Cite as In: Gordon, T. J., Florescu, E., & Glenn, J. C. (Eds.). (2017). *Identification of Potential Terrorists and Adversary Planning: Emerging Technologies and New Counter-terror Strategies* (Vol. 132). IOS Press. (NATO workshop series.)

New Technology Options and Threats To Detect and Combat Terrorism

Dr. Paul J. Werbos[1]

Abstract. New algorithms and hardware technology offer possibilities for the predetection of terrorism far beyond even the imagination and salesmanship of people hoping to apply forms of deep learning studied in the IEEE Computational Intelligence Society (CIS) decades ago. For example, new developments in Analog Quantum Computing (AQC) give us a concrete pathway to options like a forwards time camera or backwards time telegraph, a pathway which offers about a 50% probability of success for a well-focused effort over just a few years. However, many of the new technologies come with severe risks, and/or important opportunities in other sectors. This paper discusses the possibilities, risks and tradeoffs relevant to several different forms of terrorism.

Keywords. Predetection, terrorism, nuclear proliferation, cyberblitzkrieg, time-symmetric physics, GHz, deep learning, internet of things, backwards time, retrocausality

1. Strategic Tradeoffs by Type of Terrorist Threat

Because the subjects of terrorism and new technology evoke strong emotions, and because they really are important to the survival of our civilization and even our species, it is important for us to be very careful in prioritizing which threats demand more extreme measures, and in being equally serious about which new technologies pose the greatest risks of unintended consequences as serious as the threats they might be used to address.

1.1. Risks From Nuclear Terrorism and Proliferation

I am grateful to the National Science Foundation and to Homeland Security for giving me an opportunity to act as CEO for three years (FY2011-FY2013) for the interagency competition to support advanced research in all directions to address the threat of nuclear or radiological terrorism [1]. Since the end of the Cold War, most of the public has become less afraid of nuclear proliferation and terrorism, but this is more the effect

[1] The views herein are the personal views of a government retiree, and should not be construed as any kind of official view of anyone now in the US government

of habituation and psychology than of any real reduction in the threat. Herman Kahn observed long ago that the next use of nuclear or radiological weapons, wherever it may occur, is likely to set in motion a chain of further events, both of a political nature and of a system-of-systems nature, which do seriously threaten the continued existence of the human species.

There are four main ways to try to reduce the threat: (1) reduce the quantity of nuclear materials and technology produced all over the world, especially in regions of less rigorous control; (2) try to make changes in world culture which reduce the force of people who desire to become mass murderers; (3) use new technology, like advanced predetection and better sensors and border controls, to keep terrorists from reaching inside US borders; and (4) strengthen defenses against weapons carried by missiles. The fourth was beyond the scope of this workshop.

Although I spent much of my life developing and funding advanced technology for intelligent systems which could enhance predetection [2,3], it is my judgment now (to be discussed below) that the negative risks in this application outweigh the benefits [4]. More precisely, the risks of deploying more than the simple forms of deep learning and data mining already in widespread use would be life-threatening in themselves. Even the risks of excessive reliance on simpler systems like Watson or similar schemes for top-down control of the Internet of Things (IOT) are a quickly growing threat, urgently requiring a rethink of how we design IT systems for this space and for government activity in general. A safer and more promising approach to the threat of nuclear proliferation and terrorism is to redouble our efforts to reduce the world production of nuclear materials and the dissemination of nuclear technology, both military and civilian, through reducing costs and improvements in a wide range of large-capacity renewable energy sources worldwide, and to come to grips better with the schisms in world culture which lead to organized enthusiasm for mass murder. On the other hand, there are emerging new possibilities for unbreakable operating systems and communications, and quantum technology, which can be helpful at much lower levels of risk to address terrorism in general.

From the interagency research effort [1], I learned that there are many useful incremental steps which can be taken still, through the normal missions and functions of DHS, to better cope with nuclear terrorism, but that the US is likely to remain like a house with twenty doors, ten locked tight, and ten swinging wide open for all the world to see. Those normal missions are important, but no one should expect huge safe breakthroughs to change the nature of the game here. It is factor (1), sheer limits on the production of nuclear material and technology, which is the main reason why people living in major cities of the US have not already experienced a serious disaster here.

1.2. Risks From Cyberblitzkrieg on Electricity and Other Critical Infrastructure

"Cyberblitzkrieg" refers to a type of cyberattack quite different from the usual type of everyday attack well-known by now in all sectors of the economy. At the time when this workshop was held, there were good reasons why a person with common sense would not say much about the options for cyberblitzkrieg in public; however, several releases of information, widely reported in CNN and the Wall Street Journal, change the situation radically. These releases create a severe clear and present danger, and call for urgency in discussing what we would have to do to patch up the situation.

The three most important new stories were:

(1) more complete information about how the Stuxnet type of cyberattack can be used to destroy (not just shut down but destroy) large electric power generators; (2) demonstration that some adversary has obtained the general NSA tools for that kind of cyberattack [5] and offered them for use and sale to anyone; and (3) coincidentally, news of a major reorganization of NSA. Many stories about (1) and (3) can be found by a google news search on the terms "Stuxnet" and "NSA21."

How severe could the damage be, if the vulnerability is not patched in time? Roughly speaking, it is not crazy to imagine that about half the big generators in the US might be hit at the same time, in a coordinated fashion. This would make the potential damage comparable to that of a major Electromagnetic Pulse Event [6], which Congressman Trent Franks evaluated as the most frightening threat out of all the threats being evaluated by his far-reaching committee with access to classified information. The National Academy of Sciences conservatively estimated $1-2 trillion worth of damage from an EMP event resulting from a solar flare as strong as the famous Carrington event of the 1800's [7], but high-ranking speakers at Congressional events have argued that the "system of system" implications of losing half of our electricity for six months or more would actually be much worse, perhaps enough to send us back to the Stone Age and to create unmanageable warfare. Gingrich's foreword to a recent novel on this possibility [8] reflects widespread concerns. For the EMP case, the possible patches are well-known and not so expensive [6], but what about cyberblitzkrieg?

Today's large cybersecurity industry relies more and more on new tools like sandboxing and intrusion detection, making more and more use of computational intelligence methods for pattern recognition and the like. These have been good enough, barely, for the flood of the usual type of hacking attack, but would simply not do much to prevent the damage due to a well-coordinated cyberblitzkrieg using the new tools.

Fortunately, the technology does exist which could patch the situation. More precisely, in the early 1970's, a major operating system was developed, the Honeywell Multics, which was provably unbreakable by theorem [9]. By coincidence, the first implementation of the chain rule for ordered derivatives, the most important mathematical foundation of deep learning (sometimes called backpropagation [10] and sometimes called the second adjoint method [11]) was performed on that operating system. As a result, I had occasion to be present at the debugging of the PL/1 com, and to assist with some of the work at the Pentagon World-Wide Military Command and Control System (WWMCCS) using that operating system. The "orange book" of Multics security became a kind of Bible in certain circles, followed by a series of "rainbow books" [12].

For many years, this technology was discussed less and less in public, but through time both Apple and Microsoft quietly applied some aspects of the technology, involving well-defined ring brackets, to their more recent operating systems. However, even small failures to comply precisely with the requirements of the theorems creates back doors which can be used to penetrate and take over such operating systems. This led to a huge flood of viruses in Microsoft systems, requiring case-by-case "antibodies" to detect specific viruses, always days behind first use of any virus, not enough to stop a cyberblitzkrieg. In 2014, back doors were made public in the more purely Unix-like Apple system and even in the Unix systems used in embedded control chips, which already caused a window of enormous danger to critical infrastructure. Nevertheless, the US resisted systematic elimination of all such backdoors and

compliance gaps, because those same backdoors were crucial both to our own offensive cyberblitzkrieg capabilities (e.g. relevant to the threat of nuclear proliferation) and to our ability to penetrate and track networks of terrorists. The effectiveness of predetection of terrorism, even with use of computational intelligence, depends on access to the kind of data which such monitoring makes possible.

This backdoor strategy leaves the electric power industry and other critical infrastructure wide open to potential attack, as has been known for many years. (Today's power grid [13] already contains many points of vulnerability, and there is no need to elaborate here on where the worst vulnerabilities lie.) Because the potential damage from cyberblitzkrieg far exceeds the damage we see from ordinary types of terrorism, the NSA has long supported an Information Assurance activity to quietly reduce that risk. They have worked with Red Hat to create a continuously updated version of SE Linux, widely used in critical front line parts of the electric power industry. However, front line power engineers complain that they are heavily dependent on major vendors like AB&B and Siemens, and that there is always a lag between current software and full compliance. With all sources of SE Linux systems, there is concern about the continued existence of back doors. "Whose back doors do you choose? That is the only choice." As back doors become discovered more and more quickly by third parties, the gap in full compliance already created a serious large risk even before the recent leaks.

It is time to face up to the huge crisis and risk created by the recent leaks. To that end, I propose: (1) compliance with unbreakability standards and theorems should no longer be partial compliance, but full compliance, to be verified by open source compliance checking programs whose tests must be passed BEFORE new software is actually used in critical power control systems; (2) to account for the legitimate need to be able to monitor terrorists, the compliance rules should be broadened to allow inclusion of a "black subroutine" in the system whose powers over the system are provably limited but which allow a kind of "read only" capability to report back to agencies with legal compliant warrants and codes to activate such subroutines. This requires an urgent software development and demo project to accomplish this, and prove to the world that it has been accomplished, as soon as possible, by groups fully versed in rainbow book technology. This major change would eventually eliminate our cyberattack capabilities ala Stuxnet, as other nations also upgrade their critical infrastructure, but this would be worth the price if we can protect our own civilization at the same time. The cost of the change would be much less than the benefit.

Unfortunately, the news accounts about NSA21 suggest real concern that other political motives in play in the US might even reduce the limited Information Assurance activity already in place, let alone allow the dramatic change needed to patch the new urgent risk to our lives.

In further discussions, people have reminded me that the unbreakability theorems stop being valid if the physical hardware stops implementing the operating system code as written. Two major issues have emerged in that regard. First, hacks have been developed which overuse certain parts of a CPU, causing physical breakdown allowing entry to the system. It is felt that use of error-correcting codes in the hard disk and elsewhere can preserve integrity in the fact of such attacks, and even more so if new memory management procedures/rules are added to the operating systems. This could be addressed by preparing a second generation to the software compliance routines, to require even higher standards. Second, back doors can be inserted into the chips by the manufacturer, hard for a purchaser to detect; fortunately,

those would only provide access to a very small group of actors, much smaller than the circle able to perform ordinary cyberblitzkrieg [5]. To detect such back doors more reliably would tend to require "NP hard hardware devices," which now appear possible if a new level of quantum hardware technology is pursued far enough [14].

Vulnerability of communication flows, and new ways to break them and secure them, are also important, but beyond the scope of this paper.

1.3 Risks from Other Types of Terrorism

Many serious researchers believe that we also face threats in the chemical, biological and environmental realms which threaten the very existence of the human species, and that none of the hopes to develop "lifeboat" technology offer a clear alternative as yet. While I tend to agree with this viewpoint, in general, it is beyond the scope of this paper, and does materially change the tradeoffs discussed here.

Ordinary terrorism such as what we have seen in the news this year has been very successful in enraging the public, but the resulting damage remains much less than what we have seen in traffic accidents. It calls for continued efforts within present paradigms, but, as with nuclear terrorism, does not call for responses riskier than the threats themselves. Indeed, history suggests we should be very vigilant and very skeptical about threats to centralize power and reduce human freedom which try to create and exploit such public hysteria.

2. Breakthrough Technology for Prediction and Control – CI, IOT, BCI and Quist

This section will discuss the possibilities, opportunities and risks for true breakthrough technologies in the areas of computational intelligence, CI (which includes deep learning, neural networks, brain-like intelligent systems in general and allied technologies), the Internet of Things (IOT), Brain-Computer Interface (BCI) and Quantum Information Science and Technology (QuIST). Sometimes the trends in technology which seem on the surface to be riskiest and furthest off actually entail less risk , or less risk to global security to talk about; that is why, at this workshop, I asked to speak in more detail on one entertaining example from the realm of QuIST, the possibility of a forwards time camera or backwards time telegraph (BTT), for use in warning about a terrorist incident before the incident takes place.

2.1 Principles and The Way Forward for the Forward Time Camera and Backwards Time Telegraph

Here I do not want to assert that the forward-time camera or the backwards time telegraph will certainly work if you were to follow my recipe for how to do it. I estimate about a 50% risk for this specific recipe. The recipe itself also needs to be worked out in more detail. Nevertheless, this is far less speculative than discussions about singularities, about extended human lifespan, about X prize launch vehicles, and many other areas which some people have invested in. The recipe is there, and the cost of following it should be in the low millions for a few years, not the billions for decades which have been invested into some alternative lines of R&D.

The recipe is simply to carry through the following four steps, the first of which is explained in great technical detail in [14]:

(1) Using the same type of desktop machinery which created three entangled photons for the Greenberger, Horne and Zielinger (GHZ) experiment [15], replicate the stunning preliminary results achieved in 2015 on an extended experiment supporting the time-symmetric reformulation of quantum physics [14]. Because of the preliminary results so far and the strong underlying logic, I would estimate the probability of success at 80%. Note that success would also open the door to many other new technologies. Even failure would provide important new clarification about the physics and modeling requirements for advanced QuIST.

(2) Enhance the existing approach to quantum ghost imaging [16] by using that same GHz source, using two photons on the left to create the recorded image and detect when an entangled triplet is being recorded from, and the one on the right to reach into space to the object to be imaged. This is a mathematical task, mainly aimed at proving that coincidence detection can be done well enough just on the left-hand side without a need for some kind of detector out in space as required in conventional quantum ghost imaging [16]. This is where most of the risk lies in this recipe. But even in the case of failure at this stage, it seems likely that other approaches to developing BTT, exploiting the lessons from step (1), would eventually be able to work.

(3) Attach the new triphoton ghost imaging system to a powerful telescope imaging the sun, so that the third photon goes backwards into the eyepiece. If step (2) works out, this would yield an image of the sun eight minutes forwards in time, unlike the usual images which are eight minutes old by the time they reach the earth. Because the sun is highly dynamic, this should provide a very clear demonstration that we can enter a whole new era in QuIST, and it should also supply some advance warning on solar flares, of practical value in itself. If this works, the value in upgrading world culture should be immense, similar in a way to the recent discovery of exoplanets, but much larger.

(4) Then attach the triphoton ghost imaging system to long and slow optical fibers, curving the light around, to allow such capabilities as a forwards time camera or BTT for use on earth, more or less making real the kind of possibility envisioned in past in science fiction.

Following the strictest, safest procedures of science, one would normally not even talk about steps (2) through (4) in any detail, before step (1) is completed in a way which establishes more confidence and reduces legitimate uncertainty and controversy. However, for public policy, it is important to do what this workshop asks for: a look ahead to the "decision tree" of what might become possible in the future, and how to get there.

At the workshop, Karlheinz Steinmuller raised a practical question for predetection in general: "If we find a way to predetect a terrorist incident, can we intervene, legally, when we do not have evidence that it will certainly will happen?" Ah, but what if we had an actual photographs from a forwards time camera of a person performing a ghastly act? Could we not then get permission at least to prevent the ghastly outcome, quickly, and investigate further?

Ted Gordon asked about the classic science fiction paradoxes about time travel and even communication across time. In fact, in this example, what happens if we take a picture of a future ghastly event and then prevent that event from happening? Do both versions of the future actually exist in some mathematical sense, in the physics?

Unfortunately, to give a true answer to Ted's question, there is no substitute here for trying to understand the mathematics at an intuitive level, more intuitive than what people normally achieve when mechanically using (but not fully understanding) the similar but even crazier mathematics of the Feynmann path version of quantum field theory. In Feynmann path physics, and in the simpler similar models in [14], there is only one true state of the space-time continuum, which has just three dimensions of space and one of time. There is just one future. But our experience of consciousness [14,17] is like the shadows in Plato's cave – a matter of multiple "possible" scenarios for the state of space-time which interact with each other in a way similar to the interaction between universes in the better-understood "multiverse" formulation of quantum mechanics which underlies first generation QuIST [18,19]. In our example, the photograph is a photograph of a real scenario whose strength fades when we use it to prevent the incident, and it evaporates into nonexistence much as a virtual particle in a quantum loop in a Feynmann diagram evaporates with forward time.

Jerry Glenn asked a follow-on question: If we find ourselves at the future-time end of a truly horrendous event, like a nuclear missile hitting New York, and if we send a message back in time with a BTT to prevent that event, are we not committing a kind of suicide? This kind of dilemma would be a truly high-stress test of how much sanity we have achieved; however, in my view [20], a truly sane self-aware person would send that warning back in time, so that the core version of himself or herself could survive in a better world. At a fundamental level, we have evolved to survive in a cosmos governed by severe uncertainties and by quantum physics; "that is who we are."

Other possible applications of Analog Quantum Computing [14] would be many, some risky, some valuable, but difficult enough that this four-step project would not constitute public release of all of those follow-ons.

2.2. A Few Highlights from CI, IOT and BCI for predetection

Recent developments in CI, IOT and BCI require deeper, more reflective thinking than the forwards time camera does, to understand the full range of possibilities and tradeoffs. This brief chapter can at best give only a few highlights, to guide the reader towards more complete treatments [2,3,4,10,17, 20,21] and new work submitted elsewhere which extends [17]. See [4] especially for an up-to-date discussion of the big picture and tradeoffs, without explicit mention of equations but with full engagement with real-world details which tend to be grossly misunderstood by those who discuss the tradeoffs without being grounded in the mathematics.

In 2008, NSF formulated a new open, competitive research activity widely discussed within the Engineering Directorate aimed at capturing the highest possibilities but avoiding the worst risk in the general areas of neural networks and understanding the brain [3]. One of the four large grants made under this activity [21] included funding allowed Andrew Ng and Yan LeCun simply to test out basic designs for deep learning with neural networks on a host of open challenge competitions which had been well-known in artificial intelligence (AI). The basic designs here had been in

existence for decades, but for decades the traditional field of AI had been held back by decades of conventional wisdom and conservative thinking and narcissism which kept the neural network designs from being fully tested. Once LeCun performed these tests, and quickly displayed superior measurable performance in image recognition, speech recognition and natural language processing, follow-on funding quickly became available from DARPA and Google, and a massive new rediscovery of neural networks and deep learning resulted. At a plenary talk at the World Congress on Neural Networks held in Beijing in 2014, LeCun described this as the second rebirth of the neural network field, a reawakening comparable to the first reawakening in 1987/1988, quickly and dramatically breaking through decades of sleep (and pleasant dreams) in the general culture. In my own plenary talk at the same session, I agreed strongly with LeCun on all points, except that I said he was too humble in not really emphasizing how his own open-source computer work and dissemination of results was what created the reawakening.

Yet even before that, when the general culture was sleeping, the work at NSF leading up to [3] was moving ahead to more advanced designs. In leading [3], I developed the slide in figure 1 to illustrate the five most important areas of ongoing research, two of a basic and mathematical nature, and three addressing important application domains.

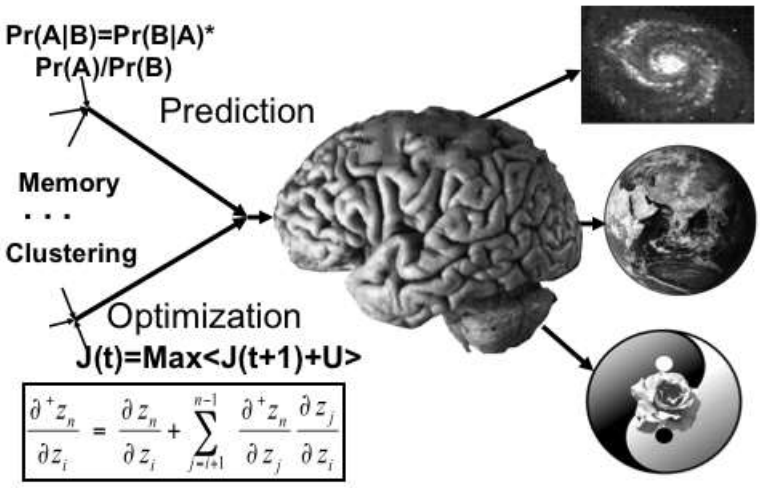

Figure 1. Iconic representation of five most important streams of research in [3] and in prior work leading to [3]

Every one of these five areas involves complex tradeoffs, and links to national security. For example, the icon on the upper right represents use of neural networks to improve our chances of success in the economically sustainable settlement of space by humans [22,23]. Prior to 2003, when Tony Tether's new emphasis on translational technology over basic discovery led to a huge general increase in proposals in the Engineering Directorate at NSF and forced some painful reallocations, we led a major effort to apply new optimization methods to the challenge of low-cost access to space (e.g. [24])

and to other aerospace applications such as defense against missiles [25]. Because of great success in the one-on-one hit-to-kill missile defense application, we began discussions with Boeing in 1991 about extending this to theater level missile defense, which terminated abruptly when one of the folks at Boeing alerted me to the reality and size of the "Terminator" risks we were beginning to develop [4].

Many salesmen seeking money for the field of AI have tried to assure us that the Terminator risk is not real, or that it can be overcome by painting happy faces on robot heads containing intelligent systems, or by inserting rules on old-style pre-intelligent expert systems. This reminds me of a great saying from Congressman Trent Franks (in his discussions of EMP): "Your folks are only worried because you do not have all the facts. If you had all the facts, you would be terrified out of your minds." After the initial, quiet Terminator discussions, some informed people asked: "Is this just another one of those advanced technologies which humans are simply not ready for yet? Should we wait until we have a better and deeper understanding of ourselves, and thus a better hold on sanity and self-control?" But in this case, the new mathematical understanding of mind and intelligence is also a prerequisite to a deeper understanding of ourselves and a more full expression of human potential [20,26], the area depicted by the icon on the lower right – the rose on the yin-yang, another huge area. Thus since 1991, we have tried to walk a tightrope, trying to push forward the positive benefits of the new mathematics with enough engineering demos to get people to pay attention, but avoiding the types of activities which increase the probability of human extinction at the hands of Terminator AI.

Notice that the activity depicted by the icon on the lower right is also important to understanding the mind of potential terrorists, a key aspect either of predicting terrorist acts or of understanding those aspects of world culture which currently breed mass murderers not only in parts of Islam but in other cultures around the world, just as dangerous.

More recently, new risks and negative trends have started to appear, just as serious as the Terminator risk, also requiring a very careful balancing act, involving the IOT and BCI [4].

The Internet of Things (IOT) is basically just the emerging new system of internet and sensors and actuators connected to the internet, worldwide, considered as a single system. It is beginning to develop features very similar to those of a brain – a massive network of elementary units, distributed in nature, connecting sensory input to action. However, the action includes virtually every controllable device in the world, from implants in people's bodies to drones to manufacturing plants and power generators. As people begin to ask how to manage the IOT as a single system, they are beginning to ask questions which we began to study in the electric power field decades ago [13].

One possible way to manage the IOT of the future more efficiently is to make it more like a massive neural network, exploiting the same general mathematics to get optimal performance. But that leads to Terminator risks. Another way, now far along in some business plans and activities quietly going ahead, is to use more "conservative" top-down control mechanisms or ad hoc mechanisms, which unfortunately may seriously undermine key concepts of freedom and diversity vital both to the happiness of the American people and to our collective intelligence in being able to cope with ever more difficult and complicated challenges [4]. (Diversity "in thought" is essential within any intelligent system; if every neuron always had the same output as every other neuron, the information processing capability of the system would be one bit.) As

we try to grope for a middle way, to move forwards and avoid the very serious risks both on the left and on the right, it would help for more people to better understand what has been learned in the electric power field [4], where secure open systems permit the development of greater freedom (and privacy protection) in large networks, empowering human beings. The development of this middle way will be essential to expanding our freedom and our capabilities and even our spiritual growth, all across this planet, in the coming decades.

References

Asterisk signifies papers also posted at www.werbos.com/Mind.htm or physics.htm.

[1] National Science Foundation, Domestic Nuclear Detection Office-National Science Foundation Academic Research Initiative, 2013. http://www.nsf.gov/pubs/2013/nsf13554/nsf13554.htm

[2] Paul J. Werbos, From ADP to the Brain: Foundations, Roadmap, Challenges and Research Priorities, in Prof. Int'l Conf. Neural Networks, 2014. IEEE. http://arxiv.org/abs/1404.0554

[3] National Science Foundation, *Emerging Frontiers in Research and Innovation 2008*, http://www.nsf.gov/pubs/2007/nsf07579/nsf07579.pdf, section on COPN

[4] Paul J. Werbos, Data Mining to Support Human Intelligence, keynote plenary talk to IEEE Computational Intelligence Society Winter School on Big Data in Computational Intelligence. Abstract posted at www.udc.edu/winter_school/IEEE_Cis_winter_school_files/Page853.htm.
Slides posted at http://www.werbos.com/Neural/ADP history/NN_BigData_2016_v2.pdf .
Video for first 13 slides posted at https://www.youtube.com/watch?v=6q1HqRd9MnA

[5] Jeff John Roberts, "A Second Snowden at the NSA: Here's What We Know," Fortune magazine, August 23, 2016

[6] http://empactamerica.org/

[7] Space Studies Board,*Severe Space Weather Events--Understanding Societal and Economic Impacts: A Workshop Report*. National Academies Press, 2008.

[8] Newt Gingrich, Foreword, in William R. Fortschen and William D. Sanders, *One Second After*, Doherty Associates, New York, 2009.

[9] http://multicians.org

[10] Paul J. Werbos, P.J., 1994. *The roots of backpropagation: from ordered derivatives to neural networks and political forecasting*. John Wiley & Sons.

[11]* Paul J. Werbos, P.J., 2006. Backwards differentiation in AD and neural nets: Past links and new opportunities. In *Automatic Differentiation: Applications, Theory, and Implementations* (pp. 15-34). Springer Berlin Heidelberg.

[12] David D. Clark and David R. Wilson, A Comparison of Commercial and Military Computer Security Policies, www.computer.org/csdl/proceedings/sp/1987/0771/00/07710184.pdf (mirror copies easily located on the web)

[13] Paul J. Werbos, Computational intelligence for the smart grid-history, challenges, and opportunities, *Computational Intelligence Magazine, IEEE* 6.3 (2011): 14-21.

[14]* Paul J. Werbos and Ludmilla Dolmatova, 2016. Analog quantum computing (AQC) and the need for time-symmetric physics. *Quantum Information Processing*, *15*(3), pp.1273-1287. Brief further discussion of the 2105 experiments is also posted at physics.htm.

[15] Dik Bouwmeester, Jian-Wei Pan, Matthew Daniell, Harald Weinfurter, and Anton Zeilinger. Observation of three-photon Greenberger-Horne-Zeilinger entanglement," *Physical Review Letters* 82, no. 7 (1999): 1345.

[16] Dmitry V. Strekalov, Baris I. Erkmen, and Nan Yu, Ghost imaging of space objects," *Journal of Physics: Conference Series*. Vol. 414. No. 1. IOP Publishing, 2013.

[17] Paul J. Werbos, How can we ever understand how the brain works? In Robert Kozma and Walter Freeman, eds, *Cognitive Phase Transition in the Cerebral Cortex: Enhancing the Neuron Doctrine by Modeling Neural Fields*. Springer International, 2016.

[18] David Deutsch, *The Fabric of Reality: The Science of Parallel Universes and Its Implications*, Penguin, 1997

[19] Cotler, Jordan, Lu-Ming Duan, Pan-Yu Hou, Frank Wilczek, Da Xu, Zhang-Qi Yin, and Chong Zu. *Experimental test of entangled histories*. *arXiv preprint arXiv:1601.02943* (2016).

[20]* Paul J. Werbos (2105), Links Between Consciousness and the Physics of Time, *International IFNA -ANS Journal "Problems of nonlinear analysis in engineering systems"*, http://www.kcn.ru/tat_en/science/ans/journals.
[21] Andrew Ng, Y. Dan, E. Boyden and Y, LeCun, EFRI-COPN Deep Learning in the Mammalian Visual Cortex, www.nsf.gov/awardsearch/showAward?AWD_ID=0835878&HistoricalAwards=false
[22] Paul J. Werbos, Towards a Rational Strategy for the Human Settlement of Space, *Futures*, Volume 41, Issue 8, October 2009
[23] Paul J. Werbos, Reviewing Space Solar Power policy, Ad Astra, Vol. 26, No.2, 2014, http://www.nss.org/adastra/volume26/ssppolicy.html
[24] Ramon Chase**,** GOALI: Neurocontrol for Design of MHD System to Enable Single Stage-To-Orbit (SSTO) Air-Breathing Propulsion System www.nsf.gov/awardsearch/showAward?AWD_ID=9814230&HistoricalAwards=false
[25] D. Han and S.N.Balakrishnan, Adaptive Critic based Neural networks for Agile Missile Control, 1998 AIAA Guidance Navigation and Control Conference and Exhibit, Boston, MA, Aug. 10-12, 1998, pp. 1803-1812.
[26] Paul J. Werbos, 2012, Neural networks and the experience and cultivation of mind. *Neural Networks* 32: 86-95.