WHAT WAS I THINKING?

Dana Charmian Angluin

I've always felt that I belonged to the After. As a card-carrying Boomer, the thing that I was after, of course, was the war—not World War II, just "the war." My family background made the war particularly salient.

My mother was born at West Point in 1914 while her father was stationed there teaching math. Her father, West Point class of 1909, is my namesake, Dana Crissy. As a little kid, I was deeply puzzled when my mother and grandmother would talk about "Dana," and not mean me.

A 1919 article in the *San Francisco Examiner* has a picture including Dana Crissy. The caption reads ". . . photos show the daring Western aviators who this morning will 'hop off' [from the Presidio] on a sensational transcontinental air race to Mineola Field, New York . . ."

"Hop off" they did, and Major Dana H. Crissy and his mechanic Sergeant Virgil Thomas were killed that first day, October 8, 1919, when their De Haviland 4 crashed on landing in Salt Lake City. The airfield at the Presidio was named Crissy Field in Dana's memory, and the widow and her two daughters were allowed to continue to live on the base for a time. My mother was five years old.

An editorial cartoon quoting the US War Department as promising "Crissy Field S.F. to be the Strongest Air Base in US" was not prophetic. After the Golden Gate Bridge was built, Crissy Field was suitable only for small planes and helicopters. It later became part of the Golden Gate National Conservancy.

Fast forward twenty years and we find my mother married to a 1929 graduate of West Point. A 1939 article in the *San Francisco Chronicle* shows my father, Captain David X. Angluin, escorting the governor of California in the review of troops at the Golden Gate International Exposition. The exposition was a kind of mini-world's fair held on Treasure Island in San Francisco Bay.

The dashing red-headed captain was of Irish stock, born in 1908 Lowell, Massachusetts, where he and his brother were raised by their widowed father. His

---

Dana Charmian Angluin received her PhD in engineering science (1976) from U.C. Berkeley under the direction of Manuel Blum. She then had postdocs with Leslie Valiant (Leeds and Edinburgh Universities) and Ronald Book (U.C. Santa Barbara). From 1979 to 2021 she was a member of the Yale Department of Computer Science in a number of roles: Gibbs Instructor, assistant professor, associate professor on term, research scientist, senior research scientist, and professor. She helped start the influential Computational Learning Theory (COLT) conference. Her teaching of Yale undergraduates was recognized by the award of the Dylan Hixon '88 Prize for Teaching Excellence in the Natural Sciences (2011), the Harwood F. Byrnes/Richard B. Sewall Teaching Prize (2020), and the Yale Phi Beta Kappa chapter's William C. Clyde DeVane Medal (2020).

father had been a steamfitter but became a school janitor after an injury. Young David took the exam for West Point and won appointment from Massachusetts, arriving at the Academy at the age of seventeen.

At West Point, Cadet Angluin graduated 226th out of his class of 299, in the bottom quartile. The courses at the time were not exactly a liberal education: two years of mathematics (only as far as calculus), a year of physics, a semester each of chemistry and electricity, two years of English, two years of French, a year of Spanish, one year of history, one year of law, one year of economics and government, and a number of military subjects—tactics (all four years), military art and engineering, military topography and graphics, ordnance, and military hygiene.

In the transcript, the tactics courses are footnoted: "Exclusive of military drills and maneuvers, equitation, physical training and compulsory athletics, which are part of the cadet's daily life throughout the entire course."

After Hitler invaded Poland, Britain declared war on Germany on September 3, 1939, and the small peacetime army of the US would not remain that for too much longer. I don't know much about my father's service in the war—like most of the war's veterans, he talked about it very sparingly. He trained in the California desert for action in North Africa and then was sent to the assault on Attu, an island in the Aleutian chain off Alaska. Later, he was assigned to supply for the assaults on Kwajalein and Leyte in the Pacific.

A Bronze Star citation for Attu says he was able to "perfect and standardize supply procedure to a highly efficient degree thereby contributing immeasurably to the success of the operation." A citation for the oak-leaf cluster to the Bronze Star for Kwajalein cites his "exceptional knowledge of logistics" and "careful planning." However, something happened during the Leyte operation that was always described as a "mental breakdown," and he was eventually retired from the Army at the rank of Colonel.

My brother David was born before the US entered the war, in September of 1940. During the war, my mother lived in Carmel, California, and took care of the little asthmatic boy and relied on a big German shepherd named Bill (and possibly a .38 revolver) for some sense of security.

We come to the After. In 1950 my family moved to "San Rafael Meadows," one of the first post-war housing developments near US 101. The land had previously been used for cattle ranching, and ranchland was still a short walk from the house during my childhood.

My best friend, David Worswick, lived across the street. He and I would roam the surrounding cattle ranches looking for skulls to collect. Two of our prizes finds were a cow skull and a horse skull, which decorated the fort we had built out of discarded railroad ties. Another pastime was catching snakes, which refused to eat in captivity and were released after a short stint in a cage. As we distantly approached puberty, David Worswick's family decided it was time for him to stop playing with girls.

Aged nine.

As for my career aspirations at this age, I was asked in school about what I wanted to be when I grew up, and I said, “a writer and a mother.” When asked about college, I couldn’t really see voluntarily extending one’s required time in school and said that maybe I would be a cabinet maker. Another option that appealed to me was taxidermist.

I made a new best friend in seventh and eighth grades, a kindred tomboy, Sharon Truett. One of our triumphs was a lemon chiffon pie produced in home arts that was the hit of the teachers’ lounge, to the point that we were asked for an encore. We might have preferred to take industrial art but could not, being girls. We did get excused from class to skin a dead skunk trapped, gassed, and brought in by our geography teacher.

After nearly completing a teaching credential, my father had decided he was not cut out for high school teaching and made a living selling insurance. However, my mother wanted the chance to live abroad and convinced him to apply to the US Civil Service. His logistics background made him suited to the new use of data processing equipment in Army supply. And so, our family was introduced to computers.

Two things came of that: after I finished eighth grade, my parents and I moved to Orléans, France for two years, and my father became convinced that computers were the wave of the future.

In France our lives centered around the US Army base in Orléans. We lived in the city, renting a three-story house whose steel shutters were still marked with bullet holes from the war. I went to Orléans American High School, with its mixture of children of officers, enlisted men, and civilian employees. I was in the outcast group who ate brown bag lunches in a utility outbuilding. Between adolescence and being uprooted from all of my childhood, I became very inward. My father wanted to expose me a little to computers, so one Saturday we went to his (very secret) workplace—a computer system that was located on train cars on a siding, so that it could be easily relocated in case of hostilities. The computer system was a major element in the system of supply for all of the US Army in Europe. He set me to transcribing some data onto sheets for the keypunch operators to key in later.

My mother was very ambitious for my education. She successfully argued for me to take sophomore biology as a freshman and somehow arranged that I would get acceleration in mathematics. As a sophomore, I took geometry and was also tutored during study hall in advanced algebra and trigonometry by my math teacher.

I had done fine in math previously but had not been accelerated. I remember as a kid in the shower working out a rule to determine how many hours there were between say 7 p.m. and 3 a.m.—that you could subtract the smaller from the larger, in this case four, and subtract that from twelve, to get the answer of eight hours. I checked the rule on a number of cases and announced it to my family. My brother (at the time in high school) recognized the rule and pointed out that I had discovered modular arithmetic.

My brother had flunked out of Reed College in his first year and subsequently worked at various jobs in San Francisco. He was also sporadically taking courses at San Francisco State and U.C. Berkeley. Informed by his unfortunate experience of abstract mathematics at Reed, he special-ordered from Blackwell's books a copy of Thurston's classic *The Number System* and sent it to me. I was particularly intrigued by the axiomatic development of group theory in that book. I scoured the tiny post library for math books and found one that covered the history of mathematics, including something on the approximation theory of Eudoxus.

At the initiative of my parents, I took the SATs toward the end of my sophomore year and achieved very good scores, including an 800 in advanced math. Somehow on the strength of the SAT scores, my parents convinced the University of California at Berkeley to admit me under their small quota of nonqualifying students for fall 1964. My father further convinced them to charge in-state tuition!

So there I was in the fall of 1964, aged sixteen, in an hours-long line to complete registration at one of several card tables set up in Harmon Gymnasium. When I got to the front of the line with my punched cards and was asked what I wanted to major in, I said, "metamathematics," and they understandably said, "What's that?" I explained that it was a kind of philosophy, which is how I got classified as a philosophy major. My brother was a philosophy major, so I figured it would be good to follow in his footsteps.

That first semester, I lived with my parents in a rental on the north side of campus and took history, speech, logic, and honors calculus. This was in an era before most high schools had calculus courses, and the honors course was intended for prospective math majors. It was taught by Leon Henkin, a celebrated logician, who paid extraordinary attention to the logical niceties of definition and proof. I still remember the pleasure I felt in taking the final exam for the first semester, which consisted of giving proofs to re-derive the foundations of the real numbers on the basis of Dedekind cuts rather than the least upper bounds that had been used in class.

His course had a profound effect on my enthusiasm for, and taste in, mathematics. I developed an abiding love of logic and proof and also a kind of blind spot for applications of analysis. I remember being quite shocked in the second semester when he very carefully drew a circle to illustrate the concept of radian measure of an angle — that was the very first picture I remember him drawing in the course.

That first semester was also the semester of Berkeley's Free Speech Movement. Berkeley's enrollment had swollen with the leading edge of the baby boom. Everything was crowded. I was unable to attend my history lecture in Wheeler Auditorium but had to watch it on closed-circuit TV in Dwinelle Hall. There was a cap of 27,500 students, but enrollment estimates had been low, and enrollments exceeded the cap. Registration began September 16, and I registered the following day.

That same day, the university administration banned the lively and crowded political solicitation tables at the Bancroft and Telegraph entrance to campus. The affected student groups pushed back, and the administration attempted various compromises to regulate the "time, place, and manner of speech." On October 2, 1964, the police tried to arrest non-student Jack Weinberg, who was manning an unauthorized political table and refused to show his identification. They placed him in a police car in Sproul Plaza, and students present reacted by sitting down around the police car, preventing its departure. Weinberg would be immobilized in the police car for thirty-two hours.

The controversy continued during the semester. Two months later on December 2, 1964, over one thousand students, led by Joan Baez singing "We Shall Overcome," occupied Sproul Hall. I watched Mario Savio's speech, watched them walk into the building, but did not join them, acutely aware that I was a minor of sixteen and that my admission to Berkeley under special conditions was quite revocable.

There was quite a lot of shock at the fact that hundreds of students were arrested on the campus. Governor Edmund Brown authorized the arrests. Ronald Reagan somewhat later made quite a successful political career out of promising "to clean up the mess at Berkeley."

A few days after the arrests, the Berkeley Academic Senate voted 824 to 115 that "the content of speech or advocacy should not be restricted by the University." By then I was sympathetic to the Free Speech Movement, and the faculty resolution restored my faith in the faculty. My parents were dismayed at these developments, quintessentially "messy" from my father's point of view. This was not the end, but just the beginning of unrest on Berkeley's and other college campuses. I recall instances of helicopters flying over the campus spraying tear gas to disperse crowds. One of my own timid acts of rebellion later in this period was to post anonymous copies of indirect criticisms of math instruction in buildings on campus.

Despite all the turmoil, I completed my first year. My father had looked around and found a summer course at Oakland City College that would teach me to program an IBM 1401 computer. The IBM 1401 was a very successful business computer, and to sell the machines, IBM had a pressing need for more people who could program it. There were aptitude tests to identify people with the mental skills to learn programming. My father convinced my brother to take one, and he passed it and started a programming course. When they got to a statement like "X = X+1," my brother, perfectly logically, said to himself, "no, X is never equal to X+1" and ended up giving up on the course. He stuck with philosophy.

There was no simple way for me to access a programming course at that time at U.C. Berkeley—it would have been in electrical engineering and require getting through courses in physics and electronics, basically ensuring that one could build a computer before being instructed in programming one. Physics and I had parted company after one semester. The advantage of the Oakland City College course was that it was intended for students who would work as programmers in businesses using the IBM 1401 and had minimal prerequisites. I attended the course. The first day was spent learning the associated tabulating equipment, using keypunches to punch cards and other machines to duplicate, sort, and print decks of cards. That night, my sleep was troubled by the clunking sounds of those machines in my brain. But we moved on to programming, and I found it challenging and fascinating.

As with most interactions with computers in those days, the programmer prepared a deck of cards and handed them to an operator and then waited for the next day to get the deck back, along with the results of running the program overnight. This meant there was a premium on very thorough understanding of what the program would be doing and developed in me the habit of simulating the program on paper to check it.

Later that same summer, my mother, father, and I went to the tennis courts at the Berkeley Rose Garden in search of a fourth for doubles. The fourth turned out to be the director of employment at Berkeley, and my father told him that I had just successfully completed a programming course and asked if there were any suitable jobs available on campus. He replied that there was an opening at a research project in Cory Hall and made some comment about the group being a bit unorthodox. But I followed up, and that is how I came to be starting work on my seventeenth birthday as a coder at Project Genie.

The goal of Project Genie was to create a time-sharing system so that a number of people could be using a computer simultaneously, but it would seem as though each one were the only person using it. The feasibility of this idea depended on the fact that people writing and running programs spent a lot of time thinking or typing, so that if the computer were quickly switched from running one person's job to another's, to another's, round-robin, then all of them could get useful work done, and each would have an interactive experience. There would no longer be the intermediary of an operator or the overnight wait to see what happened when the program actually ran.

This required that each user be provided with a terminal through which to interact with the computer, which, in the case of Project Genie, was a teletype. A teletype was an electromechanical device with a keyboard and printing mechanism that could be used to send and receive text messages over dedicated wires. With a modem to translate the electrical pulses to and from sound, the text could be sent over regular telephone wires — at the blinding speed of ten characters per second! This enabled the computer to be used not only interactively, but also remotely.

The project involved modifying the SDS 940 hardware to support time sharing and writing the operating system and all the supporting programming tools — an editor, assembler, debugger, and so on — from scratch in the machine language of the SDS 940. Project Genie was a major formative influence in my understanding of computers and their capabilities.

My initial project was to write a program to send text to the printer to be printed. I studied the manuals and wrote it in the most modular way I could. In the event, my choices made it much too slow, and Peter Deutsch ripped it apart, de-modularized it, and produced a suitably fast program.

I came to be responsible for the text editor QED, used to write programs and documents. The story went that Peter had become impatient with some other coder's slow progress writing the editor, so he took a weekend to write it. The limitations of the system are hard to convey: the code for the program needed to fit in two 2K blocks of 24-bit words, equivalent to about 12KB of text. The QED editor eventually acquired its own Wikipedia page.

Meanwhile, my parents, dismayed at developments in Berkeley, sought to return to Europe. I moved out to a rented room on the south side of campus, and they left in the summer of 1966. De Gaulle kicked the US Army out of France in March of that year, so my parents ended up at a US base in Zweibrücken, Germany. Our communication was reduced to weekly air letters, and I was pretty much on my own, which I welcomed because our relations had gotten pretty tense. But my sense of myself, and what to expect from the future, was profoundly destabilized. I read and wrote a lot, pretty randomly, always hoping for enlightenment.

In the fall of 1966, Berkeley switched to the quarter system. I made it through the first quarter, which seemed extremely rushed and stressful compared to the previous semesters, and contemplated the prospect of two more like it and dropped out. This distressed my mother no end — she refused to admit it to her friends and acquaintances. Fortunately, I could continue to support myself by working at Project Genie. I eventually went back, dropped out again, and then finished my degree in the summer of 1969. By then I had figured out that I was not cut out for philosophy and had switched to math, all math, so that I had over 100 of my 180 required credits in mathematics by the end, including some graduate courses.

You would think that a student who had gotten that many credits in math, nearly all of them As, might be encouraged to consider graduate school, but you would be wrong. As I look back, I notice now that not one of my professors, or even TAs, was a woman, in all of my Berkeley courses. It was not something to notice then; it was just how things were.

What were those As worth? Maybe not so much because the war in Vietnam and the pressing needs of the wartime draft meant that for a young man, flunking out of college would likely lead to fighting in Vietnam, and not many instructors wanted that on their consciences. Among my Berkeley male friends liable to the draft, one left for Canada and returned after the amnesty; one was too nearsighted; one documented his homosexuality; one put on excess weight for every physical; and one worked at a naval shipyard, and, when that was not enough, joined the Coast and Geodetic Survey, which was one of the seven uniformed services.

Two significant influences emerged from those years of turmoil. One was the books of Paul Goodman, starting with *Growing Up Absurd*, which acknowledged a lot of the problems that the young were seeing in society but very strongly made the point that human civilization had within it the resources to address those problems and that mathematics was among those resources. (A few readers will remember when a fat paperback cost $1.95.)

Another was a used paperback I bought in the spring of 1967—Pólya's book *How to Solve It*. Throughout my years in math courses at Berkeley, problem-solving ability was treated as a fixed in-born trait—you had it or you didn't. Pólya was a mathematician who treated problem-solving as something to be learned and taught. The book is mostly very practical advice to teachers of mathematics for how to help their students learn to solve problems. It offers a series of questions that the instructor can ask the student, and the student can learn to ask him- or herself, to help conceptualize and structure the process of problem-solving. It is based on what is now known as a "growth mindset," but it was a revelation to me at the time and continued to influence my teaching throughout my career.

One institutional response to the Free Speech Movement was an experiment devised by the famous measure theorist Michel Loève in which he recruited upper division math students (including me) each to serve as a TA to a very small number of first-year calculus students, with lectures by him. This was along the lines of Yale's "peer tutor" system and gave me a taste of teaching.

After graduating with a bachelor's in mathematics in the summer of 1969, my plan was to make my living as a programmer. I then worked at a couple of companies that would now be called startups.

There was a company called MASCOR in Sunnyvale. I mostly worked remotely and got a ride to Sunnyvale once a week or so. This was when Silicon Valley was still mostly prune orchards. I remember one time I took the bus to get there and walked the last portion past a very fragrant, freshly harvested field of onions. I stipulated a salary of three dollars per hour, with a bonus to be paid when the project was complete, but the company went broke first.

Online Decisions was located in downtown Berkeley. This was a business that used the new technology of time-sharing to sell computerized analysis and planning tools to businesses, typically in San Francisco. I was part of the back-office staff, responsible for maintenance and improvement of the system. The sales staff would go out to businesses to demonstrate the capabilities of the system, and, once they had made a sale, to customize a model of the company's operations. The most exciting times were when a salesman at a company would call to say the system had just crashed, and I needed to find and fix the underlying problem in – ideally – nothing flat. In between the exciting parts, I spent my time reading, understanding, and cleaning up the code so that there might be fewer exciting times.

This was my first actual experience with business, and I was amazed to find that companies did not make a careful rational analysis comparing the costs of our product with the costs of their current procedures. Instead, the decisions were more about the career ambitions of the junior managers who advocated adopting the system and senior managers' fears of not keeping up with the latest innovations. I ended up quitting because I felt the owner of the company was not listening to the staff. The owner offered me a raise – a rational business response, but one that I rejected.

I was hired at another computing research project on the Berkeley campus, this time writing a debugger. Eventually the professor who was my boss told me that it would be a lot easier to pay me if I were a student, so I applied to the master's program in electrical engineering and computer science and was accepted.

However, that meant that I had to start taking classes again. In the spring of the year, I took a class in automata and formal languages taught by Manuel Blum, who was covering it because the usual instructor was on leave. I became fascinated with the material, a mathematical treatment of what computers can and cannot do, and when the course ended, I asked Manuel if he would supervise me in some research. His answer was "no," that is, "no, not unless you become a PhD student." So I applied and was admitted to the PhD program in 1972 and became Manuel's student. No longer programming, I was supported on his research grant.

It was quite a time to be at Berkeley, with the excitement of analysis of algorithms, P and NP, and complexity theory more generally. One of the major influences on me at that time was Mark Gold's definition of "identification in the limit" as a formal model for learning a language, introduced in his 1967 paper.

In Gold's model, a stream of data items is provided to a computer, which responds with a sequence of guesses of rules (in the form of computer programs) to try to explain the data. The issue is then to define what it means for the computer to succeed at this task. The criterion of success is that if the stream of data eventually includes every possible example, then the guesses of the computer should stabilize after a finite number of tries and be correct. One caveat was that the computer might not be able to tell whether this had happened!

Another very exciting development at that time was the theory of NP-completeness. P denotes problems that can be solved in polynomial time, that is, by programs running in time bounded by a polynomial in the size of the input; these were generally thought of as "tractable," whereas NP denotes problems whose answers can be checked (though not necessarily found) in polynomial time. In 1972, Dick Karp presented his foundational paper about NP-completeness, showing the power of polynomial time problem reductions to suggest the computational intractability of problems like the "traveling salesman problem" that had long resisted the search for polynomial time algorithms. The foundational question of whether P is equal to NP is still open over fifty years later—I no longer expect to see it settled in my lifetime.

My first exposure to the concepts of P and NP came in a graduate class given by Gene Lawler on matroid theory. Gene had just had lunch with Dick Karp and was so excited about the concept of NP that he gave an impromptu substitute lecture about it. The lecture went right over my head, but fortunately my fellow student Len Adleman caught the ideas and later explained them to me. This episode was in my mind whenever I subsequently taught the concepts of P and NP, reminding me of how unintuitive NP could seem at first.

My mental image of the research area I was trying to pursue in graduate school was the collision of the world of inductive inference on one hand and the world of complexity theory and analysis of algorithms on the other. Hovering like a phantom was artificial intelligence, which was largely applied at the time. I was impatient with the criterion of identification in the limit, an attitude captured in the quip, "In the limit, we are all dead." (At the time I was unaware of the Keynes quote: "In the long run, we are all dead.") My hope was to apply complexity theory and analysis of algorithms to quantify how quickly learning could take place for various kinds of concepts.

In graduate school there were years of false starts: trying to prove a lower bound on Boolean matrix product, trying to give a theoretical account of "ten-second sequences," and many more. Then another advance by Mark Gold showed a path forward, when he proved that the problem of finding a finite state automaton with the minimum number of states consistent with a set of positive and negative examples is NP-hard. My thesis was essentially variations on Gold's theme and got me out of Berkeley.

My brother had ultimately graduated from Berkeley and gone off to do a DPhil at Sussex in England. As I finished my dissertation, he came across a postdoc offered by a young faculty member at Leeds that he thought sounded appropriate. I interviewed over the phone and headed off to Leeds to a postdoc with Les Valiant. When I got there, Les told me we were moving to Edinburgh in January.

The work Les and I did together concerned a polynomial time algorithm to find Hamiltonian circuits with probability approaching one in random graphs. Thus, despite being NP-hard in the worst case, the Hamiltonian circuit problem is solvable in polynomial time for most graphs. I also taught a short course in number theoretic algorithms that was overwhelmingly drawn from the hours I had spent as a graduate student listening to fellow student Gary Miller's explanations of his research.

It was then uncommon for a computer science PhD to do any postdoc, but after two years in Leeds/Edinburgh, I went on to another year as Ron Book's postdoc in

Santa Barbara. There I worked on both complexity theory and what I thought of as "concrete algorithms for inductive inference." I published a paper on "pattern languages" inspired by my less than successful work on the ten-second sequences. After a year as a postdoc at Santa Barbara, I was hired at Yale as a Gibbs Instructor in Computer Science, a kind of glorified postdoc, on a two-year appointment, with the promise of a one-year extension.

I soon acquired two wonderful graduate students: Bob Nix, a systems student whose adviser had left Yale, and Udi Shapiro, who came to Yale to work with Roger Schank and then found it impossible to get along with Roger. Bob and Udi decorated my office with balloons the day I snuck off to get married to my colleague Stan Eisenstat, which turned out to be the day that President Reagan was shot.

I was publishing papers in conferences and journals, had acquired two PhD students, and had a theoretician-sized grant from the NSF, including student support. Then in December of 1982 I was notified that I wouldn't be considered for tenure, since that required being one of the "ten best in the world" in my field. And I might not even be recommended for renewal as an untenured associate professor. Fortunately, Yale and I managed an accommodation that let me stay around for the subsequent forty years. The arrangement probably enabled me to experience more of my children's childhoods than otherwise.

I would like particularly to acknowledge how much I learned about teaching, and about Yale, from my husband, Professor Stanley Eisenstat, who passed away in December 2020. Perhaps you knew him as a regular at faculty meetings and a stickler for correctness, but he was also the most patient and sympathetic of teachers, particularly for students who struggled with the challenging material of his courses. If the student was willing to keep trying, Stan was willing to put in unlimited hours to keep helping. I owe the lion's share of the three Yale teaching prizes I was awarded over my career to Stan.

Just after I resigned my regular faculty position at Yale, I published what is probably my favorite paper, "Learning Regular Sets from Queries and Counterexamples." I think it is the first use of the terminology "Learning" in my published work, though that is what I was studying all along. The afterlife of the paper—more than thirty-five years on—is surprisingly robust: it managed 825 citations in 2022.

The terms "machine learning" and "artificial intelligence" have lately become synonymous in the popular imagination with so-called deep neural networks, particularly large language models like ChatGPT. With a lot of help from Yale colleague Bob Frank and former PhD student Sophie Hao, I have recently ventured into the exciting and somewhat terrifying domain of deep neural networks for natural language processing, after strenuously resisting neural networks for all those years.

One of the recent advances in natural language processing is a "word embedding," which maps words to high-dimensional vectors. I've been trying out ways to visualize the content of a word embedding. In the accompanying illustration I've taken the center word, “thanks,” and found the seven closest (as vectors) words to it and then the seven closest words to them. There is an edge from a word to another word if the second word is among the seven closest words to the first word. We can see that the seven closest words to “thanks” in this embedding are: “despite,” “enjoyed,” “success,” “good,” “improved,” “boosted,” and “thank.” The colors pick out clusters of more closely related words. I am grateful for your attention, hope you enjoyed this account, and always look to improve.

53

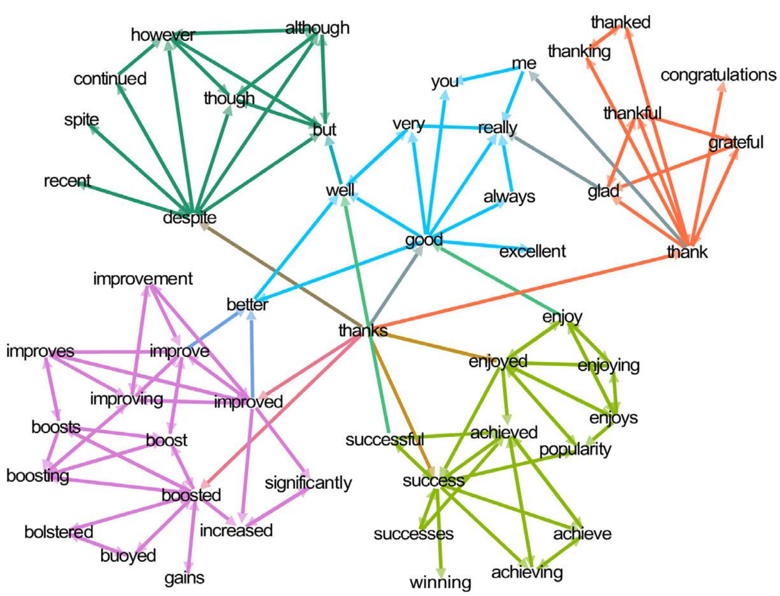

Thanks.