Public Domain Treaty Compliance Verification in the Digital Age

Implementing "Public Technical Means": The Engineering Challenges in Exploiting Satellites, Smartphones, Ubiquitous Sensors, and Connectivity for Treaty Compliance Verification.

Christopher W. Stubbs and Sidney D. Drell

1. Introduction.

We explore in this paper some of the emerging opportunities, and associated challenges, that the digital age offers for the public-domain verification of compliance with international treaties. The increase in data volume, in ever-improving connectivity, and the relentless evolution towards ubiquitous sensors all provide a rapidly changing landscape for the technical verification of international treaties. From satellites to cell phones, advances in technology afford new opportunities for verifying compliance with international agreements, on topics ranging from arms control to environmental and public health issues. We will identify some of the engineering challenges that must be overcome in order to realize these new verification opportunities.

We find it helpful to distinguish between three different approaches to treaty verification:

1.) "National Technical Means, NTM", i.e. using spy satellites and various elements of information collection carried out by nation states for verifying compliance with formal treaty agreements. This includes both overt and covert methods. The resulting data are typically held as classified information, and are made available only to professional intelligence analysts in the respective nations and (sometimes) selected allies.

2.) "Shared Technical Means, STM", by which we mean instruments and their associated data sets that are shared among participating nation states and with international compliance organizations such as the International Atomic Energy Agency (IAEA). Examples include the International Monitoring System (IMS) being implemented for the verification of the Comprehensive Test Ban Treaty, and the Open Skies Treaty. Data from these systems are shared among participating nations but access is typically controlled by governments and restricted to intelligence professionals, even if the information is unclassified.

3.) "Public Technical Means, PTM", which we define as methods that involve data, interactions, and analysis in the public domain. This includes information that is either generated *by* or is made openly accessible *to* the general public, the scientific community, the private sector, and NGOs. Examples include i) images that are produced by commercial or scientific (as opposed to NTM) satellites and made available to the public for analysis, ii) exploiting sensors that are attached to the global digital communication network, be it through mobile devices such as a smartphone or laptop, or through a desktop computer, or as a stand-alone sensor with a separate special-purpose satellite or network connection, and iii) using scientific systems such as seismic networks and remote sensing satellites.

With the rapid growth of the social media, we take an expansive definition of PTM to include "societal verification", which is evolving from its initial definition of individuals taking responsibility to report treaty violations to include aggregate activities across social networks, and the resulting "crowd" interactions. The recent political events in

North Africa should dispel any doubt about the power of social media and modern connectivity in the political arena.

These different treaty verification modes are of course inter-related, and interact with each other. The *joint* knowledge obtained from all of them, taken together, is what we can and should exploit to assess treaty compliance and adherence. This is true of nuclear arms control as well as other international agreements, which may at some stage include international carbon emission limits and other environmental agreements.

The notion of a public role in the verification process has a long heritage [1]. The principle that individuals bear a moral responsibility to speak out if they become aware of a treaty violation has been extensively discussed previously [2], and labeled as "societal verification". The focus of this paper, however, is on the data collection and analysis aspects of PTM, and not on social networking. The prospect of a technologically-empowered public participating in future verification activities has been highlighted in recent presentations [3] by high-ranking US State Department officials.

2. The PTM Concept.

The distinguishing feature of Public Technical Means is the generation and analysis of open-access data, by the public, for treaty verification. The goal is not to replace NTM or STM verification efforts, but rather to actively engage a self-selected sector of the public in verification. Our definition of PTM does not include "open source" analysis by government intelligence professionals. That process may ingest public domain information, but the output and analysis products are not typically shared openly.

2.1 Why Bother?

Given the extensive investment made in NTM and STM verification systems, why bother to empower the interested public to participate directly in treaty verification? Isn't it better and more cost-effective to leave this task to the respective governments and intelligence professionals, acting on behalf of their citizens?

The first argument against this point of view is based on the principles of free choice, and individual and collective liberty. The PTM approach allows those members of the public who have a particular interest in treaty verification to allocate their resources (both funds and time) towards PTM objectives. Just as a subset of the general population elects to purchase fuel-efficient vehicles, despite a price premium, we similarly want to empower those with a passion for compliance verification to bypass their government's resource allocation process. If people want to place a CO_2 sensor on the roof of their home, patched into their home's wireless network, they should be able to do so in a way that coordinates with the worldwide efforts of like-minded individuals. This same freedom of choice logic applies to aggregations of individuals, through NGOs or special interest groups.

A second argument in favor of PTM is competition. By providing the public with relevant data and information, they are in a position to confront and challenge either their own government or external ones. The PTM effort will pressure both STM and NTM systems towards higher efficiency, since the managers of those government programs will be aware of the PTM efforts. It also promotes intellectual integrity by providing an independent check on the NTM and STM efforts.

The third argument in favor of actively facilitating PTM is that it is inevitable. The International Monitoring System (IMS) that is being established in support of the Comprehensive Test Ban Treaty is extensive and sophisticated, having deployed and activated [4] more than 80% of its anticipated total number of 337 instrumental installations. But tens of millions of people now carry around a sophisticated computer with built-in sensors, a powerful CPU, data storage, GPS geolocation and wireless connectivity, that also happens to be a telephone. The growth of smartphone ownership is projected to continue to increase worldwide well into the future. The numbers of desktop, laptop and home computers are increasing, as the digital quality of life improves across the globe. “Ubiquitous sensing” is now an established technical phrase, with associated conferences, journals, and a critical mass of engineers and academics that consider this to be an established subfield. This observation is not in any way meant to diminish the importance of the IMS network, but rather to point out the tidal wave of latent PTM capability that is sweeping the globe.

In addition to the growth of ground-based sensors, the non-NTM remote sensing capabilities in orbit are also increasing. With the growing interest in global climate issues, we should anticipate a steady increase in the number and instrumental capability of Earth-observing satellites. These orbiting resources will produce a stream of images and data that can be brought to bear on various treaty verification issues, in the PTM context.

A fourth argument in favor of PTM is its unique ability to contribute to activity monitoring. As one looks ahead to future progress in reducing the world's nuclear arsenals, and perhaps realizing the vision of President Reagan and General Secretary Gorbachev at Reykjavik in 1986 of a world free of such weapons, the arms control agenda will inevitably expand. With the entry into force of New Start in 2011, the United States and Russia have agreed for the first time to actually counting the numbers of deployed nuclear warheads on strategic long-range delivery vehicles, in addition to their much larger and more easily identified launchers. Verification of compliance with this provision is beyond the capabilities of NTM alone and requires onsite challenge inspections. Further reductions will require additional means of access including, for example, to verify the existence and numbers of non-deployed warheads; also of the storage of special components being maintained as a potential threat to reconstitute nuclear weapons and upset a strategic balance in a relatively short time. In such situations the importance of activity monitoring will be greatly enhanced. This is because it will require active maintenance or refurbishment of a number of limited life components to maintain an effective rapid reconstitution capability for any reasonable length of time [5]. Public means of verification or societal monitoring can be a significant tool for

monitoring such activities. They can call attention to suspicious activities, providing useful targeting information to other remote NTM resources that can probe what may be going on in more detail as a basis for bringing a challenge to the appropriate treaty verification authorities. The role of societal monitoring in such cases has a future potential of high significance that needs further study.

A final argument in favor of PTM is technical agility. The PTM technical capabilities can evolve and adapt rapidly, without being constrained by the international agreement needed to implement treaty verification revisions or upgrades. Once signed, treaties tend to last a long time. But the associated verification protocols seldom keep up with technical developments [6]. By operating outside the constraints of formal international agreements, PTM systems can capitalize on technical improvements and opportunities. Furthermore, the PTM approach can implement monitoring capabilities even inside countries that have not signed on to certain international agreements.

3: PTM in the Future: Sensors Everywhere, and an Interconnected World

The amount and public accessibility of data already extends beyond that of STM which operate under official guidelines and treaties between governments that define formal limits on data acquisition and protocols for sharing with participating nations. PTM exists in the numerous academic and national seismic research institutions that supplement the IMS for detecting and identifying low-yield underground nuclear explosions covertly performed in violation of the Comprehensive Test Ban Treaty. Another example is the Institute for Science and International Security based in Washington, D.C. which analyzes unclassified data from commercial photo reconnaissance satellites to monitor activities such as nuclear fuel enrichment or missile test preparations that portend developments of new concern to arms control and nonproliferation efforts.

With the rapid advances in information technology that can be, and are, being introduced into ubiquitous mobile sensors, many more pathways for data are rapidly opening up for a less structured form of PTM – i.e. societal monitoring. In particular the rapid increase in the market for smartphones can be very useful for PTM without it being adopted universally. There is no need to require that a single smartphone provide the sensitivity and discrimination needed for treaty verification. It is the *PTM sensor network's* capability and capacity that are relevant. Redundancy and a broad spatial distribution of sensors provide both robustness and discrimination advantages. Moreover, the sensor technology insertion time frame is surprisingly short. The typical turnover/replacement time for smartphone users is 18-24 months, and so new sensors could be implemented rapidly.

The overall net effect of the rapidly evolving potential of public and societal monitoring is to make it much more difficult for governments to implement policies that violate established agreements and get away with such deception for very long. Many more people will know what is going on, around themselves and around the globe.

4. Some Challenges of Implementing PTM Treaty Verification

Here we identify and discuss some of the challenges associated with fully implementing the PTM verification concept. Addressing any of these issues in full is beyond the scope of this paper, and we hasten to point out that some of these pose significant hurdles to PTM implementation. We highlight these points as potential technical research and development areas for the technical community.

4.1 The data integrity problem: detecting spoofing and deception

An issue that arises immediately for any PTM-generated sensor data is the data integrity problem. How can the PTM system detect instances of spoofing and deception, where a nation (or an individual) modifies the data or the accompanying header information or metadata (such as sensor type, calibration information, etc)? This could arise in various ways. A nation could attempt to conceal or obscure activity through PTM information manipulation. Or a nation could attempt to modify the PTM data stream emanating from a rival country, to implant deceptive information that implies a treaty violation, or a malicious hacker might find the PTM sensor data stream to be an irresistible target.

A substantial body of cybersecurity analysis has been done [7,8] on the problem of treaty sensor data authentication, from the perspective of public-private key exchange, as well as steganography (embedding a digital watermark in the noisy portions of the data itself). PTM sensors differ from more traditional verification measurement systems in that neither the physical data collection hardware nor the acquisition software are under close supervision.

The severity of the data integrity problem depends to a large extent on the fraction of the PTM sensors that are deceitful or compromised. It would certainly be difficult to contend with the situation where a nation floods a country (either theirs or someone else's) with pseudo-PTM sensors that were fiendishly clever in generating fake data. This would essentially amount to PTM denial-of-service attack, however, and would be readily detectable. On the other hand, if the legitimate PTM sensors comprise the dominant fraction of the sensing network, we can envision identifying the deceitful ones since they would be anomalous outliers.

We must also guard against malicious data modification by individuals or groups, as opposed to nations. There are numerous examples in the press of instances where commercial or government data systems are infiltrated by hackers. In this regard we don't see the PTM data as being particularly different from other information, and employing best practices for cybersecurity is as important for PTM as it is in other domains.

There is one straightforward validation method that is intrinsic to the PTM sensor approach. Any member of the public with local access to their own sensor data can always check to make sure the data stored on a central repository are in fact identical to the local copy, which is under their control. This cross-validation could be carried out in a secure, password-protected automated certificate exchange process.

4.2 The signal discrimination problem.

Sifting through PTM data collected by a variable and dynamic group of sensors to extract signatures of interest is “challenging”, to put it mildly. But this type of analysis problem is of increasing interest in both the scientific and engineering communities.

In one regime, the sensor network is managed and calibrated under the control of an external entity, and the data are considered reliable. This is the case for the academic seismic network that complements and supplements the seismic monitoring stations being installed to verify compliance with the Comprehensive Test Ban Treaty (CTBT). But we are interested in considering cases where the sensor network is heterogeneous and not centrally managed.

The Quake Catcher Network [9] is an example of exploiting existing sensors. The QCN team collects accelerometer data from laptops and mobile devices, and provides an interesting opportunity to field-test different data collection approaches and robust analysis algorithms. If the technical PTM community tracks ongoing data mining efforts in the scientific and commercial domains, then (we suspect) PTM efforts can likely implement and adopt proven techniques, rather than having to invent new tools.

4.3 The standards problem.

Incorporating data from a diverse and evolving suite of sensors will be greatly facilitated if uniform formatting standards were applied to both data and metadata. At present, different iPhone apps generate data files in a wide diversity of formats. One approach for generating uniform data sets would be to run PTM software on each device that produces PTM data, but this in turn requires a global coordination of the PTM effort. Another approach is to have the servers that harvest and distribute PTM information take on the task of bringing all data into a common format. This is again an issue with a scope that extends far beyond the PTM context, and that must be confronted and eventually solved by the “ubiquitous sensors” engineering community.

4.4 Privacy issues: anonymity and retribution.

If PTM is to succeed, the participants must be protected from any retribution, retaliation or pressure as a consequence of their involvement. The exposure of any individual will depend on the extent to which a PTM sensor is directly associated with that individual, and on the political and social situation within their country. For example one could imagine an “adopt-a-sensor” approach, where interested individuals contribute to a fund that acquires and installs networked sensors that have no direct association with any particular individual. In that scheme, supporting PTM would be akin to directly contributing cash to a political campaign.

The more subtle issue is how to deal with sensors that are attached to or are an integral part of a mobile computing device, for which location reporting will be an essential

ingredient for properly interpreting the PTM data. This becomes an issue of individual choice, and calls for an evaluation of personal exposure vs. participatory activism. The tradeoff between the benefits of providing location information vs. privacy concerns is an ongoing topic of public discourse, and PTM is but one facet of this broader issue.

The issue of concern is the combination of location and/or other identifier information with the identity of the individual who is supplying or facilitating the generation of PTM data. The basic problem is that given the principle of open data access, the PTM sensor location information could be combined with other information (such as cell phone records) to determine which individual provided the PTM data, even if the metadata contain no explicit identifiers. The ubiquitous sensor concept is already off and running, and the PTM aspect will be a minor consideration in this broader issue. But PTM can certainly benefit from the technical solutions to this problem.

If the location information in the PTM database can be "jittered" so as to obscure the sender's precise location, while still retaining the PTM information of interest, we can imagine the PTM contributor hiding in plain sight, in the clutter of other cell phone users. The QCN does attempt to preserve privacy of its participants by fuzzing out the exact location of their sensors, for example. This requires an adequate density of non-participants. Of course in the unlikely event PTM participation hits 100% this ceases to become an issue.

Another potential approach to ameliorating this location and attribution problem would be to determine whether the raw sensed PTM data need to be permanently associated with a sensor's location, or whether some extracted and interpolated information (say the CO2 concentration at ground level) might be the publically available PTM product. We might term this "PTM anonymization by analysis".

One potentially interesting technical solution to the sensor location and attribution problem through intercepted traffic would be to provide a PTM data pathway that circumvents the local internet. A satellite communication system would accomplish this. For low-bandwidth sensors, a relatively simple relay satellite (that stores received data and then retransmits to a set of downlink receive stations) would suffice. Establishing a PTM capability within a not-entirely-trustworthy country would then simply require local citizens to obtain and install sensors that transmit their data upon being prompted by the PTM relay satellite. If these PTM sensors were solar or battery powered, they would be locally "off the grid" entirely.

4.5 Incentivizing enduring public engagement in PTM.

The Red Balloon Challenge that DARPA conducted in 2009 arguably showed [10] as much about human behavior and incentives as it did about technology. The winning teams devised strategies that successfully provided incentives for individual participation, and this aspect was arguably more important than their technical approach to the problem. But nuclear weapons aren't flagged with red balloons, and underground nuclear

tests are rare. So the PTM approach faces a substantial challenge in obtaining and sustaining public engagement in the verification arena.

On the data collection front, one approach might be to bundle together a variety of sensors that pertain to "public service" monitoring. This might include some combination of CO2 and other environmental sensors, seismic monitoring, and various biological and chemical sensors. A modular hardware standard would allow individuals to deploy whatever combination of devices best suited to their particular combination of interests and ability to invest.

On the data and image analysis side, it seems to us unavoidable that only those individuals and organizations that have a sustained interest in particular problem will invest the time and effort needed to provide a valuable PTM contribution. But the Galaxy Zoo project (discussed below) has provided an example of how a structured, supervised and reward-based image analysis framework can succeed.

5. Some Relevant Examples and Models for PTM

A number of existing systems provide relevant examples and lessons for PTM. These include:

- Public analysis and inspection of satellite images, obtained from both Earth-observing and astronomical imaging systems,
- The quake-catcher network, that exploits accelerometer data from laptops and mobile devices.

Although there are some existing interesting examples of non-governmental use of downward-looking satellite images, we point to the Galaxy Zoo project as an example of widespread supervised public engagement in image analysis and interpretation, essentially crowdsourced data analysis.

We submit that the Galaxy Zoo project is a very promising example of public participation in image analysis, with clear applicability to the verification challenges of the decades ahead. The Galaxy Zoo project [11] had the goal of using minimally trained citizen scientists to classify (with visual inspection) the characteristics of galaxies using digital astronomical images obtained with the Sloan Digital Sky Survey. The Survey obtained exquisite resolution images of the entire Northern sky, and amassed a total of over 27 TeraBytes of imaging data and object catalogs.

The response was enormous. The Galaxy Zoo project attracted 250,000 online participants, from 170 countries. This public wave of amateur astronomers succeeded in classifying over 100 million galaxies, by visual inspection of every single image. The public participation in this project is comparable to the 300,000 individuals who have helped construct Wikipedia. Numerous scientific results have been drawn from the galaxy classifications performed by this entirely volunteer community of world citizens. This demonstrates the tremendous leverage that can be attained with a guided and informed interaction of interested, connected citizens with public domain data. The

Galaxy Zoo project has been so successful that understanding its effectiveness has become a research endeavor in its own right [12].

One key to the project's success was the shrewd use of computers for those aspects of image processing where computers do well, and then presenting people with a resulting subset of the pixels in a way that allowed them to efficiently and effectively address the question of interest.

We have therefore attempted to distill some lessons from the Galaxy Zoo project, in the context of future large-scale PTM image analysis:

1. Incentives: The project must engage the public with real data, on a topic of interest, and provide effective incentives for ongoing participation.
2. Clarity: The project must provide an elegant yet effective user interface, with clearly defined tasking.
3. Education and training: Online tutorials and intuitive data access tools [13] are essential ingredients.
4. Assessment: For data analysis and interpretation tasks executed by the public, participants should pass simple proficiency tests before embarking on the project tasks. The results can then be used to assign appropriate statistical weights to the judgments rendered by each (calibrated) participant.
5. Redundancy: Multiple individuals should carry out independent analyses, across a wide geographical range. The statistical properties of these results can then be used to assess and validate the system. This also avoids any simple gaming and denial and deception attempts.
6. Enjoyable: The projects should be fun, with good visualization and interaction tools.
7. Collaborative: The project team should provide multiple mechanisms for structured and loosely supervised interaction between and among participants, thereby allowing the group's experts to mentor less experienced individuals.
8. Feedback: Accolades and appropriate publicity [14] should be showered upon individuals who make important contributions.
9. Cohesive and competent team: The Galaxy Zoo was a close collaboration between computer-adept academics, IT and web experts, and individuals with subject domain expertise.

One possible approach to greatly expanding PTM image analysis would be to implement a scheme that identifies numerous regions of interest on the surface of the Earth, and then assigns a set of participants a well-defined task for each region. One group might be assigned the task of using successive satellite images to characterize and track any new construction projects in a country of interest or concern, by first using computerized change detection between images taken at different times to identify (but not to automatically characterize) construction sites. Another group might be assigned the task of looking at multiband images of a tropical forest, monitoring deforestation and land use.

What's currently lacking is a large scale public engagement (ideally with international public participation in the hundreds of thousands) through a supervised framework that identifies verification problems of interest, and connects interested members of the public with the appropriate subset of archived and incoming image data, under dedicated expert oversight of the overall process, with a robust method that detects and suppresses attempts to "game the system".

We therefore consider the facilitation of public access to relevant images as an area where a coordinated effort in the academic and/or NGO sectors, with a modest investment in software and personnel, could yield substantial near term dividends. The IT technology and the data both exist already, and other projects have shown impressive results in projects that are very analogous to the PTM treaty verification.

6. Conclusions, Recommendations, and Suggested Next Steps.

The discussion in Section 4 illustrates some of the significant challenges PTM faces that will determine the extent to which, and the pace at which, it will make progress. We outline below some initial thoughts on next steps that might enable a substantial increase in Public Technical Means being applied to the verification of compliance with international agreements [15]. These suggestions are roughly organized by system layer, from sensor hardware at the innermost level to the public policy interface at the outermost level.

6.1 The PTM Sensor layer.

Existing sensors have a range of capabilities that are relevant to verification, but the development of cheap, reliable PTM verification devices (seismic, trace gas analysis, CO2 sensors…) that have the requisite accuracy and precision is an area for further development. This seems an area where academia could work in partnership with government, industry, NGOs and international agencies to identify both opportunities and needs and work towards the implementation of new capabilities.

6.2 PTM Software.

Many engineering schools use smartphone programming as a method to introduce their students to the basic principles of computer science. This seems a good opportunity to build some prototype smartphone verification apps and (at a more advanced level) to explore the areas of user interfaces and user feedback.

6.3 The PTM Communications Layer.

There are some data validation and verification challenges that are specific to the PTM concept. In particular the issues of retribution-free data exchange mechanisms and robust data integrity protocols need further work and experimentation by the computer science community. We also think the notion of a PTM data relay satellite merits consideration.

6.4 The PTM Data Distribution Layer.

A number of groups have undertaken the internal analysis of commercial satellite images for PTM purposes, and these efforts certainly warrant further development and refinement. There is also at least one example, hosted by the Federation of American Scientists (FAS) of an access portal where PTM images are made accessible to the public in a structured way. These efforts are all moving towards the objective of PTM image analysis, and a thoughtful assessment of these initial forays could identify best practices for additional implementations or for the evolution of the existing programs.

6.5 The PTM Analysis and Data Fusion Layers.

The distributed sensing community is wrestling with the difficult problem of extracting knowledge and understanding from the noisy and cluttered data that are generated by wide arrays of inexpensive sensors. We advocate establishing stronger linkages between the verification community and this rapidly evolving subdiscipline of engineering.

Carrying out some informative PTM exploratory projects in areas where the public is likely to be supportive and where appropriate technology has been widely adopted seems an obvious next step. One example might be to harvest all the smartphone sensor data from willing participants in one major urban area, and investigate the sensitivity, clutter and signal to noise properties of the distributed network as applied to, say, seismic sensing. This is a non-trivial endeavor but as discussed above the Quake Catcher Network provides an existing implementation of a very similar system.

A particularly important aspect of PTM analysis is to consider the mechanisms by which the PTM community can perform internal quality checks on data, on analysis products, and on interpretations. Will the PTM community self-organize and require some kind of peer review, for example? An adaptive PTM system could even modify the data collection rate and mode depending on the sensor's reported location, in the context of other information being reported in the PTM network.

6.6 People, Policy, and Governments: The Outermost PTM Interface layer.

There are a number of interesting unresolved public policy elements in the PTM concept, including issues of privacy and individual rights. The interactions between a technologically empowered public and their governments are undergoing rapid change, as evidenced by recent events in the Arab world. It is difficult to predict how the citizens of the world will access, assess and exploit both data and analyses that pertain to international agreements.

It is also interesting to contemplate the appropriate role that enlightened governments could or should play in facilitating pubic data access, both within and outside their national borders.

6.7 Closing Thoughts.

Given the rapid rate of innovation in information technology, and with ever-evolving capability for both the acquisition and the analysis of large data sets, we see the PTM technical prospects as promising. However we are also aware of the diplomatic and political obstacles between stating policy goals and actually implementing the technical framework needed for achieving them. Without political strength at home, courage among the global partners, and international organizations with the power and will to legitimize appropriate actions, Public Technical Means and societal monitoring will be of little if any value no matter how revealing the data they provide to supplement National and Shared Technical Means.

CWS is grateful to the Hoover Institution for the Annenberg Visiting Fellow appointment under which this work was carried out. We also thank Jesse Lawrence and Angela Chung of Stanford University for informative conversations about the Quake Catcher network, and Summer Tokash for editing help with the manuscript.

7. References and Further Reading.

[1] For a nice historical overview see *IPFM Global Fissile Material Report 2009: A Path to Nuclear Disarmament,* in Chapter 9 at http://www.fissilematerials.org/ipfm/site_down/gfmr09.pdf.
Mitchell's paper "*Identifying Undeclared Nuclear Sites: Contributions from Nontraditional Sources*" at http://www.ipfmlibrary.org/mit98.pdf also discusses laying tripwires through societal verification methods. Another perspective is given in the chapter "The Role of Non-Governmental Organizations in the Monitoring and Verification of International Arms Control and Disarmament Agreements," by Crowley and Persbo, in J. Borrie & V. Martin Randin (eds.), *Thinking Outside the Box in Multilateral Disarmament and Arms Control Negotiations*, Geneva: United Nations Institute for Disarmament Research (UNIDIR), 2006. Available at http://www.unidir.ch/pdf/articles/pdf-art2588.pdf.

[2] D. Deiseroth, "Societal Verification: The Wave of the Future," *Verification Yearbook 2000*, T. Findaly (eds.). London: Vertic, 2000, p. 265.

[3] R. Gottemoeller, Acting Secretary of State for Arms Control and International Security, *Arms Control in the Information Age*, remarks at Moscow State Institute of International Relations, 2012. http://www.state.gov/t/us/187159.htm

[4] http://www.ctbto.org/map/

[5] S. D. Drell and R. Jeanloz, "Nuclear Deterrence in a World Without Nuclear Weapons" in *Deterrence: Its past and Future,* G. P. Shultz, S. D. Drell, and J. E. Goodby (eds.). Stanford: Hoover Institution Press, 2011

[6] S. D. Drell and C. Stubbs, "Realizing the Full Potential of the Open Skies Treaty," *Arms Control Today*, vol. 41, July/August 2011.

[7] R.L. Craft and T.J. Draelos, *Authentication of Data for Monitoring a Comprehensive Test Ban Treaty*, Sandia document SAND96-1601 (1996), accessible at http://www.osti.gov/bridge/servlets/purl/249279-lzOBut/webviewable/249279.pdf .

[8] G.J. Simmons, *How to Insure That Data Acquired to Verify Treaty Compliance Are Trustworthy,* Proceedings of the IEEE, vol. 76, no. 5, May 1988.

[9] E. Cochran, et al., *IEEE Instrumentation and Measurement Magazine*, vol.12, no.8, 2009. See also http://qcn.stanford.edu/.

[10] J. C. Tang, et al., "Reflecting on the DARPA Red Balloon Challenge," *Communications of the ACM* , vol. 54, no.4, April 2011.

[11] See http://www.galaxyzoo.org.

[12] M. J. Raddick, et al., "Galaxy Zoo: Exploring the Motivations of Citizen Science Volunteers," http://arxiv.org/abs/0909.2925, *Astronomy Education Review*, vol. 9, no. 1, p. 010103, 2010.

[13] See http://zoo1.galaxyzoo.org/Tutorial.aspx.

[14] An example of what was very extensive press coverage was the *Economist* story "Stars in their eyes: An armchair astronomer discovers something very odd," from 26 June 2008, available at http://www.economist.com/node/11614176?source=hptextfeature&story_id=11614176.

[15] Members of the scientific and engineering communities with a potential interest in addressing technical verification issues might wish to consult the State Department web site http://www.state.gov/t/avc/vtt/c45207.htm for fellowship and funding opportunities.