AGIRI Workshop 2006

Theme: Transitioning from Narrow AI to Artificial General Intelligence

Focus: Grounding Linguistic Relationships in Nonlinguistic Reality

The field of AI is poised to make a transition from a focus on highly specialized "narrow AI" problem solving systems to confronting the more difficult issues of "human level intelligence" and more broadly "artificial general intelligence."

This workshop will focus on AGI issues in general, with a non-exclusive focus on the theme of grounding linguistic relationships in nonlinguistic reality.

Speakers will comprise individuals actively involved with the design and/or implementation of AGI systems. Additional speakers will provide insight into relevant aspects of AGI theory and application. The event will be filmed and selective proceedings published.


The term *Artificial Intelligence* was first used in 1955 by **John McCarthy** to propose the 1956 Dartmouth Summer Research Project on Artificial Intelligence. In recognition of the **50th anniversary of AI**, the **Artificial General Intelligence Research Institute (AGIRI)** hosts its first workshop May 20-21, 2006.

Novamente LLC is the exclusive sponsor:

NOVAMENTE LLC
ARTIFICIAL GENERAL INTELLIGENCE

Organizing Committee:

Dr. Ben Goertzel
ben@novamente.net

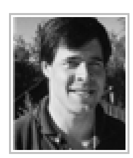

Dr. Phil Goetz
philgoetz@gmail.com

Bruce Klein
bruce@novamente.net

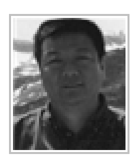

Dr. Pei Wang
dr.peiwang@gmail.com

Schedule, Sat. MAY 20th, 2006

7 - 8 AM	Coffee/Snacks	
8:00 AM	**SAT. AM - MAY 20**	**INTRO & KEYNOTE SPEAKER**
40 min	**Dr. Stan Franklin**	A Cognitive Theory of Everything: The LIDA Technology as an Artificial General Intelligence
		SESSION I: AGI BACKGROUND (Mod: Dr. Ben Goertzel)
20 min	**Eliezer Yudkowsky**	AI as a Precise Art
20 min	**Dr. Matthew Ikle'**	Imprecise Probabilities and Their Role in General Intelligence
20 min	**Richard Loosemore**	Complexity Science and Its Implications for AGI
10 min	Break	
20 min	**Cassio Pennachin**	Frontiers of Evolutionary Learning
20 min	**Stephan Vladimir Bugaj**	AGI Developmental Psychology
20 min	**David Hart**	Distributed Processing and Large-Scale System Engineering for AGI
20 min	**Jeff Medina**	Ethical and Mathematical Adventures in Coding Better Children
12 - 1:30 PM	Lunch Served	
1:30 PM		**SESSION II: CONTEMPORARY AGI PROJECTS (Mod: Cassio Pennachin)**
40 min	**Dr. Hugo de Garis**	Building an Artificial Brain for Less Than $10,000
40 min	**Dr. Sam S. Adams**	Superstition and Forgetfulness: Essentail Attributes of General Intelligence
40 min	**Dr. Eric Baum**	A Model of Thought and Implications for AGI
10 min	Break	
40 min	**Dr. Pei Wang**	From NARS to a Thinking Machine
40 min	**Dr. Nick Cassimatis**	A Cognitive Substrate for Human-Level Intelligence
10 min	Break	
40 min	**Steve Grand**	Machines like us
40 min	**Dr. Ben Goertzel**	Novamente: A Practical Architecture for Artificial General Intelligence
7:20 PM	**END**	

Note: Speakers have a TOTAL of either 20 or 40 min. in which to speak and (if desired) do Q&A. Thus, it is suggested that speakers use their first 15 min. (or 35 min.) for speaking and the remaining 5 min. for Q&A. Also, if you have one, please provide your PowerPoint to Bruce Klein at your earliest, preferably via EMAIL (bruce@novamente.net), or Bruce can load your PowerPoint onto his USB storage device during the break before your presentation. Thank you!

Schedule, Sun. MAY 21th, 2006

7 - 9:00 AM	Coffee/Snacks	
9:00 AM	**SUN. AM - MAY 21**	**SESSION III: AGI & APPLIED COMPUTATIONAL LINGUISTICS (Mod: Dr. Ben Goertzel)**
20 min	**Dr. Karl H. Pribram**	Neural Networks or Neural Webs?
20 min	**Michael Ross**	Mining Semantic Information From Unstructured Data
20 min	**Dr. Deborah Duong**	Houses of Mirrors: Deeply Adaptive Designs for Machine Cognition
10 min	**Dr. Alexei V Samsonovich**	Self-organizing Linguistic Cognitive Maps as a Key to the Human Value and Semantic Memory Systems
10 min	Break	Topic Intro by Bruce Klein (5 min)
20 min	**Dr. Thomas Rindflesch**	Knowledge-Based Semantic Interpretation for Summarizing Biomedical Text
10 min	**Dr. Ben Goertzel & Izabela Lyon Freire Goertzel**	Using Dependency Parsing and Probabilistic Inference to Extract Gene/Protein Interactions Implicit in the Combination of Multiple Biomedical Research Topics
20 min	**Moshe Looks**	Contemporary Approaches to Symbol Grounding
20 min	**Ari Heljakka**	Learning Symbol Groundings via Simulated Robotics
10 min	**Dr. William Sims Bainbridge**	About The National Science Foundation (NSF)
12 - 1:30 PM	Lunch Served	
1:30 PM		**SESSION IV: PANEL DISCUSSIONS**
Panel I 1:30 - 2:30 PM	**How are linguistic relationships grounded in nonlinguistic reality?**	Moshe Looks, Michael Ross, Dr. Thomas Rindflesch, Dr. Phil Goetz, (Mod: Dr. Ben Goertzel)
10 min	Break	
Panel II 2:40 - 3:40 PM	**How do we more greatly ensure responsible AGI?**	Eliezer Yudkowsky, Jeff Medina, Dr. Karl H. Pribram, Ari Heljakka (Mod: Stephan Vladimir Bugaj)
20 min	Break	
Panel III 4 - 6 PM	**What are the bottlenecks, and how soon to AGI?**	Dr. Stan Franklin, Dr. Hugo de Garis , Dr. Sam S. Adams , Dr. Eric Baum, Dr. Pei Wang, Dr. Nick Cassimatis, Steve Grand , Dr. Ben Goertzel (Mod: Dr. Phil Goetz)
6 PM	**END**	

Directions / Hotels

Bethesda North Marriott Hotel & Conference Center
5701 Marinelli Road
North Bethesda , Maryland 20852
Reservations: **1-800-859-8003**
(There are a limited number of hotel guest rooms ($159) available under the "**AGIRI Workshop**" name.)

If the Marriott is booked, there are two other suggested hotels located 1.2 miles north... Doubletree ($129) and Ramada Inn ($139). To reserve your room, please try: http://www.orbitz.com then click: "Hotels", click: "Search by address", enter: 5701 Marinelli Rd N. Bethesda, MD 20852, click: "Check-in" May 19, click: "Check-out" May 22, then click: "Find". Closest Airport: **Ronald Reagan Wash. National Airport** (DCA) There are two ways to travel from Reagan Airport to the Marriott:

- Taxi is least complicated.
- For the more adventurous, the Metro (Red Line) stops directly in front of the Marriott.

Both Taxi and Metro take about the same time... 50 minutes.

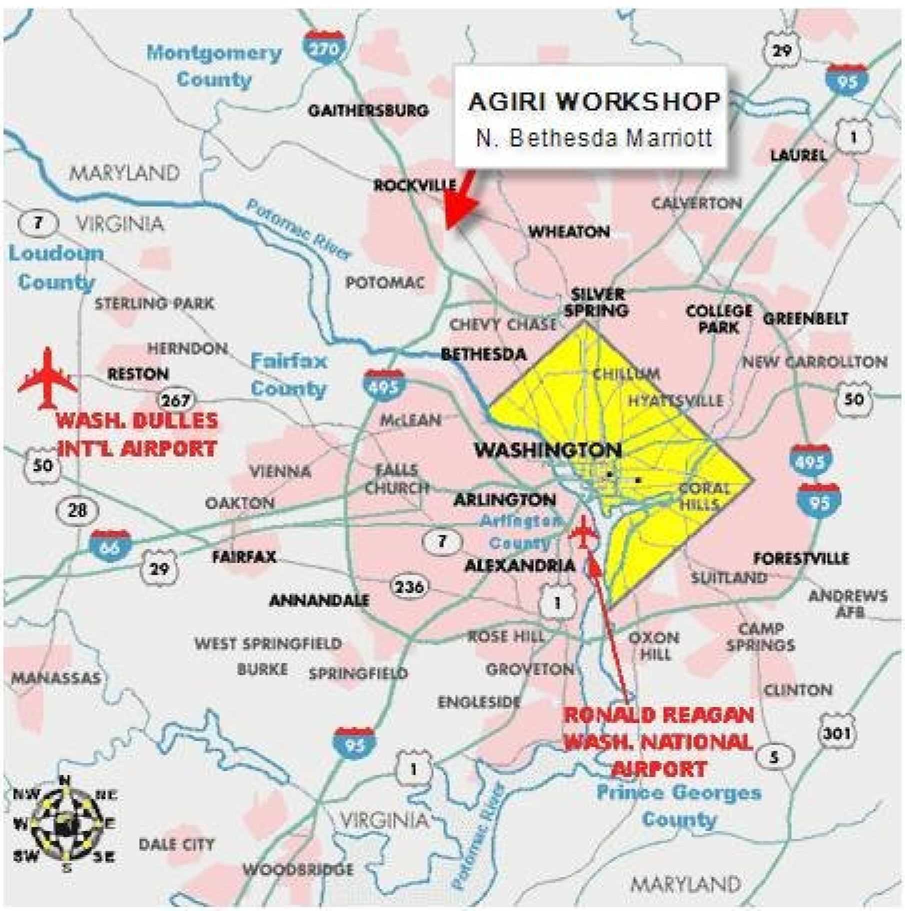

Speakers / Topics

Intro & Keynote (Sat. 8 AM - May 20)

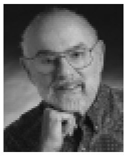

Dr. Stan Franklin, Dir. Institute for Intelligent Systems, University of Memphis
Topic: **A Cognitive Theory of Everything: The LIDA Technology as an Artificial General Intelligence**

Implementing and fleshing out a number of psychological and neuroscience theories of cognition, the LIDA conceptual model aims at being a cognitive "theory of everything." With modules or processes for perception, working memory, episodic memories, "consciousness," procedural memory, action selection, perceptual learning, episodic learning, deliberation, volition, and non-routine problem solving, the LIDA model is ideally suited to provide a working ontology that would allow for the discussion, design, and comparison of AGI systems. The LIDA technology is based on the LIDA cognitive cycle, a sort of "cognitive atom." The more elementary cognitive modules play a role in each cognitive cycle. Higher-level processes are performed over multiple cycles. This talk will give a quick overview of the LIDA conceptual model, and its underlying computational technology.

Biography: A mathematician turned computer scientist turning cognitive scientist, Stan Franklin is the W. Harry Feinstone Interdisciplinary Research Professor in the Computer Science Division at the University of Memphis. He's a recipient of the University's Board of Visitors Eminent Faculty Award, a FedEx Institute of Technology founding fellow, and Director of the Institute for Intelligent Systems. His research is motivated by wanting to know how minds work—human minds, animal minds and, particularly, artificial minds. For some years he's worked on "conscious" software agents, that is, autonomous agents modeling a psychological theory of consciousness. In addition to practical, real world applications, these agents model human and animal cognition and provide testable hypotheses for cognitive scientists and neuroscientists. This endeavor, funded by the US Navy, has been the subject of over thirty papers in scientific journals and conference proceedings. Stan's graduate degrees are from UCLA, his undergraduate degree from the University of Memphis. He has authored or co-authored almost a hundred academic papers and several edited volumes, as well as authoring a book entitled Artificial Minds published by the MIT Press, which was a primary selection of the Library of Science book club, and has been translated into Japanese and Portuguese.

Session I: AGI Background (Sat. AM – May 20)

Speakers address important conceptual, scientific and engineering issues related to AGI and the conference theme.

Eliezer Yudkowsky, Singularity Institute for Artificial Intelligence;

***Topic:* AI as a Precise Art**

If an AGI is to carry out recursive self-improvement in a stable fashion, it needs an extremely low independent probability of catastrophic error on each individual round of self-modification. In principle, the inside of a computer chip is a near-deterministic environment in which extremely low error probabilities should be achievable. However, the most powerful current AI techniques for writing programs are unsuited to deterministic self-modification because they only stochastically match a fitness criterion. The worst-case scenario is that AI techniques advance to the point of giving birth to AGI, but the techniques are intrinsically unsuited to giving that AGI a stable motivational structure. A benefit of focusing on the problem of provable stability through self-modification is that proof is antithetical to vague theories and vague expected benefits - to prove an expected benefit you must be able to state it precisely. The less we understand a process and the less we understand what we want from it, the easier it is to convince ourselves that the process will deliver good results, but the less likely it is to actually do so. By aspiring to a higher standard of proof before expecting positive results from an AGI design, we can not only avoid embarassment, but force ourselves to develop a deeper understanding of terms we would otherwise be tempted to leave vague.

Dr. Matthew Ikle', Adams State College, Department of Mathematics and Computer Science;

***Topic:* Imprecise Probabilities and Their Role in General Intelligence**

The ability to reason upon incomplete and uncertain information is clearly an important attribute of any generally intelligent system. What is, perhaps, not as obvious is the role of imprecision in this reasoning ability. However, it is being increasingly acknowledged in the research community that the capability of effectively managing uncertainty in the context of complex inferences on large amounts of information is a critical aspect of intelligence.

One popular approach to uncertain reasoning is classical Bayesian inference, which is founded upon the premise that all information, and hence uncertainty, obeys a single underlying precise probability distribution. Imprecise probability, on the other hand, is a generic term often used to represent any of a variety of less sharp, through no less rigorous, measurements of uncertainty. In his seminal work, *Statistical Reasoning With Imprecise Probabilities*, Peter Walley argues that it is more natural and consistent to represent probabilities as intervals than as a single precise value. Kurt Weichelsberger provides an alternate axiomatic approach to imprecise probabilities that provides similar interval-valued probabilities.

One practical issue with using interval probabilities in the context of probabilistic inference rules, however, is the pessimism implicit in interval arithmetic. If one takes traditional probabilistic calculations and replaces the probabilities with intervals, then one finds that the intervals rapidly expand to the interval [0,1]. Walley's, Weichelsberger's and other contemporary approaches to imprecise probabilities suffer from this problem in various forms in spite of their mathematical and statistical sophistication.

To overcome this deficiency, we have developed a new approach to handling incomplete and uncertain knowledge. Using probability distribution envelopes, we generalize both Walley's model and the standard Bayesian model. As a result, standard Bayesian inference models and imprecise probability models both reduce to special cases of a more general principle.

In the context of artificial general intelligence, our new envelope-based approach provides AI architects with unprecedented flexibility. In the context of our Probabilistic Logic Network (PLN) reasoning system (developed for use within the Novamente AI Engine), we demonstrate how our generalized notion of imprecise probabilities allow for heuristic rules, classical Bayesian inference rules, and imprecise probability rules to interoperate with ease.

Richard Loosemore (former Director of Research at Starbridge Systems)

Topic: **Complexity Science and Its Implications for AGI**

The purpose of this talk is to make three claims. (1) To argue that the Complex Systems Critique of AI has often been seen as somewhat vague, whereas in fact it can be formulated in a way that makes it precise, powerful and prescriptive. (2) To propose a formalism for describing the structure of intelligent systems that exposes the weaknesses of AI that are targeted by the Complex Systems Critique. (3) To propose a software development environment designed specifically to allow intelligent systems to be built that are consistent with the suggested formalism.

Complex Systems researchers argue that we should be careful about a trap that could be called the "global-to-local disconnect," which works as follows (1) We start out with the intention of capturing the global properties of intelligence using systems of interacting, adaptive local elements. (2) Our interacting local elements turn out to constitute a complex system. (3) Because these systems are complex, there is not necessarily any obvious or analytic connection between local and global, so our naïve local implementation of the local mechanisms is not guaranteed to yield the global behavior we desire. (4) We (perhaps unconsciously) modify our proposed local mechanisms and our research goals in order to make it seem that our systems are working better than they are.

In the first part of the talk, we look briefly at some of the ways that the global-to-local disconnect actually manifests itself in practice, and at some of the ways that we try to sidestep it.

In the second part of the talk we propose a formalism for describing the structure of complete AI (AGI) systems. At the core of the framework is a "foreground" (working memory) in which "elements" (basic units of knowledge) can instantiate and interact with one another. In general, what these elements do is engage in a process of dynamic relaxation: they impose weak constraints on one another in an attempt to form connected structures of elements that are mutually consistent. External constraints act on the foreground from a sensory input area, while an analogous motor output area allows the foreground to have effects on the world. Most of the meat of the formalism is in the specific choice of local mechanisms used to implement the dynamic relaxation: we describe some of the candidate mechanisms, and examples of how various conventional AI systems map onto this framework.

In the final part of the talk we suggest that a proper methodology for escaping the complex systems problem is to build a software development environment that allows AGI systems to be quickly built and systematically compared. The previously described conceptual framework is designed to be used as the language within which AGI systems would be built in this development environment. We outline both the structure of the proposed environment and the way it would be used as part of a concerted research paradigm. The main features of the proposed environment are its use of parallelism, the ease with which systems can be constructed, and the tools for empirical study of the behavior of candidate systems.

Cassio Pennachin, VettaLabs and Novamente LLC

***Topic:* Frontiers of Evolutionary Learning**

Classical evolutionary learning and optimization policies such as the genetic algorithm and genetic programming hope to perform well on a multitude hard problems by adapting their neighborhood structures to the problem as search proceeds, via the crossover operator. However, crossover sometimes creates more problems than it solves, by introducing biases that are difficult to understand and control. These biases lead to poor scalability. Recently a class of "competent" adaptive optimization algorithms have been introduced that replace crossover with explicit centralized probabilistic modeling and sampling to adaptively learn a problem decomposition. While such algorithms are a significant improvement over "classical" evolutionary techniques, we believe that they aren't sufficiently adaptive for the pattern and program learning needs of an AGI. It is insufficient to merely learn a problem decomposition in terms of prespecified solution parameters - the parameters of the representation themselves must be adapted as search progress. Furthermore, prior knowledge of operator semantics is required to tractably learn problem decompositions over more general representations (e.g., computer programs). We have recently developed an approach to learning, probabilistic adaptive-representational evolution (PARE), which fulfils these requirements.

Stephan Vladimir Bugaj, co-author of *The Path to Posthumanity*

***Topic:* AGI Developmental Psychology**

Distributed computing, an architecture where numerous computing resources are physically distributed at any distance where intercommunication is practical, and its logical counterpart massively parallel processing, are key components of most practical architectures for AGI systems. While AGI is theoretically achievable utilizing traditional computing and programming methods on very large and fast singular computers, this talk will discuss reasons for choosing distributed computing techniques and, as this choice influences most aspects of AGI systems design and implementation, its role in the design of Novamente as a representative AGI system. Programming for distributed computing presents numerous challenges and programming overheads, as it represents a "paradigm shift" from traditional linear computing used in nearly every computer built to date. The benefits of distributed computing however include a nearly limitless scaling potential, low costs of entry, low costs for organic growth, and the flexibility of a heterogeneous infrastructure. By way of example, numerous architectural hooks for distributed processing exist today in development versions of Novamente, and strategies for distributing Novamente's evolutionary learning component (MOSES) will be discussed.

David Hart, Atlantis Blue

***Topic:* Distributed Processing and Large-Scale System Engineering for AGI**

Distributed computing, an architecture where numerous computing resources are physically distributed at any distance where intercommunication is practical, and its logical counterpart massively parallel processing, are key components of most practical architectures for AGI systems. While AGI is theoretically achievable utilizing traditional computing and programming methods on very large and fast singular computers, this talk will discuss reasons for choosing distributed computing techniques and, as this choice influences most aspects of AGI systems design and implementation, its role in the design of Novamente as a representative AGI system. Programming for distributed computing presents numerous challenges and programming overheads, as it represents a "paradigm shift" from traditional linear computing used in nearly every computer built to date. The benefits of distributed computing however include a nearly limitless scaling potential, low costs of entry, low costs for organic growth, and the flexibility of a heterogeneous infrastructure. By way of example, numerous architectural hooks for distributed processing exist today in development versions of Novamente, and strategies for distributing Novamente's evolutionary learning component (MOSES) will be discussed.

Jeff Medina, Singularity Institute for Artificial Intelligence; Institute for Ethics & Emerging Technologies; School of Philosophy, Univ. of London

Topic: **Ethical and Mathematical Adventures in Coding Better Children**

Artificial general intelligence [AGI] is ostensibly, though not necessarily, a new variety of sentient living thing. It is, therefore, potentially worthy of moral consideration; of treatment befitting a morally valuable life form. After looking briefly at one of the stronger, decision-theoretic reasons to treat AGIs like they matter morally (as opposed to other, similarly-aimed arguments that frequently depend upon certain more contentious metaphysical premises) -- like they are more suited to the company of the "born" than the "built" -- we investigate the responsibility we, generally and as AGI engineers, have for the lives we create. Are AGI engineers in a position analogous to parents, or to gods, or neither? And what does the answer provide as further guidance for aspiring life-builders? Two of the better possibilities are considered in more detail, but still each has relatively significant problems. The most philosophically promising answer beckons us to solve a mathematical optimization problem, itself contingent upon open problems concerning the relationships between intelligence, problem-difficulty, solution-findability, resource usage, time and substrate speed, and the utility of incremental self-improvement.

Session II: Contemporary AGI Projects (Sat. PM - May 20)

Speakers review their AGI projects (where relevant including a brief discussion of how their AGI systems deal with the focus issue of grounding linguistic relationships in nonlinguistic reality).

Dr. Hugo de Garis, Computer Science Department, Utah State University, Logan, Utah, USA. International School of Software, Wuhan University, Wuhan, Hubei Province. ***Topic:*** **Building an Artificial Brain for Less Than $10,000**

Most artificial brain projects currently underway are expensive, involving the use of supercomputers or PC clusters containing 100s-1000s of processor nodes (e.g. Switzerland's "Blue Brain Project" with its IBM "Blue Gene" supercomputer, or Artificial Development's 1000 PC processor cluster brain simulation (http://www.ad.com)). This talk presents a way to build an artificial brain for less than $10,000. The approach my research teams take is the following. A 6-megagate FPGA electronic board (http://www.celoxica.com) costing about $7000 is used to evolve neural network modules via a standard genetic algorithm. The software is written in "Handel-C", a high level, C-like language, that is hardware compiled into the FPGA. Using the electronic speeds of the Celoxica board, the evolution of a single module takes minutes, as against hours to days on an ordinary PC. Evolution speed is essential if 10,000s of modules are to be evolved to build an artificial brain in a reasonable time. Each module is downloaded into the memory of a PC. Once several 10,000s of such modules (each with its own little evolved function) are downloaded, they are interconnected according to the designs of "BAs" (Brain Architects) to build artificial brains. The neural signaling of the whole brain is performed by the PC in real time (i.e. 25 Hz per neuron). 2-way radio antennas between the brain in the PC and the robot allow the brain to control the many autonomous behaviors of the robot containg a CMUCAM2 camera, a gripper, and wheels. Such a robot can be built for less than $1000, hence the total cost of brain + vehicle $10,000. At such a price, it is hoped that "Brain Building" or "Artificial Brains" as a research field, can now afford to take off.

Dr. Sam S. Adams, IBM Distinguished Engineer, IBM Research

Topic: **Superstition and Forgetfulness: Essential Attributes of General Intelligence**

Decades of effort related to Artificial Intelligence have failed to recognize two of the primary characteristics of general intelligence, namely superstition and forgetfulness. This failure was primarily due to two fundamental though unspoken principles of computer science: that only well established "truths" about a domain are relevant to a system's operation, and that the loss of any acquired data during operation is inherently a bad thing. Our research into creating a system capable of Artificial General Intelligence (AGI) has established not only the necessity of both superstitiion and forgetfulness in embodied, autonomous AGI systems, but their primacy in the process of bootstrapping and maintaining AGI within any unknown embodiment situated within any unknown environment.

This talk will define what we mean by superstition and forgetfulness in an AGI, present the logical and biological justification for our claim to their crucial role in bootstrapping general intelligence, and discuss practical implementation issues of these essential attributes within the Joshua Blue project at IBM.

Dr. Eric Baum, Baum Research Enterprises, author of What Is Thought?

Topic: **A Model of Thought and Implications for AGI**

My book What is Thought? proposed a model of what thought is and how it works. I survey the model, discuss implications for how and whether it may be possible to achieve AGI, and discuss recent work on implementation.

Dr. Pei Wang, Temple University

Topic: **From NARS to a Thinking Machine**

NARS (Non-Axiomatic Reasoning System) is a project aimed at a general-purpose intelligent system.

This project is based on the belief that "intelligence" is the capability of a system to adapt to its environment and to work with insufficient knowledge and resources. This capability can explain the difference between the human mind and machines (or animals), and such a working definition of intelligence clearly distinguishes Artificial Intelligence from other branches of science, such as computer science and psychology. Furthermore, many practical problems are "hard" exactly because they need to be solved without sufficient knowledge and resources.

NARS formalizes this theory of intelligence in the framework of a reasoning system. This framework is used because of its domain-independent nature. It also combines the justifiability of individual inference steps and the flexibility of linking these steps together in a context-sensitive manner in run time.

The development of NARS takes an incremental approach. In each phase, part of the formal model is added into the system, based on the parts already established in the previous phase. Consequently, the system at the end of each phase is more intelligent than the previous one (according to the working definition of NARS) by having a richer language and more inference rules, and being more adaptive and more efficient in using available knowledge and resources.

When NARS is fully implemented, it will completely realize the working definition of intelligence, as defined above, in a general-purpose reasoning system. However, such a system by itself will not have human-like capability and behavior, because they come from human experience. Though human experience is not required for a system to be intelligent, there are many other reasons to approximate it in an AI system. For these reasons, NARS can be augmented with additional features in various ways.

Dr. Nick Cassimatis, Rensselaer Polytechnic Institute

Topic: **A Cognitive Substrate for Human-Level Intelligence**

Making progress towards modeling and replicating human-level artificial intelligence seems to require a large number of difficult-to-integrate computational methods and enormous amounts of knowledge about the world. I provide evidence from linguistics, cognitive psychology and neuroscience for the Cognitive Substrate Hypothesis that a relatively small set of properly integrated data structures and algorithms can underlie the whole range of cognition required for human-level intelligence. Some computational principles (embodied in the Polyscheme cognitive architecture) are proposed to solve the integration problems involved in implementing such a substrate. A natural language syntactic parser that uses only the mechanisms of an infant physical reasoning model developed in Polyscheme demonstrates that a single cognitive substrate can underlie intelligent systems in superficially very dissimilar domains. This work suggests that identifying and implementing a cognitive substrate will accelerate progress towards human-level intelligence. More information, including papers and demonstrations, can be found at http://www.cassimatis.com

Steve Grand, Founder, Cyberlife Research Ltd

Topic: **Machines like us**

Computational functionalism and the paradigm of the digital computer have been entrenched and often unquestioned assumptions for most of the history of AI. At the same time, the brain sciences have taken their lead from AI, and so computational metaphors abound in neuroscience. Yet the mammalian brain is not at all like a computer. True, it can be likened to one in a metaphorical sense, but then it can be likened to many things, none of which is very helpful for trying to understand how it works. In my view, we will never be able to create artificial general intelligence until we understand how the mammalian brain works, and to do that we probably have to conceive of computational processes that are very unlike those we currently have theories for. They may not be all that profound, but they are probably hard to think about and even harder to deduce from the evidence alone.

Lucy is a long-term robotics project, targeted at trying to understand the brain by building one from the ground up. The aim is to take a creative, holistic, engineering approach, but without sidestepping the many difficult biological facts that normally get in the way of a neat theory. Many of these facts are almost certainly fundamental to the computational mechanisms at work, and should be embraced rather than ignored. By building a complete humanoid (strictly, anthropoid) robot I hope to find some of the underlying unities that explain many of the apparently disparate properties of intelligent and conscious action. So far the project is in its early stages, but some interesting common principles are already coming to light.

Dr. Ben Goertzel, Founder, CEO/CSO, Novamente LLC

Topic: **Novamente: A Practical Architecture for Artificial General Intelligenc**e

This talk will give a brief overview of the Novamente AGI architecture, which is described in more depth in the whitepapers online at www.novamente.net

The Novamente AI Engine (NAIE) is a C++ software system, currently under development, which is aimed directly at the "grand challenge" of creating robust Artificial General Intelligence at the human level and beyond.

The NAIE architecture involves a synthesis of original and familiar ideas -- including multi-agent systems, probabilistic inference, evolutionary learning, and reinforcement learning. It is grounded in a comprehensive, coherent theory of cognition based in patternist philosophy and complex systems theory.

Currently the NAIE is being developed along two parallel and interrelated courses:

- For the control of a humanoid agent within the AGI-SIM 3D simulation world
- To enable natural language processing within the Sagacity™ natural language question answering product

Both of these applications are considered as means to the end of powerful AGI, but also as important software applications in their own right.

In this talk, I will overview the NAIE architecture and the AGI-SIM and Sagacity™ applications, and will also briefly discuss my current research in the area of "Novamente developmental psychology."

Session III: AGI and Applied Computational Linguistics (Sun. 9 AM - May 21)

Speakers address what current computational linguistics methods can do, and where they fall short due to their lack of general contextual understanding.

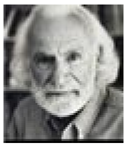

Dr. Karl H. Pribram, Distinguished Research Professor, George Mason University, School of Computational Sciences

Topic: **Neural Networks or Neural Webs?**

Neural networks are made up of circuits. In the nervous system such neuronal circuits consist of large nerve fibers connected by synapses. Nerve impulses are the the currency of communication in the networks. Before and after the connecting synapses the large fibers divide into fine fibers: teledendrons and dendrites. These fine fibers cannot sustain nerve impulses because the amplitude of an impulse is dependent on nerve fiber diameter. Rather, communication is achieved by way of oscillating hyper- and de-polarization. Within the dendritic arborizations nodes of activity are formed, nodes that have been shown to be sensitive to change with successive inputs, that is with "learning". The large fibered circuits sample these nodes -- there is not a one to one corresponence between node and large fiber (axon). Communication in the nervous system can thus be described at two levels: a neural circuit level and a neuro-nodal level. My claim is that the models developed in the study of artificial neural networks actually model a web characteristic of the neuro-nodal level of activity in the real nevous system.

Michael Ross, Science Applications International Corporation

Topic: **Mining Semantic Information From Unstructured Data**

MICCE is a general purpose algorithm, based on Clustering-By-Committee (CBC), for mining semantic data. The algorithm discovers semantic classes by grouping together items with similar features or which appear in similar contexts. Instances of ambiguous items can then be assigned to classes. MICCE is designed to work in any domain which contains structural and classification ambiguities. With domain-specific plugins, it can be applied to linguistic, visual, auditory, spatial, or mixed data. In addition, MICCE can be applied to the results of other algorithms, and make use of external ontologies. This flexibility allows it to use its own classifications or other new information in a feedback loop which can iteratively improve results.

As an example of the algorithm's behavior in a text domain, MICCE may discover two semantic classes for the noun 'plant.' These classes correspond with different meanings or senses of 'plant,' and are represented by sets of words which occur in similar contexts, i.e., {plant, flower, shrub, tree} and {plant, factory, mill, warehouse}. The senses are associated with a representations of the contexts in which they occur. Once these senses are discovered, occurrences of the word 'plant' can be disambiguated. MICCE compares an occurrence's context with each candidate sense's context representation. The occurrence can then be assigned to the best candidate. Thus, the documents (and individual words in the documents) may be tagged with semantic data. Furthermore, because MICCE can handle multiple types of context representation, the disambiguation data can be fed back into subsequent iterations of the algorithm. For instance, if an occurrence of 'plant' is disambiguated to the sense {plant, flower, shrub, tree}, this provides a more detailed description of the context for nearby words such as 'soil' or 'garden.' Such words can then be more accurately disambiguated, and syntactic parse ambiguities involving these words may be more accurately resolved. In non-text domains, such as visual/spatial data, the algorithm could similarly discover classes of objects based on their contexts and/or features, and assign instances to classes. Thus, although MICCE is currently being tested only on sequences of words, the algorithm can be understood more generally as a method for bootstrapping high-level representations from low-level data in any domain with structural and classification ambiguities.

At SAIC, MICCE is being developed to discover and disambiguate referents of names (determining whether "Michael Jordan" refers to the famous basketball player or to another individual with the same name). Interim results are presented along with a discussion of applications and future development directions.

Dr. Deborah Duong, Virginia Tech, Applied Research Laboratory for National and Homeland Security

***Topic:* Houses of Mirrors: Deeply Adaptive Designs for Machine Cognition**

Two deeply adaptive designs for cognition will be discussed: MICCE, the Mutual Information Contextual Concept Extractor, and Simulacra, an agent based coevolutionary simulation system. Both designs are "Houses of Mirrors" in which the definition of the each element of the system depends on the definition of the other elements of the system. This circular logic makes these systems deeply adaptive and flexible. They work through self fulfilling prophecy: if seeded with external elements, collocations of the elements come to exist that in turn allow the redefinition of the seed, an iterative process that continues until the seed develops into a necessary part of the system.

MICCE is a natural language processing software system, whose purpose is to extract semantic relationships from natural language text. In MICCE, the context dependence of linguistic relationships leads to feedback between various aspects of the interpretations of various relationships; and iterative feedback leads to a global consensus on the set of relationships constituting the meaning of a text. Each relationship is interpreted in the context of the other relationships (and the context of the text itself), and is in this way redefined – and this process continues iteratively until the process converges to a set of relationships coherently capturing the network of meanings implicit in the text. The operation of MICCE automatically leads to the emergence of ontologies that are data-driven rather than human-engineered like most ontologies. These ontologies are comprised of role-assignments of actors in systems, and relationships between roles in systems. At INSCOM, MICCE is applied to text related to national and homeland security, and the ontologies extracted are interpretable as social networks. Though the current implementation focuses on natural language, the MICCE architecture and many of the detailed algorithms are applicable to any kind of data comprising a large number of entities connected via a large number of relationships.

MICCE's iterative feedback process can also incorporate external information beyond the text or other data under analysis. For example, the external information could be a point of view with which to interpret the text, such as an intelligence analyst's ontology. The ontology need not be well developed: MICCE has the flexibility to adapt to it and fill in the gaps so that an ontology emerges within MICCE which is consonant with the external ontology. The emergent ontology is both data driven and hypothesis driven. The analyst's ontology adapts towards the text, and the text adapts towards it.

Another example of external information being incorporated into MICCE's feedback process is the use of information arising from visual signals to help MICCE interpret natural language texts related to those visual signals. MICCE can be applied to the visual data as well as the linguistic data, and points of correspondence can be formed between linguistic and nonlinguistic MICCE ontologies that ground the linguistic data in the non linguistic data and visa versa. Just as in the analyst's ontologies, there need only be very few explicit correspondence points for this to occur (for example, pointing at a ball and saying "ball").

Next, Simulacra is a design for a coevolutionary simulation system that incorporates several MICCEs interacting with each other and in effect grounding each other's relationships. A family of interacting MICCEs becomes an entire cognition system when Simulacra is added. The house of mirrors design pattern exists among the multiple MICCEs as well as within each MICCE: to stretch the metaphor a bit, we have in a effect a house of mirrors that are themselves houses of mirrors. Simulacra's architecture is based on the concept of a "language of thought" being grounded on simulation processes. A family of MICCEs, one linguistic and several corresponding to different non linguistic sensory modalities, can ground each other's ontologies, but in Simulacra the family of MICCEs is bound together yet more powerfully by a simulation process, which builds a model system capable of giving rise to the set of relationships detected by the MICCE family. Simulacra is designed to hold a model of process that MICCEs applied to different domains build up, fusing the data from the different domains. Simulacra interprets the roles and relationships extracted by MICCE as system specifications to guide the creation of a simulation model; and the running of this simulation model causes new relationships to be observed, which may then be fed back to MICCE for utilization in the internal (dynamically interrelating) feedback processes of the MICCE systems in the MICCE family. In order to achieve this feedback between simulation and relationship-extraction, both Simulacra and MICCE have been explicitly designed to interoperate.

Dr. Alexei V Samsonovich, George Mason University at Fairfax VA

***Topic:* Self-organizing linguistic cognitive maps as a key to the human value and semantic memory systems**

Self-organizing linguistic cognitive maps as a key to the human value and semantic memory systems. Alexei V. Samsonovich[1] and Giorgio A. Ascoli[2]. [1]Krasnow Institute for Advanced Study, George Mason University, Fairfax, VA 22030-4444. [2]Krasnow Institute for Advanced Study and Psychology Department, George Mason University, Fairfax, VA 22030-4444.

The concept of a cognitive map (O'Keefe & Nadel, 1978) is currently broadly understood as an abstract, multi-dimensional memory indexing space, presumably implemented in the medial temporal lobe. Despite the long history, the dimensions, intrinsic metrics and functional properties of cognitive maps are poorly understood. Given that direct access (e.g. Quiroga et al., Nature 435:1102-7, 2005) to brain representations of concepts is difficult due to ethical and technological reasons, an alternative approach deserves consideration. The idea is to use material extracted over millennia from billions of human brains: natural language. The present study is based on the Microsoft Word thesaurus of English synonyms and antonyms, of which only the giant connected cluster of 8,000 concepts (including words and short phrases) was selected. The cluster has a structure close to a scale-free network. In this study, each concept was represented by an abstract particle randomly allocated in a compact multidimensional manifold. Forces of attraction between synonyms and repulsion between antonyms were introduced. Thus defined dynamical system underwent simulated annealing and dimensional reduction procedures. The resultant self-organized configuration was subject to principal component analysis. Principal findings: (1) The stable, self-organized distribution of concepts in the abstract space forms a cognitive map, in that spatial coordinates of concepts reflect their semantics. (2) The distribution is polarized into two broad, fuzzy clusters of nearly equal size, separated along the main principal component that captures the evaluative dimension: “positive vs. negative”. (3) Otherwise, the map is a quasi-continuum, in which at various scales relative locations of concepts correlate with their semantic differences. (4) Interestingly, the strongest principal components roughly correspond to Osgood’s (1957, 1969) dimensions of the “affective space”: evaluation, potency, and activity. The main take-home message is that our spatial-dynamic analysis of natural language may provide a key to understanding the human system of values and semantic memory.

Moreover, we predict similar results for episodic memories or any representation system where synonym and antonym relationships can be defined. Support: DARPA IPTO BICA Grant "An Integrated Self-aware Cognitive Architecture".

Dr. Thomas Rindflesch, National Library of Medicine

***Topic:* Knowledge-Based Semantic Interpretation for Summarizing Biomedical Text**

SemRep is a natural language processing system being developed to recover semantic predications from biomedical text using underspecified syntactic analysis and structured domain knowledge. The SPECIALIST Lexicon and a stochastic tagger support the identification of simple noun phrases and verb groups, while domain knowledge is provided by components of the Unified Medical Language System (UMLS). During interpretation, simple noun phrases functioning as referring expressions are mapped to concepts in the Metathesaurus, and syntactic phenomena that "indicate" semantic relations are mapped to predicates in the Semantic Network. Syntactic constraints on argument identification are controlled by an underspecified dependency grammar. SemRep serves as the basis for several ongoing research initiatives in biomedical information management, including automatic summarization in the semantic abstraction paradigm.

Predications generated by SemRep from a set of Medline citations on a specified disorder topic are summarized into a conceptual condensate of the input text. The method relies on four principles, which transform a list of predications into a summarized version: All predications are subject to frequency of occurrence constraints (saliency principle), while relevance and connectivity determine that both core predications on the topic and additional "useful" predications should be included. Finally, novelty excludes predications that the user likely already knows.

It is increasingly challenging for researchers, health professionals, and consumers to effectively exploit the extensive textual resources provided by the National Library of Medicine. We are developing a Web application that exploits automatic summarization in cooperation with PubMed to pinpoint the most relevant information in large document sets. The summary is displayed in graphical format with links to the text that generated it as well as connections to supporting structured knowledge such as the UMLS Metathesaurus, Entrez Gene, and the Genetics Home Reference.

Dr. Ben Goertzel & Izabela Lyon Freire Goertzel, Novamente LLC, BioMind LLC

***Topic:* Using Dependency Parsing and Probabilistic Inference to Extract Gene/Protein Interactions Implicit in the Combination of Multiple Biomedical Research Topics**

I will describe a prototype software system designed to infer relationships involving protein/gene interactions from research Topics in the domain of the molecular genetics of oncology. The architecture uses a natural language processing module to extract entities, dependencies and simple semantic relationships, and feeds these features into a sophisticated probabilistic reasoning system (part of the Novamente AI architecture) which discovers complex, high-level semantic relationships that combine simpler semantic relationships found in multiple articles.

Moshe Looks, Washington University in St. Louis and Novamente LLC

***Topic:* Contemporary Approaches to Symbol Grounding**

In order to operate effectively in a given context, an intelligent agent must be able to make strong, dimensionality-reducing predictions: expectations of particular sensations, as well as general classes of sensory patterns. Information processing in complex environments corresponds to throwing out large masses of irrelevant data. In other words, strong inductive bias in correspondence with an external reality obviates the need for any artificial grounding of symbols.

Pragmatically this is best achieved in an artificial-general-intelligence context via the careful combination of preprogrammed heuristics and biases and interactive learning with a persistent agent. To illustrate, I will focus on some of the to-date and anticipated results obtained with a new program-learning technique, probabilistic adaptive-representational evolution (PARE). These will include object recognition, and learning to play simple games such as fetch and tag. A primary focus for future work is integration with Novamente's Probabilistic Logic Network (PLN) reasoning system.

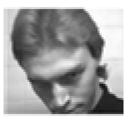

Ari Heljakka, GenMind Ltd. and Novamente LLC

***Topic:* Learning Symbol Groundings via Simulated Robotics**

Effectively interpreting and using prepositions has proved one of the biggest stumbling blocks for natural language processing systems. We present a novel approach to overcoming this problem via utilizing an AI system with a general intelligence architecture, and the ability to interpret prepositions in the context of their usage in a 3D virtual world as well as in purely linguistic contexts. This work is part of an overall project aimed at contructing artificial general intelligence with robust human-level language processing capability as well as other abilities.

Via physical or virtual embodiment, an AI can develop fundamental mental imagery sufficient for both concrete and Topic spatial reasoning. The AI Topics the underlying general cognitive constructs from its particular world-experiences by applying analogical probabilistic inference to its internal representations of experienced world phenomena. This enables grounding of the corresponding spatial natural language prepositions, and reasoning on relationships involving these prepositions.

We use an open-source 3-D virtual environment AGI-SIM for virtual embodiment of our AI, the Novamente AI Engine, which seamlessly combines evolutionary learning via a specialized Estimation of Distribution Algorithm with probabilistic inference via Probabilistic Logic Networks. We have designed a set of training sessions of increasing difficulty with the purpose of systematically teaching the AI more and more complex spatial relationships and enabling natural language interaction with humans, based on thus grounded prepositions.

Dr. William Sims Bainbridge, About The National Science Foundation (NSF)

The NSF is an independent federal agency created by Congress in 1950 "to promote the progress of science; to advance the national health, prosperity, and welfare; to secure the national defense…" With an annual budget of about $5.5 billion, we are the funding source for approximately 20 percent of all federally supported basic research conducted by America's colleges and universities. In many fields such as mathematics, computer science and the social sciences, NSF is the major source of federal backing.

Session IV: Panel Discussions (PM - May 21)

Discussions on particular issues related to AGI and grounding linguistic relationships in nonlinguistic reality:

Panel I 1:30 - 2:30 PM	**How are linguistic relationships grounded in nonlinguistic reality?**	Moshe Looks, Michael Ross, Dr. Thomas Rindflesch, Dr. Phil Goetz, (Mod: Dr. Ben Goertzel)
10 min	Break	
Panel II 2:40 - 3:40 PM	**How do we more greatly ensure responsible AGI?**	Eliezer Yudkowsky, Jeff Medina, Dr. Karl H. Pribram, Ari Heljakka (Mod: Stephan Vladimir Bugaj)
20 min	Break	
Panel III 4 - 6 PM	**What are the bottlenecks, and how soon to AGI?**	Dr. Stan Franklin, Dr. Hugo de Garis , Dr. Sam S. Adams , Dr. Eric Baum, Dr. Pei Wang, Dr. Nick Cassimatis, Steve Grand , Dr. Ben Goertzel (Mod: Dr. Phil Goetz)

Virtual Workshop (Ongoing)

Speakers describing their AGI projects are encouraged to complete the following questionnaire for publication to the Online AGIRI Workshop Forum:

www.agiri.org/workshop

- What is the end goal of the project?
- What milestones are expected to be achieved along the way?
- Roughly how much effort is expected to be required to bring the project to the point of roughly "human level" intelligence?
- What are the largest obstacles expected to be confronted along the path?
- Roughly how much computing resources are expected to be required?
- How is knowledge intended to be represented in the AGI system?
- How is learning intended to be accomplished in the AGI system?
- How is logical reasoning intended to be accomplished?
- How is perception intended to be accomplished?
- How is the control of motor systems intended to be accomplished?
- How is human language understanding intended to be accomplished?
- How is human language generation intended to be accomplished?
- How is self-modification of the system's own source code intended to be accomplished (if at all)?
- How, if at all, would one go about trying to impose a particular human-determined goal system on an instance of the AGI system?
- What (if any) practical applications do you think will be enabled by your system at an intermediary stage, before it has achieved human-level AGI?