

Measurement and Classification of Humans and Bots in Internet Chat

Steven Gianvecchio, Mengjun Xie, Zhenyu Wu, and Haining Wang
Department of Computer Science
The College of William and Mary
{*srgian, mjxie, adamwu, hnw*}@cs.wm.edu

Abstract

The abuse of chat services by automated programs, known as chat bots, poses a serious threat to Internet users. Chat bots target popular chat networks to distribute spam and malware. In this paper, we first conduct a series of measurements on a large commercial chat network. Our measurements capture a total of 14 different types of chat bots ranging from simple to advanced. Moreover, we observe that human behavior is more complex than bot behavior. Based on the measurement study, we propose a classification system to accurately distinguish chat bots from human users. The proposed classification system consists of two components: (1) an entropy-based classifier and (2) a machine-learning-based classifier. The two classifiers complement each other in chat bot detection. The entropy-based classifier is more accurate to detect unknown chat bots, whereas the machine-learning-based classifier is faster to detect known chat bots. Our experimental evaluation shows that the proposed classification system is highly effective in differentiating bots from humans.

1 Introduction

Internet chat is a popular application that enables real-time text-based communication. Millions of people around the world use Internet chat to exchange messages and discuss a broad range of topics on-line. Internet chat is also a unique networked application, because of its human-to-human interaction and low bandwidth consumption [9]. However, the large user base and open nature of Internet chat make it an ideal target for malicious exploitation.

The abuse of chat services by automated programs, known as *chat bots*, poses a serious threat to on-line users. Chat bots have been found on a number of chat systems, including commercial chat networks, such as AOL [15, 29], Yahoo! [19, 25, 26, 28, 34] and MSN [16],

and open chat networks, such as IRC and Jabber. There are also reports of bots in some non-chat systems with chat features, including online games, such as World of Warcraft [7, 32] and Second Life [27]. Chat bots exploit these on-line systems to send spam, spread malware, and mount phishing attacks.

So far, the efforts to combat chat bots have focused on two different approaches: (1) keyword-based filtering and (2) human interactive proofs. The keyword-based message filters, used by third party chat clients [42, 43], suffer from high false negative rates because bot makers frequently update chat bots to evade published keyword lists. The use of human interactive proofs, such as CAPTCHAs [1], is also ineffective because bot operators assist chat bots in passing the tests to log into chat rooms [25, 26]. In August 2007, Yahoo! implemented CAPTCHA to block bots from entering chat rooms, but bots are still able to enter chat rooms in large numbers. There are online petitions against both AOL and Yahoo! [28, 29], requesting that the chat service providers address the growing bot problem. While on-line systems are besieged with chat bots, no systematic investigation on chat bots has been conducted. The effective detection system against chat bots is in great demand but still missing.

In the paper, we first perform a series of measurements on a large commercial chat network, Yahoo! chat, to study the behaviors of chat bots and humans in on-line chat systems. Our measurements capture a total of 14 different types of chat bots. The different types of chat bots use different triggering mechanisms and text obfuscation techniques. The former determines message timing, and the latter determines message content. Our measurements also reveal that human behavior is more complex than bot behavior, which motivates the use of entropy rate, a measure of complexity, for chat bot classification. Based on the measurement study, we propose a classification system to accurately distinguish chat bots from humans. There are two main components in our

classification system: (1) an entropy classifier and (2) a machine-learning classifier. Based on the characteristics of message time and size, the entropy classifier measures the complexity of chat flows and then classifies them as bots or humans. In contrast, the machine-learning classifier is mainly based on message content for detection. The two classifiers complement each other in chat bot detection. While the entropy classifier requires more messages for detection and, thus, is slower, it is more accurate to detect unknown chat bots. Moreover, the entropy classifier helps train the machine-learning classifier. The machine learning classifier requires less messages for detection and, thus, is faster, but cannot detect most unknown bots. By combining the entropy classifier and the machine-learning classifier, the proposed classification system is highly effective to capture chat bots, in terms of accuracy and speed. We conduct experimental tests on the classification system, and the results validate its efficacy on chat bot detection.

The remainder of this paper is structured as follows. Section 2 covers background on chat bots and related work. Section 3 details our measurements of chat bots and humans. Section 4 describes our chat bot classification system. Section 5 evaluates the effectiveness of our approach for chat bot detection. Finally, Section 6 concludes the paper and discusses directions for our future work.

2 Background and Related Work

2.1 Chat Systems

Internet chat is a real-time communication tool that allows on-line users to communicate via text in virtual spaces, called chat rooms or channels. There are a number of protocols that support chat [17], including IRC, Jabber/XMPP, MSN/WLM (Microsoft), OSCAR (AOL), and YCHT/YMSG (Yahoo!). The users connect to a chat server via chat clients that support a certain chat protocol, and they may browse and join many chat rooms featuring a variety of topics. The chat server relays chat messages to and from on-line users. A chat service with a large user base might employ multiple chat servers. In addition, there are several multi-protocol chat clients, such as Pidgin (formerly GAIM) and Trillian, that allow a user to join different chat systems.

Although IRC has existed for a long time, it has not gained mainstream popularity. This is mainly because its console-like interface and command-line-based operation are not user-friendly. The recent chat systems improve user experience by using graphic-based interfaces, as well as adding attractive features such as avatars, emoticons, and audio-video communication capabilities. Our study is carried out on the Yahoo! chat network, one

of the largest and most popular commercial chat systems.

Yahoo! chat uses proprietary protocols, in which the chat messages are transmitted in plain-text, while commands, status and other meta data are transmitted as encoded binary data. Unlike those on most IRC networks, users on the Yahoo! chat network cannot create chat rooms with customized topics because this feature is disabled by Yahoo! to prevent abuses [24]. In addition, users on Yahoo! chat are required to pass a CAPTCHA word verification test in order to join a chat room. This recently-added feature is to guard against a major source of abuse—bots.

2.2 Chat Bots

The term *bot*, short for robot, refers to automated programs, that is, programs that do not require a human operator. A chat bot is a program that interacts with a chat service to automate tasks for a human, e.g., creating chat logs. The first-generation chat bots were designed to help operate chat rooms, or to entertain chat users, e.g., quiz or quote bots. However, with the commercialization of the Internet, the main enterprise of chat bots is now sending chat spam. Chat bots deliver spam URLs via either links in chat messages or user profile links. A single bot operator, controlling a few hundred chat bots, can distribute spam links to thousands of users in different chat rooms, making chat bots very profitable to the bot operator who is paid per-click through affiliate programs. Other potential abuses of bots include spreading malware, phishing, booting, and similar malicious activities.

A few countermeasures have been used to defend against the abuse of chat bots, though none of them are very effective. On the server side, CAPTCHA tests are used by Yahoo! chat in an effort to prevent chat bots joining chat rooms. However, this defense becomes ineffective as chat bots bypass CAPTCHA tests with human assistance. We have observed that bots continue to join chat rooms and sometimes even become the majority members of a chat room after the deployment of CAPTCHA tests. Third-party chat clients filter out chat bots, mainly based on key words or key phrases that are known to be used by chat bots. The drawback with this approach is that it cannot capture those unknown or evasive chat bots that do not use the known key words or phrases.

2.3 Related Work

Dewes et al. [9] conducted a systematic measurement study of IRC and Web-chat traffic, revealing several statistical properties of chat traffic. (1) Chat sessions tend to last for a long time, and a significant number of IRC ses-

sions last much longer than Web-chat sessions. (2) Chat session inter-arrival time follows an exponential distribution, while the distribution of message inter-arrival time is not exponential. (3) In terms of message size, all chat sessions are dominated by a large number of small packets. (4) Over an entire session, typically a user receives about 10 times as much data as he sends. However, very active users in Web-chat and automated scripts used in IRC may send more data than they receive.

There is considerable overlap between chat and instant messaging (IM) systems, in terms of protocol and user base. Many widely used chat systems such as IRC predate the rise of IM systems, and have great impact upon the IM system and protocol design. In return, some new features that make the IM systems more user-friendly have been back-ported to the chat systems. For example, IRC, a classic chat system, implements a number of IM-like features, such as presence and file transfers, in its current versions. Some messaging service providers, such as Yahoo!, offer both chat and IM accesses to their end-user clients. With this in mind, we outline some related work on IM systems. Liu et al. [21] explored client-side and server-side methods for detecting and filtering IM spam or *spim*. However, their evaluation is based on a corpus of short e-mail spam messages, due to the lack of data on *spim*. In [23], Mannan et al. studied IM worms, automated malware that spreads on IM systems using the IM contact list. Leveraging the spreading characteristics of IM malware, Xie et al. [41] presented an IM malware detection and suppression system based on the honeypot concept.

Botnets consist of a large number of slave computing assets, which are also called “bots”. However, the usage and behavior of bots in botnets are quite different from those of chat bots. The bots in botnets are malicious programs designed specifically to run on compromised hosts on the Internet, and they are used as platforms to launch a variety of illicit and criminal activities such as credential theft, phishing, distributed denial-of-service attacks, etc. In contrast, chat bots are automated programs designed mainly to interact with chat users by sending spam messages and URLs in chat rooms. Although having been used by botnets as command and control mechanisms [2, 11], IRC and other chat systems do not play an irreplaceable role in botnets. In fact, due to the increasing focus on detecting and thwarting IRC-based botnets [8, 13, 14], recently emerged botnets, such as Phatbot, Nugache, Slapper, and Sinit, show a tendency towards using P2P-based control architectures [39].

Chat spam shares some similarities with email spam. Like email spam, chat spam contains advertisements of illegal services and counterfeit goods, and solicits human users to click spam URLs. Chat bots employ many text obfuscation techniques used by email spam such

as word padding and synonym substitution. Since the detection of email spam can be easily converted into the problem of text classification, many content-based filters utilize machine-learning algorithms for filtering email spam. Among them, Bayesian-based statistical approaches [6, 12, 20, 44, 45] have achieved high accuracy and performance. Although very successful, Bayesian-based spam detection techniques still can be evaded by carefully crafted messages [18, 22, 40].

3 Measurement

In this section, we detail our measurements on Yahoo! chat, one of the most popular commercial chat services. The focus of our measurements is on public messages posted to Yahoo! chat rooms. The logging of chat messages is available on the standard Yahoo! chat client, as well as most third party chat clients. Upon entering chat, all chat users are shown a disclaimer from Yahoo! that other users can log their messages. However, we consider the contents of the chat logs to be sensitive, so we only present fully-anonymized statistics.

Our data was collected between August and November of 2007. In late August, Yahoo! implemented a CAPTCHA check on entering chat rooms [5, 26], creating technical problems that made their chat rooms unstable for about two weeks [3, 4]. At the same time, Yahoo! implemented a protocol update, preventing most third party chat clients, used by a large proportion of Yahoo! chat users, from accessing the chat rooms. In short, these upgrades made the chat rooms difficult to be accessed for both chat bots and humans. In mid to late September, both chat bot and third party client developers updated their programs. By early October, chat bots were found in Yahoo! chat [25], possibly bypassing the CAPTCHA check with human assistance. Due to these problems and the lack of chat bots in September and early October, we perform our analysis on August and November chat logs. In August and November, we collected a total of 1,440 hours of chat logs. There are 147 individual chat logs from 21 different chat rooms. The process of reading and labeling these chat logs required about 100 hours. To the best of our knowledge, we are the first in the large scale measurement and classification of chat bots.

3.1 Log-Based Classification

In order to characterize the behavior of human users and that of chat bots, we need two sets of chat logs pre-labeled as bots and humans. To create such datasets, we perform log-based classification by reading and labeling a large number of chat logs. The chat users are labeled in three categories: human, bot, and ambiguous.

The log-based classification process is a variation of the Turing test. In a standard Turing test [37], the examiner converses with a test subject (a possible machine) for five minutes, and then decides if the subject is a human or a machine. In our classification process, the examiner observes a long conversation between a test subject (a possible chat bot) and one or more third parties, and then decides if the subject is a human or a chat bot. In addition, our examiner checks the content of URLs and typically observes multiple instances of the same chat bot, which further improve our classification accuracy. Moreover, given that the best practice of current artificial intelligences [36] can rarely pass a non-restricted Turing test, our classification of chat bots should be very accurate.

Although a Turing test is subjective, we outline a few important criteria. The main criterion for being labeled as human is a high proportion of specific, intelligent, and human-like responses to other users. In general, if a user's responses suggest more advanced intelligence than current state-of-the-art AI [36], then the user can be labeled as human. The ambiguous label is reserved for non-English, incoherent, or non-communicative users. The criteria for being classified as bot are as follows. The first is the lack of the intelligent responses required for the human label. The second is the repetition of similar phrases either over time or from other users (other instances of the same chat bot). The third is the presence of spam or malware URLs in messages or in the user's profile.

3.2 Analysis

In total, our measurements capture 14 different types of chat bots. The different types of chat bots are determined by their triggering mechanisms and text obfuscation schemes. The former relates to message timing, and the latter relates to message content. The two main types of triggering mechanisms observed in our measurements are timer-based and response-based. A timer-based bot sends messages based on a timer, which can be periodic (i.e., fixed time intervals) or random (i.e., variable time intervals). A response-based bot sends messages based on programmed responses to specific content in messages posted by other users.

There are many different kinds of text obfuscation schemes. The purpose of text obfuscation is to vary the content of messages and make bots more difficult to recognize or appear more human-like. We observed four basic text obfuscation methods that chat bots use to evade filtering or detection. First, chat bots introduce random characters or space into their messages, similar to some spam e-mails. Second, chat bots use various synonym phrases to avoid obvious keywords. By this method, a template with several synonyms for multiple words can

lead to thousands of possible messages. Third, chat bots use short messages or break up long messages into multiple messages to evade message filters that work on a message-by-message basis. Fourth, and most interestingly, chat bots replay human phrases entered by other chat users.

According to our observation, the main activity of chat bots is to send spam links to chat users. There are two approaches that chat bots use to distribute spam links in chat rooms. The first is to post a message with a spam link directly in the chat room. The second is to enter the spam URL in the chat bot's user profile and then convince the users to view the profile and click the link. Our logs also include some examples of malware spreading via chat rooms. The behavior of malware-spreading chat bots is very similar to that of spam-sending chat bots, as both attempt to lure human users to click links. Although we did not perform detailed malware analysis on links posted in the chat rooms and Yahoo! applies filters to block links to known malicious files, we found several worm instances in our data. There are 12 W32.Imaut.AS [35] worms appeared in the August chat logs, and 23 W32.Imaut.AS worms appeared in the November chat logs. The November worms attempted to send malicious links but were blocked by Yahoo! (the malicious links in their messages being removed), however, the August worms were able to send out malicious links.

The focus of our measurements is mainly on short term statistics, as these statistics are most likely to be useful in chat bot classification. The two key measurement metrics in this study are inter-message delay and message size. Based on these two metrics, we profile the behavior of human and that of chat bots. Among chat bots, we further divide them into four different groups: periodic bots, random bots, responder bots, and replay bots. With respect to these short-term statistics, human and chat bots behave differently, as shown below.

3.2.1 Humans

Figure 1 shows the probability distributions of human inter-message delay and message size. Since the behavior of humans is persistent, we only draw the probability mass function (pmf) curves based on the August data. The previous study on Internet chat systems [9] observed that the distribution of inter-message delay in chat systems was heavy tailed. In general our measurement result conforms to that observation. The body part of the pmf curve in Figure 1 (a) (log-log scale) can be linearly fitted, indicating that the distribution of human inter-message delays follows a power law. In other words, the distribution is heavy tailed. We also find that the pmf curve of human message size in Figure 1 (b) can be well fitted by an exponential distribution with $\lambda = 0.034$ after




Figure 1: Distribution of human inter-message delay (a) and message size (b)

excluding the initial spike.

3.2.2 Periodic Bots

A periodic bot posts messages mainly at regular time intervals. The delay periods of periodic bots, especially those bots that use long delays, may vary by several seconds. The variation of delay period may be attributed to either transmission delay caused by network traffic congestion or chat server delay, or message emission delay incurred by system overloading on the bot hosting machine. The posting of periodic messages is a simple but effective mechanism for distributing messages, so it is not surprising that a substantial portion of chat bots use periodic timers.

We display the probability distributions of inter-message delay and message size for periodic bots in Figure 2. We use '+' for displaying August data and '•' for November data. The distributions of periodic bots are distinct from those of humans shown in Figure 1. The distribution of inter-message delay for periodic bots clearly manifests the timer-triggering characteristic of periodic bots. There are three clusters with high probabilities at time ranges [30-50], [100-110], and [150-170]. These clusters correspond to the November periodic bots with timer values around 40 seconds and the August periodic bots with timer values around 105 and 160 seconds, respectively. The message size pmf curve of the August periodic bots shows an interesting bell shape, much like a normal distribution. After examining message contents, we find that the bell shape may be attributed to the message composition method some August bots used. As shown in Appendix A, some August periodic bots compose a message using a single template. The template has several parts and each part is associated with several synonym phrases. Since the length of each part is inde-

pendent and identically distributed, the length of whole message, i.e., the sum of all parts, should approximate a normal distribution. The November bots employ a similar composition method, but use several templates of different lengths. Thus, the message size distribution of the November periodic bots reflects the distribution of the lengths of the different templates, with the length of each individual template approximating a normal distribution.

3.2.3 Random Bots

A random bot posts messages at random time intervals. The random bots in our data used different random distributions, some discrete and others continuous, to generate inter-message delays. The use of random timers makes random bots appear more human-like than periodic bots. In statistical terms, however, random bots exhibit quite different inter-message delay distributions than humans.

Figure 3 depicts the probability distributions of inter-message delay and message size for random bots. Compared to periodic bots, random bots have more dispersed timer values. In addition, the August random bots have a large overlap with the November random bots. The points with high probabilities (greater than 10^{-2}) in the time range [30-90] in Figure 3 (a) represent the August and November random bots that use a discrete distribution of 40, 64, and 88 seconds. The wide November cluster with medium probabilities in the time range [40-130] is created by the November random bots that use a uniform distribution between 45 and 125 seconds. The probabilities of different message sizes for the August and November random bots are mainly in the size range [0-50]. Unlike periodic bots, most random bots do not use template or synonym replacement, but directly repeat messages. Thus, as their messages are selected from a database at random, the message size distribution re-

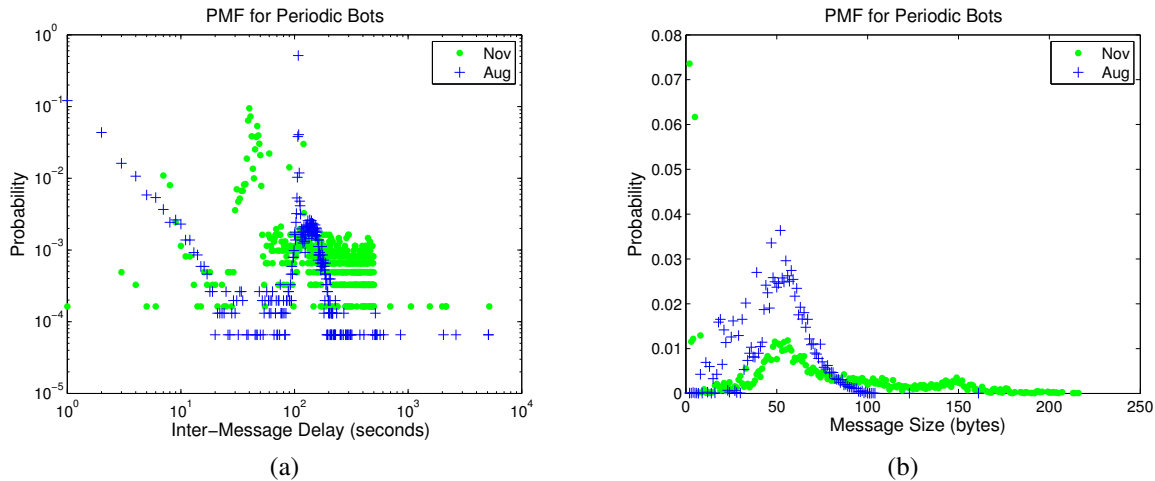


Figure 2: Distribution of periodic bot inter-message delay (a) and message size (b)

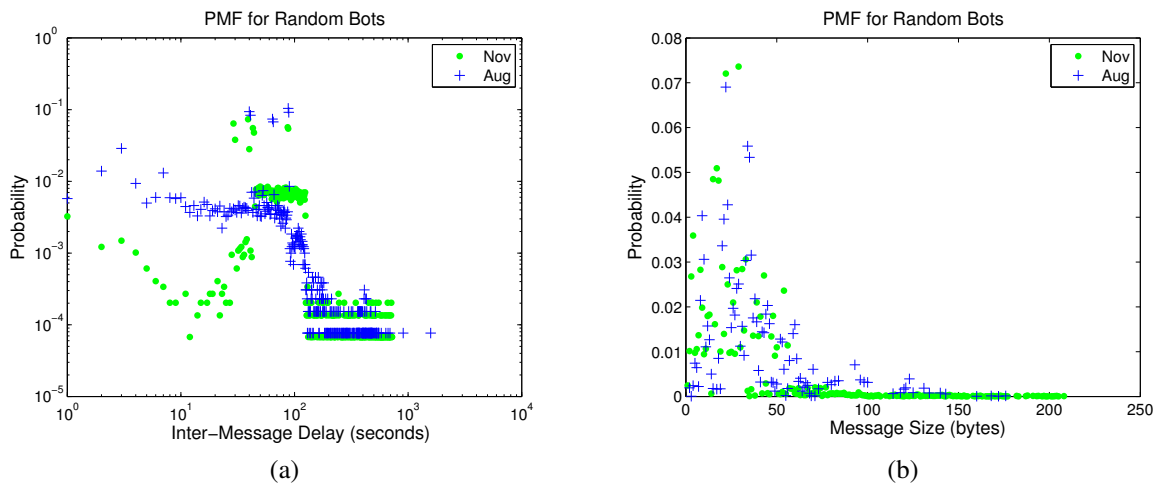


Figure 3: Distribution of random bot inter-message delay (a) and message size (b)

flects the proportion of messages of different sizes in the database.

3.2.4 Responder Bots

A responder bot sends messages based on the content of messages in the chat room. For example, a message ending with a question mark may trigger a responder bot to send a vague response with a URL, as shown in Appendix A. The vague response, in the context, may trick human users into believing that the responder is a human and further clicking the link. Moreover, the message triggering mechanism makes responder bots look more like humans in terms of timing statistics than periodic or random bots.

To gain more insights into responder bots, we managed to obtain a configuration file for a typical responder

bot [38]. There are a number of parameters for making the responder bot mimic humans. The bot can be configured with a fixed typing rate, so that responses with different lengths take different time to “type.” The bot can also be set to either ignore triggers while simulating typing, or rate-limit responses. In addition, responses can be assigned with probabilities, so that the responder bot responds to a given trigger in a random manner.

Figure 4 shows the probability distributions of inter-message delay and message size for responder bots. Note that only the distribution of the August responder bots is shown due to the small number of responder bots found in November. Since the message emission of responder bots is triggered by human messages, theoretically the distribution of inter-message delays of responder bots should demonstrate certain similarity to that of humans.




Figure 4: Distribution of responder bot inter-message delay (a) and message size (b)




Figure 5: Distribution of replay bot inter-message delay (a) and message size (b)

Figure 4 (a) confirms this hypothesis. Like Figure 1 (a), the pmf of responder bots (excluding the head part) in log-log scale exhibits a clear sign of a heavy tail. But unlike human messages, the sizes of responder bot messages vary in a much narrower range (between 1 and 160). The bell shape of the distribution for message size less than 100 indicates that responder bots share a similar message composition technique with periodic bots, and their messages are composed as templates with multiple parts, as shown in Appendix A.

3.2.5 Replay Bots

A replay bot not only sends its own messages, but also repeats messages from other users to appear more like a human user. In our experience, replayed phrases are related to the same topic but do not appear in the same chat

room as the original ones. Therefore, replayed phrases are either taken from other chat rooms on the same topic or saved previously in a database and replayed.

The use of replayed phrases in a crowded or “noisy” chat room does, in fact, make replay bots look more like human to inattentive users. The replayed phrases are sometimes nonsensical in the context of the chat, but human users tend to naturally ignore such statements. When replay bots succeed in fooling human users, these users are more likely to click links posted by the bots or visit their profiles. Interestingly, replay bots sometimes replay phrases uttered by other chat bots, making them very easy to be recognized. The use of replay is potentially effective in thwarting detection methods, as detection tests must deal with a combination of human and bots phrases. By using human phrases, replay bots

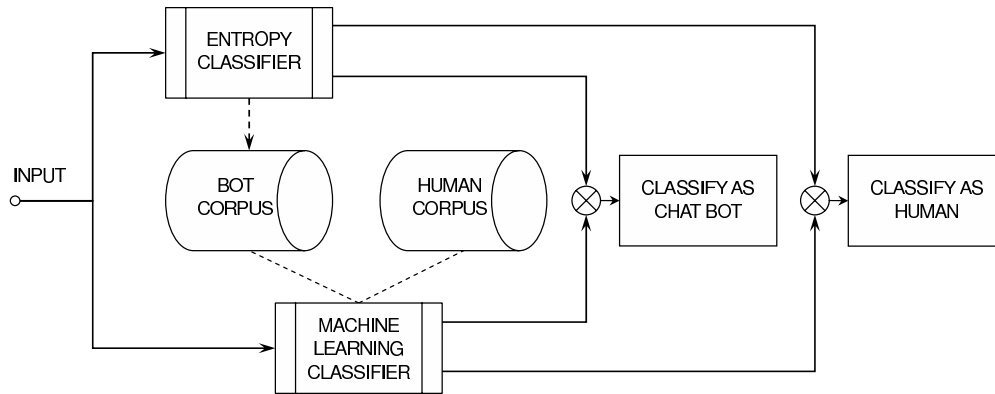


Figure 6: Classification System Diagram

can easily defeat keyword-based message filters that filter message-by-message, as the human phrases should not be filtered out.

Figure 5 illustrates the probability distributions of inter-message delay and message size for replay bots. In terms of inter-message delay, a replay bot is just a variation of a periodic bot, which is demonstrated by the high spike in Figure 5 (a). By using human phrases, replay bots successfully mimic human users in terms of message size distribution. The message size distribution of replay bots in Figure 5 (b) largely resembles that of human users, and can be fitted by an exponential distribution with $\lambda = 0.028$.

4 Classification System

This section describes the design of our chat bot classification system. The two main components of our classification system are the entropy classifier and the machine learning classifier. The basic structure of our chat bot classification system is shown in Figure 6. The two classifiers, entropy and machine learning, operate concurrently to process input and make classification decisions, while the machine learning classifier relies on the entropy classifier to build the bot corpus. The entropy classifier uses entropy and corrected conditional entropy to score chat users and then classifies them as chat bots or humans. The main task of the entropy classifier is to capture new chat bots and add them to the chat bot corpus. The human corpus can be taken from a database of clean chat logs or created by manual log-based classification, as described in Section 3. The machine learning classifier uses the bot and human corpora to learn text patterns of bots and humans, and then it can quickly classify chat bots based on these patterns. The two classifiers are detailed as follows.

4.1 Entropy Classifier

The entropy classifier makes classification decisions based on entropy and entropy rate measures of message sizes and inter-message delays for chat users. If either the entropy or entropy rate is low for these characteristics, it indicates the regular or predictable behavior of a likely chat bot. If both the entropy and entropy rate is high for these characteristics, it indicates the irregular or unpredictable behavior of a possible human.

To use entropy measures for classification, we set a cutoff score for each entropy measure. If a test score is greater than or equal to the cutoff score, the chat user is classified as a human. If the test score is less than the cutoff score, the chat user is classified as a chat bot. The specific cutoff score is an important parameter in determining the false positive and true positive rates of the entropy classifier. On the one hand, if the cutoff score is too high, then too many humans will be misclassified as bots. On the other hand, if the cutoff score is too low, then too many chat bots will be misclassified as humans. Due to the importance of achieving a low false positive rate, we select the cutoff scores based on human entropy scores to achieve a targeted false positive rate. The specific cutoff scores and targeted false positive rates are described in Section 5.

4.1.1 Entropy Measures

The entropy rate, which is the average entropy per random variable, can be used as a measure of complexity or regularity [10, 30, 31]. The entropy rate is defined as the conditional entropy of a sequence of infinite length. The entropy rate is upper-bounded by the entropy of the first-order probability density function or first-order entropy. A independent and identically distributed (i.i.d.) process has an entropy rate equal to its first-order entropy. A highly complex process has a high entropy rate, while a highly regular process has a low entropy rate.

A random process $X = \{X_i\}$ is defined as an indexed sequence of random variables. To give the definition of the entropy rate of a random process, we first define the entropy of a sequence of random variables as:

$$H(X_1, \dots, X_m) = - \sum_{X_1, \dots, X_m} P(x_1, \dots, x_m) \log P(x_1, \dots, x_m),$$

where $P(x_1, \dots, x_m)$ is the joint probability $P(X_1 = x_1, \dots, X_m = x_m)$.

Then, from the entropy of a sequence of random variables, we define the conditional entropy of a random variable given a previous sequence of random variables as:

$$H(X_m | X_1, \dots, X_{m-1}) = H(X_1, \dots, X_m) - H(X_1, \dots, X_{m-1}).$$

Lastly, the entropy rate of a random process is defined as:

$$\bar{H}(X) = \lim_{m \rightarrow \infty} H(X_m | X_1, \dots, X_{m-1}).$$

Since the entropy rate is the conditional entropy of a sequence of infinite length, it cannot be measure for finite samples. Thus, we estimate the entropy rate with the conditional entropy of finite samples. In practice, we replace probability density functions with empirical probability density functions based on the method of histograms. The data is binned in Q bins of approximately equal probability. The empirical probability density functions are determined by the proportions of bin number sequences in the data, i.e., the proportion of a sequence is the probability of that sequence. The estimates of the entropy and conditional entropy, based on empirical probability density functions, are represented as: EN and CE , respectively.

There is a problem with the estimation of $CE(X_m | X_1, \dots, X_{m-1})$ for some values of m . The conditional entropy tends to zero as m increases, due to limited data. If a specific sequence of length $m - 1$ is found only once in the data, then the extension of this sequence to length m will also be found only once. Therefore, the length m sequence can be predicted by the length $m - 1$ sequence, and the length m and $m - 1$ sequences cancel out. If no sequence of length m is repeated in the data, then $CE(X_m | X_1, \dots, X_{m-1})$ is zero, even for i.i.d. processes.

To solve the problem of limited data, without fixing the length of m , we use the corrected conditional entropy [30] represented as CCE . The corrected conditional entropy is defined as:

$$CCE(X_m | X_1, \dots, X_{m-1}) = CE(X_m | X_1, \dots, X_{m-1}) + perc(X_m) \cdot EN(X_1),$$

where $perc(X_m)$ is the percentage of unique sequences of length m and $EN(X_1)$ is the entropy with m fixed at 1 or the first-order entropy.

The estimate of the entropy rate is the minimum of the corrected conditional entropy over different values of m . The minimum of the corrected conditional entropy is considered to be the best estimate of the entropy rate from the available data.

4.2 Machine Learning Classifier

The machine learning classifier uses the content of chat messages to identify chat bots. Since chat messages (including emoticons) are text, the identification of chat bots can be perfectly fitted into the domain of machine learning text classification. Within the machine learning paradigm, the text classification problem can be formalized as $f : T \times C \rightarrow \{0, 1\}$, where f is the classifier, $T = \{t_1, t_2, \dots, t_n\}$ is the texts to be classified, and $C = \{c_1, c_2, \dots, c_k\}$ is the set of pre-defined classes [33]. Value 1 for $f(t_i, c_j)$ indicates that text t_i is in class c_j and value 0 indicates the opposite decision. There are many techniques that can be used for text classification, such as naïve Bayes, support vector machines, and decision trees. Among them, Bayesian classifiers have been very successful in text classification, particularly in email spam detection. Due to the similarity between chat spam and email spam, we choose Bayesian classification for our machine learning classifier for detecting chat bots. We leave study on the applicability of other types of machine learning classifiers to our future work.

Within the framework of Bayesian classification, identifying if chat message M is issued by a bot or human is achieved by computing the probability of M being from a bot with the given message content, i.e., $P(C = bot|M)$. If the probability is equal to or greater than a pre-defined threshold, then message M is classified as a bot message. According to Bayes theorem,

$$P(bot|M) = \frac{P(M|bot)P(bot)}{P(M)} = \frac{P(M|bot)P(bot)}{P(M|bot)P(bot) + P(M|human)P(human)}.$$

A message M is described by its feature vector $\langle f_1, f_2, \dots, f_n \rangle$. A feature f is a single word or a combination of multiple words in the message. To simplify computation, in practice it is usually assumed that all features are conditionally independent with each other for

Table 1: Message Composition of Chat Bot and Human Datasets

	AUG. BOTS			NOV. BOTS			HUMANS
	periodic	random	responder	periodic	random	replay	human
number of messages	25,258	13,998	6,160	10,639	22,820	8,054	342,696

the given category. Thus, we have

$$P(bot|M) = \frac{P(bot) \prod_{i=1}^n P(f_i|bot)}{P(bot) \prod_{i=1}^n P(f_i|bot) + P(human) \prod_{i=1}^n P(f_i|human)}$$

The value of $P(bot|M)$ may vary in different implementations (see [12, 45] for implementation details) of Bayesian classification due to differences in assumption and simplification.

Given the abundance of implementations of Bayesian classification, we directly adopt one implementation, namely CRM 114 [44], as our machine learning classification component. CRM 114 is a powerful text classification system that has achieved very high accuracy in email spam identification. The default classifier of CRM 114, OSB (Orthogonal Sparse Bigram), is a type of Bayesian classifier. Different from common Bayesian classifiers which treat individual words as features, OSB uses word pairs as features instead. OSB first chops the whole input into multiple basic units with five consecutive words in each unit. Then, it extracts four word pairs from each unit to construct features, and derives their probabilities. Finally, OSB applies Bayes theorem to compute the overall probability that the text belongs to one class or another.

5 Experimental Evaluation

In this section, we evaluate the effectiveness of our proposed classification system. Our classification tests are based on chat logs collected from the Yahoo! chat system. We test the two classifiers, entropy-based and machine-learning-based, against chat bots from August and November datasets. The machine learning classifier is tested with fully-supervised training and entropy-classifier-based training. The accuracy of classification is measured in terms of false positive and false negative rates. The false positives are those human users that are misclassified as chat bots, while the false negatives are those chat bots that are misclassified as human users. The speed of classification is mainly determined by the minimum number of messages that are required for accurate classification. In general, a high number means slow classification, whereas a low number means fast classification.

5.1 Experimental Setup

The chat logs used in our experiments are mainly in three datasets: (1) human chat logs from August 2007, (2) bot chat logs from August 2007, and (3) bot chat logs from November 2007. In total, these chat logs contain 342,696 human messages and 87,049 bot messages. In our experiments, we use the first half of each chat log, human and bot, for training our classifiers and the second half for testing our classifiers. The composition of the chat logs for the three datasets is listed in Table 1.

The entropy classifier only requires a human training set. We use the human training set to determine the cutoff scores, which are used by the entropy classifier to decide whether a test sample is a human or bot. The target false positive rate is set at 0.01. To achieve this false positive rate, the cutoff scores are set at approximately the 1st percentile of human training set scores. Then, samples that score higher than the cutoff are classified as humans, while samples that score lower than the cutoff are classified as bots. The entropy classifier uses two entropy tests: entropy and corrected conditional entropy. The entropy test estimates first-order entropy, and the corrected conditional entropy estimates higher-order entropy or entropy rate. The corrected conditional entropy test is more precise with coarse-grain bins, whereas the entropy test is more accurate with fine-grains bins [10]. Therefore, we use $Q = 5$ for the corrected conditional entropy test and $Q = 256$ with m fixed at 1 for the entropy test.

We run classification tests for each bot type using the entropy classifier and machine learning classifier. The machine learning classifier is tested based on fully-supervised training and then entropy-based training. In fully-supervised training, the machine learning classifier is trained with manually labeled data, as described in Section 3. In entropy-based training, the machine learning classifier is trained with data labeled by the entropy classifier. For each evaluation, the entropy classifier uses samples of 100 messages, while the machine learning classifier uses samples of 25 messages.

5.2 Experimental Results

We now present the results for the entropy classifier and machine learning classifier. The four chat bot types are: periodic, random, responder, and replay. The classification tests are organized by chat bot type, and are ordered by increasing detection difficulty.

Table 2: Entropy Classifier Accuracy

	AUG. BOTS			NOV. BOTS			HUMANS
	periodic	random	responder	periodic	random	replay	human
test	true pos.	true pos.	true pos.	true pos.	true pos.	true pos.	false pos.
EN(imd)	121/121	68/68	1/30	51/51	109/109	40/40	7/1713
CCE(imd)	121/121	49/68	4/30	51/51	109/109	40/40	11/1713
EN(ms)	92/121	7/68	8/30	46/51	34/109	0/40	7/1713
CCE(ms)	77/121	8/68	30/30	51/51	6/109	0/40	11/1713
OVERALL	121/121	68/68	30/30	51/51	109/109	40/40	17/1713

5.2.1 Entropy Classifier

The detection results of the entropy classifier are listed in Table 2, which includes the results of the entropy test (*EN*) and corrected conditional entropy test (*CCE*) for inter-message delay (*imd*), and message size (*ms*). The overall results for all entropy-based tests are shown in the final row of the table. The true positives are the total unique bot samples correctly classified as bots. The false positives are the total unique human samples mistakenly classified as bots.

Periodic Bots: As the simplest group of bots, periodic bots are the easiest to detect. They use different fixed timers and repeatedly post messages at regular intervals. Therefore, their inter-message delays are concentrated in a narrower range than those of humans, resulting in lower entropy than that of humans. The inter-message delay *EN* and *CCE* tests detect 100% of all periodic bots in both August and November datasets. The message size *EN* and *CCE* tests detect 76% and 63% of the August periodic bots, respectively, and 90% and 100% of the November periodic bots, respectively. These slightly lower detection rates are due to a small proportion of humans with low entropy scores that overlap with some periodic bots. These humans post mainly short messages, resulting in message size distributions with low entropy.

Random Bots: The random bots use random timers with different distributions. Some random bots use discrete timings, e.g., 40, 64, or 88 seconds, while the others use continuous timings, e.g., uniformly distributed delays between 45 and 125 seconds.

The inter-message delay *EN* and *CCE* tests detect 100% of all random bots, with one exception: the inter-message delay *CCE* test against the August random bots only achieves 72% detection rate, which is caused by the following two conditions: (1) the range of message delays of random bots is close to that of humans; (2) sometimes the randomly-generated delay sequences have similar entropy rate to human patterns. The message size *EN* and *CCE* tests detect 31% and 6% of August random bots, respectively, and 7% and 8% of November random bots, respectively. These low detection rates are again due to a small proportion of humans with low mes-

sage size entropy scores. However, unlike periodic bots, the message size distribution of random bots is highly dispersed, and thus, a larger proportion of random bots have high entropy scores, which overlap with those of humans.

Responder Bots: The responder bots are among the advanced bots, and they behave more like humans than random or periodic bots. They are triggered to post messages by certain human phrases. As a result, their timings are quite similar to those of humans.

The inter-message delay *EN* and *CCE* tests detect very few responder bots, only 3% and 13%, respectively. This demonstrates that human-message-triggered responding is a simple yet very effective mechanism for imitating the timing of human interactions. However, the detection rate for the message size *EN* test is slightly better at 27%, and the detection rate for the message size *CCE* test reaches 100%. While the message size distribution has sufficiently high entropy to frequently evade the *EN* test, there is some dependence between subsequent message sizes, and thus, the *CCE* detects the low entropy pattern over time.

Replay Bots: The replay bots also belong to the advanced and human-like bots. They use replay attacks to fool humans. More specifically, the bots replay phrases they observed in chat rooms. Although not sophisticated in terms of implementation, the replay bots are quite effective in deceiving humans as well as frustrating our message-size-based detections: the message size *EN* and *CCE* tests both have detection rates of 0%. Despite their clever trick, the timing of replay bots is periodic and easily detected. The inter-message delay *EN* and *CCE* tests are very successful at detecting replay bots, both with 100% detection accuracy.

5.2.2 Supervised and Hybrid Machine Learning Classifiers

The detection results of the machine learning classifier are listed in Table 3. Table 3 shows the results for the fully-supervised machine learning (*SupML*) classifier and entropy-trained machine learning (*EntML*) classifier, both trained on the August training datasets, and the

Table 3: Machine Learning Classifier Accuracy

	AUG. BOTS			NOV. BOTS			HUMANS
	periodic	random	responder	periodic	random	replay	human
test	true pos.	true pos.	true pos.	true pos.	true pos.	true pos.	false pos.
<i>SupML</i>	121/121	68/68	30/30	14/51	104/109	1/40	0/1713
<i>SupMLretrained</i>	121/121	68/68	30/30	51/51	109/109	40/40	0/1713
<i>EntML</i>	121/121	68/68	30/30	51/51	109/109	40/40	1/1713

fully-supervised machine learning (*SupMLretrained*) classifier trained on August and November training datasets.

Periodic Bots: For the August dataset, both *SupML* and *EntML* classifiers detect 100% of all periodic bots. For the November dataset, however, the *SupML* classifier only detects 27% of all periodic bots. The lower detection rate is due to the fact that 62% of the periodic bot messages in November chat logs are generated by new bots, making the *SupML* classifier ineffective without re-training. The *SupMLretrained* classifier detects 100% of November periodic bots. The *EntML* classifier also achieves 100% for the November dataset.

Random Bots: For the August dataset, both *SupML* and *EntML* classifiers detect 100% of all random bots. For the November dataset, the *SupML* classifier detects 95% of all random bots, and the *SupMLretrained* classifier detects 100% of all random bots. Although 52% of the random bots have been upgraded according to our observation, the old training set is still mostly effective because certain content features of August random bots still appear in November. The *EntML* classifier again achieves 100% detection accuracy for the November dataset.

Responder Bots: We only present the detection results of responder bots for the August dataset, as the number of responder bots in the November dataset is very small. Although responder bots effectively mimic human timing, their message contents are only slightly obfuscated and are easily detected. The *SupML* and *EntML* classifiers both detect 100% of all responder bots.

Replay Bots: The replay bots only exist in the November dataset. The *SupML* classifier detects only 3% of all replay bots, as these bots are newly introduced in November. The *SupMLretrained* classifier detects 100% of all replay bots. The machine learning classifier reliably detects replay bots in the presence of a substantial number of replayed human phrases, indicating the effectiveness of machine learning techniques in chat bot classification.

6 Conclusion and Future Work

This paper first presents a large-scale measurement study on Internet chat. We collected two-month chat logs for 21 different chat rooms from one of the top Internet chat service providers. From the chat logs, we identified a total of 14 different types of chat bots and grouped them into four categories: periodic bots, random bots, responder bots, and replay bots. Through statistical analysis on inter-message delay and message size for both chat bots and humans, we found that chat bots behave very differently from human users. More specifically, chat bots exhibit certain regularities in either inter-message delay or message size. Although responder bots and replay bots employ advanced techniques to behave more human-like in some aspects, they still lack the overall sophistication of humans.

Based on the measurement study, we further proposed a chat bot classification system, which utilizes entropy-based and machine-learning-based classifiers to accurately detect chat bots. The entropy-based classifier exploits the low entropy characteristic of chat bots in either inter-message delay or message size, while the machine-learning-based classifier leverages the message content difference between humans and chat bots. The entropy-based classifier is able to detect unknown bots, including human-like bots such as responder and replay bots. However, it takes a relatively long time for detection, i.e., a large number of messages are required. Compared to the entropy-based classifier, the machine-learning-based classifier is much faster, i.e., a small number of messages are required. In addition to bot detection, a major task of the entropy-based classifier is to build and maintain the bot corpus. With the help of bot corpus, the machine-learning-based classifier is trained, and consequently, is able to detect chat bots quickly and accurately. Our experimental results demonstrate that the hybrid classification system is fast in detecting known bots and is accurate in identifying previously-unknown bots.

There are a number of possible directions for our future work. We plan to explore the application of entropy-based techniques in detecting other forms of bots, such as web bots. We also plan to investigate the development of more advanced chat bots that could evade our hybrid

classification system. We believe that the continued work in this area will reveal other important characteristics of bots and automated programs, which is useful in malware detection and prevention.

Acknowledgments

We thank the anonymous reviewers for their insightful comments. This work was partially supported by NSF grants CNS-0627339 and CNS-0627340. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- [1] AHN, L. V., BLUM, M., HOPPER, N., AND LANGFORD, J. CAPTCHA: Using hard AI problems for security. In *Proceedings of Eurocrypt* (Warsaw, Poland, May 2003).
- [2] BÄCHER, P., HOLZ, T., KÖTTER, M., AND WICHERSKI, G. Know your enemy: Tracking botnets, 2005. <http://www.honeynet.org/papers/bots> [Accessed: Jan. 25, 2008].
- [3] BACON, S. Chat rooms follow-up. <http://www.ymessengerblog.com/blog/2007/08/21/chat-rooms-follow-up/> [Accessed: Jan. 25, 2008].
- [4] BACON, S. Chat rooms update. <http://www.ymessengerblog.com/blog/2007/08/24/chat-rooms-update-2/> [Accessed: Jan. 25, 2008].
- [5] BACON, S. New entry process for chat rooms. <http://www.ymessengerblog.com/blog/2007/08/29/new-entry-process-for-chat-rooms/> [Accessed: Jan. 25, 2008].
- [6] BLOSSER, J., AND JOSEPHSEN, D. Scalable centralized bayesian spam mitigation with bogofilter. In *Proceedings of the 2004 USENIX Systems Administration Conference (LISA'04)* (Atlanta, GA., USA, November 2004).
- [7] CRISLIP, D. Will Blizzard's spam-stopper really work? <http://www.wowinsider.com/2007/05/16/will-blizzards-spam-stopper-really-work/> [Accessed: Dec. 25, 2007].
- [8] DAGON, D., GU, G., LEE, C. P., AND LEE, W. A taxonomy of botnet structures. In *Proceedings of the 2007 Annual Computer Security Applications Conference (ACSAC'07)* (Miami, FL., USA, December 2007).
- [9] DEWES, C., WICHMANN, A., AND FELDMANN, A. An analysis of Internet chat systems. In *Proceedings of the 2003 ACM/SIGCOMM Internet Measurement Conference (IMC'03)* (Miami, FL., USA, October 2003).
- [10] GIANVECCHIO, S., AND WANG, H. Detecting covert timing channels: An entropy-based approach. In *Proceedings of the 2007 ACM Conference on Computer and Communications Security (CCS'07)* (Alexandria, VA., USA, October 2007).
- [11] GOEBEL, J., AND HOLZ, T. Rishi: Identify bot contaminated hosts by IRC nickname evaluation. In *Proceedings of the USENIX Workshop on Hot Topics in Understanding Botnets (HotBots'07)* (Cambridge, MA., USA, April 2007).
- [12] GRAHAM, P. A plan for spam, 2002. <http://www.paulgraham.com/spam.html> [Accessed: Jan. 25, 2008].
- [13] GU, G., PORRAS, P., YEGNESWARAN, V., FONG, M., AND LEE, W. Bothunter: Detecting malware infection through IDS-driven dialog correlation. In *Proceedings of the 2007 USENIX Security Symposium (Security'07)* (Boston, MA., USA, August 2007).
- [14] GU, G., ZHANG, J., AND LEE, W. BotSniffer: Detecting botnet command and control channels in network traffic. In *Proceedings of the 2008 Annual Network and Distributed System Security Symposium (NDSS'08)* (San Diego, CA., USA, February 2008).
- [15] HU, J. AOL: spam and chat don't mix. http://www.news.com/AOL-Spam-and-chat-dont-mix/2100-1032_3-1024010.html [Accessed: Jan. 7, 2008].
- [16] HU, J. Shutting of MSN chat rooms may open up IM. http://www.news.com/Shutting-of-MSN-chat-rooms-may-open-up-IM/2100-1025_3-5082677.html [Accessed: Jan. 7, 2008].
- [17] JENNINGS III, R. B., NAHUM, E. M., OLSHEFSKI, D. P., SAHA, D., SHAE, Z.-Y., AND WATERS, C. A study of internet instant messaging and chat protocols. *IEEE Network Vol. 20*, No. 4 (2006), 16–21.
- [18] KARLBERGER, C., BAYLER, G., KRUEGEL, C., AND KIRDA, E. Exploiting redundancy in natural language to penetrate bayesian spam filters. In *Proceedings of the USENIX Workshop on Offensive Technologies* (Boston, MA., USA, August 2007).
- [19] KREBS, B. Yahoo! messenger network overrun by bots. http://blog.washingtonpost.com/securityfix/2007/08/yahoo_messenger_network_overru.html [Accessed: Dec. 18, 2007].
- [20] LI, K., AND ZHONG, Z. Fast statistical spam filter by approximate classifications. In *Proceedings of 2006 ACM/SIGMETRICS International Conference on Measurement and Modeling of Computer Systems* (St. Malo, France, June 2006).
- [21] LIU, Z., LIN, W., LI, N., AND LEE, D. Detecting and filtering instant messaging spam - a global and personalized approach. In *Proceedings of the IEEE Workshop on Secure Network Protocols (NPSEC'05)* (Boston, MA., USA, November 2005).

- [22] LOWD, D., AND MEEK, C. Good word attacks on statistical spam filters. In *Proceedings of the 2005 Conference on Email and Anti-Spam (CEAS'05)* (Mountain View, CA., USA, July 2005).
- [23] MANNAN, M., AND VAN OORSCHOT, P. C. On instant messaging worms, analysis and countermeasures. In *Proceedings of the ACM Workshop on Rapid Malcode* (Fairfax, VA., USA, November 2005).
- [24] MILLS, E. Yahoo! closes chat rooms over child sex concerns. http://www.news.com/Yahoo-closes-chat-rooms-over-/child-sex-concerns/2100-1025_3-5759705.html [Accessed: Jan. 27, 2008].
- [25] MOHTA, A. Bots are back in Yahoo! chat rooms. <http://www.technospot.net/blogs/bots-are-back-in-yahoo-chat-room/> [Accessed: Dec. 18, 2007].
- [26] MOHTA, A. Yahoo! chat adds CAPTCHA check to remove bots. <http://www.technospot.net/blogs/yahoo-chat-captcha-check-to-remove-bots/> [Accessed: Dec. 18, 2007].
- [27] NINO, T. Linden Lab taking action against landbots. <http://www.secondlifeinsider.com/2007/05/18/linden-lab-taking-action-against-landbots/> [Accessed: Jan. 7, 2008].
- [28] PETITION ONLINE. Action against the Yahoo! bot problem petition. <http://www.petitiononline.com/> [Accessed: Dec. 18, 2007].
- [29] PETITION ONLINE. AOL no more chat room spam petition. <http://www.petitiononline.com/> [Accessed: Dec. 18, 2007].
- [30] PORTA, A., BASELLI, G., LIBERATI, D., MONTANO, N., COGLIATI, C., GNECCHI-RUSCONE, T., MALLIANI, A., AND CERUTTI, S. Measuring regularity by means of a corrected conditional entropy in sympathetic outflow. *Biological Cybernetics Vol. 78*, No. 1 (January 1998).
- [31] ROSIPAL, R. *Kernel-Based Regression and Objective Nonlinear Measures to Assess Brain Functioning*. PhD thesis, University of Paisley, Paisley, Scotland, UK, September 2001.
- [32] SCHRAMM, M. Chat spam measures shut down multi-line reporting add-ons. <http://www.wowinsider.com/2007/10/25/chat-spam-measures-shut-down-multi-line-reporting-addons/> [Accessed: Jan. 17, 2008].
- [33] SEBASTIANI, F. Machine learning in automated text categorization. *ACM Computing Surveys Vol. 34*, No. 1 (2002), 1–47.
- [34] SIMPSON, C. Yahoo! chat anti-spam resource center. <http://www.chatspam.org/> [Accessed: Sep. 25, 2007].
- [35] SYMANTEC SECURITY RESPONSE. W32.Imaut.AS worm. http://www.symantec.com/security_response/writeup.jsp?docid=2007-080114-2713-99 [Accessed: Jan. 25, 2008].
- [36] THE ALICE ARTIFICIAL INTELLIGENCE FOUNDATION. ALICE(Artificial Linguistic Internet Computer Entity). <http://www.alicebot.org/> [Accessed: Jan. 25, 2008].
- [37] TURING, A. M. Computing machinery and intelligence. *Mind Vol. 59* (1950), 433–460.
- [38] UBER-GEEK.COM. Yahoo! responder bot. <http://www.uber-geek.com/bot.html> [Accessed: Jan. 18, 2008].
- [39] WANG, P., SPARKS, S., AND ZOU, C. C. An advanced hybrid peer-to-peer botnet. In *Proceedings of the USENIX Workshop on Hot Topics in Understanding Botnets (Hot-Bots'05)* (Cambridge, MA., USA, April 2007).
- [40] WITTEL, G. L., AND WU, S. F. On attacking statistical spam filters. In *Proceedings of the 2004 Conference on Email and Anti-Spam (CEAS'04)* (Mountain View, CA., USA, July 2004).
- [41] XIE, M., WU, Z., AND WANG, H. HoneyIM: Fast detection and suppression of instant messaging malware in enterprise-like networks. In *Proceedings of the 2007 Annual Computer Security Applications Conference (ACSAC'07)* (Miami Beach, FL, USA, December 2007).
- [42] YAHELITE.ORG. Yahelite chat client. <http://www.yahelite.org/> [Accessed: Jan. 8, 2008].
- [43] YAZAKPRO.COM. Yazak pro chat client. <http://www.yazakpro.com/> [Accessed: Jan. 8, 2008].
- [44] YERAZUNIS, B. CRM114 - the controllable regex multilator, 2003. <http://crm114.sourceforge.net> [Accessed: Jan. 25, 2008].
- [45] ZDZIARSKI, J. A. *Ending Spam: Bayesian Content Filtering and the Art of Statistical Language Classification*. No Starch Press, 2005.

A Chat Bot Examples

Note that in a chat room the following example messages would be spread out over several minutes.

Example 1: Response Template

```
bot: user1, that's a damn good question.
bot: user1, To know more about Seventh-day Adventist; visit http://www.sda.org
Sabbath; http://www.sabbathtruth.com EGW; http://www.whiteestate.org
bot: user2, no! don't leave me.

bot: user1, too much coffee tonight?
bot: user2, boy, you're just full of questions, aren't you?
bot: user2, lots of evidence for evolution can be found here http://www.talkorigins.org/faqs/comdesc/
```

In the above example, the bot uses a template with three parts to post links:
[username], [link description phrase] [link].

Example 2: Synonym Template

```
bot: Allo Hunks! Enjoy Marjorie! Check My Free Pics
bot: What's happening Guys! Marjorie Here! See more of me at My Free Pics
bot: Hi Babes! I am Marjorie! Rate My Live Cam
bot: Horny lover Guys! Marjorie at your service! Inspect My Site
bot: Mmmm Folks! Im Marjorie! View My Webpage
```

In the above example, the bot uses a template with three parts to post messages:
[salutation phrase]! [introduction phrase]! [web site advertisement phrase].

Example 3: Character Padding

```
bot: anyone boredjn wanna chat?uklcss
bot: any guystfrom the US/Canada hereiqjss
bot: hiyafxqss
bot: nel hereqbored?fiqss
bot: ne guysmwanna chat? ciuneed somel to make megsmile :-)pktpps
```

In the above example, the bot adds random characters to messages.