


THE AGE OF
**ARTIFICIAL
INTELLIGENCE**

what does it mean to be human?

By Alison Main



Cogito Ergo Sum

"I think, therefore I am."

This famous quote by Rene Descartes, a 17th century French philosopher and mathematician, was originally published in his 1637 treatise *Discourse on the Method of Rightly Conducting One's Reason and of Seeking Truth in the Sciences*.

At least, that's what your brain's neural network is telling you, based on a five-word-sentence data input. You get a gold star for intelligence.

But what happens when you feed that same five-word sentence to a computer, and it accurately tells you the author, publication date, and book title based on an algorithm it was programmed to follow? Would we conclude that the machine can "think?" And if it *can* think and solve problems, does that machine exist as a sentient, intelligent being? And does the existence of artificial intelligence consequently place a neutralized value on humanity?

These are big questions without straightforward answers. But they're questions we need to be aware of and address mindfully as we usher in the age of A.I. Read some books on cognitive technology, listen to podcasts on machine learning, have enlightening conversations with digital futurists, and suddenly Netflix's *Black Mirror* (a series that explores the dark side of cutting-edge technology) makes perfect sense, and you're 79.6 percent certain that we're living in the Matrix. It's time to wake up, Neo . . .

OUR AUGMENTED FUTURE

Inventor and futurist Ray Kurzweil said, "Perfecting our biology will only get us so far. The reality is that biology will never be able to match what we will be capable of engineering."

Artificial intelligence in daily life has been little more than a technologist's moonshot pipe dream . . . until now. The 2018 Consumer Electronics Show (CES) in Las Vegas showcased an astonishing array of intelligent machines all set to emerge from the engineer's lab and into homes, roadways, schools, transit systems, offices, and public spaces around the world.

With regard to humanitarian goals, A.I. promises to save the world, offering solutions to looming global challenges and mental and physical shortcomings that have plagued humanity since inception, including disease, pollution, poverty, crime, and, ultimately, death. Silicon Valley's tech evangelists say A.I. will foster connections in ways we've yet to imagine, while making us healthier and saving lives. That sounds wonderfully utopian, an attainable reality. But what about the possible dystopian end to this scenario? Not to mention every scenario in between?

The fact is, no one knows exactly where A.I. will take us. Both positive and negative, there are infinite possible outcomes. A few likely permutations play out as follows:

[1] A peaceful and thriving coexistence between man and "friendly A.I.," with a universal basic income for all citizens, the eradication of pain and suffering, and solutions for environmental threats like climate change. In sum, an extended existence of blissful leisure, well-being, and abundance for all.

[2] The evolution of the human species through transhumanism, involving a mass-scale abandonment of our biological bodies, living forever through radically expanded cognitive ability as digital uploads of our consciousness to an artificial neural network.

[3] The unintentional annihilation of humanity by a superintelligent A.I. that had absolutely no malicious intent, but nonetheless continued to advance its goal system to the ultimate extinction of the human race (see "Paperclip Apocalypse" sidebar for one example).

Even the term "artificial intelligence" may be misleading, creating a false sense of security. Lisa Messeri, Assistant Professor of Sociocultural Anthropology at Yale University, says "One of the potential problems [with the moniker] is that it allows for the illusion that A.I. is something humans are in control of—because it's artificial. It

presumes A.I. is something less than human intelligence. Is there something even in the naming of A.I. that gives us a sense of superiority over the technology that we don't really have?"

ORIGINS

Let's dispel one myth: At this exact moment in history, A.I. is neither new nor superhuman. At a convocation of minds in the summer of 1956, Dartmouth professor John McCarthy coined the phrase by proposing "a two-month 10-man study of artificial intelligence to be carried out . . . to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."¹

In 1950, computer scientist Alan Turing designed a test to determine whether a machine can exhibit intelligence. A judge asks both a human and a computer a set of written questions. If the judge cannot tell which respondent is the computer, and which respondent is human, then the computer "wins" and can thus be considered "intelligent." But Turing himself called this the "imitation game," because the A.I. doesn't have to *think* like a human to pass the test—it needs only to *pretend to think* to trick the judge.²

The fact is, no one knows exactly where A.I. will take us. Both positive and negative, there are infinite possible outcomes.



PAPERCLIP APOCALYPSE

Nick Bostrom's famous "Paperclip Maximizer" thought experiment is a prime example of misaligned goals between A.I. and humans, despite benevolent intentions.

The parable goes like this:

An innocuous team of computer engineers designs an artificial general intelligence system to collect as many paperclips as possible. The AGI might accomplish this by manufacturing paperclips, acquiring pre-existing paperclips, earning money to buy paperclips, building new paperclip factories, etc.

Even though the AGI's goal-oriented mission of global paperclip domination is not best aligned with the interests of humanity, this fact does not make the AGI a power-hungry, malevolent monster. Rather, it simply makes the AGI the absolute best at its job.

To improve its paperclip acquisition abilities, thereby maximizing its reward/utility function, the AGI begins to improve its own intelligence, undergoing an "intelligence explosion." Further enhancing its abilities to self-improve, the AGI soon attains far-above-human intelligence (a.k.a. "superintelligence").

The consequence: The AGI transforms all of earth into paperclips and paperclip manufacturing facilities, utilizing all available resources (water, food, land, money, sunlight, etc.) and energy to further its goal, ultimately extending outward into space. The entire universe becomes one big shiny paper clip factory, leaving zero resources or inhabitable space for humans, leading to the end of the human race.

REFERENCE: Bostrom N. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press, 2014.

Keith Pine, Chief Operating Officer of the digital agency Organic, explains, "A.I. has been around since the '50s. And you're already interacting with it . . . the spam filtering on your email, the recommendations you're being served on Netflix. Those are all machine-learning algorithms that are based off your individual data, and a larger data set as a whole."

It's important, too, to understand the difference between narrow A.I. versus general A.I. What we see around us today are examples of narrow A.I.: intelligent systems that have been taught—or have learned—how to carry out a singular or specific task.³ Examples of narrow A.I. include Siri, Google Translate, NYSE and NASDAQ buying and selling, self-driving cars, and IBM's Watson.

Pine says, "A.I. is a tool to help humans do things better, just like a computer is a tool or a pencil is a tool. People worry it's going to replace all our jobs. A.I. is not meant to replace people, but to help us do our jobs better."

General A.I. (sometimes referred to as "AGI" or "artificial general intelligence") is fundamentally different, and we have not yet witnessed its inception outside of sci-fi movies. General A.I. is the type of adaptable intellect found in humans, the kind

that can learn like we do, can solve any problem (not merely a specific one as with narrow A.I.), and has the potential to not only match human intelligence but to exceed it.^{2,3}

There's a debate among scientists and engineers as to when (and whether) we will see the emergence of general A.I., and, if we do, what form it will take. Most expert surveys show a 10 percent probability we'll have AGI by 2022; a 50 percent probability by 2040; and a 90 percent probability by 2075. Other experts say we'll never get there.⁴

Leigh Christie, Director of the Isobar NowLab, says, "If you're optimizing for one particular metric, and you have a large data set with which you can train your model and your neural net, then in those cases I'm pretty confident that the vast majority of problems that are narrow in nature will be solved. But it's a very murky future with respect to A.I. when it comes to really difficult abstract problems. I'm optimistic that we will solve general A.I., but I'm not optimistic that we'll solve it in a way that's safe."

IS THE SINGULARITY NEAR?

As popularized in his 2005 book *The Singularity Is Near*, Kurtzweil predicts that once the tipping point, or singularity, has been reached, machine intelligence will be infinitely more powerful than all human intelligence combined. At that time, human intelligence will transcend biology, and machine intelligence and humans will merge, expanding outward from the planet until it saturates the entire universe.^{2,4}

In 1965, British mathematician Irving Good argued, "Let an ultraintelligent machine be defined as a machine that can far surpass all the intellectual activities of any man however clever. Since the design of the machines is one of these intellectual activities, an ultraintelligent machine could design even better machines; there would then unquestionably be an 'intelligence explosion,' and the intelligence of man would be left far behind. Thus the first ultraintelligent machine is the last invention that man need ever make, provided that the machine is docile enough to tell us how to keep it under control."⁶

Well, that's vaguely unsettling. The question, however, isn't whether a malevolent *Terminator* scenario will wipe out humanity in an unrelenting fiery blaze. Rather, it's about establishing goals that avoid harm to humanity. Max Tegmark writes in his book *Life 3.0*, "The fear of machines turning evil is another red herring. The real worry isn't malevolence, but competence. A superintelligent A.I. is by definition very good at attaining its goals, whatever they may be, so we need to ensure that its goals are aligned with ours."⁷

What happens to humanity in an environment where non-biological AGI is given free rein to both thrive and grow exponentially? This question has sparked leading tech icons and engineers around the world to vocalize concerns about our future. Eliezer Yudkowsky, co-founder at the Machine Intelligence Research Institute, has stated, "The A.I. does not hate you, nor does it love you, but you are made out of atoms which it can use for something else." Elon Musk, CEO

and chairman of Tesla, has spoken of A.I. as “our greatest existential threat.” Peter Thiel, co-founder of PayPal and a current funder of A.I. research, has pronounced, “**People are spending way too much time thinking about climate change, and way too little thinking about A.I.**” The late theoretical physicist Stephen Hawking wrote that, while achieving superintelligence would be “the biggest event in human history,” it might “also be the last, unless we learn to avoid the risks.”⁵

Messeri says, “This raises ethical questions. Because, if there’s potential for good, then there’s likely also potential for evil. I don’t think we should stop progress. But how do we deploy this technology in a meaningful and productive way, so that we’re not being blinded by some utopian good that we think is the only option for what this technology can bring about? How do we be realists in the creation of these technologies?”

So what to do about this existential risk?

According to Pine, “Now is the time to set a trajectory. We can get ahead of a problem before it happens by thinking and talking about it right now, on a global level. We need to create rules, regulations, and a synergistic knowledge base to help steer A.I. in a direction that is productive for humanity.”

Christie adds, “Do not underestimate the speed at which an A.I. can go from being worse than a human at something to being 10 times better. Only a handful of people are studying these long-term systemic existential issues with A.I. Meanwhile, marketing forces are pressing forward on A.I. Companies are making rapid progress, and there’s not enough people in these corporations raising the concern.”

Is there any national or international regulatory agency currently setting policy to safeguard the future of humanity? The answer is no. According to Christie, “There’s this ‘control problem,’ or ‘alignment problem,’ such that even a slight misalignment between humans and general A.I. could be catastrophic. How do we constrain it such that we’re able to get the benefits without the existential threats? We need to have some sensible and carefully thought-out regulation, but in a way that helps the industry thrive.”

“Now is the time to set a trajectory. We can get ahead of a problem before it happens by thinking and talking about it right now, on a global level.”

CAN WE BIOHACK HUMANITY?

“Resistance is futile.” Anyone who watched *Star Trek: The Next Generation* will remember that prophetic warning. The Borg were a collection of “cybernetic organisms,” linked in a hive mind called the Collective. They advanced through the galaxies, consuming technology and assimilating other species to ostensibly raise quality of life. Transforming individuals into hive drones by injecting bio-chips and surgically augmenting them with cybernetic components, the Borg’s ultimate goal was to achieve perfection.⁸

ORGANIZATIONS DEDICATED TO THE DEVELOPMENT OF “FRIENDLY A.I.”

MACHINE INTELLIGENCE RESEARCH INSTITUTE:

A nonprofit whose mission is to ensure that the creation of smarter-than-human intelligence has a positive impact.

[Intelligence.org/summary](https://intelligence.org/summary)

CENTRE FOR THE STUDY OF EXISTENTIAL RISK (CSER):

A multidisciplinary research center dedicated to the study and mitigation of risks that could lead to human extinction.

www.cser.ac.uk/

FUTURE OF HUMANITY INSTITUTE: A multidisciplinary research institute bringing the tools of mathematics, philosophy, and science to bear on big-picture questions about humanity and its prospects.

www.fhi.ox.ac.uk

FUTURE OF LIFE INSTITUTE: Catalyzes and supports research and initiatives for safeguarding life and developing optimistic visions of the future, including positive ways for humanity to steer its own course considering new technologies and challenges.

futureoflife.org

PARTNERSHIP ON A.I.: Established to study and formulate best practices on A.I. technologies, to advance the public’s understanding of A.I., and to serve as an open platform for discussion and engagement about A.I. and its influences on people and society.

PartnershipOnAI.org

GLOBAL CATASTROPHIC RISK INSTITUTE: A think tank leading research, education, and professional networking on global catastrophic risk.

GCRInstitute.org

80,000 HOURS: A career guide for AI safety researchers.

80000hours.org/career-reviews/artificial-intelligence-risk-research

REFERENCE: Future of Life Institute. <https://futureoflife.org/background/benefits-risks-of-artificial-intelligence>

Although some transhumanist groups are indeed working toward Borg-like life extension, we are quite far from being assimilated into a collective, cybernetic, hive mind. Current research and development into augmenting the human body with A.I. shows significant promise in advanced prosthetics for amputees, robotic heart pumps for those with heart disease, exoskeletons to enable mobility despite loss of muscular reflexes, and brain implants for those with paralysis. Internationally, bioengineers are bio-printing synthetic blood vessels, nerves, muscles, bladders, heart valves, corneas, jaws, hip implants, vertebrae, skin, capillaries, and mini-organs. There is great hope to repair and support any damaged organ or system of the body. Consequently, humans will lead longer, more productive lives, with less pain and dysfunction.⁹

But at what augmented stage is a body no longer biological? If your body, or brain, is 50 percent bionic, are you still human? What about 70 percent? One hundred percent?



"ALEXA, PLEASE SHOW ME EXAMPLES OF NARROW A.I. IN USE RIGHT NOW."

- 1. Amazon.com:** Designed a neural network to predict consumer interest based on your online behavior, creating those magical "Recommended for You" suggestions.
- 2. Facebook image recognition:** Uses a 9-layer-deep neural network with more than 120 million parameters to automatically tag users in an uploaded photo.
- 3. Netflix:** Provides highly accurate predictive technology based on a customer's reactions to films, so you can binge watch to your heart's content.
- 4. Pandora:** Recognizes songs that jive with a listener's style, based on 400 musical characteristics, often leading you to discover songs and artists you didn't know you already loved.
- 5. Alexa, Siri, Cortana, and other voice assistants:** Use machine-learning technology to get smarter and thus better able to predict and understand a user's natural-language questions and requests.
- 6. Nest:** Uses behavioral algorithms to predictively learn a user's heating and cooling preferences, anticipating and automatically adjusting the temperature in a home or office to suit your needs. Goldilocks would have loved this tool.

REFERENCE:

Adams RL. "10 Powerful Examples of Artificial Intelligence in Use Today." *Forbes*. 10 January 2017. <https://www.forbes.com/sites/robertadams/2017/01/10/10-powerful-examples-of-artificial-intelligence-in-use-today>

"The AI in Your Pocket." Intel.com. 4 June 2018. <https://www.intel.com/content/www/us/en/analytics/artificial-intelligence/ai-in-your-pocket-infographic.html>

In *To Be a Machine*, author Mark O'Connell writes, "The idea of the cyborg is no more or less than a particular way of thinking about the human, a peculiarly modern picture of a person as a mechanism for the processing of information."⁵

Given this perspective, is death just a technical problem to be solved? "Can we live forever?" Messeri posits. "I think to answer that question, you have to get to the heart of the more fundamental question: What does it mean to live?"

There is no categorical answer to that question. But the message of the cyborg, O'Connell writes, is "that we would eventually be redeemed of our human nature, of our animal selves, and that all we had to do to secure this redemption was to let technology into our mortal bodies, thereby achieving a communion with machines."⁵

Messeri contends, "We're already cyborgs. The fact that I wear contacts, and have my phone pretty much attached to my hand, and I have constant access to my laptop . . . we are already fully informed in the way we live by technology. Is this evolution or not?"

Our technology influences how we as humans live and interact in the world. From the first huts humans built as shelter, to the robot vacuums now roaming our floors, technology and humanity have formed a symbiotic relationship. Messeri says, "Technology is going to continue to change and enhance, and maybe in certain cases detract from, what it means to be human in the world, but it is incorrect to claim that we're entering some sort of new phase in which for the first time technology is integrated into our bodies in a way that it never has [been] before."

A recent conversation I had about A.I. with a precocious 11-year-old revealed a unique perspective. He sees the efficacy of using robots to help the elderly, but only if they're alone with no family or support system. He worries about a world where humans become extinct, not because of an evil A.I. uprising, but because humans *consciously choose* robots over other humans, such that there are no more human relationships. "Why do we need to be friends with robots anyway?" he asked me. "What's so bad about humans?"

That's an excellent question. And ultimately, what it means to be human will be up to you.

THE IMITATION GAME

A final thought experiment: Is there any way to prove that a human author wrote this article? Maybe an A.I. was programmed to reproduce my written syntax based on a collection of my previously published work, college thesis, and two decades of emails, online chats, and text messages. There is a sufficient amount of data for an engineer to use to create an "Alison Main Writer" program. Maybe an intelligent computer wrote this article?

How would you know?

And does it matter to you? 🤖

- Shoba G, Joy D, Joseph T, Majeed M, Rajendran R, Srinivas PS. "Influence of Piperine on the Pharmacokinetics of Curcumin in Animals and Human Volunteers." *Planta Med.* 1998 May;64(4):353-6. doi: 10.1055/s-2006-957450
- Polley KR, Jenkins N, O'Connor P, McCully K. "Influence of Exercise Training with Resveratrol Supplementation on Skeletal Muscle Mitochondrial Capacity." *Appl Physiol Nutr Metab.* 2016 Jan;41(1):26-32. doi: 10.1159/ajpn-2015-0370
- Nagarkatti P, Pandey R, Rieder SA, Hegde VL, Nagarkatti M. "Cannabinoids as Novel Anti-inflammatory Drugs." *Future Med Chem.* 2009 Oct; 1(7): 1333-1349. doi: 10.4155/fmc.09.93

How Fluctuation Supports Peak Fitness Performance

By Mark Sisson, pp 60-62

- Gwozdziński K, Pieniazek A, Brzeszczynska J, Tabaczar S, Jegier A. "Alterations in Red Blood Cells and Plasma Properties after Acute Single Bout of Exercise." *Scientific World Journal.* 2013 Dec 18;2013:168376. doi: 10.1155/2013/168376
- Nanda B, Balde J, Manjunatha S. "The Acute Effects of a Single Bout of Moderate-Intensity Aerobic Exercise on Cognitive Functions in Healthy Adult Males." *J Clin Diagn Res.* 2013 Sep;7(9):1883-5. doi: 10.7860/JCDR/2013/5855.3341
- McDonnell MN, Buckley JD, Opie GM, Ridding MC, Semmler JG. "A Single Bout of Aerobic Exercise Promotes Motor Cortical Neuroplasticity." *J Appl Physiol (1985).* 2013 May;114(9):1174-82. doi: 10.1152/jappphysiol.01378.2012
- Roig M, Skriver K, Lundbye-Jensen J, Kiens B, Nielsen JB. "A Single Bout of Exercise Improves Motor Memory." *PLoS One.* 2012;7(9):e44594. doi: 10.1371/journal.pone.0044594
- Vissing K, McGee S, Farup J, Kjølhede T, Vendelbo M, Jessen N. "Differentiated mTOR but Not AMPK Signaling after Strength vs Endurance Exercise in Training-Accustomed Individuals." *Scand J Med Sci Sports.* 2013 Jun;23(3):355-66. doi: 10.1111/j.1600-0838.2011.01395.x
- MacKenzie MG, Hamilton DL, Pepin M, Patton A, Baar K. "Inhibition of Myostatin Signaling through Notch Activation Following Acute Resistance Exercise." *PLoS One.* 2013 Jul 2;8(7):e68743. doi: 10.1371/journal.pone.0068743
- Little JP, Safdar A, Bishop D, Tarnopolsky MA, Gibala MJ. "An Acute Bout of High-Intensity Interval Training Increases the Nuclear Abundance of PGC-1 α and Activates Mitochondrial Biogenesis in Human Skeletal Muscle." *Am J Physiol Regul Integr Comp Physiol.* 2011 Jun;300(6):R1303-10. doi: 10.1152/ajpregu.00538.2010

- Moras PK, Campbell CS, Sales MM, Motta DF, Moreira SR, et al. "Acute Resistance Exercise Is More Effective than Aerobic Exercise for 24h Blood Pressure Control in Type 2 Diabetics." *Diabetes Metab.* 2011 Apr;37(2):112-7. doi: 10.1016/j.diabet.2010.08.008

Four Ways to Shed Pounds and Boost Brain Power with Better Sleep

By Cassie Bjork, pp 64-66

- Bosy-Westphal A, Hinrichs S, Jauch-Chara K, Hitz B, Later W, et al. "Influence of Partial Sleep Deprivation on Energy Balance and Insulin Sensitivity in Healthy Women." *Obes Facts.* 2008 1(5):266-73. doi: 10.1159/000158874
- Donga E, van Dijk M, van Dijk JG, Biermasz NR, Lammers GJ, et al. "A Single Night of Partial Sleep Deprivation Induces Insulin Resistance in Multiple Metabolic Pathways in Healthy Subjects." *J Clin Endocrinol Metab.* 2010 Jun;95(6):2963-8. doi: 10.1210/jc.2009-2430
- Buxton OM, Pavlova M, Reid EW, Wang W, Simonson DC, et al. "Sleep Restriction for 1 Week Reduces Insulin Sensitivity in Healthy Men." *Diabetes.* 2010 Sep;59(9):2126-33. doi: 10.2337/ab09-0699
- Sio UN, Monaghan P, Ormerod T. "Sleep on It, But Only If It Is Difficult: Effects of Sleep on Problem Solving." *Mem Cognit.* 2013 Feb;41(2):159-66. doi: 10.3758/s13421-012-0256-7
- Nedeltsheva AV, Kilkus JM, Imperial J, Schoeller DA, Penev PD. "Insufficient Sleep Undermines Dietary Efforts to Reduce Adiposity." *Ann Intern Med.* 2010 Oct 5;153(7):435-41. doi: 10.7326/0003-4819-153-7-201010050-00006
- St-Onge MP, Mikic A, Pietrolungo CE. "Effects of Diet on Sleep Quality." *Adv Nutr.* 2016 Sep 15;7(5):938-49. doi: 10.3945/ajpn.116.012336
- Held K, Antonijevic IA, Künzel H, Uhr M, Wetter TC, et al. "Oral Mg(2+) Supplementation Reverses Age-Related Neuroendocrine and Sleep EEG Changes in Humans." *Pharmacopsychiatry.* 2002 Jul;35(4):135-43.
- Herxheimer A, Petrie KJ. "Melatonin for the Prevention and Treatment of Jet Lag." *Cochrane Database Syst Rev.* 2002 Apr 22. doi: 10.1002/14651858.CD0001520
- Shieh KR, Chu YS, Pan JT. "Circadian Change of Dopaminergic Neuron Activity: Effects of Constant Light and Melatonin." *Neuroreport.* 1997 Jul 7;8(9-10): 2283-7.

The Age of Artificial Intelligence

By Alison Main, pp 68-72

- McCarthy J, Minsky ML, Rochester N, Shannon CE. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence." August 31, 1955. *AI Magazine.* 27.4:12-14 (2006).
- Barrat J. *Our Final Invention: Artificial Intelligence and the End of the Human Era.* St. Martin's Press, 2013.

- Heath N. "What is AI? Everything You Need to Know about Artificial Intelligence." ZDNet. 12 February 2018. <https://www.zdnet.com/article/what-is-ai-everything-you-need-to-know-about-artificial-intelligence/>
- Bostrom N. *Superintelligence: Paths, Dangers, Strategies.* Oxford University Press, 2014.
- O'Connell M. *To Be a Machine: Adventures Among Cyborgs, Utopians, Hackers, and the Futurists Solving the Modest Problem of Death.* Doubleday, 2017.
- Good IJ. "Speculations Concerning the First Ultraintelligent Machine." In *Advances in Computers*, edited by Franz L. Alt and Morris Rubinson, 6: 31-88. Academic Press, 1965. doi: 10.1016/S0065-2458(08)60418-0
- Tegmark M. *Life 3.0: Being Human in the Age of Artificial Intelligence.* Alfred A. Knopf, 2017.
- Star Trek Database. "Borg." 11 May 2018. http://www.startrek.com/database_article/borg
- "Our Future Body." *Popular Science.* Time Inc. Books Special Edition, 2018.

The Benefits of Roughhousing

By Darryl Edwards, pp 100-102

- Panksepp J. "Attention Deficit Hyperactivity Disorders, Psychostimulants, and Intolerance of Childhood Playfulness: A Tragedy in the Making?" *Current Directions in Psychological Science* 7.3 (1 June 1998): 91-98.
- Figueiredo R. "Today's Kids Are Three Times More Likely to End Up in Hospital Because They Fell Out of Bed Rather Than Out of a Tree." *The Daily Mail.* 5 June 2016. <http://www.dailymail.co.uk/news/article-3626687/Today-s-kids-likely-fall-bed-tree.html>
- DeBenedet AT, Cohen L. *The Art of Roughhousing: Good Old-Fashioned Horseplay and Why Every Kid Needs It.* Quirk Books, 2011.
- Panksepp J. *Affective Neuroscience: The Foundations of Human and Animal Emotions.* Oxford University Press, 2004.
- Panksepp J. "The Long-Term Psychobiological Consequences of Infant Emotions: Prescriptions for the Twenty-First Century." *Infant Mental Health Journal.* 2001 Jan-Apr;22(1-2): 132-173.
- Brown S. *Play As An Organizing Principle: Clinical Evidence And Personal Observations.* *Animal Play: Evolutionary, Comparative, And Ecological Perspectives.* Cambridge University Press, 1998.
- Moskal JR, Burgdorf J, Kroes RA, Brudzynski SM, Panksepp J. "A Novel NMDA Receptor Glycine-Site Partial Agonist, Glyx-13, Has Therapeutic Potential for the Treatment of Autism." *Neuroscience & Biobehavioral Reviews.* 2011 Oct;35(9):1982-1988. doi: 10.1016/j.neurobiorev.2011.06.006
- Hamilton J. "Scientists Say Child's Play Helps Build a Better Brain." National Public Radio. 6 August 2014. <https://www.npr.org/sections/ed/2014/08/06/336361277/scientists-say-childrens-play-helps-build-a-better-brain>