The Elusive Quest for Intelligence in Artificial Intelligence

Dr Shoumen Palit Austin Datta

MIT Auto-ID Labs
Department of Mechanical Engineering
Massachusetts Institute of Technology
shoumen@mit.edu

MDPnP Lab
Massachusetts General Hospital
Harvard Medical School
sdatta8@mgh.harvard.edu

The elusive quest for intelligence in artificial intelligence prompts us to consider that instituting human-level intelligence in systems may be (still) in the realm of utopia. In about a quarter century, we have witnessed the winter of AI (1990), being transformed and transported to the zenith of tabloid fodder, about AI (2015). The discussion, at hand, is about the elements that constitute the canonical idea of intelligence. The delivery of intelligence as a pay-per-use-service, popping out of an app or from a shrink-wrapped software or AI-as-a-Service, is in contrast to the bio-inspired view of intelligence, formed from a tapestry of events, often cross-pollinated by instances, each with its own microcosm of experiences, and learnings, which may not be discrete, all-or-none functions but continuous, over space and time. The enterprise world may not require, aspire or desire such an engaged solution, to improve its services for enabling digital transformation, through the deployment of digital twins, for example. One might ask whether the "work-flow on steroids" version of decision support may substitute for intelligence? Are we harking back to the era of rule based expert systems? The image conjured by the publicity machines offers solutions with human-level AI and preposterous claims, about capturing the "brain in a box" by 2020. Even emulating arthropods (leave alone cephalopods) may be difficult, in terms of rational AI. Perhaps we can try to focus on worms (*Caenorhabditis elegans*) which may offer what businesses may need, to *pareto-quench* its thirst, for so-called intelligent applications based on elements of AI.

The Elusive Quest for Intelligence in Artificial Intelligence

Is intelligence an illusion in artificial intelligence? This essay is an academic blog. The opinions are due to the author.

Dr Shoumen Palit Austin Datta, MIT Auto-ID Labs, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA and MD PnP Lab, Massachusetts General Hospital, Harvard Medical School

PLEASE OPEN DOCUMENT WITH ADOBE ACROBAT TO ACTIVATE THE EMBEDDED HYPERLINKS

The promise, and pessimism, about AI (artificial intelligence) charts a sinusoidal path. The questions about "intelligence" in AI persists. Many more questions are meeting with fewer answers. Intelligence in AI may be an advanced form of multi-level dynamics with/without a higher order of structured complex networks or (even better) temporal networks. The semantics depends on, and is colored by, one's view of systems or interpretation of what is intelligence, what constitutes proof of intelligence, and the nature of complexity.

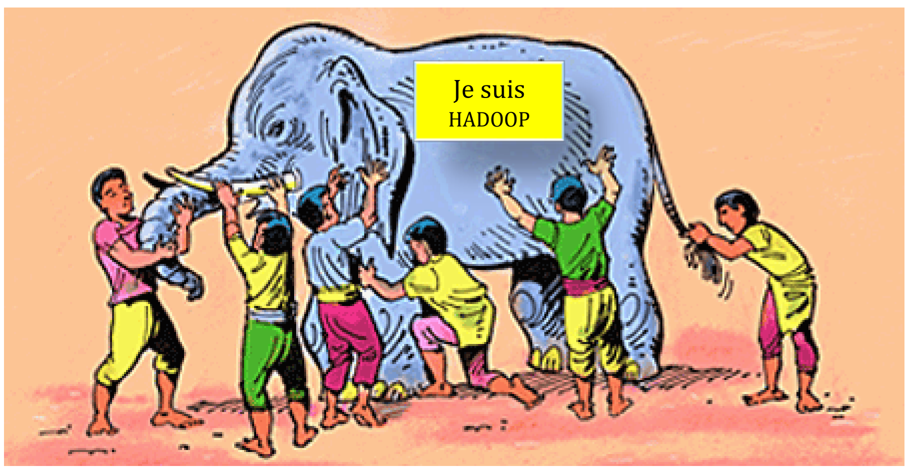

Figure 1: The Blind Men and the Elephant ▪ John Godfrey Saxe (1816-1887)

In systems thinking a common analogy is that of six blind men from Indostan (India) touching various parts of an elephant and claiming that "elephant is a tree" (man who touched the leg), "it is like a rope" said the blind man who touched the tail. "Like a snake" (man who touched the trunk). "It is like a big hand fan" (man who touched the ear). "Like a huge wall" said the man who touched the belly of the elephant. "Like a spear" said the blind man who touched the tusk. They were misguided in describing their part and wrong about the whole picture. The definition of intelligence may be colored by the professional bias of the interpreter. Intelligence is not a point. It is a fabric or an array or collective continuum of network systems which may not be boxed with human skills, yet (the latter is by far the most ludicrous and incredible claim, at the present time, under our known circumstances).

Why this mad pursuit ?

University of Cambridge, which had done best at teaching mathematics, is the one from amongst whose graduates, have come more of the English poets. While Oxford, which has specialized in the humanities, has tended to turn out writers who have attained, on the whole, a high level of mediocrity. By the time one has discussed literature with a witty and learned professor, one knows what has been achieved, and how good it is. You become respectful, and begin to wonder, who am I to do better? (Dialogues of Alfred North Whitehead (1861-1947) as recorded by Lucien Price. Little, Brown ▪ Boston, 1954)

The definitive treatise by Stuart J Russell and Peter Norvig (*Artificial Intelligence: A Modern Approach*, 3rd edition, 2015) should suffice to discourage all amateurs (myself included) to desist from testifying about the trials and tribulations with respect to AI.

Who am I to do better ?

In this neo-Norvig-ean era of human-level AI claims, we find LISP programming language, the 2nd oldest language, since FORTRAN. It is still regarded as a powerful AI tool. Figure 1 in the first LISP paper (John McCarthy, MIT) illustrates "Representation of S-Expressions by List Structure" which takes us back to the history of perceptrons elaborated in the seminal paper by Warren McCulloch and Walter Pitts (1943) *A Logical Calculus of the Ideas Immanent in Nervous Activity* in Bulletin of Mathematical Biophysics **5** 115-133 (marks the dawn of AI).

The perceptron algorithm created by Frank Rosenblatt, was championed by the US Navy, in fueling the "intelligence" controversy, to the extent, that *The New York Times* reported (in 1958) the perceptron to be "the embryo of an electronic computer that [the Navy] expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence." A decade later, Marvin Minsky and Seymour Papert observed (in Perceptrons, MIT, 1969) that such statements may be guilty of the wildest exaggeration. However, Minsky is no stranger to extending the euphoria, and feeding the frenzy with predictions about AI's rosy potential. Extending the scope of intelligence in AI has always been in vogue and continues unabated.

It is well nigh impossible to over emphasize that perceptrons, decision trees, recursion functions, etc, are good topological representations and manifestations of the anatomy, which are visible, in the context of neural connectivity, and the neural networks, which we observe, for example, in worms, eg *Caenorhabditis elegans* (Robert Horvitz, MIT; figure 3).

Abstractions of these networks, eg, cube-on-cube (The Society of Mind by Marvin Minsky, MIT, 1985) and "learning" mechanisms (The Organization of Behavior by Donald O Hebb, 1949) may be building blocks of intelligence. These processes populate fields with values from the user environment which can be selectively used (*per contra* hard coded defined sets). NEST Learning Thermostat uses input values, to tune *your preferred* temperatures.

Elements of the equation/rule-based (brittle, static) structures caused the bust of expert systems. It ended the lure of *The AI Business* (Winston and Prendergast, MIT, 1984) before the rise of artificial neural networks (ANN, popularity circa 1990). Topology and (dubious) synaptic weights, in stochastic models, are suggestive of flexible infrastructure. The intent is to improve profitability from data. But, is it really an application of intelligence in AI?

Rather than partial differential equations, exploding due to increase in state functions (due to the large number of parameters), use of Agents allow each variable to be represented as a single-function entity. Collective output from an Agency of Agents improves predictive or prescriptive precision, compared to operations research applications (see Figure 2). The behavior of Agents and Agencies using "AI" concepts originated from the principles of stigmergy (Pierre-Paul Grasse, 1959) which continues to evolve.

Think Data before Thinking of AI

The recent surge in the hype associated with "big data" has navigated profitability from analytics to the front and center. Intelligence is marketed as a commodity in this scenario.

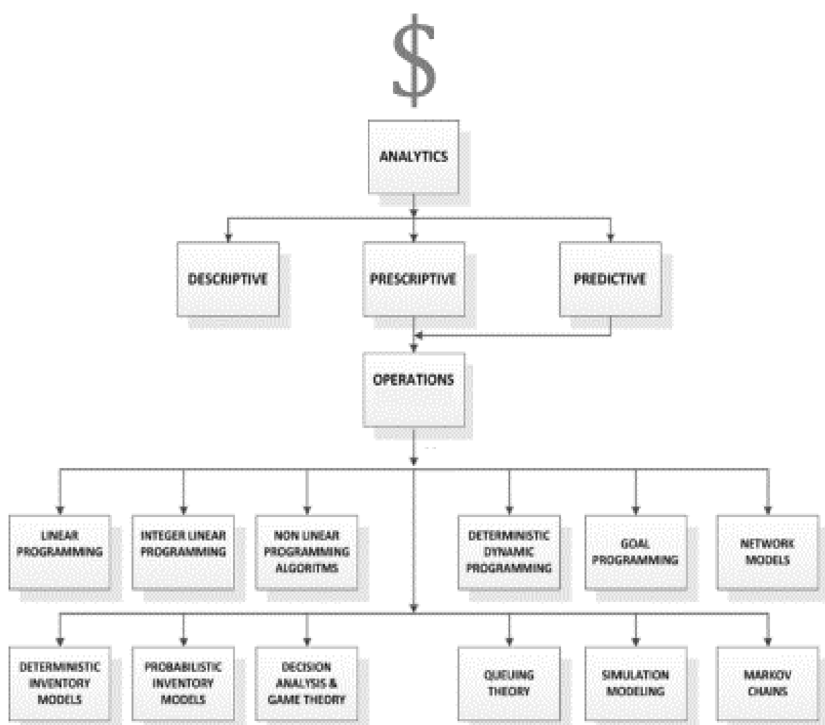

Figure 2: Pursuit of the Luminous Summit [$] Intelligence as a Service

In order to market intelligence as a service, the AI paradigm is being refurbished as a commodity (brain in a box), and touted to business and industry as an essential tool to reach the luminous summit [$]. Winning at games using ANN is advocated as intelligence. Smart and intelligence are emerging as speculative tabloid fodder. Witnessing the rapid transmutation of tabloid fodder from speculation to business truth is deeply troubling.

Claims about "original" thinking in containerization of data and processes are good ideas, but are concepts proposed almost half a century ago by Marvin Minsky (page 315 in the original book or search page 311 in this PDF version of the book). Connecting entities (containers) using IPv6 resonates with ideas suggested about a decade ago. However, it is reassuring that the concepts are not lost, but are being developed to advance the march of digital transformation (see Digital Twins https://dspace.mit.edu/handle/1721.1/104429).

Data, especially the volume of curated data, and the fundamental tools of knowledge representation (KR), are required for deploying AI and machine learning (ML) techniques. Data, AI, ML, are topics of intensive investigation, grand debates and energetic discourse, for at least half a century. However, the accelerating pace of rancor in this field, stems from marketing pundits and less informed celebrities, of a certain type.

Philosophical notions, linguistics and epistemology, in convergence with logic, ontology and computable models, fuels KR. Comprehending these principles, in combination with conventional mathematics, statistics, computer science and neurology, begins to form the foundation of AI analytics. To profit from AI applications, depth of domain knowledge must be added, to this foundation.

Those who can grasp this confluence, are converts, proselytizing about abstractions. The call for abstractions, which can hide the complexity of these principles, are essential for mass consumption of AI/ML applications. Intuitive, friendly, user interfaces (think Lego Mindstorms) which can be configured by the masses, lacking in specialized skills, are the "killer" apps, which will transform business, commerce, government, academia and *any instance, any event, any operation,* where decision may be necessary, before the next step.

Excruciatingly detailed level of minutiae are germane to KR frameworks. It remains quite inaccessible to the global population, except a few thousand computer science graduates. Minor changes in the sequence of symbols, in predicate logic (predicate calculus), can make a crucial difference in the meaning, as illustrated in the following example from John Sowa:

$$(\forall x)(\exists y)((\mathrm{woman}(x) \wedge \mathrm{dept}(x,\mathrm{MechE})) \supset (\mathrm{man}(y) \wedge \mathrm{hometown}(y,\mathrm{Boston}) \wedge \mathrm{married}(x,y))) \quad \text{[I]}$$
$$(\exists y)(\forall x)((\mathrm{woman}(x) \wedge \mathrm{dept}(x,\mathrm{MechE})) \supset (\mathrm{man}(y) \wedge \mathrm{hometown}(y,\mathrm{Boston}) \wedge \mathrm{married}(x,y))) \quad \text{[II]}$$
$$(\forall x)(\exists y)((\mathrm{man}(x) \wedge \mathrm{dept}(x,\mathrm{MechE})) \supset (\mathrm{woman}(y) \wedge \mathrm{hometown}(y,\mathrm{Boston}) \wedge \mathrm{married}(x,y))) \quad \text{[III]}$$
$$(\exists y)(\forall x)((\mathrm{man}(x) \wedge \mathrm{dept}(x,\mathrm{MechE})) \supset (\mathrm{woman}(y) \wedge \mathrm{hometown}(y,\mathrm{Boston}) \wedge \mathrm{married}(x,y))) \quad \text{[IV]}$$

[I] in English and French, respectively, may read as follows:
Every woman in department of Mech Engineering married a man who came from Boston
Chaque femme du département de génie mécanique a épousé un homme venu de Boston

[II] in English and French, respectively, may read as follows:
A man who came from Boston married every woman in department of Mech Engineering
Un homme venu de Boston a épousé toute femme dans le département de génie mécanique

[III] in English and French, respectively, may read as follows:
Every man in department of Mech Engineering married a woman who came from Boston
Tout homme dans le département de génie mécanique a épousé une femme venant de Boston

[IV] in English and French, respectively, may read as follows:
A woman who came from Boston married every man in department of Mech Engineering
Une femme venant de Boston a marié tous les hommes dans le département de génie mécanique

The difference between the meaning of [I] and [IV] leaves room for social disequilibrium. The tiny change which could ignite such an immense social cataclysm may be attributed to an almost imperceptible alteration of the order of the symbols between [I] and [IV]. The latter may be obvious, only to a handful of people, in the entire world. Hence, the clamor for abstraction.

Tools for the masses must sweep these intricacies under "big buttons" which one can drag and drop, to convey their intent, through rudimentary natural language contexts. Anyone can create programs, without the faintest idea of what is a program, or any knowledge of programming, whatsoever (Lego Mindstorms). When outcome, function, and output, can be optimized by the masses, the floodgates to profit from rational AI apps, may cease to close. It will be a part of the fabric of our daily lives (quote from Herbert Simon, CMU).

Before we reach that halcyon era, we must deal with data. Data must feed algorithms, some of which, may be based on principles of AI and ML. However, most gains arise from great features, not great ML algorithms. Often, a heuristic model can solve 80% of the problem even without ML. Feature engineering is a critical skill in data science. Incorrect features will lead to irrelevant or imperfect training for algorithms. Before using a model live or exporting a model, the performance of the model must be checked using data. Anomalies in performance is an indicator of problems either with data, features or the model or all of the above. Continuous checking of the AUC (area under ROC (receiver operating characteristic) curve) is a prudent approach to validate performance of models and their suitability.

Human-level AI ?

It is hard to improve the statement from Rodney Brooks, hence, here is a verbatim quote: "We have found a way to build fixed topology networks of our finite state machines which can perform learning, as an isolated subsystem, at levels comparable to these examples. At the moment, of course, we are in the very position we lambasted most AI workers for, earlier in this paper. We have an isolated module of a system working and the inputs and outputs have been left dangling." (Intelligence without Representation, 1991)

Learning triggers profound, sustained, often long term changes in our neural networks at many levels that we cannot even begin to understand, or grasp, its cognitive repercussions. Almost all assumptions made by McCulloch & Pitts (1943) may be violated (Appendix 1).

The "all or none" phenomena assumed by McCulloch & Pitts (1943) is relevant from a mechanical perspective if one assumes (perhaps incorrectly) that input data is supposed to transduce a signal, and the resultant action potential (neuronal activation), may be one form of a proof of learning. Neurologists may strenuously and vociferously take exception.

AI experts may wish adopt this view, about "learning" in the AI context. The neurological state of learning, cognition, and behavior is usually a continuous function, modulated by evolutionary weights, which are not subject, in the least, to the limitations of discrete-state machines. Application of machine learning models are often inconsistent and incorrect.

Discrete systems have a finite (countable) number of states, which may be described in precise mathematical models. The computer is a finite state machine which may be viewed as a discrete system. The brain is not a computer. The neural infrastructure and networks are not finite state machines. Imposing any such model (real-world continuous systems) or ill-advised abstraction or gross extrapolation (by those not so well informed) may only perpetrate great lengths of fantasy, about intelligence, and learning, related to AI systems.

"Of the vast stream of sense data that pour into our nervous systems, we are aware of few and we name still fewer. For it is the fact that even percepta are wordless. Only by necessity do we put a vocabulary to what we touch, see, taste, and smell, and to such sounds as we hear, that are not themselves words. We look at a landscape, a rich carving and majestic architecture of a cathedral, listen to the development of harmonies in a symphony, or admire special skill in games and find ourselves woefully lacking in ability to describe our percepts. Words, as we very rightly say, fail us, either to describe the plain facts of these experiences or to impart to others, our feelings." (G Jefferson CBE, FRS, MS, FRCS, Professor of Neurosurgery in *The Mind of Mechanical Man* in British Medical Journal, 25th June 1949). The author was aware of "Dr Wiener of Boston, his entertaining book *Cybernetics* (1948)."

Alan Turing was cognizant of the over-reach in claiming "intelligence" in AI and outlined potential objections including Godel's theorem (mathematical objection) and "Argument from Consciousness" which he reproduced from Professor Geoffrey Jefferson as a quote (from his Lister Oration, 1949) "Not until a machine can write a sonnet, or compose a concerto because of thoughts and emotions felt, and not by the chance fall of symbols, could we agree that machine equals brain. That is, not only write but know that it had written it. No mechanism could feel (and not merely artificially signal, an easy contrivance) pleasure at its successes, grief when its valves fuse, be warmed by flattery, be made miserable by its mistakes, be charmed by sex, be angry or depressed when it cannot get what it wants."
A. M. Turing (1950) Computing Machinery and Intelligence. Mind **49** 433-460 (PDF)

Page 452 (see Appendix 2) removes any doubt that Turing had grave doubts regarding claims of intelligence in the context of computers. Turing's suggested starting point is "the child machine" (Appendix 2). Then he proposes to add the roles or processes of "evolution" "hereditary material" "mutation" "education" and "natural selection" in order to mature "the child machine" to "imitate an adult human mind" as a path forward to intelligence. To understand even vaguely what happens after "the initial state of mind, say at birth" the reader is urged to review Patterns in the Mind by Ray Jackendoff (1966) and then take into consideration the field of linguistics and natural language development (1970, PhD thesis of Terry Winograd, MIT http://hci.stanford.edu/winograd/shrdlu/AITR-235.pdf).

For all this to happen, we must process information encoded via developmental and environmental signals. Hence, the suggestion, research and convergence *on* the concept of molecular logic gates. The complexity of the process may help deter one from concluding that we ***are*** dealing with intelligence, with respect to computers, machinery or AI systems.

However, the human spirit and the fabric of scientific research cannot step away from problems, even if all available reason suggests that something is impossible, at the time. It is with this fervor the 1956 Dartmouth Summer Research Project on Artificial Intelligence (June 17 - August 16) was proposed in 1955 by a visionary group of eminent and erudite academic scholars (www.aaai.org/ojs/index.php/aimagazine/article/view/1904/1802).

The proposal (see Appendix 3) admits it is a "conjecture" but continued "that every aspect of learning, or any other feature of intelligence, can, in principle, be so precisely described that a machine can be made to simulate it."

Great strides (Appendix 4) have been made, yet the 1956 Summer Research "conjecture" looms overhead. Progress of AI is evident from the 1145 page book by Russell and Norvig (AI - A Modern Approach, 3rd ed). We are also learning how decisions can be made without a brain in cognitive organisms (unicellular mould *Physarum polycephalum*).

Neurobiology 101

Topology and weights are the foundational underpinnings of artificial neural nets (ANN) which are the mainstay of AI systems. How reliable are these extrapolations? Are we still talking about AI? Let me reiterate what Rodney Brooks has stated but in a different vein.

Structural design of network topology aims to mimic the commonly observed *organization* of neurons. Topology based on neural organization (small world networks) may be fraught with errors, as evidenced by studies on wiring configuration and neuroanatomical analysis which reveals differences in circuit architecture and connectivity, if viewed at mesoscopic vs microscopic scale. On a mesoscopic scale, seemingly random networks exhibit consistent properties. It may be difficult, if not impossible, to extract useful/meaningful abstractions from these counter-intuitive non-linear yet dynamic structure-function complementarities.

In ANN, weights are assigned to signify connectivity strengths (the links between the perceptrons). These are arbitrary, at best, because synaptic weights between neurons and clusters are subjected to conditions not well understood. Even if one acquired neuro-physiological data related to frequency variations of action potentials (~200 Hertz), in an attempt to understand synaptic weights between neurons, the results may be inconclusive. The complexity may be compounded by the fact nerve transmissions are modified by ions, electrical threshold potential, chemical neuro-transmitters and their location in the brain.

In synaptic design, one assumes the all-or-none process (Appendix 1) and the weights are modeled based on extrapolation from "inferential changes" which are in the order of milliseconds to seconds (hence, subject to observation, data collection and extrapolation). But, the nature of the connectivity and resultant weight is also influenced by epigenetic factors (time scale – seconds to days), ontogenic factors (days to years) and phylogenetic factors, which are the result of generations or are derived from the evolutionary time scale, as noted in Appendix 2. Hence, the nature of the weight deduced from "inferential" changes (primarily sense and response mechanisms) are only the tip of the iceberg. We are almost completely in the dark about the nature of the influence from these and many other factors.

Synaptic weights ascribed in stochastic models are based on heuristics. This approach may not be sufficiently informed to account for the complex distribution of spike rates, observed synaptic weights and intrinsic excitability for neurons, in different areas of the brain, for example, auditory or visual cortex, hippocampus, cerebellum, striatum, midbrain nuclei. These processes underlie neuroplasticity and the potential for continuous configuration and reconfiguration of the brain through dynamic Hebbian Learning (Donald O. Hebb).

By [definition](), in Hebbian learning (one of the oldest learning algorithms), a synapse between two neurons is strengthened when the neurons on either side of the synapse (input and output) have highly correlated outputs. Following the analogy to an artificial system (ANN), the tap weight is increased with high correlation between two sequential neurons. Capturing this dynamic change in weights and including this change appropriately in the function of the algorithm, requires feeding the algorithm. In other words, getting the correct data with the correct features (discrete features and/or crosses) to the algorithm.

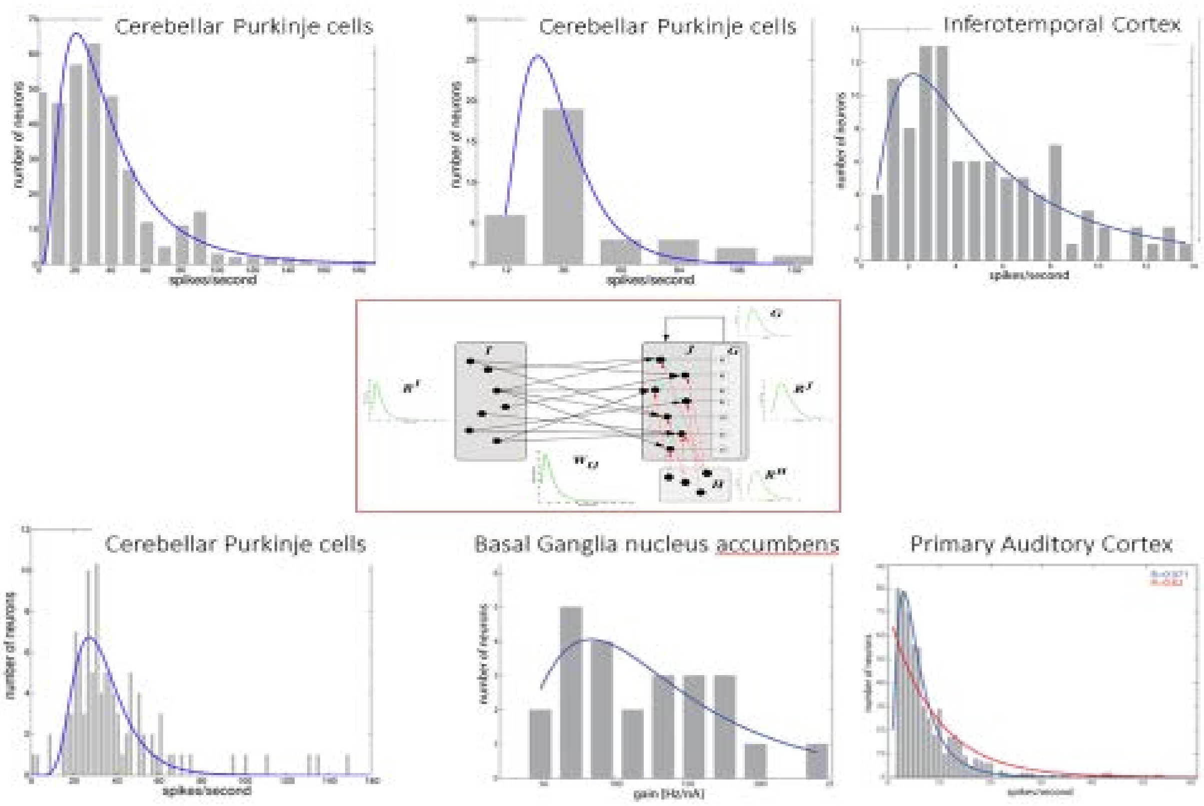

Figure 3: Distribution of neural spikes and intrinsic excitability from various parts of the brain indicate a range of responses. A representation of artificial neural network (center). Scheler, Gabriele (2017) Logarithmic distributions prove that intrinsic learning is Hebbian. *F1000Research* **6** 1222 (https://f1000research.com/articles/6-1222/v1)

Heuristic weights used in ANN (Figure 3, center) may not resemble reality, with respect to observed pattern of tap weights, in biological neural nets. Evidence of this discrepancy is based on observations of variability (Figure 3), central to biological processes, learning and neuroplasticity. This information may be useful when training artificial neural nets (ANN).

Untrained memory network is a structural representation (Figure 4, left panel) of a minuscule segment of the 100 trillion synapses in the (human) brain. When humans or animals learn certain behaviors, the synaptic relationship changes. A class of the latter type may be referred to as Hebbian learning (a synapse between two neurons is strengthened when the neurons on either side of the synapse, input and output, have highly correlated outputs). But, learning is not an unimodal action. Animals and humans are multi-sensory, multi-modal and multi-dimensional learners. Hence, an event cannot be mapped to change in one synapse. A simple learning, may impact a neural network, no matter how small.

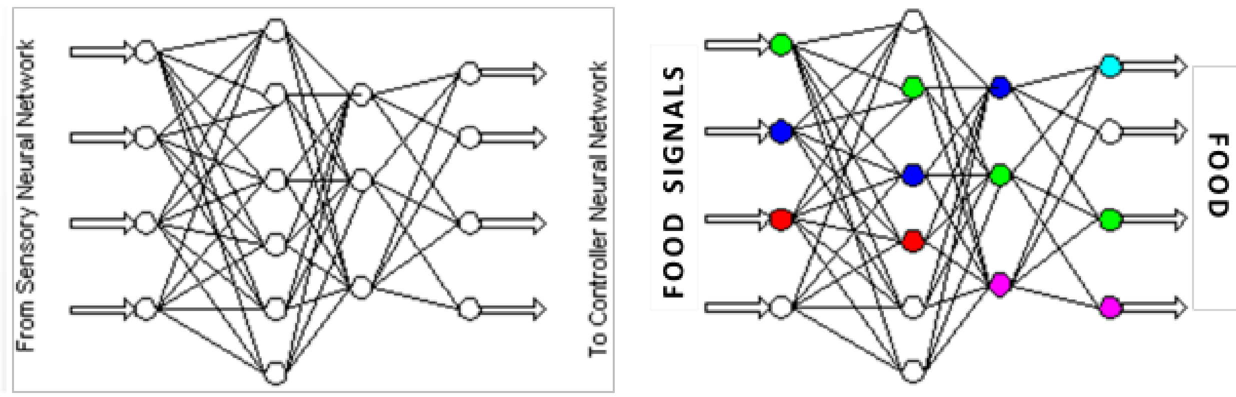

Figure 4: Untrained memory network (left panel). Trained memory network (right panel)

Re-visiting the classic example of Ivan Pavlov's dog (Spot) the "food" signal is a complex multi-sensory input. For simplicity we dissect it as follows: [1] auditory signal (Spot, Food!), [2] olfactory signal (smell of food) and [3] optical signal (sight of food). Assuming that Spot did not understand Russian or English, using the Hebbian learning principle, calling out "Spot, Food!" is just a sound, albeit emphatic, but still a sound, but it follows a specific form, a pattern, which repeats (reinforcement learning). In course of these experiments, Ivan Pavlov's dogs may have "learned" to differentiate between arbitrary background chatter and the call "Spot, Food!" by their owner. Because of the natural language ambiguity, the blue input (Figure 4, right panel) manifests as a turquoise output. The choice of the closely related colors may indicate reduced synaptic weight, and perhaps insinuate a different strength of learning in this memory circuit. However, the auditory signal (Spot, Food!), in ***combination***, with other signals, sends the correct message (sounds like food) to the neural controller. The combinatorial and synergistic effects, are weighted differently, and are important determinants in memory circuits.

The blue input signal synapses with the green input signal, as if to validate the incoming olfactory signal (green cluster), since the sense of smell is a signature stimuli, for dogs. The output from the olfactory signal (input, green) is definitive (output, green) and strengthens the output message (food). The sight of food (input, red) adds to the combined output but it may be weak (output, pink, related color to red, input) because ***humans*** are unclear if dogs are capable of ***visually*** differentiating between chicken tikka masala v beef stroganoff.

> *I digress to make the point about "humans" understanding or not understanding. Logic, ontology and programming of computable models, is colored by the syntax and semantics of humans, their natural languages and their habitats. The same problem approached by an English, Chinese and Indian person, may produce three, related, but different outputs. This underscores the critical importance of semantics and interoperability between standards and interoperability between platforms. Hence, the role of interfaces is key to abstraction.*

Returning to neural nets and weights, the point of this explanation is to convey that neural networks, naturally, lead to memory circuits. In this case, the dog's memory circuit gets trained, over time (hence, these are linked to time series) to sight, sound and smell. The training creates patterns, corresponding to the input signals feeding the network. The dynamic nature of this training is represented by colors of the neural cell clusters, with colors representing strength of connections.

The complexity of this very simple event illustrates how learning, synaptic weights and modifiers, influences memory. In this case, you can execute associative control logic by having an integrator/arbitrator neural cluster (sight and sound). The degree to which we can abstract these principles and create ANNs closely resembling natural neural nets remains an open question.

But, it should not discourage humans from trying to use bio-inspired models in developing and using ANNs to support "intelligent" decision support systems. No matter how intricate and multifactorial natural neural nets may be, humans are persevering in their attempt to create ANNs and use it as effectively, as possible, even if "intelligence" remains a mirage.

For example, you can plug in the memory network for the stimulus represented by green neural cell cluster (Fig 4) and track the memory of an event, executing the associative logic.

Similarly, we can integrate associative control clusters and create a complex of associative control logic. In Figure 5 (right), two associative control circuits are integrated. Each circuit feeds a gray memory circuit above it. The two memory circuits tie in to a control circuit (black node) that feeds into a memory circuit at the apex (gray node). This topology can execute a complex set of behavioral controls in response to a complex set of multi-source stimuli. It has an internal memory of the actions taken in response to the situation. The latter is reminiscent of human action, that is, the ability to remember actions in a given set of circumstances. The ability of humans to show restraint and the ability to "forget" are important elements which may be driven by higher level cognition. Pragmatic applications of AI are keen to capture part of this "human problem solving and processing ability" due to neural networks. It is certainly worthy of mimicry.

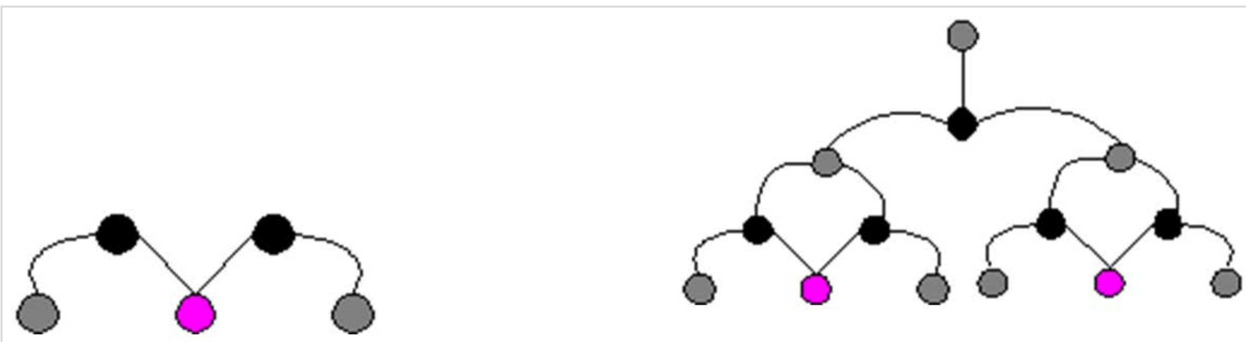

Figure 5: Simple associative control (L) and association complex with internal memory (R)

Taken together, perhaps we shall invest more thought in the design of ANNs by re-visiting assumptions which are generally incorrect because we are still significantly uninformed. Having said that, one must hasten to add, that, no matter how approximately correct the synthetic weights, may be, it may not be impossible to conceive building better ANNs with partially unsure topologies and unqualified numbers (Figure 3) using the Pareto Principle. Perhaps we can achieve 80% of the goal, using the approximations, at 20% of the cost.

Clever integration of tools and techniques (for example, Luenberger Observer, back propagation algorithms) may allow us to tune and re-tune AI systems in a dynamic data-driven manner and improve the learning, as discussed above, to generate actionable information. Over-fitting the model may cause harm, for example, in collision avoidance. Understanding the principles that govern these processes (neurobiology) may improve the thinking of analysts to use better design principles while constructing dynamic models.

Scholars continue to discuss new ways of using robots to make robots, create self-healing intelligent machines and adaptive machines to optimize up-time. Thinkers are conjuring up ways to harness the developmental foundations of neurons – *neurogenesis*. Emulation of *neural development* using computational AI systems can incorporate characteristics of natural neural systems into engineering design. Scientists are claiming that rather than *designing* neural networks, emulation of neurogenesis shall enable us to *generate* neural networks to serve dynamic and even more complex systems of the future. This emerging field of programmable artificial neurogenesis appears to call for a meta-design paradigm which may begin with components (objects?). It aims to build higher order intelligent systems which will adapt (to demands, environment, resources) without re-programming component level entities. When components are updated, the changes will be propagated, via appropriate "learning" functions, up/down hierarchies.

The great desire to emulate the grand vision latent in intelligence, cognition and the brain, works almost as an aphrodisiac. The immense powers of biology, and the ability to distil and capture even an iota of that potential, from bio-inspired systems, through convergence with computation, will continue to be a Holy Grail. Here is one example. We have about 3 billion base pairs (A-T, G-C) in the human genome ($3x10^9$) which codes for about 10,000 – 20,000 genes resulting in a human body with 100 trillion cells ($1x10^{14}$). At least a third of the approximately 20,000 different genes that make up the human genome are active (expressed) in the brain. We have about $8.5x10^{10}$ neural cells (there are an equivalent number of glial cells). Each neural cell connects on an average with 1,000 other neural cells to create about $1x10^{14}$ neural connections (100 trillion synapses). This is the natural neural network which makes us human, brews intelligence and cognition. In terms of compression ratio, the ratio approaches 10^{11} (7,000 genes creating $1x10^{14}$ connections). From the non-biological world, the most effective compression algorithm CMIX doesn't even come close. The illegal 42.zip bomb which unfolds to 4.5 petabytes (pb) from a 42 kilobytes (kb) single symbol zip, approaches a compression ratio of 10^{11} in an artificial circumstance devoid of intelligence. The human compression of 10^{11} offers sustainable, real, life-long intelligence.

Conclusion – These 'intelligent' machines may never be intelligent in a human sense (p339)

A quantum leap, still cryptic within the unknown unknowns, may unleash what is intelligence within AI, in the future. We must continue to explore far and wide, emulate insects and think about the Octopus. "We can only see a short distance ahead, but we can see plenty there that needs to be done." Convergence of tools (statistics, math) with data curation (noise vs signal) is replete with promise and profitability even if it lacks (human-level) intelligence. We must continue to explore ways to profit from AI (Appendix 6).

APPENDIX – 1

Warren S McCulloch and Walter H Pitts (1943) *A Logical Calculus of the Ideas Immanent in Nervous Activity*. Bulletin of Mathematical Biophysics **5** 115-133
http://www.cse.chalmers.se/~coquand/AUTOMATA/mcp.pdf

A Logical Calculus of Ideas Immanent in Nervous Activity

adequate. But for nets undergoing both alterations, we can substitute equivalent fictitious nets composed of neurons whose connections and thresholds are unaltered. But one point must be made clear: neither of us conceives the formal equivalence to be a factual explanation. *Per contra!*—we regard facilitation and extinction as dependent upon continuous changes in threshold related to electrical and chemical variables, such as after-potentials and ionic concentrations; and learning as an enduring change which can survive sleep, anaesthesia, convulsions and coma. The importance of the formal equivalence lies in this: that the alterations actually underlying facilitation, extinction and learning in no way affect the conclusions which follow from the formal treatment of the activity of nervous nets, and the relations of the corresponding propositions remain those of the logic of propositions.

The nervous system contains many circular paths, whose activity so regenerates the excitation of any participant neuron that reference to time past becomes indefinite, although it still implies that afferent activity has realized one of a certain class of configurations over time. Precise specification of these implications by means of recursive functions, and determination of those that can be embodied in the activity of nervous nets, completes the theory.

THE THEORY: NETS WITHOUT CIRCLES

We shall make the following physical assumptions for our calculus.

1. The activity of the neuron is an "all-or-none" process.
2. A certain fixed number of synapses must be excited within the period of latent addition in order to excite a neuron at any time, and this number is independent of previous activity and position on the neuron.
3. The only significant delay within the nervous system is synaptic delay.
4. The activity of any inhibitory synapse absolutely prevents excitation of the neuron at that time.
5. The structure of the net does not change with time.

APPENDIX – 2

[Page 452](#)

A. M. Turing (1950) Computing Machinery and Intelligence. Mind **49** 433-460
www.csee.umbc.edu/courses/471/papers/turing.pdf

In the process of trying to imitate an adult human mind we are bound to think a good deal about the process which has brought it to the state that it is in. We may notice three components.

(a) The initial state of the mind, say at birth,

(b) The education to which it has been subjected,

(c) Other experience, not to be described as education, to which it has been subjected.

Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's? If this were then subjected to an appropriate course of education one would obtain the adult brain. Presumably the child brain is something like a notebook as one buys it from the stationer's. Rather little mechanism, and lots of blank sheets. (Mechanism and writing are from our point of view almost synonymous.) Our hope is that there is so little mechanism in the child brain that something like it can be easily programmed. The amount of work in the education we can assume, as a first approximation, to be much the same as for the human child.

We have thus divided our problem into two parts. The child programme and the education process. These two remain very closely connected. We cannot expect to find a good child machine at the first attempt. One must experiment with teaching one such machine and see how well it learns. One can then try another and see if it is better or worse. There is an obvious connection between this process and evolution, by the identifications

Structure of the child machine = hereditary material

Changes of the child machine = mutation,

Natural selection = judgment of the experimenter

APPENDIX – 3

1955 Dartmouth Summer Research Proposal
http://www.aaai.org/ojs/index.php/aimagazine/article/view/1904/1802

A Proposal for the

DARTMOUTH SUMMER RESEARCH PROJECT ON ARTIFICIAL INTELLIGENCE

June 17 - Aug. 16

We propose that a 2 month, 10 man study of artificial intelligence be carried out during the summer of 1956 at Dartmouth College in Hanover, New Hampshire. The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it. An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. We think that a significant advance can be made in one or more of these problems if a carefully selected group of scientists work on it together for a summer.

The following are some aspects of the artificial intelligence problem:

1) Automatic Computers

If a machine can do a job, then an automatic calculator can be programmed to simulate the machine. The speeds and memory capacities of present computers may be insufficient to simulate many of the higher functions of the human brain, but the major obstacle is not lack of machine capacity, but our inability to write programs taking full advantage of what we have.

2) How Can a Computer be Programmed to Use a Language

It may be speculated that a large part of human thought consists of manipulating words according to rules of reasoning

Photo courtesy Dartmouth College.

Page 1 of the Original Proposal.

APPENDIX – 4

1943 Warren S McCulloch and Walter H Pitts *A Logical Calculus of the Ideas Immanent in Nervous Activity* in Bulletin of Mathematical Biophysics **5** 115-133
http://www.cse.chalmers.se/~coquand/AUTOMATA/mcp.pdf

1949 Geoffrey Jefferson *The Mind of Mechanical Man* in British Medical Journal June 25
http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2050428/pdf/brmedj03683-0003.pdf

1950 Alan M. Turing *Computing Machinery and Intelligence* in Mind 49 433-460
www.csee.umbc.edu/courses/471/papers/turing.pdf

1955 Proposal for Dartmouth Summer Research Project on Artificial Intelligence
http://www.aaai.org/ojs/index.php/aimagazine/article/view/1904/1802

1960 Marvin Minsky *Steps Toward Artificial Intelligence* in Proceedings of the IRE
http://web.media.mit.edu/~minsky/papers/steps.html
http://www.cs.utexas.edu/~jsinapov/teaching/cs378/readings/W2/Minsky60steps.pdf

1966 Ray Jackendoff *Patterns in the Mind*
http://ling.umd.edu/~omaki/teaching/Ling240_Summ2007/Jackendoff94_Ch1.pdf

1970 Natural Language – PhD Thesis of Terry Winograd, MIT
http://hci.stanford.edu/winograd/shrdlu/AITR-235.pdf

1974 Marvin Minsky *A Framework for Representing Knowledge* in MIT AI Lab Memo 306
www.u-picardie.fr/~furst/docs/Minsky_Frames_1974.pdf

1984 Patrick H Winston *Artificial Intelligence: A Perspective*
https://mitpress.mit.edu/sites/default/files/titles/content/9780262570770_sch_0001.pdf

1984 Winston & Prendergast *The AI Business* https://mitpress.mit.edu/books/ai-business

1986 Kay Rodgers – US Library of Congress http://files.eric.ed.gov/fulltext/ED271107.pdf

1992 P Winston *Artificial Intelligence* 3rd http://courses.csail.mit.edu/6.034f/ai3/rest.pdf

2015 Stuart J Russell and Peter Norvig *Artificial Intelligence - A Modern Approach* (3rd ed)
www.amazon.com/Artificial-Intelligence-Approach-Russell-2015-01-01/dp/B0182Q542W

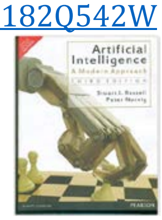

APPENDIX – 5 – What business needs for sustainable long term profitability

As an AI optimist, one cannot leave "AI" hanging without some form of a path forward to deploy the treasure of principles. The discussion about lack of "intelligence" in AI should not indicate the author's lack of support for the *ideas that are represented* through AI. Informed individuals are likely to agree that human level intelligence is still utopian. A cursory reading of Chapter 1 of Russell and Norvig's seminal work (2015) will enlighten even a pre-teen student that the flight of an air plane is not *artificial flight* which mimics birds and *artificial intelligence* or AI is not similar to human intelligence. The introduction of the term "AI" by John McCarthy during the summer of 1956 was a genius stroke of public relations and marketing coup completely unbeknownst to John McCarthy or others in 1956.

Russell & Norvig (2015) *Artificial Intelligence: A Modern Approach* (3rd ed). Page 17
Perhaps "computational rationality" would have been more precise and less threatening, but "AI' has stuck. At the 50th anniversary of the 1956 Dartmouth conference (in 2006), McCarthy stated that he resisted the terms "computer" or "computational" in deference to Norbert Weiner, who was promoting analog cybernetic devices (in the 1950's) rather than digital computers. (Section 1.3.2)

Definitions are often useless and narrow but if one must re-define "AI" then consider this: The science and art of designing engineering tools, systems and platforms to perform tasks that generally requires human input, as a gradient of responses (segmented by simplicity or complexity) based on rational, contextual, and optimal attributes of human reasoning.

Those seeking to distill the "value proposition" of AI may consider this verbose statement: The value of AI is in performance of tasks, using one or more tools or systems, repeatedly and accurately, often, at a very high frequency, often, with immense sets of structured or even unstructured data, in various possible combinations of the features, of the data, using a dynamic selection of algorithms and solvers, from the system or platforms or several other sources (cloud, fog), in near real-time, or with negligible latency which the outcome can tolerate. The delivery of the value of AI is inextricably linked to the ***strategic*** value of data and/in connected networks. The value of AI analytics may be reaped at the edge or core or in sub-systems, if the outcome or output can effectively aid decision support systems, within bounded latency, to improve performance and/or actuate/execute events, without further human input, in near real time, or at the correct time, to optimize the scenario or improve profitability from the decision or outcome based on AI data analytics.

Organizations and corporations are perhaps most interested in the AI economy. Hence, the application and deployment of AI tool kits. To achieve these goals, at least in part, diffusion of the principles and practice of AI tools must precede the P&L quest for profitability. The most critical catalyst for breaking new ground is *information arbitrage*. AI is no exception.

The industry needs skilled communicators with content knowledge, those who can simplify but refrain from being simplistic, explain without dilution or downsizing the material and reach the executive core as well as practitioners without making it a journey of "billable hours" consulting. The communicators must also understand strategy and management.

Strategic communication of the science and engineering, related to the principles and practice of AI, requires a synthesis of logic, ontology and computation. Taken together, the trinity, is the foundation of knowledge representation (KR). The abstraction of KR may be industry agnostic for certain classes and industry specific in many instances. Widespread understanding of these elements and the participation of the industry specific knowledge experts in their development, is the Holy Grail, which is essential for systems to benefit and profit from AI. The *strategic* value of this modus operandi must be in balance rather than serving AI as a one shoe fits all or a panacea for all the woes of decision support systems.

To be driven by data, AI tools cannot exclude the principles of data handling, feature engineering, data curation and how to use dynamic high volume data to extract *actionable information*. These are not apps which teenagers can build in a few hours. A most sincere concerted effort is a pre-requisite. This is a field where certain tools and standards exist. Inclusion, interoperability between standards and integration may be essential, in certain domains. The ability to cross-pollinate existing and available tools from different sectors with new, novel or unique capabilities, requires a dynamic platform, agile in its ability to adopt and adapt. These attributes are missing in certain vendors who promote their "all-in-one" package or install cryptic boundaries (example: closed data dictionaries or semantic incompatibilities between knowledge representation modules). In a volatile economy, any restriction may become a detriment to global businesses and an anathema for profitability.

These roadblocks are nothing new and are usually organizational. The general disregard for experts and academic-industry collaborations, fuels the "illusion of expertise" which often retards innovation. Evolution of digital transformation can easily fall prey to market forces which manufactures these veneers or silos. One role of the visionary leader or the entrepreneurial corporation is to recognize the mirage created by illusions of completeness and prevent the delusion from asphyxiating the transformation of vision into reality.

Businesses need intuition and imagination to inculcate vision. The leadership may invest in building the capacity to execute in a manner that catalyzes entrepreneurial innovation. The path to AI driven data analytics is calling out for leaders. Also, AI calls for confluence, even more than other broad areas. Context of AI applications are pivotal in order to reduce the *garbage-in/garbage-out* syndrome. Cloud, cybersecurity, computation and communication must converge for the data to be available, secure, analyzed and used *before* the value of the information perishes. These and other factors are critical for success of AI applications.

For businesses to stay in business it must look beyond business, for networks of innovation which will depend on the supply chain of global talent. The excitement about new tools and technologies may die a premature death in a manner similar to the hypothetical demise of the industrial revolution, if the world ran out of coal. Data is not the new coal, or new oil or new goal. Talent and education was and still remains the coal to ignite our imagination. The wealth of nations must be invested to cultivate talent through education. The US ought to be deeply concerned by the diminishing number of STEM graduates from its institutions.

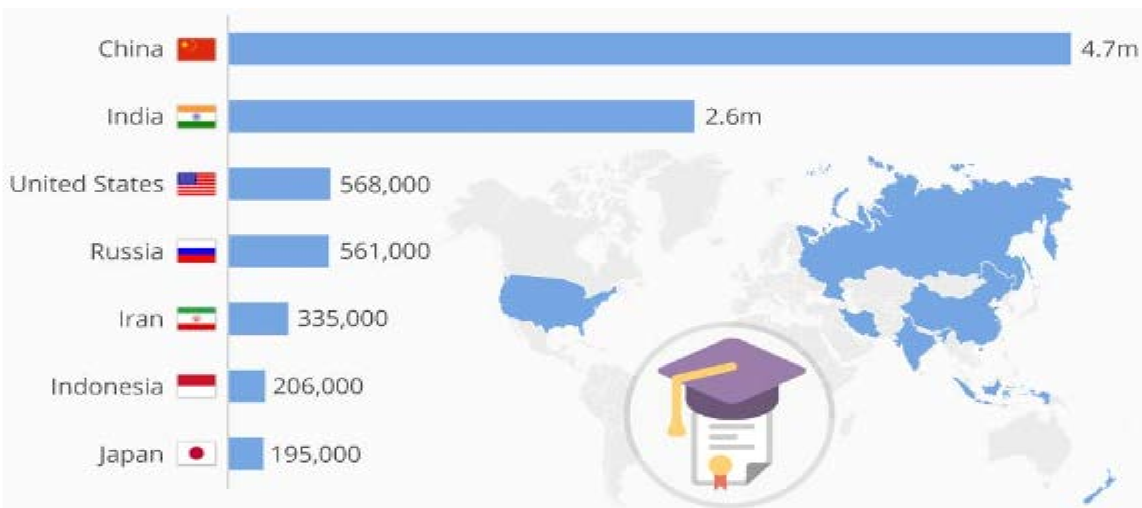

According to the World Economic Forum, US produced 568,000 STEM graduates in 2016. If 1% of these students (5,680) are in the general area of AI, and if 1% of that number (568) are proficient in communicating about AI, and if 1% of the strategic experts may choose to work for your business, then the human capital you may have to promote AI education, strategy and usability in your business, amounts to less than 6 employees, if you are lucky! The supply chain of talent is the ultimate rate limiting step. Education needs an avenue to amplify its reach, increase the pool of talent, globally. The new path is digital, it is digital learning. For corporations with global ambitions, sponsorship of education, digital learning and academic-industry partnerships, are sign posts on the road, to long term profitability.

APPENDIX 6 – Prescription: Profit from AI

Please see attached op-ed article (Letter), thanks!

↓

JIM 5, 2 (2017) 3-7
HANDLE: http://hdl.handle.net/10216/105431

Prescription for Profit

Dr. Shoumen Palit Austin Datta

Senior Member, MIT Auto-ID Labs
Research Affiliate, Department of Mechanical Engineering
Massachusetts Institute of Technology, Room 35-203, Cambridge, Massachusetts 02139, USA
Senior Scientist, Medical Interoperability Program, MDPnP Laboratory
Massachusetts General Hospital, Harvard Medical School
Partners MGH Research Building, 65 Landsdowne Street, Cambridge, MA 02139, USA
`shoumen@mit.edu, sdatta8@mgh.harvard.edu`

Letter from Academia

AI is being marketed as a panacea solution or a complex "black box" by the PR spin doctors. How to profit from AI in business applications is still unclear. The lack of understanding of knowledge representation, data structures and feature engineering, are a few of the core underlying problems, devoid of easy solutions. This short guide is a note on strategy with respect to the use of AI tool kits. What is necessary for rational use and integration of AI tools with business are humans.

Keywords. AI, Digital Twins, Rational Agents, Logic, Ontology, Knowledge Representation, Data, Analytics, Computational Rationality, Neural Networks, Cognition, Machine Learning, ERP, Syntax, Semantics, Synapse, Speech, Vision, Neurobiology, Innovation, Management.

1 Pragmatic use of Artificial Intelligence

Pragmatic use of Artificial Intelligence (AI), which can catalyze corporations to profit from applications of AI, is the ultimate goal for business and industry. Academia could help industry achieve this goal, albeit in part, the rational part.

The probability of bursting the public relations bubble, and the hype about the promises made on behalf of AI, are increasing. It may be reduced if industry and management better understood why the world must opt to lower its expectations of "intelligence" as an outcome from AI tools.

The suggestions in this letter does not distract from the rational possibilities of using the principles of AI in data analytics, decision support, and even, in automation.

During 1955-56, the term "AI" rather than "computational rationality" was used to describe a "new" and "emerging" field. The difference between the terms is a matter of states because "intelligence" is continuous (core attribute of many biological processes) while "computational rationality" is a discrete process. The latter is explained by the boundaries of limited rationality which systems generate, based on computable models.

The cognitive glue, necessary to bond discrete events to form the continuum, may be a cherished objective, but remains an illusion for science, and delusion for engineers, at this time. This brings to mind the pithy words of John Searle "brains cause minds" as

ISSN 2183-0606
http://www.open-jim.org

if to say that a mere collection of cells (neurons, glia) will lead to thought, action and consciousness. It is true, the brain is a collection of these cells, but does that suffice to serve as a platform, to extrapolate the brain to the scope of the human mind?

To illustrate the issue on page 7 (see https://dspace.mit.edu/handle/1721.1/108000), I refer to (https://link.springer.com/content/pdf/10.1007%2Fs11633-017-1093-8.pdf) a recent paper. It starts with states (each binary bit has two states, 0 and 1) and memory storage capacity. A human brain has $1x10^{14}$ neural cells (100 trillion synapses) with approximately $2x10^{15}$ states, equivalent to a storage capacity of 500 terabytes (assume 4.6 bits of information stored by each synapse). Hence, $2x10^{15}$ is the number of synaptic connections in a human brain. Is this, then, the capacity of the human mind?

In our mind, cognition allows us to read, write, create, and understand language, simple and complex. Our vision can distinguish topology of objects, colors, size, depth, shades. Our five senses, in combination, can respond to an array of input, to produce a vast (unknown) number, and variety, of output. If artificial neural nets claim to "capture" the brain and if we can scientifically describe this *capture* as "brain in a box" then, this network, if etched on a "neuromorphic chip" is the sum total of "intelligence" that we may rely on, for all our activities. By this rationale, "intelligence" is governed by the maximum number of states of our synapses. That number is about $2x10^{15}$ and that is, by this account, the total number of instances or the magnitude of combinations of our thoughts. Is this a true statement?

As Danko Nikolic points out, an English speaker's vocabulary has about 15,000 words which consists of 5% adverbs, 20% adjectives, 20% verbs and 55% nouns (750, 3000, 3000, 8250 words in each of the four categories, respectively). From those numbers, we can calculate the number of all combinations, of sentences, of different lengths. For four word sentences, consisting of a noun, followed by a verb and ending with a noun plus an adjective, we obtain 8250×3000×8250×3000 or about $2x10^{18}$ combinations. We have not even pondered about the semantic boundaries of the syntax in the four-word sentences. This number ($2x10^{18}$) is already bigger than the limit that is posed by the total number of synapses in the brain ($2x10^{15}$). By this reasoning, there isn't enough memory in our brains to generate a different response even for sentences with 4 words!

At this stage, we have only considered "speech" and limited our expression to 4-word sentences. Limited by the storage capacity of the total number of synapses in our brain ($2x10^{15}$) we will not be able to see, hear, taste or touch, among other things. Do we still wish to continue, and support marketing campaigns, suggesting that deep neural nets are equivalent to biological intelligence, which powers AI? Hence, is there intelligence in AI?

The fact that humans possess at least five senses, and do much more than what $2x10^{15}$ synapses may allow, is due to the fact that this number is an anatomical representation of the number of discrete connections. This is the ***structure*** of the organizational aspect of the anatomy and topology of the human brain. Structure is ***not*** the same as ***function*** in the same manner that anatomy (human skeleton) is not equivalent to physiology and physiological function (human organism), even though the skeleton (structure) is quintessential for physiology (function).

It is ***function*** that generates the amorphous quality of intelligence and makes humans intelligent. The numbers in the structure are discrete. The numbers matter, of course.

With 302 neurons, potential structural relationships in *Caenorhabditis elegans* (worms) may not qualify to provide intelligent functions or even pattern recognition.

The function of intelligence is best perceived as a *continuous* fabric, inextricably linked with data, rules, patterns, experiences, knowledge and learnings to inform or support decisions.

The almost unlimited number of connected continuity, the underpinning of intelligent human action, is a result of $2x10^{15}$ synapses which are being formed, and re-formed, connected and disconnected, re-connected and re-configured, in an asynchronous, dynamic manner, in response to signals, perceived, received, in processing or being transmitted. Signals may originate from diverse sources (internal, external, autocrine, endocrine) or may be presented to sensory interfaces in a multitude of shapes or forms.

Continuity is not an attribute of a computable model. The term AI was less appropriate than "computational rationality" in 1956 and it is even less appropriate, today. But, we may use the term AI, for the sake of posterity, its magnetic image and public imprint.

The term intelligence is supposed to present a mental image relating evolution of words, objects, ideas, in terms of meaning and context. It is not a discrete, structural, one to one syntax, which can be translated. It is an *interpretation*, based on semantics, and by extension, logic, and ontology. The fact that intelligence may not be amenable to simple syntactic translation was demonstrated by the almost abject failure of the Russian to English translation during 1960s (prior to use of convolutional or recursive neural nets).

The "artificial" architecture of intelligence may have literary roots. Perhaps, a reference to *Leviathan* by Thomas Hobbes (1651) or similar, from that school of thought. Hobbes argued for "artificial animal" based on observation that the heart is a spring, nerves are strings and joints are wheels. Attempts to mimic birds and develop "artificial flight" did not lead to aviation. The right approach by the Wright Brothers was to view flight as a function of aerodynamics, which gave birth to the airline industry. Reality of flying, for human use, was not a reproduction of the fantasy of viewing birds in flight.

2 From Taylorisms to Terabytes

The movement from Taylorisms to terabytes needs AI, and its tools. Hence, AI, despite its limitations and a handicapped terminology, presents opportunities for companies to automate business processes. But, fantasy driven scenarios, about winning at GO or poker, may not suffice for integrating AI or ML applications, in the real world. While ERP implementations enhanced competitiveness, several companies also uncovered nightmares. The promised opportunity from ERP never came to fruition, for some. Do we have a sense of *déjà vu* with AI? The rain on the AI parade falls mostly on input data and the output/outcome. Unless reliably automated, the outcome requires people to do something with the information. Is it actionable? AI analytics cannot help if input data is noisy or corrupt. How do you know the data or the outcome is of poor quality?

AI and ML can augment performance. In case of AI (more than ERP) those changes create highly skilled tasks which require education, prudence and domain expertise, *from humans*. Businesses are forever in an elusive quest for "low hanging fruit" without gaining the wisdom from repeated failures. The pursuit of "low hanging fruits" require only low level skills. That *modus operandi* may not help, at all, to profit from AI.

Generating value from AI by recruiting more data scientists is an amorphous escape clause. Several domains converge under the umbrella of data science, which makes it impossible to ascribe the term data scientist, to any one individual. Data science is a team sport. Bringing the talent together, and synthesizing the unpacked problems, are tasks that few companies can execute because companies do not have, or rarely employ, strategic *cube-on-cube* thinkers.

Companies do not even know, that they do not know, that they lack trans-disciplinary cross-pollinators. Companies and HR are unable to comprehend that they need people with broad spectrum of knowledge "cubes" and a matrix of experiences, unlike those that can fit in a box. "Thinking different" is not a principle that HR departments can practice. Hence, the clamor for data scientists but lack of jobs describing the need for out-of-the-box thinkers, followed by an absence of zeal, to pursue the road not taken.

Thinkers are pivotal to assist teams to dissect problems into components, to identify the confluence of domains, and underpinnings of potential solutions. Creative thinkers are key to assist the leaders to move the fulcrum and mentor the rank and file to frame the correct questions. Hiring and allowing *cube-on-cube* thinkers to form agile, case-dependent teams, staffed with vertical experts, across silos (network of business units), may be the first step to profitability, from advanced applications, which are fueled by convergence, such as, AI, analytics, robotics and nanotechnology.

Data science must start with data. Data must be acquired, processed and curated to serve the business needs. Hence, the critical demand for domain experts, and field knowledge providers, who must help identify the obvious, common, and uncommon "features" that businesses are seeking. Then, add non-obvious relationship analyses, and garnish with unconventional wisdom. To harvest the latter, perhaps crowd sourcing may be useful.

Organized data, using the principles of knowledge representation and application of logic and ontology, is a starting point, to construct computable models/structures of the domains of interest (agnostic of industry, vertical or horizontal). In the computational phase, we can use algorithms and tools from AI including ML, DL, ANN, CNN, RNN.

The trinity of out-of-the-box thinkers, who can connect the cubes, with field knowledge providers, and computational experts, is the "secret sauce" which must be continuously stirred, shaken, configured and re-configured, to blend the correct team, case by case, to profit from AI, and use the ability of AI, in problem solving. This approach and grasp of the extended fabric, is lacking in businesses and absent within corporate leadership.

The marketing hype, which is furiously polishing the chrome, on the AI engine, may help to explode the bubble and trigger a second AI winter. Global warming will be essential to thaw the AI ice age. But, before we boil the ocean, let us try to warm up to what may be necessary, the prerequisites, what is missing, how deep is the abyss and how education may bridge the chasm. Let us imagine, we have managed to fast forward to the spring of AI. Assume, AI in the tool kit is generating probabilistic output.

As pointed out by Jeanne Ross, an AI application indicates that a lead has a 95% chance of converting into a sale, while another has a 60% chance. Should we assume the salesperson knows what to do with that information?

ML applications may help lawyers identify appropriate legal precedents, help vendor management teams ensure compliance with contracts, assist financial institutions to gauge risk. These systems use ML to perform mundane tasks. Systems can learn to develop spreadsheets and search databases for relevant information. But, in order to generate competitive advantage from ML (AI), we may need skilled humans to process the outcome. Hence, companies must redesign accountabilities, motivate employees to deploy ML tools, when they believe it may enhance outcomes. The educated workforce of the future must possess higher order skills, capable of consuming intelligence, and trigger actions to benefit, and hopefully, profit, from the deployment of AI tool kits.

Hence, these AI tools must be capable of use by general employees. The tools may be drag and drop interfaces representing abstractions. The employee may not need to deal with the computational complexities, programming principles and Boolean operators. To use these abstracted tools and intuitive interfaces, the educated workforce, in future, must possess skills which catalyzes the *consumption of intelligence, the outcome*. The educated consumer is the best customer for the future of AI and to profit from AI tools.

This raises a critical issue concerning K-12 education and how learning must be adapted to deal with the imminent socio-economic disequilibrium. Education must address the changing face of the supply chain of talent as well as the ingredients which are necessary to future-proof workforce preparation and make the workforce future-ready.

The future is not about apocalyptic reduction of employment. It is about a refresh of skills, which must be updated and upgraded, for the humans in the loop, to play relevant roles. We need the AI tool kit to help reduce uncertainty, and better manage, volatility.

To achieve that goal, executives need to appreciate the principles of data analytics and AI. Leaders must support education, inculcate insight and remain eager to learn, before they leap to manage. Leaders must institute internal education and external learning liaisons, where thinkers are viewed as assets and not as cost centers. Leaders must stress on understanding how AI works rather than blindly purchasing “black box” solutions.

The digital world will still need to serve analog communities. AI may lead to profit, if allowed to offer reliable computational assistance to the workforce, customers, and the global ecosystem of consumers, seeking credible, rational, near real-time, and perhaps predictive, decision support.

All things considered, the path to profitability rests with the imagination and the vision of the executive management and their counterparts in academia, and government.

Corporate leaders must evolve in their leadership roles. Leaders must assume the risk of leadership. Leaders must engage to provide broader guidance, bring parties to the table (competitors) and advocate for interoperability of architectures, to enable digital connectivity. Without security, digital transformation could be annihilated. Without connectivity, without data from different systems and ecosystems, without knowledge of what is *beyond the boundary*, the ability of AI, analytics and tools such as blockchain applications, will be curtailed. The outcome will be less valuable, less actionable, less profitable.

Hence, leaders must champion digital transformation by leading, and inspiring global teams, and navigating businesses to lift many boats, not just their personal yachts.

Acknowledgments. 1) This “letter” also appears as APPENDIX 6 in the essay “03.AI” available from the MIT Library – see folder “CHAPTERS” https://dspace.mit.edu/handle/1721.1/106496. 2) *The Fatal Flaw of AI Implementation* by Jeanne Ross, MIT SMR (14 July 2017) served as a source of a few examples and the reference to enterprise resource planning.

EXPLORE ADDITIONAL CHAPTERS

MIT LIBRARY – zipped folder "CHAPTERS" – https://dspace.mit.edu/handle/1721.1/106496

- 00.STUDENTS
- 01.Value Networks
- 02.DEX
- 03.Intelligence_AI
- 04.Digital Twins
- 05.Cybersecurity
- 06.UTRI_preBlockchain
- 07.Healthcare_2016
- 08.Healthcare_2008
- 09.SCM RFID DATA
- 10.SCM RISK DATA
- 11.SCM Innovation in FRENCH
- 11.SCM Innovation
- 12.SCM AGENTS
- 13.AGENTS _ Review
- 14.DIGITAL DIFFUSION
- 15.Transport
- 16.Disequilibrium in FRENCH
- 16.Disequilibrium in SPANISH
- 16.Disequilibrium
- 17.Translate
- 18.Commencement
- 19.ENERGY
- 20.CACS
- 21.STEM
- 22.OP-ED