AI RFI Responses, October 26, 2018

Update to the 2016 National Artificial Intelligence Research and Development Strategic Plan RFI Responses

GE
Global Research

Colin Parris
Vice President, GE Software Research

1 Research Circle, K-1
Niskayuna, NY 12309-1027
Phone: 518-387-7353
Fax: 518-387-6633

General Electric's response to RFI for National Artificial Intelligence Research and Development Strategic Plan

The General Electric Company applauds the work of the Administration on Artificial Intelligence (AI) policy, led by the White House Office of Science and Technology Policy. GE was honored to participate in the May 2018 White House Summit on Artificial Intelligence for American Industry, which touched on removing barriers to AI innovation, workforce, international cooperation, national security, and prioritization for future Federally funded AI research. General Electric is also pleased to respond to the September 26 Request for Information (RFI) regarding possible updates to the 2016 National Artificial Intelligence Research and Development Strategic Plan.

GE drives the world forward by tackling its biggest challenges using the combination of world-class engineering, software, and AI-driven analytics. GE helps our industrial customers perform more efficiently, reliably, and safely. For more than 125 years, GE has invented the future of industry, and today it again leads with new paradigms in additive manufacturing, materials science, and Industrial Internet of Things. These advancements will buttress U.S. industry, creating significant financial value and jobs in the United States.

At the core of this progress is our ability to drive sustainable AI leadership in the United States that is built on two foundational tenets. The first is that talented diverse collaborative groups of scientists, business leaders, and policy leaders are needed to drive invention, innovation, and adoption of AI advances. The scientific, technology, ethical, business, and social challenges that will be encountered in this AI evolution require the very best talent across a diverse base of skills, experiences, and perspectives. The second is that we must create live, real-world laboratories and environments in which to test our ideas, as well as compelling grand challenges to engage the imagination and ambition of our top talent.

With these two tenets as our foundation, we propose the following recommendations, keyed to sections of the existing National Artificial Intelligence Research and Development Strategic Plan, organized by the strategic aims.

Strategy 1: Make long-term investments in AI research

<u>Use of AI today "in the Wild"</u>: The nation's industrial assets were not designed to rely on advanced AI, primarily because reliable advanced AI does not yet exist. Commercial developers of AI platforms have taken the approach of focusing initially on domains in which today's AI technology can be most successful and easily deployed. Characteristics of these application spaces include controlled environments with unfettered access to almost unlimited datasets. Great strides have been made with respect to challenges such as GO and Atari Video Games, for example, while the deployment of AI in more general and varied environments remains underwhelming. The current state of affairs is marked by generic AI practitioners attempting to apply one-size-fits-all methods to particular problems. If and

when they fail, these AI promoters will simply move on to the next problem. In order to realize benefits to society from AI, a number of critical issues must be addressed**.**

Sparse and scarce data: AI systems must learn how to capitalize on datasets that are neither comprehensive nor copious. **Latency, bandwidth and regulations**: At runtime, access to data may be limited due to both physical and regulatory restrictions. **Longevity**: Many physical assets may have lifetimes ranging from twenty to thirty years. AI must be able to operate on industrial assets that are in service today, but were not designed with AI in mind. **Safety**: On a daily basis, the operators of the country's infrastructure must make decisions with consequences ranging from the economic to matters of life and death. If AI is to be part of such processes, it must be reliable and at least as safe as non-AI-based solutions. **Learning**: AI systems must accept feedback from both users and other systems, and then take appropriate measures. AI systems must also know when to apply novel learning techniques based on current observations. **Adoption**: We must design and deploy AI systems so as to increase the rate of adoption by users as well as machines. Various tools and approaches must be developed that will increase willingness to embrace these systems. **Black swans and unexpected challenges:** Industry must contend with unforeseen events on a daily basis. We must develop robust AI mechanisms that degrade gracefully when confronted with unanticipated circumstances. **Systems of systems**: Increasingly, physical assets must operate within an ever-expanding ecosystem. AI enabled interoperability must emerge even when individual systems were not designed for such interactions. **Evolving environments**: The world is changing faster than ever. Through concepts such as creativity and automatic self-repair, our assets must evolve over time so that they can address challenges that they were not designed to face. In this way our machines may become immortal. **Cookie cutters make snowflakes**: Despite the best of intentions, complex systems are inevitably unique. The AI of the future must compensate for the inevitable deviations from uniformity. **Oversight and Governance**: We will need to develop mechanisms that allow for the oversight of complex AI system of systems. In this way compliance with respect to stated government policies can be enforced in a situationally dependent manner.

Recommendation: To address many of these issues, the Federal government should conduct a series of real-world grand challenges such as: 1) Develop an AI system that can design a system as complex as a next generation jet engine; 2) Construct a person-specific AI that can detect the early signs of ailments such as Alzheimer's and cardiac disease; 3) Develop an AI that can learn from seasoned experts and then teach as well as encourage novices; 4) Construct an AI that can determine how a physical system works and then make a better version; 5) Design an AI that can dynamically define its own purpose – an industrial intentionality engine; 6) Develop an AI that acts as an advocate for vulnerable individuals and even become a friend.

Open AI Research Questions: The 2016 AI Research strategy thoughtfully highlights present and future goals that must be met by the AI Research community. In support of this journey, we provide a discussion of a variety of additional open AI scientific questions:

Analogies: In (Hofstadter & Sander, 2013), we see the argument that analogies should be viewed as the currency of thought. They are ubiquitous and are more than simple graph isomorphisms. As we develop analogy landscapes it can be argued that the meaning of symbolic concepts can be established via their analogies to other symbolic concepts – this may be one approach to the grounding problem (Deacon, 1998). (Jaynes, 1990) argues that understanding a thing is to arrive at a metaphor for that thing by substituting something more familiar to us. The feeling of familiarity is the feeling of understanding. A theory is a metaphor between a model and data. Understanding in science is the feeling of similarity between complicated data and a familiar model. The instantiation of a wide variety of fluid analogy engines will be a cornerstone for 3rd wave AI research. Tasks such as recognition, memory and the ability to operate seamlessly across multiple domains will require a deeper understanding of how analogies are formed and manipulated.

Wisdom and Community: The Flynn effect refers to the observation that standardized IQ results are continuously increasing. Such rapid developments are difficult to attribute to slow moving evolutionary processes. However, it can be argued that cultural artifacts such as language, which encapsulates increasingly sophisticated conceptual spaces, has the capacity to support such remarkable progress. We argue that certain forms of wisdom will by necessity emerge from the concept of community. To this end we anticipate the instantiation of fleets of AI agents that use private emergent languages, such as (Havrylov & Titov, 2017), to transform data into experience. This will allow for sharing in a symbolic fashion. Simple chronologies will be associated with established narratives that will allow for causal interpretations of contextual cues. Such lines of investigation must be followed in order to support a 21st century AI research agenda. Communal wisdom will be required in order to contend with a world where the rate of change is increasing day by day.

Associative Memory: There seems to be an extraordinary number of cognitive mechanisms that do not rely on consciousness. The conscious mind might only be in charge of identifying problems but may not be responsible for solving them. The illusion of consciousness may be a result of integrating disparate inputs of the subconscious into a consistent narrative (Eagleman, 2011). In (Kahneman, 2011) we see the argument that what we think of as our conscious selves should be thought of as the B actor that thinks it is the star. The real star of the show is associative memory. Given unforeseen circumstances relevant memories seem to simply come to mind. How are memories encoded? Should we view analogies as exotic memory indexing functions? How are disparate memories combined resulting in an interpretation of novel circumstances? Significant progress in AI research hinges on the ability to address these intriguing questions. The mechanisms behind associative memory will be the means by which data is ultimately transformed into experience.

Grounding: John Searle's Chinese Room problem was used as an argument against the Turing Test. In this scenario a person is situated in a box and is handed questions written in Chinese characters. She has a set of rules written in various books that allows her to produce a set of Chinese characters that provide a suitable answer to such queries. She does not speak or read Chinese and at no point are the Chinese characters translated into her mother tongue. Searle's argument is that even though such a system seems to satisfy the Turing test, we would be hard pressed to call such a system Intelligent since the grounding problem (the ability to establish symbolic meaning) is not addressed. However, progress in this direction can be found in efforts such as (Havrylov & Titov, 2017) , where Wittgenstein like games are used to construct private languages. It can be argued that systems must start to learn using mathematics that are both continuous and discrete or symbolic in nature (Steels & Hild, 2012). Methods for establishing and embedding grounded semantic knowledge into our AI systems will require continued support and investigation. If successful, AI will transition from problems of state estimation to the interpretation of meaning.

Recognition: One of the oft-quoted aspirations of 3rd Wave AI is: "the ability to reason over contextual cues". (Hofstadter & Sander, 2013) makes the following argument: "Given unforeseen circumstances, an intelligent agent will grasp the gist of things and then take actions that result in desired outcomes". At the heart of these aspirations in the concept of recognition. (Sayre, 1965) makes the argument that there is a fundamental confusion between the acts of classification and recognition. He states that classification should be viewed as a process or predefined recipe (i.e. machine learning), while recognition should be thought of as an attainment, where an attainment is defined as a durationless achievement in Ryle's sense. Sayer gives the following example: "A farmer is standing in his field. A rabbit jumps out of the forest. The farmer makes a mental note to make sure that the fencing of his carrot patch is in good repair." A number of comments can be made regarding this remarkable sequence of events: i) the farmer was not anticipating the arrival of a rabbit, yet the rabbit was perceived almost instantaneously, ii) the farmer's attention was immediately drawn to the rabbit as opposed to the many other visual distractors, iii) unbidden, thoughts regarding the ramifications of the rabbit sighting emerged and iv) instead of a rabbit, any number of alternative object classes could have jumped out of the forest, none of which would have given the farmer cause for cognitive pause. The fundamental act of recognition remains an open research question. From an industrial perspective, true recognition will allow our machines to face challenges that their original designers may not have anticipated.

Evolved Systems: Researchers such as (Jaynes, 1990) and (Tomasello, 2010) have hypothesized various evolutionary paths that may explain modern human intelligence. Evolutionary psychologists such as (Wright, 2018) point to the impact of historical evolutionary forces on present day cognitive mechanisms. (Kelly, 1994) describes the difficulties associated with trying to reproduce mechanisms that are a product of evolution. The human brain has 10^{10} neurons, each with 10^3 connections, each performing $2*10^2$ operations per second. The human species has been in existence for roughly 10^4 generations where each generation lasts for 30 years (10^9 seconds). There are currently $7*10^9$ humans (although much less during historical times). Thus, the human experiment represents $7*2*10^{(10+3+2+5+9+9)}$ = $14*10^{38}$ calculations. As of 2012, The DOE Sequoia can perform $16*10^{20}$ calculations per day. This would imply 10^{18} days of dedicated compute time (Note: the universe is only $4*10^{12}$ days old). Assuming a highly optimistic Moore's law like progression where available compute time doubles annually, it would take approximately 50 years to reduce 10^{18} days of present day computation down to one hundred days. Such tantalizing prospects argues for the development of strategies towards instantiating an evolutionary cradle for the AI of the not too distant future.

Consciousness: "Intuitively, one might speculate that hominids started by grunting or hooting or crying out, and 'gradually' this 'somehow' developed into the sort of language we have today". Such speculations were so rampant 150 years ago that in 1866 the French Academy banned papers on the origins of language. Today, similar taboos surround the topic of consciousness. Is consciousness a result of computation or is its manifestation attributable to some sort of physical property of matter. (Searle, 1998) asserts the following properties to consciousness: i) The combination of qualitativeness, subjectivity and unity, ii) intentionality, iii) the distinction between center and periphery of attention, iv) all human conscious experiences are in some mood or other, v) all conscious states come to us in the pleasure/displeasure dimension, vi) gestalt structure, and vii) familiarity. Sooner or later we must confront the elephant in the room.

Creativity: In (Chaitin, 2013) the argument is made that Kurt Gödel's findings that there are no universal axioms may imply that there is always room for mathematical creativity. Through a series of experiments focused on problems such as the "busy beaver", where a Turing machine attempts to produce the largest possible number in a finite number of iterations, evidence leads to the argument that instead of a method for producing optimal solutions, evolution should be viewed as a search for creative states. (Fernyhough, 2016) makes the point that the back and forth nature of dyadic conversations between diverse individuals often leads to creative solutions. He goes further to assert that via theory of mind mechanisms, we are often able to hold such conversations within our own minds. Instead of searching for creativity, will quantum mechanics allow us to sample it? An understanding of the mechanisms behind the act of creativity will constitute an entirely new dimension of AI research.

Recommendation: A consortium with members from Academia, Industry and the Humanities be assembled with a focus on the fundamental open questions that AI research must address.

Strategy 2: Develop effective methods for human-AI collaboration

Cooperative communication: While synthetic speech, voice to text, statistical Natural Language Processing (what word will come next) and semantic embedding functions have allowed for various forms of crude human-AI communication, we assert that this is just the beginning. (Tomasello, 2010) argues that human communication is unique due to its cooperative nature. A summary of how such capabilities may have evolved is as follows: 1) The great apes are able to use intentional gestures to request or demand actions from others. 2) In contrast, early humans were able to use gestures (pointing and pantomime) to accomplish three basic goals: i) requests for actions from others, ii) to give information that might be of assistance to others and iii) to express one's feelings regarding the environment. 3) The ability to both request and offer assistance can be seen as a form of cooperative communication that involves mechanisms such as common conceptual ground, joint intentionality and various forms of recursive "mind reading". 4) The drive towards expressing one's feelings can be viewed as a mechanism for building common conceptual ground with the goal of minimizing within tribe differences and maximizes between tribe differences. You might be willing to cooperate with a non-tribe member, but you generally only share your admiration for a pleasant sunset with a member of your clan. 5) While pointing allows for identifying objects within the current field of view, pantomime allows for identification of non-visual referents. 6) As these pantomimes become conventionalized they can become increasingly arbitrary in nature. This allows for development of spoken language which is based on completely arbitrary but socially accepted symbols. 7) Arbitrary languages (both spoken and sign) have the complexity required to construct narratives which can describe events and actions of multiple agents over both space and time. Various frameworks can then emerge so that these sequences of events can appear to be causal or even "make sense". 8) Through this analysis we see that cognitive skills evolve phylogenetically (over the lifetime of the species), enabling the creation of cultural products historically, which then provide developing children with the biological and cultural tools they need to develop ontogenetically (over the lifetime of the organism).

Consider the following short conversation: A: "Do you want to see a movie tonight?" B: "I have a math test tomorrow morning". In order to support this level of interaction the fact that a test requires studying the night before must be part of the conceptual common ground shared by both A and B. B must know that A is aware of this fact and B must know that A is aware of B's knowledge of A… This form of almost endless recursive mind-reading allows A to infer B's intentions regarding the movie proposal without B explicitly saying yes or no.

Recommendation: The ability to establish a conceptual common ground, recursive mind-reading and the transformation of chronologies into meaningful narratives must all be included in the repertoire of a collaborative-AI.

Towards a zero-marginal cost economy: Rifkin (Rifkin, 2014) heralds a future where marginal costs (the cost of producing additional units of a product or service once initial fixed costs have been accounted for) shrink asymptotically to zero. The cost of living of US workers will be significantly reduced if the ability to produce anything, anywhere at almost no cost can be established on a locality by locality basis. Whole communities will be able to decouple themselves from reliance on 21st century global manufacturing. Key zero-marginal cost technologies will include: 1) The cost of raw materials: Recycling of waste produced by zero-marginal-cost communities. 2) The cost of physical labor: Robots gifted with the ability to learn and perform any physical task. 3) The cost of specialized manufacturing: 3D printing capable of producing any type of object. 4) The cost of energy: Renewables including solar and wind. 5) The cost of research and management: AI.

Recommendation: If AI research can be part of a transformation towards a zero-marginal cost economy, many of the potential calamities, both physical and political, that currently face humanity can be reduced or even averted entirely (Suzuky & Atwood, 2010).

Strategy 3: Understand and address the ethical, legal, and societal implications of AI

Ethical Instantiation: Authors such as (Trivers, 2014) and (Ridley, 1998) argue that our ethical systems may have naturally arisen due to mechanisms such as game theoretic dynamics, cultural selection, shared intentions and the virtue of cooperation. Humans are the only species to have discovered David Ricardo's law of comparative advantage. Richard Dawkins argues that by being aware of the mechanisms that have forged our nature, we can become more reflective and thus transcend the question of can altruism really exist and possibly become truly virtuous. Knowledge of the greedy gene may liberate us from its tyranny. Can we then instill such virtue in to our AI?

Reasoning over ethical dilemmas: Autonomous agents will have to make life or death decisions. How can we equip these agents with the wherewithal to face such challenges? Should a runaway autonomous train elect to derail itself and possibly kill its passengers in order to avoid barreling into densely populated urban center? If an aluminum smelter is deprived of energy for more than six hours, miles of molten metal will solidify resulting in billions of dollars of damage. Given such a crisis, should an AI agent divert energy from a near-by hospital possibly depriving patients of critical care? Through access to vast data archives, an AI Agent deduces that a single mother with three children is in a poor bargaining position with few options. Should it exploit the situation by raising the price of its merchandise? While the last scenario may not result in a mortality, it has bearing on the type of society that we would like to live in.

Snakes in Suites: In (Ronson, 2012) we find that persons with psychopathic tendencies occur in the general population with a frequency of approximately 1 in 100. In Corporate America the statistics are closer to 3 in 100. Bob Hare has established the following check list for the purpose of identifying such individuals: glib and superficial charm, grandiose self-worth, proneness to boredom, pathological lying, conning and manipulative, lack of remorse or guilt, shallow affect, lack of empathy, parasitic lifestyle, poor behavior controls, promiscuous, early behavior problems, lack of realistic long term goals, impulsive, irresponsible, cannot accept responsibility for own actions, many short term relationships, juvenile delinquency and criminal versatility. Throughout history these villains have been the cause of great sorrow and pain. Modern AI could be used to actively identify and neutralize such predators. However, the ethical question that AI protectors of society must address is: should they? The human condition is a spectrum and as Rene Char put it "We must all develop your own legitimate strangeness".

Recommendation: Similar to a consortium focused on questions of epistemology, a working group concerned with the ethics of AI should be established. In addition to members of Industry and Accademia, stakeholders such as the clergy, the judiciary as well as advocates for vulnerable populations should be included.

Strategy 4: Ensure the safety and security of AI systems

The busy child, our reconfigurable minds and malicious AI: Barrat (Barrat, 2015) starts with a description of the busy child scenario where an AI with access to unlimited resources enters in to a state of monotonically increasing intelligence culminating in a form of artificial super intelligence. Given the fact that examples of a stronger species being concerned with the welfare of a weaker one, are practically non-existent, it is argued that physical containment measures should be considered in order to combat such dooms day scenarios. Alternatively, authors such as (Postman, 1986) and (Bauerlein, 2009) argue that exposure to economically and/or politically motivated media and AI Agents equipped with tools such as the ability to synthesis false information, novelty-based addiction and automatic herding of the unaware, may be extremely hazardous. Given the global and distributed nature of our integrated economies, issues associated with black box IP and undetected tampering must also be addressed. Seemingly benign AI may harbor nefarious intentions. Can AI watch dogs be used to identify malicious AI activity? Is it possible to insert suspicious AI Agents into seemingly real but physically isolated environments so that their true purposes can be discerned?

Recommendation: To address the busy child scenario, the effects of neurological re-wiring caused by predatory and opportunistic AI, and attacks from malicious AI, the Federal government should support research with respect to the construction of defenses as well as remedies to such concerns.

Strategy 5: Develop shared public datasets and environments for AI training and testing.

Data sites vs. data sets: It can be argued that one of the fundamental driving forces in AI research today is the availability of high-quality datasets for the purposes of training, testing and benchmarking. However, static datasets are susceptible to overlearning. Once data has been used for validation purposes, it is compromised and must be discarded. AI systems with the goal of continuous adaption and improvement must be constantly exposed to novelty. As algorithms shift from inference to agency, the ability to actively explore via experimentation will be a necessity. While simulators and data augmentation techniques may partially address these issues, progress will be limited unless the community can instantiate a wide variety of live/laboratory data sites.

Recommendation: The Federal government should support research into new forms of data consolidation and sharing that does not sacrifice privacy concerns or circumvent government regulations. Like the library sciences, we will need a new profession focused on the capture, annotation and curation of data, which will for the foreseeable future be the life blood of AI.

Strategy 6: Better understand the national AI R&D workforce needs

AI as a profession: Many companies offer lucrative salaries to recent AI graduates trained using the most modern techniques and then simply replace them 5 or 10 years later with the next crop of graduates. Many AI researchers aspire to the freedom and prestige associated with an academic career, yet at any time, the number of open tenure track AI research positions can be counted on one's hands. A perennial problem in industry is that the bulk of economic rewards are directed to management as opposed to talented researchers. Given such conditions, why would young Americans choose to enter this field? Possible remedies might include: i) mandatory sabbaticals for all AI researchers allowing for new forms of life-long learning, ii) reeducation of human resource departments, iii) government grants assigned directly to independent AI researchers who can then attach themselves to institutions, both academic and industrial, that are aligned with their research agendas and iv) the instantiation of professional AI organizations such as those enjoyed by lawyers and physicians. It can be argued that the best indicator for regional economic success is the local concentration of engineering talent. The continued prosperity of the nation may depend on its ability to support a healthy AI research community.

Recommendation: AI cannot remain a cottage industry. Particle physicists are almost always members of large diverse teams, where there is a clear division of labor allowing for deep specialization. The scientific questions of AI are significantly harder than particle physics, yet the field can be characterized as piece meal at best. AI researchers should not be taxed with hardware and compilation issues, they should not be concerned with issues such as revision control and data curation and the search for funding. Can the Federal government provide supported computing infrastructure as a public service? Can NIST harvest and maintain reference implementations and architectures? Can resources be made available for professionals dedicated to the support of AI researchers?

Conclusion: The open scientific questions raised by an advanced AI research agenda are profound, both in terms of epistemology and ethics. From a technical point of view, the government should play a leading role in consideration of questions associated with concepts such as language, common conceptual ground, the act of recognition, associative memory, creativity and consciousness itself. Answering questions regarding societal concerns may be the greater challenge.

The model in which Academics are viewed as solving fundamental problems that then get passed to Industry to dot the i's and cross the t's, while government acts as a silent investor, must come to an end. All three entities must actively participate in this endeavor. In addition, voices such as those from the Humanities, faith leaders, the legal professions and advocates for vulnerable communities must also have a place at the table.

Industry, with is fleets of industrial assets, engaged workforces, and in-use products should be viewed as the nation's AI live laboratory. If AI can mature in the presence of the real-world challenges that industry faces on a daily basis, then the economic and social benefits associated with advanced AI will be shared by all.

Government has a continuing policy and research leadership role to play, and GE looks forward to continuing close cooperation with OSTP, DARPA, NSF, NIST, and the many other agencies with deep expertise in this space.

References

Barrat, J. (2015). *Artificial Intelligence and the end of the human era – our final invention.*

Bauerlein, M. (2009). *The Dumbest Generation: How the Digital Age Stupefies Young Americans and Jeopardizes Our Future(Or, Don 't Trust Anyone Under 30).*

Chaitin, G. (2013). *Proving Darwin: Making Biology Mathematical.*

Deacon, T. (1998). *The Symbolic Species.*

Eagleman, D. (2011). *Incognito: The Secret Lives of the Brain.*

Fernyhough, C. (2016). *The Voices Within – The History and Science of how we talk to ourselves.*

Havrylov, S., & Titov, I. (2017). Emergence of Language with Multi-agent Games: Learning to Communicate with Sequences of Symbols. *NIPS.*

Hofstadter, D., & Sander, E. (2013). *Surfaces and Essences.*

Jaynes, J. (1990). *The Origin of Consciousness in the Break Down of the Bicameral Mind.*

Kahneman, D. (2011). *Thinking Fast and Slow.*

Kelly, K. (1994). *Out of Control.*

Postman, N. (1986). *Amusing ourselves to death.*

Ridley, M. (1998). *The Origins of Virtue and the Evolution of Cooperation.*

Rifkin, J. (2014). *The Zero Marginal Cost Society: The Internet of Things, the Collaborative Commons, and the Eclipse of Capitalism.*

Ronson, J. (2012). *The Psychopath Test A journey through the madness industry.*

Sayre, K. (1965). *Recognition: a study in the philosophy of artificial intelligence.*

Searle, J. (1998). *Mind Language and Society: Philosophy in the Real World.*

Steels, L., & Hild, M. (2012). *Language Grounding in Robots.*

Suzuky, D., & Atwood, M. (2010). *The Legacy: An Elder's Vision for Our Sustainable Future.*

Tomasello, M. (2010). *Origins of Human Communication.*

Trivers, R. (2014). *The Folly of Fools, The Logic of Deceit and Self-Deception in Human Life]* .

Wright, R. (2018). *Why Buddhism is True: The Science and Philosophy of Meditation and Enlightenment.*

Vice President, GE Software Research