

Search and Discovery Through Human Computation

Albert Yu-Min Lin, Andrew Huynh, Luke Barrington, and Gert Lanckriet

Introduction

Machines are good at handling huge amounts of data, but they lack the flexibility and sensitivity of human perception when making decisions or observations. To understand human perception, we look toward what defines being human. To sense, observe, and make sense of the world around us, we combine our biological receptors (eyes, ears, etc.) with our cognitive faculties (memory, emotion, etc.). But the memory banks that we pull from to create comparative reasonings are unique from individual to individual. Thus, we each see things in slightly different ways, i.e. what

A.Y.M. Lin (✉)

California Institute for Telecommunications and Information Technology,
UC San Diego Division, 9500 Gilman Drive, La Jolla, CA 92093-0436, USA

National Geographic Society, Washington, D.C., USA
e-mail: a5lin@ucsd.edu

A. Huynh

California Institute for Telecommunications and Information Technology,
UC San Diego Division, 9500 Gilman Drive, La Jolla, CA 92093-0436, USA

Computer Science and Engineering Department, University of California at San Diego,
La Jolla, CA, USA
e-mail: a5huynh@cs.ucsd.edu

L. Barrington

Digital Globe Corporation, Longmont, Colorado, USA
e-mail: lukeinusa@gmail.com

G. Lanckriet

Electrical and Computer Engineering Department, University of California at San Diego,
9500 Gilman Dr, La Jolla, CA 92093, USA
e-mail: gert@ece.ucsd.edu



is beautiful to one person may not be to another. However, there are trends that emerge among our collective human consciousness and efforts to tap a consensus of human perception, i.e. crowdsourcing, depend upon these trends to scale up analytical tasks through massively parallel networks of eyes and minds. This concept of crowd based computing has become an important approach to the inevitable “data avalanches” we face.

The Modern Age of Human Information Processing: More than one quarter of the world’s population has access to the Internet (Internet World Stats 2009), and these individuals now enjoy unprecedented access to data. For example, there are over one trillion unique URLs indexed by Google (Google Blog 2008), three billion photographs on Flickr, over six billion videos viewed every month on YouTube (comScore 2009), and one billion users of Facebook, the most popular social networking site. This explosion in digital data and connectivity presents a new source of massive-scale human information processing capital. User generated content fills blogs, classifieds (www.craigslist.org), and encyclopedias (www.wikipedia.org). Human users moderate the most popular news (www.reddit.com), technology (www.slashdot.org), and dating (www.plentyoffish.com) sites. The power of the internet is the power of the people that compose it, and through it we are finding new ways to organize and connect networks of people to create increasingly powerful analytical engines.

Breaking up the Problem: To combine the large-scale strength of online data collection with the precision and reliability of human annotation, we take a creative approach that brings the data collection process close to humans, in a scalable way



Fig. 1 Ultra-high resolution imagery of Mongolia displayed on the HiperSpace visualization facility at UC San Diego

that can motivate the generation of high quality data. Human computation has emerged to leverage the vast human connectivity offered by the Internet to solve problems that are too large for individuals or too challenging for automatic methods. Human computation harnesses this online resource and motivates participants to contribute to a solution by creating enjoyable experiences, appealing to scientific altruism, or offering incentives such as payment or recognition. These systems have been applied to tackle problems such as image annotation (von Ahn and Dabbish 2004), galaxy classification (www.galaxyzoo.org), protein folding (Cooper et al. 2010), and text transcription (von Ahn et al. 2008). They have demonstrated that reliable analytics can be produced in large scales through incremental contributions from parallel frameworks of human participation.

One approach to human computation motivates participants by creating enjoyable, compelling, engaging games to produce reliable annotations of multimedia data. Markus Krause's chapter (in this book) on gamification provides a brilliant investigation of this specific topic. These "games with a purpose" (von Ahn 2006) have been applied to classify images (von Ahn and Dabbish 2004; von Ahn 2006), text (von Ahn et al. 2006) and music (Mandel and Ellis 2007; Barrington et al. 2012b; Law and vonAhn 2009). In general, these games reward players when they agree on labels for the data and, in turn, collect information that the consensus deems reliable. The goal of these games has been to collect data on such a massive scale that all the available images, text or music content could be manually annotated

by humans. Although simple and approachable online games – “casual games” – have broadened the video gaming demographic (International Game Developers Association 2006), designing a human computation game that meets these data collection goals while being sufficiently attractive to players in massive volumes remains a challenge.

In this chapter we describe several efforts to produce game like frameworks that take on a needle-in-a-haystack problems, often when the needle is undefined. Specifically, we explore innovative networks of human computation to take on the ever expanding data challenges of satellite imagery analytics in search and discovery. We describe frameworks designed to facilitate peer directed training, security through the partitioning and randomization of data, and statistical validation through parallel consensus. In each case it is clear that careful architecture of information piping is a determinate in the success of parallel human computation. We begin with an overview of our initial efforts in satellite remote sensing for archaeology, followed by subsequent experiences in disaster assessment, and search and rescue.

Case Study: Archaeological Remote Sensing

In 2010 we launched “Expedition: Mongolia” as the satellite imagery analytics solution for the *Valley of the Khans Project* (VOTK), an international collaboration between UC San Diego, the National Geographic Society, and the International Association for Mongol Studies to perform a multidisciplinary non-invasive search for the tomb of Genghis Khan (*Chinggis Khaan*). We turned to massively parallel human computation out of frustration from the inability to effectively survey the vast quantity of imagery data through automated or individual means.

Since the invention of photography, aerial images have been utilized in archaeological research to provide greater understanding of the spatial context of ground features and a perspective that accentuates features which are not otherwise apparent (Riley 1987; Bewley 2003; Deuel 1969; Lyons 1977). Buried features can produce small changes in surface conditions such as slight differences in ground level, soil density and water retention, which in turn induce vegetation patterns (cropmarks), create variability in soil color (soilmarks) or even shadows (shadowmarks) that can be seen from above.

The introduction of earth sensing satellites has further contributed to the integration of remote sensing in archaeology (Fowler 1996; Parcak 2009). The ability of detecting features on the ground from space is largely dependent upon the ratio of feature size to data resolution. As sensor technologies have improved, the potential to utilize satellite imagery for landscape surveys has also improved (Wilkinson et al. 2006; Lasaponara and Masini 2006; Blom et al. 2000). In September of 2008 the GeoEye-1 ultra-high resolution earth observation satellite was launched by GeoEye Inc. to generate the world’s highest resolution commercial earth-imaging (at the time of launch) (Madden 2009). Generating 41 cm panchromatic and 1.65 m multispectral data this sensor further expanded the potential of satellite based

archaeological landscape surveys. However, the massive amount of data that is collected each day by these sensors has far exceeded the capacity of traditional analytical processes. Thus, we turn to the crowds to scale human computation towards a new age of exploration.

We construct a massive parallel sampling of human perception to seek and survey the undefined. Specifically, we aim to identify anomalies in vast quantities of ultra-high resolution satellite imagery that represent archaeological features on the ground. Because these features are unknown we are searching for something we cannot predefine. Our internet-based collaborative system is constructed such that individual impact is determined by independent agreement from the “crowd” (pool of other participants who have observed the same data). Furthermore, the only direction that is provided to a given participant comes from the feedback in the form of crowd generated data shown upon the completion of each input. Thus, a collective perception emerges around the definition of an “anomaly”.

The Framework

Ultra-high resolution satellite imagery covering approximately 6,000km² of landscape was tiled and presented to the public on a National Geographic website¹ through a platform that enabled detailed labeling of anomalies.

Within the data interface participants are asked to annotate features within five categories: “roads”, “rivers”, “modern structures”, “ancient structures”, and “other”. For each image tile, participants were limited to create no more than five separate annotations. This limitation was designed to limit the influence that any single individual could have on a given section of imagery (see Fig. 2).

Image tiles (with georeference meta data removed) were distributed to participants in random order. By providing segmented data in random order a collection of participants (or participant with multiple registrations) could not coordinate a directed manipulation of any given location. This was designed to both secure the system against malicious data manipulation as well as to protect the location of potential sites from archaeological looters.

At the onset of the analysis, ground truth information did not exist to provide an administrative source of feedback of the accuracy of analysis to participants. Thus we depend upon peer feedback from data previously collected by other randomly and independent observers of that image tile to provide a consensus based reference to position ones input in relation to the “crowd” (see Fig. 3).

The semi-transparent feedback tags provide a reference to gauge one’s input to the perceptive consensus of a crowd. This reference information cannot be used to change the input provided to that particular image tile, however is designed to influence the participant for the following image tiles. Basing training on an evolving

¹<http://exploration.nationalgeographic.com/mongolia>

NATIONAL GEOGRAPHIC CELEBRATING 125 YEARS +

Connect: [f](#) [t](#) [g+](#) Search Sign In Join

Home Video Photography Animals Environment Travel Adventure Television Kids **Subscribe** Shop

Daily News The Magazine Maps Science Education Games Events Blogs Movies Explorers Apps Trips

Field Expedition: Mongolia

HOME/MAP HOW TO VIDEO BLOG HISTORY AND CULTURE ABOUT THE EXPEDITION BEHIND THE SCIENCE **LOGOUT**

Drag & drop the icons on the map.

- Road
- River
- Modern Structure
- Ancient Structure
- Other

I'M DONE MARKING OBJECTS

Hello albert Yu-Min Lin

Your Rank: Novice 3 Next Rank: Intermediate 1

Images to next rank:

Satellite imagery made available by GeoEye Foundation

JOIN

Join the EXPEDITION

Sign up today and become a part of the expedition.

SIGN UP

SHARE

Share the ADVENTURE

Spread the word to your friends and more.

[f](#) [t](#) [g+](#) [st](#) [su](#) [p](#)

BLOG

"Valley of the Khans" Experts Meet in D.C.
Jun 4, 2012 by Dr. Albert Yu-Min Lin
Two of the world's greatest scholars of Mongol history joined their collaborators NG Emerging...

Heading Out
Aug 15, 2011 by Dr. Kostas Stamatou
My horse was breathing heavily climbing the steep muddy trail. Rocks rolled under his feet and...

A 21st Century Approach to Ground Surveying
Jul 26, 2011 by Andrew Huynh
Our tech for the field increases exponentially every year we come back to Mongolia.

VIEW ALL BLOGS

ABOUT

MISSION STATEMENT

This study aims to utilize modern non-invasive tools in the search for the tomb of Genghis Khan, thus shedding light on Mongolia's rich historical heritage and enabling conservation and education of this rapidly changing landscape.

Learn how to find ANCIENT STRUCTURES

Watch the video for instructions, tips, and tricks on tagging maps.

Fig. 2 User interface for online participants to identify anomalies within randomly presented sub-sectioned satellite imagery (Presented on <http://exploration.nationalgeographic.com/mongolia>)

peer generated data set we allow a form of emergent collective reasoning to determine the classifications, an important design element when searching for something that cannot be predefined.

The emergence of “hotspots” of human agreement also provide a form of validation through agreement among independent observers (a multiply parallel blind test). The mathematical quantification of agreement is the basis for extracting insight from the noisy human data. A detailed investigation of this framework and the role of collective reasoning will be reported in a forthcoming manuscript (Lin et al. 2013).



Fig. 3 Peer based feedback loop (Presented on <http://exploration.nationalgeographic.com/mongolia>)

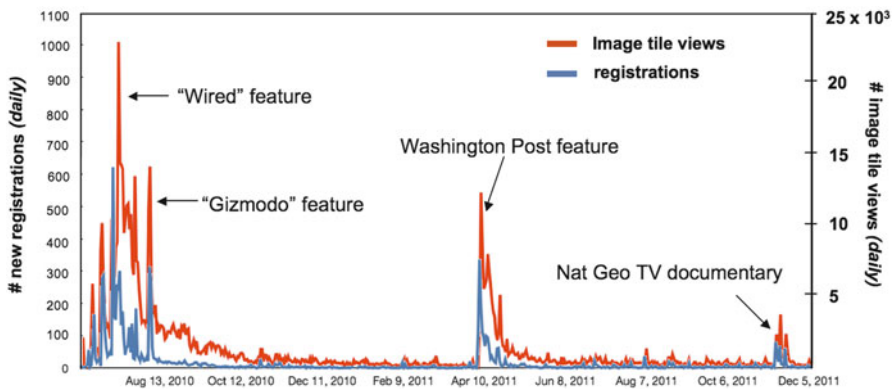


Fig. 4 Registration (blue) and image view (red) statistics across the duration of the experiment

Opening the Flood Gates

Since its launch over 2.3 million annotations from tens of thousands of registered participants were collected. Recruitment was facilitated through public media highlights, i.e. news articles and blogs. These highlighting events provide observable spikes of registration/participation, as seen in Fig. 4. We show this trend to articulate the importance of external communities to drive participation in crowdsourced initiatives.

Overlaying this huge volume of human inputs on top of satellite imagery creates a complex visualization challenge (Huynh et al. 2013) a subset of which is depicted in Fig. 5. While independently generated human inputs are inherently noisy, clusters




Fig. 5 Human generated tags overlaid on satellite imagery showing emergent agreement around features. Tag categories “road” and “ancient” are represented in *red* and *yellow*, respectively. We have explored methods of clustering to define linear features through tags (roads and rivers, Huynh and Lin (2012))

of non-random organization do emerge. Categorical filtering highlights road networks, rivers, and archeological anomalies, respectively.

Guided by this global knowledge of public consensus, we launched an expedition to Mongolia to explore and groundtruthed locations of greatest convergence (defined mathematically through kernel density estimations). From three base camp locations along Mongolia’s Onon River Valley we were restricted to a proximity boundary based upon 1 day’s travel range and limitations associated with extreme inaccessibility. This created an available coverage distance of approximately 400 square miles. Within these physical boundaries we created and explored a priority list of the 100 highest crowd rated locations of archaeological anomalies. The team applied a combination of surface, subsurface geophysical (ground penetrating radar and magnetometry), and aerial (UAV based) survey to ground truth identified anomalies (Lin et al. 2011). Of those 100 locations, over 50 archaeological anomalies were confirmed ranging in origins from the Bronze age to the Mongol period (see example in Fig. 6).

Case Study: Christchurch Earthquake Damage Mapping

Born out of the success of “Expedition:Mongolia” Tomnod Inc. was formed in 2011 to explore broader application of human computation in remote sensing. While search targets varied, the computation challenge was consistent. The methodology of large scale human collaboration for earth satellite imagery analytics was quickly applied in the aftermath of a 6.3 magnitude earthquake that devastated the city of Christchurch, New Zealand in February 2011.



Fig. 6 Rectangular burial mound (identified through our human computation network) from early to late Bronze Age origins (Allard and Erdenebaatar 2005; Jacobson-Tepfer et al. 2010)

Once again, a website was developed to solicit the public's help in analyzing large amounts of high-resolution imagery: in this case 10cm aerial imagery (Barrington et al. 2012a). Users were asked to compare imagery taken before and after the quake and to delineate building footprints of collapsed or very heavily damaged buildings. The interface was designed to be simple and intuitive to use, building on widespread public familiarity with web-mapping platforms (Google Maps, Google Earth, Bing Maps, etc.), so that more of the user's time is spent analyzing data versus learning how to use the interface. Using a simple interface that runs in a web browser, rather than an 'experts-only' geographic information system (GIS) platform, opens the initiative to a larger group of untrained analysts drawn from the general Internet public (Fig. 7)

After just a few days, thousands of polygons outline areas of damage were contributed by hundreds of users. The results are visualized in Fig. 8 below where areas of crowd consensus can be clearly identified by densely overlapping polygons. The crowd's results were validated by comparison to ground-truth field surveys conducted in the days immediately following the earthquake. The field surveys marked buildings with red (condemned), yellow (dangerous) or green (intact) tags, indicating the level of damage. Ninety-four percentage of the buildings tagged by the crowd were actually reported as damaged (red or yellow) by the field survey (Foulser-Piggott et al. 2012).



Fig. 7 Tomnod Disaster Mapper Interface in the Christchurch GEOCAN effort

Case Study: Peru Mountain Search & Rescue

The previous case studies demonstrated the capability of large networks of distributed human analysts to identify undefined features and apply visual analytics to remote sensing datasets on a massive scale. The final application of crowdsourced remote sensing we discuss highlights the timeliness that can be achieved when hundreds of humans help search through imagery and rapidly identify features of interest. On July 25, 2012, two climbers were reported to be lost in the Peruvian Andes. Missing in a remote, inaccessible region, the fastest way for their friends in the US to help find them was to search through satellite images. DigitalGlobe's WorldView-2 satellite captured a 50cm resolution image and, once again, Tomnod launched a crowdsourcing website to facilitate large scale human collaboration. Friends, family and fellow climbers scoured the mountain that the climbers were believed to have been ascending. The crowd tagged features that looked like campsites, people, or footprints and, within hours, every pixel of the entire mountainside had been viewed by multiple people (Fig. 9).

One of the first features identified within just 15 min of launching the website showed the 3-man rescue team making their way up the glacier in search of the climbers. Over the next 8h, consensus locations were validated by experienced



Fig. 8 Results of the crowd-contributed damage outlines and highlights of two destroyed buildings. *Red* = completely destroyed, *orange* = heavy damage, *yellow* = light damage




Fig. 9 Comprehensive crowdsourcing maps an entire mountain in just a few hours. The crowd identified possible footsteps (*orange*), people (*green*), campsites (*blue*) and avalanche regions (*red*)




Fig. 10 Fresh foot tracks in the snow outlined through crowdsourcing analytics of near real time ultrahigh resolution satellite imagery

mountaineers and priority locations were sent to the rescue team on the ground (e.g., footprints in the snow, Fig. 10).

The search ended the next morning when the climbers bodies were discovered where they had fallen, immediately below the footprints identified by the crowd. While this case study has a tragic ending, the story highlights the power of human collaboration networks to search a huge area for subtle clues and, in just a few hours, go from image acquisition to insight. Furthermore, we observe that in times of need, humans want to help, and when channeled in appropriate collaborative pipelines can do so through computation.

Next Step: Collaborating with the Machine

While we have shown three examples of scalable human analytics, it would be a challenge for human computation alone to analyze every image on the web, every galaxy in the sky or every cell in the human body. However, human computation systems can produce well-labeled examples in sufficient volume to develop machine learning methods that can tackle such massive problems autonomously (Barrington et al. 2012b; Snow et al. 2008; Novotney and Callison-Burch 2010). By integrating machine intelligence systems with human computation, it is possible to both focus the human effort on areas of the problem that can not yet be understood by machines and also optimize the machine's learning by actively querying humans for labels of examples that currently confound the machine.

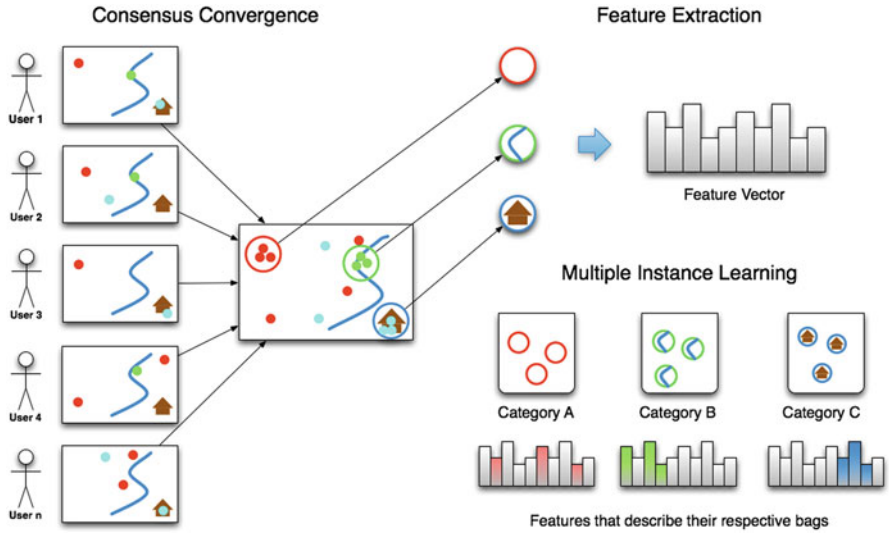


Fig. 11 Three phase approach to combine machine learning with search and discovery human computation: consensus convergence; feature extraction; and multiple instance learning

The detection of anomalies within an image is a difficult problem: we know that they may be located in regions of the image, but we don't know exactly where. We believe the application of multiple instance learning (Babenko et al. 2006; Maron and Lozano-Pérez 1998; Maron and Ratan 1998; Zhang et al. 2002) would be best suited for the problem at hand. Unlike the classical approach to learning, which is based on strict sets of positive and negative examples, multiple instance learning uses the concept of positive and negative bags to address the nature of fuzzy data. Each bag may contain many instances, but while a negative bag is comprised of only negative instances, a positive bag is comprised of many instances which are undetermined. While there may be negative examples in the positive bag due to noisy human input, the majority the positive examples will tend to lie in the same feature space, with negative examples spread all over. Multiple instance learning is able to rely on this insight to extrapolate a set of features that describes the positive bag. This is very appropriate for our data since a single image patch may contain many alternative feature vectors that describe it, and yet only some of those feature vectors may be responsible for the observed classification of the patch. A schematic of a proposed workflow for combining human computation and multiple instance learning (a machine based method) is outlined in Fig. 11.

If we are able to pool human perception to identify and categorize hard to define anomalies, we can begin applying this approach. From each of the many instances in a given category bag (i.e. ancient structure) we extract a set of image feature vectors. Since not every instance in the bag truly represents the labeled concept, some of these features will describe random image details, while others may be drawn

from an actual ancient structure and will, for example, exhibit a certain rectangular shape. As we iterate through all the instances in multiple bags, the aim is that the features that describe an anomaly will become statistically significant. As the signal from multiple positive instances emerges from the uniformly distributed background noise, we can identify the features that best describe certain classes of anomaly. Thus even with multiple, noisy, weakly-labeled instances from our training set, applying multiple-instance learning will result in a set of features that describe each anomaly and which we can apply to new data to find anomalies therein.

Conclusions

The idea of collecting distributed inputs to tap the consensus of the crowd for decision making is as old as the democratic function of voting, but in this digital age, networks of individuals can be formed to perform increasingly complicated computational tasks. Here, we have described how the combined contribution of parallel human micro-inputs can quickly and accurately map landscapes and features through collective satellite imagery analytics.

In “Expedition:Mongolia” we designed a system of peer based feedback to define archaeological anomalies that have not been previously characterized, to leverage a collective human perception to determine normal from abnormal. Participants without pre-determined remote sensing training were able to independently agree upon image features based on human intuition, an approach avails of the flexibility and sensitivity of human perception that remains beyond the capability of automated systems. This was critical in our search for the “undefined needle in a haystack”.

While this initial effort focused on an archaeological survey, applications of crowdsourced remote sensing exist across domains including search & rescue and disaster assessment. This was demonstrated through the efforts of Tomnod Inc., a group born out of the experiences in Mongolia to tackle the data challenges of the commercial satellite imaging industry through crowdsourced human computation. In the Christchurch disaster mapper effort we observe a remarkable 94% accuracy to ground truth. This result opens new possibilities for human computation and remote sensing in the assessment and ultimately recovery of disaster events. The Peruvian Mountain search & rescue operation demonstrated the remarkable speed with which insight could be gained from pooling human effort for large scale data analytics, suggesting that a combination of networked human minds and fast data pipelines could actually save lives.

Each example demonstrates the potential of online communities to mine unbounded volumes of digital data and catalyze discovery through consensus-based analytics. We have shown how human perception can play a powerful role when seeking unexpected answers in noisy unbounded data.

However, while our approach depends upon emergent trends of agreement as the validating principle of actionable information, we observe this inherently does not

capture the value of outliers (independent thinkers). Future work may identify mechanisms to reward “out of the box” thinking, possibly through a more detailed understanding and utilization of the individual human variables that contribute to a distributed human computation engine.

Finally, we observe that the natural next step in the evolution of human centered computation will be the collaboration between human and automated systems. This synergy will likely be required as we face the increasingly overwhelming data avalanches of the digital world.

Acknowledgements We thank N. Ricklin, S. Har-Noy of Tomnod Inc. as well as the entire Valley of the Khans (VOTK) project team; S. Bira, and T. Ishdorj of the International Association for Mongol Studies and F. Hiebert of the National Geographic Society for co-leadership in field expeditions; D. Vanoni, K. Ponto, D. Lomas, J. Lewis, V. deSa, F. Kuester, and S. Belongie for critical discussions and contributions; S. Poulton and A. Bucci of National Geographic Digital Media; the Digitaria team; Ron Eguchi and ImageCat Inc.; Digital Globe. This effort was made possible by the support of the National Geographic Society, the Waitt Institute for Discovery, the GeoEye Foundation, and the National Science Foundation EAGER ISS-1145291 and HCC IIS-1219138.

References

- Allard F, Erdenebaatar D (2005) Khirigsuurs, ritual and mobility in the bronze age of mongolia. *Antiquity* 79:547–563
- Babenko B, Dollár P, Belongie S (2006) Multiple instance learning with query bags, pp 1–9. vision.ucsd.edu
- Barrington L, Ghosh S, Greene M, Har-Noy S, Berger J, Gill S, Lin AYM, Huyck C (2012) Crowdsourcing earthquake damage assessment using remote sensing imagery. *Ann Geophys* 54(6):680–687
- Barrington L, Turnbull D, Lanckriet G (2012) Game-powered machine learning. *Proc Natl Acad Sci* 109(17):6411–6416
- Bewley RH Aerial survey for archaeology. *Photogramm Rec* 18:273–292 (2003)
- Blom RG, Chapman B, Podest E, Murowchick R (2000) Applications of remote sensing to archaeological studies of early Shang civilization in northern China. In: *Proceedings IEEE 2000 international geoscience and remote sensing symposium, IGARSS 2000, vol 6, pp 2483–2485 Honolulu, Hawaii*
- comScore. http://www.comscore.com/Press_Events/Press_Releases/2009/3/YouTube_Surpasses_100_Million_US_Viewers, Mar 2009
- Cooper S, Khatib F, Treuille A, Barbero J, Lee J, Beenen M, Leaver-Fay A, Baker D, Popović Z, Players F (2010) Predicting protein structures with a multiplayer online game. *Nature* 466(7307):756–760
- Deuel L (1969) *Flights into yesterday: the story of aerial archaeology*. St. Martin’s Press, New York
- Foulser-Piggott R, Spence R, Brown D (2012). 15th World Conference on Earthquake Engineering (Lisbon) following February 2012 The use of remote sensing for building damage assessment following 22 nd February 2011 Christchurch earthquake: the GEOCAN study and its validation
- Fowler MJF (1996) High-resolution satellite imagery in archaeological application: a Russian satellite photograph of Stonehenge region. *Archaeological Prospection*, 70; pp 667–671 *Antiquity* Portland Press
- Google Blog. <http://googleblog.blogspot.com/2008/07/we-knew-web-was-big.html>, July 2008

- Huynh A, Lin AYM (2012) Connecting the dots: triadic clustering of crowdsourced data to map dirt roads. In: Proceedings of 21st international conference on pattern recognition, Tsukuba
- Huynh A, Ponto K, Lin A Yu-Min, Kuester F (2013) Visual analytics of inherently noisy crowdsourced data on ultra high resolution displays. Aerospace conference proceedings IEEE. 1–8
- International game developers association (2006) Casual games white paper
- Internet world stats. <http://www.internetworldstats.com/stats.htm>, Aug 2009
- Jacobson-Tepfer E, Meacham JE, Tepfer G (2010) Archaeology and landscape in the Mongolian Altai: an Atlas. ESRI, Redlands
- Lasaponara R, Masini N (2006) Identification of archaeological buried remains based on the normalized difference vegetation index (NDVI) from Quickbird satellite data. *Geosci Remote Sens Lett IEEE* 3:325–328
- Law E, vonAhn L (2009) Input-agreement: a new mechanism for collecting data using human computation games. In: ACM CHI, Boston
- Lin AY, Novo A, Har-Noy S, Ricklin ND, Stamatou K (2011) Combining geoeye-1 satellite remote sensing, uav aerial imaging, and geophysical surveys in anomaly detection applied to archaeology. *IEEE J Sel Top Appl Earth Obs Remote Sens* 4(4):870–876
- Lin AYM, Huynh A, Lanckriet G, Barrington L (2013) Crowdsourcing the Unknown: The Search for Genghis Khan (in preparation)
- Lyons TR (1977) Remote sensing: a handbook for archeologists and cultural resource managers. Cultural Resources Management Division, National Park Service, US. Department of the Interior: For sale by the Supt. of Docs., U.S. Govt. Print. Off.
- Madden M (2009) Geoeye-1, the world's highest resolution commercial satellite. In: Conference on lasers and electro-optics/international quantum electronics conference, OSA technical digest (CD). San Jose, CA, USA
- Mandel M, Ellis D (2007) A web-based game for collecting music metadata. In: ISMIR, Vienna
- Maron O, Lozano-Pérez T (1998) A framework for multiple-instance learning. In: Advances in neural information processing systems. Madison, Wisconsin, USA, pp 579–576 (Citeseer)
- Maron O, Ratan AL (1998) Multiple-instance learning for natural scene classification. The fifteenth international conference on machine learning, San Francisco, pp 341–349, April 1998
- Novotney S, Callison-Burch C (2010) Cheap, fast and good enough: automatic speech recognition with non-expert transcription. In: Human language technologies: 11th conference of the North American chapter of the association for computational linguistics (NAACL HLT), Los Angeles
- Parcak SH (2009) Satellite remote sensing for archaeology. Routledge, London/New York
- Riley DN (1987) Aerial photography in archaeology. University of Pennsylvania Press, Philadelphia, PA
- Snow R, O'Connor B, Jurafsky D, Ng AY (2008) Cheap and fast—but is it good? Evaluating non-expert annotations for natural language tasks. In: 13th conference on empirical methods in natural language processing (EMNLP), Waikiki
- von Ahn L (2006) Games with a purpose. *IEEE Comput Mag* 39(6):92–94
- von Ahn L, Dabbish L (2004) Labeling images with a computer game. In: ACM CHI, Vienna
- von Ahn L, Kedia M, Blum M (2006) Verbosity: a game for collecting common-sense facts. In: ACM CHI, Montral, Montréal, Canada
- von Ahn L, Maurer B, McMillen C, Abraham D, Blum M (2008) reCAPTCHA: human-based character recognition via web security measures. *Science* 321(5895):1465–1468
- Wilkinson KN, Beck AR, Philip G (2006) Satellite imagery as a resource in the prospection for archaeological sites in central Syria. *Geoarchaeology* 21:735–750
- Zhang Q, Goldman SA, Yu W, Fritts JE (2002) Content-based image retrieval using multiple-instance learning. In: Machine learning-international workshop then conference- number 2, pp 682–689 (Citeseer). San Jose, CA, USA