А.И. Матяшевская, Е.В. Тиден

THE POWER OF ALGORITHMS:

part 3

Учебное пособие

Саратов

2019

Составители - А.И. Матяшевская, Е.В. Тиден

The power of algorithms: part 3: Учебное пособие по иностранному языку для студентов /Сост. А.И. Матяшевская, Е.В. Тиден. — Саратов, 2019. — 83 с.

Рецензент:

Кандидат философских наук Шилова С.А.

Table of Contents

PREFACE

Настоящее учебное пособие включает актуальные тексты (2018-2019гг.) учебно-познавательной тематики для студентов механико-математического факультета (направления 02.03.01 «Математика и компьютерные науки», 01.03.02 «Прикладная математика и информатика», 38.03.05 «Бизнес-информатика»). Целью данного пособия является формирование навыка чтения и перевода научно-популярных текстов, а также развитие устной речи студентов (умение выразигь свою точку зрения, дать оценку обсуждаемой проблеме).

Пособие состоит из 5 разделов, рассматривающих значение информационных технологий в современном мире. Каждый из них содержит аутентичные материалы (источники: *Quanta Magazine, The Atlantic, Gizmodo, Aeon, Vox, Logic magazine, Bloomberg*) и упражнения к ним.

Раздел “Supplementary reading“ служит материалом для расширения словарного запаса и дальнейшего закрепления навыков работы с текстами по специальности. Пособие может успешно использоваться как для аудиторных занятий, так и для внеаудиторной практики.

1. New Theory Cracks Open the Black Box of Deep Learning

Exercise I.

Say what Russian words help to guess the meaning of the following words: algorithms, video, principle, systems, operates, neurons, signals, photo, experts, neuroscientist

Exercise II.

Make sure you know the following words and word combinations.

coarse-graining, drawn-out, distinct, plausibility, feat, to glean, discrete, to traverse, salient, to squeeze

New Theory Cracks Open the Black Box of Deep Learning

A new idea called the "information bottleneck" is helping to explain the puzzling success of today's artificial-intelligence algorithms — and might also explain how human brains learn

Even as machines known as "deep neural networks" have learned to converse, drive cars, beat video games, dream, paint pictures and help make scientific discoveries, they have also confounded their human creators, who never expected so-called "deep-learning" algorithms to work so well. No underlying principle has guided the design of these learning systems, other than vague inspiration drawn from the architecture of the brain (and no one really understands how that operates either). Like a brain, a deep neural network has layers of neurons — artificial ones that are figments of computer memory. When a neuron fires, it sends signals to connected neurons in the layer above. During deep learning, connections in the network are strengthened or weakened as needed to make the system better at sending signals from input data — the pixels of a photo of a dog, for instance — up through the layers to neurons associated with the right high-level concepts, such as “dog.” After a deep neural network has “learned” from thousands of sample dog photos, it can identify dogs in new photos as accurately as people can. The magic leap from special cases to general concepts during learning gives deep neural networks their power, just as it underlies human reasoning, creativity and the other faculties collectively termed “intelligence.” Experts wonder what it is about deep learning that enables generalization — and to what extent brains apprehend reality in the same way. Naftali Tishby, a computer scientist and neuroscientist, presented evidence in support of a new theory explaining how deep learning works. Tishby argues that deep neural networks learn according to a procedure called the “information bottleneck” . The idea is that a network rids noisy input data of extraneous details as if by squeezing the information through a bottleneck, retaining only the features most relevant to general concepts. Striking new computer experiments reveal how this squeezing procedure happens during deep learning, at least in the cases they studied.

Tishby’s findings have the AI community buzzing. “I believe that the information bottleneck idea could be very important in future deep neural network research,” said Alex Alemi of Google Research, who has already developed new approximation methods for applying an information bottleneck analysis to large deep neural networks. The bottleneck could serve “not only as a theoretical tool for understanding why our neural networks work as well as they do currently, but also as a tool for constructing new objectives and architectures of networks,” Alemi said.

Some researchers remain skeptical that the theory fully accounts for the success of deep learning, but Kyle Cranmer, a particle physicist at New York University who uses machine learning to analyze particle collisions at the Large Hadron Collider, said that as a general principle of learning, it "somehow smells right." Geoffrey Hinton, a pioneer of deep learning who works at Google and the University of Toronto, emailed Tishby after watching his Berlin talk. "It's extremely interesting," Hinton wrote. "I have to listen to it another 10,000 times to really understand it, but it's very rare nowadays to hear a talk with a really original idea in it that may be the answer to a really major puzzle." According to Tishby, who views the information bottleneck as a fundamental principle behind learning, whether you're an algorithm or a conscious being, that long-awaited answer "is that the most important part of learning is actually forgetting." Tishby began contemplating the information bottleneck around the time that other researchers were first mulling over deep neural networks, though neither concept had been named yet. It was the 1980s, and Tishby was thinking about how good humans are at speech recognition — a major challenge for AI at the time. Tishby realized that the crux of the issue was the question of relevance: What are the most relevant features of a spoken word, and how do we tease these out from the variables that accompany them, such as accents, mumbling and intonation? In general, when we face the sea of data that is reality, which signals do we keep? "This notion of relevant information was mentioned many times in history but never formulated correctly," Tishby said in an interview last month. "For many years people thought information theory wasn't the right way to think about relevance, starting with misconceptions that go all the way to Shannon himself." Claude Shannon, the founder of information theory, in a sense liberated the study of information starting in the 1940s by allowing it to be considered in the abstract — as 1s and 0s with purely mathematical meaning. Shannon took the view that, as Tishby put it, "information is not about semantics." But, Tishby argued, this isn't true. Using information theory, he realized, "you can define 'relevant' in a precise sense." Imagine X is a complex data set, like the pixels of a dog photo, and Y is a simpler variable represented by those data, like the word "dog." You can capture all the "relevant" information in X about Y by compressing X as much as you can without losing the ability to predict Y. In their 1999 paper, Tishby and co-authors formulated this as a mathematical optimization problem. It was a fundamental idea with no killer application. "I've been thinking along these lines in various contexts for 30 years," Tishby said. "My only luck was that deep neural networks became so important."

Though the concept behind deep neural networks had been kicked around for decades, their performance in tasks like speech and image recognition only took off in the early 2010s, due to improved training regimens and more powerful computer processors. Tishby recognized their potential connection to the information bottleneck principle in 2014 after reading a surprising paper by the physicists David Schwab and Pankaj Mehta. The duo discovered that a deep-learning algorithm invented by Hinton called the "deep belief net" works, in a particular case, exactly like renormalization, a technique used in physics to zoom out on a physical system by coarse-graining over its details and calculating its overall state. When Schwab and Mehta applied the deep belief net to a model of a magnet at its “critical point,” where the system is fractal, or self-similar at every scale, they found that the network automatically used the renormalization-like procedure to discover the model’s state. It was a stunning indication that, as the biophysicist Ilya Nemenman said at the time, “extracting relevant features in the context of statistical physics and extracting relevant features in the context of deep learning are not just similar words, they are one and the same.”

The only problem is that, in general, the real world isn’t fractal. But Tishby realized that both deep learning and the coarse-graining procedure could be encompassed by a broader idea. He hypothesized that deep learning is an information bottleneck procedure that compresses noisy data as much as possible while preserving information about what the data represent. New experiments with deep neural networks reveal how the bottleneck procedure actually plays out. In one case, the researchers used small networks that could be trained to label input data with a 1 or 0 (think “dog” or “no dog”). They then tracked what happened as the networks engaged in deep learning with 3,000 sample input data sets. Each time the training data are fed into the network, a cascade of firing activity sweeps upward through the layers of artificial neurons. When the signal reaches the top layer, the final firing pattern can be compared to the correct label for the image — 1 or 0, “dog” or “no dog.” Any differences between this firing pattern and the correct pattern are “back-propagated” down the layers, meaning that, like a teacher correcting an exam, the algorithm strengthens or weakens each connection to make the network layer better at producing the correct output signal. Over the course of training, common patterns in the training data become reflected in the strengths of the connections, and the network becomes expert at correctly labeling the data, such as by recognizing a dog, a word, or a 1. In the experiments, Tishby tracked how much information each layer of a deep neural network retained about the input data and how much information each one retained about the output label. The scientists found that, layer by layer, the networks converged to the information bottleneck theoretical bound: a theoretical limit derived in Tishby original paper that represents the absolute best the system can do at extracting relevant information. At the bound, the network has compressed the input as much as possible without sacrificing the ability to accurately predict its label. Tishby also made the intriguing discovery that deep learning proceeds in two phases: a short "fitting" phase, during which the network learns to label its training data, and a much longer "compression" phase, during which it becomes good at generalization, as measured by its performance at labeling new test data. As a deep neural network tweaks its connections, at first the number of bits it stores about the input data stays roughly constant or increases slightly, as connections adjust to encode patterns in the input and the network gets good at fitting labels to it. Some experts have compared this phase to memorization. Then learning switches to the compression phase. The network starts to shed information about the input data, keeping track of only the strongest features — those correlations that are most relevant to the output label. This happens because more or less accidental correlations in the training data tell the network to do different things. This randomization is effectively the same as compressing the system's representation of the input data. As an example, some photos of dogs might have houses in the background, while others don't. As a network cycles through these training photos, it might "forget" the correlation between houses and dogs in some photos as other photos counteract it. It's this forgetting of specifics, Tishby argues, that enables the system to form general concepts. Indeed, their experiments revealed that deep neural networks ramp up their generalization performance during the compression phase, becoming better at labeling test data. (A deep neural network trained to recognize dogs in photos might be tested on new photos that may or may not include dogs, for instance.)

It remains to be seen whether the information bottleneck governs all deep-learning regimes, or whether there are other routes to generalization besides compression. Some AI experts see Tishby's idea as one of many important theoretical insights about deep learning to have emerged recently. Andrew Saxe, an AI researcher and theoretical neuroscientist at Harvard University, noted that certain very large deep neural networks don't seem to need a drawn-out compression phase in order to generalize well. Instead, researchers program in something called early stopping, which cuts training short to prevent the network from encoding too many correlations in the first place. Tishby argues that the network models analyzed by Saxe and his colleagues differ from standard deep neural network architectures, but that nonetheless, the information bottleneck theoretical bound defines these networks' generalization performance better than other methods. Questions about whether the bottleneck holds up for larger neural networks are partly addressed by Tishby most recent experiments. The scientist saw the same convergence of the networks to the information bottleneck theoretical bound; they also observed the two distinct phases of deep learning, separated by an even sharper transition than in the smaller networks. “I’m completely convinced now that this is a general phenomenon,” Tishby said. The mystery of how brains sift signals from our senses and elevate them to the level of our conscious awareness drove much of the early interest in deep neural networks among AI pioneers, who hoped to reverse-engineer the brain’s learning rules. AI practitioners have since largely abandoned that path in the mad dash for technological progress, instead slapping on bells and whistles that boost performance with little regard for biological plausibility. Still, as their thinking machines achieve ever greater feats — even stoking fears that AI could someday pose an existential threat — many researchers hope these explorations will uncover general insights about learning and intelligence. Brenden Lake, an assistant professor of psychology and data science at New York University who studies similarities and differences in how humans and machines learn, said that Tishby’s findings represent “an important step towards opening the black box of neural networks,” but he stressed that the brain represents a much bigger, blacker black box. Our adult brains, which boast several hundred trillion connections between 86 billion neurons, in all likelihood employ a bag of tricks to enhance generalization, going beyond the basic image- and sound-recognition learning procedures that occur during infancy and that may in many ways resemble deep learning. For instance, Lake said the fitting and compression phases that Tishby identified don’t seem to have analogues in the way children learn handwritten characters, which he studies. Children don’t need to see thousands of examples of a character and compress their mental representation over an extended period of time before they're able to recognize other instances of that letter and write it themselves. In fact, they can learn from a single example. Lake and his colleagues' models suggest the brain may deconstruct the new letter into a series of strokes — previously existing mental constructs — allowing the conception of the letter to be tacked onto an edifice of prior knowledge. "Rather than thinking of an image of a letter as a pattern of pixels and learning the concept as mapping those features" as in standard machine-learning algorithms, Lake explained, "instead I aim to build a simple causal model of the letter," a shorter path to generalization. Such brainy ideas might hold lessons for the AI community, furthering the back-and-forth between the two fields. Tishby believes his information bottleneck theory will ultimately prove useful in both disciplines, even if it takes a more general form in human learning than in AI. One immediate insight that can be gleaned from the theory is a better understanding of which kinds of problems can be solved by real and artificial neural networks. "It gives a complete characterization of the problems that can be learned," Tishby said. These are "problems where I can wipe out noise in the input without hurting my ability to classify. This is natural vision problems, speech recognition. These are also precisely the problems our brain can cope with." Meanwhile, both real and artificial neural networks stumble on problems in which every detail matters and minute differences can throw off the whole result. Most people can't quickly multiply two large numbers in their heads, for instance. "We have a long class of problems like this, logical problems that are very sensitive to changes in one variable," Tishby

said. “Classifiability, discrete problems, cryptographic problems. I don’t think deep learning will ever help me break cryptographic codes.” Generalizing — traversing the information bottleneck, perhaps — means leaving some details behind. This isn’t so good for doing algebra on the fly, but that’s not a brain’s main business. We’re looking for familiar faces in the crowd, order in chaos, salient signals in a noisy world.
Adapted from Quanta Magazine

Exercise III.

Fill in the gaps.

1) Once a major traffic ____________, it still suffers from congestion at peak times.

2) To ____________ the stereotypes, the campaign decided to challenge them directly.

3) ________________ research slowed until computers achieved greater processing power.

4) Officers used websites such as Craigslist to identify and ___________ the suspects.

5) And for that purpose, it comes with all kinds of _____________ technical apparatus.

6) Is it simply that we find such superhuman spectacles too bizarre to ________________?

7) The ________________ on this device, though, is its terrific traffic navigation.

8) Once an error is generated, it will generally ____________ through the calculation.

9) Knowing that sweeping change is very stressful, ____________ that by anticipating.

10) Scientists show _______________ of new pathway to life's chemical building blocks.

Exercise IV.

Make up sentences of your own with the following word combinations: deep neural network, to mull over, to zoom out, on the fly, to apprehend, to buzz, to contemplate, to stumble on problems, to throw off the whole result, to multiply large numbers in one'heads

Exercise V.

Match the words to the definitions in the column on the right:

bottleneck	park (a vehicle) in a depot
confound	act against (something) in order to reduce its force or neutralize it
crux	breed specimens of (a plant, animal, etc.) by natural processes from the parent stock
fractal	twist or pull (something) sharply
to encompass	(of several people or things) come together from different directions so as eventually to meet
to propagate	relating to or of the nature of a curve or geometric figure, each part of which has the same statistical character as the whole
to converge	the decisive or most important point at issue
to tweak	cause surprise or confusion in (someone), esp. by acting

	against their expectations
to shed	the neck or mouth of a bottle
to counteract	surround and have or hold within

Exercise VI.

Identify the part of speech the words belong to: figment, extraneous, creators, algorithms, inspiration, figments, connections, magic, special, generalization

Exercise VII.

Match the words to make word combinations:

computer	inspiration
vague	memory
killer	success
scientific	bottleneck
video	application
human	Box
deep	discoveries
information	games
puzzling	learning
Black	brains

Exercise VIII.

Summarize the article "New Theory Cracks Open the Black Box of Deep Learning".

2. Artificial Intelligence Shows Why Atheism Is Unpopular

Exercise I.

Say what Russian words help to guess the meaning of the following words: atheism, president, minimizing, resources, policy, international, philosophers, mimic, attributes, project

Exercise II.

Make sure you know the following words and word combinations. seamlessly, to overstate, refugees, malaise, violence, yield, secularization, noble, xenophobic anxiety, assumptions

Artificial Intelligence Shows Why Atheism Is Unpopular

Imagine you're the president of a European country. You're slated to take in 50,000 refugees from the Middle East this year. Most of them are very religious, while most of your population is very secular. You want to integrate the newcomers seamlessly, minimizing the risk of economic malaise or violence, but you have limited resources. One of your advisers tells you to invest in the refugees' education; another says providing jobs is the key; yet another insists the most important thing is giving the youth opportunities to socialize with local kids. What do you do? Well, you make your best guess and hope the policy you chose works out. But it might not. Even a policy that yielded great results in another place or time may fail miserably in your particular country under its present circumstances. If that happens, you might find yourself wishing you could hit a giant reset button and run the whole experiment over again, this time choosing a different policy. But of course, you can't experiment like that, not with real people. You can, however, experiment like that with virtual people. And that's exactly what the Modeling Religion Project does. An international team of computer scientists, philosophers, religion scholars, and others are collaborating to build computer models that they populate with thousands of virtual people, or "agents." As the agents interact with each other and with shifting conditions in their artificial environment, their attributes and beliefs—levels of economic security, of education, of religiosity, and so on—can change. At the outset, the researchers program the agents to mimic the attributes and beliefs of a real country's population using survey data from that country. They also "train" the model on a set of empirically validated social-science rules about how humans tend to interact under various pressures. And then they experiment: Add in 50,000 newcomers, say, and invest heavily in education. How does the artificial society change? The model tells you. Don't like it? Just hit that reset button and try a different policy. The goal of the project is to give politicians an empirical tool that will help them assess competing policy options so they can choose the most effective one. It's a noble idea: If leaders can use artificial intelligence to predict which policy will produce the best outcome, maybe we'll end up with a healthier and happier world. But it's also a dangerous idea: What's "best" is in the eye of the beholder, after all. “Because all our models are transparent and the code is always online,” said LeRon Shults, who teaches philosophy and theology at the University of Agder in Norway, “if someone wanted to make people more in-group-y, more anxious about protecting their rights and their group from the threat of others, then they could use the model to figure out how to ratchet up anxiety.”

The one that focuses most on refugees, Modeling Religion in Norway(modrn), is still in its early phases. Led by Shults, it’s funded primarily by the Research Council of Norway, which is counting on the model to offer useful advice on how the Norwegian government can best integrate refugees. Norway is an ideal place to do this research, not only because it’s currently struggling to integrate Syrians, but also because the country has gathered massive data sets on its population. By using them to calibrate his model, Shults can get more accurate and fine-grained predictions, simulating what will happen in a specific city and even a specific neighborhood. Another project, Forecasting Religiosity and Existential Security, examines questions about nonbelief: Why aren’t there more atheists? Why is America secularizing at a slower rate than Western Europe? Which conditions would speed up the process of secularization—or, conversely, make a population more religious? Shults’s team tackled these questions using data from the International Social Survey Program. “We were able to predict from that 1998 data—in 22 different countries in Europe, and Japan—whether and how belief in heaven and hell, belief in God, and religious attendance would go up and down over a 10-year period. We were able to predict this in some cases up to three times more accurately than linear analysis," Shults said, referring to a general-purpose method of prediction that prior to the team's work was the best alternative. Using a separate model, Future of Religion and Secular Transitions (forest), the team found that people tend to secularize when four factors are present: existential security (you have enough money and food), personal freedom (you're free to choose whether to believe or not), pluralism (you have a welcoming attitude to diversity), and education (you've got some training in the sciences and humanities). If even one of these factors is absent, the whole secularization process slows down. This, they believe, is why the U.S. is secularizing at a slower rate than Western and Northern Europe. "The U.S. has found ways to limit the effects of education by keeping it local, and in private schools, anything can happen," said Shults's collaborator, Wesley Wildman, a professor of philosophy and ethics at Boston University. "Lately, there's been encouragement from the highest levels of government to take a less than welcoming cultural attitude to pluralism. These are forms of resistance to secularization." Another project, Mutually Escalating Religious Violence (merv), aims to identify which conditions make xenophobic anxiety between two different religious groups likely to spiral out of control. As they built this model, the team brought in an outside expert: Monica Toft, an international-relations scholar with no experience in computational modeling but a wealth of expertise in religious extremism. "They brought me in so I could do a reality check—like, do the social-science assumptions behind this model make sense? And then to evaluate whether this tracks with case studies in reality," Toft told me. At first, she said, "I was a little skeptical with this stuff. But I think what surprised me was how well it modeled onto the case." It shows that mutually escalating violence is likeliest to occur if there's a small disparity in size between the majority and minority groups (less than a 70/30 split) and if agents experience out-group members as social and contagion threats (they worry that others will be invasive or infectious). It's much less likely to occur if there's a large disparity in size or if the threats agents are experiencing are mostly related to predators or natural hazards. This might sound intuitive, but having quantitative, empirical data to support social-science hypotheses can help convince policymakers of when and how to act if they want to prevent future outbreaks of violence. And once a model has been shown to track with real-world historical examples, scientists can more plausibly argue that it will yield a trustworthy recommendation when it's fed new situations.

To that end, the next step is getting others interested in trying out the models. But that's proven difficult. The team is building an online platform that will allow people with zero programming experience to create agent-based models. Still, Wildman is pessimistic about his own ability to get politicians interested in such a new and highly technical methodology. "Whenever there's bafflement, you've got a trust problem, and I think there will be a trust problem here," he said. "We're modelers, sociologists, philosophers—we're academic geeks, basically. We're never going to convince them to trust a model." But he believes that policy analysts, acting as bridges between the academic world and the policy world, will be able to convince the politicians. "We're going to get them in the end." Even harder to sway may be those concerned not with the methodology's technical complications, but with its ethical complications. As Wildman told me, "These models are equal-opportunity insight generators. If you want to go militaristic, then these models tell you what the targets should be." When you build a model, you can accidentally produce recommendations that you weren't intending. Years ago, Wildman built a model to figure out what makes some extremist groups survive and thrive while others disintegrate. It turned out one of the most important factors is a highly charismatic leader who personally practices what he preaches. "This immediately implied an assassination criterion," he said. "It's basically, leave the groups alone when the leaders are less consistent, but kill the leaders of groups that have those specific qualities. It was a shock to discover this dropping out of the model. I feel deeply uncomfortable that one of my models accidentally produced a criterion for killing religious leaders." The results of that model have been published, so it may already have informed military action. "Is this type of thing being used to figure out criteria for drone killings?" Wildman said. "I've come to assume that on the secret side they've pretty much already thought of everything we've thought of... But it could be that this model actually took them there. That's a serious ethical conundrum." The other models raise similar concerns, he said. "The modrn model gives you a recipe for accelerating secularization—and it gives you a recipe for blocking it. That keeps me up at night." According to Neil Johnson, a physicist who models terrorism and other extreme behaviors that arise in complex systems, "That's an overstatement of the power of the models." There's no way that removing one factor from a society can reliably be counted on to slow or stop secularization, he said. That may well be true in the model, but "that's a cartoon of the real world." A real human society is so complex that "all the things may be interconnected in a different way than in the model." Although Johnson said he found the team's research useful and important, he was unimpressed by their claim to have outperformed previous predictive methods. He cautioned that we should be skeptical about the word prediction in relation to this type of model. "It's great to have as a tool," he said. "It's like, you go to the doctor, they give an opinion. It's always an opinion, we never say a doctor's prediction. Usually, we go with the doctor's opinion because they've seen many cases like this, many humans who come in with the same thing. It's even more of an opinion with these types of models, because they haven't necessarily seen many cases just like it—history mimics the past but doesn't exactly repeat it." The silver lining here is that if the power of the models is being overstated then so, too, is the ethical concern. Nevertheless, just like Wildman, Shults told me, "I lose sleep at night on this. It is social engineering. It just is—there's no pretending like it's not." Instead, he and Wildman believe the answer is to do the work with transparency and simultaneously speak out about the ethical danger inherent in it.

Adapted from The Atlantic

Exercise III.

Fill in the gaps.

1) Britain may not use the euro, but it isn't immune to the continent-wide ______________.

2) Despite his novice status as an analyst he slipped into comment mode ________________.

3) At the story's _______________, he doesn't even comprehend that there is a world outside.

4) Yet Boeing itself serves as a _______________ in how globalization can cut both ways.

5) Surgeons are often deaf to patients'compelling desire for less _______________ options.

6) With the opening of schools next week, illness and _______________ are sure to worsen.

7) Chances are, the _______________ between the two reports was mostly statistical noise.

8) Scientific finding is a value judgment based on evidence and __________________.

9) This very simple process hardly seems to have anything that can cause ___________________.

10) On the other hand, it would be hard to _______________ the significance of this panel.

<u>Exercise IV</u>.

Make up sentences of your own with the following word combinations:

to ratchet up, case study, human outcomes, to invest in, to work out, under present circumstances, to run the whole experiment, at the outset, to mimic the attributes and beliefs of a real country's population

<u>Exercise V</u>.

Match the words to the definitions in the column on the right:

slate	confusion resulting from failure to understand

malaise	apparent validity
attribute	(esp. of plants or a disease) tending to spread prolifically and undesirably or harmfully
outset	the communication of disease from one person to another by close contact
insight	a great difference
invasive	the start or beginning of something
contagion	a general feeling of discomfort, illness, or uneasiness whose exact cause is difficult to identify
disparity	a quality or feature regarded as a characteristic or inherent part of someone or something
plausibility	a fine-grained gray, green, or bluish metamorphic rock easily split into smooth, flat pieces
bafflement	the capacity to gain an accurate and deep intuitive understanding of a person or thing

Exercise VI.

Identify the part of speech the words belong to. religious, population, violence, socialize, local, miserably, particular, present, circumstances, experiment

Exercise VII.

Match the words to make word combinations:

economic	data

limited	security
real	button
artificial	tool
economic	idea
survey	East
noble	malaise
empirical	environment
reset	people
Middle	resources

<u>Exercise VIII</u>.

Summarize the article "Artificial Intelligence Shows Why Atheism Is Unpopular".

3. The Future of Online Dating

Exercise I.

Say what Russian words help to guess the meaning of the following words: profiles, information, realize, journalist, massive, indicators, volunteer, reaction, version, professor

Exercise II.

Make sure you know the following words and word combinations.
allegedly, faded, mining, gimmick, dating app, to compute, allegedly better, to misrepresent, to fill out, to link their social media accounts

The Future of Online Dating

Loveflutter, a Twitter-themed dating app from the UK, is paired with the language processing company Receptiviti.ai to compute the compatibility between me and its user base using the contents of our Twitter feeds. Is this good matchmaking or a gimmick? Dating apps promise to connect us with people we're supposed to be with allegedly better than we know ourselves. Sometimes it works out, sometimes it doesn't. But as machine learning algorithms become more accurate and accessible than ever, dating companies will be able to learn more precisely who we are and who we "should" go on dates with. How we date online is about to change. The future is brutal and we're halfway there. Today, dating companies fall into two camps: sites like eHarmony, Match, and OkCupid ask users to fill out long personal essays and answer personality questionnaires which they use to pair members by compatibility (though when it comes to predicting attraction, researchers find these surveys dubious). Profiles like these are rich in information, but they take time to fill out and give daters ample incentive to misrepresent themselves (by asking questions like, “How often do you work out?” or “Are you messy?”). On the other hand, companies like Tinder, Bumble, and Hinge skip surveys and long essays, instead asking users to link their social media accounts. Tinder populates profiles with Facebook friends and likes, and Instagram photos. Instead of matching users by “compatibility,” these apps work to provide a stream of warm bodies as fast as possible. It’s true that we reveal more of ourselves in Twitter posts, Facebook likes and Instagram photos than we realize.

We give dating apps access to this data and more: when one journalist from The Guardian asked Tinder for all the information it had on her, the company sent her a report 800 pages long. Sound creepy? Maybe. But when I worked as an engineer and data scientist at OkCupid, massive streams of data like these made me drool. In the future, apps like Tinder may be able to infer more about our personalities and lifestyles through our social media activity than an eHarmony questionnaire ever could capture. Researchers already think they can predict whether or not we’re depressed from our Tweets and the filters we choose on Instagram, and how intelligent, happy, and likely to use drugs we are from our Facebook likes. What’s more, the relationship between our online behavior and what it implies about us is often unintuitive. One study from Cambridge University that analyzed the connection between Facebook likes and personality traits found the biggest predictors of intelligence were liking “Science” and “The Colbert Report” (unsurprising) but also “Thunderstorms” and “Curly Fries.” That connection might defy human logic, but what does that matter if you're feeding a personality algorithm into a matchmaking algorithm?

Because indicators of our personality can be subtle, and we tend not to curate our activity on Facebook as closely as we might a dating profile, perhaps there's more integrity to this data than what users volunteer in survey questions. "My initial reaction to online dating is that people might present a version that's unrealistic," said Chris Danforth, Flint professor of Mathematical, Natural, and Technical Sciences who's studied the link between Instagram, Twitter, and depression. "But what seems to be revealed every time one of these studies comes out is that it looks to be the case that we reveal more about ourselves than we realize, maybe not as much in surveys but in what we do. Someone's likes on Facebook could be a better predictor of whether they would get along with someone than survey answers." The data could also be used to keep users honest when they're making their accounts. "I think it would be interesting if OkCupid called you out as you're filling out your profile," said Jen Golbeck, a researcher who studies the intersection of social media and information. "It could say something like, 'I analyzed your likes and it looks like maybe you are a smoker. Are you sure you want to choose that answer?'" A more jaded dating app could instead alert the person viewing the profile that their match might be lying. They could also ban users who display personality traits that don't work well in relationships. eHarmony, for example, rejects applicants who've been married four or more times or those whose survey responses indicate they might be depressed. Algorithms could also use our online behavior to learn the real answers to questions we might lie about in a dating questionnaire. One of OkCupid's matching questions, for example, asks "Do you work out a lot?" But MeetMeOutside, a dating app for sporty people, asks users to link their Fitbits and prove they're physically active through their step counts. This type of data is harder to fake. Or, rather than ask someone whether they're more likely to go out or chill on a Friday night, a dating app could simply collect this data from our GPS or Foursquare activity and pair equally active users. It's also possible that computers, with access to more data and processing power than any human, could pick up on patterns human beings miss or can't even recognize. "When you're looking through the feed of someone you're considering, you only have access to their behavior," Danforth says. "But an algorithm would have access to the differences between their behavior and a million other people's. There are instincts that you have looking through someone's feed that might be difficult to quantify, and there may be other dimension we don't see… nonlinear combinations which aren't easy to explain." Just as dating algorithms will get better at learning who we are, they'll also get better at learning who we like—without ever asking our preferences. Instead of asking questions about individuals, we work purely on their behavior as they navigate through a dating site. Rather than ask someone, 'What sort of people do you prefer? Ages 50-60?' we look at who he's looking at. If it's 25-year-old blondes, our system starts recommending him 25-year-old blondes. OkCupid data shows that male users tend to message women significantly younger than the age they say they're looking for, so making recommendations based on behavior rather than self-reported preference is likely more accurate. Algorithms that analyze user behavior can also identify subtle, surprising, or hard-to-describe patterns in what we find attractive—the ineffable features that make up one's "type." Or at least, some app makers seem to think so. If you look at the recommendations we generated for individuals, you'll see they all reflect the same type of person—all brunettes, blondes, of a certain age. Naturally, we might not like the patterns computers find in who we're attracted to. When I asked Justin Long, founder of the AI dating company Bernie.ai, what patterns his software found, he wouldn't tell me: "Regarding what we learned, we had some disturbing results that I do not want to share. They were quite offensive." I'd guess the findings were racist: OkCupid statistics show that even though people say they don't care about race when choosing a partner, they usually act as if they do. "I personally have thought about whether my swiping behavior or the people I match with reveal implicit biases that I'm not even aware that I have," said Camille Cobb, who researches dating tech and privacy. "We just use these apps to find people we're interested in, without thinking. I don't think the apps are necessarily leaking this in a way that would damage my reputation—they're probably using it to make better matches—but if I wish I didn't have those biases, then maybe I don't want them to use that."

Even if dating companies aren't using our data to damage our reputations, they might be using it to make money. "It's sketchy to think what type of information they could give advertisers, especially if it's information we don't even know about ourselves… I don't smoke but maybe if I swipe right on a lot of guys who like cigarettes in my pictures, it reveals I think cigarettes make you look cool." An advertiser could learn what products we find subconsciously interesting and show us targeted ads. Yet these types of tailored recommendation algorithms all seek to make us swipe more. As apps truly get better at learning who we like and who we are, they may render swiping, liking, and messaging obsolete. This was the thought Canadian engineer Justin Long had when he built a "personal matchmaker assistant" called Bernie.ai. Frustrated by how much time he spent swiping and messaging compared to going on actual dates, he decided to build a bot to do the work for him. His app, Bernie, asked users to link their existing Tinder accounts and then watched them swipe, meanwhile modeling users' individual tastes. Then Bernie started swiping on Tinder for them. If the AI encountered a mutual match, it would start a conversation with the opening line, "Do you like avocados?" Tinder eventually forced Long to cease operation, but Long thinks personal dating assistants like Bernie are the future of dating tech. Instead of spending time swiping and messaging, we'll give our digital matchmakers access to our calendars and GPS locations and let them deal with logistics on our behalves. Then, "my Bernie will talk to your Bernie," says Long, and organize dates automatically. When algorithms are so good that we trust their decisions, perhaps we won't mind giving them more control of our love lives. As algorithms get better, they'll need to collect data not just on whose profile photos we like but also who we feel chemistry with in person. Not a single dating app (that I'm aware of) asks users for the outcomes of actual dates. When I asked OkCupid's Director of Engineer Tom Jacques why, he cites bias: "It's a tricky issue because there is a very steep drop-off in what information people will volunteer, and we can only keep track of interactions between members while they are using the site.

At some point, they will take their connection to the real world, and very few people who go on a date (successful or not) will tell us.” Yet we volunteer more than enough information for apps to be able to deduce how our dates went. They could use our GPS coordinates to watch who we go on dates with, how long those dates last, and whether they lead to a second date. The dating app Once even let daters monitor their heart rates on dates through their Fitbits to tell how much they found their date arousing. Today, dating apps don’t (openly) mine our digital data as nearly much as they could. Maybe they think we’d find it too creepy, or maybe we wouldn’t like what they learned about it. But if data mining were the key to the end of the bad date, wouldn’t it be worth it? I’m still on the fence, but as much as I like the idea of a hyper-intelligent, perceptive dating algorithm, I think I’ll delete my Loveflutter account.

Adapted from Gizmodo

Exercise III.

Fill in the gaps.

1) He declined to respond to critics who call him a joke, a ____________, a distraction.

2) When that man tried to call 911, Farah _____________ grabbed his phone and threw it.

3) He's just trying to get to know her in a way she'll find either sweet or _____________.

3) The thing is, a micromachine would ___________ some sort of automation being required.

4) As they talk about what once was, _____________ memories become new, living friendships.

5) It is not possible for us to ________________ the potential effects of such adjustments.

6) Such human encounters, momentary though they are, enhance life in an ______________ way.

7) Instead, the researchers used a test that measured levels of _____________ prejudice.

8) The judge said he would wait until Friday morning to _________________ a decision.

9) The fundamental laws of physics are in essence _____________ and need to be reformed.

10) Concept gadgets rule, editors _______________, and wireless, hands-free tech takes over.

Exercise IV.

Make up sentences of your own with the following word combinations:

to connect with, to go on dates with, to date online, to be about to change, to fall into, to fill out , to pair members by compatibility, to be rich in information, to take time, to misrepresent themselves

Exercise V.

Match the words to the definitions in the column on the right:

gimmick	a pear-shaped fruit with a rough leathery skin, smooth oily edible flesh, and a large stone
creepy	too great or extreme to be expressed or described in words
to drool	implied though not plainly expressed
to infer	no longer produced or used; out of date

to quantify	provide or give (a service, help, etc.)
ineffable	express or measure the quantity of
implicit	deduce or conclude (information) from evidence and reasoning rather than from explicit statements
to render	causing an unpleasant feeling of fear or unease
obsolete	a trick or device intended to attract attention, publicity, or business
avocado	saliva falling from the mouth

Exercise VI .

Identify the part of speech the words belong to. compatibility , contents, accurate, accessible, brutal, personal, personality, attraction, researchers, misrepresent

Exercise VII.

Match the words to make word combinations:

Dating	dating
user	University
personal	logic
Cambridge	algorithm
human	apps
personality	base
Instagram	essays
Twitter	matchmaking
good	photos

online	feeds

Exercise VIII.

Summarize the article “The Future of Online Dating”

4. Seduction, Inc

Exercise I.

Say what Russian words help to guess the meaning of the following words: industry, conference, instruction, techniques, forums, commercial, products, services, interest, style

Exercise II.

Make sure you know the following words and word combinations.
settling, contained, compelling, gimmicky, ostentatiously, to delineate, to harness, conjure, demise, to foster

Seduction, Inc

The pickup industry mates market logic with the arts of seduction – turning human intimacy into hard labour

Striding from the back of the conference room, the trainer calls for our attention. He asks everyone to explain why they’re here. The first student stands up: ‘I’ve come to get hands-on experience.’ Another says he has no problem meeting women, but for some reason he never manages to date the kind he really wants. The trainer nods with recognition: ‘Settling is the worst thing you can do. Because every time you see a guy with a hotter girl, you think: “I wish I was him.”’ One of the last men to introduce himself – visibly uncomfortable, shifting in his chair – starts to explain that he's a 'decent guy'. The trainer interrupts him: 'The problem is, you're not the guy that's going to take them home… We need to get you to be that guy.' On any given weekend, events such as this one in London take place in cities around the world – from New York to Tel Aviv, Stockholm to Mumbai. The attendees, largely in their 20s and 30s, receive detailed instruction in the so-called 'art of seduction': learning and rehearsing techniques to meet and seduce women. Commonly known as 'pickup' or 'game', the seduction industry first took shape in the United States in the early 2000s. What began as a few online forums soon gave rise to commercial products and services. Some of those with a personal interest in seduction began to style themselves as professionals, offering practical training and personal development for men who wanted greater choice and control in their intimate lives. While deploying the language of artistry – with terms such as 'pickup artist' (or 'PUA') – seduction trainers frame their activities as quasi-scientific endeavours, involving the development of hypotheses, strategic field-testing and the cultivation of expertise. Their thinking is often shaped by evolutionary psychology and management theory, particularly of the pop-sci and self-help variety, and comprises a suite of techniques that men can use to navigate their interactions with women. The basic precept is that male-female relations are subject to certain underlying principles that, once understood, can be readily manipulated. A typical training session might include instruction in female psychology and body language, alongside lessons in mindset and motivation. Students receive detailed guidance about how to approach and ‘open’, as well as about general conversation patterns. Further direction might include for dealing with resistance. With an emphasis on experiential learning, virtually all live events encompass an ‘in-field’ component where men approach women on the streets, in shops and cafés, at pubs and clubs. As well as choreographing men’s interactions with women, trainers observe and give feedback. Some use covert devices so as to watch interactions without the women’s awareness. In short, men are taught how to walk, talk, stand, speak, think and feel. Through seduction training, all aspects of the self are made available for assessment and improvement. The aim is not simply to impart a discrete skill set, but to inculcate deeper dispositions of body and mind based on a particular conception of what it means – and what it takes – to be a man.

Unsurprisingly for an industry that promises men ‘mastery’ with women, seduction has attracted a good deal of feminist commentary and criticism. Websites publish articles challenging its underlying assumptions and raising concerns about the propensity of its teachings to promote harassment, coercion and violence. Prominent seduction trainers are subject to campaigns seeking to restrict the availability of their products and services, and to limit their ability to travel internationally to teach. Feminists and others have good reasons for attacking the industry. Yet much existing commentary tends to parcel it off as anomalous – a subcultural oddity that, already contained, can be easily eliminated. In this way, those who participate in this sphere are framed as readily recognisable and uniquely deplorable – an ‘army of sleazebags and weirdos’, in the words of Hadley Freeman at The Guardian.But this underestimates the popular resonance of seduction techniques. While it's true that a relatively small number of men attend live training events, the sector's reach is large. Mailing lists of some seduction-training companies (often small enterprises in terms of staff and turnover) easily reach tens of thousands of readers. Online forums attract even larger numbers of commentators and browsers. Instructional videos posted on social media can accrue hundreds of thousands of views. Outside observers have paid scant attention to what makes seduction so compelling to so many men right now. What leads them to seek out this form of expertise? What kinds of problems are they hoping to address? What is it that they want to realise or achieve? To be clear, I'm not saying we should avoid criticisms of seduction – far from it. But I'm wary of how the tendency to isolate the industry – to section it off as an egregious subculture – prevents us from examining what its existence and appeal might reveal about contemporary patterns of lust and love. During my extensive ethnographic fieldwork, I came to see that seduction training is far more complex and disturbing than mainstream commentary suggests. Rather than being an anomaly, the seduction sector is evidence of how neoliberalism – as an economic system and cultural rationality – embeds market logic into the most intimate dimensions of our lives. Within a neoliberal context, the logic of competitive individualism has come to dominate spheres such as education and employment. Framing attraction as a skill that can be acquired, the industry channels this logic into the private realm. Men are told that they can achieve the kind of relationships they aspire to, provided they are willing to put in the necessary time, energy and, crucially, money. It's often assumed that men who seek out seduction training lack experience. However, the relationship histories of the men I interviewed varied widely. What they all had in common, however, was that they were uniformly dissatisfied with their lives The sociologist Eva Illouz demonstrates how romantic disappointment – generally perceived as a uniquely personal experience – is culturally patterned and commercially managed. Illouz argues that 'culturally induced desires create ordinary forms of suffering, such as chronic dissatisfaction, disappointment, and perpetual longing'. Her insight suggests that, in order to understand what brings men to seduction, we need to consider what's feeding their disappointment at a cultural level. Listening to men talk about what they wanted from relationships, what was most striking was the relentlessly aspirational ethos. Past partners were routinely held up and found wanting, often on aesthetic grounds, while many men enumerated detailed physical criteria to which women should adhere. Danny, one of the trainers, crisply articulated this preoccupation with attaining a higher 'calibre' of partner: 'The reason why a lot of guys want to do game is so that they can attract higher-value women. So they might be dating, say, fives and sixes, and they actually want to have a girl who's a 10 in terms of looks.' The market mentality underpinning this situation was not lost on Danny, who acknowledged: 'It's an exchange of values – "What can I get for what I'm offering?" It becomes very economical.'

In societies in thrall to market metrics, in which women's bodies are constantly held up to scrutiny, it's both lamentable and predictable that a philosophy of value-exchange should pervade how men relate to their partners. The desire to access 'high-value' women shows how the architecture of desire is being remade by consumer culture. As in so many other aspects of life, the ‘upgrade’ logic has taken hold. Over the years, techniques originating in the seduction industry have gained wider social purchase. The most well known also tend to be the most gimmicky, as with ‘peacocking’ (dressing ostentatiously to attract attention) and ‘negging’ (making backhanded compliments). Yet, on the whole, seduction is a more sophisticated enterprise than these examples suggest – and all the more insidious for it. Sublimated in all seduction practices are the unwritten ‘feeling rules’ – a phrase first used by the sociologist Arlie Russell Hochschild in The Managed Heart to delineate the social norms that shape how people try to feel (or not feel) in a given situation. Hochschild uses the example of the bride on her wedding day, who – knowing it is supposed to be the happiest day of her life – tries to feel happy. But precipitating the moment of marriage are the myriad socially regulated and culturally enforced norms about how a woman should feel when given a compliment, how she ought to respond when told she is desirable. Seduction methods tap into these emotional patterns. One popular model known as ‘daygame’ exhorts users to harness the power of ideals of romance, specifically romantic comedy films – a genre overwhelmingly aimed at and consumed by women. The website explains: ‘Once you learn how to strike up a conversation with a beautiful woman during the day, you’ll play into her fantasy of randomly meeting the guy of her dreams just like in the movies. She’ll believe that YOU are the guy of her dreams.’ What matters, of course, is not whether men really are ‘the guy of her dreams’, but rather than they can seem like they are – at least temporarily. Tom Torero, a noted proponent of this method, gives further details in his self-published book Daygame – the 'incredible story of Tom's journey from Oxford nerd to top street seducer', according to his website. Torero approaches women using the same basic opening line, continuously adapted so as to appear unique. Dates, too, are conducted according to a predefined script as Torero goes to a set venue, where he tells the same stories, makes the same jokes, asks the same questions. These interactions are not devoid of emotion. Rather, emotion is deployed tactically as a means to an end. Significant effort goes into affecting attraction, orchestrating desire, conjuring trust.

In books and blogs, seduction trainers document their transformation: from a past in which they were supposedly lonely and unpopular, to a present in which they enjoy near-constant access to beautiful women, plus an enviable lifestyle of world travel, financial independence and male camaraderie. Inevitably, not all those who use seduction methods actually find the relationships they're looking for. Many men I interviewed freely admitted that 'success with women' continued to elude them, even after months or years of training. Yet when they failed to master the 'art of seduction', these men almost always framed it as a personal failing – a testament to the depth of their own deficiency, or evidence that they were simply not trying hard enough. The tendency to blame oneself persisted even when seduction techniques had resolutely negative consequences: 'She just said I'd changed, and that she didn't know me anymore' he said, blinking in an effort to hold back the tears. 'She was the only thing I really cared about.' How did this make him feel about seduction? 'I'm mad and angry,' Anwar said. 'But not at pickup, I'm angry at me. Because it's my fault. It's a bit like you give me a set of tools and, if I didn't know how to use those tools properly, I'm going to make a mistake.' Anwar thus remained attached to the promise that, properly administered, seduction skills could furnish his desires. Anwar's insistence on blaming himself might seem illogical. Yet it's only by locating the fault within that he can sustain the fantasy that seduction will, eventually and effortfully, enable him to attain the relationship he desires. His engagement with seduction is a form of what the cultural theorist Lauren Berlant described in Cruel Optimism, when something we desire becomes an obstacle to our flourishing. For some men, seduction can become a consuming, even compulsive, pursuit. Derek had been a client of the industry for more than a year and had spent several thousand pounds on training courses. As we were waiting to buy coffee before our interview, I observed him talking to the woman behind the counter. His whole persona transformed, as he became suddenly playful and teasing. Later, he explained that he'd brought a different woman here every day this week, and wanted to be sure that the attractive barista took notice. Sitting down at a window table, Derek told me how seduction had changed his life for the better. He declared himself finally confident with women, and described the sense of freedom this gave him. But his tone shifted as he recalled the events of the previous evening: 'I was out last night, and I'm just walking around the streets at 11 o'clock at night, on my own, and I'm just like: "What am I doing out here?" It started out with some real motivation, desire to be good at this, but now it's almost fear of losing this ability.' In this, he exhibited symptoms of what the late cultural theorist Mark Fisher terms 'depressive hedonia'. Where depression is typically characterised by anhedonia, an inability to experience pleasure, depressive hedonia denotes 'an inability to do anything else except pursue pleasure'. As an affective condition, it exemplifies the profound insecurity that neoliberalism fosters at the level of subjectivity itself. For Derek, the need to maintain the skills he'd worked so hard to develop has become an end in itself. This industry's promise of control is itself seductive for many men – such that they find themselves continuing to invest even when the system does not serve their needs and interests. It might be referred to as 'game', but clearly seduction is not a recreational pastime or entertaining diversion. It's a form of labour that requires ongoing investment, often at considerable cost. The pursuit of relationships becomes a form of work: the work of seduction. By promoting an entrepreneurial solution to the problem of finding and forging intimate relationships, the seduction industry shows up some of the most dubious tendencies of modern culture. Individual self-work is prescribed as the solution for problems that are culturally and socially shaped. Labour-intensive and profit-orientated modes of socialising end up eclipsing other forms of being and relating. Ethical concerns are cast aside in favour of personal promotion and unencumbered self-interest. While seduction training is often framed as a deviant subculture, the men involved are entirely ordinary. If their desires and discontents strike us as strange, perhaps we should look more closely at the context in which they have been formed. To ask what makes seduction so compelling for those drawn into its folds is not to dispense with critique; rather, it is to insist on the necessity of posing difficult questions and having uncomfortable conversations.

Adapted from Aeon

Exercise III.

Fill in the gaps.

1) Born in Poland, she pursued a singing career before _____________down in California.

2) You are absolutely right that the XBox and PS3 do not make good ______________ computers.

3) A great documentary does it with an ______________ that rivals the best fiction films.

4) In other words, this study represents only the first stage of _______________.

5) Researchers also will conduct a second ______________ to study how far bees travel.

6) The term electrical engineering may or may not ______________electronic engineering.

7) Is a group's __________________ to click on ads inversely related to its tech savvyness?

8) You can't stop your partner being on the internet, it makes you look like a _________________.

9) As _______________ as that witness behavior was, they can begin to set things right.

10) According to an admittedly _______________ online source, her spirit animal is the hawk.

Exercise IV.

Make up sentences of your own with the following word combinations:

to parcel off, to impart, to inculcate, to stride, to encompass, to eliminate, to dispense, to adhere, to thrall, to forge

Exercise V.

Match the words to the definitions in the column on the right:

artistry	(esp. of a smell) spread through and be perceived in every part of
quasi	not openly acknowledged or displayed
precept	a person whose dress or behavior seems strange or eccentric
mindset	barely sufficient or adequate
covert	critical observation or examination
weirdo	a general rule intended to regulate behavior or thought
scant	having some resemblance
scrutiny	creative skill or ability
to pervade	the ideas and attitudes with which a person approaches a situation, esp when these are seen as being difficult to alter.

Exercise VI.

Identify the part of speech the words belong to: propensity, coercion, anomalous, deplorable, egregious, perpetual, relentlessly, aspirational, insidious, preoccupation

Exercise VII.

Match the words to make word combinations:

commercial	interest
scientific	experience
field	room
found	forums
personal	products
online	wanting
detailed	intimacy
hands-on	test
human	instruction
conference	endeavor

Exercise VIII .

Summarize the article “Seduction, Inc.”

SUPPLEMENTARY READING

The Shallowness of Google Translate

The program uses state-of-the-art AI techniques, but simple tests show that it's a long way from real understanding.

One Sunday, at one of our weekly salsa sessions, my friend Frank brought along a Danish guest. I knew Frank spoke Danish well, since his mother was Danish, and he, as a child, had lived in Denmark. As for his friend, her English was fluent, as is standard for Scandinavians. However, to my surprise, during the evening's chitchat it emerged that the two friends habitually exchanged emails using Google Translate. Frank would write a message in English, then run it through Google Translate to produce a new text in Danish; conversely, she would write a message in Danish, then let Google Translate anglicize it. How odd! Why would two intelligent people, each of whom spoke the other's language well, do this? My own experiences with machine-translation software had always led me to be highly skeptical about it. But my skepticism was clearly not shared by these two. Indeed, many thoughtful people are quite enamored of translation programs, finding little to criticize in them. This baffles me. As a language lover and an impassioned translator, as a cognitive scientist and a lifelong admirer of the human mind's subtlety, I have followed the attempts to mechanize translation for decades. When I first got interested in the subject, in the mid-1970s, I ran across a letter written in 1947 by the mathematician Warren Weaver, an early machine-translation advocate, to Norbert Wiener, a key figure in cybernetics, in which Weaver made this curious claim, today quite famous: When I look at an article in Russian, I say, "This is really written in English, but it has been coded in some strange symbols. I will now proceed to decode."

Some years later he offered a different viewpoint: "No reasonable person thinks that a machine translation can ever achieve elegance and style. Pushkin need not shudder." Whew! Having devoted one unforgettably intense year of my life to translating Alexander Pushkin's sparkling novel in verse Eugene Onegin into my native tongue (that is, having radically reworked that great Russian work into an English-language novel in verse), I find this remark of Weaver's far more congenial than his earlier remark, which reveals a strangely simplistic view of language. Nonetheless, his 1947 view of translation-as-decoding became a credo that has long driven the field of machine translation.

Since those days, "translation engines" have gradually improved, and recently the use of so-called "deep neural nets" has even suggested to some observers (see "The Great AI Awakening" by Gideon Lewis-Kraus in The New York Times Magazine, and "Machine Translation: Beyond Babel" by Lane Greene in The Economist) that human translators may be an endangered species. In this scenario, human translators would become, within a few years, mere quality controllers and glitch fixers, rather than producers of fresh new text.

Such a development would cause a soul-shattering upheaval in my mental life. Although I fully understand the fascination of trying to get machines to translate well, I am not in the least eager to see human translators replaced by inanimate machines. Indeed, the idea frightens and revolts me. To my mind, translation is an incredibly subtle art that draws constantly on one's many years of experience in life, and on one's creative imagination. If, some "fine" day, human translators were to become relics of the past, my respect for the human mind would be profoundly shaken, and the shock would leave me reeling with terrible confusion and immense, permanent sadness.

Each time I read an article claiming that the guild of human translators will soon be forced to bow down before the terrible swift sword of some new technology, I feel the need to check the claims out myself, partly out of a sense of terror that this nightmare just might be around the corner, more hopefully out of a desire to reassure myself that it's not just around the corner, and finally, out of my longstanding belief that it's important to combat exaggerated claims about artificial intelligence. And so, after reading about how the old idea of artificial neural networks, recently adopted by a branch of Google called Google Brain, and now enhanced by "deep learning," has resulted in a new kind of software that has allegedly revolutionized machine translation, I decided I had to check out the latest incarnation of Google Translate. Was it a game changer, as Deep Blue and AlphaGo were for the venerable games of chess and Go?

I learned that although the older version of Google Translate can handle a very large repertoire of languages, its new deep-learning incarnation at the time worked for just nine languages. (It's now expanded to 96.) Accordingly, I limited my explorations to English, French, German, and Chinese.

Before showing my findings, though, I should point out that an ambiguity in the adjective "deep" is being exploited here. When one hears that Google bought a company called DeepMind whose products have "deep neural networks" enhanced by "deep learning," one cannot help taking the word "deep" to mean "profound," and thus "powerful," "insightful," "wise." And yet, the meaning of "deep" in this context comes simply from the fact that these neural networks have more layers (12, say) than do older networks, which might have only two or three. But does that sort of depth imply that whatever such a network does must be profound? Hardly. This is verbal spinmeistery.

I am very wary of Google Translate, especially given all the hype surrounding it. But despite my distaste, I recognize some astonishing facts about this bête noire of mine. It is accessible for free to anyone on earth, and will convert text in any of roughly 100 languages into text in any of the others. That is humbling. If I am proud to call myself "pi-lingual" (meaning the sum of all my fractional languages is a bit over 3, which is my lighthearted way of answering the question "How many languages do you speak?"), then how much prouder should Google Translate be, since it could call itself "bai-lingual" ("bai" being Mandarin for 100). To a mere pilingual, bailingualism is most impressive. Moreover, if I copy and paste a page of text in Language A into Google Translate, only moments will elapse before I get back a page filled with words in Language B. And this is happening all the time on screens all over the planet, in dozens of languages.

The practical utility of Google Translate and similar technologies is undeniable, and probably it's a good thing overall, but there is still something deeply lacking in the approach, which is conveyed by a single word: understanding. Machine translation has never focused on understanding language. Instead, the field has always tried to "decode"—to get away without worrying about what understanding and meaning are. Could it in fact be that understanding isn't needed in order to translate well? Could an entity, human or machine, do high-quality translation without paying attention to what language is all about? To shed some light on this question, I turn now to the experiments I made.

I began my explorations very humbly, using the following short remark, which, in a human mind, evokes a clear scenario:

In their house, everything comes in pairs. There's his car and her car, his towels and her towels, and his library and hers.

The translation challenge seems straightforward, but in French (and other Romance languages), the words for "his" and "her" don't agree in gender with the possessor, but with the item possessed. So here's what Google Translate gave me:

Dans leur maison, tout vient en paires. Il y a sa voiture et sa voiture, ses serviettes et ses serviettes, sa bibliothèque et les siennes.

The program fell into my trap, not realizing, as any human reader would, that I was describing a couple, stressing that for each item he had, she had a similar one. For example, the deep-learning engine used the word "sa" for both "his car" and "her car," so you can't tell anything about either car-owner's gender. Likewise, it used the genderless plural "ses" both for "his towels" and "her towels," and in the last case of the two libraries, his and hers, it got thrown by the final "s" in "hers" and somehow decided that that "s" represented a plural ("les siennes"). Google Translate's French sentence missed the whole point.

Next I translated the challenge phrase into French myself, in a way that did preserve the intended meaning. Here's my French version:

Chez eux, ils ont tout en double. Il y a sa voiture à elle et sa voiture à lui, ses serviettes à elle et ses serviettes à lui, sa bibliothèque à elle et sa bibliothèque à lui.

The phrase "sa voiture à elle" spells out the idea "her car," and similarly, "sa voiture à lui" can only be heard as meaning "his car." At this point, I figured it would be trivial for Google Translate to carry my French translation back into English and get the English right on the money, but I was dead wrong. Here's what it gave me:

At home, they have everything in double. There is his own car and his own car, his own towels and his own towels, his own library and his own library.

What?! Even with the input sentence screaming out the owners' genders as loudly as possible, the translating machine ignored the screams and made everything masculine. Why did it throw the sentence's most crucial information away?

We humans know all sorts of things about couples, houses, personal possessions, pride, rivalry, jealousy, privacy, and many other intangibles that lead to such quirks as a married couple having towels embroidered “his” and “hers.” Google Translate isn’t familiar with such situations. Google Translate isn’t familiar with situations, period. It’s familiar solely with strings composed of words composed of letters. It’s all about ultrarapid processing of pieces of text, not about thinking or imagining or remembering or understanding. It doesn’t even know that words stand for things. Let me hasten to say that a computer program certainly could, in principle, know what language is for, and could have ideas and memories and experiences, and could put them to use, but that’s not what Google Translate was designed to do. Such an ambition wasn’t even on its designers’ radar screens.

Well, I chuckled at these poor shows, relieved to see that we aren’t, after all, so close to replacing human translators by automata. But I still felt I should check the engine out more closely. After all, one swallow does not thirst quench.

Indeed, what about this freshly coined phrase “One swallow does not thirst quench” (alluding, of course, to “One swallow does not a summer make”)? I couldn’t resist trying it out; here’s what Google Translate flipped back at me: “Une hirondelle n’aspire pas la soif.” This is a grammatical French sentence, but it’s pretty hard to fathom. First it names a certain bird (“une hirondelle”—a swallow), then it says this bird is not inhaling or not sucking (“n’aspire pas”), and finally reveals that the neither-inhaled-nor-sucked item is thirst (“la soif”). Clearly Google Translate didn’t catch my meaning; it merely came out with a heap of bull. “Il sortait simplement avec un tas de taureau.” “He just went out with a pile of bulls.” “Il vient de sortir avec un tas de taureaux.” Please pardon my French—or rather, Google Translate’s pseudo-French.

From the frying pan of French, let’s jump into the fire of German. Of late I’ve been engrossed in the book Sie nannten sich der Wiener Kreis (They Called Themselves the Vienna Circle), by the Austrian mathematician Karl Sigmund. It describes a group of idealistic Viennese intellectuals in the 1920s and 1930s, who had a major impact on philosophy and science during the rest of the century. I chose a short passage from Sigmund’s book and gave it to Google Translate. Here it is, first in German, followed by my own translation, and then Google Translate’s version. (By the way, I checked my translation with two native speakers of German, including Karl Sigmund, so I think you can assume it is accurate.)

Sigmund:

Nach dem verlorenen Krieg sahen es viele deutschnationale Professoren, inzwischen die Mehrheit in der Fakultät, gewissermaßen als ihre Pflicht an, die Hochschulen vor den “Ungeraden” zu bewahren; am schutzlosesten waren junge Wissenschaftler vor ihrer Habilitation. Und Wissenschaftlerinnen kamen sowieso nicht in frage; über wenig war man sich einiger.

Hofstadter:

After the defeat, many professors with Pan-Germanistic leanings, who by that time constituted the majority of the faculty, considered it pretty much their duty to protect the institutions of higher learning from “undesirables.” The most likely to be dismissed were young scholars who had not yet earned the right to teach university classes. As for female scholars, well, they had no place in the system at all; nothing was clearer than that.

Google Translate:

After the lost war, many German-National professors, meanwhile the majority in the faculty, saw themselves as their duty to keep the universities from the “odd”; Young scientists were most vulnerable before their habilitation. And scientists did not question anyway; There were few of them.

The words in Google Translate’s output are all English words (even if, for unclear reasons, a couple are inappropriately capitalized). So far, so good! But soon it grows wobbly, and the further down you go the wobblier it gets.

I’ll focus first on “the ‘odd.’” This corresponds to the German “die ‘Ungeraden,’” which here means “politically undesirable people.” Google Translate, however, had a reason—a very simple statistical reason—for choosing the word “odd.” Namely, in its huge bilingual database, the word “ungerade” was almost always translated as “odd.” Although the engine didn’t realize why this was the case, I can tell you why. It’s because “ungerade”—which literally means “un-straight” or “uneven”—nearly always means “not divisible by two.” By contrast, my choice of “undesirables” to render “Ungeraden” had nothing to do with the statistics of words, but came from my understanding of the situation—from my zeroing in on a notion not explicitly mentioned in the text and certainly not listed as a translation of “ungerade” in any of my German dictionaries.

Let’s move on to the German “Habilitation,” denoting a university status resembling tenure. The English cognate word “habilitation” exists but it is super-rare, and certainly doesn’t bring to mind tenure or anything like it. That’s why I briefly explained the idea rather than just quoting the obscure word, since that mechanical gesture would not get anything across to anglophonic readers. Of course Google Translate would never do anything like this, as it has no model of its readers’ knowledge.

The last two sentences really bring out how crucial understanding is for translation. The 15-letter German noun “Wissenschaftler” means either “scientist” or “scholar.” (I opted for the latter, as in this context it was referring to intellectuals in general. Google Translate didn’t get that subtlety.) The related 17-letter noun “Wissenschaftlerin,” found in the closing sentence in its plural form “Wissenschaftlerinnen,” is a consequence of the gendered-ness of German nouns. Whereas the “short” noun is grammatically masculine and thus suggests a male scholar, the longer noun is feminine and applies to females only. I wrote “female scholar” to get the idea across. Google Translate, however, did not understand that the feminizing suffix “-in” was the central focus of attention in the final sentence. Since it didn’t realize that females were being singled out, the engine merely reused the word “scientist,” thus missing the sentence’s entire point. As in the earlier French case, Google Translate didn’t have the foggiest idea that the sole purpose of the German sentence was to shine a spotlight on a contrast between males and females.

Aside from that blunder, the rest of the final sentence is a disaster. Take its first half. Is "scientists did not question anyway" really a translation of "Wissenschaftlerinnen kamen sowieso nicht in frage"? It doesn't mean what the original means—it's not even in the same ballpark. It just consists of English words haphazardly triggered by the German words. Is that all it takes for a piece of output to deserve the label "translation"?

The sentence's second half is equally erroneous. The last six German words mean, literally, "over little was one more united," or, more flowingly, "there was little about which people were more in agreement," yet Google Translate managed to turn that perfectly clear idea into "There were few of them." We baffled humans might ask "Few of what?" but to the mechanical listener, such a question would be meaningless. Google Translate doesn't have ideas behind the scenes, so it couldn't even begin to answer the simple-seeming query. The translation engine was not imagining large or small amounts or numbers of things. It was just throwing symbols around, without any notion that they might symbolize something.

It's hard for a human, with a lifetime of experience and understanding and of using words in a meaningful way, to realize how devoid of content all the words thrown onto the screen by Google Translate are. It's almost irresistible for people to presume that a piece of software that deals so fluently with words must surely know what they mean. This classic illusion associated with artificial-intelligence programs is called the "Eliza effect," since one of the first programs to pull the wool over people's eyes with its seeming understanding of English, back in the 1960s, was a vacuous phrase manipulator called Eliza, which pretended to be a psychotherapist, and as such, it gave many people who interacted with it the eerie sensation that it deeply understood their innermost feelings.

For decades, sophisticated people—even some artificial-intelligence researchers—have fallen for the Eliza effect. In order to make sure that my readers steer clear of this trap, let me quote some phrases from a few paragraphs up—namely, "Google Translate did not understand," "it did not realize," and "Google Translate didn't have the foggiest idea." Paradoxically, these phrases, despite harping on the lack of understanding, almost suggest that Google Translate might at least sometimes be capable of understanding what a word or a phrase or a sentence means, or is about. But that isn't the case. Google Translate is all about bypassing or circumventing the act of understanding language.

To me, the word "translation" exudes a mysterious and evocative aura. It denotes a profoundly human art form that graciously carries clear ideas in Language A into clear ideas in Language B, and the bridging act not only should maintain clarity, but also should give a sense for the flavor, quirks, and idiosyncrasies of the writing style of the original author. Whenever I translate, I first read the original text carefully and internalize the ideas as clearly as I can, letting them slosh back and forth in my mind. It's not that the words of the original are sloshing back and forth; it's the ideas that are triggering all sorts of related ideas, creating a rich halo of related scenarios in my mind. Needless to say, most of this halo is unconscious. Only when the halo has been evoked sufficiently in my mind do I start to try to express it—to "press it out"—in the second language. I try to say in Language B what strikes me as a natural B-ish way to talk about the kinds of situations that constitute the halo of meaning in question.

I am not, in short, moving straight from words and phrases in Language A to words and phrases in Language B. Instead, I am unconsciously conjuring up images, scenes, and ideas, dredging up experiences I myself have had (or have read about, or seen in movies, or heard from friends), and only when this nonverbal, imagistic, experiential, mental "halo" has been realized—only when the elusive bubble of meaning is floating in my brain—do I start the process of formulating words and phrases in the target language, and then revising, revising, and revising. This process, mediated via meaning, may sound sluggish, and indeed, in comparison with Google Translate's two or three seconds per page, it certainly is—but it is what any serious human translator does. This is the kind of thing I imagine when I hear an evocative phrase like "deep mind."

Of course I grant that Google Translate sometimes comes up with a series of output sentences that sound fine (although they may be misleading or utterly wrong). A whole paragraph or two may come out superbly, giving the illusion that Google Translate knows what it is doing, understands what it is "reading." In such cases, Google Translate seems truly impressive—almost human! Praise is certainly due to its creators and their collective hard work. But at the same time, don't forget what Google Translate did with these two Chinese passages, and with the earlier French and German passages. To understand such failures, one has to keep the ELIZA effect in mind. The bailingual engine isn't reading anything—not in the normal human sense of the verb "to read." It's processing text. The symbols it's processing are disconnected from experiences in the world. It has no memories on which to draw, no imagery, no understanding, no meaning residing behind the words it so rapidly flings around.

A friend asked me whether Google Translate's level of skill isn't merely a function of the program's database. He figured that if you multiplied the database by a factor of, say, a million or a billion, eventually it would be able to translate anything thrown at it, and essentially perfectly. I don't think so. Having ever more "big data" won't bring you any closer to understanding, since understanding involves having ideas, and lack of ideas is the root of all the problems for machine translation today. So I would venture that bigger databases—even vastly bigger ones—won't turn the trick.

Another natural question is whether Google Translate's use of neural networks—a gesture toward imitating brains—is bringing us closer to genuine understanding of language by machines. This sounds plausible at first, but there's still no attempt being made to go beyond the surface level of words and phrases. All sorts of statistical facts about the huge databases are embodied in the neural nets, but these statistics merely relate words to other words, not to ideas. There's no attempt to create internal structures that could be thought of as ideas, images, memories, or experiences. Such mental etherea are still far too elusive to deal with computationally, and so, as a substitute, fast and sophisticated statistical word-clustering algorithms are used. But the results of such techniques are no match for actually having ideas involved as one reads, understands, creates, modifies, and judges a piece of writing.

Despite my negativism, Google Translate offers a service many people value highly: It effects quick-and-dirty conversions of meaningful passages written in language A into not necessarily meaningful strings of words in language B. As long as the text in language B is somewhat comprehensible, many people feel perfectly satisfied with the end product. If they can "get the basic idea" of a passage in a language they don't know, they're happy. This isn't what I personally think the word "translation" means, but to some people it's a great service, and to them it qualifies as translation. Well, I can see what they want, and I understand that they're happy. Lucky them!

I've recently seen bar graphs made by technophiles that claim to represent the "quality" of translations done by humans and by computers, and these graphs depict the latest translation engines as being within striking distance of human-level translation. To me, however, such quantification of the unquantifiable reeks of pseudoscience, or, if you prefer, of nerds trying to mathematize things whose intangible, subtle, artistic nature eludes them. To my mind, Google Translate's output today ranges all the way from excellent to grotesque, but I can't quantify my feelings about it. Think of my first example involving "his" and "her" items. The idealess program got nearly all the words right, but despite that slight success, it totally missed the point. How, in such a case, should one "quantify" the quality of the job? The use of scientific-looking bar graphs to represent translation quality is simply an abuse of the external trappings of science.

Let me return to that sad image of human translators, soon outdone and outmoded, gradually turning into nothing but quality controllers and text tweakers. That's a recipe for mediocrity at best. A serious artist doesn't start with a kitschy piece of error-ridden bilgewater and then patch it up here and there to produce a work of high art. That's not the nature of art. And translation is an art.

In my writings over the years, I've always maintained that the human brain is a machine—a very complicated kind of machine—and I've vigorously opposed those who say that machines are intrinsically incapable of dealing with meaning. There is even a school of philosophers who claim computers could never "have semantics" because they're made of "the wrong stuff" (silicon). To me, that's facile nonsense. I won't touch that debate here, but I wouldn't want to leave readers with the impression that I believe intelligence and understanding to be forever inaccessible to computers. If in this essay I seem to come across sounding that way, it's because the technology I've been discussing makes no attempt to reproduce human intelligence. Quite the contrary: It attempts to make an end run around human intelligence, and the output passages exhibited above clearly reveal its giant lacunas.

From my point of view, there is no fundamental reason that machines could not, in principle, someday think, be creative, funny, nostalgic, excited, frightened, ecstatic, resigned, hopeful, and, as a corollary, able to translate admirably between languages. There's no fundamental reason that machines might not someday succeed smashingly in translating jokes, puns, screenplays, novels, poems, and, of course, essays like this one. But all that will come about only when machines are as filled with ideas, emotions, and experiences as human beings are. And that's not around the corner. Indeed, I believe it is still extremely far away. At least that is what this lifelong admirer of the human mind's profundity fervently hopes.

When, one day, a translation engine crafts an artistic novel in verse in English, using precise rhyming iambic tetrameter rich in wit, pathos, and sonic verve, then I'll know it's time for me to tip my hat and bow out.

Adapted from The Atlantic

Style Is an Algorithm

No one is original anymore, not even you.

The camera is a small, white, curvilinear monolith on a pedestal. Inside its smooth casing are a microphone, a speaker, and an eye-like lens. After I set it up on a shelf, it tells me to look straight at it and to be sure to smile! The light blinks and then the camera flashes. A head-to-toe picture appears on my phone of a view I'm only used to seeing in large mirrors: me, standing awkwardly in my apartment, wearing a very average weekday outfit. The background is blurred like evidence from a crime scene. It is not a flattering image.

Amazon's Echo Look, currently available by invitation only but also on eBay, allows you to take hands-free selfies and evaluate your fashion choices. "Now Alexa helps you look your best," the product description promises. Stand in front of the camera, take photos of two different outfits with the Echo Look, and then select the best ones on your phone's Echo Look app. Within about a minute, Alexa will tell you which set of clothes looks better, processed by style-analyzing algorithms and some assistance from humans. So I try to find my most stylish outfit, swapping out shirts and pants and then posing stiffly for the camera. I shout, "Alexa, judge me!" but apparently that's unnecessary.

What I discover from the Style Check™ function is as follows: All-black is better than all-gray. Rolled-up sleeves are better than buttoned at the wrist. Blue jeans are best. Popping your collar is actually good. Each outfit in the comparison receives a percentage out of 100: black clothes score 73 percent against gray clothes at 27 percent, for example. But the explanations given for the scores are indecipherable. "The way you styled those pieces looks better," the app tells me. "Sizing is better." How did I style them? Should they be bigger or smaller?

The Echo Look won't tell you why it's making its decisions. And yet it purports to show us our ideal style, just as algorithms like Netflix recommendations, Spotify Discover, and Facebook and YouTube feeds promise us an ideal version of cultural consumption tailored to our personal desires. In fact, this promise is inherent in the technology itself: Algorithms, as I'll loosely define them, are sets of equations that work through machine learning to customize the delivery of content to individuals, prioritizing what they think we want, and evolving over time based on what we engage with.

Confronting the Echo Look's opaque statements on my fashion sense, I realize that all of these algorithmic experiences are matters of taste: the question of what we like and why we like it, and what it means that taste is increasingly dictated by black-box robots like the camera on my shelf.

In his 2017 book Taste, the Italian philosopher Giorgio Agamben digs up the roots of the word. Historically, it is defined as a form of knowledge through pleasure, from perceiving the flavor of food to judging the quality of an object. Taste is an essentially human capacity, to the point that it is almost subconscious: We know whether we like something or not before we understand why. "Taste enjoys beauty, without being able to explain it," Agamben writes. He quotes Montesquieu: "This effect is principally founded on surprise." Algorithms are meant to provide surprise, showing us what we didn't realize we'd always wanted, and yet we are never quite surprised because we know to expect it.

Philosophers in the 18th century defined taste as a moral capacity, an ability to recognize truth and beauty. "Natural taste is not a theoretical knowledge; it's a quick and exquisite application of rules which we do not even know," wrote Montesquieu in 1759. This unknowingness is important. We don't calculate or measure if something is tasteful to us; we simply feel it. Displacing the judgment of taste partly to algorithms, as in the Amazon Echo Look, robs us of some of that humanity.

Every cultural object we aestheticize and consume — "the most everyday choices of everyday life, e.g., in cooking, clothing or decoration," Pierre Bourdieu writes in his 1984 book Distinction: A Social Critique of the Judgement of Taste — is a significant part of our identities and reflects who we are. "Taste classifies, and it classifies the classifier," Bourdieu adds. If our taste is dictated by data-fed algorithms controlled by massive tech corporations, then we must be content to classify ourselves as slavish followers of robots.

We might say that "taste" is the abstract, moralized knowledge, while "style" is its visual expression. Fashion makes taste easily visible as style, in part because its distinctions between color or cut in clothing are so specific and yet so random ("rules which we don't even know"). In the past, a whimsical consensus among elites dictated fashion culture; a royal court or an echelon of magazine editors imposed a certain taste from the top of society, down. Roland Barthes noticed this arbitrariness in his 1960 essay Blue Is in Fashion This Year. Barthes scrutinizes a fragment of text from a fashion magazine — "blue is in fashion this year" — to see where its thesis, that a particular color is particularly tasteful right now, comes from. His conclusion is that it doesn't come from anywhere: "We are not talking about a rigorous production of meaning: the link is neither obligatory nor sufficiently motivated." Blue is not in fashion because it is particularly functional, nor is it symbolically linked to some wider economic or political reality; the statement has no semantic logic. Style, Barthes argues, is an inexplicable equation (a faulty algorithm). Further evidence of the artificial and hierarchical nature of style in the past can be found in that scene from the 2006 film The Devil Wears Prada, in which Meryl Streep (as magazine editor and Anna Wintour facsimile Miranda Priestly) tells her assistant played by Anne Hathaway that the chunky blue sweater she is wearing was, in essence, chosen for her. “That blue represents millions of dollars and countless jobs, and it’s sort of comical how you think you made a choice that exempts you from the fashion industry when, in fact, you’re wearing a sweater that was selected for you by the people in this room from a pile of stuff,” Streep says. In other words, blue is in fashion this year because some people decided it was. You, the non-tastemaker, have no choice in the matter.

Is it possible that instead of this artificial fashion language, algorithms like those powering Alexa could create a more systemic, logical construction of fashion aesthetics built on data? Blue is in fashion this year because 83.7 percent of users purchased (or clicked like on) blue shirts, the Amazon Echo Look algorithm says, therefore it is in fashion, therefore businesses should manufacture more blue shirts, and you, the customer, will buy and wear them. No human editors needed.

I’m not sure if this technology-derived algorithmic facticity of taste is better or worse than Meryl Streep-Anna Wintour deciding what I wear, which might be the core concern of this essay.

When modes of tastes change, there is a certain fear: Am I in or out? Do I understand the new or am I stuck in the old? In 1980, the New Yorker published George W.S. Trow’s essay describing this feeling under the title of “Within the Context of No Context,” from which I took the epigraph and structure for this piece. Trow’s essay came out as a book in 1981 and again in 1997. In the appended introduction to the 1997 edition, he uses the phrase “collapsing dominant” to describe a situation in which an older, established mode of cultural authority, or a taste regime, is fading and being replaced by a newer one. These regimes have two parts: the subjects of taste and the way taste is communicated.

Today we are seeing the collapse of the dominant regime that Trow originally observed emerging, mass-media television, which had previously replaced the moralistic mid-century novels of New England WASPs. Now, we have Instagram likes, Twitter hashtags, and Google-distributed display advertising spreading taste values. Instead of the maximalist, celebrity-driven, intoxicant culture of ‘70s television — Nixon, Star Wars, shag rugs, cocaine, nuclear bombs — we now have the flattened, participatory, somehow salutary aesthetic of avocado toast, Outdoor Voices leggings, reclaimed wood, Sky Ting yoga classes, and succulents in ceramic planters.

That we are in the midst of this shift in taste might help explain our larger mood of instability and paranoia (or is it just me?). We can’t figure out what might be sustainable to identify with, to orient our taste on. The algorithm suggests that we trust it, but we don’t entirely want to. We crave a more “authentic,” lasting form of meaning.

In 2009, a designer named Ben Pieratt, now living in Massachusetts, launched Svpply. It was a kind of online social network based on shopping, where invitation-only members could curate selections of products from elsewhere on the internet and users could follow their favorite tastemakers. Eventually, any user could become a curator. I remember it from the time as a calm, limpid pool in the midst of so much internet noise. The site presented only cool clothes, bags, and accessories, all chosen by individual humans, since algorithmic feeds weren't widely deployed at the time. On Svpply you could find the melange of signifiers of a certain class of early-adopter design-bro: minimalist sneakers, fancy T-shirts, Leica cameras, and drop-crotch sweatpants. In 2012, eBay acquired the company and quickly shut it down. In 2014, Pieratt launched a Kickstarter for Very Goods, a Svpply replacement that's still active. Today he sees Svpply as a cautionary tale about the limits of human curation on the internet. Over the phone, we talk about how taste doesn't really scale. The bigger a platform gets, the harder it is to maintain a particular sense of style. By opening the platform, Pieratt had tried to "convert from a human-driven community into a machine," he explains. "When we lost the exclusivity, people didn't really care anymore." Svpply's innate sense of uniqueness didn't survive: "If everyone's editing Vogue, it wouldn't be Vogue."

Another question: How good of a tastemaker can a machine ultimately be?

I worry that we are moving from a time of human curation (early Svpply) to a time in which algorithms drive an increasingly large portion of what we consume (the Facebook feed). This impacts not only the artifacts we experience but also how we experience them. Think of the difference between a friend recommending a clothing brand and something showing up in targeted banner ads, chasing you around the internet. It's more likely that your friend understands what you want and need, and you're more likely to trust the recommendation, even if it seems challenging to you.

Maybe it's a particularly shapeless garment or a noisy punk track. If you know the source of the suggestion, then you might give it a chance and see if it meshes with your tastes. In contrast, we know the machine doesn't care about us, nor does it have a cultivated taste of its own; it only wants us to engage with something it calculates we might like. This is boring. "I wonder if, at the core of fashion, the reason we find it fascinating is that we know there's a human at the end of it," Pieratt says. "We're learning about people. If you remove that layer of humanity from underneath, does the soul of the interest leave with it?"

Pieratt makes a further distinction between style and taste. Style is a superficial aesthetic code that is relatively simple to replicate, whereas taste is a kind of wider aesthetic intelligence, able to connect and integrate disparate experiences. Algorithms can approximate the former — telling me I should wear a blue shirt — but can't approximate the latter because the machine can't tell me why it thinks I should wear a blue shirt or what the blue shirt might mean to me. When a machine has taken over the exploration of taste, the possibility of suddenly feeling something from a surprising object is narrowed to only what the machine decides to expose. "I don't think there's such a thing as machine taste yet," says Pieratt.

Of course, he and I might just be part of the fading regime, our "collapsing dominant." The dystopian babies raised on algorithmic Spiderman-slash-Frozen YouTube videos may have different appetites in the future.

The threat of banality (or the lack of surprise) implicit in full machine curation reminds me of the seemingly random vocabulary meant to improve SEO on Craigslist posts. As one chair listing I encountered put it: "Goes with herman miller eames vintage mid century modern knoll Saarinen dwr design within reach danish denmark abc carpet and home arm chair desk dining slipper bedroom living room office."

Imagine the optimized average of all of these ideas. The linguistic melange forms a taste vernacular built not on an individual brand identity or a human curator but a freeform mass of associations meant to draw the viewer in by any means necessary. If you like this, you'll probably like that. Or, as a T-shirt I bought in Cambodia a decade ago reads, "Same same but different." The slogan pops into my mind constantly as I scroll past so many content modules, each unique and yet unoriginal.

Algorithms promise: If you like this, you will get more of it, forever. This experience is leaking from the internet of Google ads for the bag you just bought into the physical world. Look to the artist Jenny Odell's investigation of "free watch" offers on Instagram for an example. The watches appear, at a minimum, stylish, with small variations on minimalist faces and metal bands. But they are not the result of an enlightened sense of taste, per Pieratt's definition. The brands that sell them are thin fictions whipped up in Squarespace and the actual products are the result of Alibaba manufacturing and Amazon drop-shipping, in which a product moves directly from manufacturer to consumer having never entered a store. The phantom watches are empty fashion language, objects without content.

Other ways in which our experiences are warped by algorithmic platforms include Spotify possibly commissioning original music from "fake" artists to match the latent content desires of its audience, as Noisey noticed; delivery restaurants that are only virtual, conjuring a digital brand out of a shadowy group kitchen and serving food via Uber Eats; the surreal kids' YouTube videos, which exist because they are rewarded with views by the feed algorithm and thus earn their creators advertising profit; and the globalized visual vernacular of Airbnb interior decorating, which approximates a certain style emerging from the platform itself. Having analyzed the data from some platform or another, these are things the machine thinks you want, and it can serve them up immediately and infinitely.

We find ourselves in a cultural uncanny valley, unable to differentiate between things created by humans and those generated by a human-trained equation run amok. In other words, what is the product of genuine taste and what is not. (This lack of discernibility also contributes to the problems of fake news, which algorithmic feeds promote like any other content, however inaccurate.)

Spotify's fake artists aren't fake, per se; they're a kind of muzak created by a Swedish production company that just so happens to have the same investors as Spotify. That the simple possibility of non-genuine music fed to us by an algorithmic platform without our knowledge created a media frenzy speaks to our fundamental fear — a possibly irrational or at least abstruse 21st-century anxiety — of an algorithmic culture.

In 1935, Walter Benjamin observed that the work of art in the 20th century was undergoing a change during the advent of photography and film. The newfound reproducibility of the individual work of art through these technologies meant that art was deprived of its "aura": "the here and now of the original" or "the abstract idea of its genuineness," as Benjamin writes.

Photography, as Benjamin observed, could reproduce a singular work of art. Algorithmic machine learning, however, can mimic an entire stylistic mode, generating new examples at will or overlaying a pre-existing object with a new style unrelated to its origins. In 2015, researchers released a paper in which they turned a photograph of Tübingen, Germany into a van Gogh painting, then overlaid the style of Munch and Kandinsky in turn. The system "achieves a separation of image content from style," the researchers write (a disconnect that contributes to our anxiety).

So it's not just an individual work which can be reproduced, but rather an artist's entire aesthetic. The resulting lack of aura devalues unique style, or changes our experience of it, just as photography once challenged painting. "The reproduced work of art is to an ever-increasing extent the reproduction of a work of art designed for reproducibility," Benjamin writes. Another cultural crisis is looming as we realize that "new" or popular styles will be increasingly optimized for their algorithmic reproducibility (in other words, designed to spread meme-like over digital platforms) instead of their originality.

Want another Picasso, Gucci, Gehry, Glossier, Beyoncé? Just push the button. It'll be close enough. There's already an Instagram influencer with over 700,000 followers, Miquela, who appears to be a 19-year-old model dressing up in clothes from Chanel, Proenza Schouler, and Supreme. Her vibe is Kylie Jenner, with her malevolent-cherub face and embrace of streetwear. Except Miquela is actually a virtual character her designers rendered by computer, as if produced by a Kardashian-fed AI. Unlike Jenner, Miquela is a style that can be reproduced cheaply and infinitely.

Every platform, canvassed by an algorithm that prioritizes some content over other content based on predicted engagement, develops a Generic Style that is optimized for the platform's specific structure. This Generic Style evolves over time based on updates in the platform and in the incentives of the algorithm for users.

When we encounter the Generic Style in the world, we feel a shiver of fear: We have entered the realm of the not-quite-human, the not-quite-genuine. Did we make an independent decision or do the machines know us better than we know ourselves? (This anxiety might just be an iteration of the debate between free will and fate.)

What do we do, then, about this shift from human to digital taste? It's possible to consciously resist the algorithm, like someone might buck the current fashion trend — wearing bell-bottoms and tie-dye, say, instead of trim, blank basics. I might only read books I stumble across in used bookstores, only watch TV shows on local channels, only buy vinyl, only write letters, forsake social media for print newspapers, wear only found vintage. (Etsy is already algorithmic, with its own faux-folksy Generic Style.) I could abstain from algorithmic culture like the Luddites who resisted the automation of textile factories in the 19th century by destroying machines. It would be so organic. Cool! Obscure! Authentic!

But as soon as something Cool, Obscure, and Authentic gets put back on the internet, it is factored into the equation, maybe it goes viral, and soon enough it's as omnipresent. In this way, algorithmic culture is not encouraging of diversity or the coexistence of multiple valid viewpoints and identities. If a stylistic quirk is effective, it is integrated into the Generic Style as quickly as possible; if it is ineffective, it is choked of public exposure. So you'd also have to keep your discoveries analog. Put an air gap between your brain and the internet.

I grew up in the early 2000s during the beginning of the social internet, when there were no smart feeds or adaptive algorithms to sort content. The primary ways I discovered new things were through forums, where members suggested which shoes to buy or bands to listen to, and through digital piracy, which gave me a relatively unfiltered list of possible cultural artifacts to consume on Kazaa or BitTorrent, which did not come with "You May Also Like This" recommendations. (I did not live in a city and the local comprehensive bookstore was a Borders 45 minutes away.) These services were the digital equivalent of used vinyl shops: You take what you find, either you like it or not, and then you try again, constantly refining an image of what you want and (thus) who you are. Since those were formative teenage years, I derived a good part of my identity as a cultural consumer from DIY piracy. Still, the results were neither exceptional nor original. I downloaded a lot of Dave Matthews Band concert bootlegs and sought out American Apparel in the mall after seeing it online. But at least these things felt like mine? Or at least the assemblage aggregated into something I might have called personal taste.

Now YouTube tells me which videos to watch, Netflix serves me TV shows, Amazon suggests clothes to wear, and Spotify delivers music to listen to. If content doesn't exist to match my desires, the companies work to cultivate it. The problem is that I don't identify as much with these choices as what I once pirated, discovered, or dug up. When I look at my Spotify Discover playlists, I wonder how many other people got the exact same lists or which artists paid for their placement. I feel nostalgic for the days of undifferentiated .rar files loading slowly in green progress bars. There was friction. It all meant something.

To be fair, this content consumption was also extremely unethical. And it's not like I don't like Netflix shows or Spotify playlists. Like cigarettes or McDonald's, they were designed for me to like them, so of course I like them. It's just that I don't always like that I like them.

Yet there are an increasing number of legal alternatives to these mainstream platforms. We're seeing a profusion of smaller platforms with different brand images, the equivalent of a Reformation instead of a J.Crew or Glossier instead of Clinique. If Gap is a mainstream platform for fashion basics, then Everlane, with its transparent manufacturing and minimalist branding, and now Scott Sternberg's Entireworld, which purports to offer a utopian clothing system, are its more niche, though no less generic, hipster equivalents.

FilmStruck, for example, streams "critically acclaimed classic movies, hard-to-find gems, and cult favorites" like those in the Criterion Collection, while MUBI selects "cult, classic, independent and award-winning films from around the world." The full-bleed, black-and-white stills on their websites differentiate them as far hipper than Netflix or cable — you might feel safer about identifying your taste with them ("I don't watch TV; I only watch FilmStruck," a platform hipster says). Instead of Spotify, there's The Overflow, with vetted Christian worship music, or Primephonic, with high-definition classical recordings. Quincy Jones launched the "Netflix of jazz."

Digital platforms exist for non-digital products, too. The start-up Feather will rent you a "hip bedroom" bundle of faux-mid century side tables and bed frame for $109 a month in a kind of minimally stylish pre-packaged taste kit, a thinly reproduced aesthetic lacking any aura. Similarly, fashion companies like Gustin and Taylor Stitch crowdfund their new products, counting pre-orders before manufacturing anything. These are different from traditional brands in that they are driven from the bottom-up by the actions of users rather than the diktats of auteur creative directors. And, like the drop-shipped generic watches, they are extremely boring, releasing wave after wave of artisanal fabrics turned into rustic, vaguely outdoorsy gear.

What these businesses suggest is that you can have the benefits of a digital platform and an algorithmic feed while still feeling self-satisfied, pretentious, and exclusive in the knowledge that your content has been carefully curated by humans. Or, you could hire a tastemaker of your own. As The Verge reported, a musician named Deb Oh freelances as a Spotify curator through her service Debop, making custom playlists for $125. She culls from the "the symphony of algorithms," as she beautifully puts it, and comes back with something more manageable, more human.

Oh's services present original curation as a luxury good. It costs money to step off the consumption rails so conveniently laid out for us by tech companies and their advertisers. In the future, taste will be built on allegiances to platforms as much as individual creators or brands. Are you more of an Amazon, Apple, WeWork, Airbnb, or Facebook person? Unless you go off-platform, there are no other choices. Not just for your technology, but for your culture: fashion, furniture, music, art, film, media.
Platformization is something the fashion industry is already familiar with, of course: Each major brand is its own platform, expanding in a profusion of seasonal lines and accessories meant to cater to your every need within a single taste-system. LOT2046 is a smaller, independent algorithmic platform for fashion that I subscribed to last year and I haven't looked back. Its thesis is simple: Your clothing desires can be reduced to a series of signifiers that the service automates and adapts to you. Shipments of all-black clothing and accessories arrive every month; the only customizations are a few stylistic choices — short socks or long, crew-neck or V-

neck — and that the items come with your name emblazoned on them, like a black duffle bag I recently received that says KYLE CHAYKA in raised black thread. LOT is pro-algorithm. “Any technology should know what you need and want more than you know,” its founder Vadik Marmeladov, a Russian designer who prefers to stay behind the scenes, told me. “Platforms will be telling you what you want before you want it.” He feels that machines should not just suggest things, but make decisions for us, from planning a weekend trip to a morning coffee order. In other words, they should supplant our taste entirely.

Surrendering to LOT is a kind of freedom to stop thinking about fashion, freeing the mind for loftier things — like contemplating mortality, Marmeladov suggests. Its promise is that by drastically narrowing the variables, perhaps an algorithm can actually help you achieve individuality, not just through clothing but induced existentialism. I don’t wear LOT’s clothes all the time, but I find its ethos seeping into how I think about my consumption in the algorithm age more generally. If our decisions about what we consume don’t seem to communicate much about ourselves anymore, why not just choose to not make them?

The promise of algorithms is that they will show you yourself, refining an image of your tastes that should be identical to what you would have chosen on your own. The current reality is that these feeds silo you in homogenizing platforms, calculating the best-fitting average identity. That these average identities come in increasingly minute shades does not mean that they are unique.

A better mode of resistance might be to use the algorithms’ homogenizing averageness against them, adapting their data for productive disruption. We can take advantage of the clash between multiple algorithmic ideals, or between an algorithm’s vision of the world and reality, creating a glitch-based aesthetic. What would be error could be art.

As culture has changed to accommodate every other technological innovation, so our ideas about algorithms will change. “Eventually we may opt to shift our definition of art in order to make accommodation for the creativity of artificial intelligence,” says Marian Mazzone, an art history professor at the College of Charleston who worked on a project in which AI created original styles of painting (they mostly look like mash-ups of Impressionism, Fauvism, and Cubism).

Oscar Sharp is the director of Sunspring, a short sci-fi film with a script generated by a machine-learning algorithm trained on episodes of The X-Files, Star Trek, and Futurama. The result is something spiky, mostly non-narrative — it doesn’t make much sense, but it is compelling and unique. The film doesn’t try to fool the viewer into thinking it’s 100 percent human-made. Rather, the actors strain to adapt to the aesthetics of the machine and discover something new in the process.

“It’s like you’re working on a big TV show with a very powerful showrunner who has written the episode, and the showrunner got drunk last night, passed out, and you couldn’t not make the episode,” Sharp says. “You have to do everything within your power to make the episode as it was written.” The challenge was generative:

“Augmented creativity is much more interesting than a replacement of creativity,” he says.

The automated clothing service Stitch Fix, kind of a preppy version of LOT, uses algorithmic help to optimize their new original designs to increase sales and address gaps in the market, what they call “Hybrid Design”: customers like ruffles and plaid, so why not plaid ruffles? But we could instead go in the opposite direction, making clothes no one wants — yet. Algo-clash clothing would be more like the artist Philip David Stearns’s glitch textiles, unique fabrics generated from software gone intentionally awry, the discordant pattern of pixels made into a Baroque style. Fashion is always one step ahead, though. The triple-waistband jeans recently released by ASOS already look like a glitched algorithm designed them.

It is not just that artists can collaborate with algorithms; there is always a person at the end of the machine — like the man behind the curtain in Oz — regulating what it does. The majority of these are currently Silicon Valley engineers. And we human consumers are still on the other side of the algorithm, with our freedom to decide what we consume or to opt out. Our decisions shape what is popular in the present as well as what is preserved into the future. “Let’s not forget the audience has a major role to play in determining what will matter and what will not, what is liked and what is not,” Mazzone says. In the long term, this is slightly comforting.

I leave the cyclopic Amazon Echo Look on a shelf in my living room, where it glares at me every time I walk past, not stopping for it to evaluate my outfit. It yearns to assign inexplicable percentages, and yet I am more comfortable judging for myself. It takes fine pictures, but like a mirror, it mostly shows me what I already know. And the device is trying to match me to some universalized average, not my individual style, whatever that might be. It doesn’t know me at all — it can’t tell what kind of clothes I’m comfortable in nor how the clothes I wear will function as symbols outside, in the place I live, in the contexts of class or gender. All-black doesn’t play the same in Kansas City as it does in New York, after all. This is the kind of social, aesthetic intelligence, the sense of taste, that our algorithms are missing, for now at least.

Amazon says the Look is for achieving your best style, but its ulterior motives aren’t hard to spot. When I asked the machine about my plaid shirt, an ad popped up on the app’s feed showing me a few other, similarly colored plaid shirts — none particularly stylish or different enough from the one I own, bereft of brand name — that I could buy on Amazon. In fact, Amazon is already using the data it collects to manufacture its own clothing lines, and the results are about what you’d expect from a robot: wan imitations of whatever is currently popular, from the “globally inspired” Ella Moon to the cool-French-girl knockoff Paris Sunday. Training on millions of users’ worth of data and images from the Look showing what we actually wear could make the in-house brands slightly less uncanny. Then again, imagine a potential leak, not of credit card data but an extensive cache of your outfits.

It's up to us whether or not we care about the shades of distinction between human and machine choice, or indeed if we care about fashion at all. Maybe taste is the last thing separating us from the Singularity; maybe it's the first thing we should get rid of. "I don't think the consumer cares, as long as it works," one Stitch Fix executive said of its algorithmically designed clothes.

But if we do want to avoid displacing or reassigning our desires and creativity to machines, we can decide to become a little more analog. I imagine a future in which our clothes, music, film, art, books come with stickers like organic farmstand produce: Algorithm Free.

Adapted from Vox

JavaScript is for Girls

Decades ago, men kicked women out of the programming profession just as it was taking off. Now that women are fighting their way back in, men are finding new ways to protect their status.

Technology has a gender problem, as everyone knows. The underrepresentation of women in technical fields has spawned legions of TED talks, South by Southwest panels, and women-friendly coding boot camps. I've participated in some of these get-women-to-code workshops myself, and I sometimes encourage my students to get involved. Recently, though, I've noticed something strange: the women who are so assiduously learning to code seem to be devaluing certain tech roles simply by occupying them.

It's not always obvious to outsiders, but the term "technology sector" is a catch-all for a large array of distinct jobs. Of course there are PR, HR, and management roles. But even if we confine ourselves to web development, technical people often distinguish among "front-end," "back-end," and "full-stack" development. The partition between the two "ends" is the web itself. There are people who design and implement what you see in your web browser, there are people who do the programming that works behind the scenes, and there are people who do it all. In practice, the distinction is murky: some developers refer to everything user-facing as the front-end, including databases and applications, and some developers use front-end to mean only what the user sees. But while the line shifts depending on who you're talking to, most developers acknowledge its existence.

I spoke to a number of developers who confirmed something I'd sensed: for some time, the technology industry has enforced a distinct hierarchy between front-end and back-end development. Front-end dev work isn't real engineering, the story goes. Real programmers work on the back-end, with "serious" programming languages. Women are often typecast as front-end developers, specializing in the somehow more feminine work of design, user experience, and front-end coding.

Are women really more likely to be front-end developers? Numbers are hard to pin down. Most studies consider the tech sector as a single entity, with software engineers lumped together with HR professionals. A survey showed that front-end jobs—"Designer," "Quality Assurance," and "Front-End Web Developer"—were indeed the top three titles held by women in the tech industry, although that survey itself has some problems.

We need better numbers, as feminist developers have been saying for years, but it also doesn't seem like a huge stretch to take developers at their word when they say that front-end development is understood to occupy the girlier end of the tech spectrum.Front-end developers, importantly, make about $30,000 less than people in back-end jobs like "DevOps" engineers, who work on operations and infrastructure, according to the salary aggregation site Glassdoor.

The distinction between back and front wasn't always so rigid. "In the earliest days, maybe for the first ten years of the web, every developer had to be full-stack," says Coraline Ada Ehmke, a Chicago-based developer who has worked on various parts of the technology stack since 1993. "There wasn't specialization."

Over time, web work professionalized. By the late 2000s, Ehmke says, the profession began to stratify, with developers who had computer science degrees (usually men) occupying the back-end roles, and self-taught coders and designers slotting into the front.

For many people who are teaching themselves to code, front-end work is the lowest-hanging fruit. You can "view source" on almost any web page to see how it's made, and any number of novices have taught themselves web-styling basics by customizing WordPress themes. If you're curious, motivated, and have access to a computer, you can, eventually, get the hang of building and styling a web page.
Which is not to say it's easy, particularly at the professional level. A front-end developer has to hold thousands of page elements in her mind at once. Styles overwrite each other constantly, and what works on one page may be disastrous on another page connected to the same stylesheet. Front-end development is taxing, complex work, and increasingly it involves full-fledged scripting languages like JavaScript and PHP.

"Serious" developers often avoid acknowledging this by attributing front-end expertise not to mastery but to "alchemy," "wizardry," or "magic." Its adepts don't succeed through technical skill so much as a kind of web whispering: feeling, rather than thinking, their way through a tangle of competing styles.
"There's this perception of it being sort of a messy problem that you have to wrangle with systems and processes rather than using your math-y logic," says Emily Nakashima, a full-stack developer based in San Francisco. That's not true, of course; nothing on a computer is any more or less logical than anything else. But perhaps it's easier to cast women in a front-end role if you imbue it with some of the same qualities you impute to women.

The gendered attributes switch as you travel to the back of the stack. At the far end, developers (more often "engineers") are imagined to be relentlessly logical, asocial sci-fi enthusiasts; bearded geniuses in the Woz tradition. Occupations like devops and network administration are "tied to this old-school idea of your crusty neckbeard dude, sitting in his basement, who hasn't showered in a week," says Jillian

Foley, a former full-stack developer who's now earning her doctorate in history. "Which is totally unfair! But that's where my brain goes."

The brilliant but unkempt genius is a familiar figure in the history of computing—familiar, but not immutable. Computing was originally the province of women, a factinnumerable articles and books have pointed out but which still seems to surprise everyone every time it's "revealed." The bearded savant of computer science lore was the result of the field's professionalization and increasing prestige, according to the computing historian Nathan Ensmenger.

"If you're worried about your professional status, one way to police gender boundaries is through educational credentials," says Ensmenger. "The other way, though, is genius. And that's something I think nerd culture does really well. It's a way of defining your value and uniqueness in a field in which the relationship between credentials and ability is kind of fuzzy." And "genius," of course, is a strongly male-gendered attribute—justlook at teaching evaluations.

When programming professionalized, women got pushed out. Marie Hicks, a computing historian who's looked closely at this phenomenon, explains that as programming came to be viewed as more important to national and corporate welfare, hiring managers began associating it with a specific set of skills. In the British case, Hicks's specialty, a good programmer was supposed to be the ultimate systems-thinker, able to see and synthesize the big picture. In the United States, as Ensmenger and others have documented, the best programmers were purportedly introverted chess nerds, obsessed with details, logic, and order. (There's very little evidence that these characteristics actually make a good programmer.)

The traits of a "good programmer" differed by country, but they were universally male-gendered, enforced by hiring managers and other programmers who sought to replicate their own characteristics—not consciously, for the most part, but simply because the jobs were important. Hiring managers wanted to bet on qualities everyone agreed were indicators of success. "The people with more prestige in a culture are favored for all sorts of things, including jobs," says Hicks. "If you have a job that you want to fill, you want to get the best worker for it. So in more prestigious fields, employers are looking for those employees that they think are the best bet. This tends to attract men who are white or upper-class into these more desirable jobs."

People often think that as a profession matures it gets more complex, and thus edges women out because it demands higher-level skills. But "historically, there's very little to bear that out," says Hicks, who has uncovered multiple incidents of women programmers training, and then being replaced by, their male counterparts.
The case of the female front-end developer is flipped in the other direction—it's a feminizing subfield, rather than a masculinizing one. But it's governed by many of the same market forces that edged women out of programming in the first place: prestige accrues to labor scarcity, and masculinity accrues to prestige. Front-end jobs are easier for women to obtain, and feminized jobs are less prestigious. In turn, the labor market generates its own circular logic: women are front-end developers

because they're well-disposed to this kind of labor, and we know this because women are front-end developers.

No one says any of this explicitly, of course, which is why the problem of women in technology is thornier than shoehorning women onto all-male panels. The developers I spoke to told me about much more subtle, very likely unconscious incidents of being steered toward one specialization or another. Two different women told me about accomplished female acquaintances being encouraged to take quality assurance jobs, currently one of the least prestigious tech gigs. Ehmke told me about a friend who applied for a back-end developer position. Over the course of the interview, the job somehow morphed into a full-stack job—for which Ehmke's friend was ultimately rejected, because she didn't have the requisite front-end skills.

And everyone can rattle off a list of traits that supposedly makes women better front-end coders: they're better at working with people, they're more aesthetically inclined, they care about looks, they're good at multitasking. None of these attributes, of course, biologically inhere to women, but it's hard to dispute this logic when it's reinforced throughout the workplace.

Once you're cast as a front-end developer, it can be challenging to move to different parts of the stack, thus limiting the languages and development practices you're exposed to. "Particularly in Silicon Valley, there's a culture of saying developers should always be learning new things," says Nakashima, the San Francisco-based full-stack developer. Front-end specialization "can be a place that people go to and don't come back from. They're working on these creative projects that are in some ways very interesting, but don't allow them to move to an area of the stack that's becoming more popular."

Viewed from one angle, the rise of get-girls-to-code initiatives is progressive and feminist. Many people involved in the movement are certainly progressive feminists themselves, and many women have benefited from these initiatives. But there are other ways to look at it too. Women are generally cheaper, to other workers' dismay. "Introducing women into a discipline can be seen as empowerment for women," says Ensmenger. "But it is often seen by men as a reduction of their status. Because, historically speaking, the more women in a profession, the lower-paid it is." An influx (modest though it is) of women into the computing profession might be helping to push developers to make distinctions where they didn't exist before. "As professions are under threat, stratification is very often the result," says Ensmenger. "So you take those elements that are most ambiguous and you push those, in a sense, down and out. And down and out means they become more accessible to other groups, like women." But these roles are also markedly distinct from the main work of software engineering—which is safely insulated from the devaluing effect of feminization, at least for the time being.

Hicks, the computing historian, can't stand it when people tout coding camps as a solution to technology's gender problem. "I think these initiatives are well-meaning, but they totally misunderstand the problem. The pipeline is not the problem; the meritocracy is the problem. The idea that we'll just stuff people into the pipeline assumes a meritocracy that does not exist."

Ironically, says Hicks, these coding initiatives are, consciously or not, betting on their graduates' failure. If boot camp graduates succeed, they'll flood the market, devaluing the entire profession. "If you can be the exception who becomes successful, then you can take advantage of all the gatekeeping mechanisms," says Hicks. "But if you aren't the exception, and the gatekeeping starts to fall away, then the profession becomes less prestigious."

My students are always so excited that they're "learning to code" when I teach them HTML and CSS, the basic building blocks of web pages. And I'm happy for them; it's exhilarating to see, for the first time, how the web is built. Increasingly, though, I feel the need to warn them: the technology sector, like any other labor market, is a ruthless stratifier. And learning to code, no matter how good they get at it, won't gain them entrance to a club run by people who don't look like them.

Adapted from Logic magazine

McDonald's CEO Wants Big Macs to Keep Up With Big Tech

Steve Easterbrook is giving the Golden Arches a data makeover, but franchisees are balking at the cost.

Three years ago, Steve Easterbrook ran out of patience. Before flying home to Chicago for the Christmas holidays, he stopped in Madrid to meet with Spanish executives from McDonald's. In a conference room at the company's local office off the A6 highway, the mood soured as managers lamented heavy losses on the evenings when FC Barcelona and Real Madrid C.F. competed. Diners were staying home and ordering from archrival Burger King for delivery—a service McDonald's didn't offer.

Conceding to Burger King in any circumstance is an indignity, but losing hundreds of thousands of customers to the enemy's modernized tactics during one of Spain's most important weekly fixtures was the final straw. It represented everything that was defective at the business Easterbrook had been running for 22 months—McDonald's Corp. was just too analog. A week before he was named chief executive officer, the company announced it had suffered one of its worst years in decades as dejected U.S. customers abandoned the brand for Chipotle burritos and Chick-fil-A sandwiches. In the U.K. hundreds of artisanal burger competitors had appeared seemingly overnight on the food-delivery mobile app Deliveroo, which indulged the couch potato demographic with an unprecedented ease of access that felled the appeal of McDonald's drive-thrus. The time had come to address a weakness that stretched far beyond the company's Iberian territories.

"He looked at me and said, 'We're not going to go through the traditional market pilot and study delivery for six months. We're just going to do it,' " says Lucy Brady, who oversees McDonald's global strategy and business development teams. He instructed her to get every country manager on a conference call on Monday morning.

Brady cautioned him that it might be difficult to reach some managers who'd already left for the holiday; Easterbrook said everyone could spare a half-hour. He would command each manager to nominate their best executive to the task of building an online delivery business that would aim to be fully operational by the beginning of January—in two weeks' time. When Brady suggested they target delivery from 3,000 restaurants by July 1, he told her he would be disappointed if they didn't get to 18,000—about half of McDonald's locations around the globe.

Management's compensation would be tied to the speed and breadth of the rollout, and the only limiting factor Easterbrook would accept would be the number of couriers in cars, on bikes, and on foot that their delivery partners could supply. For the widest possible deployment, McDonald's teamed with Uber Eats. The partnership was so significant that Uber Technologies Inc. devoted two full pages to its then-exclusive delivery agreement with McDonald's in a roadshow prospectus ahead of Uber's initial public offering in May. Easterbrook now regularly uses the service while traveling on business to gauge its quality.

"I'm a Quarter Pounder guy," he says with a calculated slowness not unlike Daniel Day-Lewis's "I'm an oil man" in There Will Be Blood. The 52-year-old British CEO has the tall, broad frame of a rugby player, with thick waves of black hair and piercing blue eyes. He's described as an inscrutable blend of mild manners and obsessive competition by members of his fresh-faced leadership team. (Upon taking the top job in 2015, Easterbrook fired or let go 11 of the 14 most senior executives he inherited.) He expects the delivery business to account for about $4 billion in sales by the end of this year. Catching up to Burger King on delivery would be the first item on a long list of improvements Easterbrook already had in mind for McDonald's. Broadly, he wants to reconfigure his restaurants into enormous data processors, complete with machine learning and mobile technology, essentially building the Amazon of excess sodium. Franchisees have balked at the costs of implementing his vision, which includes drive-thrus equipped with license-plate scanners (the better to recall one's previous purchases) and touchscreen kiosks that could ultimately suggest menu items based on the weather.

So far, the strategy has proved compelling: Only a handful of other companies in the S&P 500, almost all of them California technology suppliers such as semiconductor giant Advanced Micro Devices and chipmaker Nvidia, have outperformed McDonald's returns since 2015. The gains have generously rewarded institutional investors like BlackRock Inc. and Vanguard Group Inc., who've long been among the chain's largest backers. Easterbrook wants to reclaim the company's image as a beacon of innovation, a designation McDonald's hasn't enjoyed since roughly the Truman administration.

In 1940 brothers Dick and Mac McDonald redesigned and rebuilt their modest hot dog drive-in, in the shadows of California's San Bernardino mountain range, into McDonald's Bar-B-Q, which sold 25 items. By 1948 they dropped "Bar-B-Q" from the name and streamlined the menu to offer only the most profitable foods: hamburgers, cheeseburgers, potato chips, coffee, soft drinks, and apple pie. The restaurant seated about a dozen customers on outdoor stools and sold 15¢ hamburgers, which were bagged within 30 seconds of being ordered thanks to the pioneering Speedee service system.

The system had begun with the brothers sketching out life-size kitchen blueprints on a tennis court with chalk and having employees act out cooking and serving tasks. After settling on the fastest method, they contracted kitchen equipment companies to build machinery to support the choreography. Breakthroughs included custom-made saucing guns for the buns and curved steel ramps on which burgers would slide down into cashiers' hands to pass on to diners. At the time, only a handful of burger chains were using similarly bespoke hardware.

No one was as passionate about McDonald's potential for expansion as Ray Kroc, a struggling milkshake machine salesman from Illinois who met Dick and Mac at their restaurant in 1954 on a business trip. McDonald's had by far the most efficient kitchen he'd ever seen, and he immediately lobbied the brothers to let him franchise the business. In 1961 he bought out the co-founders for $2.7 million, and in 1965 he took the company public. Today, McDonald's is the world's most recognizable restaurant empire and a formidable real estate venture—its franchising model has earned the company a fortune by acquiring and subsequently leasing the land beneath stores to their operators.

For a half-century, McDonald's greased its way onto every continent except Antarctica. It stayed ahead of scores of copycats, but the baby boomer loyalty that propped it up has steadily waned. It's also become something of a cultural laggard. The suitability of McDonald's in a looming Age of Kale was aggressively pondered in Super Size Me, the 2004 documentary film in which director Morgan Spurlock attempts to subsist on the restaurant chain's food for a month. He cast the company as an abhorrent peddler of heartburn and substandard bowel movements. There's also the inevitable discomfort of being one of the world's largest purchasers of beef and poultry. Younger generations concerned about the environmental cost of industrialized meat are opting for plant-based alternatives such as Beyond Meat and the Impossible Burger, which is now available at Burger King. Animal-rights activists regularly erect giant inflatable chickens with bereaved expressions on the sidewalk outside McDonald's new head office in downtown Chicago.

The company boasts a market valuation of $159 billion and an immense global reach, feeding about 1% of the human population daily. But even in the fast-food realm it dominates, its share of the U.S. market has shrunk to 13.7% from 15.6% in 2013, according to data from Euromonitor International, ceding ground to Pret a Manger and Panera Bread Co. In the burger wars, it's been besieged by cooler competitors with cult followings, including Shake Shack, Five Guys, and In-N-Out. Earnings began to stagnate at McDonald's in 2013 and crashed by almost a fifth, to $4.7 billion, the following year as diners deserted. Four months before stepping down in March 2015, Don Thompson, Easterbrook's predecessor, lamented that the company had failed to evolve "at the same rate as our customers' eating-out expectations." As insurgents claimed an ever-growing share of the market

McDonald's had created, the morale at the old headquarters in Oak Brook—a tranquil if uninspiring 1970s amalgam of gray cubicles set in a parkland in Illinois—began to sap. The company's strategic quagmire took on a superstitious quality when the estate itself became a hive of bad omens, with parts of the office complex flooding on an annual basis.

Easterbrook became global chief brand officer in 2013. The following year, he traveled to Cupertino, Calif., to sit down with Tim Cook, Apple Inc.'s CEO, to discuss being a launch partner for the Apple Pay mobile payment system. The card readers McDonald's used lacked the necessary technology, so Easterbrook had a digital add-on installed on every machine at its 14,000 locations in the U.S.

Easterbrook first joined McDonald's in the finance department in London in 1993, and spent the majority of his career there. After graduating with a natural sciences degree from Durham University, where he played competitive cricket alongside the future England captain, he worked as an accountant for the partnership that would become PwC. He later worked as a restaurant manager for McDonald's before being named to head its U.K. division, which he turned around in the 2000s after years of waning sales. In that role, he mounted a defense against fast-food critics by debating them on live television. He revitalized the company's image as a family-friendly outlet by introducing organic milk, cutting the fries' salt content, and offering free Wi-Fi. He also tried unsuccessfully to get the Oxford English Dictionary to amend its definition of "McJob," a slang term used since at least 1986 that denoted "an unstimulating, low-paid job."

In the fall of 2014, McDonald's went public with "Experience of the Future," an initiative Easterbrook had been shepherding. It reimagined the store entirely, from how orders were placed to what services were offered. In the upgraded restaurants, diners can use touchscreen kiosks to customize their burgers into millions of permutations, such as adding extra sauce and bacon to a Big Mac. The thinking was that giving customers more say over their orders would result in them paying more for tailored items. Some franchisees have benefited so much that their restaurants' sales are now growing at a double-digit rate. But others have banded together in open rebellion and forced the company to slow the program's full rollout two years past its original target. They object to the enormous costs of the project, which, for owners of several locations, can run into tens of millions of dollars, even with McDonald's offering to subsidize 55% of the capital for the remodels.

From a business perspective, the enhancements are achieving what they set out to do—annual profits have inched higher since Easterbrook's appointment, and McDonald's posted its fastest global sales gain in seven years last quarter. Initiatives such as all-day breakfast, which includes the staple McMuffin, and new products like doughnut sticks are also credited with bringing customers back even as the expanded menu hampers the classic McDonald brothers' efficiency.

The company has also introduced a curbside pickup system. An order placed through the McDonald's app automatically appears on the store's order list when the diner's phone is within 300 feet of the property. The food is prepared and delivered to the curb by floor employees. The workers and franchisees who've long complained about low hourly wages and poor working conditions in campaigns such as Fight for $15 have generally taken a dim view of Easterbrook's overhaul. Westley Williams, a Floridian in his early 40s, says the initiatives and the chaos caused by mobile app orders, new items, and self-order kiosks riddled him with so much anxiety that he defected to nearby burger chain Checkers. "It's more stressful now," said Williams, who added that he didn't get a raise for doing more work. "When we mess up a little bit because we're getting used to something new, we get yelled at."

Concerns about staff welfare have become a major issue for McDonald's in the U.S., where the median pay for food and beverage service workers is $10.45 an hour. Accusations of coercion soared this year after workers filed a total of 25 claims and lawsuits alleging endemic sexual harassment. The complaints have since become a national conversation and part of the political fabric: In June a group of eight senators led by Democrat Tammy Duckworth of Illinois and including 2020 Democratic presidential candidates Bernie Sanders of Vermont, Elizabeth Warren of Massachusetts, Kamala Harris of California, and Amy Klobuchar of Minnesota sent a letter to Easterbrook decrying "unsafe and intolerable" conditions and "unacceptable" behavior in the chain's restaurants.

Carlos Mateos Jr., whose family owns 21 stores near Washington, D.C., says Easterbrook's modernization has succeeded in attracting new customers to his restaurants, but revamping everything simultaneously was a burden. About a quarter of his franchises still need to be remodeled. "There's training that's involved. We have to get the employees ready for it—mobile order and pay and Uber Eats and kiosks. All these different things are happening at the same time, and it really took a toll on us."

Adding an Uber Eats counter for delivery, touchscreen kiosks, modern furniture, and power outlets to charge mobile phones means franchisees incur additional costs from $160,000 to $750,000 per restaurant, McDonald's has said. Blake Casper, a Tampa-based franchisee who operates more than 60 McDonald's and founded the National Owners Association last fall to resist Easterbrook's amelioration plan, would theoretically have to fork over at least $5 million to make the CEO's dream a reality.

"I would like to make the kitchen as stress-free as it possibly can be," says Eli Asfaw, who operates seven franchises in the Denver area. For a start, scaling back rollouts mandated by the company, such as all-day breakfast, would "make it easier for us to keep people and make our people happy." Asfaw also says the remodeling plan has heaped pressure on owners, from financial headwinds to the tight window in which the company wants the upgrades to be completed.

The resistance from a faction of franchisees to Easterbrook's mandated remodels—in some cases drastic enough to require a restaurant to be razed and rebuilt—reached a breaking point in January. The National Owners Association wrote in a letter to its 400 members then (it now counts more than 1,200) that the changes should be halted amid concerns about eroding profits and the costs of implementing Experience of the Future. “To put it bluntly,” the letter read, “stop everything that is not currently in the works.”

Easterbrook concedes his rollout hasn’t been perfect. “We were just going so hard at it, it proved to be a bit of a handful,” he says of introducing the features in the U.S., many of which had already been phased in years before in France and Australia. While franchisees were right to put off remodeling to ensure they weren’t distracted from efficiently running their restaurants, the domestic business was in dire need of a significant revamp, he says. The number of customers visiting U.S. stores had been declining in the last half of his predecessor’s tenure. “It was pretty obvious we were operating and moving slower than the outside world, and customers were voting with their feet.”

In November, McDonald’s said it was slowing the pace of remodels in the U.S. The conversations are often fraught. When Easterbrook invited eight franchisees to break bread with a group of McDonald’s executives at a steakhouse in Washington, D.C., in April, one operator accused him of saddling stores with impossible demands. For longtime managers who back Easterbrook’s goal and enjoy the internal energy it’s created, the prospect that his plans could fall through is unthinkable.

“When they ask a question that’s a bit of an attack, I sit there and get a little pissed, because I’m ready to lean in,” says Charlie Strong, a 66-year-old McDonald’s executive who oversees more than 5,700 restaurants across the western U.S. He affixes a lapel pin of the letter “M” in the style of the golden arches logo to a navy Brooks Brothers blazer, and his right pinkie is weighed down by a 14-karat yellow gold ring inset with five diamonds, onyx, and the golden arches. The company gave it to him to celebrate his 25th anniversary with McDonald’s. He expects to receive another for his 50th in two years.

Strong says one of Easterbrook’s key qualities is that he doesn’t take any criticism of his strategy personally. “He just rolls with it and swings it back to what’s important about the business, what’s important about the vision, and to not get bogged down with these little things along the way.”

Easterbrook’s strategy so far has been vindicated by the numbers. That tailwind is breathing new life into the business. Strong drives 40 miles from his home in Aurora, Ill., every morning to be at his desk by 6 a.m., where he and a handful of other masochistic early risers blast rousing tunes by Journey or Adele on a Bose sound system to get the day going. It’s a routine they began after moving into the new head office, a $250 million building replete with sofa pods in the red and yellow McDonald’s color scheme, an amphitheater, rooftop terraces, and thousands of antique and modern Happy Meal toys locked inside cased glass like priceless museum specimens. Easterbrook opened the office in June of last year in a bid to attract young, tech-forward talent.

In March, McDonald’s acquired artificial intelligence startup Dynamic Yield, headquartered in New York and Tel Aviv, for $300 million—the company’s largest acquisition in 20 years. The burger chain had been testing the machine learning software on drive-thrus at four restaurants in Florida, where screens automatically updated with different items based on the time of day, restaurant traffic, weather, and trending purchases at comparable locations. That technology has been deployed at 8,000 McDonald's and counting, with plans to be in almost all drive-thrus in the U.S. and Australia by the end of the year, Easterbrook says. The deal signaled an ambition to align the chain with the same predictive algorithms that power impulsive purchasing on Amazon.com or streaming preferences on Netflix. In April, McDonald's acquired a minority stake in New Zealand-based mobile app vendor Plexure Group Ltd., which helps restaurants engage with diners on their phone with tailored offerings and loyalty programs. The effort falls into the consumer-goods industry's wider trend toward micromarketing, which has proved effective in driving sales.

In early September, McDonald's said it was buying Silicon Valley startup Apprente Inc., a developer of voice-recognition technology. The idea is to help speed up lines by eventually having a machine, instead of a person, on the other side of the intercom to relay orders to kitchen staff. The deal for Apprente is McDonald's third such investment in a technology business in the past six months as the company shakes off a tamer takeover strategy that for decades had focused on buying and selling restaurants from or to operators. McDonald's is pursuing this new business model even as the latest burger trends steal the buzz from its offerings. Beyond fashionable vegan patties, a new and daunting foe is the fried chicken sandwich at Popeyes Louisiana Kitchen (a Miami-based chain owned by the same company that controls Burger King), which became a national obsession when it was introduced in the U.S. in August.

McDonald's has leased space in a discreet industrial complex more than an hour away from headquarters, where a gray building about the size of an aircraft hangar, with a single column painted yellow and dotted with sesame-seed stencils, has become a testing ground for putting Easterbrook's thoughts into practice. But for all the technological breakthroughs, the deals, and the jousting with franchisees, the company's guiding light has barely changed. Inside a room beyond a corridor stamped with the word "innovate" in block capital letters, the hum of computers and data processing towers is drowned out by a cacophony of test-kitchen staff running trials on secret processes that aim to shave seconds off a Big Mac's assembly, much like in the old days, when McDonald's first upended the food industry. "In old-school business logic, the big eats the small," Easterbrook says. "In the modern day, the fast eats the slow."

Adapted from Bloomberg

See No Evil

Software helps companies coordinate the supply chains that sustain global capitalism. How does the code work—and what does it conceal?

Trawling a hotel minibar one night while on a work trip to Amsterdam, I found a piece of chocolate with an unusual name: Tony's Chocolonely. I giggled at how apt the name was—who eats minibar chocolate unless they are, indeed, a little lonely?—and, on a whim, plugged it into Google.

The results were more sobering than I'd expected. The founder of Chocolonely, Teun (Tony) van de Keuken, founded the company with the goal of making the first (the "lonely only") chocolate bar produced without labor exploitation. According to the company, this goal actually landed them in legal trouble: Bellissimo, a Swiss chocolatier, sued Chocolonely in 2007, allegedly claiming that "slave-free chocolate is impossible to produce."

I had heard similar claims about other industries. There was the Fairphone, which aimed at its launch in 2013 to be the first ethically produced smartphone, but admitted that no one could guarantee a supply chain completely free from unfair labor practices. And of course one often hears about exploitative labor practices cropping up in the supply chains of companies like Apple and Samsung: companies that say they make every effort to monitor labor conditions in their factories.

Putting aside my cynicism for the moment, I wondered: What if we take these companies at their word? What if it is truly impossible to get a handle on the entirety of a supply chain?

The thing that still confused me is how reliable supply chains are, or seem to be. The world is unpredictable—you've got earthquakes, labor strikes, mudslides, every conceivable tragedy—and yet as a consumer I can pretty much count on getting what I want whenever I want it. How can it be possible to predict a package's arrival down to the hour, yet know almost nothing about the conditions of its manufacture?

In the past twenty years, popular and academic audiences have taken a growing interest in the physical infrastructure of global supply chains. The journalist Alexis Madrigal's Containers podcast took on the question of how goods travel so far, so quickly. The writer Rose George traveled the world on a container ship for her book Ninety Percent of Everything. And Marc Levinson's The Box startled Princeton University Press by becoming a national bestseller. Most recently, Deborah Cowen's The Deadly Life of Logistics offered a surprisingly engrossing history of that all-important industry.

These books help us visualize the physical infrastructure that makes global capitalism possible. But the data infrastructure has yet to be explored. How does information travel through the supply chain in such a peculiar way, so that I know to wait impatiently at my door at the exact moment my new iPhone will arrive—but no one really seems to know how it has gotten to me?

I set out to find the answer, and what I found surprised me. We consumers are not the only ones afflicted with this selective blindness. The corporations that make use of supply chains experience it too. And this partial sight, erected on a massive scale, is what makes global capitalism possible.

The industry of supply-chain management (or SCM, to its initiates) is both vast and secretive. It's one of the most rapidly growing corporate fields, and the subject of reams of books, journal articles, and blog posts. You can even get a degree in it.

But most companies are leery about revealing too much about their own logistics operations. It's not only because they are afraid of exposing what dark secrets might lurk there. It's also because a reliable, efficient supply chain can give a company an invaluable edge over its competitors.

Take Amazon: it's not so much a retailer as a supply chain incarnate. Its advantage lies in the high speed and the low price with which it can get a set of bath towels to your door. No wonder the retailer is famously tight-lipped about its supply-chain infrastructure. Few people outside of Amazon know much about the software that Amazon uses to manage its logistics operations.

In the supply-chain universe, there are large, tech-forward companies like Amazon and Apple, which write and maintain their own supply-chain software, and there's everyone else. And most everyone else uses SAP. SAP—the name stands for Systems, Applications, and Products—is a behemoth, less a single piece of software than a large, interlocking suite of applications, joined together through a shared database. Companies purchase SAP in "modules," and the supply-chain module interlocks with the rest of the suite. Among people who've used SAP, the reaction to hearing its name is often a pronounced sigh—like all large-scale enterprise software, SAP has a reputation for being frustrating.

Nevertheless, SAP is ubiquitous, with modules for finance, procurement, HR, and supply-chain management. "A very high percentage of companies run SAP for things like finance," says Ethan Jewett, an SAP consultant and software developer who helps companies implement SAP modules. "And so, if you're running it for one part of your business, you'll default to running it for another part of your business."

Leonardo Bonanni is the founder and CEO of a company called Sourcemap, which aims to help companies map their own supply chains. Bonanni suspects that companies' inability to visualize their own supply chain is partly a function of SAP's architecture itself. "It's funny, because the DNA of software really speaks through," said Bonanni. "If you look at SAP, the database is still actually written in German. The relations in it are all one-link. They never intended for supply chains to involve so many people, and to be interesting to so many parts of the company."

This software, however imperfect, is crucial because supply chains are phenomenally complex, even for low-tech goods. A company may have a handle on the factories that manufacture finished products, but what about their suppliers? What about the suppliers' suppliers? And what about the raw materials?

"It's a staggering kind of undertaking," said Bonnani. "If you're a small apparel company, then you still might have 50,000 suppliers in your supply chain. You'll have a personal relationship with about 200 to 500 agents or intermediaries. If you had to be in touch with everybody who made everything, you would either have a very small selection of products you could sell or an incredible margin that would give you the extra staff to do that."

We call them "supply chains," but that image is misleading. They really look more like a network of waterways, with thousands of tiny tributaries made up of sub-suppliers trickling into larger rivers of assembly, production, and distribution.

Bonanni explained that while workplace abuses get a lot of attention when they take place in the supply chains of large, prestigious companies like Apple and Samsung, working conditions are actually most opaque and labor abuse is most rampant in other industries, like apparel and agriculture. “Apparel, every quarter they have 100 percent turnover in the clothing that they make, so it’s a whole new supply chain every season. And with food, there’s millions of farmers involved. So in these places, where there’s way too many nodes for anyone to see without a computer, and where the chain changes by the time you’ve monitored it—those are the places where we see a lot of problems and instability.”

The picture that many of us have of supply chains involve state-of-the-art factories like those owned by Foxconn. In reality, the nodes of most modern supply chains look much less impressive: small, workshop-like outfits run out of garages and outbuildings. The proliferation and decentralization of these improvisational workshops help explain both why it’s hard for companies to understand their own supply chains, and why the supply chains themselves are so resilient. If a fire or a labor strike disables one node in a supply network, another outfit can just as easily slot in, without the company that commissioned the goods ever becoming aware of it. It’s not like there’s a control tower overseeing supply networks. Instead, each node has to talk only to its neighboring node, passing goods through a system that, considered in its entirety, is staggeringly complex. Supply chains are robust precisely because they’re decentralized and self-healing. In this way, these physical infrastructures distributed all over the world are very much like the invisible network that makes them possible: the internet.

By the time goods surface as commodities to be handed through the chain, purchasing at scale demands that information about their origin and manufacture be stripped away. Ethan Jewett explained the problem to me in terms of a theoretical purchase of gold: In some sense all gold is the same, so you just buy the cheapest gold you can get. But if you look at it in another way, it matters how it was mined and transported. And then all of the sudden, every piece of gold is a little bit different. And so it becomes very difficult to compare these things that, in terms of your actual manufacturing process, are almost exactly the same. To be traded as a commodity, in other words, gold must be gold.

As Jewett described this state of affairs, I felt a jolt of recognition. The system he was outlining was, in a word, modular: a method of partitioning information that’s familiar to every computer programmer and systems architect. Modular systems manage complexity by “black-boxing” information; that is, they separate code or information into discrete units. A programmer need only know about the module with which she is working, because managing the complexity of the entire system would be too much to ask of any single individual. Modularity is the method we’ve devised to manage complexity at a time when we’re drowning in information.

The computing historian Andrew Russell told me that “black-boxing reduces all kinds of cognitive and informational overhead, because you just know what the box spits out; you don’t need to know anything about what’s going on in there.”

Modularity, as Russell has documented, emerged as a term in architecture, and then spread to the military, where it was picked up to describe Project Tinkertoy, a post-World War II program to design interchangeable, self-contained parts for electronics. From there, the notion of modularity proliferated wildly, as a way of thinking about and structuring everything from organizations to economics to knitting. “It’s kind of a characteristic of modernity,” Russell said.

Supply chains are highly modular by design. Think of the shipping container. It wasn’t revolutionary because it was a box; it was revolutionary because it was a standardized, interchangeable box that could be locked in and transported. It makes globalization possible—it makes global scale possible—because of what it obscures. One doesn’t need to know what’s in the box, just where it needs to go.

How do you manage the complexity of a system that procures goods from a huge variety of locations? You make it modular: when you black-box each component, you don’t need to know anything about it except that it meets your specifications. Information about provenance, labor conditions, and environmental impact is unwieldy when the goal of your system is simply to procure and assemble goods quickly. “You could imagine a different way of doing things, so that you do know all of that,” said Russell, “so that your gaze is more immersive and continuous. But what that does is inhibit scale.” And scale, of course, is key to a globalized economy.

On the one hand, this all seems very logical and straightforward: to manage complexity, we’ve learned to break objects and processes into interchangeable parts. But the consequences of this decision are wide-ranging and profound.

It helps explain, for one thing, why it’s so hard to “see” down the branches of a supply network. It also helps explain why transnational labor organizing has been so difficult: to fit market demands, workshops have learned to make themselves interchangeable. It sometimes seems as though there’s a psychological way in which we’ve absorbed the lessons of modularity—although the world is more connected than ever, we seem to have trouble imagining and articulating how we’re linked to the other denizens of global manufacturing networks.

If technology enables a selective blindness that makes the scale of global supply chains possible, can technology also cure the problem of disavowal? Can software, having created the black box, help crack it open?

Recently, there’s been a lot of buzz about blockchain and the Internet of Things (IoT) among SCM practitioners. IoT technology would attach transmitters to components, so that their locations could be traced and monitored. With blockchain technology, each component that passes through a supply chain could have a unique, traceable ID number, and a log that registers every time it changes hands. Their proponents say that these technologies could bring radical transparency and unprecedented safety to global supply chains.

Blockchain is the technology that underlies bitcoins. The idea is that at each “stop” along a chain of users, a database associated with a particular coin (or component) updates to register the change of hands. The identity of each user could be cryptographically concealed, or it could be recorded transparently. Either way, the record of transactions is available to everyone along the chain, and it's near-impossible to forge.

Blockchain is security for the age of decentralization, and it could, in theory, make it possible for companies to verify the safety, composition, and provenance of manufactured goods. Supply Chain 24/7, an industry newsletter, calls blockchain a "game-changer" that "has the potential to transform the supply chain."

IoT is a different technology that addresses a similar problem. A company somewhere along a supply chain embeds a small transmitter, like an active RFID tag, in a component, allowing a monitor to see its location and status in real time. With sensors, a company could also keep track of the component's environment, checking on things like temperature and humidity. It sounds like a solution custom-fitted for the problem at hand: with these tiny trackers, companies could finally get the visibility they say they're after.

But the supply-chain specialists I talked to were skeptical. To make blockchain meaningful, Bonnani told me, you'd need to get every vendor to agree to disclose information about its practices; otherwise you'll just see a string in a database. "If you get suppliers to agree to be transparent, then blockchain is a way to verify that the thing you receive actually came from the person who sent it to you, and it's extremely valuable in that respect," said Bonnani. "But if you don't get them to opt in, then all you know is, you got what you asked for. They're not going to tell you who they got it from, or who that person got it from."

IoT lends itself to the same problems. Without genuine buy-in from suppliers, IoT "becomes one more technology to counterfeit," said Bonnani. "You're basically not improving the current problem, which is a lack of visibility." Given the pressure on suppliers to move quickly and flexibly, it's hard to imagine anyone volunteering more information than necessary.

One could imagine a system in which IoT and blockchain enable detailed information on labor conditions and safety, but the reality of global capitalism suggests that IoT is more likely to bring us smart toasters than socially responsible supply chains.

SCM innovation continues to thrive, but it's not trending toward the kind of visibility that Tony's Chocolonely is looking for. The newest technology that logistics professionals find exciting is machine learning, which involves creating algorithms that are capable of making predictions or decisions by "learning" from a set of data.

Machine learning is already in heavy use on the consumer side, where companies like Target use it to wager that a shopper who purchases unscented lotion, vitamins, hand sanitizer, and soft furniture might be getting ready to have a baby. But in the SCM world, machine learning could make it much easier to discover which suppliers and routes will deliver goods most quickly and reliably. A company could "predict the performance of each supplier, carrier, forwarder, port, lane, road, manufacturing facility, warehouse, etc. within the extended supply chain, under varying conditions," according to the SCM analytics company Transvoyant.

In a machine-learning scenario, companies could use historical data about manufacturers and goods to assign suppliers risk scores. Weather, changes in a supplier's political climate, or economic factors could all reassign risk scores, causing the supply network to automatically reconfigure itself to favor less risky suppliers. It's a stunning idea: supply-chain networks dynamically rerouting themselves in response to global risk factors, just the way Google Maps sends you down surface streets when the freeway is clogged.

This would increase efficiency, but at the cost of making it even more impossible to identify the supplier of your smartphone's LCD screen. It would aggravate, not alleviate, the problem of selective blindness that's already so deeply embedded in global supply chains.

In reality, the prospect of using machine learning on the manufacturing end of supply chains remains mostly speculative. When a company doesn't even know the most basic facts about its suppliers, it's hard to imagine how it would assemble the data necessary to develop efficient machine-learning models.

But its attraction for SCM specialists is notable, because it points to the kind of visibility that companies are talking about when they call for supply-chain transparency: not the kind of information that would help a consumer see where her candy comes from, but the kind of information that would get it into her hands faster and cheaper.

The challenges are political as well as technical, in other words. And the political challenges are immense. In the absence of real efforts to create democratic oversight of supply chains, we've come to see them as operating autonomously—more like natural forces than forces that we've created ourselves.

In 2014, the Guardian reported that Burmese migrants were being forced into slavery to work aboard shrimp boats off the coast of Thailand. According to Logan Kock of Santa Monica Seafood, a large seafood importer, "the supply chain is quite cloudy, especially when it comes from offshore." I was struck by Kock's characterization of slavery as somehow climatological: something that can happen to supply chains, not just something that they themselves cause.

But Kock was right, supply chains are murky—just in very specific ways. We've chosen scale, and the conceptual apparatus to manage it, at the expense of finer-grained knowledge that could make a more just and equitable arrangement possible.

When a company like Santa Monica Seafood pleads ignorance of the labor and environmental abuses that plague its supply chains, I find myself inclined to believe it. It's entirely possible to have an astoundingly effective supply chain while also knowing very little about it. Not only is it possible: it may be the enabling condition of capitalism at a global scale.

It's not as though these decentralized networks are inalterable facts of life. They look the way they do because we built them that way. It reminded me of something the anthropologist Anna Tsing has observed about Walmart. Tsing points out that Walmart demands perfect control over certain aspects of its supply chain, like price and delivery times, while at the same time refusing knowledge about other aspects, like labor practices and networks of subcontractors. Tsing wasn't writing about data, but her point seems to apply just as well to the architecture of SAP's supply-chain module: shaped as it is by business priorities, the software simply cannot absorb information about labor practices too far down the chain.

This peculiar state of knowing-while-not-knowing is not the explicit choice of any individual company but a system that's grown up to accommodate the variety of goods that we demand, and the speed with which we want them. It's embedded in software, as well as in the container ships that are globalization's most visible emblem.

We know so much about the kinds of things we can get and when we can get them. But aside from the vague notion that our stuff comes from "overseas," few of us can really pin down the stations of its manufacture. Is a more transparent—and more just—supply chain possible? Maybe. But, as the Chocolonely lawsuit demonstrates, it could mean assimilating a lot of information that companies have become very good at disavowing—a term that, in its Freudian sense, means refusing to see something that might traumatize us.

Adapted from the Logic magazine