewise, ADAPT successfully adapt to the concept drift and eventually follows the same trend and performance of the oracle. According to Table. II, ADAPT was the most effective when all timestamps are considered. In country T, ADAPT has outperformed the oracle in 2016–2017 period. In the next period, ADAPT performed closely to the

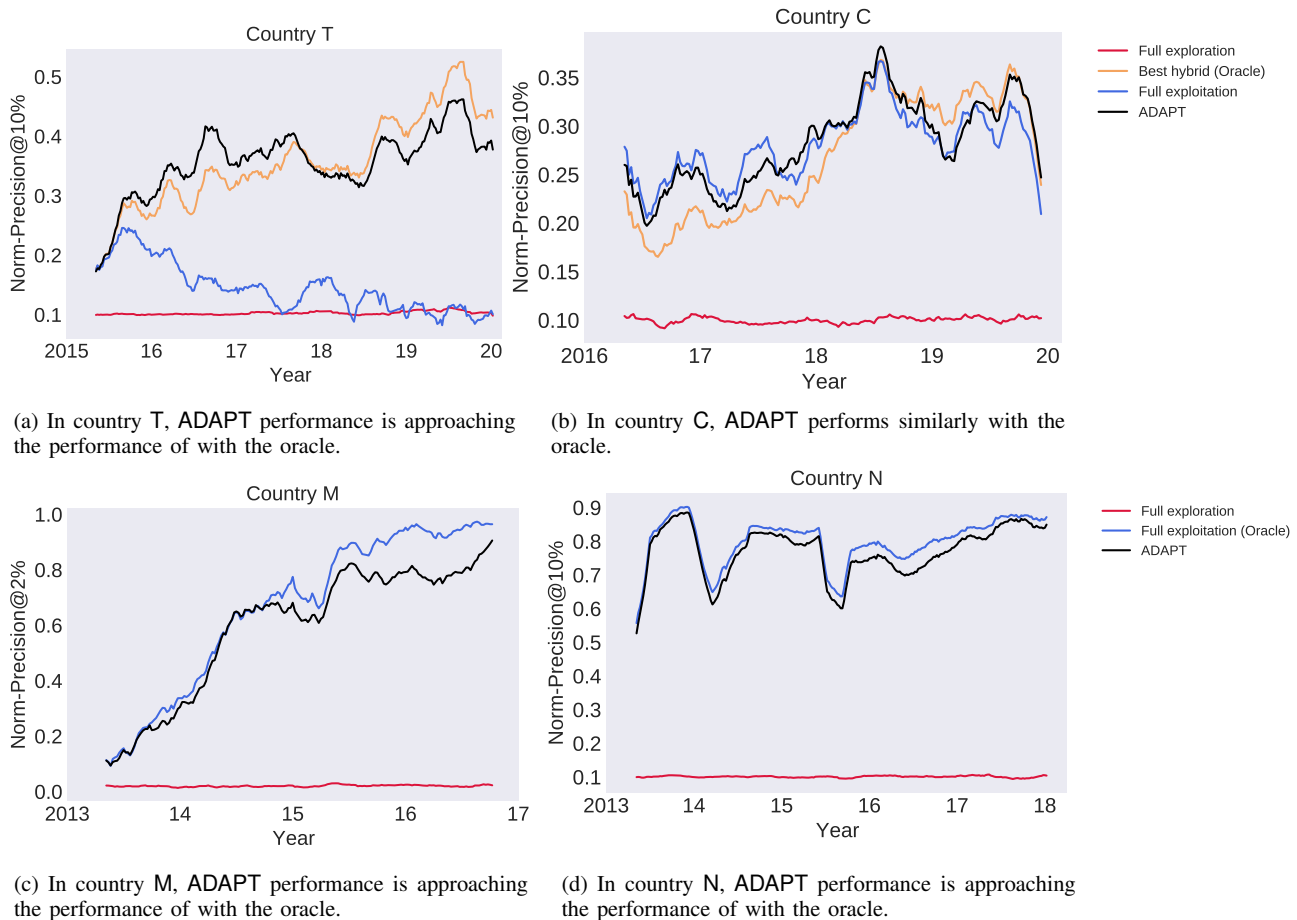


Fig. 4: **Performance of ADAPT in 4 countries (T, C, M, N)** compared to the oracle, fully exploration and fully exploitation. In all countries, ADAPT performance approaches that of the oracle.

oracle with similar trends. On average, ADAPT has similar average precision and revenue with the best hybrid strategy, with more stable performance through more than five years. ADAPT outperforms 9 out of 11 hybrid models. In country M and N, ADAPT's performance approaches that of the best hybrid. Out of 11 hybrids, ADAPT outperform 7 models in country M and 8 models in country N.

From this experiment, we find that ADAPT makes the informed decision for every timestamp so that it achieves comparable performance to the best hybrid method at each point. Compared to the best performing hybrid, oracle, ADAPT even showed higher performance in some time periods. With the accumulated data, it was possible to experiment with multiple hybrids with various setups, but it is infeasible to do this in real-world scenarios. In comparison, ADAPT can adaptively update the ratio based on new data immediately and does not require any hyperparameter tuning, showing a better practice for online setting.

3) *Concept drift analysis (Q3)*: We analyze the concept drift trends of each country. The EMD-based concept drift

score is shown in Figure 5. As evidenced from our analysis of Q1, the amount of concept drift seems to be higher in country T and country C, where injecting exploration boosts the performance compared with full exploitation. In general, our EMD-based concept drift score reflects this trend, with the minimal concept drift scores for country M hovering around 0.15, while the score for T hovers around 0.3. For country C, the full exploitation strategy performs the best before 2018, then the model with 0.2 exploration takes the lead. This suggests that the concept drift starts becoming substantial in 2018. Our score also captures this, with concept drift abruptly soars from around 0.05 before 2018 and plateaus at around 0.2. The EMD-based score for country N peaks at some points then decreases to 0.2.

We conduct further analysis to show the correlation of concept drift score with the model performance. The gist of this analysis is that the performance of the fully-exploitation model will decrease if concept drift occurs. We investigate two datasets with concept drift: T and C (after 2018), and the results are reported in Fig. 6. Pearson correlation test




Fig. 5: Concept drift score measured by optimal transport on the data from four countries.




Fig. 6: A strong negative correlation is observed between concept drift in the data (blue line) and performance of full exploitation model (red line).

reveals a strong negative correlation between the two trends (correlation coefficient for T: -0.55 , for C: -0.23) with very high confidence (p-value for country T: 1.82×10^{-20} , for C: 0.024). This indicates concept drift indeed harms the typical fraud detection model run by exploitation mechanism.

4) *Effectiveness of adaptive strategies in adapting novelties (Q4)*: To evaluate the model’s ability to discover new fraud patterns via exploration, we measure the performance on a subset of items declared by new importers. We investigate two countries having concept drift and report the results in Fig. 7.


ADAPT outperforms fully exploitation in both cases. In country C after 2018, ADAPT secure 39.5% of the revenue on average, compared with 34.6% secured by full exploitation. In country T, ADAPT captures 13.2% of possible revenue, significantly outperform full exploitation which secured merely 1.1% of revenue. In some periods, full exploitation completely fails to pick up any fraud from new importers, while ADAPT still performs robustly. This result supports that ADAPT has explored and inspected the trades declared by new importers, and reused the data successfully to update the model.

5) *Ablation study (Q5)*: To validate the contribution of each component of ADAPT, we conduct an ablation study by examining the performance after removing each component. Precision trends between these three methods are shown in Fig. 4 and the averaged results are in Table. II.


- **ADAPT**: Use both performance signals and concept drift signals.
- **w/o concept drift signal (APT)**: EXP3.S-based method without filtering (ignore input from concept drift scorer)
- **w/o performance signal (ADA)**: Directly use the concept drift score as the exploration rate (ignore performance signal)

In all four countries, the performance of ADAPT is better than APT and closest to the best hybrid strategy, indicating that the concept drift successfully guides the APT arm selector to concentrate on the correct range of ratios.

In country M and country C, precision and revenue of ADAPT are better than ADA. In the other two countries, they perform very similarly. This suggests that the MAB-based model with the performance signal could improve upon ADA by considering neighboring ratios with a history of good



(a) In country T, ADAPT completely dominates full exploitation strategy. Full exploitation completely fails to pick up illicit trades by new importers.



(b) In country C, during duration between 2018 and 2019, two models perform approximately the same. Starting from 2019, ADAPT has risen above GBDT.

Fig. 7: Performance of two strategies on the set of items declared by new importers.

performance. In cases with concept drift or not, ADAPT is expected to be more robust than ADA.

VI. CONCLUSION

This paper examines the importance of the exploitation-exploration ratio and proposes an algorithm, namely ADAPT, that combines two signals for dynamically deciding this value.

Country	Method	Average of Norm-Precision				Average of Norm-Revenue			
		All time	Last 2 years	Last 1 year	Last 0.5 year	All time	Last 2 years	Last 1 year	Last 0.5 year
C	ADAPT	0.2803	0.3106	0.2985	0.2979	0.2829	0.3192	0.3348	0.3660
	ADA	0.2785	0.3061	0.2956	0.2851	0.2768	0.3087	0.3336	0.3556
	APT	0.1806	0.1931	0.1884	0.1799	0.1937	0.2119	0.2386	0.2439
	<i>Full exploration</i>	0.1002	0.1007	0.1020	0.1022	0.0996	0.1030	0.1034	0.1062
T	ADAPT	0.3520	0.3827	0.4154	0.3906	0.4385	0.4848	0.5442	0.5651
	ADA	0.3502	0.3794	0.4127	0.3927	0.4345	0.4801	0.5435	0.5713
	APT	0.3145	0.3214	0.3589	0.3284	0.3934	0.4072	0.4705	0.4692
	<i>Full exploration</i>	0.1023	0.1035	0.1049	0.0998	0.0987	0.0974	0.0967	0.0977
M	ADAPT	0.5873	0.7721	0.8221	0.8703	0.5615	0.7304	0.8087	0.8859
	ADA	0.5001	0.6638	0.6977	0.7557	0.4704	0.6222	0.6860	0.7439
	APT	0.3804	0.5554	0.6413	0.8017	0.3469	0.5147	0.6651	0.8382
	<i>Full exploitation (Oracle)</i>	0.0204	0.0218	0.0203	0.0225	0.0168	0.0167	0.0167	0.0178
N	ADAPT	0.7658	0.7953	0.8435	0.8567	0.7099	0.7355	0.7408	0.7500
	ADA	0.7656	0.7953	0.8431	0.8570	0.7079	0.7317	0.7345	0.7388
	APT	0.5208	0.5098	0.5348	0.5100	0.4787	0.4894	0.4954	0.4644
	<i>Full exploitation (Oracle)</i>	0.1008	0.1015	0.1008	0.1001	0.0999	0.0972	0.0912	0.0871

TABLE II: Precision and revenue of our proposed methods and baselines. Performance over the last k years are averaged and reported.




Fig. 8: Performance of ADAPT with its variants APT and ADA. ADAPT outperforms APT in all countries and showed similar or higher performance than ADA.

By conducting experiments in online settings with four different countries' customs data, we show that ADAPT can perform commensurately with the oracle. Moreover, since the proposed algorithm is based on the shifts in data distribution and the preceding model performance, it does not require extensive tuning and careful result observation over a long period. We also show the importance of concept drift detection in avoiding model failures over time when changes in data distribution occur, which are frequent in customs trades. In the customs setting, the potential concept drift, indicated by ADAPT, can alert customs officers to new fraud patterns that require more cautious inspection. We will open-source the code so that ADAPT can be employed by the customs administrations to ease and improve the trade inspection process.

ACKNOWLEDGMENT

This work was supported by the Institute for Basic Science (IBS-R029-C2, IBS-R029-Y4). We thank World Customs Organization and their partner countries to support their datasets.

REFERENCES

- [1] J. J. Filho, "Artificial intelligence in the customs selection system through machine learning (SISAM)," *Receita Federal do Brasil*, 2015.
- [2] K. Mikuriya and T. Cantens, "If algorithms dream of customs, do customs officials dream of algorithms? a manifesto for data mobilisation in customs," *World Customs Journal*, vol. 14, no. 2, 2021.
- [3] S. Kim, T.-D. Mai, T. N. D. Khanh, S. Han, S. Park, K. Singh, and M. Cha, "Take a chance: Managing the exploitation-exploration dilemma in customs fraud detection via online active learning," *arXiv preprint arXiv:2010.14282*, 2020.
- [4] J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang, "Learning under concept drift: A review," *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, no. 12, pp. 2346–2363, 2020.
- [5] T. Bayoumi, "Changing patterns of global trade," *International Monetary Fund*, 2011.
- [6] J.-Y. Audibert, R. Munos, and C. Szepesvári, "Exploration–exploitation tradeoff using variance estimates in multi-armed bandits," *Theoretical Computer Science*, vol. 410, no. 19, pp. 1876–1902, 2009.
- [7] C. Han and R. Ireland, "Performance measurement of the KCS customs selectivity system," *Risk Management*, vol. 16, no. 8, pp. 25–43, 2014.
- [8] N. Houthby, F. Huszár, Z. Ghahramani, and M. Lengyel, "Bayesian active learning for classification and preference learning," 2011.
- [9] O. Sener and S. Savarese, "Active learning for convolutional neural networks: A core-set approach," in *ICLR*, 2018.
- [10] J. T. Ash, C. Zhang, A. Krishnamurthy, J. Langford, and A. Agarwal, "Deep batch active learning by diverse, uncertain gradient lower bounds," in *ICLR*, 2020.
- [11] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *KDD*, 2016, pp. 785–794.
- [12] J. Vanhoeyveld, D. Martens, and B. Peeters, "Customs fraud detection: Assessing the value of behavioural and high-cardinality data under the imbalanced learning issue," *Pattern Analysis and Applications*, vol. 23, 2020.
- [13] S. Kim, Y.-C. Tsai, K. Singh, Y. Choi, E. Ibok, C.-T. Li, and M. Cha, "DATE: Dual attentive tree-aware embedding for customs fraud detection," in *KDD*, 2020, pp. 2880–2890.
- [14] J. Gama, P. Medas, G. Castillo, and P. Rodrigues, "Learning with drift detection," in *Advances in Artificial Intelligence – SBIA 2004*, A. L. C. Bazzan and S. Labidi, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 286–295.
- [15] N. Lu, G. Zhang, and J. Lu, "Concept drift detection via competence models," *Artificial Intelligence*, vol. 209, pp. 11–28, 2014.
- [16] L. Du, Q. Song, and X. Jia, "Detecting concept drift: An information entropy based method using an adaptive sliding window," *Intelligent Data Analysis*, vol. 18, 03 2014.
- [17] C. Alippi and M. Roveri, "Just-in-time adaptive classifiers — Part I: Detecting nonstationary changes," *Neural Networks, IEEE Transactions on*, vol. 19, pp. 1145 – 1153, 08 2008.
- [18] H. Wang and Z. Abraham, "Concept drift detection for streaming data," in *JCNN*, 2015.
- [19] P. Ball, J. Parker-Holder, A. Pacchiano, K. Choromanski, and S. Roberts, "Ready policy one: World building through active learning," in *ICML*, 2020, pp. 591–601.
- [20] A. Slivkins, "Introduction to multi-armed bandits," 2019.
- [21] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. Schapire, "Gambling in a rigged casino: The adversarial multi-armed bandit problem," *Proceedings of IEEE 36th Annual Foundations of Computer Science*, Aug 1998.
- [22] P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire, "The non-stochastic multiarmed bandit problem," *SIAM Journal on Computing*, vol. 32, no. 1, p. 48–77, 2002.
- [23] R. Allesiardo and R. Feraud, "EXP3 with drift detection for the switching bandit problem," *2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, 2015.
- [24] S. Bubeck and A. Slivkins, "The best of both worlds: Stochastic and adversarial bandits," *JMLR: Workshop and Conference Proceedings vol 23*, 2012.
- [25] C. Villani, *Optimal Transport: Old and new*. Springer-Verlag Berlin An., 2016.
- [26] R. Flamary, N. Courty, A. Gramfort, M. Z. Alaya, A. Boisbunon, S. Chambon, L. Chapel, A. Corenflos, K. Fatras, N. Fournier, L. Gautheron, N. T. Gayraud, H. Janati, A. Rakotomamonjy, I. Redko, A. Rolet, A. Schutz, V. Seguy, D. J. Sutherland, R. Tavenard, A. Tong, and T. Vayer, "POT: Python optimal transport," *Journal of Machine Learning Research*, vol. 22, no. 78, pp. 1–8, 2021.
- [27] N. Bonneel, M. Panne, S. Paris, and W. Heidrich, "Displacement interpolation using lagrangian mass transport," *ACM Transactions on Graphics (TOG)*, vol. 30, no. 6, 2011.