

Artificial Intelligence for Everyone


Pedro Gambus, M.D., Steven L. Shafer, M.D.

THIS month's issue of ANESTHESIOLOGY contains a landmark article applying machine learning to model the interaction of remifentanyl and propofol on processed electroencephalogram.¹ It's a big deal.

The June 2017 issue of MIT Technology Review claims "Google stakes its future on a piece of software."² The software, TensorFlow, was written by Google Brain less than 2 yr ago. TensorFlow is the software behind AlphaGo, a machine learning program that plays the ancient Chinese board game Go. There are approximately 10^{170} Go board positions,³ vastly exceeding the measly 10^{82} atoms in the observable universe.⁴ After a few months of training, AlphaGo readily defeated human Go champions who had trained for a lifetime.

Google is revamping its entire business around TensorFlow and machine learning. In 2016, Google replaced Google Translate with sequence-to-sequence neural networks based on TensorFlow.⁵ Google Photos is applying machine learning to your photographs. Several months ago, one of us (S.L.S.) saw *Hamilton* with Ted Eger and his wife, Lynn Spitzer. Shortly after taking a group picture below the *Hamilton* poster, Google asked if we would like to share the picture with Ted and Lynn. How did Google know who was in the photograph? TensorFlow.

The four most valuable companies in the world, Apple, Alphabet, Microsoft, and Amazon, are heavily investing in machine learning. Along with Facebook and IBM, they have formed a "Partnership in AI"⁶ (artificial intelligence) to help ensure that artificial intelligence serves the interests of society. That isn't a given. A few weeks ago Vladimir Putin said on Russian television, "Whoever becomes the leader in this sphere will become the ruler of the world."⁷ Elon Musk quickly agreed: "Competition for AI superiority at national level is the most likely cause of WW3."⁸



"...machine learning to model the interaction of remifentanyl and propofol on processed electroencephalogram.... Is this the end of clinical pharmacology?"

In 1943, Warren McCulloch and Walter Pitts¹² suggested that our slow, binary neurons manage the subtle elegance of human cognition through organization into networks. Twenty years later, their ideas were implemented in the first computer capable of learning, adjusting digital synapses to recognize a pattern.¹³ More complex neural networks followed, built with tensors and implemented in software rather than hardware.

Tensors are clusters of data with an arbitrary number of dimensions, *e.g.*, a scalar (zero dimensions), a vector (one dimension), a matrix (two dimensions), a cube (three dimensions), a set of cubes (four dimensions), a set of cubes in a bunch of balloons (five dimensions), and so on. In neural networks, elements in one tensor flow into elements of the next tensor based on adjustable parameters, which function as synapses. Hundreds of thousands of parameters control the flow of information from one tensor to the next using tensor algebra. Tensor algebra (addition, subtraction, multiplication, division) is just basic algebra applied over many dimensions. Each step yields incremental feature recognition and extraction. The tensors grow to recognize more features

What is artificial intelligence and machine learning? Why is Putin praising AI, and Elon Musk fearful? What does AI have to do with neural networks, Tensorflow, and the article by Lee *et al.*?

Artificial intelligence is the capability of a computer to respond in a manner resembling human intelligence.⁹ It is nearly impossible to write programs that appear smart. Instead, computers are programmed to learn from experience, just like we do. Although some data analysis techniques are described as "learning,"¹⁰ machine learning involves "programming a digital computer to behave in a way which, if done by human beings or animals, would be described as involving the process of learning."¹¹ As practical matter, all useful artificial intelligence is built on machine learning, and nearly all machine learning is built on neural networks.

Image: J. P. Rathmell.

Corresponding article on page 492.

Accepted for publication October 10, 2017. From the Anesthesiology Department, Hospital CLINIC de Barcelona, and Neuroimmunology Research Group, Institut d'Investigacions Biomèdiques August Pi i Sunyer (IDIBAPS), Barcelona, Spain (P.G.); the Department of Anesthesia and Perioperative Care, University of California at San Francisco, San Francisco, California (P.G.); the Department of Anesthesiology, Perioperative and Pain Medicine, Stanford University, Stanford, California (S.L.S.); and the Department of Bioengineering and Therapeutic Sciences, University of California at San Francisco, San Francisco, California (S.L.S.).

Copyright © 2018, the American Society of Anesthesiologists, Inc. Wolters Kluwer Health, Inc. All Rights Reserved. Anesthesiology 2018; 128:431-3

and shrink to remain computationally tractable. The final step shrinks the tensor to the shape of the output: perhaps a state (awake *vs.* asleep), a phrase in another language, or the name of a person in a picture. Training neural networks to recognize simple things, like handwritten digits from zero to nine, requires tens of thousands of passes over different representations of the digits and hundreds of millions of calculations. Until very recently the coding complexity and computational requirements limited neural networks to an academic exercise.

Video games changed that. Realistic video games require graphics coprocessors capable of millions of parallel calculations per second. Modern graphics processing units have thousands of cores. Tensor algebra is inherently parallel. Neural network training that would take days a central processing unit with four or eight cores runs in minutes on a graphics processing unit.

With programs like TensorFlow and Keras (a front end for TensorFlow), assembling a working neural network is simply assembling tensors in sequence, like sticking LEGO blocks together. Once you have assembled the LEGO blocks, if you feed enough data into the neural network, it is likely to work. Using TensorFlow, nearly any programmer can write neural network software. Soon neural networks themselves will code neural networks, freeing us of that drudgery. It gets a bit spooky.

In their article in this issue of ANESTHESIOLOGY, Lee *et al.*¹ use machine learning to link target controlled infusion rates of propofol and remifentanyl to the bispectral index (BIS) in 231 individuals. Their learning set, which they have graciously provided as a digital supplement, contains more than 2 million data points: one sample every second of the propofol infusion rate, the remifentanyl infusion rate, and the BIS. Their model knows nothing about pharmacokinetics or pharmacodynamics. It simply trains a neural network to predict the BIS based on the infusion rates of propofol and remifentanyl. It does exceptionally well—better than the published models of propofol pharmacokinetics, remifentanyl pharmacokinetics, and their pharmacodynamic interaction, using generally accepted metrics of anesthetic drug performance.¹⁴

Is this the end of clinical pharmacology? Will neural networks, assembled from tensor LEGO blocks requiring zero knowledge of pharmacology replace conventional pharmacokinetic and pharmacodynamic models? Will the aging cadre of clinical pharmacologists in our specialty be forced to retire and drive for Uber? Probably, but not yet.

We know a lot about drug behavior. We know when we double the dose, we double the concentration. We know if we give a dose now, and another dose in 10 min, the concentration after the second dose is the sum of the contributions from each dose given alone. We know drug concentrations peak shortly after we give the drug and rapidly decrease as the drug is transported into other tissues. We know clearance permanently eliminates drug from the body. We know drug effect increases with increasing amounts of drug. We

know there is synergy between propofol and remifentanyl. We know that the body has physiologic limits on drug effect.

We know a lot! Standard pharmacokinetic models incorporate our knowledge. This is why models based on brief infusions, like the Schnider model for propofol¹⁵ and the Minto model for remifentanyl,¹⁶ were able to control drug infusions for many hours with reasonable accuracy in the study by Lee *et al.*¹ This is how response surface models predict the combination of any ratio of two drugs.¹⁷ Our knowledge gives us physiologic insight into the processes of drug distribution and clearance. Our models readily calculate infusion rates to achieve the desired responses.

Neural networks start with nothing, nada, zip. They learn by matching the input (*e.g.*, propofol and remifentanyl infusion rates) to the output (*e.g.*, BIS). The resulting model knows nothing about clearance, volume of distribution, or drug potency. The neural network trained by Lee *et al.* likely cannot predict what happens when the dose of either drug is doubled, because that was not part of the training data. The model of Lee *et al.* may not work when the ratio of propofol to remifentanyl is very different from the ratio used in the training data or even when extra boluses of propofol are given when the patient moves during surgery.

These are serious shortcomings. Our fundamental knowledge enables extrapolation from what we know to new situations. Neural networks offer a radically different approach: continuing to learn from data until extrapolation is not necessary. What would happen if computers fed the BIS data during *every* propofol and remifentanyl infusion worldwide into a continuously learning neural network? The resulting neural network, based on millions of patients, would have deeper understanding of the relationship between propofol, remifentanyl, and BIS than any human. We would not need extrapolation. The “serious shortcomings” would be irrelevant. That day may not be far off. Google continuously learns the habits and interests of nearly every human outside China. Medicine will be fundamentally transformed by machine learning once Epic, Cerner, or a major healthcare provider successfully applies machine learning to their massive clinical databases.

The article by Lee *et al.*¹ is a crack in this door, applied to the narrow discipline of anesthetic pharmacology. We still need pharmacokinetic and pharmacodynamic models but perhaps not for long. Machine learning offers a more robust approach than the esoteric modeling clinical pharmacologists have used for decades. For those of us who consider ourselves experts in modeling drug behavior, Uber is hiring.¹⁸

Competing Interests

The authors are not supported by, nor maintain any financial interest in, any commercial activity that may be associated with the topic of this article.

Correspondence

Address correspondence to Dr. Shafer: steven.shafer@stanford.edu

References

1. Lee H-C, Ryu H-G, Chung E-J, Jung C-W: Prediction of bispectral index during target-controlled infusion of propofol and remifentanyl: A deep learning approach. *ANESTHESIOLOGY* 2018; 128:492–501
2. Simonite T: Google stakes its future on a piece of software. *MIT Tech Rev* 2017; 120:58–60
3. Number of legal go positions. Available at: <http://tromp.github.io/go/legal.html>. Accessed September 14, 2017
4. Villanueva JC: How many atoms are there in the universe? Available at: <http://www.universetoday.com/36302/atoms-in-the-universe/>. Accessed September 14, 2017
5. Turok S: The magic behind Google Translate: Sequence-to-sequence models and TensorFlow. Available at: <http://blog.altoros.com/the-magic-behind-google-translate-sequence-to-sequence-models-and-tensorflow.html>. Accessed September 14, 2017
6. Partnership on AI website. Available at: <http://www.partnershiponai.org/>. Accessed September 14, 2017
7. Gigova R: Who Vladimir Putin thinks will rule the world. *CNN website*. Available at: <http://www.cnn.com/2017/09/01/world/putin-artificial-intelligence-will-rule-world/index.html>. Accessed September 14, 2017
8. Musk E. Available at: Musk Elon@elonmusk Twitter, 2:33 AM - September 4, 2017. Accessed September 14, 2017
9. Artificial intelligence. Available at: http://en.wikipedia.org/wiki/Artificial_intelligence. Accessed September 14, 2017
10. Gambús PL, Jensen EW, Jospin M, Borrat X, Martínez Pallí G, Fernández-Candil J, Valencia JF, Barba X, Caminal P, Trocóniz IF: Modeling the effect of propofol and remifentanyl combinations for sedation–analgesia in endoscopic procedures using an Adaptive Neuro Fuzzy Inference System (ANFIS). *Anesth Analg* 2011; 112:331–9
11. Samuel AL: Some studies in machine learning using the game of checkers. *IBM J* 1959; 3:210–29
12. McCulloch WS, Pitts WH: A logical calculus of the ideas immanent in nervous activity. *Bull Mathematical Biophys* 1943; 5:15–133
13. Widrow B, Hoff ME: Adaptive switching circuits. Available at: <http://www-isl.stanford.edu/~widrow/papers/c1960adaptive-switching.pdf>. Accessed September 14, 2017
14. Varvel JR, Donoho DL, Shafer SL: Measuring the predictive performance of computer-controlled infusion pumps. *J Pharmacokinetic Biopharm* 1992; 20:63–94
15. Schnider TW, Minto CF, Gambus PL, Andresen C, Goodale DB, Shafer SL, Youngs EJ: The influence of method of administration and covariates on the pharmacokinetics of propofol in adult volunteers. *ANESTHESIOLOGY* 1998; 88:1170–82
16. Minto CF, Schnider TW, Egan TD, Youngs E, Lemmens HJ, Gambus PL, Billard V, Hoke JF, Moore KH, Hermann DJ, Muir KT, Mandema JW, Shafer SL: Influence of age and gender on the pharmacokinetics and pharmacodynamics of remifentanyl: I. Model development. *ANESTHESIOLOGY* 1997; 86:10–23
17. Bouillon TW, Bruhn J, Radulescu L, Andresen C, Shafer TJ, Cohane C, Shafer SL: Pharmacodynamic interaction between propofol and remifentanyl regarding hypnosis, tolerance of laryngoscopy, bispectral index, and electroencephalographic approximate entropy. *ANESTHESIOLOGY* 2004; 100:1353–72
18. Uber Advanced Technologies Group website. Available at: <http://www.uber.com/info/atg/>. Accessed September 14, 2017