The Other Side: Algorithm as Ritual in Artificial Intelligence

Kieran Browne

Australian National University Canberra ACT 0200, Australia kieran.browne@anu.edu.au Ben Swift Australian National University Canberra ACT 0200, Australia ben.swift@anu.edu.au

Abstract

Our cultural and scientific understandings of neural networks are built on a set of philosophical ideas which might turn out to be superstitions. Drawing on methodologies of defamiliarisation and performance art which have been adopted by HCI, we present an analog apparatus for the ritualistic performance of neural network algorithms. The apparatus draws on the interaction modes of the Ouija board to provide a system which involves the user in the computation. By recontextualising neural computation, the work creates critical distance with which to examine the philosophical and cultural assumptions embedded in our conception of AI.

Author Keywords

Neural networks; Ritual; Séance; Defamiliarisation; Performance art.

ACM Classification Keywords

H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI'18 Extended Abstracts, April 21–26, 2018, Montreal, QC, Canada Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-5621-3/18/04...\$15.00. https://doi.org/10.1145/3170427.3188404 "It is impossible to give any date for the early appearances of external intelligent power of a higher or lower type impinging upon the affairs of men."

Arthur Conan Doyle, *The History of Spiritualism* Vol 1, 1926 [9]

"The Navy revealed the embryo of an electronic computer today that it expects will be able to walk, talk, see, write, reproduce itself and be conscious of its existence."

'NEW NAVY DEVICE LEARNS BY DOING,' *New York Times*, July 8, 1958, Page 25. [23]

"There's something magical about Recurrent Neural Networks".

Andrej Karpathy (Director of AI at Tesla), *The Unreasonable Effectiveness of Recurrent Neural Networks* (http://karpathy.github.io/2015/ 05/21/rnn-effectiveness/)

Introduction: superstition, science and AI

The history of science & technological development is more entwined with superstition than the narratives of postenlightenment rationalism would suggest. Many great scientific minds have engaged in mystic or superstitious inquiry; notably the alchemical convictions of Sir Isaac Newton [19], the Vedic mysticism of Nikola Tesla [15], and the mystic references in the analytic psychology of Carl Jung [8].

The historical development of artificial intelligence (AI) is similarly full of philosophical assumptions, both examined and unexamined (see Dreyfus 2007 [10]).

Notions of the digital as immaterial, intelligence as emergent, and neural networks as simulations of human brains abound. Many of the significant figures in the development of the digital computer—including Babbage, Lovelace, Turing, and Von Neumann—have shown interest in and contributed to the development of machine intelligence.

However, the separation between "reasonable" ideas like AI and "unreasonable" ideas which we might call superstition is less clear than one might expect. In *Computing machinery and intelligence*, Alan Turing considers the use of a "telepathy-proof room" to protect the integrity of his imitation game from players exhibiting extrasensory perception (see [22]). This may cause us to cringe in hindsight—it's uncomfortable to imagine heroes of science believing such unlikely things. But good science demands open-mindedness and the courage to challenge accepted truths (for another example, consider the apparent nonsensicality of quantum mechanics). Researchers are in a difficult position, expected to dismiss "silly" ideas like telepathy and yet take seriously the idea that bits of metal and silicon might become intelligent if you program them the right way. In this paper, we lean in to this tension between science and superstition to call into question the assumptions upon which AI is based. We present an interactive analog apparatus for the ritualistic computation of neural networks, performing the algorithm as a séance. This performance draws on the language of the modern spiritualist movement in ways we discuss shortly. We frame this neural network computation as a ritual, with the aim of communicating with emergent intelligence (rather than the spirits of the dead).

The apparatus is a type of analog computer, based on a slide rule but extruded radially and expanded to the size of a table. This allows for collaborative, interactive computation. The performance blurs historical and contemporary concepts of the computer; as a human vocation, as a physical analogy and as an abstract information processing machine. We attempt to make physical and observable the processes which are normally obscured both by the esoterica of computer code and their enaction as momentary and microscopic paths of charge through a computer chip far beyond the reach of human senses.

Computer programming has long held a mystique as a particularly indecipherable practice to the public. But deep neural networks represent a new frontier in academic esoterica and obscurantism beguiling expert and non-expert alike.

This is hardly the first paper to point out the difficulty of defining "machine intelligence" [17], or the gulf (both in form and function) between the biological neuron and the perceptrons at the heart of an artificial neural network, or the problems with our lack of understanding of how modern machine learning and AI techniques actually *work* [16]. Furthermore, the motives of this paper are not cynical; we do not wish to mock or to criticise scientists and engineers, although we are fully aware of the negative connotations of superstition among these groups. Our position can be sum-

marised as such; any or none of the philosophical assumptions underlying AI may turn out to be true, but they are significantly underexamined. This paper adopts a methodology of defamiliarisation via performance art. In doing so, we create a space for critical engagement with the strangeness at the core of AI and provoke a reexamination of the assumptions inherent in neural networks as AI.

Neural computing as AI: the unreasonable effectiveness of neural networks

Comparisons to magic are common in the discussions of neural networks, for example Jitendra Malik, as quoted in the Nature feature article *The Learning Machines*: "Neural nets were always a delicate art to manage. There is some black magic involved" [14, p147]. Of course, most researchers don't actually believe that neural networks *are* magic, instead it is acknowledged that they are more effective than seems reasonable and we can't really explain why. This might be shrugged off as a case of Clarke's third law, "any sufficiently advanced technology is indistinguishable from magic" [7], but we will argue that there is more at the heart of this comparison.

In recent times, neural network techniques under the brand, "deep learning," (see [21] for an overview) have been extremely successful in a range of machine learning/AI endeavours. Indeed, when the MIT Technology Review asks "Who is winning the AI race?", they answer it by measuring the amount of Deep Learning research papers produced by various institutions [1]. These neural networks algorithms seem to intensify the superstitious notions in AI. Unlike expert systems and symbolic AI, neural networks are often indecipherable even to those who engineer them. There is plenty of air-time given to the neural network's black-box problem, but this is primarily interpreted as a barrier to the adoption of the technology rather than as a flaw in the technology itself. Part of the mysticism of neural networks is that they offer answers but not explanations. And for the most part, we don't care—so long as the algorithm performs well on the test set. We consult neural networks like the ancients consulted oracles. The process may be esoteric and opaque, but if it gets results that's good enough.

Al is a chimera. It lives on even as its underlying technologies lose favour or prove untenable. Expert systems and neural networks have both at various times marched under the banner of Al. Though it has existed in many forms, the ideas and questions at the core of Al remain the same. Questions of agency, singularity, apocalypse are well established in public mind through both science fiction and the breathless reporting on the promise (and dangers!) of the latest breakthroughs in Al in the wider media (see, for example, the reporting on Google's 2017 victory in the game of Go [12]).

The language used to speak about AI is also well established. Assertions that a machine "learned", "discovered", "outsmarted" presuppose agency and often imply consciousness. Even treating the machine as something that deserves to occupy the subject of a sentence deserves examination.

At its core, AI is built upon an ontology which separates mind from material reality. This can be seen in Norbert Wiener's assertion that "information is information, not matter or energy" [24, p155]. Wiener and other cyberneticists characterise intelligence as a pattern of information processing separated entirely from its physical embodiment. The digital is taken to be immaterial and otherworldly. This is reflected in the physical/virtual dichotomy, and appears again in the metaphor of "the cloud", which characterises the internet as an etherial stream of information raining down from the heavens and not an acre of coal-fired servers in Bangalore or tendrils of copper cables snaking under the pacific, occasionally attacked by sharks.

Methodology

This paper employs defamiliarisation via performance art to critically examine neural computing as AI.

Defamiliarisation originated in literary theory and has been adopted as a methodology by HCI researchers [3]. Concepts with which we are overly familiar come to seem natural, obvious and unquestionable. It's easy to forget that ideas emerge in particular historical contexts which shape their development. Defamiliarisation is a technique for creating critical distance between an audience and a concept. By presenting familiar concepts in strange and unfamiliar ways, viewers are asked to reconsider their ideas. We propose to borrow the strangeness of spiritualist ritual in order to defamiliarise the neural network as AI and question the its assumptions therein.

On Ritual

There are many competing theories of ritual from which a definition could be derived. Most obvious examples of ritual come from the canon of rites of a given culture [2], however for the purposes of this paper we wish to invoke a particular sense of ritual as applicable to the modern spiritualist movement.

Modern spiritualism emerged in the mid 19th century and gathered surprising momentum, particularly in the United States [18]. The movement seems at first to be something of a paradox, arising at a time supposedly dominated by empirical and scientific thought, but crucially, spiritualism attempted to prove itself within the bounds of empiricism; many leading spiritualists believing that science was the only path to knowledge [18, p476-77].

Spiritualism quickly developed its own set of rituals through which spirits could be channeled. Primary amongst these is the séance, which took many forms but was often performed in dimly lit parlours by touring mediums for payment. Spiritualism's attempt to prove the existence of an immaterial spirit world within a scientific framework is rather bizarre and compelling. Mediums would invite skeptics to their séances, challenging them to explain the phenomena without recourse to spirits. Of course, many mediums were eventually shown to be fraudulent, notably the Fox sisters who are often used to mark the beginning of the movement [18].

Modern spiritualism is an interesting case study in the cohabitation of scientific and superstitious thinking.

With the proposed performance, we borrow the strangeness of the spiritualist séance to defamiliarise the neural network. We present an analogy between the neural network as AI and the ritual of the modern spiritualist movement. Both frame the mind as a metaphysical entity outside the physical world. For the spiritualists, who inherited Descartes' dualism of mind and body, these entities were the souls of the dead living on in the spirit realm. For contemporary AI, which has inherited immaterial notions of intelligence from cybernetics, this mind is emergent from patterns of information processing separate from material reality. The algorithm, too, can be related to occult ritual where the right combination of words chanted, or the right set of symbols scrawled on parchment can bridge the gap between this world and the next.

The séance

We propose to perform a séance *live* at alt.chi, to collectively (with members of the alt.chi audience) perform the mathematics of a neural network in calculating the "exclusive OR" (XOR) function in the ritual and "enchanted space" [13] of the séance. In this new context those present are forced to consider the nature of the neural network's intelligence. If the network is intelligent where does this intelligence reside? In the users? In the apparatus? In the immaterial pattern with which information is organised by the ritual?

Here we will outline the steps of the séance before expanding on specifics of the computing device (hereafter referred to as "the apparatus").

The apparatus, covered in markings and divided in concentric discs sits in the centre of a large round table. Presenters enter dressed in academic robes. One carries a large book and a string of wooden beads to a lectern, the other takes a seat at the table. At the lectern, the presenter welcomes the audience and explains that they will be seeking to communicate with an immaterial intelligence. Attendees are offered the places at the table. When these are filled, the remainder observe from the audience. Reading aloud from the book, the presenter instructs those seated to place their hands on the outmost disc of the apparatus. The presenter lifts the string of beads to their forehead briefly and flicks a few beads back or forward; then calls out the increment to which the outmost disc must be shifted. Those seated rotate the disc. Next the presenter announces the increment to which the second outmost disc must be shifted and flicks beads back and forth. Those seated rotate the disc. The presenter asks of the outmost disc "what does it read?" Those seated respond. The presenter announces the increment to which the third outmost disc must be shifted. Those seated move the disc. The presenter instructs those seated to hold the two inner most discs in alignment and shift these so that the third outmost disc returns to point at 0. Those seated move the discs.

The presenter asks of the inmost disc "what does it read?" Those seated respond. This process is repeated for every neuron in the network. Upon reaching the final calculation, the presenter asks of those seated; "what does it read?" The presenter declares that the intelligence has answered our query and thanks it. The presenter then, with reference to the language of the spiritualists, gives closing remarks.

The apparatus

Our vehicle for defamiliarisation in this paper is ritualistic performance which centres around an analog computer. This apparatus is designed after a Ouija board, a game closely associated with occult ritual due to appearances in popular culture, notably the 1973 film The Exorcist. The Ouija board is an oracle-it provides answers to questions. Usually multiple persons place their hands on a cursor at a time. Someone asks a question then the cursor, seemingly of its own volition (or that of a spirit), moves across the board to spell out the answer. The design of the Ouija board exploits ideomotor actions; ie. behaviours that are unconsciously initiated [11]. This is a likely explanation for the imagined agency in the Ouija board. Researchers Gauchou et al. [11] showed that the ideomotor interactions of the Ouija board allow users to express nonconscious knowledge, answering questions that they believed they did not know the answer to with significantly better results. This trick of psychology provides a useful metaphor for the attribution of agency and sometimes blame to the neural network. The revelation that machine learning algorithms trained on human data learn human-like biases [5] should have come as no surprise. We first animate a neural network with human data then gasp when it reflects our biases back at us.

Functionally, the apparatus works like a slide rule, using a physical representation of mathematical functions to calcu-



Figure 1: © Kieran Browne 2018. The computational apparatus works much like a slide rule, but is extruded radially. It is large enough that a number of people may use it at a time.

late results. Bivariate functions like multiplication and addition are calculated by moving concentric discs into correct alignment, while univariate functions like relu do not require alignment.

Unlike a conventional slide rule, the apparatus is extruded radially and scaled up to the size of a table so as to invite group interaction like the Ouija board.

The apparatus works because neural network forward propagation only requires a small set of operations; multiplication, summation and activation.

The two outmost discs enable multiplication. The ticks are distributed logarithmically, this in essence converts addition to multiplication. For example, to multiply 2 by 3 first slide the outmost disc so that it's one aligns with the adjacent inner disc's "2" then find the "3" on the outer disc and read the inner disc's value. This interaction will be second nature to anyone who went to school in the 70s. The slide rule deals in values normalised to between 1 and 10, thus it is easiest to think of the input values in normalised scientific notation and treat zero as a special case. The apparatus like other slide rules does not place the decimal. This is calculated by the presenter using a string of wooden beads. These reference rosary or prayer beads in appearance and interaction but are without identifying iconography. They allow the presenter to keep track of the decimal point. This is done as follows;

The presenter begins holding the central, black bead which represents an exponent of zero. For both operands; if \geq 10 take one bead for every numeral except the first before the decimal point, else if < 1 drop one bead for every zero before the first significant figure. Finally, if the result from the slide rule is \geq 10 take one more bead.

The two inmost discs allow for summation and activation. Summation is normally not a part of the slide rule as addition is fairly simple to do mentally, but we have included it as it is a fundamental part of neural network mathematics. Beginning with both discs indicating zero; align the outer disc so that the number to be added points to the triangular indicator. Next, holding both discs in alignment, rotate the outmost disc so that its zero again aligns with the indicator. This process can be repeated any number of times after which the result is designated by the number on the inner disc aligned with the inner indicator. Note that for the inner disc there is a trail of zeros on the left, this represents the relu activation function.

Like the slide rule, the apparatus is a form of analog computer. Analog computing-so called due to the analogical relationship to the phenomena being computed—has a long history which predates and coincides with the digital computer. Analog computers model one physical or mathematical system with another analogous physical system [6]. Although analog computing and digital computing is in some sense multiple realisations of the same computing "process", analog computation is not so much an automatic device as an "interactive visualisation". The concern of analog computing is on cognitive support rather than automation of cognitive activity. Because of this, instead of distancing the user from the computation, it draws them in [6]. Ritual places the same importance on action—the point of ritual is to partake, to make it happen. In this apparatus, the point is to involve the user in the work of the neural network. In the séance, the apparatus supports (rather than supplants) the cognition of the participants.

This proposed ritual has echoes of the notion of "computeras-job-description" which arose in the 1930s. Here, the "computer" was someone who would perform simple mechanical calculations under the supervision of a mathematician. Computing was thus lesser than mathematics, requiring neither deep insight nor complex knowledge of the calculations being performed. The participants in the séance are in this sense human computers, carrying out the mechanical ritual which they need not understand.

XOR

The choice of the "exclusive or" (XOR) function is a deliberate one. It is a function with specific historical significance in the history of AI and artificial neural network algorithms in particular. It was central to the historical rivalry between symbolic AI and neural computing in the 50s and 60s. The controversy around the limitations of neural networks and the subsequent disputed history is examined at length by Mikel Olazaran [20]. In 1969 Marvin Minsky and Seymour Papert published Perceptrons which included a mathematical proof that the artificial neuron was incapable of computing the XOR function, and this seemed to imply that neural computing itself was a deadend, although this is contested. The fall out from the book's publication is often claimed to have caused the subsequent AI winter [25, p14-15]; [4, p74]. The thesis in *Perceptrons* is something of a straw man. The Minsky and Papert first define perceptrons sufficiently narrowly, and contra to the typical perceptron as used by those in the field, then prove its insufficiencies. [20, p631-634]. Although single perceptrons are incapable of computing XOR, even small multi layer networks are able to. For this reason, we consider this to be the simplest true neural network problem.

Conclusion

This paper describes an apparatus and ritual for performing neural network computation as a séance, with the goal of foregrounding the underexamined philosophical basis of AI through a process of defamiliarisation.

REFERENCES

- 1. 2017. Who Is Winning the AI Race? (27 June 2017). https://www.technologyreview.com/s/608112/ who-is-winning-the-ai-race/
- 2. Catherine M Bell. 1997. *Ritual: Perspectives and dimensions*. Oxford University Press on Demand.
- 3. Genevieve Bell, Mark Blythe, and Phoebe Sengers. 2005. Making by making strange. *ACM Transactions on Computer-Human Interaction* 12, 2 (2005), 149–173. DOI:http://dx.doi.org/10.1145/1067860.1067862
- John Mark Bishop. 2015. Computational Intelligence. Eolss Publishers, Chapter History and Philosophy of Neural Networks, 22–96.
- 5. Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356, 6334 (2017), 183–186.
- Charles Philip Care. 2008. From analogy-making to modelling: the history of analog computing as a modelling technology. Ph.D. Dissertation. University of Warwick Department of Computer Science, Coventry, CV4 7AL, UK.
- Arthur C. Clarke. 1999/2013. *Profiles of the Future*. Orion, London, UK, Chapter Hazards of Prophecy: The Failure of Imagination, 25–33. 'Third law' appears in endnotes on p. 233.
- 8. Harold Coward. 1979. Mysticism in the analytical psychology of Carl Jung and the yoga psychology of Patanjali: a comparative study. *Philosophy East and West* 29, 3 (1979), 323–336.
- 9. Arthur Conan Doyle. 1926. *The history of spiritualism*. Vol. 1. The Spiritual Truth Press.

- Hubert L Dreyfus. 2007. Why Heideggerian AI failed and how fixing it would require making it more Heideggerian. *Artificial Intelligence* 171, 18 (2007), 1137–1160.
- Hélène L Gauchou, Ronald A Rensink, and Sidney Fels. 2012. Expression of nonconscious knowledge via ideomotor actions. *Consciousness and cognition* 21, 2 (2012), 976–982.
- Elizabeth Gibney. 2016. Google AI algorithm masters ancient game of Go. *Nature News* 529, 7587 (Jan. 2016), 445. DOI:http://dx.doi.org/10.1038/529445a
- Julian Holloway. 2006. Enchanted Spaces: The Séance, Affect, and Geographies of Religion. Annals of the Association of American Geographers 96, 1 (March 2006), 182–187. DOI:http: //dx.doi.org/10.1111/j.1467-8306.2006.00507.x
- 14. Nicola Jones. 2014. The learning machines. *Nature* 505, 7482 (2014), 146.
- 15. Thomas Lee Kelley. 1997. *The enigma of Nikola Tesla: a cultural studies analysis of his legacy.* Master's thesis. Arizona State University.
- 16. Will Knight. 2017. The Dark Secret at the Heart of Al. *TECHNOLOGY REVIEW* 120, 3 (2017), 54–61.
- 17. John McCarthy and Patrick J. Hayes. 1969. Some philosophical problems from the standpoint of artificial intelligence. *Readings in artificial intelligence* (1969), 431–450.
- R Laurence Moore. 1972. Spiritualism and Science: Reflections on the First Decade of the Spirit Rappings. *American quarterly* 24, 4 (1972), 474–500.

- 19. William R Newman. 2015. The problem of alchemy. *The New Atlantis* 44 (2015), 65–75.
- 20. Mikel Olazaran. 1996. A Sociological Study of the Official History of the Perceptrons Controversy. Social Studies of Science 26, 3 (1996), 611–659. http://www.jstor.org/stable/285702
- 21. Jürgen Schmidhuber. 2015. Deep learning in neural networks: An overview. *Neural networks* 61 (2015), 85–117.
- 22. Alan M. Turing. 2009. Computing machinery and intelligence. *Parsing the Turing Test* (2009), 23–65.
- 23. Unknown. 1958. NEW NAVY DEVICE LEARNS BY DOING; Psychologist Shows Embryo of Computer

Designed to Read and Grow Wiser. The New York Times. (8 July 1958). Retrieved January 13, 2018 from http://www.nytimes.com/1958/07/08/archives/ new-navy-device-learns-by-doing-psychologist-shows-embryo-of. html.

- 24. Norbert Wiener. 1948. *Cybernetics: Control and communication in the animal and the machine*. Wiley New York.
- 25. Neha Yadav, Anupam Yadav, and Manoj Kumar. 2015. An Introduction to Neural Network Methods for Differential Equations. Springer, Chapter History of Neural Networks, 13–15.

Commentary

For alt.chi paper The Other Side: algorithm as ritual in artificial intelligence

Michael A. DeVito

Northwestern University 2240 Campus Drive Evanston, IL 60208, USA devitom@u.northwestern.edu The familiarization of terms such as "algorithm" or "machine learning system" was inevitable. For us, HCI researchers, educators, and practitioners, this resulted in what Gillespie calls "algorithm as synecdoche" – a simplified representation with which to discuss larger issues without the need for specific expertise in algorithmic formalisms. For the user, this resulted in what Gillespie calls "algorithm as talisman" – the algorithm as a powerful, objective, somewhat unknowable, authoritative decision-making force. These understandings are not reality; they are conveniences, but conveniences that rob us of our critical distance.

For us, those that treat "algorithm" and "machine learning" as synecdoche, this theatrical defamiliarization via a mystically-tinged physicalization is a way to get that critical distance. We treat algorithms, particularly machine learning algorithms, almost as mystical constructs. By taking this to an absurd extreme we cannot help but reconsider this treatment. Despite the mystical atmosphere, use of the device as described reveals a core truth: there's nothing magical about neural networks, just a whole hell of a lot of clever math. This is useful for all of us to remember, especially those of us that don't, in the course of our research, ever actually interact with the math. Here, the mystic can demystify and in turn serve to educate. This setup will likely make some uncomfortable, and puzzle others; both of these outcomes are good, as we need prompting to question the complexities of systems we now so casually rely on. This exact dynamic makes the séance an excellent contrast and companion to the many papers on algorithms in this

year's main papers track. Those authors, this commenter included, need reminders such as this.

Perhaps more important, though, is how this séance can give us the distance to address our own curse of knowledge. To the average end user, most algorithms, let alone something as complex as a neural network, may as well be a Quija board. End users have fragmentary, mostly-wrong folk theories of how these complex systems work. It is essential for designers and researchers to remember that. For all the distance in understanding between the parts of the CHI community that are debuting novel machine learning algorithms and those of us who focus primarily on the human/social components of HCI, the gap between all of us and the end user is larger. Within this gap lies the potential for, even unintentionally, robbing users of agency and awareness of key processes that affect their world. This séance might help us remember that, to the user, algorithms are essentially a very, very convenient set of spells with little mechanism for critical interrogation or the assertion of user agency.

The spiritualists of the 19th century took advantage of the "black box" of unexplained phenomena to profit off keeping people in the dark. Where people sought answers, the spiritualists chose to obfuscate. This séance functions to point out that we, the modern mediums and translators of complex, difficult-toexplain phenomena, have our own choice in front of us. We must decide: when this particular veil is parted, will users see a field prepared to help, educate, and advocate, or just more obfuscating charlatans selling convenient, feel-good explanations of complex computational phenomena?