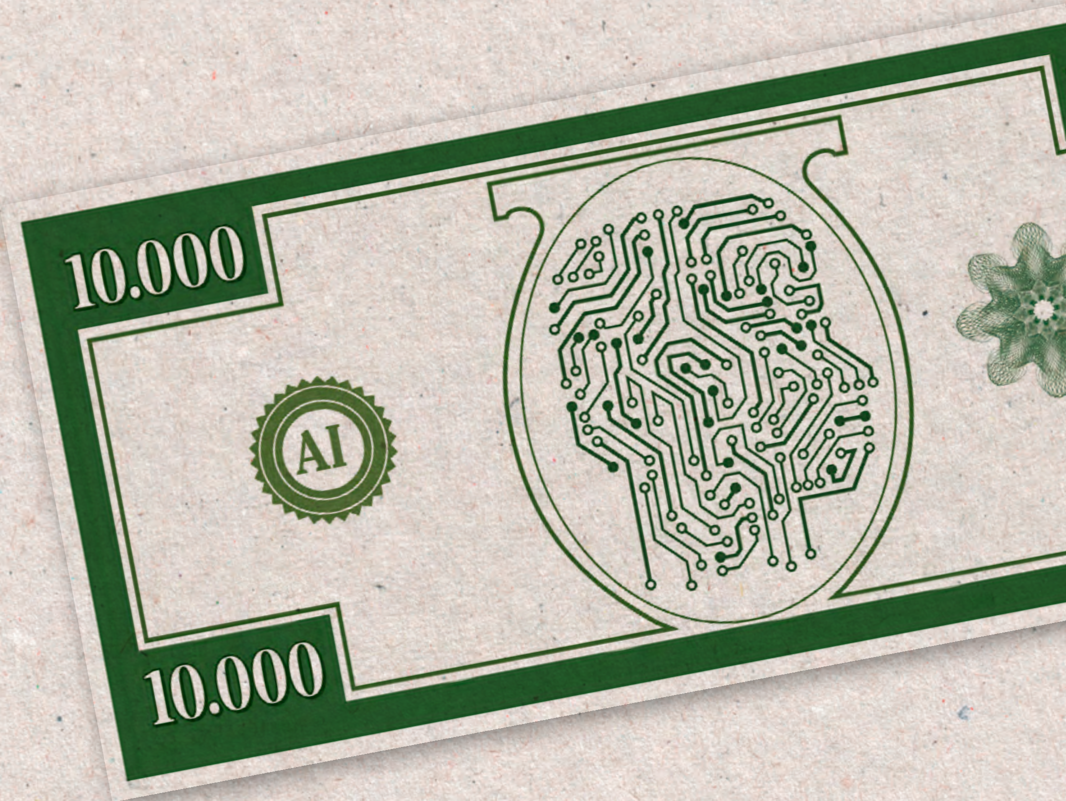

Newsletter

of the Icelandic Institute for Intelligent Machines, Reykjavik

Volume Six
Issue One, May 2017

COVER

Since the beginning of trade, people have been coming up with financial instruments. The money bill itself is a financial instrument created to store value and make the exchange of goods more convenient. As artificial intelligence and extensive automation invade our daily lives, they affect our economic system in various ways. Besides making new financial novelties, these technologies also can be used to help us understand better the very system in which they operate.

ILLUSTRATIONS

Regardless of whether we give much thought to the financial system in our daily lives, we are all interconnected through it and heavily dependent on its smooth operation. The global financial system is a world-spanning network, which creates a causal relationship between people in separate corners of the world. A disruption of the system can have cataclysmic consequences for ordinary citizens and spark social unrest. The illustrations in this issue depict demonstrations from the 2008 banking crisis in Iceland, reminding us of the fragile relationship of financial prosperity and peace.

Our industry collaborators include:

CM LABS Powering the AI Revolution™ · nox medical · Activity Stream · ECHORD++ · Technische Universität München · SVARMI · ALDIN · COSTNER · Ríkislögreglustjórinn · ÖSSUR · Mesher · 3Z · RÖGG

Our research partners include:

DFKI German Research Center for Artificial Intelligence · IDSIA Istituto Dalle Molle di studi sull'intelligenza artificiale · Reykjavik University · CADIA · rannís · Háskóli Íslands

DR. KRISTINN R. THÓRISSON

FROM THE DIRECTOR

WHAT PSYCHOANALYSIS & PHYSICS CAN TEACH ECONOMICS

In the 19th and beginning of the 20th century psychoanalysis became the ingenious theory that promised to make sense of the perplexing, often incongruous, and seemingly impenetrable nature of the human mind. It brought apparent structure and high-level organization to all topics mental, and it was framed in the rigorous language of medicine.

A scientific theory must refer to measurable and quantifiable factors. The litmus test for a scientific theory is asking it to make predictions that can be verified by experimentation. But there is something else that any good scientific theory must do: it must explain. And an explanation doesn't really explain anything unless it identifies relevant causal factors. Science is fundamentally about cause and effect in the real world – if factor A causes B, which in turn causes C, then tweaking A will change C – and A and C are clearly correlated. But remove B and no amount of tweaking A will have any effect on C, which demonstrates that correlation is not enough to explain how the world works.

Psychoanalysis was not a good scientific theory in part because it failed to identify the atomic elements at play in mental phenomena. Sure, it was indeed an insightful and creative theory, but it was the wrong kind of creative (the kind that accountants do not want to be caught doing). It spurred philosophers and scientists from Feynman to Feyerabend to contemplate how real scientific work can be distinguished form pseudo-science dressed up in the language of science.

The inability of psychoanalytic theory to reliably say for instance whether a person might cry or get angry when threatened was due to a lack of an "atomic theory of mind". In contrast, physics doesn't say that when metal heats it "may either expand or shrink". No, it can tell us the ways in which metal will respond because it has uncovered the actual causal relations in its substrate.

One problem with modern economic theory is its lack of an

"atomic theory" of the substrate on which it runs. Some might say economics is in essence impossible because "human minds cannot be modeled" (as a high-ranking director at a national bank recently tried to convince me of), or that "only mathematics can get to the fundamentals of economics". Neither is in fact true. Thinking that only mathematics can get to the core of economic mechanisms confuses *method* with *subject*: Math is a method. But fundamentally the subject of economics is not a mathematical phenomenon but a (man-made) *natural phenomenon*. Its essence lies in physics, through the intervening layers of human biology and the human mind. So, yes, economics is defined by human behavior, but the aspects of it relating to finance are no more irregular than many other things we already have good models for. And therein lies the second problem with modern eocnomic theory: Ignorance of causal relations. Actual causal relations must be part of any model if we are to trust it to tell us anything about the phenomenon it models.

This requirement certainly makes economics a non-trivial subject matter, but still one that is amenable to *non-mathematical modeling* (as demonstrated by the water-driven 2-meter high Moniac Computer exhibited at the Reserve Bank of New Zealand). We don't use water computers today – we have electronic ones. Modern computational models can help answer what-if questions that would be too complex or costly to do in the physical world. But for such what-if simulations to give answers about our economy they must capture actual causal relations found in the physical world: Otherwise a modification to factors A, B, or C may not mean what we think it means.

An Economics based fundamentally on actual and relevant causal relations of its substrate would give it a solid scientific foundation. With such a model we could produce empirically grounded answers to highly complex questions, instead of the current continual feeding of ungrounded speculation and emotion-driven argumentation.

IIIM's own Threadneedle project was motivated by the simple goal to accurately model the financial system at an "atomic level": double-entry bookkeeping. The Threadneedle framework allows experiments to be performed on simple and complex artificial economies. Threadneedle is not an end, it is a beginning: For those who want to take a scientific approach to economics, it is the start of a long journey, to a society with a deeper understanding of the fundamentals of economic systems. The code for Threadneedle is now open source for you to scrutinize on Github. So you can see for yourself how it works, and if you want to help move the needle towards a true science of economics, please join!

DR. JACKY MALLETT

Dr. Jacky Mallett graduated from MIT in 2005. She has over 2 decades' of industry experience, designing, building and troubleshooting distributed systems, with extensive experience in real-time critical systems, high performance computing, signal processing, and wide area networking. She returned to university and completed her Ph.D. at the MIT Media Lab in 2005 with work on the problem of co-ordinating large groups of autonomous cameras in real time. She became interested in the banking system during the 2007 credit crisis, and is currently working on building simulations of Basel regulated banking frameworks, with the goal of integrating an economic and distributed systems understanding of these frameworks' behavior, and their influence on macro- economic stability.

SIMULATING MODERN BANKING

Why are our economies hit by booms and depressions, catching virtually everyone by surprise? Could it be that our system is flawed in some way? Might it be that we lack the right tools to fix it? What is needed to better understand the behavior of the economy? What can we do to predict the future with more accuracy?

To try to shed light on those questions, IIIM research scientist Jacky Mallett has spent the last 4 years working on a banking and economic simulation system, which has the ability to build detailed simulations of banks with a mixture of financial instruments and regulatory controls.

The simulation software, which is called Threadneedle, is based on a double entry book-keeping engine which reproduces the same monetary transactions used in modern banking. Threadneedle is the world's first 'atomic level' simulation of the banking and financial system, and it can help us to understand how the underlying mechanisms of the monetary system influence the behavior of the whole economy.

As well as allowing the behavior of the banking system and its regulatory frameworks to be explored as an isolated system, Threadneedle can also be used to provide a key building block towards our ultimate goal of creating simulated economies with realistic financial systems that can be used to test regulatory and policy changes.

Threadneedle is consequently designed to provide a flexible building-block system into which long term simulations of the different components of the financial system (for example pension funds, housing funds, insurance, taxation policies, capital markets and derivatives, and foreign exchange relationships) can be progressively introduced, tested, and their impact on the underlying economic system thereby explored and understood.

> **"Today's purely mathematical models offer no reliable way to analyse the reasons for the observations they are based on."**

What does Threadneedle do differently to current economic models?

Today's economic models are purely mathematical models. They are built from mathematical formulas derived from observation but there is actually no proof that they represent how the economy in fact works. As they are purely mathematical they also offer no reliable way to analyse the reasons for the observations they are based on. One way to describe the issue is that a short enough ruler will fit any curve. Threadneedle allows us to look at the economy under a microscope. We can see the exact series of events – at the level of every individual doing what they do in their normal life, such as making loan repayments, getting a salary, buying things – and we can see how the combination of these actions from a lot of different agents leads to the phenomena called an economic crash or banking crisis.

How important is a better understanding of the system?

We, like many others, think it is critical that we do understand this system and learn to control it properly. There is a disturbing background music to history where financial crises precede devastating social crises such as revolutions and war. This was certainly what happened in the 20th century - with a direct causal chain leading from the Great Depression to the Nazi Dictatorship. We can't afford that in this century, we're going to have enough other problems as it is. But this understanding has to be public and democratic. As citizens, we all participate in our democracies, we contribute to the decisions made about them, and so we have to understand what the consequences of those decisions can be. It can't be something confined to a small elite that can't explain how it works to the rest of us.

RECENT PUBLICATIONS & TECH REPORTS

2015

Thórisson, K. R., Bieger, J., Schiffel, S., & Garrett, D. (2015). **Towards flexible task environments for comprehensive evaluation of artificial intelligent systems and automatic learners.** In J. Bieger, B. Goertzel, & A. Potapov (Eds.), *Proceedings of Artificial General Intelligence (AGI-15),* 187–196. Berlin, Germany: Springer-Verlag.

Perkin, S., Garrett, D., & Jenssen, P. (2015). **Optimal wind turbine selection methodology: A case-study for Búrfell, Iceland.** *Renewable Energy, 75*(3), 165–172.

Mallett, J. (2015). Threadneedle: **An experimental tool for the simulation and analysis of fractional reserve banking systems.** Article presented at the Eastern Economic Association's 41st Conference in 2015, New York City.

Mallett, J. (2015). **General disequilibrium: The hidden conflict between fractional reserve banking and economic theory.** *Cosmos+Taxis, Studies in Emergent Order and Organization, 2*(2), 18–33.

Bieger, J., Thórisson, K. R., & Wang, P. (2015). **Safe Baby AGI.** In J. Bieger, B. Goertzel, & A. Potapov (Eds.), *Proceedings of Artificial General Intelligence (AGI-15)* 46–49. Berlin, Germany: Springer-Verlag.

2016

Thórisson, K. R., Bieger, J., Thorarensen, T., Sigurdardottir, J. S., & Steunebrink, B. R. (2016). **Why artificial intelligence needs a task theory – and what it might look like?** In B. Steunebrink et al. (Eds.), *Proc. 9th International Conference on Artificial General Intelligence (AGI-16)*, July 16-19, New York City.

Steunebrink, B. R., Thórisson, K. R., Kremelberg, D. & Schmidhuber, J. (2016). **Growing Recursive Self-Improvers.** In B. Steunebrink et al. (Eds.), *Proc. 9th International Conference on Artificial General Intelligence (AGI-16)*, July 16–19, New York City.

Thórisson, K. R., Kremelberg, D., Steunebrink, B. R., & Nivel, E. (2016). **About Understanding.** In B. Steunebrink et al. (Eds.), *Proc. 9th International Conference on Artificial General Intelligence (AGI-16)*, July 16–19, New York City.

Pereira, G. (2016). **A Multi-Layer Non-Newtonian Model of Cardiovascular Inflammation.** *J Biomed Eng Med Device, 1*(3), 1–6.

Pereira, G. (in press). A multiscale haemorheological computer-based model of chronic inflammation: an in-depth investigation of erythrocytes-driven flow characteristics in atheroma development:The application of the threeIB method. In S. Malik et al. (Eds.), *Biotechnology and production of anti-cancer compounds.* Berlin: Springer-Verlag.

Pereira, G. (in press). Genomics and artificial intelligence working together in drugs discovery and repositioning: The advent of adaptive pharmacogenomics in glioblastoma and chronic arterial inflammation therapies. In S. Malik et al. (Eds.), *Biotechnology and Production of Anti-Cancer Compounds.* Berlin: Springer-Verlag.

VITVÉLASTOFNUN ÍSLANDS SES

Vitvélastofnun Íslands ses er sjálfseignarstofnun með það markmið að brúa bilið á milli iðnaðar og háskólarannsókna og að hraða nýsköpun í hátæknliiðnaði á Íslandi. Náið samstarf stofnunarinnar við Tölvunarfræðideild Háskólans í Reykjavík tryggir tengsl við fremstu vísindamenn landsins á helstu tæknisviðum, svo sem stærðfræði, fræðilegri tölvunarfræði, verkfræði og gervigreind.

Rannsóknir Vitvélastofnunar eru að miklu leyti knúnar áfram af þörfum iðnaðarins og niðurstöðurnar hafa nýtingarmöguleika á mörgum. Þar má nefna framleiðslu, tölvuleiki, þjálfun með aðstoð tölvutækni, lífupplýsingafræði, orkukerfi og stjórn vélmenna.

Vitvélastofnun leggur áherslu á að bæta gæði hugmynda og auka samskipti og flæði upplýsinga milli samstarfsaðila sinna. Markmiðið er að flýta fyrir árangri og hjálpa fyrirtækjum að sjá lengra inn í framtíðina, breikka sjóndeildarhringinn og auka möguleika þeirra á að koma hátæknivörum fyrr á markað.

IIIM THE ICELANDIC INSTITUTE FOR INTELLIGENT MACHINES

The Icelandic Institute for Intelligent Machines (IIIM) is a nonprofit research institute that catalyzes innovation through a focused exchange of ideas, people, projects, and intellectual property. Through close affiliation with Iceland's strongest technological academic department, Reykjavik's School of Computer Science, we bridge the gap between industrial engineering needs and academic research results.

Our work is driven by the needs of industry, and has relevance to a wide range of application areas. To name just a few: Computer-based training, bioinformatics, computer games, energy system, virtual and augmented realities, robotics, artificial intelligence, machine learning, and data manipulation, IIIM's software tools, methods, and systems help companies see further into the future, bring high technology to their product lines, and produces more advanced products faster.

CONTACT

IIIM is located on the 2nd floor of Reykjavik University's new Millennium building in Nautholsvik, within unique outdoor areas and near the country's only artificial beach.

Icelandic Institute for Intelligent Machines
Menntavegur 1, Venus, 2nd fl. IS-101 Reykjavik, Iceland

info@iiim.is
+354.552.1020 (voice)
+354.872.0026 (fax)

JORDI BIEGER, REYKJAVÍK UNIVERSITY

Jordi Bieger is a PhD student at the Center for Analysis and Design of Intelligent Agents at Reykjavík University (CADIA). He is originally from the Netherlands where he received a bachelor's and master's degree in artificial intelligence from Radboud University in Nijmegen. He came to Iceland in the Spring of 2013 to pursue his dream of studying artificial general intelligence. Before starting his PhD in the fall of that year under the supervision of Dr. Kristinn R. Thórisson, Jordi briefly worked at the Icelandic Institute for Intelligent Machines. Bieger is interested mainly in artificial general intelligence (AGI) and was involved in the organization of the 2015 conference on AGI (AGI-15) as co-chair of the program committee. His research focuses on the teaching of artificial learning systems in an area he calls Artificial Pedagogy (AP).

ARTIFICIAL PEDAGOGY – TEACHING AI TO LEARN ON ITS OWN

One of the things that separates man from machines is our ability to learn new things. So-called *narrow* AI systems are capable of outperforming humans in some specialized domains such as chess and poker, but they lack what we think of general intelligence, which among other things allows humans and animals to transfer their prior knowledge to new tasks and figure out how to complete a wide variety of novel and often complex plans in a wide variety of dynamically changing environments.

However, we do not do this alone. Human children are trained, educated, and raised to help them grow cognitively and gain the knowledge that allows them to perform the wide range of complex tasks that the world throws at them.

As far back as 1950, Alan Turing mentioned the idea of developing a "child machine" and teaching it, rather than trying to program an "adult" directly. Since then, however, research in AI has focused mainly on developing algorithms with different learning mechanisms and much less on how to teach artificial systems. Bieger believes that a shift in perspective from learner to teacher could help us develop new methods for making the most of current and future AI systems. He has spent the last three years working on and researching "Artificial Pedagogy" (AP), the science of how to teach artificial intelligence, with an emphasis on systems that aspire to reach or surpass human-level intelligence. He answered some of our questions about the bringing up of artificial systems.

What is the optimal teaching strategy for artificial systems?

This is essentially the central question of AP. The answer depends on many factors including the teacher, the learner, the task environment, and constraints on the teaching process such as budget. Possible strategies include providing rewards, teaching components of a task separately, and slowly increasing the difficulty of the task. We can also vary the interaction method from teaching by explanation to demonstrating the correct behavior or modifying the environment to facilitate learning (e.g., by removing obstacles or leaving clues).

Can't we just program AI systems to do what we want?

We can make software for many different tasks. To call these systems “artificial intelligence,” we typically require that they do something that previously only humans could. Many of these tasks come so naturally to us that we have trouble describing how we do them in the level of detail and precision that programming requires. Furthermore, an autonomous system might encounter tasks that its programmers did not anticipate.

Even if this were not the case, there is a limit to how much a team of scientists can feasibly program in an acceptable amount of time. It is likely that the complexity required to reach (super)human levels of general intelligence is well beyond this reach. Therefore, it seems to make sense to take the more holistic approach of having the system largely self-construct and self-organize through interaction with the environment after starting with only a small amount of “seed” knowledge or “bootstrap code.” Moreover, if we cannot program the necessary knowledge for performing at an “adult” level into an AI system, it will need to be acquired by the system itself: It must learn.

Why should we study Artificial Pedagogy?

We know from experience with humans, other animals, and AI that teaching can significantly enhance the learning process. Nevertheless, relatively little effort is being made to further our understanding of teaching, with the result that every research team must reinvent the wheel for any new combination of learner, task, and environment research.

“It is likely that the complexity and amount of information required to reach human intelligence in machines is well beyond what can be programmed by a team of experts. Systems should be able to learn on their own.”

There are now systems such as the Autocatalytic Endogenous Reflective Architecture (AERA)—developed in large part by IIIM’s Kristinn R. Thórisson and Eric Nivel as part of the HUMANOBS project—that have both the ambition and potential to operate in highly complex environments, as well as the capability to make use of the sophisticated teaching techniques required to learn this from relatively little prior knowledge. While it is unclear how far we are from developing true AGI, the demand for a theory of teaching will continue to grow as the field of AI advances and ever more systems are built that have the need and support for different kinds of teaching. It is therefore high time to start studying artificial pedagogy seriously.

IIIM ADVISED PASSING OF THE BILL

Althingi, the Icelandic Parliament, requested IIIM's comments on a Parliamentary Resolution for an international ban on the manufacture and use of autonomous weapons. IIIM sent a letter to Althingi in support of such a ban and offered its assistance in its execution. IIIM's director, Dr. Kristinn R. Thórisson, was subsequently invited to a meeting with the Foreign Affairs Committee on the subject.

K. R. Thórisson © 2009

BAN ON KILLER ROBOTS: ICELAND'S PARLIAMENT FIRST

On June 2nd 2016 the Icelandic Parliament, Althingi, became the world's first to pass a resolution for a support of a global ban of the manufacturing and usage of autonomous weapons - a.k.a. killer robots.

The mainly positive effects of great enhancements in the field of artificial intelligence and automation in the last decade are shadowed by the increasing emphasis on armament around the world. Research into the utilization of this technology for military purposes may lead to a high-tech arms race, which would be a threat to peace and prosperity and would increase the likelihood of misuse of power over the general public. An international ban would slow down and even hinder this development.

This cause has gained a significant following recently, especially from within the scientific community. Renowned scientists such as Stephen Hawking (as well as 20,000 others) have supported and signed a letter penned by the Future Life Institute calling for an international ban of killer machines. The UN Human Rights Council has also put forward a resolution for a ban of such weapons, which nine countries have already supported. In addition, the international organization ICRAC (International Committee for Robot Arms Control) has fought for this cause since 2009.

IIIM has been very vocal on these issues. By publishing its Ethics Policy for Peaceful R&D in 2015, IIIM became the first research lab in the world to ban any funding and collaborations whose results have the potential to harm citizens.

Iceland is in a good position to offer support to the ban due to its lack of military forces and an unusually peaceful history for all of its 1100 years. IIIM has offered Althingi its assistance on projects that might be initiated following the passing of the bill, such as defining what kind of automation the ban would apply to, how inspection could be conducted, and what Iceland's further involvement in this cause could be. IIIM's sister institution and Iceland's only academic AI research laboratory, CADIA at Reykjavik University, also sent a letter to Althingi in support of the resolution.

IIIM & CADIA'S AI FESTIVAL

The annual AI Festival, Gervigreindarhátíðin, was held on November 11th 2016 at Reykjavík University (RU). Over 150 attendees flocked to learn about the latest advancements and the progress that artificial intelligence has made this year.

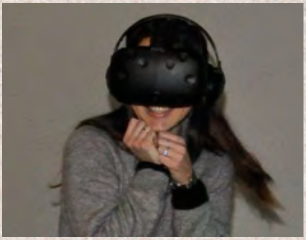

ARTIFICIAL INTELLIGENCE AND THE ECONOMY

The aim of the AI Festival is to raise awareness of AI and high-tech companies among both the general public and industry players, and at the same time, to support the cluster of high-tech companies in Iceland.

The focus of this year's Festival was how the invasion of artificial intelligence and extensive automation into our daily lives will affect the economic system. It also covered some ground breaking new approaches in economic research applying agent-based simulation.

The guest of honor was Prof. Dr. Doyne Farmer, a Professor at the Mathematical Institute at the University of Oxford and Director of the Complexity Economics program at the Institute for New Economic Thinking (INET). Dr. Farmer has advocated for reforms in economic research and theory and, amongst others, led a team of scientist who worked on an agent-based model of the economy and financial system funded by the European Union.

During this year's Festival, IIIM Senior Research Scientist Dr. Jacky Mallett introduced her economic software 'Threadneedle', which is a causal-relational agent-based simulation of the Icelandic banking system built to test the implications of different financial activity and regulations on banking, and its impact on the economy.

Other speakers included Dr. Magnús Torfason, Assistant Professor at the University of Iceland School of Business, as well as IIIM's Founding Director, Dr. Kristinn R. Thórisson, Professor at Reykjavik University School of Computer Science and CADIA's director, Dr. Jón Guðnason, Assistant Professor at RU's School of Science and Engineering.

The lectures were followed by a party and poster session where attendees had the opportunity to have a glimpse at some of the most advanced products, services and research in the field of AI, automation, and high-tech in Iceland.

The festival is a collaboration between IIIM and CADIA and was sponsored by Marel, a leading global provider of advanced food processing systems and services.