

**From:** John Brockman <[REDACTED]>

**To:** Paul Allen <[REDACTED]>, Sergey Brin <[REDACTED]>, Jeff Bezos  
<[REDACTED]>, Ginevra Elkann <[REDACTED]>, Brian Eno  
<[REDACTED]>, John Doerr <[REDACTED]>, Jeffrey Epstein  
<[REDACTED]>, Tony Fadell <[REDACTED]>, Peter Gabriel  
<[REDACTED]>, Bill Gates <[REDACTED]>, "W. Daniel Hillis"  
<[REDACTED]>, Maja Hoffmann <[REDACTED]>, Jennifer Jacquet  
<[REDACTED]>, Bill Joy <[REDACTED]>, Dean Kamen  
<[REDACTED]>, Daniel Kahneman <[REDACTED]>, Salar Kamangar  
<[REDACTED]>, Vinod Khosla <[REDACTED]>, Ian McEwan  
<[REDACTED]>, Milner Yuri <[REDACTED]>, Dave Morin  
<[REDACTED]>, Elon Musk <[REDACTED]>, Craig Mundie  
<[REDACTED]>, Nathan Myhrvold <[REDACTED]>, Hans Ulrich Obrist  
<[REDACTED]>, Pierre Omidyar <[REDACTED]>, Nick Pritzker  
<[REDACTED]>, Larry Page <[REDACTED]>, Lucy Page/Southworth  
<[REDACTED]>, Lori Park <[REDACTED]>, Sean Parker  
<[REDACTED]>, Jean Pigozzi <[REDACTED]>, Muccia Prada  
<[REDACTED]>, "J.E. Safra" <[REDACTED]>, Ricardo Salinas Pliego  
<[REDACTED]>, Charles Simonyi <[REDACTED]>, Jeff Skoll  
<[REDACTED]>, Sebastian Thrun <[REDACTED]>, Tom Stoppard  
<[REDACTED]>, Jimmy Wales <[REDACTED]>, Evan Williams  
<[REDACTED]>, Anne Wojcicki <[REDACTED]>, Mark Zuckerberg <[REDACTED]>

**Subject:** George Dyson - NSA: The Entscheidungsproblem (The Decision Problem)

**Date:** Fri, 26 Jul 2013 15:59:47 +0000

---

NSA: THE DECISION PROBLEM

George Dyson

<http://www.edge.org/conversation/nsa-the-decision-problem>]

"The ultimate goal of signals intelligence and analysis is to learn not only what is being said, and what is being done, but what is being thought. With the proliferation of search engines that directly track the links between individual human minds and the words, images, and ideas that both characterize and increasingly constitute their thoughts, this goal appears within reach at last. "But, how can the machine know what I think?" you ask. It does not need to know what you think—no more than one person ever really knows what another person thinks. A reasonable guess at what you are thinking is good enough."

GEORGE DYSON, Science Historian, is the author of TURING'S CATHEDRAL: THE ORIGINS OF THE DIGITAL UNIVERSE, and DARWIN AMONG THE MACHINES.

[ED. NOTE: George Dyson's piece was commissioned by Frank Schirrmacher, co-publisher of the national German newspaper FRANKFURTER ALLGEMEINE ZEITUNG (FAZ), where he is Editor of the Feuilleton, cultural and science pages of the paper. First published by FAZ on July 26, 2013.]

---

NSA: THE DECISION PROBLEM

Shortly after noon, local time, on 19 August 1960, over the North Pacific Ocean near Hawaii, a metal capsule about the size and shape of a large kitchen sink fell out of the sky from low earth orbit and drifted by parachute toward the earth. It was snagged in mid-air, on the third pass, by a C-119 "flying boxcar" transport aircraft from Hickam Air Force base in Honolulu, and then transferred to Moffett Field Naval Air Station, in Mountain View, California—where Google's fleet of private jets now sit parked. Inside the capsule was 3000 feet of 70mm

Kodak film, recording seven orbital passes over 1,650,000 square miles of Soviet territory that was closed to all overflights at the time.

This spectacular intelligence coup was preceded by 13 failed attempts. Secrecy all too often conceals waste and failure within government programs; in this case, secrecy was essential to success. Any reasonable politician, facing the taxpayers, would have canceled the Corona orbital reconnaissance program after the eleventh or twelfth unsuccessful launch.

The Corona program, a joint venture between the CIA, the NSA, and the Department of Defense, was coordinated by the Advanced Research Projects Agency (ARPA) and continued, under absolute secrecy, for 12 more years and 126 more missions, becoming the most productive intelligence operation of the Cold War. "It was as if an enormous floodlight had been turned on in a darkened warehouse," observed former CIA program director Albert D. Wheelon, after the operation was declassified by order of President Clinton in 1995. "The Corona data quickly assumed the decisive role that the Enigma intercepts had played in World War II."

The resources and expertise that were gathered to support the Corona program, operating under cover of a number of companies and institutions centered around Sunnyvale, California (including Fairchild, Lockheed, and the Stanford Industrial Park) helped produce the Silicon Valley of today. Google Earth is Corona's direct descendant, and it is a fact as remarkable as the fall of the Berlin wall that anyone, anywhere in the world, can freely access satellite imagery whose very existence was a closely guarded secret only a generation ago.

PRISM, on the contrary, has been kept in the dark. Setting aside the question of whether wholesale, indiscriminate data collection is legal—which, evidently, its proponents believed it was—the presumed reason is that for a surveillance system to be effective against bad actors, the bad actors have to be unaware that they are being watched. Unfortunately, the bad actors to be most worried about are the ones who suspect that they are being watched. The tradecraft goes way back. With the privacy of houses came eavesdropping; with the advent of written communication came secret opening of mail; with the advent of the electric telegraph came secret wiretaps; with the advent of photography came spy cameras; with the advent of orbital rocketry came spy satellites. To effectively spy on the entire Internet you need your own secret Internet—and Edward Snowden has now given us a glimpse into how this was done.

The ultimate goal of signals intelligence and analysis is to learn not only what is being said, and what is being done, but what is being thought. With the proliferation of search engines that directly track the links between individual human minds and the words, images, and ideas that both characterize and increasingly constitute their thoughts, this goal appears within reach at last. "But, how can the machine know what I think?" you ask. It does not need to know what you think—no more than one person ever really knows what another person thinks. A reasonable guess at what you are thinking is good enough.

Data mining, on the scale now practiced by Google and the NSA, is the realization of what Alan Turing was getting at, in 1939, when he wondered "how far it is possible to eliminate intuition, and leave only ingenuity," in postulating what he termed an "Oracle Machine." He had already convinced himself of the possibility of what we now call artificial intelligence (in his more precise terms, mechanical intelligence) and was curious as to whether intuition could be similarly reduced to a mechanical procedure—although it might (indeed should) involve non-deterministic steps. He assumed, for sake of argument, that "we do not mind how much ingenuity is required, and therefore assume it to be available in unlimited supply."

And, as if to discount disclaimers by the NSA that they are only capturing metadata, Turing, whose World War II work on the Enigma would make him one of the patron saints of the NSA, was already explicit that it is the metadata that count. If Google has taught us anything, it is that if you simply capture enough links, over time, you can establish meaning, follow ideas, and reconstruct someone's thoughts. It is only a short step from suggesting what a target may be thinking now, to suggesting what that target may be thinking next.

Does this not promise a safer world, protected not only from bad actors attempting to do dangerous things, but from bad actors developing dangerous thoughts? Yes, but at what cost? There's a problem, and it's the problem

that Alan Turing was trying to answer when he first set us down this path. Turing delivered us into the digital age, as a 24-year-old graduate student, not by building a computer, but by writing a purely mathematical paper, "On Computable Numbers, with an Application to the Entscheidungsproblem," published in 1936. The Decision Problem, articulated by Göttingen's David Hilbert, concerned the abstract mathematical question of whether there could ever be any systematic mechanical procedure to determine, in a finite number of steps, whether any given string of symbols represented a provable statement or not.

The answer was no. In modern computational terms (which just happened to be how, in an unexpected stroke of genius, Turing framed his argument) no matter how much digital horsepower you have at your disposal, there is no systematic way to determine, in advance, what every given string of code is going to do except to let the codes run, and find out. For any system complicated enough to include even simple arithmetic, no firewall that admits anything new can ever keep everything dangerous out.

What we have now is the crude equivalent of snatching snippets of film from the sky, in 1960, compared to the panopticon that was to come. The United States has established a coordinated system that links suspect individuals (only foreigners, of course, but that definition becomes fuzzy at times) to dangerous ideas, and, if the links and suspicions are strong enough, our drone fleet, deployed ever more widely, is authorized to execute a strike. This is only a primitive first step toward something else. Why kill possibly dangerous individuals (and the inevitable innocent bystanders) when it will soon become technically irresistible to exterminate the dangerous ideas themselves?

There is one problem—and it is the Decision Problem once again. It will never be entirely possible to systematically distinguish truly dangerous ideas from good ones that appear suspicious, without trying them out. Any formal system that is granted (or assumes) the absolute power to protect itself against dangerous ideas will of necessity also be defensive against original and creative thoughts. And, for both human beings individually and for human society collectively, that will be our loss. This is the fatal flaw in the ideal of a security state.

When the creation of the U.S. Department of Homeland Security was announced, Marvin Minsky, one of Turing's leading disciples, responded that "what we need is a Department of Homeland Arithmetic." He was right. This sounds depressing. What do we have to do, turn all the computers off? No, we just need to turn off the secrecy, and conduct our data collection and data mining in the open, where it belongs. Ordinary citizens can tell the difference between regular police and secret police, and should be trusted to make the choice.

Consider the use of security cameras, for example in the UK. They are ubiquitous, visible, and used openly by the police under rules that have been defined in open court. Similarly, reasonable people might well support the maintenance of a global Internet memory buffer for law enforcement purposes, with access to the repository controlled under open rules by an open court.

There will always be illicit spying, but it should be kept within reasonable bounds. It is disturbing if laws had to be broken to conduct the PRISM surveillance program, but, if laws didn't have to be broken, that's worse. Edward Snowden has brought this matter before the public, and the path that led from Corona to Google Earth, through Silicon Valley, demonstrates that a secret program can be brought into the open, to the benefit of all, without necessarily being brought to a halt.

This is much bigger than the relative merits of national security vs. the fourth amendment to the U.S. Constitution, or any of the other debates by which the Snowden revelations have been framed. We are facing a fundamental decision (as Turing anticipated) between whether human intelligence or machine intelligence is given the upper hand. The NSA has defended wholesale data capture and analysis with the argument that the data (and metadata) are not being viewed by people, but by machines, and are therefore, legally, not being read. This alone should be cause for alarm.

And what of the current obsession with cyberterrorism and cyberwar? We should deliberately (and unilaterally if need be) abandon the weaponization of codes and the development of autonomous weapons—two different approaches to the same result. They both lead us into battles that can never be won. A good example to follow is

the use of chemical and biological weapons—yes, they remain freely available, but we have achieved an almost universal consensus not to return to the horrors of poison gas in World War I. Do we have to repeat the mistake? We are currently taking precisely the wrong approach: fast-tracking the development of secret (and expensive) offensive weapons instead of an open system of inexpensive civilian-based defense.

Fourteen years ago, I spent an afternoon in La Jolla, California with Herbert York, the American physicist of Mohawk ancestry who became Eisenhower's trusted advisor and one of the wisest and most effective administrators of the Cold War. York was appointed founding scientific director of ARPA and was instrumental both in the development of the hydrogen bomb and its deployment, in a few short years, by a working fleet of Intercontinental Ballistic Missiles, or ICBMs. He was sober enough to be trusted with the thermonuclear arsenal, yet relaxed enough about it that he had to be roused out of bed in the early morning of July 6, 1961, because he had driven someone else's car home by mistake.

York understood the workings of what Eisenhower termed the military-industrial complex better than anyone I ever met. "The Eisenhower farewell address is quite famous," he explained to me over lunch. "Everyone remembers half of it, the half that says beware of the military-industrial complex. But they only remember a quarter of it. What he actually said was that we need a military-industrial complex, but precisely because we need it, beware of it. Now I've given you half of it. The other half: we need a scientific-technological elite. But precisely because we need a scientific-technological elite, beware of it. That's the whole thing, all four parts: military-industrial complex; scientific-technological elite; we need it, but beware; we need it but beware. It's a matrix of four."

We are much, much deeper in a far more complicated matrix now. And now, more than ever, we should heed Eisenhower's parting advice. Yes, we need big data, and big algorithms—but beware.

---

JB

John Brockman

Mobile [REDACTED]

---

President [REDACTED]

Edge Foundation, Inc. tel: [REDACTED]

fax: [REDACTED]

New York, NY 10001

---

Visit EDGE <http://edge.org/> ("The World's Smartest Website", The Guardian) <http://edge.org/edge-in-the-news>