Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media

Modeling Destructive Group Dynamics in Online Gaming Communities

Akshay Patil[1], Juan Liu[2], Bob Price[2], Hossam Sharara[3], and Oliver Brdiczka[2]
[1] Stony Brook University, Stony Brook, NY
[2] Palo Alto Research Center, Palo Alto, CA
[3] University of Maryland, College Park, MD

Abstract

Social groups often exhibit a high degree of dynamism. Some groups thrive, while many others die over time. Modeling destructive dynamics and understanding whether/why/when a person will depart from a group can be important in a number of social domains. In this paper, we take the World of Warcraft game as an exemplar platform for studying destructive group dynamics. We build models to predict if and when an individual is going to quit his/her guild, and whether this quitting event will inflict substantial damage on the guild. Our predictors start from in-game census data and extract features from multiple perspectives such as individual-level, guild-level, game activity, and social interaction features. Our study shows that destructive group dynamics can often be predicted with modest to high accuracy, and feature diversity is critical to prediction performance.

Introduction

As the complexity of online activity has increased, formal group structures have come to play an increasingly important role in the experience and effectiveness of an individual's online life. While formal groups can be seen in many task-oriented online communities, the use and impact of formal groups in online role playing games is among the most highly developed at this time. In these games, players join groups known as guilds, some of which include hundreds of members, to share resources, plan strategy and execute large scale attacks on opposing forces. The effectiveness of these groups can be undermined when group members depart, taking with them, experience, resources and possibly other group members. The ability to predict imminent departure of group members and the probable impact of their departure is highly desirable, as it offers insights on factors that affect online group effectiveness. It also provides practical guidance to tasks such as risk management and customer retention.

Instability in on-line role-playing game groups has been well documented (Ducheneaut et al. 2006; 2007). Specific factors identified as contributing to member departures include guild leadership style (Williams et al. 2006), disengagement from the online community itself (Yee 2006), and internal conflicts between members. Most previous works (Ducheneaut et al. 2006) (Williams et al. 2006) (Yee 2006) have investigated guild membership dynamics using qualitative case studies and small sample surveys of selected players. These studies have provided us with many qualitative insights, but have not yet reached the level of practical mathematical predictors that can be deployed online to make predictions based on actual game data.

Automated analysis has been applied to the discovery of networks that support undesirable activities such as gold farming (Keegan et al. 2011) (Heeks 2008) and the trade of contraband items (Ahmad et al. 2011a). Network analysis has also be used to investigate trust among players (Ahmad et al. 2011b). We also exploit properties of the player's interaction with the community to make predictions, however we are unaware of prior work specifically addressing the prediction of departures from groups and the estimation of associated damage. Perhaps the closest work to ours is the research on churn prediction in MMOGs (Borbora et al. 2011). Player churn occurs when players stop playing the game (a more extreme event than switching guilds). The authors employ machine learning classifiers directly on game data to predict game departure. In churn prediction, the player's overall game satisfaction (as measured by achievement) is the key driver. In contrast, after a guild departure, the player is often still highly engaged in the game and continues play with a different guild. The motivational structure is therefore quite different. As a consequence, the features for prediction are also different. Churn detection primarily rely on game achievement features, while our guild quitting prediction weighs heavily on features regarding the quality of social interactions within a specific guild.

In addition to identifying players likely to quit their guild, we also wish to predict when the quitting event will happen. Models such as the dynamic influence model (Asavathiratham et al. 2001; Dong et al. 2007; Pan et al. 2011) have been used to represent changing patterns of social interactions. This Bayesian model uses a detailed description of the conditional dependence between each player's current state at time t and the previous states of all players at time $t-1$. While powerful, the model is computationally expensive, and requires detailed modeling of social interactions between actors making it difficult to scale to large networks.

Statistic on server Eitrigg	
Number of Characters	51,224
Number of Guilds	2906
Number of Edges	2,447,577
Average Collaboration Time (hrs.)	1.73 ± 1.09
% Characters changing Guild	26.53

Table 1: Overall Network Statistics

The paper is organized as follows. It starts with an overview of WoW data and guild-quitting statistics, followed by a brief investigation of whether social interaction plays any role in guild-quitting decisions. We then move on to define the impact of an imminent quitting event. The later sections present predictive models to predict (1) the potential impact and (2) if and when a quitting event will happen. We conclude with a discussion of the important factors for predictions and an outline of future work.

Though the work presented in this paper is focused on the specific problem of guild quitting dynamics in WoW, we hope that similar ideas and approaches can be generalized to other social domains. We note the recent work on unfollow dynamics in Twitter (Kivran-Swaine, Govindan, and Naaman 2011; Kwak, Chun, and Moon 2011), where social features are investigated with respect to tie breaking events. The key difference is that unfollow is between dyads (pairwise relationship) while we focus on group dynamics. Group destruction is a common phenomenon in real world as well as in virtual world, for instance, an employee quitting job in a corporate environment may often have co-workers following him/her to join a new company. Likewise, market studies have shown that customer loyalty to a product can be weakened if the customer's close friend opt out to a new product. It is desirable to gain the insights on social group dynamics and the capability of prediction will be highly valuable.

Overview of WoW Data and Guild-quitting Events

To explore guild quitting dynamics, we use data from a previous WoW study (Ducheneaut et al. 2007). A web-based crawler was deployed to log in-game activities based on the API specified by Blizzard Entertainment, the producer of WoW. The crawler periodically issues "/who" requests every 5 to 15 minutes, depending on server load, to get a list of characters currently being played on a given server. Over six months of data are logged, from November 2010 to May 2011. The data is sometimes referred to as the WoW census. Three types of servers are logged: player-vs-environment (PvE), player-vs-player (PvP), and role playing (RP). The servers may present players with different game tasks, but are otherwise identical in terms of game organization and support. Overall we observed more than 470,000 unique characters forming over 15000 guilds, scattered on three servers: Eitrigg (a PvE server), Cenarion Circle (a RP server), and Bleeding Hollow (a PvP server).

Social interaction may be an important influencing factor in guild-quitting events. First, we define a friendship network among guild members, where nodes are characters, and edges indicate co-occurrence within gaming zones — if two characters were observed in the same game location (zone in WoW), an edge is added between the corresponding nodes. One legitimate concern for social analysis is that in WoW a single player can have multiple characters, and social interaction among the characters belonging to the same player may be different than characters from different players. However, this concern is eliminated by the construction of the co-occurence network, because only one character can be logged on from a player at any given point of time (except rare work-arounds). We use the co-occurrence network as the platform of our study of social interactions, with the underlying assumption that if characters co-occur in a gaming zone, it is highly likely that the characters are collaborating on a gaming activity. Two possible limitations are noted: (1) there are some gaming zones not necessarily associated with any gaming activity, for instance, characters are often left "AFK" (Away from keyboard) in the game's main cities before or at the end of a play session. In this case, the geographic proximity does not necessarily reflect any kind of joint activity. In our data logger, we remove such ambiguous zones from the co-occurrence criteria. (2) Characters may co-occur by chance. This is treated as noise in the social network graph. The basic assumption is that with large amount of accumulated gaming data, the ties between characters driven by real social interaction will dominate.

Secondly, we add a membership network to indicate the affiliation between characters and guilds. Nodes fall into two categories: (1) guild nodes, and (2) character nodes. If a character is observed appearing in a guild, an affiliation edge is added. The overall network is the super-imposition of the friendship and the affiliation networks. It is an undirected multi-graph i.e. it allows for multiple edges between any two nodes in the network.

The social network graph above has a temporal dimension: each edge has a weight indicating the duration of social interaction, and is tagged with a timestamp. We use a graph summarization approach as described in (Sharan and Neville 2008) to simplify representation into temporal snapshots. Given link weights $W_1, W_2 \ldots W_t$, the exponential kernel can be computed as follows,

$$W_t^S = \begin{cases} (1-\theta)W_{t-1}^S + \theta W_t & \text{if } t > t_0 \\ \theta W_t & \text{if } t = t_0 \end{cases} \tag{1}$$

where t_0 is defined as the initial time. The exponential kernel weighs the recent past highly and decays the weight exponentially as time passes, with the decaying rate defined by the parameter θ. The summarized graph will be used for feature computing in the later sections.

Table 1 lists some statistics in the raw social network on Eitrigg. Guild quitting events are fairly common — around 26% of characters quit from a guild at least once in our observation period. Similar guild quitting statistics are observed on Cenarion Circle and Bleeding Hollow.

Are Quitting Events Correlated?

A common perception in WoW and other MMOGs is that social interactions between players influence a player's decisions about joining and quitting guilds. When a quitting

Statistic on server Eitrigg	
Number of Guilds	2906
Number of Guilds with > 30 quits	181 (6.23%)
Guilds following Poisson	23 (12.71%)
Guild NOT following Poisson	158 (87.29%)

Table 2: Guild Quitting Events Analysis

event happens, the quitting player may pull friends out of the original guild. In this section, we examine the data formally to test this hypothesis. In WoW, the question to ask is, does data suggest that social interaction leads to correlated guild departures, or that on the contrary, departures are independent? This is the "sanity-check" question that should be addressed before diving into the social group analysis. If quitting events are independent, then social group analysis will be irrelevant and should not be the topic of our study.

To test the correlated guild-quitting hypothesis, we model the quitting events as a Poisson process, a well-known probabilistic model for discrete events arrivals. Under the following assumptions, an arrival process is Poisson (Stark and Woods 1994): (a) the probability that one arrival occurs between t and $t + T$ is proportional to T (the proportion λ is known as the arrival rate); (b) the number of arrivals in non-overlapping intervals are statistically independent; (c) the probability of two or more arrivals is negligible when T is small. Assumptions (a) and (c) are intuitive and reasonable, and we are interested in validating assumption (b) via a Poisson test. If the guild-quitting events are Poisson, then (b) is validated. The other way is also true: if guild-quitting events are non-Poisson, then under assumptions (a) and (c), this implies that guild-quitting events are dependent.

It is known that, if the events follow a Poisson process, the inter-arrival time, defined as the time between two consecutive events, would be exponentially distributed with probability density function

$$f(T) = e^{-\lambda T}, \tag{2}$$

where λ is the arrival rate that can be estimated from the observation data. This gives us a method to test the validity of Poisson models. To determine whether the probabilistic model (2) fits the observation data, we use the standard chi-square goodness-of-fit test (Nikulin 1973). To ensure statistical reliability, we limit our analysis to guilds with 30 or more departures. A guild is considered Poisson if the model fits with a confidence of 0.95 or higher. Table 2 summarizes the result of goodness-of-fit test. A large fraction (87%) do *NOT* follow a Poisson Process. This indicates that quitting events are not independent, but rather correlated. Factors that may have caused the correlation are explored in the later sections.

Potential Damage of a Quitting Event

In WoW, the decision to quit a guild is likely influenced by social factors — unhappy players convince friends to join them in quitting, or a circle of friends reaching consensus to leave as a group. Modeling the details of this influence

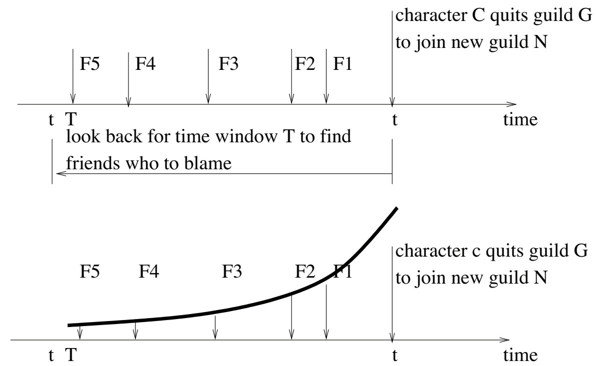

Figure 1: Damage assessment

process is difficult without direct observation of player interactions. We adopt an abstract model in which a character's quitting decision is influenced by all preceding quitting events among his or her friends (as defined by the social graph) in proportion to the time elapsed. The immediate corollary is that a player who quits is responsible for, or shares the blame for every friend that subsequently quits the guild. We use the term *damage score* to refer to a character's aggregate share of the blame for subsequent quitting events. Essentially, damage score is an empirical notion of impact.

Figure 1 illustrates the computation of damage score. For a given quitting event (character C quitting guild G at time t to join new guild N), we first look back in time to identify friends who may have caused this quitting event. In Figure 1a, we note that a number of C's friends $(F1, F2, F3, F4, F5)$ all quit G to join N recently, hence C may be under their influence. We then assign a "blame" score $b_{F,C}$ to the friends $F \in \mathcal{F} = \{F1, F2, F3, F4, F5\}$. Quitting events in the recent past receives high blame, while quitting events in the distant past is assumed to have little impact and receives little blame. In Figure 1b, $F1$ receives the highest blame, while $F5$ receives the lowest. Mathematically we use a normalized exponentially decay function:

$$b_{F,C} = \frac{e^{\alpha(|t_F - t|}}{\sum_{i \in \mathcal{F}} e^{\alpha(|t_i - t|}}, \tag{3}$$

where the parameter α controls how fast the blame function decays over time. The denominator normalizes the blame score so that the total blame across $\mathcal{F}$ sums up to 1.

Given the blame assignment, we now compute the damage of any given character X's quitting as the sum of blame scores X receives from all its followers $C \in \mathcal{C}$, i.e.,

$$d_X = \sum_{C \in \mathcal{C}} d_{X,C} \tag{4}$$

For simplicity in implementation, $\mathcal{C}$ is defined over a time window T after X's quitting. Going further is unnecessary, as the quitting events are too far apart to assign or receive substantial blame. In our experiment, we use $\alpha = 1.0$ and $T = 15$ days.

In our observation data, the overall damage statistics is the following: among 23014 quitting events documented in the

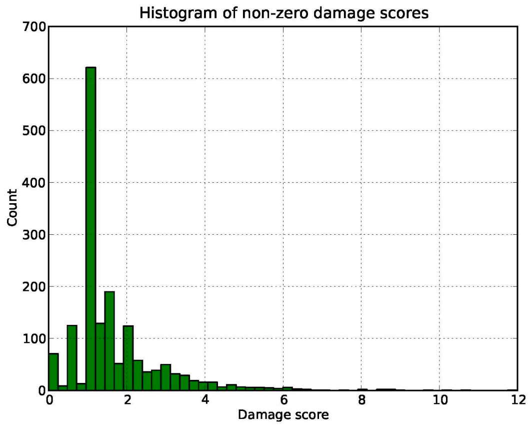

Figure 2: Histogram of non-zero damage scores

WoW Eitrigg server (PvE), only 1700 quitting events have a non-zero damage score. This indicates that most quitting events are isolated. Figure 2 shows the histogram of non-zero damage scores. Overall statistics are as follows: mean= 0.1232, standard deviation = 0.5581, min = 0, and max= 11.9911. Now the question is, from the observation data of a character's activity in WoW census, can we predict the potential damage if he/she leaves the guild?

Predicting Potential Damage

Features

For damage prediction, we take a supervised learning approach, where a feature set is generated and a mapping from the feature set onto the potential damage score is learned from a training set. Damage is inflicted by a character's quitting event (say, X quitting guild G) to the guild. Several categories of features may be relevant: (1) individual player X's profile, (2) the guild G's profile, (3) game activity of X, in the WoW space as well in the guild G, and (4) social interactions and structural importance. The features are listed in Table 3.

Most feature names are self-explanatory. A few are explained below. For guild profile features (the second feature block in the table), two clustering coefficients are evaluated: the clustering coefficient of the guild, and the clustering coefficient of the entire friendship network. Clustering coefficient is a topological concept, measuring the degree to which nodes in a graph tends to be clustered together. Formally it is measured as the ratio between the number of closed triplets over the number of connected triples of vertices. It takes value 0 in a star topology and value 1 if the graph is fully connected. Through these features, we would like to investigate whether structure balance has an impact on potential damages of quitting events.

In the game statistics features (the third feature block in the table), in addition to playing time, our method logs collaboration time, the time spent by character X in playing with other characters (in the entire WoW space and within the guild). In addition, we have a few features indicating the playing style of a given character. Collaboration coefficient is defined as the ratio between the collaboration time and the overall playing time. A high collaboration coefficient indicates a social playing style while a low value implies a "lone wolf" player. Loyalty coefficient is defined as the ratio of X's time in guild G and X's total playing time. This measure indicates how loyal X has been to the guild G (prior to quitting).

Social features measures the relative importance of the given character X in the game space and the guild. Some friends have a lot of interactions. "Played with excessively" here is defined as those with collaboration time exceeding 2 standard deviations above the mean collaboration time among all pairs of characters. These are considered close friends. Weighted degree is a measure of centrality. It is computed from the temporally-summarized graph. We would expect that more central X is, more damaging his/her departure would be.

Feature Importance

Table 3 also contains the correlation coefficient between each feature and the damage score (the third column). The correlation coefficient can be considered as a rough measure of the feature's importance in predicting damage. If the damage is independent from a feature (hence the feature is useless to prediction), the correlation coefficient will be 0. Sign of the correlation indicates when a feature is positively correlated (large feature value implies large damage) or negatively correlated (otherwise). Significant correlation coefficients (with absolute value exceeding 0.25) are marked in bold-face fonts.

Among individual player X's profile features, character level appears important. High level players are more likely to cause significant damage. This is not surprising. The number of guilds that X has joined prior to the quitting event may be related to X's willingness to join new guilds, but it does not appear to be indicative of how much damage X's departure may cause to the original guild.

Guild profile features do not seem to be very strong features. Both size and topology (measured by the clustering coefficients) seems to be weak features.

Game activity features are of high importance. Total playing time, in the entire WoW space and within the guild, are both strongly positively correlated with damage. This agrees with the intuition — the more X plays, the more important he/she becomes in the guild, and hence X quitting the guild is likely to cause more damage. Collaboration time in the WoW space and within the guild are also strongly positively correlated with damage. The relative ratios, collaboration coefficient and loyalty coefficient, are not as strong. This can also be justified — being loyal to the guild simply means it is hard to quit, however it does not necessarily imply low damage.

Social features are very important. The centrality measures (weighted degree features, number of friends, etc) are strongly correlated with damage. This agrees with our intuition that the departure of a central node can be very damaging because it can cause a snowballing effect. Furthermore, the number of friends of X who have already quit the guild

Feature category	Feature description	Corr coeff with damage score	Corr coeff with binary damage label
individual profile	character level	0.1557	0.1591
	number of guilds joined	-0.0423	-0.0457
guild profile	guild size	0.0499	0.0558
	clustering coefficient in guild	0.0523	0.0669
	overall clustering coefficient	0.0369	0.0442
game stats	playing time	**0.2651**	**0.3022**
	playing time within guild	**0.3581**	**0.4081**
	collaboration time	**0.2849**	**0.3274**
	collaboration time within guild	**0.3470**	**0.3991**
	collaboration coefficient	0.1601	0.1747
	collaboration coefficient in guild	0.1393	0.1521
	loyalty coefficient	0.0770	0.0871
social features	number of friends	0.2383	**0.2752**
	number of friends already having quit guild	**0.3918**	**0.4436**
	percentage of guild members played with	0.1789	0.1925
	percentage of guild members played with excessively	0.1781	0.1800
	weighted degree	**0.2673**	**0.3189**
	weighted degree in guild	**0.2618**	**0.3091**

Table 3: Correlation coefficient between features and damage score. Significant correlation coefficients (with absolute value exceeding 0.25) are marked in bold-face fonts.

is very important. This may be justified from a social psychology perspective – if many of X's friends have already quit the guild, then there is a tendency that more friends will join in the quitting streak. This boosts up the potential, and when X quits, the snowballing effect can accelerate.

Prediction of Damage Significance

We first formulate damage prediction as a classification problem — is X's quitting going to cause substantial damage to its original guild G? The class label in this case is binary: substantial if the potential damage score exceeds a pre-determined threshold (with value 1 in our experiment) and non-substantial otherwise. The last column of Table 3 lists the correlation coefficient between the features and the class label. Qualitatively it does not differ much from the previous column. Game activity and social interaction remain the strongest feature categories.

Server	accuracy	Precision	Recall	F-meas.
Eitrigg	82.50%	0.825	0.825	0.825
Cenarion Circle	81.23%	0.812	0.823	0.813
Bleeding Hollow	79.9%	0.799	0.799	0.799

Table 4: Damage classification. Results are reported on 10-fold cross validation over the training set.

One problem to note is that the observation data set contains an overwhelmingly large non-substantial damage class (93% of all samples) and a small substantial damage class. Out of the 23014 quitting instances, only around 1700 has a non-zero damage value. If we randomly select training samples, the predictive model will overfit in the small damage value interval, and may fail to predict high damage quitting events. To avoid overfitting, our approach takes a balanced sampling approach to generate a training set containing roughly an equal amount of samples from each class.

To predict the damage class label, we have experimented with a number of classification methods. Tree-based classification methods work well. Table 4 reports classification results using a Weka implementation (Hall et al. 2009) of the random forest method, which builds a library of decision trees from random feature subsets and predicts class labels by voting or averaging. The overall classification accuracy is around 80%for all WoW servers. This is significantly better than random guessing (with around 50% accuracy). This result indicates that our feature set is predictive of future damage, and the classifier can be used to predict damage labels reliably.

Regression For Damage Prediction

Damage prediction can also be formulated as a regression problem, i.e., constructing a mapping from the feature set to the continuous damage value. Regression performance is measured using common metrics such as mean absolute error (MAE) and mean squared error (MSE). The correlation coefficient between the predicted value and the true target value can also be considered as a metric of accuracy, with high value indicating good prediction performance.

Table 5 lists regression results using a Weka implementation (Hall et al. 2009) of Bagging method (Hastie, Tibshirani, and Friedman 2001), which builds a library of regression trees and averages their prediction results. For Eitrigg, the MAE is 0.6836 (in a dynamic range of 0 to 11.99). and the MSE is about 1 (in a dynamic range of 0 to 143.7). This indicates that regression accuracy is quite reasonable. Similar regression accuracy has been observed on Cenarion Circle and Bleeding Hollow as well.

We conclude that the feature set combining individual profile, guild profile, game activity statistics, and social interactions can support damage prediction, and that our approach can predict damages associated with quitting events with reasonable prediction and classification accuracy.

Server	MAE	MSE	Corr Coef
Eitrigg	0.6836	1.0830	0.7046
Cenarion Circle	0.6032	0.9140	0.6729
Bleeding Hollow	0.6853	1.0798	0.6495

Table 5: Regression from feature set to damage (evaluated on the training set via 10-fold cross validation)

Predicting Guild Quitting Events

In this section, we examine the feasibility of predicting if and when a character might quit his or her guild. Prediction of guild quitting behavior has both practical and theoretical value. From a theoretical point of view, identification of factors that predict guild quitting events gives us insight into factors affecting guild stability. From a practical point of view, we can build a useful detector for identifying characters who might potentially quit and abscond with guild resources or otherwise damage the effectiveness of the guild. We might then be able to intervene before adverse outcomes are experienced.

The prediction problem is formulated as the following: given the game trace up to current time, predict whether a character will quit from the guild within a specified future interval. In WoW, there are two notions of time: (1) calendar time, measured in hours and days, and (2) game time, measured as the count of game sessions and events. Here we unify these two notions by grouping game events within a guild by day and call it a "day record". Usually there is one record per day, but if the player switches guilds, there may be more than one day record for a specific calendar date, and if the player does not play at all, there will be no day record. The day record notion is flexible and can accommodate variation in playing style, ranging from players with only occasional game activities to more serious players. Based on this notion, we define a prediction window in terms of day records (7 in our experiment). If this window of future day records includes observations of the character in another guild, we set the class label *target_guild_has_changed* to true to indicate a quitting event has occurred. This class label is our prediction target.

Features

Prediction of quitting event is inherently dynamic. Game activities unfolds as time advances, and prediction should be updated accordingly. To accommodate the dynamic nature, we keep a running window of features. A personal history window of 14 day records is used to generate features summarizing the character's playing history, as shown in Figure 3. The features are listed in Table 6 in italicized fonts in the first feature block. Basically, these features are designed to measure game engagement and achievement within the history window. The guild count represents the character's past tendency to quit his or her guild. The character level and level change attributes are intended to capture the character's sense of progress. The underlying theory being that a character that is making progress will be content with their current guild. The number of game events in the window and duration of the window are designed to help the classifier address special cases that occur for many characters when we are predicting at the beginning of their histories. Average event duration is designed to capture how dense the character's play is. Time since last event measures the gap since the last historical game session.

Feature	Corr. coeff
guild_count	0.183
time_since_last_event	0.064
event_count	-0.056
level_begin	-0.051
level_end	-0.050
level_change	0.041
avg_event_duration	0.014
window_duration	0.001
numberOfGuildMembership	0.170
clusteringCoefficientWithinGuild	-0.151
playingTimeWithinGuild	-0.122
collaborationTimeWithinGuild	-0.107
weightedDegreeWithinGuild	-0.105
overallClusteringCoefficient	-0.096
overallWeigtedDegree	-0.089
playingTime	-0.089
collaborationTime	-0.089
percentageOfMembersPlayedWithExcessively	-0.082
percentageOfMembersPlayedWithBefore	-0.078
numberOfFriends	-0.062
numberOfGuildMembers	-0.025
numberOfFriendsAlreadyHavingQuitGuild	0.007

Table 6: Correlation between features and quitting events. Personal history features are shown italicized, and social history features are in regular fonts.

From a social psychology point of view, a person's decision to leave a guild may be primarily due to the unsatisfaction towards the guild. Hence prediction should not only involve individual character's personal history, but also features regarding the guild, especially social relationship with guild members. Table 6 contains these features as well (second feature block, in regular fonts). Details regarding these features are explained in the last section and hence are omitted here. The social features are computed from the temporally summarized graph using an exponentially decaying summarization kernel. This is equivalent to a "fading" social history, as illustrated in Figure 3.

Table 6 sorts the two feature categories according to relative feature importance, measured as the absolute value of correlation coefficient. *Guild_count* (and equivalently the number of guild membership) seems to be the strongest feature. This is intuitive — past switching events are indicative of a character's lack of loyalty, and implies an easy tendency to depart from his/her current guild as well. The feature *time_since_last_event* is positively correlated with quitting, i.e., the longer a player stays activity free, the more likely he/she will quit the guild. Among the social history features, clustering coefficient (measuring structure balance), playing time and collaboration time (measuring engagement within the guild), and weighted degree (measuring importance) are all negatively correlated with guild departure. This is also in-

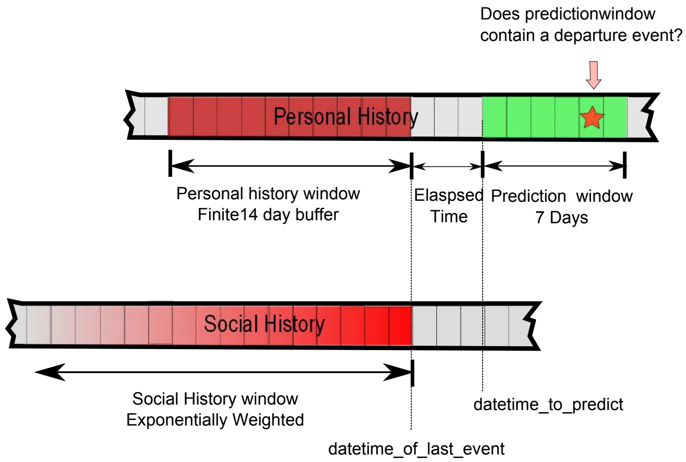

Figure 3: Personal and social history windows and the prediction window

tuitive. A person heavily engaged and playing a central role in a well-balanced guild is less likely to leave the guild.

We note the related work in (Borbora et al. 2011) which primarily relies on individual features to detect churn events, and the influence model (Asavathiratham et al. 2001) which only considers social context. The novelty of our approach here is that we have identified factors from both categories and their relative importance. As Table 6 shows, important features comes from both categories and should be unified for prediction. We have tried Weka's linear forward greedy feature selection to find an informative feature subset. The subset returned contains two personal features (guild_count, time_since_last_event) and five social features (playingTimeWithinGuild, collaborationTimeWithinGuild, numberOfGuildMembers, numberOfGuildMembership, clusteringCoefficientWithinGuild).

Method and Results

In this prediction problem, the prediction events for a server are not completely independent. Successive predictions for a given character will be correlated. This can lead to overfitting where the system learns to identify specific characters who quit instead of generic features of characters which are useful for predicting quitting events. To guard against this, we developed distinct training and test sets. History and prediction windows in the training set are generated from a different set of characters than the history and prediction windows in the test set.

Guild quitting prediction is further complicated by the unbalanced class problem. In most of the observation snapshots, characters continue to play in their current guild, and quitting events are rare by comparison. To avoid overfitting the classifier to non-quitting events, we use a random sampling method to balance the training set so that it contains approximately equal proportions of non quitting and quitting events. Classification results were similar for a number of prediction models. Table 7 reports the classification performance for a random forest with 10 trees and unlimited depth and feature counts. Guild quitting prediction classifiers are built separately for 3 WoW servers: Eitrigg, Cenarion Circle, and Bleeding Hollow. Classification performance does not vary much on servers, indicating that our approach is generally applicable in the WoW space.

We conclude that it is possible to predict with modest precision and recall and that past loyalty, guild stability, and social engagement are key predictive factors for quitting.

Conclusion

Our analysis has shown that destructive group dynamics models can be constructed from WoW in-game census data. We have build predictors to predict (1) the potential impact of an imminent guild departure and (2) if and when the depart will happen in a prediction window. The predictors have reasonably accuracy (best predictors in the 80-90% accuracy range). It is important to choose features from multiple perspectives, such as features regarding the individual, guild features, game activity statistics, and social interaction and topological features. Combining diverse features is essential to the predictor performance.

Our future work will likely include extending group dynamics modeling and prediction to constructive dynamics to understand how a player joins a guild and how a guild grows. More importantly, we would also like to extend the group dynamics analysis to other social groups, for instance, real-world social interactions, and online social networks.

Acknowledgment We would like to acknowledge contributions from our colleagues: Nicholas Ducheneaut, Nick Yee, and Les Nelson (PARC) for providing WoW game census data and insights about game dynamics, Jie Gao (Stony Brook University) for help on social network analysis, and Jianqiang Shen and Rui Zhang for algorithmic suggestions.

References

Ahmad, M.; Keegan, B.; Sullivan, S.; Williams, D.; Srivastava, J.; and Contractor, N. 2011a. Illicit bits: Detecting and analyzing contraband networks in massively multiplayer online games. *Proceedings of SocialCom-11*.

Ahmad, M.; Keegan, B.; Williams, D.; Srivastava, J.; and Contractor, N. 2011b. Trust amongst rogues? a hypergraph approach for comparing clandestine trust networks in mmogs. *Proceedings of Fifth International AAAI Conference on Weblogs and Social Media (ICWSM 2011)*.

Asavathiratham, C.; Roy, S.; Lesieutre, B.; and Verghese, G. 2001. The influence model. *Control Systems, IEEE* 21(6):52–64.

Borbora, Z.; Hsu, K.; Srivastava, J.; and Williams, D. 2011. Churn prediction in mmorpgs using player motivation theories and ensemble approach. *Proceedings of SocialCom-11*.

Dong, W.; Lepri, B.; Cappelletti, A.; Pentland, A. S.; Pianesi, F.; and Zancanaro, M. 2007. Using the influence model to recognize functional roles in meetings. *Proceedings of ICMI'07* 271–278.

Server	Overall	Non Quitting Event			Quitting Event		
	Accuracy	Precision	Recall	F Measure	Precision	Recall	F Measure
Eitrigg	82.7432	0.878	0.926	0.901	0.389	0.268	0.317
Cenarion Circle	89.0973	0.917	0.967	0.941	0.342	0.164	0.222
Bleeding Hollow	79.8396	0.855	0.91	0.881	0.396	0.276	0.325

Table 7: Results on balanced training set tested on disjoint character IDs

Ducheneaut, N.; Yee, N.; Nickell, E.; and Moore, R. 2006. Alone together? exploring the social dynamics of massively multiplayer online games. *Proc. of CHI 2006* 407–416.

Ducheneaut, N.; Yee, N.; Nickell, E.; and Moore, R. J. 2007. The life and death of online gaming communities: A look at guilds in world of warcraft. *Prof. of CHI*.

Hall, M.; Frank, E.; Holmes, G.; Pfahringer, B.; Reutemann, P.; and Witten, I. 2009. The weka data mining software: an update. *ACM SIGKDD Explorations Newsletter* 11(1):10–18.

Hastie, T.; Tibshirani, R.; and Friedman, J. 2001. *The Elements of Statistical Learning*. Springer.

Heeks, R. 2008. Current analysis and future research agenda on gold farming: Real world production in developing countries for the virtual economies of online games. *Institute for Development Policy and Management, University of Manchester*.

Keegan, B.; Ahmad, M.; Williams, D.; Srivastava, J.; and Contractor, N. 2011. What can gold farmers teach us about criminal networks? *ACM Crossroads* 17(3):11–15.

Kivran-Swaine, F.; Govindan, P.; and Naaman, M. 2011. The impact of network structure on breaking ties in online social networks: unfollowing on Twitter. *CHI*.

Kwak, H.; Chun, H.; and Moon, S. 2011. Fragile online relationship: A first look at unfollow dynamics in Twitter. *CHI*.

Nikulin, M. S. 1973. Chi-squared test for normality. *Proc. of International Vilnius Conference on Probability Theory and Mathematical Statistics* 2:119–122.

Pan, W.; Dong, W.; Cebrian, M.; Kim, T.; and Pentland, A. 2011. Modeling dynamical influence in human interaction. *MIT Technical Report TR-661*.

Sharan, U., and Neville, J. 2008. Temporal-relational classifiers for prediction in evolving domains. In *Data Mining, 2008. ICDM'08. Eighth IEEE International Conference on*, 540–549. IEEE.

Stark, H., and Woods, J. W. 1994. *Probability, Random Processes, and Estimation Theory for Engineers,* second edition. Upper Saddle River, NJ: Prentice Hall.

Williams, D.; Ducheneaut, N.; Xiong, L.; Zhang, Y.; Yee, N.; and Nickell, E. 2006. From treehouse to barracks: The social life of guilds in world of warcraft. *Games and Culture* 1(4):338–361.

Yee, N. 2006. The labor of fun: How video games blur the boundaries of work and play. *Games and Culture* 1(1):68–71.