Ubiquitous Human Computing

Citation

Jonathan Zittrain, Ubiquitous Human Computing, 366 Phil. Transactions of the Royal Soc'y 3813 (2008).

Published Version

doi:10.1098/rsta.2008.0116

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:10611809>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

LEGAL RESEARCH PAPER SERIES

Paper No 32/2008

June 2008

Ubiquitous Human Computing

JONATHAN ZITTRAIN

This paper can be downloaded without charge from the
Social Science Research Network electronic library at:
<http://ssrn.com/abstract=1140445>

An index to the working papers in the
University of Oxford Legal Research Paper Series is located at:
<http://www.ssrn.com/link/oxford-legal-studies.html>

Ubiquitous Human Computing

BY JONATHAN ZITTRAIN[†]

*University of Oxford, Oxford Internet Institute
1 St Giles, Oxford, OX1 3JS, United Kingdom*

Ubiquitous computing means network connectivity everywhere, linking devices and systems as small as a thumb tack and as large as a worldwide product distribution chain. What could happen when *people* are so readily networked? This essay explores issues arising from two possible emerging models of ubiquitous human computing: fungible networked brainpower and collective personal vital sign monitoring.

Keywords: human computing; ubiquitous computing; law; ethics; predictions

1. Introduction

We are nearing a world of cheap and plentiful sensors, fast processors, and high-speed saturating wireless networks. Together these may allow computing to be far from the people initiating it and using its results. They also may allow much human *thinking* to be as far as we like from the people initiating it and using its results. Networks connect people as well as devices, and when they are cheap and easy to use it means that those intellectual tasks more efficiently performed elsewhere by other people can be broken out and distributed.

Sensors can connect people's bodies as well as their brainpower to the network. The sensors we currently use to remotely monitor industrial equipment, automobile traffic flows, and livestock provide a model for deployment to people, recording and relaying a spectrum of individuals' vital signs for a variety of purposes. Both forms of human computing – distributed networks of people who think in small, discrete slices for others, and distributed networks of people comprising a collective pulse for a particular geographic region or polity – raise ethical issues. Although we cannot predict exactly the issues that will arise, if we can forge a coherent philosophy of what we want and what we cannot accept in these areas, we will find these networks easier to regulate as they come about.

[†] Professor of Internet Governance and Regulation, University of Oxford. I thank James Edelman, Elisabeth Oppenheimer, and participants in the Oxford Internet Institute research seminar and World Economic Forum 2008 Annual Meeting workshop, "How Science Will Redraw the Business Landscape of the 21st Century," for helpful discussions about this essay.

2. Distributed human brainpower

We are in the initial stages of distributed human computing that can be directed at mental tasks the way that surplus remote server rackspace or Web hosting can be purchased to accommodate sudden spikes in Internet traffic (von Ahn 2005; Hewlett Packard (HP) 2008) or PCs can be arranged in grid computing configurations, each executing code in an identical virtual environment (International Business Machines (IBM) 2006). At some fast food drive-throughs, the microphone and speaker next to the marquee menu are patched through to an order-taker thousands of miles away. That person types up the requested order and dispatches it back to a screen in the food preparation area of the restaurant while the car is idling (Richtel 2006). Services like LiveOps recruit workers for such mental contracting tasks (LiveOps 2008a). Applicants to LiveOps navigate a fully automated hours-long vetting system that tests their skills and suitability. Out of 2,000 applicants per week, roughly 40 emerge for a second round of interviews by LiveOps managers (LiveOps 2008b).

Those who succeed and become contractors for firms like LiveOps encounter an unusual combination of freedom and control. They can work whenever they like, wherever they like, for as much or as little time as they like. When they log in to work they choose from a menu of assignments tailored to their skill and reputation levels. These might include taking pizza orders, placing sales calls, lobbying for a political candidate, or handling customer service inquiries. Then there is the control: every call and transaction is measured and recorded. Interactions can be monitored live by fellow LiveOps mentors or official LiveOps managers, or pulled up later as part of a larger assessment of contractors' work. Judgments are developed and recorded about contractors' performance, such that an incoming pizza order can be routed to the best pizza-order-taker – who may not be the same as the best political campaigner (Hornik 2007). Contractors can be de-accredited at any time.

The rise of the mental contractor factotum is not limited to the disaggregation of human-to-human consumer transactions, which might be viewed as merely an extension of phenomena like the outsourcing of medical X-ray analyses to less costly overseas radiologists (Nyberg & Lanzieri 2006-07). Cheap networks mean that nearly any mental task can become unbundled, no matter how minor. Marketplaces like Amazon's Mechanical Turk enumerate "HITs" – human intelligence tasks – for sale one unit at a time, from as low as \$0.01 to as high as \$10.00. Recent HITs include using one's mouse to highlight the specific pavement area of bus stops on pictures of city streets; leaving positive topical comments on particular blogs; and collecting case write-ups from MBA students (Amazon Mechanical Turk 2008). On the upper end of the scale are marketplaces like InnoCentive, where bounties are placed on answers to minor scientific problems, and people from around the world compete to be the first to answer, netting the bounty in exchange for giving all intellectual property rights in the solution to the purchaser. One can visualize in the near future a subway car packed with people, each far less attuned to the local environment and to each than with today's distractions of newspapers and iPods. Instead they will stare into screens even for just a few minutes and earn as much money in that time as their respective skills and stations allow.

Other human intelligence tasks take place without payment to, or even the awareness of, the people performing the tasks. Von Ahn and Dabbish’s (2004) pioneering ESP Game comprises a Web site where people are shown images and then guess labels for them. When their guesses match those of other anonymous players, they score points and the images are labeled with their tags. The points have no extrinsic value. Tens of thousands of people have played the game, many for 40 hours a week. They created more than 10 million labels within a few months (von Ahn 2006). Google has licensed von Ahn’s technology for use in sustaining its image search (Google Image Labeler 2008). Von Ahn has also created reCAPTCHA, a system that rides the coattails of the standard captcha, where Web surfers are asked to type in a word to match a difficult-to-read graphic image of a word in order to prove that they are people and not robots. reCAPTCHA provides two words instead of one in random order: one is the captcha test word, and the other is a word from a scanned text that has defied attempts at optical character recognition. Through reCAPTCHA, human brains are employed one word at a time to complete the accurate scanning of old texts somewhere in the world (Captcha 2008).

Systems like these can encompass myriad tasks such as language translation, household or border monitoring (“How many people do you see in this video, and what are they doing?”), and classifying over one million newly-discovered galaxies (Zittrain 2008; Distributed Computing 2008). Some may be created for purposes likely at odds with the values of their contributors. For example, spammers have sought ways to fill out captchas that do not require hiring humans. They have created Web sites and software that offer free pornography. A user is prompted to answer a captcha before he or she can proceed – with the captcha’s image drawn from an unrelated site where the spammer wishes to post a comment. When the user answers the captcha, the answer is transmitted back to the spammer’s computer, which can then provide the answer to the target site providing the initial image (British Broadcasting Corporation (BBC) 2007).

3. Distributed human sensors

One form of ubiquitous human computing – the distributed human brainpower described above – stands to dissolve some of the longstanding networks of the traditional firm and its physical workplace and career workforce. Another form – distributed human sensors – will create a new kind of network to govern and monitor citizens. We can contribute to a fungible set of disembodied pulses as well as to one of disembodied mindpower.

Today we have remote sensors to monitor the location and status of personal and corporate assets, including cars, shipping containers, and cattle (Radio Frequency Identification (RFID) 2008). Early theft deterrent systems designed to locate and disable stolen vehicles have diversified into full-fledged traffic monitoring schemes. With enough cars reporting location and speed to each other or to a central coordinator, traffic snarls can be instantly identified – and nearby cars advised to re-route both to help individual drivers and to ease overall congestion (Johnson 2007). Today several states place radio frequency identifiers into passports, allowing them – and their holders – to be counted and identified at short distances (U.S. Department of State 2007). Mobile phone providers use cellular tower triangulation to count the location,

number, and identity of gatherings of people, knowing how many (and perhaps who) attended a sporting event or a political rally (Taipei Times 2008). A variety of needs could leverage these programs into a ubiquitous weather map of human vital signs, inferring patterns of illness and mass emotion through basic physical indicators. Here the human sensor networks would be notable for precisely for the variations from one person or group to the next, rather than human computing’s implicit aim to make mental cycles an undifferentiated commodity.

Initial motivating factors include older citizens seeking continued independence from retirement homes by having their vital signs continuously monitored for signs of trouble (Dulay *et al.* 2005; Komninos and Stamou 2006); health insurance that would trade such monitoring in exchange for discounts on premiums; and a desire for early warning against (and epidemiological tracing of) either naturally- or artificially-sourced communicable disease ranging from bird flu to hemorrhagic fever. Indeed, some systems not only monitor vital signs but allow for the remotely-administered injection of drugs into people connected to the network (Brown and Adams 2008). In addition to citizens who choose to subscribe to such monitoring, coordinated and mandated deployment could take place within military or law enforcement circles to monitor their soldiers and officers, or public health authorities insisting that front line medical professionals be wired up.

Once achieving critical mass, a human sensor net with variables as modest as personal identity, pulse rate, and temperature could become vital for many uses, especially in states that strike a balance towards collective rather than individual privacy when the two are in tension – or that are outright authoritarian. Law enforcement will take readings on the public or officers on patrol for warning of riots or other anxiety-producing trouble, and first responders could use the sensor net to assure the evacuation of buildings and to identify people in need of assistance in the run up to or wake of a mass casualty event.

As we become accustomed to less controversial uses of vital sign mapping, other variations could be incrementally introduced. If unique identifiers were integrated into individual telemetry, legal privileges could be authoritatively associated with each person – with extra scrutiny for those who fail to have a transponder at all. Biometric identifiers are entering into widespread use in public databases (Arena & Cratty 2008; Ignatius 2008). Just as police today can scan automatically the license plates of a row of cars to discover if any are stolen or not insured (Farivar 2005), drivers will be scanned at checkpoints to see if any are unlicensed, and states wishing to crack down on undocumented immigration could require a check of someone’s status as part of any financial transaction, or entry to any public establishment. The venerable practice of putting out an “all points bulletin” for a person of interest will be transformed into asking millions of distributed scanners to check for a particular identity and summon police if it is found. The parallel diffusion of radiofrequency identifiers into household products will allow researchers and plaintiffs’ lawyers to run very large scale statistical inquiries associating particular products with illness or death.

4. Issues arising from ubiquitous human computing

The efficiencies gained by treating distributed brainpower the same as distributed computing power are tempered by the prospect of profound if more inchoate dignitary harms.

First, when tasks and taskmasters can change or be changed by the minute, mental contractors cultivate no particular sense of affiliation to an employer or to a team, and they gain no sense of the larger enterprise for which they have been asked to perform just one small step. To be sure, LiveOps itself experiences very low turnover (LiveOps 2008c), indicating satisfaction – or at least the absence of desirable alternatives – among its contractors. But the disconnection between task and project, or task and project commissioner, can be even more severe with brainpower-harnessing tools like the ESP Game that need not innately signal that its use is performing epiphenomenal work for someone else. Alienation from the gestalt of one’s work is not a new phenomenon; the assembly line is the hallmark of an industrial economy. But distributed human computing risks extending the assembly line from the mechanical to the intellectual, spawning a new class of knowledge workers whose work lives are fully atomized, an existence blinkered even from fellow assembly line mates. The challenge will be to ensure that the technologies that supplant full time employment still enable human relationships to help the work stay meaningful and fulfilling, and to facilitate a form of career development among workers whose skills grow.

Second, human computing’s freedom to work when and where desired creates new incentives and justifications for monitoring and control of that work beyond that of the traditional workplace. With the line between home and work fully transformed from physical to functional, those who commission the work can rightly claim to inspect all work-related activities. As each action and transaction is monitored and analyzed, the level of scrutiny we might apply to sensitive jobs like policing or day care can also be cost-effectively devoted to run-of-the-mill business transactions and communications. Such monitoring will be qualitatively different than that of the traditional workplace. It will likely be based upon metrics collected and analyzed by a higher force that has never met or engaged with the worker, and those metrics assessed through Bayesian-style statistics may not provide a narrative rationale for a judgment – only a statistical claim about its truth. Psychologists and political scientists agree that some amount of nonaccountability – the ability to tell a white lie, or to refuse a request with minimal explanation – actually facilitates social interaction, but ubiquitous monitoring can eliminate this slack space. (Iachello et al. 2005,) Gains in service quality through such strictures are thus offset to an unknown degree by the crush of the panopticon (Foucault 1979; Boyle 1997.) In discovering how to farm out through the Net the tasks that computers cannot do, we may find ourselves treating – and making – people more like computers.

Third, disembodied HITs can deprive people of the chance to make judgments about the moral valence of their work. One can weigh whether to embark on a career with, say, a tobacco company, and even where such career choices are economically compelled (or legally compelled, as with military conscription), there are signals to relate one’s work to its effect on the world. For those HITs that are epiphenomenal – gleaned as a by-product of people’s activities rather

than because they aim to perform them – an absence of disclosure deprives people of the freedom to choose the goals that their intelligence will advance. Defense lawyers may be professionally obliged to assist the culpable with avoiding law’s consequences, but even in circumstances where they cannot choose (or cease representing) their clients, they can choose their profession overall, and they can invoke the crucial role their loyalty to clients plays in the larger system (Weiss 2005). For those whose game playing or Internet usage may be directly deployed for groups or causes they find repugnant, there is no such comfort.

How and if these harms arise may depend on whether human computing itself is lessened as artificial intelligence improves, enlarging the set of tasks that can be fully automated by computers alone. In the meantime, human computing can develop humanely.

First, harvesters of human mindpower can be encouraged – or perhaps required – to disclose their activities to those who benefit them. While no law specifically bans or regulates the practice of surreptitious HITs, restrictions found in many states on more general “unfair or deceptive trade practices” could be applied (Federal Trade Commission (FTC) 2006). Whether captcha-solving is deployed to scan an old book for a university library or to support the work of a particular religious institution or political party or spammer, the captcha-solver ought to know. Such knowledge will help develop and reinforce norms about what kinds of borrowed mental cycles are acceptable.

Second, users might be given the opportunity to opt out. Google and other search engines use the corpus of webmasters’ linking behavior to help rank Web sites, but they also respect a robots.txt designation asking that they not harvest links from particular sites at all, and Google’s “#nofollow” designation provides a way for someone publishing a link to indicate that the link’s existence should not be marked as approval of its destination (Koster 2008; Google 2005). Thus people can link to their enemies without elevating them, and more generally can see their contributions to collective computing as a choice freely given rather than one deviously or unavoidably extracted.

Third, those who work in the new HIT markets ought to be able to take their reputations with them, building portfolios that can be offered to competing HIT markets and employers. In short, the constraints associated with physical work environments are absent when mental contractors work wherever they choose, but this need not moot the ideals motivating legal protections aimed to prevent exploitation in the workplace. Minimum wage, maximum working hours, unionization (or at least the ability to know and contact one’s co-workers) may be revisited to see if and how they should be adapted to distributed human computing, still knowing that each lessens freedom of contract. Each state could formulate its own policies, and firms like LiveOps could choose to recruit contractors from various jurisdictions depending on whether they wish to accede to that jurisdiction’s rules. A trickier question is how to treat those people who contribute without payment. Minimum wage requirements could kick in at a penny for a thought, but when they are given freely – whether as Wikipedia edits or as idle play of the ESP game – it would be significantly more intrusive to limit the transaction in the name of paternalism. Moreover, facilitating markets in ubiquitous human computing could affect the

prospects for freely given mental cycles and vice-versa. If Mechanical Turk had preceded Wikipedia, potential contributors to the latter might have been more primed to wonder why they should give away for free what otherwise commands a price, just as some worry that a market in human organs would stymie organ donations (Satel 2007).

For the system of distributed human vital sign sensors, its undesirability might be thought to speak for itself, at least the moment that unique identifiers are introduced. But ubiquitous, unobtrusive sensors fool many of our natural protective mechanisms: we lack the normal social cues that remind us how much information we are broadcasting, and the sensors are easy to forget (Bellotti & Sellen 1993). More important, our views of privacy are malleable (Taylor 2003; Vascellaro 2008), and often depend on the benefit we perceive in return for sharing information (Iachello et al. 2005). Public uptake of the foundational technologies for other reasons, combined with compelling use cases for the deployment into collective vital sign mapping, could bring such a system about. Will people be able to opt, temporarily or permanently, out without fear of extra scrutiny? Will we be able to design feedback systems that remind people of their participation in the system? Many salutary uses, such as early warning for contagious illness, do not require ubiquitous uptake, but others – such as license or status verification – depend on each person emanating his or her digital passport.

Human sensor nets to detect and deal with disease outbreaks, such as *E. coli* from a tainted food source, could also create new vulnerabilities. They could facilitate more tightly coupled systems and thus less margin for error should the sensor net break down. For example, pressure to decentralize meat processing would lessen if outbreaks could be more efficiently contained than prevented in the first instance. Human vital sign monitoring, in addition to helping us respond to problems, could encourage a host of new ones.

As these sensor systems emerge, their ranges of uses and abuses will be largely determined by who runs them – and how widely available and secure their data are. The methods we have developed for computer security today may not work with so many devices broadcasting information over wireless and ad hoc networks (Stajano 2002). Yet solutions that keep such data available only to the few will reduce its generative potential. We might wish anonymized information from the sensor net to be treated as many states treat raw weather data, gathered at government expense and distributed in real time to any who want to review it. The broader the availability of the anonymized data, the more uses can be made of it – and the more the public itself can choose to respond to what is found or synthesized with it. More broadly, we can project from the competing trajectories of today’s computing environment to frame the possibilities for tomorrow’s human one. The Internet faced down proprietary networks like CompuServe and America Online, and the PC largely supplanted dedicated information appliances (Zittrain 2008). But the open and generative need not always trump the sterile and proprietary. Ubiquitous human computing – both as fungible brainpower and as collective physical monitoring – might arise in a proprietary fashion, with a small circle of public and private entities as gatekeepers to its use. Or it could unfold more organically, the way that the Internet’s byways are not owned in traditional ways. A ubiquitous human computing

infrastructure that is “unowned” in the way the current Internet and its protocols are unowned will offer significantly different vectors of abuse and control than one run by a limited set of firms or sovereigns.

5. Conclusion

Over the next fifteen years some of the most notable advances in computing will be in its relationship to people: distributing human mindpower to solve problems both large and tiny, and monitoring and ultimately altering people’s bodies and actions in ways previously impossible. These are not phenomena to be avoided so much as they are to be organized and perhaps regulated so that their ubiquity will enhance rather than debase the human condition.

Shaping them will require an informed and widespread debate with tools drawn from many disciplines. Philosophers will attempt to construct a utilitarian calculus (Anderson 1991). Computing security professionals will ask what information we want to protect, and then seek to construct a security system (Stajano 2002). Without being able to foresee every problem ahead, it makes sense to reflect on the values we consider most important, such as autonomy, privacy, and health, and the case studies to place them appropriately in tension, so we can build and refine systems sensitive to them.

References

- Anderson, B. The Ethics of Research into Invasive Technologies. Rank Xerox Research Center, Technical Report EPC-1991-107 (1991).
- Amazon Mechanical Turk 2008, Amazon Mechanical Turk Welcome.
<http://www.mturk.com/mturk/welcome>.
- Arena, K. & Cratty, C. FBI wants palm prints, eye scans, tattoo mapping. *CNN.com*, February 4, 2008. <http://edition.cnn.com/2008/TECH/02/04/fbi.biometrics/index.html>.
- Bellotti, V. and Sellen, A. Design for privacy in ubiquitous computing environments. In Proceedings of the Third Conference on European Conference on Computer-Supported Cooperative Work (Milan, Italy, September 13 - 17, 1993). G. de Michelis, C. Simone, and K. Schmidt, Eds. ECSCW. Kluwer Academic Publishers, Norwell, MA, 77-92. 1993.
- Boyle, J. Foucault in Cyberspace: Surveillance, Sovereignty, and Hard-Wired Censors, 66 *University of Cincinnati Law Review* 177 (1997).
- British Broadcasting Corporation (BBC). PC stripper helps spam to spread. *BBCNews.com*, October 30, 2007. <http://news.bbc.co.uk/2/hi/technology/7067962.stm>.
- Brown, I. and Adams, A. The Ethical Challenges of Ubiquitous Healthcare. *International Review of Information Ethics* (2008).
- Captcha 2008. The Official CAPTCHA Site. <http://www.captcha.net>.
- Distributed Computing 2008, Distributed Computing—Active Projects—Human.
<http://distributedcomputing.info/ap-human.html>.
- Dulay, N., Heeps, S., Lupu, E., Sharma, O., Sloman, M., and Sventek, J., Autonomic Management for Ubiquitous e-Health Systems, Proceedings of the UK e-Science Programme All Hands Meeting (2005).
- Farivar, C. Grand Theft Auto Meets Robocop. *Wired*, June 17, 2005.
<http://www.wired.com/cars/energy/news/2005/06/67864>.
- Federal Trade Commission (FTC) 2006. Federal Trade Commission Act, Incorporating U.S. SAFE WEB amendments of 2006, 15 U.S.C. Sec. 45. <http://www.ftc.gov/ogc/ftcact.shtm>.
- Foucault, M. *Discipline and Punish: The Birth of the Prison* (Alan Sheridan ed. & trans., 1979).
- Google, Official Google Blog: Preventing Comment Spam. January 18, 2005.
<http://googleblog.blogspot.com/2005/01/preventing-comment-spam.html>.
- Google Image Labeler 2008. Google Image Labeler. <http://images.google.com/imagelabeler>.

Hornik, D. LiveOps: From Next Available Agent to Best Available Agent. June 21, 2007. http://ventureblog.com/articles/2007/06/shameless_selfp_2.php.

Hewlett Packard (HP) 2008. On Demand Solutions from HP. <http://h20219.www2.hp.com/services/cache/10615-0-0-224-121.html>.

Iachello, G., Smith, I., Consolvo, S., Chen, M., and Abowd, G. D. Developing privacy guidelines for social location disclosure applications and services. In Proceedings of the 2005 Symposium on Usable Privacy and Security (Pittsburgh, Pennsylvania, July 06 - 08, 2005).

International Business Machines (IBM), Grid Computing, Past, Present, and Future. June 2006. http://www-03.ibm.com/grid/grid_literature.shtml.

Ignatius, D. Learning to Fight a War. *Washington Post*, February 10, 2008. <http://www.washingtonpost.com/wp-dyn/content/article/2008/02/08/AR2008020802559.html>.

Koster, M. 2008. A Standard for Robot Exclusion. http://www.robotstxt.org/wc/no_robots.html.

Johnson, J. Getting a Read on Congestion. *DC Velocity*, August 2007. http://www.dcvelocity.com/articles/?article_id=1043.

Komninos, A. and Stamou, S. HealthPal: An Intelligent Personal Medical Assistant for Supporting the Self-Monitoring of Healthcare in the Ageing Society. Proceedings of UbiHealth 2006: The 4th International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications, Copenhagen 2006.

LiveOps 2008a. LiveOps On Demand, Virtual Call Center Platform. www.liveops.com.

LiveOps 2008b. LiveOps Agents. http://www.liveops.com/liveops_agents_who.html.

LiveOps 2008c. LiveOps About Us. http://www.liveops.com/news/press_07_12_07.html.

Nyberg, E. & Lanziere, C. American Diagnostic Radiology Moves Offshore, 20 *Journal of Law and Health* 254 (2006-07).

Radio Frequency Identification 2008. RFID Journal—RFID FAQs. <http://www.rfidjournal.com/faq>.

Richtel, M. The Long Distance Journey of a Fast Food Order. *New York Times*, April 11, 2006. <http://www.nytimes.com/2006/04/11/technology/11fast.html>.

Satel, S. Desperately Seeking a Kidney. *New York Times*, December 16, 2007. <http://www.nytimes.com/2007/12/16/magazine/16kidney-t.html>.

Stajano, F. Security for Ubiquitous Computing. 2002.

Taipei Times. China Mobile stuns Davos Forum with private data claims. *Taipei Times*, January 28, 2008.

<http://www.taipeitimes.com/News/worldbiz/archives/2008/01/28/2003399233>.

Taylor, H. Most People Are “Privacy Pragmatists” Who, While Concerned About Privacy, Will Sometimes Trade It Off for Other Benefits. Mar. 19, 2003.

http://www.harrisinteractive.com/harris_poll/index.asp?PID365.

U.S. Department of State. Final Rule: Card Format Passport; Changes to Passport Fee Schedule. 72 Fed. Reg. 74169 (December 31, 2007).

Vascellaro, J. Thinking About Tomorrow. *Wall Street Journal*, January 28, 2008.

von Ahn, L., and Dabbish, L. Labeling Images With a Computer Game, *ACM CHI* 2004.

<http://www.cs.cmu.edu/~biglou/ESP.pdf>. See also The ESP Game: Labeling the Web.

<http://www.espgame.org>.

von Ahn, L. *Human Computing*. 2005. <http://reports-archive.adm.cs.cmu.edu/anon/2005/abstracts/05-193.html>.

von Ahn, L. Games with a Purpose. *IEEE Computer Magazine*, June 2006.

Weiss, M. *Public Defenders: pragmatic and political motivations*. 2005.

Zittrain, J. *The Future of the Internet—And How to Stop It*. 2008.