

©IEEE. Cite as in In Proc. Int't Joint Conf. Neural Networks, 2018.  
Per IEEE copyright agreement, this was the final accepted version before formal  
final submission. The final official version will appear in IEEE xPlore.

# AI Intelligence for the Grid 16 Years Later: Progress, Challenges and Lessons for Other Sectors

Paul J. Werbos  
National Science Foundation (NSF), retired  
IntControl, LLC, partner  
werbos@ieee.org

**Abstract**--How could the “new AI” based on neural networks and deep learning be applied to the electric power grid, so as to get maximum benefit from the new technology, and serve as a model for how to organize the new Internet of Things (IOT) in general? The first of these questions was already assessed in great detail in workshops organized jointly by NSF and the Electric Power Research Institute (EPRI) in 2002 [1], drawing on new technologies which included today’s deep learning but also more advanced technologies in the same family [2]. The NSTC (White House) Smart Grid policy of June 2011 cited [1] in stating: “NSF is currently supporting research to develop a ‘4th generation intelligent grid’ that would use intelligent system-wide optimization to allow up to 80% of electricity to come from renewable sources and 80% of cars to be pluggable electric vehicles (PEV) without compromising reliability, and at minimum cost to the Nation.” This paper gives some highlights of the progress made, the open challenges, and important connections to the larger needs of humanity, in that order. The synergy between new intelligence, new technology for cybersecurity [3] and new physical hardware [4] is essential to maximum success, and even to the very survival of our endangered species. Lessons from the power grid are essential to better understanding of urgent challenges central to the IOT in general [5].

**Keywords** -- *intelligent grid, RLADP, internet of things, IOT, cybersecurity, renewables, orange book, ISO, terminator*

## 1 CORE ISSUES IN THE INTELLIGENT GRID

This section will give a relatively quick, informal review of progress with the truly intelligent grid since the more thorough 2011 review [1]. In fact, any accurate picture of policy issues in Washington DC requires a great deal of informality, because the process itself is not entirely consistent with the highest traditional standards of IEEE.

In 2009, I had a chance to serve as a Brookings Legislative Fellow in the office of Senator Specter, assigned to the Environment and Public Works (EPW) Committee, with responsibilities for climate, energy, defense technologies and space, in that order. This gave me a chance to discuss how a new kind of intelligent grid, making full use of reinforcement learning and

approximate dynamic programming (RLADP) [6] with neural networks, could substantially reduce the total cost of renewable energy to end users in the US [7].

Those who track the popular literature on “the new AI” should note that ALPHA-GO was built upon one of the simpler forms of neural network RLADP discussed in [6]. Older forms of optimization assumed that the world is linear, or without uncertainty, or even static; RLADP addresses the general case of how to make optimal decisions in an environment which may (or may not) be so simple, with anticipation or foresight built-in as part of the designs. After all, to beat the best human player of Go, the system had to anticipate future moves, explicitly or implicitly.

Roughly speaking [4], world electricity generation now costs about \$2 trillion per year if 10 cents per kwh is assumed; large-scale renewable energy can cost anything from 3 cents per kwh unsubsidized to 50 cents or more cost to the end user (due to add-on costs related to control). Better, more intelligent optimization could thus be worth trillions of dollars, and could decide whether we humans can make a profitable transition to renewables before the new worst case climate events literally kills us all [4,7].

Thanks in part to IEEE USA, new legislation called for use of more modern, powerful optimization technology as a core component of the intelligent grid – but efforts to implement the legislation in the executive branch have at times been almost comic. At one of the annual IEEE Smart Grid conferences, a leading advisor of certain programs said: “Of course we understand that this word ‘optimization’ is just a metaphor. No one does real optimization in the power grid. It’s just a fashion statement, like the ‘smart grid,’ which really just means we are smart people and we do our best.” Progress in upgrading the grid has been slowed by such gross misconceptions; see [1] for links to presentations at the Federal Energy Regulatory Commission (FERC), which explain in great detail how 70% of the electricity market in the US is managed by massive optimization programs located at a handful of Independent System Operators, such as PJM, which gives specific orders to every

generator from Chicago to New Jersey and Virginia whether to turn on, turn off or ramp up or down every five minutes. Even in the power engineering community, some policy leaders are aware of the “balancing authorities” (a different set of interstate systems which aim to prevent gross instability and collapse on a second-by-second basis, WITHIN the 5-to-15 minute window controlled by the ISOs), but not about what the ISOs do.

Up until my retirement from NSF in 2015, the most important promising work on the intelligent grid as such was roughly divided into three areas: (1) improvements in the transmission system, to convert renewables from a liability to an asset in stabilizing the grid and to allow planning which exploits opportunities to build new transmission and generation together to reduce costs; (2) improvements in the distribution system, motivated at times by rooftop solar and the microgrid movement, crucial to our ability to accommodate massive numbers of plug-in hybrid cars at the distribution level and to national security [8]; (3) use of household intelligent agents (like the Mannheim Project in Germany [1]) to shift demand from bad times to good times, essential to the economics of renewable energy in the US and EU. I tried to allocate funds roughly 40-40-20 to these three areas, but it was hard to find good proposals in demand response (item 3) because many US researchers were more interested in studying conservative US systems which had much less benefit than the best German experiments, which call for real intelligent systems (RLADP).

At the transmission level, work by Haibo He, Ron Harley and Ganesh Venayagamoorthy has led to exciting new results, though policy changes since 2014 have had the effect of moving progress faster in China and slower in the US [2]. Modern wind farms and solar panels rely on fast power electronics; RLADP with neural networks allow optimization of decisions made at fractions of a second, fully exploiting what power electronics and neural network chips [2] can do, unlike the less powerful optimization now used on the less advanced, less parallel processors at ISOs which require 5 minutes cycle time. Fast optimization of the power electronics, responding to fast signals from the grid, is what can turn these energy sources from an asset to a liability. Harley and Divan of Georgia Tech have estimated that this can reduce costs to end-users by a factor of five or more (again, worth many trillions of dollars in time) when coupled with Divan’s new technology for low-cost power switching, useful both at the transmission and distribution level. New control and new control authority naturally have huge synergies with each other. Joint optimization of transmission and generation investments in planning is also important, because the greatest cost reductions require both together [4].

It is really crucial to remember that power markets vary enormously by nation and by time of day. High levels of AI are not so important in Latin America as in the US and EU [4], because more reliable

renewables and energy storage are already available, and because the Brazilian power grid is uniquely huge already. However, new work in market design aimed at a new Argentina-Brazil-Chile deal would be essential, to stimulate the necessary expansion in two-way HVDC transmission systems.

In the EU, efficient harmonization of markets (see [nss.org/EU](http://nss.org/EU)) linked to expanding the grid is essential, along with more grid intelligence. Korea, Russia and Scandinavia have great limitations in renewable resources, suggesting they might have a reason to take leadership in overcoming the sheer political barriers to affordable switchable energy from space [10] and to advanced quantum technology, while also building more transmission lines to better endowed neighbors. Neural network RLADP technology, used to enhance the “teleautonomy” design for controlling teams of robots (see <http://www.penguinasi.com/>), could be essential to the actual assembly of such minimum-cost power stations in space.

At the distribution and microgrid level, almost all researchers in power engineering agree that we need more use of “meshed networks,” both for security and for sustainability. Distribution operators tell me we already have more switching at the distribution level than most researchers know, and that the problem has been with systems to fully exploit even the switching we already have. (FERC has done studies illustrating the same issue at the transmission level.) Exciting results have been reported by several teams (e.g. [11]) showing how RLADP can substantially improve performance in such applications. Transmission companies such as Duke Power have reported that plug-in hybrid cars would be easy to accommodate at the transmission and power generation level, even if every household in the US had a pluggable car, but local distribution systems need substantial upgrades to make this possible.

Electricity also has its own crusading used-car salesmen, just as oil companies do. Many have shown great excitement about using intelligent systems in cars to sell power back to the grid. I was very amused to observe a deeply informative NSF panel where a proposal in that area was evaluated favorably by grid people, and then torn to pieces logically by an advanced researcher from the automotive industry. Battery lifetime is a really essential component of the utility functions for use in RLADP both for the grid (when batteries are used) and for vehicles. It is one of the key drivers of the costs of flexible vehicles [9]. Even now, lack of effective open-access crossdisciplinary research connecting electrochemists with RLADP engineers, for better battery modeling, is a crucial need for the future. It is unfortunate how policy changes since 2013 or 2014 have made it more difficult to meet this need effectively by government funding in the US – but holes in one area do imply new opportunities for others.

### 1.1. Related policy issues

Many policy analysts rightly ask us to say more about possible unintended consequences, especially when the “new AI” (neural networks) might be used to displace humans. Most of the applications discussed in this section would not displace humans, because they involve optimization at time intervals too fast for humans to control directly (milliseconds), or at levels of complexity already handled by older optimization programs (like economic dispatch), or in household energy systems where humans do not really want to make specific decisions every five minutes of the day. The choice of what to optimize comes from humans. Humans make the basic decision about values coming down from longer time intervals.

Nevertheless, there would be room for standards for home energy systems, to make sure that humans have full control, that their behavior is transparent, and that they obey security and privacy standards to be discussed in section 2.1. Economists rightly debate the issues about larger values to be input to these systems; for example, they debate how much carbon tax should be applied (if any) to the consumption of fossil fuels, which is an important input to the optimization. However, that is not anything new to the power industry, and it would not change the technical requirements on optimization systems as such.

Applications to the planning process are somewhat trickier. In the past, the human-based stakeholder debate systems and the algorithm-based optimization process have both been unable to fully capture what new energy projects, combining new transmission and generation together, could offer us. Inputs of that kind would help human stakeholder meetings to receive better options and inputs. Even in the most algorithm-based regions of the US, human investors decide in the end which projects to fund, regardless of what the planning process comes up with.

In truth, the greatest barrier to improved efficiency and greater use of low cost renewables in the US, as in Latin America [4], is the difficulty of building long-distance power lines between different regions of the US. Under the Interstate Commerce Clause in the US Constitution, the Federal Energy Regulatory Commission (FERC) has long had authority to cut through a myriad of local politics, and approve interstate natural gas pipelines. Edison Electric has long argued, rightly, that fair competition and efficiency demand a similar author at FERC to approve long distance transmission lines, such as lines proposed by Pickens to give him permission to build large new solar farms or wind farms in Texas to sell into the high-cost peak power market on the East Coast.

## 2. CRUCIAL ISSUES IN OUR PARTNER TECHNOLOGIES IN THE BIG PICTURE WITH ENERGY

Because of the synergies between intelligence and physical technology, we should pay attention to some of our partner technologies for electricity which affect the value of what we do.

### 2.1. Cybersecurity Issues

Perhaps the most urgent and important of these is the technology for cybersecurity [3]. When I was asked to look into cybersecurity for Specter’s office, and later when it fell into my mandate at NSF to coordinate interagency research related to nuclear terrorism, I was amazed at how shallow the knowledge is in most policy making in that area, more so than for the intelligent grid. There are good reasons for that; in the past, back doors both in operating systems and in chips were a key part of US national security strategy. But as of now, more and more massive leaks ([3] and beyond) make the old strategy unsustainable. Mainstream power engineering already has massive and important efforts to prevent damage from the ordinary forms of hacking which we see every day [12,13], but has left us naked before a larger, growing new risk.

AI is often used in intrusion detection, but brains alone are not enough to protect us from fatal viruses or malware; we urgently need a new kind of immune system as discussed in [3]. There is growing risk that an organized enemy could launch a massive attack, totally different from what the usual hacking security statistics pick up, enough to shut down half the US power grid, as bad as Electromagnetic Pulse (EMP) attack but not so easy to protect against. Critical power functions already use Security-Enhanced Linux, commonly provided by the software suppliers they rely on (like ABB or Siemens), often a slightly lagging version of what Red Hat provides with help from the Information Assurance Division of NSA (which was recently abolished as part of a new emphasis on AI and offensive capabilities). The most urgent need/opportunity is to upgrade that process, by still allowing back-doors but only read-only backdoors (useful for monitoring but not recontrol) and automated machine verification using formal methods (like those used in the development of the old Multics operating system) for use not only by electric utilities but all critical infrastructure. The lags must be eliminated.

The very best work in mainstream cybersecurity for power does recognize that unbreakable operating systems are important for the power system, and that major new efforts would be needed to make us safe again [13]. But for historical and political reasons, it tends to underestimate both how urgent and how practical such efforts would be, as discussed for example by the world’s leading authority in unbreakable operating systems [14], whose work is reflected in many sources available on the web internationally [15,16,17]. Automated open-source vetting of software prior to deployment is already

available [18], and making it scalable is no more difficult in principle than compiler design itself, which can cope with arbitrary levels of complexity in conforming systems. It would make sense to have a two-tiered system, in which big systems like android telephones or servers still may have read-only backdoors, but very small systems like chips used in heart pacemakers have no backdoors whatsoever.

Again, it should be emphasized that these new capabilities would not replace other urgent improvements in security which are now ongoing. However, they can make some of those other improvements easier and more effective. For example, in the crucial area of authentication, unbreakable operating systems can make it impossible for any but the most trusted users to have access to the parts of the system which could be used to create worst case damage like burned out generators.

Longer-term, security can be enhanced by using both new levels of quantum technology [20,21,22] which can be used with neural network designs either to enhance security or break codes in communication, or to help detect backdoors in chips. At a recent hearing of the Senate Intelligence Committee on global threats in general [23], senators and heads of intelligence agencies put special stress on machine learning, cybersecurity and new quantum technology. Many in Congress rightly give great emphasis as well to the threat from natural or malicious Electromagnetic Pulses (EMP), but low-cost technology to plug that hole in the security of the grid are already well known. (See <http://empactamerica.org>) I often recall the words of Congressman Trent Franks: "You are only worried because you don't have all the facts. If you had all the facts, you would be terrified out of your mind."

## 2.2. New Renewable Electricity Sources

Section 1 already discussed some of the key new issues in renewable electricity generation, discussed further in [4]. Here I will briefly mention two new details.

First, the competition between solar farms based on solar cells (PVs) and solar farms based on solar thermal power has become more challenging. Chile has accepted a Purchase Power Agreements (PPA) at under 3 cents per kwh from a PV solar farm, but still relies much more on solar thermal "power towers" at about 8 to 9 cents (even in the world's best, lowest cost site), because of the storage issue. New connections to Brazil would allow them to use the lower cost solar power [4], but it has not happened, in part because of legal/political risks but in part because many investors would not trust the 3 cents cost to remain in force after the Chinese suppliers of solar panels start to charge more. Leading solar companies in East Asia and West Asia and Europe have made intense confidential efforts to estimate the real long-term costs from China, and advised that we assume 9 to 10 cents long-term.

If the Chinese themselves have a basis for greater optimism, THEY may have basis for making the

investment, but other than that our best hope is for more aggressive new solar thermal technologies discussed in [4]. It is an extremely important unmet opportunity, clearly well-grounded in a proper understanding of Carnot's laws and new materials, but not moving ahead as yet as fast as it could be. Years ago, the company STM demonstrated solar dish technology, using Stirling engines with 30% efficiency to convert heat to electricity; we know that efficiency >50% can be attained using new engines mass-produced in existing engine factories, which would imply costs like 6 cents per kwh even in Texas, but the new Chinese owners of STM have focused their efforts more on marketing the existing proven designs, and Johansson has shown no interest in moving to China help them improve it. Al Sobey, former Division Director of General Motors for advanced products [24], has broken with Johansson, and led the development of a new engine design which he claims is more reliable and has efficiency >50%; more precisely, his new company, PDT LLC, claims to have created five generations of prototypes, all of which exceeded design and performance objectives, after years of effort, only now reaching out to external funding. The solid state JTEC technology for converting heat to electricity, which promises even higher efficiency in theory, has been used recently in a successful pilot demonstration for NASA, for use in space to increase the electricity output of small nuclear reactors.

Soon after the first draft of this paper was sent out for review, the US instituted anti-dumping rules against Chinese solar panels. This seems consistent with the best information available to the US, as described above. Yet even the best information is uncertain and incomplete, in the real world. If China has other information suggesting that they actually can produce solar panels at the low cost they now sell them for, in volume, without loss, then China's special knowledge would allow them to make trillions of dollars of profit by betting their money on this possible fact [4]. This would be fully consistent with China's wise new emphasis on low cost renewable energy based on solar farms in the best sites linked to the expansion of modern, low-cost transmission technology [19]. But even if they do, they would be well-advised to hedge their bets by advancing the best solar thermal technology in parallel with this.

Second, more radical relevant breakthroughs may also be possible via high-risk research in new quantum technologies [20-22,25]. There are times when I see uncanny parallels between the developing technology situation here and the classic novel *Atlas Shrugged* -- but not with the bit about draining the earth's magnetic field.

## 2.3. New Power Electronics, Motors and Fuel Flexibility

Most real energy experts understand that transportation fuel security is a more urgent issue than electricity supply as such, in most nations of the world, including the US [9]. Concerns about oil have grown larger and larger as a great complicating factor in international relations [3]. Greater fuel security and

diversity could be crucial in helping the nations of the world find more of a win-win, Pareto optimal resolution to the current life or death conflicts [3,9]. Electricity use in cars is part of that, but so also is greater ability to use liquid fuels in a more effective (and profitable, sustainable long-term) way.

IEEEUSA has already reviewed the most current options for alternate liquids [9]. (The new engines reviewed in [4] could also assist fuel-based operation in cars.) For pluggable cars, new technologies for power conversion in cars, for recharge stations at homes and in parking lots, and new technologies for electric motors and batteries all offer huge new unmet opportunities likely to succeed with the right kind of advanced research.

For example, oil industry spokesmen gave great publicity a few years ago to the fear that China's monopoly on rare earths needed for electric motors would cause the US to have a dependency problem worse than what we have with oil, if we relied too much on electric cars. (I say this as a witness, a fellow plenary speaker at an NDU conference on energy security in DC.) This was already a red herring, because China's advantage was in the process control used to extract raw earths (which RLADP can assist) and resources elsewhere have been put in service since. But Switched Reluctance Motors (SRM) and perhaps Induction Motors (IM) already offered greater whole-cycle efficiency than the older motors, without rare earths; the problem was with nonlinear control, which RLADP can handle but which simpler nonlinear controls already handle in advanced locations.

Batteries were already mentioned above. Universal on-board power converters, like those developed by Alireza Khaligh, combined with fast-recharge batteries like those proven on the road by BYD in China, could allow us to reduce the cost of pluggable vehicles and fast recharge stations substantially. RLADP control could help, particularly in highly flexible cars offering more options and security (and lifetime) to the user. More radical improvements in batteries are also available, to those not afraid to support the kind of research which NSF excelled in under Joe Bordogna and others of his kind, with full support from Congress for the original vision of Vannevar Bush. In practice, though Bordogna cared passionately about the broader impacts of research, there were options available to enhance those aspects... but the current version of the stakeholder system (aka "the swamp") dramatically reduces the effective throughput for intelligent decision-making, in my view. Failure to fund Excellatron's proven technology for rechargeable lithium-air batteries [26] was a bad omen for the US government in general (not NSF in particular). (DOE did provide funding for a time to Argonne, which had performed the first cycle tests proving that Excellatron's design had great promise, but not to Excellatron itself. Excellatron tells me that Argonne's efforts were shut down when they could not find a suitable electrolyte to go further, but Excellatron itself has now done so on its own funding.) Funding systems are

themselves a very important part of IT infrastructure, and a warning for the coming IOT.

### 3. THE POWER GRID AS A TEST PLATFORM FOR THE IOT IN GENERAL

The "new AI" is just one element of a massive transformation moving very quickly in the information technology (IT) industry, and in the industries affected by it (essentially the entire world economy). For several years now, there has been general agreement that we are moving from the old Internet to a new Internet of Things (IOT), which will control every vehicle (civilian or military), every factory, every generator, every household, every building and even every implant in every body on earth, in an integrated way. Major companies have been spending many billions of dollars to support their diverse, clashing visions of what this new IOT will look like [27-32]. People have grown up with the habit of assuming that IT is just one of many economic sectors controlled by the financial system, but we are moving quickly into a situation where the financial systems themselves are just pieces of the larger IOT which will control them, more and more. Banks have been spending a lot of money on traditional cybersecurity, but are far behind the power sector (section 2.1) in hardening their systems against coming massive attacks. There is also growing concern that current trends even in the existing Internet pose a massive new threat to freedom all over the earth, and that we need a massive rethink of our designs to avoid the very worst [33].

The six overview slides in [5] (developed for a high-level meeting in Silicon Valley on the future of IT) summarize my view of the coming challenges – above all, the challenge to those of us who have the technical capacity to develop a new paradigm here. See [2] for more explanation of the key slides. There is a serious risk that companies rushing to the market with a nonsustainable but simple product will create a terrible situation "on the ground", like what has happened at times in the past in parts of the IT industry, but this time with huge consequences harder to fix. Ironically, the core problem here is the local minimum problem, the risk that developers following local gradients to make money in the short term will move collectively to a situation which is not the best, and may even be nonsustainable, as I will discuss in section 4.

Some companies will simply ignore the issues of values and efficiency altogether, exposing us to the risks of what I call "Artificial Stupidity" (AS). AS already exists in our world, but the vast expansion and strengthening of the phenomenon could be as risky as "Terminator" style AI. Various others aim to control the world as a single intelligent system, maximizing values laid down from the top, unintentionally but systematically supporting a vast accelerating of inequalities and corruption eroding and disempowering human freedom and intelligence in a way leading to many other deep threats. The challenge to us is to grope as coherently as

we can towards a new type of general computer platform, somehow integrating and expanding all the five core positive developments shown in my “new paradigm” slide: (1) software lessons learned in the electric power sector, where values come directly from a huge array of human players, interacting via a rational and efficient market design, designed to avoid the abuses and irrationalities which early naïve power markets fell into (e.g. Enron); (2) the Penguin teleautonomy paradigm for controlling teams of robots, mentioned in section 1, which allow the use of RLADP in a weaker, more stable and safer form in individual robots responding to humans; (3) unbreakable operating systems and less breakable communications, as discussed in section 2.1; (4) advanced quantum technology to support all this; and (5), hardest but most important, IT specifically designed to elicit and mobilize and train the highest levels of natural human potential [35].

The power market is a perfect example of how “horizontal” IT companies could strengthen their position by developing new “vertical” products which do more complete justice to “specialized” issues (like dynamic market design and harder cybersecurity in electric power) which can actually be useful in larger markets after that. For example, if an upgraded version of the Red Hat type of product appeared, with apps for effective market design, what financial institution would want to be the last to offer harder security to its major clients? The challenge is to fully appreciate just how important the electric power system is as a prototype for the IOT of the future, and to upgrade it in a way which we would be ready to transfer to other sectors as soon as possible.

#### 4. FURTHER TECHNICAL ISSUES FOR THE IOT IN GENERAL

Reviewers of this paper have asked for more discussion of the larger dilemmas mentioned in section 3, going beyond what we can learn from the electric power example. Because I do not claim to have a complete answer, and because the issues implicitly involve very complex mathematics, I will have to be even more informal in this section than in the rest.

Above all, I worry more and more that the entire development of new IT technology is now following a kind of incremental path, descending to a kind of local minimum which is so dysfunctional that the human species may not even survive the outcome. The human species is now facing a variety of concrete life or death threats which require a very high level of collective intelligence to address effectively [2,5]. For example, there are serious threats of a “new” kind of climate change, involving emission of H<sub>2</sub>S from the oceans which, in the past, has led to conditions which would have killed every human on earth, if humans had been present at the time [4,38]. Coping with these threats, and developing more sustainable and more human-friendly core IT platforms, is basically an exercise in solving large local minimum problems in our collective intelligence. Of course, these local minimum problems are basically the

same as the nonconvexity problems familiar in general mathematics, and which include the barriers to entry and chicken and egg problems familiar in practical market economics.

These kinds of nonconvexity problems do not solve themselves. In intelligent RLADP systems [6], it is necessary to learn predictive capabilities able to simulate the future regions of state space which may entail survival or death, to use those capabilities in exploring those regions of space state and adapting value functions accordingly, and to develop a kind of cognitive map of the space of possibilities. Mammal brains do have that kind of capability hardwired into individual brains[34], but human societies (with or without IT) show little sign at present either of making full use of the relevant capabilities of individual human brains, or of building organizations with enough collective intelligence. The traditional organizational systems which drive human societies, from DNA and from money, lack aspects essential to achieving such capability, but, in principle, we could build IT platforms which do. (Market feedback cannot match correct modulated backpropagation training signals, simply because of the way it must be conserved.) Design for an Integrated Market Platform (IMP), integrating products from electricity to control of IOT devices to content, could be approached from this perspective, but there are difficulties even there.

Could human societies be organized somehow to be like a neural network system, capable of implementing a kind of collective intelligence which is more than the sum of its components? In fact, when I first developed backpropagation circa 1970, there was a serious chance that the PhD thesis would not be accepted, because neural networks were seen as such heresy in those days. It was possible, politically, to get the general algorithm accepted in a Harvard thesis, only because of the great support of my advisor, Prof. Karl Deutsch, whom I had chosen to work for because of his vision of how a human society could actually function like a neural network system [36]. (And also because of his important role behind the scenes in envisioning and guiding the European Union.) Dual Heuristic Programming [6], which is more powerful than the type of decision system used in the popular ALPHA-GO system, actually outputs value signals,  $\lambda_i$ , which are correct price signals. Nevertheless, human societies do not automatically compute these kinds of correct signals for an uncertain and nonlinear world. What is needed for building a hybrid system of humans controlling machines, to arrive at real collective intelligence?

A more effective integrated foresight capability would be one of the requirements. This by itself calls for a discussion more detailed than this entire paper. Every one of the traditional methods – neural network predictive models using massive time-series data, econometric style models, quantitative models with subjective inputs, artificial betting markets, or deep dialogue systems (like the real-time Delphis of [www.themp.org](http://www.themp.org)) -- has significant limitations, and would have to be enhanced somehow, and combined with others. IT platforms for

human collaboration, marshalling and cultivating the deepest latent ability of humans to work together [23], could help. The best experience from well-run review panels of the “old NSF” (before stakeholder politics and corruption effects became more powerful) give some indication of how to make this work for the bottom-up aspects of decisions, while an improved version of futurism could be integrated with these for a more effective strategic dialogue.

Another key challenge is how to manage the definition of global utility functions (whether implicit or explicit). Today’s literature on “cyberdemocracy” is woefully inadequate for the task. Even with social “penalty functions,” a world based purely on IT markets is likely to be dysfunctional in some ways, but even so much better than the more haphazard or centrally controlled systems now expanding rapidly, fostering centralization of power and diminishment of dialogue and effective collective intelligence. There are limits to how far cyberdemocracy can create a happy world, in any case; Malthusian effects have not gone away, and the revival of the nepotism-based revision of Confucianism in China (dating to about 1000AD) is no more promising than the money-based descendant of the original culture and social contract (attempted Pareto optimum) of the United States. There is an urgent need to work for a more integrated worldview and education, building on the deeper principles from both cultures but without the recent barnacles which get in the way of such synthesis, and also integrating the hard core rational and realistic culture of objective science with the deep spiritual sensitivities and vision which more enlightened cultures have cultivated for millennia [35,37].

Even though Malthusian effects and selection pressures will never go away for long in human life, open and semi-transparent and secure IMPs could at least allow more honorable competition, with less physical direct pain, and more utilization of human potential. New types of ledger systems, more secure and scalable than today’s cryptocurrencies, designed for linkage to IMPs, could be very useful at this time, for international cooperation to avoid leakage of funds from major economies to well-funded disruptive and criminal groups of all kinds which lie behind many of the current political problems across the world, sowing irrational and unnecessary conflicts and decay of actual human freedoms. (It is important to remember, as the Cambridge Analytics scandal shows us, that freedom can be reduced not only by governments but by unfettered private organizations, and that misbehavior in governments can often be traced back to well-heeled corrupt influences from outside.)

For the issue of death by H2S in particular, the new technologies discussed in sections 2 and 3 should help, but more is needed. If we had true collective intelligence here, we would understand the need for serious, honest focused efforts to learn what we can from aquarium-level research (using AI for effective assay of sulfanogenic archaea) about the precise chemical and physical conditions which allow proliferation of those

archaea, linked to greater foresight regarding ocean chemistry, as well as the development of options for geoengineering to try to restore the thermohaline currents of the Antarctic [38], and related measures. The threat of extinction by nuclear conflict (or Terminator AI or other negative syndromes) is also important, and also a great test case for collective intelligence, but beyond the scope of this paper.

## REFERENCES

- [1] Werbos, P. J. (2011). Computational intelligence for the smart grid-history, challenges, and opportunities. *IEEE Computational Intelligence Magazine*, 6(3), 14-21.  
<http://www.werbos.com/E/GridAI.pdf>
- [2] Werbos, P.J. (2018), The New AI: Basic Concepts, and Urgent Risks and Opportunities in the Internet Of Things, in R. Kozma, C. Alippi, Y. Choe and F.C. Morabito (eds), *Artificial Intelligence in the Age of Neural Networks and Brain Computing*, Elsevier 2018
- [3] Werbos, P. J. (2017). New technology options and threats to detect and combat terrorism. In *Identification of Potential Terrorists and Adversary Planning: Emerging Technologies and New Counter-Terror Strategies*, NATO/IOS, 132, 34.  
[www.werbos.com/NATO\\_terrorism.pdf](http://www.werbos.com/NATO_terrorism.pdf).
- [4] Werbos, P. J., & IntControl, L. L. C. (2016). Urgent Investment Opportunity To Scale Up Solar Industry in Atacama. [www.werbos.com/Atacama.pdf](http://www.werbos.com/Atacama.pdf), prepared for Solar Energy Research Consortium of Chile and Enersol 2016.
- [5] [www.werbos.com/IT\\_big\\_picture.pdf](http://www.werbos.com/IT_big_picture.pdf)
- [6] Lewis, F. L., & Liu, D. (Eds.). (2013). *Reinforcement learning and approximate dynamic programming for feedback control*(Vol. 17). John Wiley & Sons (IEEE Press).
- [7] Werbos, P. J. (2009, June). Putting more brain-like intelligence into the electric power grid: What we need and how to do it. In *Neural Networks, 2009. IJCNN 2009. International Joint Conference on* (pp. 3356-3359). IEEE.
- [8] Pereira, S., Oxygen Is Disappearing From The World’s Oceans At An Alarming Rapid Pace, *Newsweek*, Jan 15, 2018
- [9] IEEEUSA (2015), Transforming Transportation by Diversifying Energy Sources, policy position at <https://ieeusa.org/wp-content/uploads/2017/02/Transportation1015.pdf>
- [10] P. Werbos (2014), Reviewing Space Solar Power policy, *Ad Astra*, Vol. 26, No.2, 2014,  
<http://www.nss.org/adastra/volume26/ssppolicy.html>
- [11] Li, S., Fairbank, M., Johnson, C., Wunsch, D. C., Alonso, E., & Proao, J. L. (2014). Artificial neural networks for control of a grid-connected rectifier/inverter under disturbance, dynamic and power converter switching conditions. *IEEE transactions on neural networks and learning systems*, 25(4), 738-750.

- [12] He, H., & Yan, J. (2016). Cyber-physical attacks and defences in the smart grid: a survey. *IET Cyber-Physical Systems: Theory & Applications*, 1(1), 13-27. <http://digital-library.theiet.org/content/journals/10.1049/iet-cps.2016.0019>
- [13] Landwehr, C. E., & Valdes, A. (2017). *Building Code for Power System Software Security*. Technical Report. IEEE Computer Society, March. <http://www.landwehr.org/2017-05-building-code-for.pdf>
- [14] Marcus J. Ranum (2018), Building a Secure Operating System With Roger R. Schell, TechTarget, SearchSecurity, [searchsecurity.techtarget.com/opinion/Building-a-secure-operating-system-with-Roger-R-Schell](http://searchsecurity.techtarget.com/opinion/Building-a-secure-operating-system-with-Roger-R-Schell)
- [15] Stuart, B. L. (2008). *Principles of operating systems: design & applications*. Cengage Learning EMEA.
- [16] Fisch, E. A., & White, G. B. (1999). *Secure computers and networks: Analysis, design, and implementation*. CRC Press.
- [17] Wikipedia (2018), Common Criteria, [https://en.wikipedia.org/wiki/Common\\_Criteria](https://en.wikipedia.org/wiki/Common_Criteria)
- [18] J. Voas and K. Schaffer (2016), Whatever Happened to Formal Methods for Security?, *IEEE Computer Magazine* (Volume: 49, Issue: 8, Aug. 2016 )
- [19] Charlie Xu (2016), *A New Energy Network: HVDC Development in China* <http://cleanandsecuregrid.org/2017/01/02/a-new-energy-network-hvdc-development-in-china/>
- [20] Song, C., Xu, K., Liu, W., Yang, C. P., Zheng, S. B., Deng, H., ... & Zhang, P. (2017). 10-qubit entanglement and parallel logic operations with a superconducting circuit. *Physical review letters*, 119(18), 180511. <https://arxiv.org/pdf/1703.10302>
- [21] Werbos, P. J. (2014, May). Example of lumped parameter modeling of a quantum optics circuit. In *Quantum Information and Computation XII* (Vol. 9123, p. 91230A). International Society for Optics and Photonics. <https://arxiv.org/abs/1309.6168>
- [22] Werbos, P. J., & Dolmatova, L. (2016). Analog quantum computing (AQC) and the need for time-symmetric physics. *Quantum Information Processing*, 15(3), 1273-1287. [www.werbos.com/triphoton.pdf](http://www.werbos.com/triphoton.pdf)
- [23] U.S. Senate Select Committee on Intelligence, Open Hearing on Worldwide Threats, Feb. 13, 2018, <https://www.intelligence.senate.gov/hearings/open-hearing-worldwide-threats-hearing-1> and <https://www.c-span.org/video/?440888-1/national-security-leaders-testify-world-wide-threats&live>
- [24] Marina von Neumann Whitman (2016) , How GM Has Changed in 25 Years, *Detroit Free Press*, Jan. 14, 2016 <https://www.freep.com/story/opinion/contributors/2016/01/14/gm-electric-vehicles/78764942/>
- [25] M. Scully, INSPIRE: Photonic Quantum Heat Engines Including: Lasers without Inversion, Photo-Carnot Engines, Quantum PV Cells and Quantum Coherence Effects in Photosynthesis, [https://www.nsf.gov/awardsearch/showAward?AWD\\_ID=1241032](https://www.nsf.gov/awardsearch/showAward?AWD_ID=1241032)
- [26] L. Johnson (2010), The Viability of High Specific Energy Lithium Air Batteries, <http://www.rexresearch.com/johnsonliarbatty/Session4-350-Johnson.pdf>. See also <https://patents.google.com/patent/US20130130131A1/en>, [www.excellatron.com](http://www.excellatron.com), and cycle testing by Argonne.
- [27] Mike Kavis, The internet of things will radically change your big data strategy, *Forbes* magazine, June 26, 2014.
- [28] James Manyika, Michael Chui, Jacques Bughin, Richard Dobbs, Peter Misson and Alex Marrs, Disruptive technologies: advances that will transform life, business, and the global economy. [www.mckinsey.com/insights/business\\_technology/disruptive\\_technologies](http://www.mckinsey.com/insights/business_technology/disruptive_technologies)
- [29] M2M, Internet of things (IoT) companies gear up for \$6 trillion market demand, <http://mesh-systems.com/news/iot-companies>
- [30] Brad Bech, James Jamison, Ling Shao and Glenn Wightwick, The interconnecting of everything, IBM Academy of Technology, RedBooks, 2013
- [31] Michael Chui, Markus Loffler and Roger Roberts, The Internet of Things, McKinsey, March 2010, [www.mckinsey.com/insights/high\\_tech\\_telecoms\\_internet/the\\_internet\\_of\\_things](http://www.mckinsey.com/insights/high_tech_telecoms_internet/the_internet_of_things)
- [32] Simona Jankowski, James Covello, Heather Bellini, Joe Ritchie and Daniela Cost, *IoT primer: the Internet of things: making sense of the next mega-trend*, Global Investment Research, Goldman Sachs <http://www.goldmansachs.com/our-thinking/outlook/internet-of-things/iot-report.pdf>
- [33] Bowman, W. M., & Camp, L. J. (2012). Protecting the internet from dictators: technical and policy solutions to ensure online freedoms, *The Innovation Journal: The Public Sector Innovation Journal*, 18(1), 2013, article 3. [www.innovation.cc/scholarly-style/warigia\\_camp\\_bowman5edits18vi1a3.pdf](http://www.innovation.cc/scholarly-style/warigia_camp_bowman5edits18vi1a3.pdf)
- [34] P. Werbos, (2009). Intelligence in the brain: A theory of how it works and how to build it. *Neural Networks*, 22(3), 200-212.
- [35] Werbos, P. J. (2012). Neural networks and the experience and cultivation of mind. *Neural Networks*, 32, 86-95.
- [36] Deutsch, K. W. *The nerves of government; models of political communication and control* (No. 04; JA73, D4.) Free Press, 1963.
- [37] P. Werbos (2010), Links Between Consciousness and the Physics of Time, *International IFNA -ANS Journal* 741 "Problems of nonlinear analysis in engineering systems", [http://www.kcn.ru/tat\\_en/science/ans/journals](http://www.kcn.ru/tat_en/science/ans/journals). Also [www.werbos.com/Mind\\_in\\_Time.pdf](http://www.werbos.com/Mind_in_Time.pdf)
- [38] P. Werbos, Reply to public comments from Freeman Dyson expressing skepticism about climate change, March 12, 2018. <http://drpauljohn.blogspot.com/2018/03/reply-to-public-comments-from-freeman.html>