How Smart is Anything?

by J. Wunderlich PhD

This document available on-line at:
http://users.etown.edu/w/wunderjt/ITALY_2009/PUBLICATION_SUBMITPAPERdefininglimitsREVISED16b.pdf

Assignment:

1) Read the following publication:

   [1] Wunderlich, J.T. (2004). **Top-down vs. bottom-up neurocomputer design.** In *Intelligent Engineering Systems through Artificial Neural Networks, Proceedings of ANNIE 2004 Int'l Conference, St. Louis, MO.* H. Dagli (Ed.): Vol. 14. (pp. 855-866). ASME Press. ["Novel Smart Engineering System Design Award," 2nd runner-up best paper from over 300 submissions]
   Available on-line:
   http://users.etown.edu/w/wunderjt/ITALY_2009/PUBLICATION_SUBMIT_FINAL_ANNIE2004_WUNDERLICH_61_TO_PRINT_fixed_after.pdf

2) Read the attached revised version of this publication:

   [2] Wunderlich, J.T. (2003). **Defining the limits of machine intelligence.** In *Proceedings of IEEE SoutheastCon, Ocho Rios, Jamaica,* [CD-ROM]. IEEE Press.

3) Skim through this PPT lecture:
   http://users.etown.edu/w/wunderjt/ITALY_2009/TALK_MACHINE_INTELLIGENCE.pdf

4) Watch PBS Frontline Video "Digital Nation" (http://www.pbs.org/wgbh/pages/frontline/digitalnation/view/) and question Mental Abilities *#20* ***"MULTITASKING,"*** *#30* ***"SELECTIVE AWARENESS (FILTERING),"*** and *#42* ***"GROUP PSYCHOLOGY, SOCIAL NETWORKING, and LIVING IN THE CLOUD(s)"***

The original publication [2] above was published in 2003 and was later revised for lectures in 2003 by adding two Mental Abilities, #41 ***Awareness of Mortality****, and #42* ***Group Psychology****. I again updated this work in 2013 to include research in several new publications, and new survey data on mental abilities from recent iterations of my courses -- including more research on autonomy, and socially-networked engineering and architectural design. This revised paper is shown below.*

Defining the Limits of Machine Intelligence (updated in 2003 and 2013)

Joseph T. Wunderlich, Ph.D.
Elizabethtown College

Abstract - Machine Intelligence can be defined as encompassing all of the developments in both symbolic artificial intelligence and artificial neural networks. Traditional symbolic AI uses programmed heuristics and forms of knowledge representation to produce results in a seemingly more intelligent way than typical computer programs. Artificial neural networks are a form of connectionist computer architecture where many simple computational nodes are connected in an architecture similar to that of biological brains for the purpose of solving problems which require rapid adaptation or solutions where underlying governing equations are not known or cannot be easily computed. This paper explores the limits of machine intelligence by comparing the potential of these man-made systems to the "mental ability" of two common biological life forms; namely humans and bugs. The discussion begins with a study of basic animal abilities such as adaptation, self-preservation, motor-coordination, and processing complex sensory information. More advanced abilities are then explored including tool-manipulation, creativity, emotions, group psychology, and autonomy.

I. Introduction

The following discussion is organized around five simple questions:

1) What can a human do?
2) What can a bug do?
3) What can a conventional computer program do?
4) What can a symbolic AI program do?
5) What can an artificial neural network do?

The human in question is one of average mental ability and the "bug" is one with simple predatory and self-protection capabilities (e.g., a spider). The "conventional computer program" is assumed to be running on a typical uni-processor von Neumann type architecture machine. The last two questions assess the limits of machine intelligence. Traditional "symbolic" Artificial Intelligence (AI) programs use heuristics, inference, hypothesis testing, and forms of knowledge representation to solve problems. This includes "Expert Systems" and programming languages such as PROLOG and LISP, with the knowledge contained in the logic, algorithms, and data structures [2]. An artificial neural network is a form of connectionist computer architecture (hardware or software) where many simple computational nodes are connected in an architecture *similar* to that of a biological brain. The typical network is trained (i.e., learns) by changing the strength (weight) of inter-neuron connections such that multiple input/desired-output pairs are satisfied simultaneously; the final set of network weights represents the compromises made to satisfy multiple constraints simultaneously [1 to 5].

II. Discussion

An attempt to answer the five *simple questions* above is made for 20 different "basic animal abilities" and 22 "complex abilities." Each of the 42 is grouped by abilities which are often related.

(1) Acquire and retain knowledge, and (2) Solve problems: These are often assessed for humans by standardized tests such as the SAT exam for college entrance. Although the "verbal" and "quantitative" section of the SAT would likely be incomprehensible to a spider, a spider can solve simple problems such as where to place its web. It also needs knowledge of its environment and prey. All man-made computational devices, intelligent or not, can solve problems and retain knowledge; they only differ in memory capacity, method of storage, method of solving, and class of solvable problems.

(3) Learn and adapt: Both humans and spiders can easily learn and adapt to new environments and stimuli, and do so in both real-time and evolutionary time. Conventional computers have great difficulty with this. A human programmer is almost always needed to modify the programs. Traditional symbolic AI is somewhat adaptable to new input, however artificial neural networks are much better at this -- with an ability to generalize when presented new inputs. They can also learn very quickly when embedded in hardware [6, 7].

(4) Motor coordination, (5) Acquire energy, (6) Protect self: These have been referred to as "Mobility", "Acquisition", and "Protection" [8] and are essential for the survival of most animals. These have been somewhat implemented by conventional and *intelligent* machines (e.g., robotic motor control, power supplies, firewalls).

(7) Sensory processing, (8) Real-time thought, (9) React instinctively, (10) Anticipate, (11) Predict: Most animals sense their surroundings and think quickly and often instinctually what to do. They therefore can anticipate outcomes. They can also predict by extrapolating known information. With the exception of instinct, conventional and *intelligent* machines can also do these things; however neural networks outperform symbolic AI when dealing with new stimuli and can be much faster (especially if embedded in hardware [6,7]).

(12) Communication: Animals, conventional computers, and *intelligent* machines all communicate. However nothing comes close to what humans can do with natural language processing. Traditional symbolic AI has been attempted this for decades, however neural networks have had more recent success in speech recognition including the difficult understanding of "context"[1 to 5].

(13) Generalize: Generalize is *"to derive or induce a general principle or concept from particulars"[9].* Animals do this well. Conventional computers don't; they give very specific responses to very specific inputs. Symbolic AI can only do this to the extent the program has been built with variations to consider. Neural networks are very good at this; with the ability to generalize such that outputs are produced which "best fit" (i.e., classify) a set of inputs (even when they differ from what the network was trained with).

(14) Associate, (15) Recognize patterns: All animals do this well; however no animal surpasses the human's ability to associate concepts and memories. Conventional computers do this in a very limited sense; they can associate by correlating data and can recognize the simple encoded patterns of bit-streams input by humans and other machines. Symbolic AI programs do this better, but are still limited by the fixed structure (i.e., the "*state-space*" is fixed regardless of how efficiently it is searched). Neural networks are very good at association – with an ability through generalization to associate patterns such as never-seen hand-written characters to recorded ASCII representations. Neural networks are widely used for recognizing image and speech patterns [1,4].

(16) Robust under partial failure: Evolution has insured that animals can often continue to function when one or more subsystems fail (including parts of the brain). Conventional computers can't do this to any significant degree; even a simple one-bit error in program execution can sometimes cause a system to "lock-up." Symbolic AI programs running on conventional computers (or even super-computers) are also likely to not function when the underlying computer system fails. Neural networks are very robust under partial failure and have the ability to partially function when some neurons or inter-neuron connections fail [1 to 5].

(17) Autonomous thought: Most animals are free to make their own decisions. Conventional computers and symbolic AI programs are not autonomous unless they the software developer creates the code to run without any human intervention or oversight. Neural networks, with their ability to learn, generalize, deal with never-seen input, and think in a distributed fashion do have the potential to become entirely autonomous. Therefore, all computer and forms of machine intelligence can be given autonomy if the humans so choose; but how tightly should we hold the leash? [20, 21, 22, 25].

(18) Drive to reproduce: With the exception of programming dictated by genes (including the drive to reproduce), many animals, and humans, are free to make their own decisions including suppressing the urge to reproduce. All machines are nowhere close to wanting to reproduce (unless someone programs this). But it's not beyond the realm of possibility that someday far in the future *intelligent* machines could decide to reproduce.

***(19) Stability, Repeatability, Predictability*:** There is a definite degree of uncertainty associated with all animal behavior. "*physics has managed to incorporate uncertainty into its prospectuses, and there is no reason to believe that the scientific study of behavior can not successfully incorporate a "biobehavioral uncertainty principle" as well.....Intrinsic variability not only removes the spectra of absolute predictability, but may provide a basis for admitting more fully into scientific discourse the concept of free will.... behavior is fundamentally exploratory" [10]*. Conventional computers and symbolic AI don't have this problem (or virtue). They simply respond in a pre-programmed way. Neural networks however can produce unexpected results; especially when dealing with never-seen input.

(20) Multitask: The evolution of most biological life has led to brains with multiple subsystems working in a coordinated fashion; some performing basic system regulation (e.g., pulmonary, respiratory, temperature, and motor control), some pre-processing information before relaying it to higher reasoning centers (e.g., visual cortex), and some performing higher reasoning [13]. Conventional computers are becoming better at this, with subsystems performing tasks simultaneous to the functioning of the CPU (Central Processing Unit). Examples are DMA (Direct Memory Access), and graphics-board processors [11,12]. It's important to note that "multitasking" in computer industry nomenclature often implies time-sliced use of the CPU and not *true* simultaneous, parallel functionality. This is one reason to be careful when comparing human performance with typical uni-processor computer performance. When discussing brain performance, one must consider the brain's high degree of parallelism and pre-processing. Multitasking is typically only found in symbolic AI programs when written for multi-processor machines. Multitasking is however a significant part of artificial neural network learning. Recent research shows that humans immersed in technology reach a limit where multitasking adversely effects other mental abilities [24].

(21) Abstraction, (22) Intuition, (23) Common sense: Abstract is: "*having only intrinsic form with little or no attempt at pictorial representation or narrative content" [9].* Intuition is: *"Knowing without conscious reasoning" [9].* Combining these definitions can yield insight into the more complex workings of the human brain (i.e., partially defined or disconnected thoughts could lead to higher reasoning). Conventional computers and symbolic AI programs simply respond in a pre-programmed way. The ability of neural networks to learn by repeatedly modifying inter-neuron connection weights until a compromise is reached could be a form of abstraction. Common sense is: *"Sound and prudent but often unsophisticated judgment" [9].* Some very analytical people are sometimes said to not have common sense; perhaps the need for logic and "sophisticated judgment" to prove hypotheses may hinder the ability to temporarily think in a disconnected fashion – even if an abstract, intuitive, and somewhat unsophisticated thought could lead to a common sense answer. New research is showing that advances in computing will soon yield machines that can exhibit these qualities by drawing from the equivalent of multiple brain centers simultaneously [23].

(24) Manipulate tools: Although a spider can design and construct elaborate webs, it is not likely to envision extensions of its appendages (i.e., tools) to do so. Manipulating tools is exclusive to more evolved animals and arguably can be attributed to humans becoming bipedal; allowing our front "feet" to become hands for manipulating tools [14]. Conventional and *intelligent* computational systems can also manipulate tools by creating signals to send to actuators (e.g., motors, etc.), which in turn position and orient tools. This is a definition of robotic-arm control. Not only what a robotic arm holds, but the arm itself can be considered a tool for the computer to realize manipulation of the physical world around it.

(25) Heuristics, (26) Inference, (27) Hypotheses testing: Most animals don't consider every possible way to react to a situation before acting (i.e., an exhaustive search); they instead apply heuristics to more efficiently select an action. They also recognize when one scenario infers another, and can solve problems by testing multiple hypotheses to result in one solution. Conventional computer programs only somewhat do this. Symbolic AI programs (especially "Expert Systems") can do all of these things [2]. Most neural networks however are not well suited for the step by step process needed to apply heuristics or hypothesis test, but can somewhat infer results for given input data (including never-seen input).

(28) Self-discipline & Impulse control, (29) Ethical behavior: Despite genetic, instinctual, programmed animal "drives," humans can override their programming to maintain a level of self-restraint, and can even develop a set of rules (i.e., ethics and values) to maintain civilization. Bugs seem to act purely instinctually and show no signs of ethical imperatives. Conventional computer programs are incapable of these things; however symbolic AI programs can incorporate all of the rules (and therefore ethics and values) of a given human. Also, you could train a neural network to respond "ethically" to given situations.

(30) Selective awareness (filtering): Most animals have the ability to focus on a task while ignoring distractions such as extraneous noise or motion around them. They are also able to find images semi-obscured by camouflage or clutter. Conventional computer programs and symbolic AI programs can achieve this through pre-processing of input data by using signal and image processing techniques. Also, several types of neural networks, with their ability to generalize and deal with never-seen input, can perform very well when given "fuzzy" input [1 to 5]. Recent research shows that humans immersed in technology reach a limit where multitasking adversely effects their ability to focus [24].

(31) Open to inspection: Despite many years of scientific advances in understanding both the biological and behavioral function of animals brains, tracing mental thoughts is still less "exact" then tracing the execution of a conventional or symbolic-AI program. Neural networks are less open to inspection than programs because of the many compromises made in changing inter-neuron weight values during the training (learning) phase (i.e., to satisfy many input/desired-output pairs simultaneously).

(32) Emotions, (33) Imagination, (34) Creativity, (35) Passion, (36) Playfulness: The ability to feel, to imagine and create, to have passions and ambitions, and to experiment through playful curiosity are still primarily human traits; and although other animals may exhibit these abilities, it is unclear what a spider can think in these regards. Play seems to have also played an important role in evolution: *"Given that the adaptiveness of behavior itself derives from an evolutionary process in which variability and play are absolutely essential .....playfulness is indeed not only to be enjoyed but to be accorded high value for its fundamental role in the success of all organisms, including human" [10].* No man-made device is yet capable of these things. New research is showing that advances in computing will soon yield machines that can exhibit these qualities by drawing from the equivalent of multiple brain centers simultaneously [23].

(37) Empathy, (38) Courage, (39) Leadership: The ability to empathize with the feelings of others, to take risks including self-sacrifice for the benefit of others, and to display leadership qualities (e.g., vision, compassion, motivation of others), are still primarily human traits; and although other animals may possess these mental abilities, it is unlikely a spider does. No man-made device is yet capable of these things. However, simple programmed responses to perceived human emotion are now possible [15].

(40) Self-Awareness, (41) Awareness of mortality: It is unlikely that a spider could recognize itself in the mirror or could clearly recognize impending doom. However, humans can see themselves, their lives, their influence on others, their influence on the future, and their mortality. Conventional computer programs can't do these things. Also, it seems unlikely (but not impossible) that *intelligent* machines could ever become self-aware. However, it is very likely they could achieve immortality as long as there is an ample supply of replacement parts.

(42) Group Psychology, *Social Networking, and Living in the Cloud(s)**:*** Humans can play, work, raise children, and wage war as teams. They can also collectively share beliefs. Although some bugs work in a collective (e.g., ants, bees, etc.), most spiders appear to be isolated thinkers. Networked conventional computer programs and *Intelligent* machines, especially if implemented with parallel processing architectures, have the potential to implement the equivalent of group psychology, and new research in Social Networking and Crowd Sourcing shows that humans can collectively achieve as teams of virtual avatars in semi-realistic simulated worlds [26].

III. Intelligent Machine Platforms and Devices

Typical predictions of when computer performance will reach that of the human brain employ Moore's Law to extrapolate increases in computing speed or number of transistors per chip:

$$Qnew = Qold * (2 \wedge (n/1.5)) \qquad (1)$$

Where Qold is today's computing speed (or chip density), and Qnew is the computing speed (or chip density) expected n years in the future (i.e., speed and chip density double every 18 months). Although this law remains valid to-date, it must eventually break down. If we look less than a hundred years into the future, assuming a present day Qold speed of 3Ghz and a chip density of 10 million transistor per chip, Moore's Law predicts a Qnew that would require electricity to travel through a transistor faster than the speed of light and more transistors on a chip than the number of atoms that would fit in that volume. This type of prediction can also be misleading if one doesn't consider the relative degree of parallel processing that occurs in many biological and man made *intelligence* systems. A significant problem to solve is multitasking manmade subsystems as efficiently and elegantly as the human brain. To explore the limits of multitasking in machine intelligence, combine the understanding of mental abilities as discussed above (and summarized in Table 1.) with an understanding of the "Levels of Computing" as defined in Table 2. The degree of parallelism (DOP) [11] needed to be comparable to a human brain is simply not found in PC's, Workstations, or even mini-computers. Only in some supercomputers does the parallelism begin to become close to what might be required. Some embedded systems may however be able to achieve these goals by having many simple devices working independently [7]; however most embedded systems lack the computation power (and precision) of even the simplest PC [16]. They have the DOP but not the processing power.

Multitasking is typically only found in symbolic AI programs when written for multi-processor machines. Multitasking is however a significant part of artificial neural networks where learning occurs between the many simple computational nodes. This can be compared to MPP (Massively Parallel Processing) supercomputers. If an MPP machine could be built with billions of nodes (like the human brain), instead of just thousands (to-date), it could possibly implement an artificial neural network to rival all of the functionality of the human brain.

Another hurdle to overcome for those hoping to build *intelligent* machines that rival human brains is choosing an architecture that is either:

1. structurally similar to, or
2. merely produces results in a *similar fashion* to

the human brain (i.e., "bottom-up" vs. "top-down" design). Most artificial neural networks are top-down designs which learn and can be trained to react to external stimuli such that they mimic certain biological brain function. They learn by repeatedly applying mathematics to change inter-neuron connection strengths (weights) until the outputs converge to desired tolerances [1,3,4]. The network is trained (i.e., learns) by changing the strength of connections such that multiple input/desired-output pairs are satisfied simultaneously; the final set of weights represents the compromises made to simultaneously satisfy the constraints. A major problem in implementing this is that these computations require matrix and vector manipulations, but are often run on von-Neumann type uni-processor machines that have a "bottle-neck" forcing non-parallel computations. SMP (Symmetric Multi-Processing) machines can improve performance; however the best machines for these calculations are MPP or vector-register supercomputers, or embedded, application-specific, highly parallel systems – especially those which can provide learning in real-time. The all-digital vector-register neural network processor (with on-chip learning) proposed by Wunderlich in [7] is one example of this.

The "bottom-up" approach is to build a man-made system which functions like a biological brain at the circuit-level. The theory in [17] is to build artificial dendritic trees as RC analog circuit elements (i.e., built with resistors and capacitors) that produce signals close to those propagating through the dendritic tree inter-neuron connections of the human brain. Fig. 1 is a VLSI chip built by Wunderlich [18] to implement this theory. It has 64 neurons built from approximately 10,000 transistors on a 2mm x 2mm die.

Even though the semiconductor industry continues to find ways to increase the number of transistors per unit area, the chip-area required to include billions of neurons (like that of the human brain) would need to be millions of times larger than a typical chip. One reason for this is that our brains are three-dimensional whereas integrated circuits are mostly two-dimension (despite multiple levels of layerization). Another problem is connecting all of these neurons since the wire routing would be in mostly two dimensions. Even with several layers of metallization (for wires), it would be extremely difficult to connect billions of neurons (with each requiring thousands of connections to other neurons). Perhaps the most difficult problem to overcome with this type of implementation is mimicking human learning where

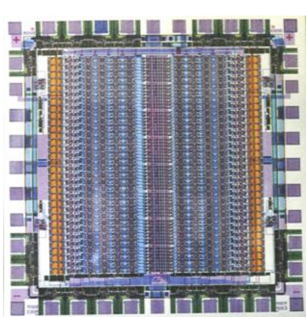

inter-neuron connections are not only strengthened or weakened during learning, but are often grown. Wires on chips need to be fixed, or at-best of variable resistance, and considering the required extensive connectivity between billions of neurons, would likely take many years to be realized.

Figure 1. Neural network chip by Wunderlich [18, 19].

Table 1. Mental Ability Matrix

		Can human do?	Can bug do? (spider)	Can Conventional Computer Program do?	Can Symbolic AI Program do?	Can Artificial Neural Network do?	Comments
	BASIC ANIMAL ABILITIES:						
1	Acquire and retain knowledge	yes	yes	yes	yes	yes	
2	Solve problems	yes	yes	yes	yes	yes	
3	Learn and adapt	yes	yes	no	somewhat	yes	Evolution
4	Motor coordination	yes	yes	somewhat	somewhat	somewhat	Survival
5	Acquire energy	yes	yes	somewhat	somewhat	somewhat	Survival
6	Protect self	yes	yes	somewhat	somewhat	somewhat	Survival
7	Sensory processing	yes	yes	yes	yes	yes	
8	Real-time thought	yes	yes	yes	yes	yes	
9	React instinctively	yes	yes	no	not yet	not yet	
10	Anticipate	yes	yes	yes	yes	yes	
11	Predict	yes	yes	yes	yes	yes	
12	Communicate	yes	yes	yes	yes	yes	
13	Generalize	yes	yes	no	somewhat	yes	
14	Associate	yes	yes	somewhat	somewhat	yes	
15	Recognition patterns	yes	yes	somewhat	somewhat	yes	
16	Robust under partial failure	yes	yes	no	no	yes	
17	Autonomous thought	yes	yes	if programmed	somewhat	soon	How tightly to hold the leash?
18	Drive to reproduce	yes	yes	no	not yet	not yet	
19	Stability, repeatability, predictability	somewhat	somewhat	yes	yes	somewhat	Uncertainty
20	Multitask	to a point	yes	yes	no	yes	
	COMPLEX ABILITIES:						
21	Abstraction	yes	unlikely	no	no	somewhat	
22	Intuition	yes	unlikely	no	not yet	soon	
23	Common sense	yes	yes	no	not yet	soon	
24	Manipulate tools	yes	no	yes	yes	yes	Evolution
25	Heuristics	yes	yes	somewhat	yes	no	
26	Inference	yes	yes	somewhat	yes	somewhat	
27	Hypothesis testing	yes	somewhat	somewhat	yes	no	
28	Self-discipline, impulse-control	yes	unlikely	no	somewhat	no	
29	Ethical behavior	yes	unlikely	no	somewhat	somewhat	If coded/trained
30	Selective awareness (filtering)	to a point	yes	yes	yes	yes	
31	Open to inspection	somewhat	somewhat	yes	yes	somewhat	
32	Emotions	yes	unlikely	no	not yet	soon	
33	Imagination	yes	unlikely	no	not yet	soon	
34	Creativity	yes	unlikely	no	not yet	soon	
35	Passion	yes	unlikely	no	not yet	soon	
36	Playfulness	yes	unlikely	no	not yet	soon	Evolution
37	Empathy	yes	unlikely	no	not yet	soon	
38	Courage	yes	unlikely	no	not yet	soon	
39	Leadership	yes	unlikely	no	not yet	not yet	
40	Self awareness	yes	unlikely	no	not yet	not yet	
41	Awareness of mortality	yes	unlikely	immortal?	Immortal?	Immortal?	Replaceable parts
42	Group psychology, Social Networking, and Living in the Cloud(s)	yes	unlikely	somewhat	somewhat	somewhat	Networking, Crowd-sourcing, Socially-networked design

IV. Conclusions

Machine Intelligence can be defined as encompassing all of the developments in both symbolic artificial intelligence and artificial neural networks. This paper explores the limits of machine intelligence by comparing the potential of these man-made systems to the "mental ability" of two common biological life forms; namely humans and bugs; this also led to a discussion of various machine platforms and devices for implementing *machine intelligence*, and to the observation that perhaps complex mental abilities like emotions and creativity, which are now known to be processes that draw on many parts of the brain [23], can be compared to the Group psychology of Social Networking,, which in a macro way mirrors the complex collaborative micro processes in the brain for emotions and creativity. Designing over the internet using crowd-sourcing and socially-networking is discussed in [26] and shown in Figures 2.

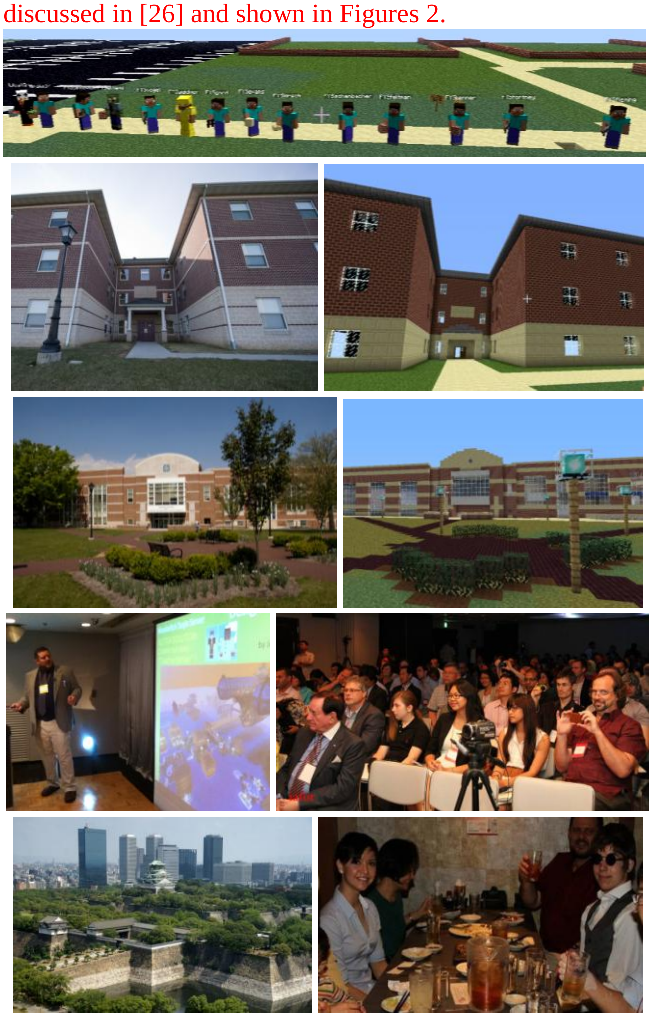

Figure 2. 2012 Elizabethtown College Architectural crowd-sourcing projects; and related 2013 key-note talk and paper in Osaka, Japan [26].

References

[1] S. Haykin, *"Neural networks, a comprehensive foundation."* 2nd ed. Upper saddle River, NJ: Prentice-Hall, 1999.

[2] G. Luger and W. Stubblefield, *"Artificial intelligence."* 2nd edition, Benjamin/Cummings, 1992.

[3] J. Hertz, A. Krogh, and R. G. Palmer, *"Introduction to the theory of neural computation."* Redwood City, CA: Addison-Wesley, 1991.

[4] B. Soucek, *"Neural and concurrent real-time systems."* New York: John Wiley & Sons, 1989.

[5] B. Soucek and M. Soucek, *"Neural and massively parallel computers."* New York: John Wiley & Sons, 1988.

[6] J. T. Wunderlich, "Simulation vs. real-time control; with applications to robotics and neural networks," in Proc. of ASEE Nat'l Conf., 2001, Albuquerque, NM.

[7] J. T. Wunderlich, "A vector-register neural-network microprocessor with on-chip learning," Masters thesis, Pennsylvania State University, Jan. 1992.

[8] Hasslacher, Brosl, and Tilden, "Living machines", Los Alamos National Laboratory Report, Jan 3, 1996

[9] *Webster's Ninth new collegiate dictionary*, Merriam Webster, 1983.

[10] Paul Grobstein, *"The encyclopedia of human behavior, volume 4",* Academic Press, 1994 (pp 447-458)

[11] K. Hwang, *" Advanced computer architecture: parallelism, scalability, programmability ,"* Boston, MA: McGraw-Hill, 1993.

[12] Patterson and J. Hennessy, *"Computer organization & design,"* 2nd. ed., Morgan Kaufmann, 1998.

[13] Sherman and Sherman, *"Biology, a human approach,"* New York: Oxford University Press, 1983.

[14] C. Ward, (2003, February 1), "*Big mystery: what made us stand tall,*" [online], Available: http://www.msnbc.com/news/213934.asp

[15] Breazeal. C., (2003, February 1), "*Kismet,"* [online], Available: www.ai.mit.edu/projects/humanoid-robotics-group/kismet/kismet.html

[16] J. T. Wunderlich, "Focusing on the blurry distinction between microprocessors and microcontrollers, " in Proc. of ASEE Nat'l Conf., 1999, Charlotte, NC.

[17] J. G. Elias, "Artificial dendritic trees," Neural Computation, vol. 5, pp. 648-663, 1993.

[18] J. T. Wunderlich, et al., "Design of an artificial dendritic tree VLSI microprocessor," Univ. of Delaware research report, Jan. 1993

[19] Wunderlich, J.T. (2004). Top-down vs. bottom-up neurocomputer design. In *Intelligent Engineering Systems through Artificial Neural Networks, Proceedings of ANNIE 2004 Int'l Conference, St. Louis, MO.* H. Dagli (Ed.): Vol. 14. (pp. 855-866). ASME PressWunderlich,

[20] J.T. (2004) Simulating a robotic arm in a box: redundant kinematics, path planning, and rapid-prototyping for enclosed spaces. In *Transactions of the Society for Modeling and Simulation International*: Vol. 80. (pp. 301-316). San Diego, CA: Sage Publications.

[21] Painter J. and Wunderlich, J.T. (2008). Wunderbot IV: autonomous robot for international competition. In *Proceedings of the 12th World Multi-Conference on Systemics, Cybernetics and Informatics: WMSCI 2008, Orlando, FL*: (pp. 62-67).

[22] Coleman, D. and Wunderlich, J.T. (2008). O^3: an optimal and opportunistic path planner (with obstacle avoidance) using voronoi polygons. In *Proceedings of IEEE the 10th Int'l Workshop on Advanced Motion Control, Trento, Italy*. vol. 1, (pp. 371-376). IEEE Press.

[23] Minsky, Marvin (2007), *The Emotional Machine, Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind,* Simon & Schuster.

[24] PBD Frontline Video "Digital Nation" (2008) available on-line at: http://www.pbs.org/wgbh/pages/frontline/digitalnation/view/

[25] Wunderlich, J.T. (2011). Designing robot autonomy: how tightly should we hold the leash? *The 5th Int'l Conference on Design Principles and Practices*, Rome, Italy.

[26] Wunderlich, J.T. and Wunderlich, J.J. (2013). Green architecture and environmental design using rapid-prototyping social-networking sandbox tools, followed by professional architectural software. *Asian Conference on Sustainability, Energy & the Environment* (ACSEE 2013), June 6-9, Osaka, Japan.

Table 2. Levels of Computing

LEVEL	TYPICAL APPLICATION	CHARACTERISTICS	HARDWARE And DEVICES	OPERATING SYSTEMS	FOR MACHINE INTELLIGENCE?
Embedded	Real-time control, automobiles, appliances, factory automation	Cheap, small, and often fast	**Microcontroller:** (Intel, Motorola, PIC's) **Microprocessor: (**Intel, Motorola, PowerPC) **ASIC:** (Application Specific IC's)	None or custom	Good for high-speed real-time-learning neural network applications. Not usually used for symbolic AI programming
PC	General-purpose "low-end computing"	Usually faster than embedded, but otherwise relatively slow, < ~$5000	Microprocessor (Intel, Motorola, PowerPC)	Windows, DOS, MAC OS, B, Linux	Acceptable for neural network simulations and symbolic AI programming
PC Server or Workstation	LAN server for ~100 people, 3-D simulations, VLSI circuit design (e.g., "Cadance")	Fast, ~$3000 to ~$20,000	Multiple microprocessors (Intel, Motorola, PowerPC) Silicon Graphics , SUN or IBM RS6000 workstations	Windows NT, UNIX, AIX	Good for neural network simulations and symbolic AI programming
Mini-Computer	LAN server for ~500 people	Fast, ~$100,000	IBM AS400, Amdahl, HP, Hitachi	UNIX, MVS, VMS, OS 390	Good for neural network simulations and symbolic AI programming
Super-Computer	**SMP:** LAN, WAN, or Internet server for 1000's of people, Air traffic control, NYSE **MPP:** Grand challenge applications, Chess **Vector:** Matrix-intensive grand challenge applications	Extremely fast, ~$1,000,000 to ~$10,000,000	**SMP**: IBM S/390, **MPP**: IBM SP2 (e.g., "Deep Blue"), **Vector**: CRAY	**SMP**: UNIX, MVS, VMS, OS 390 **MPP**: custom distributed OS **Vector**: custom vector OS	Very good for neural network simulations and symbolic AI programming

JOSEPH T. WUNDERLICH

Dr. Wunderlich is an Associate Professor of Engineering at Elizabethtown College. Previously, he worked for Purdue University as an Assistant Professor and for IBM as a researcher and hardware development engineer. Dr. Wunderlich received his Ph.D. in Electrical and Computer Engineering from the University of Delaware, his Masters in Engineering Science/Computer Design from The Pennsylvania State University, and his BS in Engineering from the University of Texas at Austin. He has been involved with Robotics and Machine Intelligence for over 22 years.