

Embedded Noninteractive Continuous Bot Detection

ROMAN V. YAMPOLSKIY AND VENU GOVINDARAJU
University at Buffalo, Buffalo

Multiplayer online computer games are quickly growing in popularity, with millions of players logging in every day. While most play in accordance with the rules set up by the game designers, some choose to utilize artificially intelligent assistant programs, a.k.a. bots, to gain an unfair advantage over other players. In this article we demonstrate how an embedded noninteractive test can be used to prevent automatic artificially intelligent players from illegally participating in online game-play. Our solution has numerous advantages over traditional tests, such as its nonobtrusive nature, continuous verification, and simple noninteractive and outsourcing-proof design.

Categories and Subject Descriptors: H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—*Artificial, augmented, and virtual realities*

General Terms: Design, Experimentation

Additional Key Words and Phrases: Games, agents, bots, CAPTCHA, reverse Turing test

ACM Reference Format:

Yampolskiy, R.V., and Govindaraju, V. 2008. Embedded Noninteractive continuous bot detection. *ACM Comput. Entertain.* 5, 4, Article 7 (March 2008). 11 Pages. DOI = 10.1145/1324198.1324205 <http://doi.acm.org/10.1145/1324198.1324205>

1. INTRODUCTION

Multiplayer online computer games are quickly growing in popularity, with millions of players logging in every day. While most play in accordance with the rules set up by the game designers, some choose to utilize artificially intelligent assistant programs, a.k.a. bots, to gain an unfair advantage over other players. With the growth in the economic and social importance of the virtual game worlds, use of such forbidden game bots is becoming increasingly problematic. Use of bots by some players makes the game less interesting for the unaided players [Mowbray 2007]. As a result, it costs thousands of dollars to game designers in lost revenues from disillusioned players who stop participating and in resources used for preventing different forms of cheating, including use of bots.

In this article we propose a solution to the problem of unauthorized bots playing in online games, particularly games with real financial outcomes such as poker. We begin with the introductory section on the use of unauthorized artificially intelligent assistants,

Author's address: Roman V. Yampolskiy: Center for Unified Biometrics and Sensors and IGERT in GIS, University at Buffalo, 2145 Monroe Ave. #4, Rochester NY, 14618; email: rvy@buffalo.edu. Venu Govindaraju: Center for Unified Biometrics and Sensors, University at Buffalo, 520 Lee Entrance, Suite 202, UB Commons Amherst, NY 14228; email: govind@buffalo.edu

Permission to make digital/hard copy of part of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date of appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Permission may be requested from the Publications Dept., ACM, Inc., 2 Penn Plaza, New York, NY 11201-0701, USA, fax: +1 (212) 869-0481, permissions@acm.org
©2008 ACM 1544-3574/08/0300-ART7 \$5.00 DOI = 10.1145/1324198.1324205
<http://doi.acm.org/10.1145/1324198.1324205>

followed by an overview of previous work aimed at preventing bots from obtaining resources and participating in games. Finally, we present our solution and its experimental evaluation.

2. UNAUTHORIZED GAME BOTS

Artificially intelligent (AI) programs are quickly becoming a part of our everyday lives. Virtual assistants, shopping bots, and smart search engines, to give just some examples, are used daily by millions of people. Such automated intelligent assistants are known as bots, which is a shortened version of *robots*. In the context of computer games, many different bots are known to exist, but all can be classified into one of the three major categories:

- Bots designed to enhance a user's intellectual abilities such as chess-playing programs, which can be consulted to defeat a human opponent who would otherwise be able to defeat the cheating player.
- Bots aimed at improving a user's physical abilities such as hand-eye coordination. An example is an aimbot used in first-person shooter games to augment the user's reflexes to the point of perfection.
- There is also a large number of bots designed to automate a tedious repetitive task such as resource gathering in games like the *World of Warcraft*.

Bots can also be categorized on the amount of human participation they require:

- *Nonautonomous*: Game guides, calculators, statistical tables, and other non-interactive sources of help fall into this category. Typically such assistance is not considered to be a form of cheating, and so is not the main target of bot-detection or prevention research.
- *Simiautonomous*: Bots capable of automatically performing certain sets of repetitive tasks, but require human assistance for at least a part of game interaction.
- *Fully autonomous*: Bots capable of playing the complete game without human intervention. They require no input from a human supervisor to either interact with the game software, select between different game options, or to terminate play. Some may be equipped with antidetection capabilities such as simulating a simple verbal interaction in the form of a chat.

Two additional types of bots are worth mentioning for the sake of completeness:

- *Bot networks*: A number of bots in the same game space can be connected to create an information-sharing network in which all bots actively help each other to win either by simple information sharing or via active action assistance. Additionally, an even bigger edge can be obtained if such a bot network has access to an external database of player profiles [Poker-edge.com. 2006; Pokerprophecy.2006].
- *Nonplayer characters* (NPCs): Probably the best researched type of bots [Byl 2004; Namee et al. 2003; Doyle 1999; Kline and Blumberg 1999; Moyer 2007]. They serve as opponents to human players in the games, and ideally are supposed to model the intelligent behavior of human players to make the game as interesting and realistic as possible. Current research concentrates on

creating human-like NPCs with respect to emotion [Freeman 2004]; intelligence [Laird 2001]; skills [Laird and Duch 2000; Enrique et al. 2003]; and overall look and feel [Blau 2002].

Techniques aimed at counteracting bot participation may be used to enforce one of the desirable properties: human presence or human play, which are defined by Golle and Ducheneaut [2005a; 2005b], as follows:

Human presence implies that a bot can't play completely unsupervised and that a human being is present for at least some interaction with the game software. This somewhat weak condition of human presence precludes a numerical explosion of participating bots in a game, by limiting the number of active bots to some function of human beings actually participating in the game. The value of such function depends on a number of bots a single human player may supervise at the same time. The property of human presence guarantees that a human being is investing at least some amount of time into playing the game, and so any resources obtained by the bot are not completely free, as time is money.

Human play is a much stronger property, which is probably not realistic to achieve, as it requires that all interaction with the game comes from a human, without any involvement from the bot.

2.1 Bot Detection and Counter-Detection Methods

There is very little published on the subject of game-bot detection, perhaps due to the inherent difficulty of the problem. Here we present a short overview of methods known to be used by online casinos and other online game operators. To detect bots the game software may check a number of conditions:

- *Running processes*: Which software is running on the system and what network connections are active? These questions are asked to find out whether well-known commercial bots are being run by the user [Winholdem 2006].
- *Reaction time*: Bots may exhibit a predetermined reaction time as measured from the appearance of stimulus to the making of an action.
- *Duration of play*: Bots may be run for unreasonably long periods of time without any breaks. Human beings, and even professional players, are not likely to play for over 12 hours straight.
- *Consistency of behavior*: Bots are often utilized to accomplish repetitive tasks within games, and so may use exactly the same set of commands to accomplish their goals: for example, always clicking on exactly the same pixel within the image. This is something a human is unlikely to do or may even be incapable of doing with a high degree of accuracy.
- *Network traffic*: One of only a few papers to address bot detection in games "Identifying MMORPG Bots: A Traffic Analysis Approach" by Chen et al. [2006] suggests that a traffic-level detection system is possible. Bot-generated traffic differs from human-generated traffic with respect to the regularity of the release time of commands, the trend and magnitude of traffic bursts in multiple time scales, and the sensitivity of interaction to network responsiveness.

If the game software has the capacity for interplayer chat (which most do), engaging the player in a conversation may reveal his true nature. However, chat bots exist and

become increasingly better at mimicking an interhuman conversations [Shawar and Atwell 2005; Crews 2006]. They are often incorporated into game bots as an antidetection measure. Additional approaches to avoiding bot detection can be clearly seen from analyzing bot- detection methods. Bots should be run in a process with a randomly generated name and always for short periods of time, not to exceed a few hours. Bots' actions should be randomized both in terms of commands used and spatial and temporal decisions made.

3. PREVENTING BOT PARTICIPATION

With the steady increase in the popularity of games and services via the Internet, the problem of securing such services from automated attacks became apparent. In order to protect limited computational resources against the growing number of human-impersonating bots, a methodology became necessary to discriminate between bots and people [Pope and Kaur 2005].

In 1950, Alan Turing published his best-known paper "Computing Machinery and Intelligence" in which he proposed evaluating the abilities of an artificially intelligent machine on how closely it could mimic human behavior [Turing 1950]. The test, which is now commonly known as the Turing test, is structured as a conversation and can be used to evaluate multiple behavioral parameters, such as agent's knowledge, skills, preferences, and strategies [French 2000]. In essence it is the ultimate multimodal behavioral biometric, postulated to make it possible to detect differences between man and machine.

The theoretical platform for an Automated Turing Test (ATT) was developed by Naor [1996]. The following properties were listed as desirable for the class of problems that can serve as an ATT:

- Many instances of a problem can be automatically generated along with their solutions.
- Humans can solve any instance of a problem quickly and with a low error rate. The answer should be easy to provide, either by a menu selection or by typing a few characters.
- The best-known artificial intelligence (AI) programs for solving such problems fail a significant percentage of times, despite full disclosure as to how the test problem was generated.
- The test problem specification needs to be concise in terms of description and area used to present the test to the user.

Since the initial paper by Naor, a great deal of research has been done in the area, with different researchers frequently inventing new names for the same concept of human/machine disambiguation [Baird and Popat 2002; Sampson 2006]. In addition to ATT, the developed procedures are known by such names as "reversed Turing test" (RTT) [Coates et al. 2001]; "human interactive proof" (HIP) [Chellapilla et al. 2005]; "mandatory human participation" (MHP) [Xu et al. 2003; and the "completely automated public Turing test to tell computers and humans apart" (CAPTCHA) [Ahn 2004; Ahn et al. 2004]. In this article we often refer to tests aimed at telling bots and humans apart as CAPTCHAs due to the recent popularity of the term.

As ongoing developments in AI research allow some tests to fail [Chellapilla and Simard 2004; Mori and Malik 2003; Aboufadel et al. 2005; Moy et al. 2004], research continues to take place on developing more secure and user friendly ways of telling

machines and humans apart [Rui et al. 2005; Chellapilla et al. 2005a; 2005b; Wang et al. 2006; May 2005; Lopresti 2005]. Such tests are always based on as yet unsolved problem in AI [Ahn et al. 2003]. Frequent examples include pattern recognition, in particular character recognition [Bentley and Mallows 2006; Baird and Riopka 2005; Baird et al. 2005a; 2005b; Chew and Baird 2003; Simard et al. 2003; Liao and Chang 2004]; or image recognition [Chew and Tygar 2004; Liao 2006; Dailey and Namprempre 2004]; some CAPTCHAs are based on recognition of different biometrics such as faces [Misra and. Gaj 2006; Rui and Liu 2003a; 2003b]; voices [Kochanski et al. 2002; Chan 2003]; and handwriting [Rusu and Govindaraju 2004; 2005]; the following types of tests have been experimented with as well [Hall 2006]:

- *reading* a password displayed as a cluttered image;
- *identifying* complex shapes;
- *rendering spatial* text images from 3D models;
- *quizzing* visual or audio puzzles or trivia questions;
- *matching* common themes for a set of related images;
- *navigating* virtual reality in a 3D world;
- *using* media files collected from the real world, particularly the web naturally; and
- *incorporating* an implicit test into the web page navigation system [Baird and Bentley 2005].

4. EMBEDDED BOT DETECTION

This work was inspired by the idea of developing implicit human-machine disambiguation procedures and expands on it to provide seamless embedded non-interactive and continuous testing. In particular, we developed tests for game environments in which the distractive nature of typical tests is particularly detrimental. We are most interested in applying our techniques to card games such as poker, in which bots have been shown to pose the greatest threat to the integrity of the game [Yampolskiy 2007].

A classical CAPTCHA algorithm can be summarized as follows:

- (1) Computer generates a test instance
- (2) Test is shown to the human/bot
- (3) Human/bot attempts to solve the instance of the test
- (4) Human/bot reports supposed solution to the computer
- (5) Computer evaluates the submitted solution
- (6) Computer reports the result of evaluation to the human/bot and allows or blocks access to a resource based on the result

Figure 1 provides a visual representation of the testing procedure.

We propose integrating the testing procedure as part of the card reading step performed by the player during the game. The identification of the card itself becomes a test that distinguishes bots from legitimate human players. Figure 2 demonstrates an embedded test which, if properly solved, reveals that the card is a king of hearts. Any well- developed text distortion technique employed in traditional CAPTCHA tests can be utilized in our testing procedure.




Fig. 1: Typical human-machine testing algorithm.




Fig.2. Left: Test embedded in a playing card; right: solution – king of hearts.

Our proposed embedded noninteractive test works as follows:

- (1) computer generates a test instance with solution corresponding to the card dealt;
- (2) test is shown to the human/bot;
- (3) human/bot attempts to solve the test; and
- (4) future game decisions of human/bot are shaped by information obtained while solving the test.




Fig. 3. A selection of different test-cards.

If the test is solved correctly, future decisions of the human/bot can be intelligent, as the human/bot has all the necessary information to make decisions. Otherwise the information on which decisions are based is faulty, and hence decisions are not optimal and the human/bot is essentially acting unintelligently. At no point does a human have to explicitly state his perceived solution for the test-reducing amount of distraction via reduction in the interaction. Figure 3 demonstrates a number of different test-cards utilized in our experiments. Both private and community cards can be encoded using the proposed methodology.

Our methodology has a number of advantages such as

- (1) There is less distraction from the task at hand, as the test is embedded in the application and is not a separate task requiring a human to perform an unrelated activity while taking a forced break from the main application.
- (2) The bot does not know if the test was passed or failed, and so can't learn from its mistakes to improve its performance for future presentations of the test.
- (3) Bots are not just prevented from obtaining a resource, but are actually punished for trying to access a resource, making it less likely that future attempts to obtain the resource will follow.
- (4) A human who realizes that he or she has solved the test incorrectly can resolve it correctly at a future point.

(5) Since no solution for a test is provided, outsourcing the test becomes infeasible, as any answer generated by the solver-for-hire will be accepted but would not be verifiable as accurate.

Figure 4 shows a full ring poker table in an online casino with the players' private cards encoded as tests for telling humans/bots apart. Players have numerous opportunities to solve the test, and can even double-check their answers to be completely sure.

The testing can be made continuous for the duration of the poker hand by making all of the community cards encoded as tests as well. Figure 5 demonstrates the same poker room interface with the three flop cards encoded to prevent bots from playing or collecting information about the game. The two still unrevealed community cards (turn and river) provide additional tests, ensuring continuous human participation for the duration of the game.



Fig. 4: The players' private cards with embedded tests.




Fig.5. Private and community cards with embedded tests.

Experimental results obtained with our methodology are encouraging. A group of volunteers achieved a very high rate of correct card recognition; they also happened to be enthusiastic online poker players. After an initial learning curve of 15 minutes, all of our 5 volunteers were able to identify a card correctly in just 1 second with 99% accuracy. The other 1% corresponded to the most difficult tests and required an additional second or two to identify the card correctly. As time progressed our volunteers also reported that the testing procedure started to become less noticeable, to the point of being nonobtrusive. Encoding, which our test uses to convert a card's suite and rank to a test format, can be done using any well-developed CAPTCHA methodology such as text morphing. As a result, the performance of human subjects on our tests is equivalent to that on a prototype CAPTCHA test, which is known to be acceptable to wide groups of human users [Rui et al. 2005].

5. CONCLUSIONS

Use of bots, in particular as assistants to human players, is becoming very popular across multiple game genres, from board games such as chess to first-person shooters such as *Doom*. While some argue that bots are like digital steroids for cyber-athletes, human-bot teams can be beneficial to the game. Bots can be thought of as a feature, and not a problem, since they enhance the level of play and make the game more interesting. As long as all players have an equal opportunity to enhance their play, it should only make the game more competitive, not less interesting, for human players.

However, participation of independent bots in most games is undesirable and should be limited to for-bots-only servers run by bot-development enthusiasts. Our solution for preventing independent bots from participating in human game networks works particularly well in most card games. With our methodology bots become a beneficial feature of the game, as they lose money to real human players including the inexperienced beginners, as opposed to acting as emotionless and tireless predators on the weak.

In this article we have demonstrated how an embedded noninteractive test can be used to prevent automatic artificially intelligent players from illegally participating in online game-play. Our solution has numerous advantages over traditional tests, such as its non-obtrusive nature, continuous verification, and simple noninteractive and outsourcing-proof design. However as with all methods that depend on limitations of current technology, a day will come when artificially intelligent machines will be able to perform at a level that is indistinguishable from their human counterparts.

REFERENCES

- ABOUFADEL, E. F., OLSEN, J., AND WINDLE, J. 2005. Breaking the Holiday Inn priority club CAPTCHA. *College Mathematics J.* (March).
- AHN, L. V. 2004. Utilizing the power of human cycles. Thesis proposal. Carnegie Mellon Univ., Pittsburgh, PA, May.
- AHN, L. V., BLUM, M., AND LANGFORD, J. 2004. How lazy cryptographers do AI. *Communications of the ACM* (Feb.).
- AHN, L. V., BLUM, M., HOPPER, N., AND LANGFORD, J. 2003. CAPTCHA: Using hard AI problems for security. In *Proceedings of the 2003 Eurocrypt Conference*.
- BAIRD, H. S. AND BENTLEY, J. L. 2005. Implicit CAPTCHAs. In *Proceedings of the SPIE/IS&T Conference on Document Recognition and Retrieval XII* (San Jose, CA, Jan.).
- BAIRD, H. S., MOLL, M. A., AND WANG, S.-Y. 2005a. A highly legible CAPTCHA that resists segmentation attacks. In *Human Interactive Proofs*. Springer Verlag, New York.
- BAIRD, H. S., MOLL, M. A., AND WANG, S.-Y. 2005b. ScatterType: A legible but hard-to-segment CAPTCHA. In *Proceedings of the Eighth International Conference on Document Analysis and Recognition* (Aug. 29 - Sept. 1), 935- 939.

- BAIRD, H. S. AND POPAT, K. 2002. Human interactive proofs and document image analysis. In *Proceedings of the 5th International Workshop on Document Analysis Systems* (Aug.19-21), 507-518.
- BAIRD, H. S. AND RIOPKA, T. 2005. ScatterType: A reading CAPTCHA resistant to segmentation attack. In *Proceedings of the SPIE/IS&T Conference on Document Recognition and Retrieval XII* (San Jose, CA, Jan.).
- BENTLEY, J. AND MALLOWS, C. L. 2006. CAPTCHA challenge strings: Problems and improvements. In *Proceedings of the Conference on Document Recognition & Retrieval* (Jan. 18-19).
- BLAU, H. 2002. The human nature of the bot: A response to Philip Auslander. *J. Performance and Art* 24, 1 (Jan.), 22-24.
- BYL, P. B.-D. 2004. *An Overview of Non-Player Characters in Games, Programming Believable Characters for Computer Games*. Charles River Media.
- CHAN, T.-Y. 2003. Using a text-to-speech synthesizer to generate a reverse Turing test. In *Proceedings of the 15th IEEE International Conference on Tools with Artificial Intelligence (ICTAI '03)*, 226.
- CHELLAPILLA, K., LARSON, K., SIMARD, P., AND CZERWINSKI, M. 2005a. Designing human friendly human interaction proofs (HIPs). In *Proceedings of the Conference on Human Factors in Computing Systems*, ACM, New York.
- CHELLAPILLA, K., LARSON, K., SIMARD, P., AND CZERWINSKI, M. 2005b. Computers beat humans at single character recognition in reading based human interaction proofs (HIPs). In *Proceedings of the Second Conference on Email and Anti-Spam* (July 21-22).
- CHELLAPILLA, K. AND SIMARD, P. 2004. Using machine learning to break visual human interaction proofs (HIPs). In *Advances in Neural Information Processing Systems 17*, MIT Press, Cambridge, MA.
- CHEN, K.-T., JIANG, J.-W., HUANG, P., CHU, H.-H., LEI, C.-L., AND CHEN, W.-C. 2006. Identifying MMORPG bots: A traffic analysis approach. In *Proceedings of the ACM SIGCHI Conference* (Los Angeles, CA, June), ACM, New York.
- CHEW, M. AND BAIRD, H.S. 2003. Baffletext: A human interactive proof. In *Proceedings of the SPIE-IS&T Conference on Electronic Imaging, Document Recognition and Retrieval, X* (Jan.), 305-316.
- CHEW, M. AND TYGAR, J. D. 2004. Image recognition CAPTCHAS. In *Proceedings of the 7th International Information Security Conference* (Sept.), Springer, New York, 268-279.
- COATES, A. L., BAIRD, H. S., AND FATEMAN, R. J. 2001. Pessimist print: A reverse Turing test. In *Proceedings of the Sixth International Conference on Document Analysis and Recognition* (Seattle, WA, Sept.), 1154--1158.
- CREWS, P. 2006. Protochat: An exploration of natural language processing. In *Proceedings of the 2006 CCEC Symposium*. Available at: http://symposium.ccec.unf.edu/cd/papers/Protochat_PCrews.pdf, 2006.
- DAILEY, M. AND NAMPREMPRE, C. 2004. A text graphics character CAPTCHA for password authentication. In *Proceedings of the IEEE Region 10 Conference (TENCON, Nov. 21-24)*, 45- 48.
- DOYLE, P. 1999. Virtual intelligence from artificial reality: Building stupid agents in smart environments. In *Proceedings of the AAAI '99 Spring Symposium on Artificial Intelligence and Computer Games* (March).
- ENRIQUE, S., WATT, A., MADDOCK, S. C., AND POLICARPO, F. 2003. Using synthetic vision for autonomous non-player characters. *Inteligencia Artificial, Revista Iberoamericana de Inteligencia Artificial* 21 (2003), 19-25.
- FREEMAN, D. 2004. Creating emotion in games: The craft and art of emotioneering. *ACM Computers in Entertainment* 2, 3 (July).
- FRENCH, R. 2000. *The Turing test: The first fifty years*. *Trends in Cognitive Sciences* 4, 3, 115-121.
- GOLLÉ, P. AND DUCHENEAUT, N. 2005a. Preventing bots from playing online games. *ACM Computers in Entertainment* 3,3 (July).
- GOLLÉ, P. AND DUCHENEAUT, N. 2005b. Keeping bots out of online games. In *Proceedings of the ACM SIGCHI International Conference on Advances in Computer Entertainment Technology* (Valencia; Spain, June 15-17).
- HALL, R.V. 2006. CAPTCHA as a Web security control. www.richhall.com/isc4350/captcha_20051217.htm. Retrieved Oct. 26, 2006.
- KLINE, C. AND BLUMBERG, B. 1999. The art and science of synthetic character design. In *Proceedings of the Symposium on AI and Creativity in Entertainment and Visual Art* (Edinburgh).
- KOCHANSKI, G., LOPRESTI, D., AND SHIH, C. 2002. A reverse Turing test using speech. In *Proceedings of the International Conferences on Spoken Language Processing* (Denver, CO), 1357-1360.
- LAIRD, J. AND DUCHI, J. 2000. Creating human-like synthetic characters with multiple skill levels: A case study using the soar quakebot. In *Proceedings of the 2000 AAAI Fall Symposium: Simulating Human Agents*, M. Freed (ed.).
- LAIRD, J. E. 2001. Using computer games to develop advanced AI. *Computer* 34, 7, 70-75.
- LIAO, W.-H. 2006. A Captcha mechanism by exchange image blocks. In *Proceedings of the 18th International Conference on Pattern Recognition (ICPR'06)*, 1179-1183.

- LIAO, W.-H. AND CHANG, C.-C. 2004. Embedding information within dynamic visual patterns. In *Proceedings of the IEEE International Conference on Multimedia and Expo (ICME '04, June 27-30)*, 895- 898.
- LOPRESTI, D. 2005. Leveraging the CAPTCHA problem. In *Proceedings of the Second HIP Conference*.
- MAY, M. 2005. Inaccessibility of CAPTCHA. Alternatives to visual Turing tests on the Web. W3C Working Group Note. www.w3.org/TR/turingtest/, Nov. 2005.
- MISRA, D. AND GAJ, K. 2006. Face recognition CAPTCHAs. In *Proceedings of the International Conference on Telecommunications, Internet and Web Applications and Services (AICT-ICIW '06, Feb. 19-25)*, 122.
- MORI, G. AND MALIK, J. 2003. Recognizing objects in adversarial clutter: Breaking a visual CAPTCHA. In *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition (June 18-20)*, I-134- I-141.
- MOWBRAY, M. 2002. Ethics for bots. Tech. Rep. HPL-2002-48R1, HP Labs 2002. <http://www.hpl.hp.com/techreports/2002/HPL-2002-48R1.html>. Retrieved Jan. 10, 2007.
- MOY, G., JONES, N., HARKLESS, C., AND POTTER, R. 2004. Distortion estimation techniques in solving visual CAPTCHAs. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2004, June 27-July 2)*, II-23- II-28.
- MOYER, C. 2007. How intelligent is a game bot, anyway? <http://www.tcnj.edu/~games/AIGames/papers/Moyer.html>, Retrieved Jan. 9, 2007.
- NAMEE, B. M., DOBBYN, S., CUNNINGHAM, P., AND O'SULLIVAN, C. 2003. Simulating virtual humans across diverse situations. In *Proceedings of the Conference on Intelligent Virtual Agents, Lecture Notes in AI*, Springer Verlag, New York, 159 - 163.
- NAOR, M. 1996. Verification of a human in the loop or identification via the Turing test. http://www.wisdom.weizmann.ac.il/~naor/PAPERS/human_abs.html, 1996. Retrieved Oct.7, 2006.
- POKER-EDGE.COM. 2006. Stats and analysis. <http://www.poker-edge.com/stats.php>. Retrieved June 7, 2006.
- POKERPROPHECY. 2006. <http://www.pokerprophecy.com>. Retrieved Sept. 26, 2006.
- POPE, C. AND KAUR, K. 2005. Is it human or computer? Defending E-commerce with Captchas. *IT Professional (March/April)*, 43-49.
- RUI, A. Y. AND LIU, Z. 2003a. ARTIFACIAL: Automated reverse Turing test using FACIAL features. In *Proceedings of the 11th ACM International Conference on Multimedia (Berkeley, CA)*, ACM, New York, 295-298.
- RUI, Y. AND LIU, Z. 2003b. Excuse me, but are you human? In *Proceedings of the 11th ACM International Conference on Multimedia (Berkeley, CA)*, ACM, New York, 462-463.
- RUI, Y., LIU, Z., KALLIN, S., JANKE, G., AND PAYA, C. 2005. Characters or faces: A user study on ease of use for HIPs. In *Proceedings of the 2nd International Workshop on Human Interactive Proofs (Lehigh University, Bethlehem, PA, May 18-20)*.
- RUSU, A. AND GOVINDARAJU, V. 2005. A human interactive proof algorithm using handwriting recognition. In *Proceedings of the 8th International Conference on Document Analysis and Recognition (Aug. 29-Sept. 1)*, 967- 971.
- RUSU, A. AND GOVINDARAJU, V. 2004. Handwritten CAPTCHA: Using the difference in the abilities of humans and machines in reading handwritten words. In *Proceedings of the 9th International Workshop on Frontiers in Handwriting Recognition (IWFHR-9 2004, Oct. 26-29)*, 226-231.
- SAMPSON, R. M. 2006. Reverse Turing tests and their applications. <http://www-users.cs.umn.edu/~sampa/research/ReverseTuringTest.PDF>. Retrieved Oct. 8, 2006.
- SHAWAR, B. A. AND ATWELL, E. 2005. A chatbot system as a tool to animate a corpus. *ICAME J.* 29, 5-24.
- SIMARD, P. Y., SZELISKI, R., BENALOH, J., COUVREUR, J., AND CALINOV, I. 2003. Using character recognition and segmentation to tell computers from humans. In *Proceedings of the 7th International Conference on Document Analysis and Recognition (Aug.)*
- TURING, A. 1950. Computing machinery and intelligence. *Mind*, 433-460.
- WANG, S.-Y., BAIRD, H. S., AND BENTLEY, J. L. 2006. CAPTCHA challenge tradeoffs: Familiarity of strings versus degradation of images. In *Proceedings of the 8th International Conference on Pattern Recognition (Aug. 20-24)*, 164-167.
- Windholdem detection avoidance. 2006. <http://www.winholdem.net/antidetct.html>, Retrieved Nov. 26, 2006.
- XU, J., LIPTON, R., ESSA, I., SUNG, M., AND ZHU, Y. 2003. Mandatory human participation: A new authentication scheme for building secure systems. In *Proceedings of the 12th International Conference on Computer Communications and Networks (ICCCN 2003, Oct. 20-22)*, 547- 552.
- YAMPOLSKIY, R. Online Poker Security: Problems and Solutions. *North American Simulation and AI in Games Conference (GAMEON-NA2007)*. Gainesville, Florida. September 10-12, 2007.

Received January 2007; accepted June 2007