2012 ASE/IEEE International Conference on Social Computing and 2012 ASE/IEEE International Conference on Privacy, Security, Risk and Trust

# The Ones That Got Away: False Negative Estimation Based Approaches for Gold Farmer Detection

Atanu Roy*, Muhammad Aurangzeb Ahmad*, Chandrima Sarkar*, Brian Keegan† and Jaideep Srivastava*
*Department of Computer Science and Engineering
University of Minnesota, Minneapolis, MN 55455
Email: (atanu, mahmad, sarkar, srivasta)@cs.umn.edu
†School of Communication
Northwestern University, Evanston, IL 60208
Email: bkeegan@northwestern.edu

***Abstract*—The problem of gold farmer detection is the problem of detecting players with illicit behaviors in massively multiplayer online games (MMOs) and has been studied extensively. Detecting gold farmers or other deviant actors in social systems is traditionally understood as a binary classification problem, but the issue of false negatives is significant for administrators as residual actors can serve as the backbone for subsequent clandestine organizing. In this paper we address this gap in the literature by addressing the problem of false negative estimation for gold farmers in MMOs by employing the capture-recapture technique for false negative estimation and combine it with graph clustering techniques to determine "hidden" gold farmers in social networks of farmers and normal players. This paper redefines the problem of gold farming as a false negative estimation problem and estimates the gold farmers in co-extensive MMO networks, previously undetected by the game administrators. It also identifies these undetected gold farmers using graph partitioning techniques and applies network data to address rare class classification problem. The experiments in this research found 53% gold farmers who were previously undetected by the game administrators.**

***Index Terms*—False Negative Estimation, Gold Farming, MMO, Graph Partitioning**

## I. Introduction

Virtual worlds constitute a class of online environments where millions of people can share a persistent virtual space and interact with one another. Given the many degrees of freedom accorded to players because of the richness of this domain a large number of behaviours which one observes in the offline world are also observed in these virtual environments [27]; these behaviours can be both positive and negative with sufficient similarity to their offline counterparts. Massively Multiplayer Online Games (MMOs) are a rich class of online games which are analogous to structured virtual worlds. Like the offline world, MMOGs are also used for illicit purposes [8]. One well-known illicit behavior found in MMOGs is gold farming (or real money trade) which involves doing in-game repetitive activities to create virtual assets in the form of virtual items, characters with superior skills, virtual gold etcetra and then selling these virtual assets to legitimate players. Game administrators argue that gold farming upsets the in-game economy, undermines the meritocracy of the game, contributes to anti-social behavior, and leads to compromised and stolen accounts which is why farmers are banned in many games. Gold farming is not a minor economic oddity but has grown to be a massive market estimated to employ hundreds of thousands of workers in China, Southeast Asia, and Eastern Europe while generating US$3 billion annually [20].

The problems associated with catching gold farmers in MMOGs is analogous to identifying criminal activity in the offline world [16]. Consider the ideal case where the justice system follows Blackstone's formulation (better that ten guilty persons escape than that one innocent suffer) and every jailed person was a true criminal (true positive) but some true criminals are not jailed (false negative) to ensure no innocent citizens are jailed (false positive). We argue that this ideal assumption holds in the context of game administrators' banning of gold farmers because the costs of removing false farmers are sufficiently high (e.g., disgruntled customers encouraging their friends to unsubscribe) to incent administrators to minimize false positives. As a result, the set of true positives is very conservative [2]. The primary challenge in criminal justice, gold farming administration, and machine learning generally is to reduce the number of false negatives without adversely increasing the number of false positives. Because there may be similarities in the behavior of the false negatives and true positives, alternative rules could be proposed to identify these similarities. In their initial work Ahmad *et. al.* [2] noted that as much as half of the players which are gold farmers may have evaded detected from the game administrators. However to the best of our knowledge no one has tried to estimate the number of gold farmers evading detection by the game administrators i.e., false negative estimation of gold farmers.

In this paper we use a dataset from an MMO called EverQuest II (EQ2) and address the problem of false negative estimation of gold farmers from a capture-recapture perspective and apply a variety of techniques to study this problem. We use information from the social networks of gold farmers as well their activity patterns to determine which gold farmers

978-0-7695-4848-7/12 $26.00 © 2012 IEEE
DOI 10.1109/SocialCom-PASSAT.2012.125

may have avoided detection. The results that we obtain are in line with previous research in this area [2], [10], [13], [15] but at the same time it opens up new ground in this field which has not been explored before.

To summarize the paper makes the following contributions:

1) It shows detecting gold farmers from MMO networks can be re-defined as a false negative estimation problem and uses capture recapture technique to find a maximum likelihood estimate of the undetected gold farmers.
2) It uses a graph partitioning algorithm over a co-extensive MMO network to trace the undetected gold farmers using the capture recapture technique's maximum likelihood estimates.
3) The paper applies network data to address the rare class classification problem.

The remainder of the paper is organized as follows. Section II provides a list of the related works. Section III discusses in detail the various problems addressed in the paper. Section IV explains the approach used in this research. Sections V and VI are devoted to dataset, experimental setup, results and discussion. The paper is concluded and future works are put forth in section VII.

## II. RELATED WORKS

Gold farming operations within MMOGs operate under many of the same constraints as other offline clandestine organizations like drug traffickers [15], [16]. Game administrators actively try to ban gold farming and rely on reports from other players of farming activity, sting operations undertaken by the game administrators, as well as automatic detection methods [14]. Because of these similarities, the richness of player behavior logged in game databases can be developed into a variety of predictive models which can in turn be mapped back to the less data-rich offline world to predict behavior there [27]. However, it should be noted at the outset that gold farmers also differ from offline clandestine operations given the vastly different penalties for being identified. Farmers who are banned may be temporarily sidelined for a few weeks until they can re-establish the necessary characters, but this relapse rate happens much more rapidly and involves far less risk than many criminal activities which can lead to incarceration or violent reprisal.

One of early studies in gold farmer detection used various classification algorithms to predict the gold farmers already identified by the game administrators. Most of the variables fail to muster a decent accuracy score across all the test except for the age and language of accounts. It was found that the old accounts seldom belonged to gold farmers and most of the gold farmers identified as Chinese speakers [2]. While gold farmers can be distinguished from typical players based these features, they also did not differ substantially from many elite players in terms of the hours played or NPCs killed [2]. This suggested that gold farmers were not a homogenous set of actors engaging in similar behaviors, but employed role differentiation where various groups in the gold farming network would specialize in particular behaviors such as *gatherers* to collect the items, *bankers* to store them, *mules* to traffick them, *dealers* to interact with the customer, and *barkers* to spam their services over channels [10], [16].

Alternative interactions have also been proposed which reflect the complex relationships of game items being associated with particular characters and these characters being associated with particular accounts. Ahmad *et. al.*[4] employed a frequent pattern mining method on hypergraphs of tripartite relationships among housing permissions among characters and character affiliations among player accounts to identify the extent to which gold farmers grant trusted permissions to each other to a greater extent or with different configurations than the general player population. The relationships among the items being exchanged, or *contraband*, also reveals distinctive patterns in the frequency with which certain items are purchased by farmers as well as the tendency for farmers to purchase particular clusters of items. For example Lee *et. al.* [18] used the trade network to identify the bankers whereas Ahmad *et. al.*[3] used the trade network to identify the most commonly purchased and sold items by the farmers. Gold farmers have a tendency to buy many low-end items to craft higher-end items and have a tendency to sell many high-end items which typical players lack the patience to acquire.

Previous work suggested that because gold farming networks under analogous conditions as other clandestine operations, both types of organizations would employ similar structural configurations to evade detection and removal [15], [16]. In particular, networks of exchange among identified gold farmers and drug traffickers both exhibit a tendency for well-connected nodes to be connected through poorly-connected intermediaries rather than to each other. Gold farmers also occupied different positions in the network (as measured by centrality) and relied to a great extent on unlabeled intermediaries who interacted with both farmers and typical players. Employing a naive *guilt-by-association* label propagation heuristic which differentiates any users interacting even once with a gold farmer from the rest of the non-farmer poulation reveals that these affiliates have structural characteristics which are distinct from both identified farmers and the rest of the player population. These affiliates are obvious candidates for being unidentified farmers (false negatives). Ahmad *et. al.* [2] hypothesize that almost half of the gold farmers evade detection owing. Moreover, most of the gold farmers create a new game subscription once their accounts are banned by the game administrators. We extend this prior research on classification of clandestine networks and use statistical measures to estimate this hidden population of gold farmers within MMOGs and propose new methods for identifying them.

Game administrators actively try to ban gold farming accounts due to reasons discussed in [14]. Some of the gold farmers are caught by their fellow players by bringing them to the notice of the administrators whereas a few of the farmers are caught by the sting operations undertaken by the game administrators [14] to rid the game of such nuisance. But Ahmad *et. al.* [2] in their initial work hypothesizes that almost

TABLE I: Confusion matrix of an ideal classifier

| | | Actual Class | |
|---|---|---|---|
| | | *True* | *False* |
| **Predicted Class** | *True* | TP | FP |
| | *False* | FN | TN |

TABLE II: Confusion matrix of a partial classifier

| | | Actual Class | |
|---|---|---|---|
| | | *True* | *False* |
| **Predicted Class** | *True* | TP | FP |
| | *False* | FN + TN | |

half of the gold farmers evade detection. Moreover most of the gold farmers create a new game subscription once their accounts are banned by the game administrators. To the best of our knowledge, for the first time, this research uses statistical measures to estimate the population of gold farmers hidden in the MMO co-extensive networks and uses graph based techniques to identify them.

## III. Problem Statement

### A. False Negative Estimation Problem [1]

In binary classification, a binary classifier classifies each instance of the dataset into one of the two classes {*True, False*}. Given the labels of the dataset, we can further infer whether a sample has been correctly classified or misclassified by the classifier. The confusion matrix in table I shows an ideal confusion matrix. An instance correctly classified as true or false by the classifier increases the frequency of the true positive (TP) or the true negative (TN) cell respectively. If an instance actually belonging to the true class is predicted false by the classifier, it is called a false negative and the vice versa, a false positive. Thus a misclassified instance increases the frequency of either the false positive (FP) or the false negative (FN) cell.

There are instances of real life datasets that do not have the true labels incorporated in them. Due to the absence of the true labels, the results predicted by the classifier accuracy can not be verified. A *gold standard test* is able to assign correct labels to the instances. However it is not always possible to conduct a gold standard test over the entire dataset due to the following reasons: [22]

- skewed class distribution
- high cost associated with labelling data (this typically involves manual labour)
- real time analysis of data (associated with domains where a huge amount of information pours in a short time)

Due to the high number of unlabelled instances and the high cost of performing gold standard test on each of the predicted instances, the users tend to to perform the gold standard test only on the class of interest [24] (generally the rare class in the skewed distribution) and ignore the results predicted by the classifier as belonging to the second class, henceforth referred to as the *false* class. This inability prevents the user from determining the classifier performance of the false class. Table II illustrates the problem where the user is unable to determine the number of instances incorrectly classified false by the classifier and is only aware of the sum of the false negative and the true negative classes. The problem of *false negative estimation* estimates the number of false negative and true negatives from their sum in table II with minimum manual effort.

[1]The materials in this section are derived from classical literature in population estimation. The respective papers are cited in the section.

### B. Gold Farming as a False Negative Estimation Problem

The discussions in the previous sections highlights the fact that the game administrators do not have any automatic gold farmer detection techniques. Most of the techniques used thus far by the administrators rely on considerable human intervention. For example the sting operations by the administrators or fellow players reporting suspicious activities. Once these incidents are brought to notice, the administrators manually verify the accused accounts for suspicious activities. If it is ascertained that the account in question incontrovertibly belongs to a gold farmer the account is banned [2], [26]. The administrators scrutinizing a reported account can be compared to a gold standard test. If we consider the reporting or the sting operations as a classifier, the gold standard test is performed only on those instances which are reported as belonging to the gold farming class, henceforth referred to as the *true* class. The gold standard test helps in differentiating between the true positives and the false positives. But the unreported instances are not subjected to manual scrutiny by the game administrators and hence the false negatives are not dissociated from the true negatives as demonstrated in table II. The gold farmers who have evaded detection are the false negatives. These instances should ideally be labelled as gold farmers (positive). Due to classification error these instances are classified as non gold farmers.

## IV. Approach

### A. Capture Recapture

The first part of our research deals with the estimation of the number of false negatives in a MMO co-extensive network. The problem of estimation of a population from a sample is known as the *missing cases problem* [22]. One of the popular techniques used for population size estimation used in diverse applications is known as the *capture-recapture* technique [6], [7], [12], [17], [19], [21], [22], [28].

The classical example of a capture-recapture method comes from the realm of ichthyology [9]. In this example, the problem is to estimate the number of fishes present in a pond. Capture recapture technique is divided into two steps; the *capture* step and the *recapture* step. In the first step shown in figure 1a, a set of fishes is captured from the pond, marked distinctively and are released into the pond. After enough time has elapsed to ensure that the fishes marked in the first sample has mixed randomly with the entire population, the second phase is embarked upon. The second phase as demonstrated in figure 1b, consists of recapturing a sample of fishes from the pond and examine the number of marked fishes in the new sample. Using the *Petersen estimator* [11], the maximum likelihood estimate for the number of fishes in the pond ($\hat{N}$) can be defined as

$$\hat{N} = \left( \frac{n_{01} \times n_{10}}{n_{11}} \right) \tag{1}$$

The technique is demonstrated in figure 1

In equation 1, $n_{01}$ refers to the total number of fishes caught in the first step, $n_{10}$ refers to the number of fishes caught in the second step and $n_{11}$ refers to the number of fishes caught in both the steps. The underlying assumptions [23] of this method are:

1) The probability of an individual being in one sample is independent of its probability being in the second sample.
2) The probability of all individuals being identified in a list is equal.
3) The population is a closed population which implies no births, no deaths and no migration in the population.

For this study we will assume that the above mentioned assumptions hold.

*1) Estimation of Gold Farmer False Negative population using Capture-Recapture:* Goldberg *et. al.* in [11] provided an intuitive way of applying capture-recapture for false negative estimation. Figure 2 demonstrates the setup for estimating false negatives in a population using capture-recapture method. Table III shows the contingency table which summarizes the variables available and the ones to be estimated. Let $n_{01}$ be the number of true positives after application of the classifier 1 followed by a gold standard test. This is analogous to the capture step of the ichthyology example presented in the last section. Let $n_{10}$ be the number of true positives after the application of classifier 2 followed by a gold standard test. This is analogous to the recapture step of the last example. Let $n_{11}$ be the number true positives common in both the method. The total number of undetected positives ($\widehat{n_{00}}$) in the population can be estimated using equation 2. Thus according to equation 1, $\widehat{n_{00}}$ can be written as

$$\widehat{n_{00}} = \frac{n_{10} \times n_{01}}{n_{11}} \tag{2}$$

Estimated total number of positives present in the population ($\widehat{TotP}$) can be defined as

$$\widehat{TotP} = n_{10} + n_{01} - n_{11} + \widehat{n_{00}} \tag{3}$$

In figure 2, it is assumed that a gold standard test exists which can assign true labels to the instances in a dataset. In the domain of detection of clandestine networks in MMO, various families of classifiers can be used [2] to classify the accounts into belonging to gold farmers (*true*) or not (*false*). Once the classification is performed the *true* instances predicted by the classifier can be verified by the game administrators to differentiate them into true positives and false positives. The last step is analogous to the gold standard test discussed in the previous subsection.

TABLE III: Contingency matrix for FN estimation using capture recapture technique

| | | **Classifier 2** | | |
|---|---|---|---|---|
| | | *True* | *False* | |
| **Classifer 1** | *True* | $n_{11}$ | $n_{01}$ | |
| | *False* | $n_{10}$ | $n_{00}$ | $\Leftarrow$ unknown |

### *B. Label Propagation*

*1) Problem setup:* Zhu *et. al.* in [31] defines the problem of label propagation as follows.

Given a labelled dataset with $X_L = \{x_1, x_2, \ldots, x_m\}$ instances where each instance $x_k$ having a corresponding label $y_k$ associated with it. The set of labels $Y_L = \{y_1, y_2, \ldots, y_m\}$ belongs to $\{1, \ldots, N\}$ class labels. Given a set of unlabelled instances $X_U = \{x_{m+1}, x_{m+2}, \ldots, x_{m+u}\}$, the problem is to predict the labels $Y_U = \{y_{m+1}, y_{m+2}, \ldots, y_{m+u}\}$ of the unlabelled dataset.

*2) Relation to Gold Farmer Problem:* This paper has already discussed about the techniques the game administrators use to ban suspected gold farming accounts. The game administrators banning these accounts is considered as the gold standard test in this research. The instances predicted as gold farmers in the dataset are the real gold farmers who got their accounts banned by the game administrators. Thus the gold farmer label can be considered as the ground truth. It is already discussed that there are undetected gold farmers present in the entire population. Thus the label *non gold farmer* is not entirely a sanguine label. Using the definition of label propagation from the previous subsection, we can assume that the non gold farmer labels are the unlabelled instances whereas the gold farmer labels are the labelled instances.

### *C. Network Based Approach*

The discussion in the previous section have used the statistical method of capture recapture to estimate the population of gold farmers not caught by the existing classifiers. In this section graph based machine learning techniques are proposed to detect the previously undetected gold farmers.

To use graph based algorithm various co-extensive networks that exist in the MMOs is utilized. These vary from game to game but in this research, three networks [14], [18] will be investigated further .

*1) Trading Network:* An important part of the game play is trade between players where they can exchange virtual items within the game for the in-game virtual currency. Thus a trade network can be constructed by creating an edge between two players if they have traded with one another.

*2) Grouping Network:* There are certain activities and quests in the game which are too difficult for individual players to complete while playing solo and thus to complete these activities they have to group together with other players to

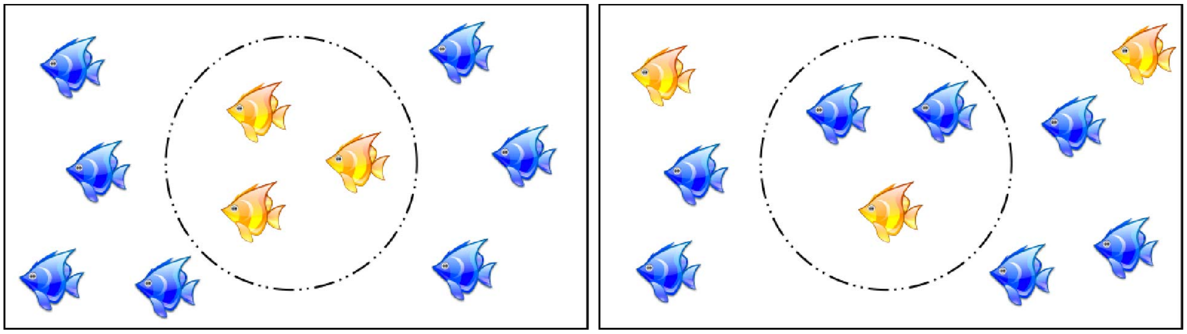

(a) Capture Step (b) Recapture Step

Fig. 1: Illustration of the Capture Recapture technique

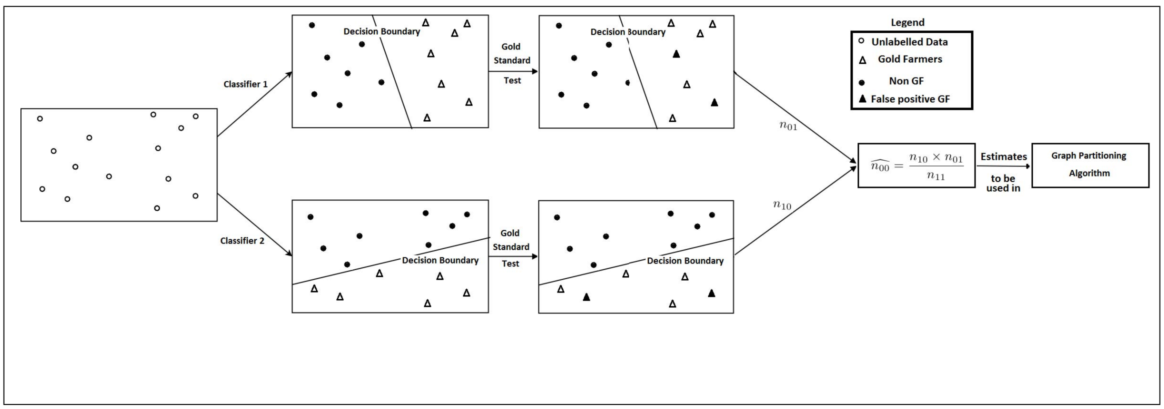


Fig. 2: Flow diagram of Capture Recapture technique for estimation of the gold farmer problem

complete these tasks. The resultant network is the grouping network.

*3) Mentor Network:* Mentoring is an in-game feature where more experienced players can *mentor* less experienced players to get them more familiar with the game. While mentoring the mentor gains achievement points and the apprentice gains experience points at an accelerated rate and thus the relationship is mutually beneficial.

### D. Graph based models

*1) Graph Partitioning:* Given a graph $G = (V, E)$ where $V$ is the set of vertices and $E$ is the set of edges in the graph, the problem of graph partitioning is defined as partitioning the graph $G$ into smaller sub-graphs which fulfil certain desired properties. In this research, a $k$-graph partitioning is used, which is defined as breaking a graph into $k$ sub graphs so that the connections (edge weights) between the sub graphs are minimized. Andreev *et. al.* in [5] proved that the problem of uniform graph-partitioning is NP-complete.

*2) Guilt by Association:* Ahmad *et. al.* in [4] and Keegan *et. al.* in [14] hypothesized that there are certain co-extensive networks in which gold farmers interact heavily with other gold farmers whereas in others they rarely interact with one of their kind. Mentor network is one such network where gold farmers are believed [1] to interact heavily with other gold farmers. The advantages of being a mentor to a fledgling gold farmer includes the inexperienced player learning the game along with its tricks and nuances quicker. But the most valuable commodity is experience for the apprentice which ensure that he levels up much faster. The experienced gold farmer on the other hand also earns vital achievement points for himself.

This research tries to detect the previously undetected gold farmers from the mentor network using the above mentioned assumptions. Keegan *et. al.* captures these assumptions in [14], where they hypothesize that a player who spends the majority of his time socializing with other gold farmers is more likely to be a gold farmer.

In graph partitioning problems, the objective function is to minimize the number of edges between the partitions. Following the assumptions from the previous sections, it can be hypothesized that in a network like the mentor network the gold farmers will form particular clusters or partitions amongst themselves and will remain detached from the bigger population. Application of robust graph partitioning algorithms [25], [29], [30] will ensure that these partitions are identified where the members of a given partition have some sort of natural affinity between one another depending on the application domain. In our current domain this affinity will be captured by the constituent members of some of these partitions being gold farmers. Thus from [14] it follows that if such partitions have a majority of gold farmers as their members then the players who are members of these partitions but are not labelled as gold farmers are in fact gold farmers who have evaded detection.

Capture recapture technique discussed in the previous section provides a maximum likelihood estimate of the possible number of undetected gold farmers. The maximum likelihood estimate thus found is used as a baseline to interpret the results from the graph partitioning algorithm as discussed in the subsequent sections.

*3) Cluto:* The experiments designed for this research uses CLUTO [25], [29], [30] as the graph partitioning algorithm. Cluto is a software package which can partition highly dense and sparse graphs using various types of cluster method and similarity functions. Cluto helps in clustering data sets arising from diverse application ranging from information retrieval to web sciences.

## V. Dataset & Experiments

The dataset used in this research is a gaming log data from EverQuest II developed and maintained by Sony Online Entertainment. EverQuest II is a fantasy based MMO role playing game (MMORPG) released in the last quarter of 2004. The data used in this research is completely anonymized and the records can not be linked to individuals in the real world. The data is collected from a PvE (player versus environment) server "*guk*" over a period of eight months. The gold farmer labels are assigned to those accounts which are banned by the game administrators on suspicion of gold farming activities. The dataset consists of approximately 668,645 unique account identifiers out of which 6,651 are banned on suspicion of gold farming activities constituting roughly 1% of the entire population. This is assumed as the ground truth in this research.

In this research, the mentoring network is used as the medium to identify the gold farmers hidden in the MMO network. Keegan *et. al.* in [14] and Ahmad *et. al.* in [1] hypothesized that the gold farmers tend to mentor other gold farmers so that both can benefit from this symbiotic relation. It was also hypothesized in [4], [14] gold farmers rarely interact with their kind in the trade network.

The mentor network is a relatively small network with 93,115 edges in it for the eight month period in *guk* server. The data available has the unique account identifiers of the mentor, the account identifiers of the apprentice along with their edge weights. The edge weights refer to the number of times a mentor has mentored a particular apprentice. For the hidden gold farmer estimation, using capture recapture technique, only the account identifiers present in the mentor graph was used for analysis. The others were weeded out from the above mentioned unique account identifier list.

### A. Experimental Setup

All of the experiments in this paper are done using Java programming language over flat files and MySQL database. All results are average of 3 runs and were performed on a cold database and buffer cache.

TABLE V: Estimation of $\widehat{n_{00}}$ using various combination of classifiers

| | BayesNet | MLP | kNN | AdaBoost | Jrip | J48 |
|---|---|---|---|---|---|---|
| **BayesNet** | x | 40 | 25 | 24 | 25 | 27 |
| **MLP** | | x | 18 | 31 | 25 | 22 |
| **kNN** | | | x | 37 | 23 | 20 |
| **AdaBoost** | | | | x | 39 | 44 |
| **JRip** | | | | | x | 20 |
| **J48** | | | | | | x |

The mentor network consists of 9,320 unique account identifiers. This population consists of 141 identified and banned gold farmer accounts and the rest non-gold farmer accounts. The first phase of experiments focusses on the problem of false negative estimation with the use of capture recapture technique. The technique is used to infer the false negative population which in the current scenario refers to an estimation of gold farmers who evade detection from the game authorities. The previously mentioned 9,320 unique account identifiers have been divided into training and test set for the classifier. $\frac{2}{3}$rd of the population was assigned to the training set whereas the remaining $\frac{1}{3}$rds was used as a test set. The test set consists of a total of 3,107 unique account identifiers with 41 banned gold farmer accounts.

Using the estimation generated from the technique, graph partitioning algorithm *Cluto* [25] is used to partition the graph into set number of partitions. The algorithm is iterated in order to create everything between 2 and 32 clusters. Out of the 141 gold farmers previously identified by the game administrators, we randomly choose 47 gold farmers ($\frac{1}{3}$rd of the gold farmer population) and tamper their labels to non-gold farmers. This was done on the assumption that graph partition based techniques will be able to recognize these tampered accounts and will be able to re-label them as gold farmers thus providing a performance evaluation of the technique.

## VI. Results & Discussion

### A. Results

False negative estimation techniques have been used in various domains to estimate the size of a population. This paper tries to estimate the population of gold farmers who have evaded detection from the game administrators. Here representative classifiers are selected from each of the popular families of classifiers. Naive Bayes classifier is chosen as a representative of the Bayes family of classifiers. Multi Layered Perceptron (MLP) is chosen from the function based family of classifiers, *k* nearest neighbour (kNN) from the non parametric family, AdaBoost as a meta learner. JRip and J48 are chosen to represent the families of rule based and tree based classifiers respectively.

The first experiment estimates the false negatives in the network using a pair of classifiers. Every combination $\binom{6}{2}$ of the above mentioned classifiers are used. The detailed

TABLE IV: Performance of various classifiers in finding the *True Positive* values. Individual cells in table is represented as a triple $\langle a, b, c \rangle$ where $a$ denotes the true positive rate of the classifier mentioned in the first column of the table, $b$ denotes the true positive rate of the classifier mentioned in the first row of the table. The true positive instances that were classified correctly by both the classifiers is denoted by $c$

| | BayesNet | MLP | kNN | AdaBoost | JRip | J48 |
|---|---|---|---|---|---|---|
| **BayesNet** | x | 20, 14, 7 | 20, 10, 8 | 20, 11, 9 | 20, 14, 11 | 20, 16, 12 |
| **MLP** | | x | 14, 10, 8 | 14, 11, 5 | 14, 14, 8 | 14, 16, 10 |
| **kNN** | | | x | 10, 11, 3 | 10, 14, 6 | 10, 16, 8 |
| **AdaBoost** | | | | x | 11, 14, 4 | 11, 16, 4 |
| **JRip** | | | | | x | 14, 16, 11 |
| **J48** | | | | | | x |

TABLE VI: Estimation of total positive $\left(\widehat{TotP}\right)$ instances in the population using various combination of classifiers

| | BayesNet | MLP | kNN | AdaBoost | Jrip | J48 |
|---|---|---|---|---|---|---|
| **BayesNet** | x | 67 | 47 | 46 | 48 | 51 |
| **MLP** | | x | 34 | 51 | 45 | 42 |
| **kNN** | | | x | 55 | 41 | 38 |
| **AdaBoost** | | | | x | 60 | 67 |
| **JRip** | | | | | x | 39 |
| **J48** | | | | | | x |

results are shown in table IV. Individual cells in table IV are represented as a triple $\langle a, b, c \rangle$ where $a$ denotes the true positive rate of the classifier mentioned in the first column of the table henceforth referred to as classifier 1. $b$ denotes the true positive rate of the classifier mentioned in the first row of the table and henceforth referred to as classifier 2 with reference to figure 2. The true positive instances that were classified correctly by both the classifiers are denoted by $c$. Comparing the triple with equation 2, returns $a = n_{01}, b = n_{10}$ and $c = n_{11}$. Using the values from the triple and substituting them in equation 2, $\widehat{n_{00}}$ is estimated. The average of the false negative estimates over three runs is demonstrated in table V.

Equation 3 estimates the total number of positive instances in the population $\left(\widehat{TotP}\right)$. Using the results from table IV & table V and substituting them in equation 3, it is possible to estimate the total number of positive instances (gold farmers) present in the population, which are subsequently calculated in table VI.

Now that a maximum likelihood estimate of the total number of gold farmers present in the network is obtained, the second set of experiment uses graph partitioning algorithms to identify the individual gold farmers who evaded detection from the game administrators. This experiment proposes to evaluate the performance of graph partition based approaches. As discussed in the previous section, labels of $\frac{1}{3}$rd of the gold farmer population have been changed to non-gold farmers, henceforth known as the transformed gold farmers. This experiment tries to find those gold farmers whose label has been changed manually to non gold farmers.

Figure 3 demonstrates the results of graph partitioning on the mentor network. The *x-axis* denotes the number of partitions in the graph, which is an input to the algorithm. The *y-axis* denotes the number of gold farmers predicted by our

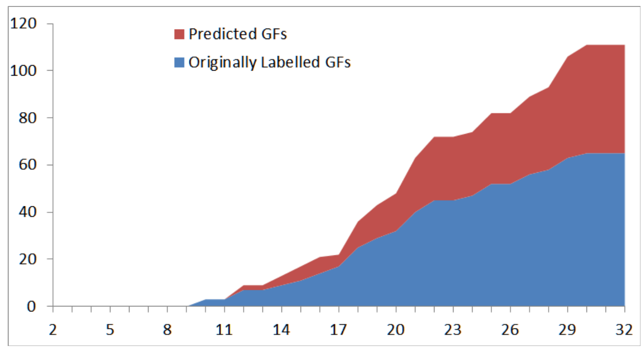


Fig. 3: Original and predicted gold farmers in mentor network after application of graph partitioning algorithms. $x-axis$ denotes the number of partitions whereas $y-axis$ denotes the total number of gold farmers, originally thus labelled as such and also the ones predicted by the approach.

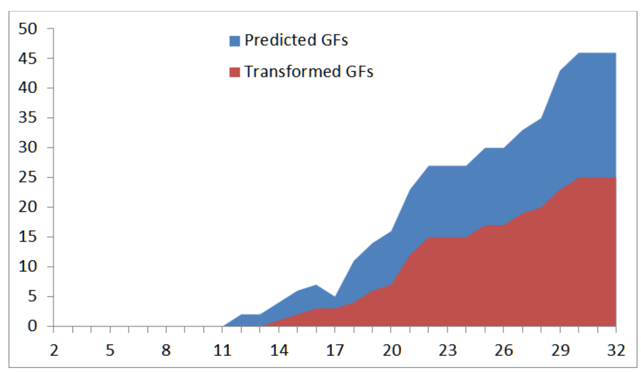


Fig. 4: Predicted and transformed gold farmers in mentor network after application of graph partitioning algorithms. $x-axis$ denotes the number of partitions whereas $y-axis$ denotes the total number of gold farmers, predicted by our approach and also the transformed gold farmers. Transformed gold farmers refer to those accounts which were originally labelled as gold farmers but have been changed to non gold farmers by us for this experiment.

TABLE VII: Original and predicted gold farmers in mentor network along with the number and percentage of label transformed gold farmers discovered using this approach

| Partitions | Labelled GFs | Predicted GFs | Transformed GFs found | Percentage Discovered |
|---|---|---|---|---|
| 2 | 0 | 0 | 0 | 0 |
| 3 | 0 | 0 | 0 | 0 |
| 4 | 0 | 0 | 0 | 0 |
| 5 | 0 | 0 | 0 | 0 |
| 6 | 0 | 0 | 0 | 0 |
| 7 | 0 | 0 | 0 | 0 |
| 8 | 0 | 0 | 0 | 0 |
| 9 | 0 | 0 | 0 | 0 |
| 10 | 3 | 0 | 0 | 0 |
| 11 | 3 | 0 | 0 | 0 |
| 12 | 7 | 2 | 0 | 0 |
| 13 | 7 | 2 | 0 | 0 |
| 14 | 9 | 4 | 1 | 2.13 |
| 15 | 11 | 6 | 2 | 4.26 |
| 16 | 14 | 7 | 3 | 6.38 |
| 17 | 17 | 5 | 3 | 6.38 |
| 18 | 25 | 11 | 4 | 8.51 |
| 19 | 29 | 14 | 6 | 12.8 |
| 20 | 32 | 16 | 7 | 14.9 |
| 21 | 40 | 23 | 12 | 25.5 |
| 22 | 45 | 27 | 15 | 31.9 |
| 23 | 45 | 27 | 15 | 31.9 |
| 24 | 47 | 27 | 15 | 31.9 |
| 25 | 52 | 30 | 17 | 36.2 |
| 26 | 52 | 30 | 17 | 36.2 |
| 27 | 56 | 33 | 19 | 40.4 |
| 28 | 58 | 35 | 20 | 42.6 |
| 29 | 63 | 43 | 23 | 48.9 |
| 30 | 65 | 46 | 25 | 53.2 |
| 31 | 65 | 46 | 25 | 53.2 |
| 32 | 65 | 46 | 25 | 53.2 |

approach. A partition is identified as a gold farmer partition if at least half the members of the partition are originally labelled as gold farmers. Once a gold farmer partition is identified, all the members in the partition are marked as gold farmers. The second and the third attribute in table VII shows the cumulative count of originally labelled gold farmers detected across all the *gold farmer partition* and the newly detected gold farmers in those partitions. The transformed gold farmer accounts which are the ones that are chosen randomly from the mentor network and subsequently received a label change to non gold farmers are pitted against the total number of predicted gold farmers in figure 4. As discussed in the experimental setup, the labels of random $\frac{1}{3}$rds (or 47) of the gold farmers are changed to non gold farmers. The 4th & the 5th column in table VII shows the cumulative count and percentage of those accounts identified using this technique.

*Number of partitions* is the input parameter that has been varied in this experiment. It has been varied from 2 partitions to 32. Creating a single partition will be a trivial solution since it will encompass the whole graph. Thus the experiment is designed to start from 2 partitions. After 32 partitions a majority of the partitions is formed of very small number of instances which are dominated by the non gold farmers. Moreover it is found that there is no considerable increase in the detected gold farmer population once we increase the number of partitions beyond 32.

### B. Discussion

This research addresses the problem of rare class classification using statistical measures and network data. The statistical measures estimate the total undetected rare class population in the entire population whereas graph based techniques implemented on co-extensive networks detect those undetected rare class instances.

Table V demonstrates, for all combination of classifiers, the false negative estimate technique estimates an average of 28 false negative instances per classifier combination. Adding the estimates with the true gold farmers predicted by the classifier results in the estimation of the total number of actual positives present in the network. As demonstrated in table VI, the average estimation of actual positives present in the network for all possible combination of the six previously mentioned classifier is approximately $45$. Thus the false negative estimation technique estimates that $60\%$ of the gold farmer population evades detection. This agrees with the original hypothesis put forward by Ahmad *et. al.* in their initial research [2] that more than half the population of gold farmers evade detection from the authorities. Note these set of experiments were performed with a third of the actual mentor network containing $3107$ unique account identifiers and $41$ banned gold farmer accounts. Classifiers which follow similar techniques for classification has a high $\frac{n_{01}}{n_{11}}$ ratio, indicating that the estimates provided by it may not be accurate. For example, classifiers like JRip and J48 falls in this category and it can be seen that for both the classifiers the $\frac{n_{01}}{n_{11}}$ ratio is very high. Whereas classifiers whose classification algorithms differ from each other and use very different techniques does not suffer from this deficiency.

Hierarchical clustering algorithm was used in this research since one is able to control the number of clusters formed in this approach.

In the prior sections, it was mentioned that the labels of a third of the account marked gold farmers are changed to non gold farmers. The second experiment (table VII) shows that at 30 clusters, $CLUTO$ is able to detect approximately $53\%$ of these gold farmers and report them. While identifying a partition as a *gold farmer partition*, the primary requisite for our approach is, at least half of the account identifiers in the partition should belong to gold farmer class. $50\%$ is a heuristic used for this research based on the hypothesis by Ahmad *et. al.* in [2] that half the gold farmers evade detection. This paper hypothesizes that the rest of the accounts detected as gold farmers are those accounts who have evaded detection by th game administrators, but we do not have ground truth to verify our claims. Ideally we would have liked a dataset where detection of gold farmers is done in two phases and thus we could have used the data from the second set of detection as our validation set.

## VII. Conclusion & Future Work

The problem of gold farmer detection have been previously studied by a number of researchers. In the context of classification all the previous works, to the best of our knowledge have assumed the labels associated with an account as the ground truth and based their classification models on this assumption. But it is known [2], [14] that more than half the population of gold farmers evade detection. This research further delves into this hypothesis and proposes to estimate and identify these gold farmers who evade detection. Well known statistical measure, capture recapture is used to estimate the undetected gold farmer population. To identify these undetected instances graph partitioning algorithm is used over a co-extensive MMO network.

This research has treated the classification of gold farmers as a binary classification problem, where an account either belongs to a gold farmer or not. In reality as Keegan et. al. [14] and Lee *et. al.* [18] suggested there are multiple types of gold farmer and each group shows high level of expertise in certain metrics. It will be exciting to consider these groups as separate labels and use appropriate co-extensive networks to identify them. One more exciting direction of research points to relabelling the the instances identified by this research as gold farmers and use the modified graph for further investigations. The false negative estimates found using the first part of the research can act as a maximum limit to the number of instances that are relabelled. This research only uses hierarchical clustering as the only partitioning algorithm to detect partitions in the co-extensive networks. One can also use different partitioning algorithms like community detection techniques and graph cut algorithms. The inability to control the number of clusters in these approaches discouraged us from using these algorithms in this research.

## Acknowledgment

The research reported herein was supported by the Air Force Research Laboratory by contract number FA8650-10-C-7010. The data used for this research was provided by Sony Online Entertainment. We also thank members of the *Data Mining and Management Research Group* at the University of Minnesota, Department of Computer Science & Engineering for their support and feedback.

## References

[1] M.A. Ahmad, D. Huffaker, J. Wang, J. Treem, D. Kumar, M.S. Poole, and J. Srivastava, *The many faces of mentoring in an mmorpg*, Social Computing (SocialCom), 2010 IEEE Second International Conference on, IEEE, 2010, pp. 270–275.

[2] M.A. Ahmad, B. Keegan, J. Srivastava, D. Williams, and N. Contractor, *Mining for gold farmers: Automatic detection of deviant players in mmogs*, Computational Science and Engineering, 2009. CSE'09. International Conference on, vol. 4, IEEE, 2009, pp. 340–345.

[3] M.A. Ahmad, B. Keegan, S. Sullivan, D. Williams, J. Srivastava, and N. Contractor, *Illicit bits: Detecting and analyzing contraband networks in massively multiplayer online games*, Privacy, Security, Risk and Trust (PASSAT), 2011 IEEE Third International Conference on and 2011 IEEE Third International Confernece on Social Computing (SocialCom), IEEE, 2011, pp. 127–134.

[4] M.A. Ahmad, B. Keegan, D. Williams, J. Srivastava, and N. Contractor, *Trust amongst rogues? a hypergraph approach for comparing clandestine trust networks in mmogs*, Proceedings of the 5th International AAAI Conference on Weblogs and Social Media, 2011, pp. 10–17.

[5] K. Andreev and H. Räcke, *Balanced graph partitioning*, Proceedings of the sixteenth annual ACM symposium on Parallelism in algorithms and architectures (New York, NY, USA), SPAA '04, ACM, 2004, pp. 120–124.

[6] A. Chao, *Estimating the population size for capture-recapture data with unequal catchability*, Biometrics (1987), 783–791.

[7] D.G. Chapman, *Some properties of the hypergeometric distribution with application to zoology census*, University of California Publications in Statistics **1** (1951), 131–160.

[8] M. Consalvo, *Cheating: Gaining advantage in videogames*, The MIT Press, 2007.

[9] J.N. Darroch, *The multiple-recapture census: I. estimation of a closed population*, Biometrika **45** (1958), no. 3/4, 343–359.

[10] A. Fujita, H. Itsuki, and H. Matsubara, *Detecting real money traders in mmorpg by using trading network*, Seventh Artificial Intelligence and Interactive Digital Entertainment Conference, 2011.

[11] J.D. Goldberg and J.T. Wittes, *The estimation of false negatives in medical screening*, Biometrics (1978), 77–86.

[12] E. B. Hook and R. R. Regal, *Capture recapture methods in epidemiology: methods and limitations*, Epidemiology Reviews **17** (1995), no. 2, 243–264.

[13] H. Itsuki, A. Takeuchi, A. Fujita, and H. Matsubara, *Exploiting mmorpg log data toward efficient rmt player detection*, Proceedings of the 7th International Conference on Advances in Computer Entertainment Technology, ACM, 2010, pp. 118–119.

[14] B. Keegan, M.A. Ahmad, S. Sullivan, D. Williams, J. Srivastava, and N. Contractor, *Untangling dark webs:theories, methods, and models for a computational social science of clandestine networks*, IEEE Third International Confernece on Social Computing (SocialCom), IEEE, 2011.

[15] B. Keegan, M.A. Ahmed, D. Williams, J. Srivastava, and N. Contractor, *Dark gold: Statistical properties of clandestine networks in massively multiplayer online games*, Social Computing (SocialCom), 2010 IEEE Second International Conference on, IEEE, 2010, pp. 201–208.

[16] B.C. Keegan, M.A. Ahmed, D. Williams, J. Srivastava, and N. Contractor, *Sic transit gloria mundi virtuali? promise and peril in the computational social science of clandestine organizing*, (2011).

[17] R.E. LaPorte, D. McCarty, G. Bruno, N. Tajima, and S. Baba, *Counting diabetes in the next millennium: application of capture-recapture technology*, Diabetes care **16** (1993), no. 2, 528–534.

[18] E. Lee, J. Lee, and J. Kim, *Detecting the bank character in mmorpgs by analysis of a clustered network*, The 3rd International Conference on Internet (2011).

[19] S.M. Lee and A. Chao, *Estimating population size via sample coverage for closed capture-recapture models*, Biometrics (1994), 88–97.

[20] V. Lehdonvirta and M. Ernkvist, *Converting the virtual economy into development potential: Knowledge map of the virtual economy. infodev*, The World Bank, Washington, DC (2011).

[21] S. Mane, J. Srivastava, S.Y. Hwang, and J. Vayghan, *Estimation of false negatives in classification*, Data Mining, 2004. ICDM'04. Fourth IEEE International Conference on, IEEE, 2004, pp. 475–478.

[22] S.V. Mane, *False negative estimation: Theory, techniques and applications*, ProQuest, 2008.

[23] L. Papoz, B. Balkau, and J. Lellouch, *Case counting in epidemiology: limitations of methods based on multiple data sources*, International Journal of Epidemiology **25** (1996), no. 3, 474–478.

[24] R.M. Pfeiffer and P.E. Castle, *With or without a gold standard*, Epidemiology **16** (2005), no. 5, 595.

[25] M. Rasmussen and G. Karypis, *gCLUTO: An Interactive Clustering, Visualization, and Analysis System*, Tech. report, TR# 04020, Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN.

[26] R. Thawonmas, Y. Kashifuji, and K.T. Chen, *Detection of mmorpg bots based on behavior analysis*, Proceedings of the 2008 International Conference on Advances in Computer Entertainment Technology, ACM, 2008, pp. 91–94.

[27] D. Williams, *The mapping principle, and a research framework for virtual worlds*, Communication Theory **20** (2010), no. 4, 451–470.

[28] J.T. Wittes, T. Colton, and V.W. Sidel, *Capture-recapture methods for assessing the completeness of case ascertainment when using multiple information sources*, Journal of chronic diseases **27** (1974), no. 1, 25.

[29] Y. Zhao, G. Karypis, and U. Fayyad, *Hierarchical clustering algorithms for document datasets*, Data mining and knowledge discovery **10** (2005), no. 2, 141–168.

[30] Ying Zhao and George Karypis, *Evaluation of hierarchical clustering algorithms for document datasets*, Proceedings of the eleventh international conference on Information and knowledge management (New York, NY, USA), CIKM '02, ACM, 2002, pp. 515–524.

[31] X. Zhu and Z. Ghahramani, *Learning from labeled and unlabeled data with label propagation*, Tech. report, Technical Report CMU-CALD-02-107, Carnegie Mellon University, 2002.