Commentaries on publication

A Biographical Background to our Papers

Brian Gaines, October 2026

In introducing the web site, we promised to provide commentaries on our publications, but Mildred has died and I am very ill, so this is a last-minute attempt to embed the commentaries in a short biography. I have only been able to provide a partial biography of Mildred, so she is sadly under-represented.

I realized that much of our later careers depended on far earlier events in our lifes, so this starts in our early childhoods. The aim is not so much to provide a full biography as to provide a setting for our publications, so that an associated AI query system might be able to answer questions about the context in which a paper was written as well as those about its content.

Mildred's Early Years

Mildred was born in Bradford on 27th February 1945. She was a late developer. She told me that she dreamed her way through school with little awareness and that her primary interest was poetry. Her personal book collection was mainly poetry books acquired in those years. Her father was a bank manager, moving from town to town and there was little continuity to her early life. When she finished school she went to Southlands College, University of London for teacher training from 1963 to 1966 and found a career that engaged her interest. She always had a great love of facilitating students' learning.

After college, she taught at a school in Woodford, taking the lower-level C stream and getting the less intellectual children engaged in mathematics and science. She also taught physical education. She furthered her own education whilst a teacher, taking A-level[1] physics, then a degree in mathematics at Birkbeck College from 1967 to 1971, an MSc in computer science at the Institute of Computer Science from 1971 to 1972, both at London University. Her MSc thesis was a critical assessment of Bill Broderick's historic introduction of computers into school education (A Critical Survey of Computer Education in Schools in the London Borough of Havering, 1972). Then she joined Middlesex Polytechnic, London from 1972 to 1982 as Senior Lecturer in Education/Mathematics and

[1] In our era an O-level exam was taken at 15, and one could leave the school system, and an A-level at 17 as a qualification for higher education

Computer Science. While she was there, she took a PhD in Psychology at Brunel University advised by Laurie Thomas and Gordon Pask, developing computer programs for eliciting and analysing George Kelly's conceptual grids for modelling a person's conceptual structures (On Becoming a Personal Scientist: Interactive Computer Elicitation of Personal Models of the World, 1978). All her qualifications were obtained part-time while in full-time employment.

I first met her as her doctoral examiner and again later at various conferences, fell in love, got married on 9th February 1982, and moved to Toronto in August 1982. Our biographies merge thereafter.

Brian's Early Years

I was born on 25th October 1940 during the Second World War. My parents, Frank and Ruby, lived in Wandsworth, London where the firebombing was intense and my mother was evacuated to a nursing home in Cirencester to have me. My earliest memories are of a train in flames on the railway bridge near our house as I walked along the street in my pyjamas and dressing gown to the nearby Anderson shelter where the lights were on all night and people were talking and I learned to sleep under any conditions. That has been helpful throughout my life, as colleagues have often remarked!

When I was two my mother, her elder sister Aunt Billie and I were evacuated to the gamekeeper's cottage of the Crosswaite Ayres' country estate in the New Forest. My grandmother was the housekeeper of their London home, and I assume they were fulfilling the expectation that the landed gentry would take in refugees from the war zone by doing this for their servants' relatives. We were guests of Algy the gamekeeper, his wife Anne and their daughter my age Susan.

It was a very pleasant time living in the peace and tranquility of the forest with ferns taller than I was away from the horrors of the bombing. My aunt took it upon herself to educate me in reading, writing and arithmetic and behaving as a young gentleman. I delighted in the lessons, experimenting with numbers and the strange behaviours of non-terminating decimal fractions, experiencing the joy of getting information from books and being able to write down my ideas.

We returned to London in May 1945 when the war ended. I remember touring London in a bus to watch the VE Day[2] celebrations. Our house had survived but my toy cupboard had been scorched by the heat. Shortly afterwards we were moved to a new council home in a

[2] Victory in Europe, still celebrated in the allied nations

block of flats, Comins House, at Clapham Common. I commenced school there in October 1945.

The school found I was already numerate and literate and parked me in the school library. It was only a small room but contained more books than I'd ever seen before and was paradise to me. They also involved me in coaching the other children and I found the experience was very satisfying and facilitated my own learning. I learnt little from the lessons as it was so much faster to read the relevant books. I learnt more from science and technology comics like *Eagle*[3] where I first came across nuclear reactors, computers and (fictitious) spacecraft.

When I was ten, we moved to a new council estate at the boundary of Eltham, in London and Sidcup in Kent and I was able to persuade my parents to join libraries in both jurisdictions and let me have the tickets. This let me borrow 24 books from the adult sections of the library and enormously increased the scope of my reading. I began to study advanced mathematics, chemistry and physics.

School Years

When I was 11, I transferred to Shooters Hill Grammar School after an interview by the headmaster, Mr Hanks, with me and my parents. I was ill-suited to a school that aped the public school system with a Head Boy, prefects, a house system targeted on sports, Rugby, etc – I ignored it all. I enjoyed the curriculum with its rigorous standards, particularly because there was a textbook handed out at the start of each course that often had additional interesting material that one could study, for example, Heath's remarks on the history of Euclid's geometry in his exposition of the *Elements*.

I particularly enjoyed the methodological teaching, such as how to write up an experiment, which I have found useful throughout my career and generalizes to many formal writing styles. When I get involved in a new discipline, I have learned to identify its methodologies to understand its rhetoric as early as possible.

Hobbies

For a while my favourite hobby was micro-analytic chemistry because I was fascinated by the idea of taking a small quantity of any substance and finding its chemical composition. I was able to persuade our local chemist to provide me with some reagents like barium nitrate after he warned me that they were poisonous, but I became concerned when I found out that one of the essential reagents was hydrocyanic acid which scared me. I tried

[3] https://en.wikipedia.org/wiki/Eagle_(British_comics)

macro-analytic chemistry, but I had no fume cupboard and my father objected to the smell of hydrogen sulfide – so I turned to electronics.

I had long been servicing radio sets, with little knowledge at first, but eventually reading Wireless World[4] and Practical Wireless[5] regularly and gaining much practical knowledge such as the colour coding of resistors. There were several war surplus shops in the Charing Cross area that sold electronic components and equipment at a low price. I started building audio amplifiers for myself and my friends, including ones for guitars with built in loudspeakers in elaborate cases because I also liked carpentry and French polishing.

In 1953 Practical Wireless published an article on how to build an oscilloscope and I bought a rocket test set from a war surplus store with a VCR97 cathode ray tube and EF50 pentode valves, designed my own 2KV power supply based on a Cockcroft–Walton multiplier, and had an oscilloscope that served me well for several years. I took it to a parents' open day at the school in 1954 and fascinated people with the waveforms of their voices and musical instruments.

Later School Years

The school placed great emphasis on exams which I found easy and often placed top of the class. This was of little interest to me because I was not a competitive person—I saw it largely because of having a good memory. It became a problem in the 5th year because the school based its end of year prizes on the exam results and I came top in all the subjects that interested me, adding woodwork, English language and English literature to my usual science prizes. The school introduced a new rule that no one could have more than 5 prizes but did not give me any choice in what ones to return, taking away the English ones which I valued most.

In the sixth form I specialized in mathematics, physics and chemistry, finding the applied mathematics syllabus boring as all the problems could be solved simply by Newton's Laws, such as force = mass x acceleration, and some trigonometry. The courses were nominally 2 years, but the school put me in for the pure math A-level after 1 and I passed with 98%. This had the advantage that the following year I had a free period during pure math and I was parked in the library where the headmaster was teaching a course on Plato to a small group of students. I was able to listen in and became very interested in Plato's ideas, an interest that has lasted for the rest of my life.

[4] https://www.worldradiohistory.com/Wireless_World_Magazine.htm
[5] https://en.wikipedia.org/wiki/Practical_Wireless

It appeared that the school library had been populated by donations from the masters and held a miscellanea of interesting books. I read Norbert Wiener's *I am a Mathematician*, which profoundly affected me and motivated me to become one also. There were also Littlewood's *A Mathematician's Miscellany* and Hardy's *A Mathematician's Apology* which further confirmed my choice. These also led to my view of mathematics as an abstract system independent of reality which later attracted me to cybernetics, system theory and Popper's notion of *three worlds*.

At the beginning of my final year in the upper sixth a very embarrassing event occurred that still seems absurd to me. Our family had to take our annual holiday at a time designated by my father's employer and often in the school term. I missed the last week of term and the final announcements. I met a friend during the holidays who told me it had been announced that I was the Head Boy for the coming year. The notion was so absurd I did not believe it but when I returned to school the headmaster took me to his office and told me I was now Head Boy. What was even more absurd was that I was automatically head of my house in which I had played no part. Another boy was appointed as deputy head—I felt very sorry for him doing all the work but carrying a lesser title. It introduced me to the incomprehensible stupidities of bureaucratic life. I vowed never to propagate them—whether I have succeeded is for others to judge.

I couldn't attend Cambridge immediately because they required me to have an O-level in Latin for entry. I had studied it during my initial year at the Grammar school and had greatly enjoyed it as it was a very precisely defined language but was transferred to German when I was streamed to the science curriculum. The Latin master accommodated me by offering to teach me for an hour on Saturday mornings, and the textbooks were excellent, so I was able to pass O-level Latin with a score of 98%.

My First Experience of Industrial Research

That left me weekdays free and I was able to obtain an intern position for 1959-60 at ITT/STC's[6] semiconductor labs at Foots Cray, a 10 min cycle ride from my home, after an intense interview with Bernie Mills, the Head of the Lab. On my first day at work, I was greeted by the other intern, Michael Grimwood, from Canterbury, where his father owned a musical instrument shop. Michael became a long-term friend and played a strong role in several of my later industrial activities. I also worked in his father's shop tuning and repairing instruments which I greatly enjoyed.

[6] The US conglomerate International Telephone and Telegraph and its UK subsidiary Standard Telephones and Cables

When I was first at the labs, the associated model production line was producing mesa transistors using thin slices of crystalline germanium doped with boron with a thin layer of phosphorus diffused into the surface. I was assigned to Dr John K Arthurs as Project Leader and given various tasks such as etching the surface of the slice and measuring the density of the etch pits using a 3D microscope to determine the quality of the crystal lattice.

I was also allocated to Dr Charlie Sandbank to help with production of tunnel diodes, and later with setting up the silicon planar process as we switched to the production of integrated circuits. These last presented new problems of quality control as the leakage current which had been microamps with the mesa process because of the exposed junction dropped to nano amps with the sealed junction of the planar process, and there was no commercial equipment capable of measuring such low currents so I took it upon myself to design such a device.

STC had a valve factory next door, and I was able to procure some triode valves that had been on life test for 15 years and whose grid current had dropped to a very low level. I was also able to purchase some 100M Ohm high voltage glass-encapsulated resistors from a government surplus store to go into the grid circuit and a milliammeter to go into the cathode circuit, and we had a nano-ammeter. It was very sensitive to stray currents from static—the operator had to wear an antistatic hairnet. I was very surprised when visiting STC's 74N series integrated circuit production line in the same building some five years later to see my nano-ammeter still in use as a quality control tool!

One of the major advantages of the job was that a number of professional journals were circulated around the laboratory every week, from the McGraw-Hill trade journal, *Electronics*[7], through IEEE Transactions, to physics journals that published extensively on semiconductor physics such as *Physical Review* and *Journal of Applied Physics*. This enabled me to gain an extensive education in the physics and mathematics of semiconductor devices, and of their burgeoning applications.

I particularly enjoyed studying the mathematical foundations of the work we were doing in the lab such as the quantum mechanical theories underlying the behavior of the tunnel diode. In a journal, I found the theory of the four-probe jigs that were used to measure the conductivity of the germanium slices and noticed that it was based on rectangular shaped slices and recomputed it for the circular slices that we were measuring. There was a difference in the computation but not enough to affect our use of it, which was reassuring. When it was found that I was comfortable with mathematics I was given more and more mathematical tasks to do.

[7] https://en.wikipedia.org/wiki/Electronics_(magazine)

I purchased a set of electronics stencils and Indian ink pens and taught myself how to draw professional circuit diagrams, with the ambition of contributing to some of the journals I read—my first paper was for *Practical Wireless* and another early one to *Electronics*. While I still saw myself as a mathematician my immediate interest was its application to electronics.

University Years

Going up to Trinity College, Cambridge was the same kind of culture shock as the grammar school. I was surrounded by upper crust students who had been groomed for the role and with whom I had nothing in common. Michael had been accepted at Queen Mary College and I would go down to London at weekends where I was more comfortable and stay at his apartment. We formed a company called Semiconductor Associates and made electronic musical instruments and amplifiers for sale to our friends.

Richard Gregory

The first thing I did in Cambridge was to look for somewhere where I would have access to professional quality electronics components and equipment. I went to the Clubs Fair and told the Chairman of the wireless society, Steve Salter, I had worked in ITT semiconductor labs and had brought with me a number of interesting new devices I had made such as transistors and tunnel diodes and asked him if he knew of anyone who would be interested in experimenting with them. It turned out that he was Richard Gregory's technical assistant, and he suggested that Richard might be interested.

Steve had been an apprentice at Saunders-Roe and was highly skilled at instrument making but knew less of electronics. I became electronics technician for him and Richard for the next seven years and started out with a generous budget to equip an electronics laboratory which became my real home at Cambridge. Richard and Steve became close personal friends, and I spent many hours with Richard discussing the possible neural mechanisms for the illusion phenomena he was studying. I also made many friends in the psychology department when it was discovered that I knew mathematics. I became an advisor on statistical calculations, largely advocating nonparametric techniques because the data generally had little claim to being normally distributed.

One of my first projects in Richard's lab was to build a 3D oscilloscope. I had been fascinated by the book Flatland about 2D creatures and wanted to explore how 3D creatures might perceive a 4D space. I set up oscillators to trace the image of a tesseract, or hypercube, in 3D space with potentiometers that allowed it to be rotated. It worked but was completely meaningless, linked cubes that moved in incomprehensible ways.

John Andreae, Standard Telecommunications Labs (STL)

In my second year another opportunity to undertake some interesting studies came up. ITT's European research lab, STL, was at Harlow only an hour by coach or train from Cambridge, and in 1961 I saw an advert in Nature from John Andreae, Project Leader for a new project on learning machines, one looking for a mathematician with knowledge of algebra and topology and other topics that intrigued me, as did the project. I applied for a part-time position, was accepted and commenced in Summer 1962, working on the STeLLA learning machine in university vacations. This led to many of my early papers and my own explorations of AI topics.

Something I welcomed at Cambridge was that there were no compulsory lectures which enabled me to undertake these other activities. Lecture courses were the responsibility of the university and attendance was not monitored. The College provided a Tutor and another student, Roger North, and I were assigned to Donald Laming, a Mathematical Psychologist, who allocated us a past exam paper each week, and then went over our answers with us the following week. I got on well with both Donald and Roger and it was a pleasant way of preparing for the exams. This pattern continued for 3 years with different tutors and partners until I took the Part II Mathematics Tripos examination and graduated Senior Optime, 2(i), in 1963.

I went to some University courses, not necessarily associated with mathematics. Second hand copies of textbooks were cheaply available through the Cambridge booksellers, and one could read up the necessary math for the exams, and also interesting math such as topology and universal algebra that was not taught until Part III of the tripos. I was particularly fascinated by projective geometry which turned geometry into algebra, a much better direction, I felt, than Descartes turning it into arithmetic. Two courses I vividly remember were Paul Roubiczek's critical appraisal of existentialism, which provided insights into Zen Buddhism that had long fascinated me, and Fred Hoyle's on electromagnetics (Maxwell's equations). I had read Hoyle's science fiction books and followed his ongoing debate with Martin Ryle on whether the universe was in a steady state or expanding and was fascinated to find out how he would present Maxwell's equations. At first, I thought deriving Ohm's simple phenomenological law from them was a waste of time but then Hoyle demonstrated that one could also derive the failure of the law at radio frequencies when the current flowed differently. The notion of resistance as an intrinsic property of a conductor was flawed—it was situational—the beauty of what Maxwell had achieved was a universal theory that applied to direct and alternating current. Hoyle illustrated this again when in one of his lectures he used it to explain magnetic bottles, a phenomenon only recently discovered in space.

Much later when I read the history of Faraday's meticulous experiments, of Maxwell's derivation of his equations from a fictional structure of vortices, and then Heinrich Hertz's derivation and demonstration of the propagation of radio waves in free space I became very interested in how the human mind constructed such theories.

This interest was further aroused by some visiting lecturers. Cambridge attracted many distinguished visitors who would often give guest lectures, and I attended many of these. Two that had a major effect on my thinking were Murray Gell-Mann and Richard Feynman expounding and musing on their theories of elementary particles. I was strongly positivist and empiricist and felt that theories should be grounded in data and mathematically sound. I went away from their lectures saying these guys are inventing this as they go along from intuition and a desire for elegance. But they were great scientists that I highly respected and I realized that science was as much a matter of human psychology as of hard data and mathematical rigor.

I was also working on machine learning at STL and had read Norbert Wiener's work on cybernetics as the study of control and communication in the animal and the machine, and Ross Ashby's work on the design of a brain that showed that phenomena of animal intelligence could be emulated by simple mathematical systems, so my stance was increasingly one of treating cognitive processes in animals and machines from a systemic perspective, drawing on both psychological and cybernetic studies.

Doctoral Studies in Psychology with Richard Gregory

When I graduated in mathematics in 1963 my plan was to take Part III of the Math Tripos since that encompassed mathematics that greatly interested me and was highly relevant to my research at STL where I expected to make my career. However, Richard Gregory and Peter Donaldson in the Department of Physiology had been offered a contract to take over a study at the Institute of Aviation Medicine on *adaptive training* of perceptual-motor skills, and Richard suggested I undertake the study as his doctoral student. The contract was well funded and would enable me to build an analogue computer and study human and machine learning while remaining in Richard's Lab, continuing my liaison with John Andreae at STL and attending whatever Math Tripos III courses I wished.

I accepted, but when Richard went to Oliver Zangwill, the Professor and Head of Department, he made my entry contingent on passing the Psychology Tripos II. That would give me a second degree, professional status in psychology, and not change my plans very much. Richard asked Alan Watson, an animal behaviourist to be my tutor for Psychology. Alan would set me an essay each week on a topic relevant to the Tripos, and I would research it from the books and journals in the Department library. Alan was fascinated by

my positivist approach (I was deeply into Carnap and Wittgenstein by then), we had a lot of interesting discussions, I learnt a lot of psychology in a short time and graduated with a 2(i) which qualified me to enter the psychology doctoral program.

First Publications

I started writing papers in 1964 for publication and submission to conferences. The first was for *Practical Wireless* as a payback for all the information I had found there over the years (*Photoconductors: The Cadmium Sulphide photo-conductive cell in practice*, 1964) detailing a large number of applications of a photocell that had become available on the government surplus market. The next was on the STL learning machine research (*Some experiments on machine learning in various environments*, 1964) which came about when I noticed the announcement of the *British Computer Society Conference on Artificial Intelligence and the Simulation of Behaviour*, the first UK conference on Artificial Intelligence (AI)[8].

I said to John we should be represented there and he said organize it, so I contacted the organizer and submitted an abstract which was accepted. I presented the paper—it later was extended and published at an IFAC[9] conference in Paris (*A learning machine in the context of the general control problem*, 1966).

Sequence Prediction Research

An interesting event occurred after I presented the paper. People had said that a problem was too easy to test STeLLA so John had Michael Grimwood make a hand-held problem box that let a person try the problem themself. It had 3 buttons under 3 lights. When you pressed a button, a light came on, and you had to learn to press the button under the light that would come on (*Workshop Manual for the Game-Playing Machine SToooP,* 1965). I demonstrated the box as part of my presentation, and Jim Doran, Donald Michie's grad student working on his graph traverser program came up to me afterwards and asked to apply the algorithm to the box. Jim figured that because the box only had 8 states the problem would be a simple one for the algorithm, but he spent an hour covering pages with working and could not solve it. What he, and other people who had tried, did not realize was that the box had a small random element in its behaviour and could not be modelled as a finite state automaton. John required all the tasks set for STeLLA have such an element because the real world is not determinate and a learning machine that could not cope with some indeterminacy was useless.

[8] The term was coined by John McCarthy in 1956 in the runup to the Dartmouth Summer Research Project on Artificial Intelligence https://en.wikipedia.org/wiki/Dartmouth_workshop
[9] International Federation for Automatic Control

That phenomenon set me off on a long line of research on how people coped with indeterminacy and whether there was an algorithm that could do better. At certain points in this narrative, it becomes appropriate to group the papers by technologies rather than eras and this is the case for sequence prediction. The earliest record I could find of that research is an internal report on research conducted by Ian Witten and me at Essex University and John Andreae and John Cleary at the University of Canterbury, Christchurch, New Zealand collaborating to study the problem (*Inference from sequential data*, 1974). Initially we were all taking a heuristic approach based on analysing dependencies between the next token in a sequence and different length preceding sequences—at which neural nets later succeeded but at great cost and dependence on advances in computer speed, memory, and parallelism.

I decided to focus on finding a mathematically sound procedure even though I knew it would be computationally intractable. This gave great insights into the problem and its underlying issues but little help in suggesting a useful approximate technique (*Approximate identification of automata*, 1975, *Behaviour/structure transformations under uncertainty,* 1976, *System identification, approximation and complexity*, 1977). I extended the procedure to other types of automata, fuzzy and multi-stack (*Sequential fuzzy system identification*, 1977). However, I realized that there is no universal solution, a 2-stack stochastic automaton could not be modelled by a model for a 1-stack one, and so on. There were no useful consequences except the realization of the increase in intractability and the impossibility of a general-purpose learning machine. I stopped further research. Maybe neural net research will eventually lead to a more fruitful conceptualization of the problem once the basis of the success of neural nets is better understood.

Returning to 1965, I also noticed an Institution of Mechanical Engineers conference on *Advances in Automatic Control* and submitted a paper that presented the STL research from a control engineering perspective (*Nonlinear simplifications of the general control problem*, 1965). We had presented the STL research on learning machines to the UK AI and control engineering communities. John also felt the need for a non-specialist presentation to visitors to the lab and we wrote (*An introductory description of the STeLLA learning machine*, 1966).

Back at Cambridge I submitted my first contract report to the Ministry of Defence which set out my methodological assumptions and background knowledge in undertaking the research (*Automated feedback trainers for perceptual-motor skills*, 1965). I built a 10-amplifier analogue computer that could simulate a wide range of control dynamics and a finite state automaton to simulate a difficult user interface to create scenarios that required good training aids. I also built an adaptive trainer that measured task performance

and adjusted task difficulty to keep performance constant. I gave a seminar at the Department Tea Club on some peculiar phenomena I had noticed in tracking tasks (*Discontinuities in tracking strategies*, 1965). I later expanded this into a detailed critique of the misleading outcomes of modelling discontinuous human movement as linear (*Linear and nonlinear models of the human controller*, 1969) which I later generalized to the issues of fitting an inappropriate class of models to any phenomena.

Alexander Luria & John Gedye

Alexander Luria was a friend of Zangwill and a frequent visitor to the Department. I went to his seminars and was impressed by his demonstrations of the influence of verbal input on perceptual motor performance. I had frequent discussions with him on the topic and made investigation of the impact on pilot performance of verbal instructions part of my experimental design. I also programmed a simulation of my task controlled by a simple learning machine, a Rosenblatt *Perceptron*, to determine whether adaptive training was effective with AI systems.

I had befriended John Gedye, a medical doctor and psychologist who had joined the department as a lecturer in medical psychology and had previously been at the Institute of Aviation Medicine. He was instrumental in arranging subjects for my experiment from among the pilots at RAF Oakington near Cambridge, a relevant constituency for the research. They came to my office in the psychology department in full uniform on time and were extremely cooperative and interested in the research. I was able to gather the data I needed very rapidly and move on to analysis.

Stochastic Computing

In June 1965 when I visited STL I found that John had commenced an all-digital reconstruction of the STeLLA learning machine, and Peter Joyce his electronics technician had developed a digital circuit approximating the analog running averager that was central to STeLLA's learning (*Adaptive element*, 1966). Some years before I had been discussing with Richard how the eyes derived depth information from the neural input from the corresponding hemifields of the two eyes, had suggested that some form of spatial correlation might be used, and had speculated that the multiplications might easily be implemented if luminance was represented by the random pulse rate of a neuron in that a coincidence detector between two neurons would multiply their pulse rates. I simulated this and found that it worked, added additional operations, and called it a *stochastic computer*. I saw that the same technique could be used to implement a digital running average with fewer components and greater accuracy and sketched it out for John and Peter.

John asked me to present this to the Laboratory and I wrote some notes as a handout presenting it as a digital emulator of an analog computer (*A Stochastic Computer: Some notes on the application of digital circuits to the operations of arithmetic and differential calculus by means of a probabilistic representation of quantities*, 1965). The presentation was well-received, Peter commenced a full implementation, and I started developing claims for a patent (*Stochastic Computing Arrangement*, 1966).

Publications on Adaptive Training

I submitted my first paper on my adaptive training research to a conference on Aspects of Educational Technology (*Teaching machines for perceptual-motor skills*, 1966). The British Chamber of Commerce had offered the Czech government the choice of any paper from the conference to be presented in their country, and they chose mine (*Vyučovací stroje pro operátory*, 1966). My second paper was presented at an IEE colloquium on Human Operator Studies (*The evaluation and control of human skil*l, 1966).

Visits to US High Technology Centres

Back at STL John Andreae had left to take up an appointment at the University of Canterbury in New Zealand and I was asked to become project leader on a part-time basis. Peter White joined as Manager of the Advanced Projects Division and came to see me because, as part of the process of assessing productivity in his division, he had noted that I was highly ranked in terms of patent production but was on the books as an intern. He said he thought I had been underpaid and offered me either £500 in books or a trip to the research labs of high technology companies in the US. I chose the latter because that had long been my ambition. The STL Travel Department routed me to Xerox, IBM Yorktown Heights, Bell Labs Murray Hill, and, at my request, various universities.

I also took the opportunity to present the first published paper on stochastic computing at the Spring Joint Computer Conference in Atlantic City (*Stochastic computing*, 1967). It was well received with Walter Karplus starting the comments by saying that he had never expected to see any more innovations in analog computing, and Granino Korn being equally complimentary, inviting me back to Tucson to see his work on hybrid computers, and writing an excellent review for IEEE Transaction on Electronic Computers (appended to paper). Again, it is appropriate to group the stochastic computing papers by technologies rather than eras. I published a paper in the widely read magazine *Electronics* (*Stochastic computer thrives on noise*, 1967) which specialized in drawing attention to new technologies, and another (*Techniques of identification with the stochastic computer,* 1967) at an IFAC conference on system identification emphasizing steepest descent techniques and Bayesian modelling. By this time I had realized that simulating stochastic

computers with existing digital technology was not cost-effective and applications depended on the emergence of new technologies that were naturally stochastic, so I wrote a long paper for *Advances in Computers* (*Stochastic computing systems*, 1969) that captured all I had done and I ceased research on the topic, with later papers being requested by the organisers of particular events (*Stochastic computers*, 1969, *Trends in stochastic computing*, 1976, *Uncertainty as a foundation of computational power in neural networks*, 1987).

There was still great interest in STL in how we might apply the new technology, particularly in digital navaids. However, in each case I could see a more effective way of doing it. In many cases it was a deterministic variant of the stochastic computer in which two types of pulse train were used, each with the same computational cycle, with one having the pulses grouped together and the other as evenly spaced as possible. Multiplication only required an AND gate, etc, and I called such a collection a *phase computer* and published several papers on it (*Phase computers*, 1967, *A modular programmed DDA for real-time computation*, 1968). It was too close to past pulse rate computers to patent as a whole but we were able to initiate patents for particular applications, such as a digital Decca Navigator. I was asked to give the keynote address at *the IFAC Symposium on Pulse Rate and Pulse Number Signals in Automatic Control* and chose to compare the merits of general-purpose computers and special-purpose ones (*Varieties of computer - their applications and inter-relationships*, 1968).

Mid 1967 marked the completion of the contract of my study for the Ministry of Defence (*Automated Feedback Trainers for Perceptual-Motor Skills: Final Report to MoD*, 1967), my dissertation examination and approval with minor corrections (*The Human Adaptive Controller*, 1971) which took some time because it had to be re-typed on a manual typewriter which was tedious and low priority! I published a number of papers on the research, some on the experiments (*Training the human adaptive controller*, 1968, *The control of human learning*, 1968, *Adaptively controlled instruction for a tracking skill*, 1969), others on the underlying theory (*Adaptive control theory: the structural and behavioural properties of adaptive controllers*, 1969, *The learning of perceptual-motor skills by men and machines and its relationship to training*, 1972, *Axioms for adaptive behaviour*, 1972, *Training, stability and control*, 1974).

Choice of Initial Career

It was also time to develop my career, and this went in a different direction from what I expected. I had been offered the chance to manage a new division of STL focused on computer research and development to support ITT companies that already had existing computer products or planned to become involved in their design or use and had already

established collaborative links with the new Electrical Engineering Science (EES) Department of the University of Essex at Colchester only an hour drive from Harlow. However, just before I was due to take the job, Peter White was appointed head of a new flight simulator division of STC and told me I was coming with him as Chief Engineer of that division. I expect he thought I would be delighted as it was a more senior and highly paid position, but I was horrified as I saw it more as a management than a research position, and all that really attracted me was research.

I visited Professor Barrie Chaplin, the Head of EES at Essex University, and explained the change and my distaste for it. He asked whether I would be interested in reversing roles, and joining his department as a faculty member while being encouraged to establish links with industry. I said yes, and there ensued a bidding war during which ITT doubled its initial offer and Essex said it could not quite match that as it was over the top of the Lecturer scale but would write me a contract encouraging me to take on consulting work. I accepted that offer and ITT promptly offered me a consultancy at almost their original full-time salary offer. So, I continued with my research and project leader role at STL, making my very bright grad student, Ray Shemer, Project Leader of the phase computer design, which he was undertaking for his University of London industry-based doctorate. Thus, I commenced an academic career rather than an industrial one, and never regretted it, although I occasionally transferred back to industry while retaining strong academic links.

Department of Electrical Engineering Science

The Department when I joined it comprised Professor Chaplin who had been a computer hardware researcher at Manchester and Chief Scientist at Plessey, the Departmental Secretary, Janet Noyes, who was a superb administrator, Guy Delafontaine, who was a very knowledgeable electronic engineer, John Sparkes, who was Reader and a superb teacher and designed the undergrad labs, and John Turner from Manchester whose research was on cathode ray tubes. For its size it was very well balanced and the core of the Department for many years as it grew rapidly from its initial 30 students.

International Journal of Man-Machine Studies

I was funded for an undergrad lab with 3 PDP8s and a research lab with a well-configured PDP9, later replaced with a PDP11. I was interested in research and development in the newly emerging technology of interactive computing, albeit with ASR33 teleprinters as the user interfaces! The University Computer was a PDP10 and a major asset for interactive computing research. I formed the Man-Machine System Laboratory and recruited several graduate students to undertake research there. John Gedye and I persuaded Academic Press to start a new journal, *The International Journal of Man-Machine Studies* (*IJMMS*

Editorial and Board - first issue, 1969), to provide the first forum for the research, commencing in January 1969. John Gedye joined me in 1968 because Barrie wanted to develop medical electronics research and we continued our very fruitful research partnership, but he left a few years later to take up a position at Baylor College in the US. I continued to edit IJMMS for 37 years, growing it to be highly respected and the largest journal in the field, eventually renaming it as *The International Journal of Human-Computer Studies*.

BASYS, QUASIC, MINSYS, AIMS-11 Interactive Languages

Our priority was to set up a man-machine systems laboratory for experiments on human-computer interaction. Peter Facey, a grad student of mine who was an excellent systems programmer, wrote a round robin scheduler for the PDP9 and I developed a programming language called BASYS (*BASYS - a language for processing interaction*, 1977), similar to Dartmouth Basic, that ran under it. A major extension in the language was string handling based on the Brooker-Morris compiler-compiler which made it simple to decompose natural language and recompose it in the style of Eliza. I ported the language to the PDP8 (*QUASIC Manual*, 1971) and MINIC (*MINSYS Manual*, 1974) and Peter to the PDP11 (*AIMS-11 Manual*, 1977). The different names reflected different companies as primary users. The versions for DEC computers took advantage of their dynamic linking loaders to allow segments of codes in other languages to be pre-compiled and loaded as run time so that speed could be achieved as needed (*A mixed-code approach to commercial microcomputer applications*, 1978).

This software technology enabled us to take low-cost machines designed for real time industrial instrumentation and control and repurpose them for interactive business, financial and educational applications. It left us with many 'firsts', first successful hospital administrative system, first successful stock exchange brokerage system, first minicomputer to replace an IBM mainframe at a bank. To us it was an opportunity to take interactive computing into the real world and report the lessons learned (*An interactive display-based system for gilt-edged security broking*, 1974, *Minicomputers in security dealing*, 1976, *Some experience in interactive system development and application*, 1975).

Questel and Automated Psychological Testing

John Gedye was interested in automating psychological testing and I designed a terminal for him that appeared to the computer as an ASR33 teleprinter but controlled a random-access microfilm projector that moved very rapidly to any frame specified by the computer (*Medical applications of programmable audio-visual displays*, 1967). John saw the potential for the system to be used to manage the clinical trials of L-dopa, a new drug for

treating elderly patients. He worked with a company in the City that raised money for startups, and they raised £400K for Questel, a company to conduct drug trials with John designing the tests and Peter Facey and I providing the software for a TSS8 in the basement of Questel's offices. We ported BASYS (as Quasic) to the TSS8 and I wrote the software to administer and manage the clinical trials.

One interesting side-effect of Questel was that the company funded a newsletter, *Behavioural Technology,* that John Gedye edited. He asked me to write a regular column of a lighter but thought-provoking nature, and I did so for a while under the pseudonym, #. Some of the articles were quite prescient of what was to come in computing. For example, the first one was written while we were developing an early time-sharing system and the terminals tended to become detached from their owners so messages could appear on the wrong terminal. I imagined a world in which the bug had become a feature and users solved problems by passing them on to others on the network, *"who better to answer a question on abstruse Chinese history than an abstruse Chinese historian"*—shades of the WWW and Reddit! (*Through a teleprinter darkly*, 1970). In another I envisioned the power of computers as statistical tools that could increase the significance of results by various dubious techniques which one suspects have been rediscovered many times! (*Interaction and significance*, 1970). I terminated the articles with a final apology to those who had not realized that they had been written by a computer, giving some technical indication of how it was done and noting that next time they would not find out—ChatGPT your time has come! (*Hash under a bushel*, 1971).

Laurie Milbank and Stock Exchange Computing

Unfortunately, Questel was a failure and the funding consortium wanted to close it down but the manager, in order to raise revenue, had taken a contract with a stockbroker, Laurie Milbank, to supply the software for a computer system they had already purchased for their government bond (gilt edged) brokerage. They did not want to let down a fellow city firm and asked whether I would take over the contract. I went to see Major Everett, the Senior Partner at the brokerage, and he asked if it was possible to do so in three months. I said I didn't know but that if I couldn't do it then no one else could—Peter Facey and I did so using BASYS/AIMS-11. Peter became Technical Director of Arbat Consultants who specialized in banking systems and installed PDP11s and AIMS-11 (the PDP11 version of BASYS) in many banks.

Southend Hospital Computing System

Another medical project that John Gedye initiated was an administrative system for Southend Hospital in-patient management. The hospital catered for a population of

300,000, had some 1,200 beds, and averaged about 200 admissions and discharges daily. Again, we used a PDP11/ AIMS-11 combination. This project was one of the only two out of eighty hospital computing initiatives funded by the Ministry of Health that was rated as successful by the Auditor General of the British Government.

As a result, I was appointed a member of the Advisory Committee to the Chief Scientist of the Department of Health and Social Security and Chief Scientist's Representative on the Computer Research & Development Committee of the Department of Health and Social Security. One unwelcome duty was that I had to assess any new application to the Department for funding a computer. I say 'unwelcome' because one reason for the lack of cost effectiveness of the projects was the government wanted to support the British Computer Industry and had imposed a 'buy British' constraint on most projects, but most British computer manufacturing was focused on batch processing mainframes. It was not a fair comparison. I was pleased to have a good excuse to resign from these positions a year or so later.

A detailed account of these interactive systems was provided in a long paper of the *Proceedings of the IEEE on interactive systems* (*Some experience in interactive system development and application*, 1975). This was also the first publication in which we consolidated our *lessons learned* into 'dialogue programming rules.' These were developed over the years (*Programming interactive dialogue*, 1978, *The technology of interaction--dialogue programming rules*, 1981, *Dialog engineering*, 1983, *Dialog shell design*, 1984), including into a book (*The Art of Computer Conversation*, 1984). The recommendations were couched in abstract terms to continue to apply as interface technologies changed but needed reviewing occasionally (*From timesharing to the sixth generation: the development of human-computer interaction. Part I*, 1986, *Foundations of dialog engineering: the development of human-computer interaction. Part II*, 1986). The customer-demand for the capability to undertake the same tasks on a mobile phone as a laptop or desktop is the greatest challenge to achieving good human factors at the computer interface nowadays.

MINIC Microcomputer, Microcomputer Systems Ltd

On the hardware front, shortly after I started at Essex University, the Air Force provided me with £18K funding to continue my adaptive training research and I went ahead with my plans to convert my analogue computer to a hybrid one as I had seen in Granino Korn's Lab, with a small digital computer controlling the experiments and data logging and the analog one providing the experimental environment. I had already designed a very low cost instrumentation computer, MINIC, that was micro programmable so that one could extend the instruction set easily for special-purpose applications. It had an 8-bit word size and 8 and 16 bit instructions (*MINIC I Manual*, 1969, *MINIC - an 8-bit microprogrammed*

minicomputer, 1971). Plessey manufactured a low-cost 1KB core store, and a Sperry 64KB drum so a complete system could be built for about £1,000.

I was surprised one day when my liaison in the Air Force phoned to say that they no longer wanted any reports from me but would continue to fund the contract. There was a new head of the Air Force who was shocked to hear they were funding research on improving flight simulators and said, "our boys came to fly real aeroplanes, I want to hear no more about it."

So when Dave Seale and Tony Maine visited to sell us a Sperry Drum, asked about the status of MINIC and said they would be interested to form a company to manufacture it, I was free to give them the rights in exchange for some shares. George Kent, the instrumentation company put up the initial seed funding for the company, *Microcomputer Systems*, in exchange for a buyout clause after 5 years. MINIC sold for £2,000 and found a market niche with Alfred Herbert, the numerically controlled machine tool company and hit the press because it enabled them to sell an NC controller at half the cost of their rival, Plessey. It also found many other uses, such as the first critical care monitoring system for 6 beds at University College Hospital which I programmed with 1KB memory plus 64KB drum—which seems a miracle nowadays! Peter Facey programmed an emulator on the PDP10 which made such developments possible (*Some experience in interactive system development and application*, 1975).

MINIC itself was highlighted in the press as the only successful British computer company that received no government funding. As a result, we were given a contract to develop a new more powerful computer which I called the Minic S. I targeted this on the market for ultra reliable, ultra secure computers with a mean time between failures of 40 years or more (Design objectives for a descriptor-organized minicomputer, 1974, Integration of protection and procedures in a high-level minicomputer, 1974, A high-level minicomputer, 1974). Minic S was taken to production prototype stage with a complete suite of software, operating system, assembler debugger and compilers, and was about to go into production when George Kent was taken over by Brown Boveri & Cie who cancelled the project because they were already committed to the PDP11. It was the last computer I designed.

Analogies and Virtual Machines

I became interested in formalizing the notion of analogy by an unusual route, not from studying its role in human reasoning or artificial intelligence but by seeing it as the natural way to formalize the notion of a 'virtual machine' in software engineering and computer architectures. I had studied the IBM360 family of computers and was fascinated how a whole family of computers with identical functionality but varying widely in performance

and price could be created by using software in the lower cost machines to emulate the hardware of the higher cost machines such that the former was a 'virtual machine' with the functionality of the latter.

I realized that I could regard the two machines as analogous where an analogy was a partial structural identity that 'forgot' some differences, such as performance and price, and focused on similarities like computational capabilities. I had this in mind when designing MINIC. It had primitive arithmetic capabilities but was designed to allow that to be extended very easily through subroutines, and the microcode at the lowest level ran on a completely different architecture. What was presented as the 'computer' was itself a virtual machine that was used to create other virtual machines.

It also struck me that this was the foundation of structured software architectures, that software 'libraries' created virtual machines for vector and matrix computations, database management, and so on. It also seemed that the variety of types of formal structures that could be represented in a virtual machine required the generality of a category-theoretic formulation. I had first come across category theory in a 1958 paper by Robert Rosen on the representation of biological systems and found its level of abstraction and generality very attractive. I together with Ian Witten, Ron Atkin, and some of our grad students had studied MacLane's *Categories for the Working Mathematician* together. Joe Goguen, a colleague and close friend had sent me his draft paper applying it to human concepts.

I saw the possibility of modeling the analogy relationship between virtual machines and programmers within a category theoretic framework (*Analogy categories, virtual machines, and structured programming*, 1975, *Analysing analogy*, 1982). I also wrote it specifically from a human-computer interaction perspective (*The role of the behavioural sciences in programming*, 1979).

My grad student, Debbie Leishman, decided for her Master's research to implement an analogy generation algorithm based on the theory and test it on examples in the literature (Analogy In knowledge acquisition: Description of an analogical tool, 1989, Analogy as a constrained partial correspondence over conceptual graphs, 1989). We also wrote up the general theory in simple algebraic framework (*A minimal framework for knowledge representation sufficient to support analogical reasoning and other forms of inductive inference*, 1989).

Despite this early interest in formalizing and implementing analogical reasoning, it played no part in our later research on knowledge acquisition. The interest had risen in an era when I was focused on designing computers and programming languages and associated libraries, and moving into human factors. In knowledge acquisition, analogy may be a

source of creative ideas from appropriating models of one domain to another, and we were not tackling that kind of problem.

It was not clear how useful it might be. The history of science is rife with spectacular failures of analogies: caloric, heat as fluid; ether as a transmission medium for electromagnetic waves; Newton's speculation that gravitation is analogous to electricity. Some analogies such as that between basic logic and basic mereology are mathematical isomorphisms, the same patterns of inference with different names, but others are misleading, and it is difficult to tell between them without an exhaustive investigation. Category theory itself may be seen as a theory of analogies between various mathematical domains that are logically well-founded.

Uncertain Reasoning, Probability Theory and Fuzzy Sets

An event in 1975 triggered off a line of research on uncertain reasoning that persisted for decades. Its origins are in 1965 in a proposal by Lotfi Zadeh, a control and system engineer whose research on notions such as *state* and *adaptivity* I greatly admired and used in my own research, that we needed a science of *cloudy* or *fuzzy* systems that could not be precisely specified. I agreed with the need but thought that his proposed solution in terms of sets with numeric degrees of membership introduced a false appearance of precision and forgot about it. However, my own research on modeling the structure of non-deterministic systems from their behaviour had given me new insights into the variety of forms of uncertainty.

Meanwhile, Lotfi continued to develop the notion, particularly in relation to system specification in terms of graded linguistic hedges. Abe Mamdani, a friend and colleague from Queen Mary College decided to replicate my work on training, priming and coding learning machines in the more realistic engineering task of steam engine control. Sedrak Assilian, his graduate student, used linguistic hedges to prime the learning system and had it translate this into a control policy using Lotfi's theory. Much to everyone's surprise this generated a very effective control policy that only worsened through learning. I was the external examiner and invited those responsible to submit a paper to IJMMS for rapid publication. This aroused great interest and was seen by some as a vindication of Lotfi's theory of fuzzy sets. I saw it as more of a vindication of his suggestion of using linguistic hedges for system specification and wrote a paper showing that the same result could be derived using probability theory (*Stochastic and fuzzy logics*, 1975).

Then I published some papers on uncertain reasoning that addressed the general question of representing uncertainty and the role of probability theory and fuzzy sets, showing that neither was adequate and suggesting that the notion of eventuality also needed to be

addressed (*A calculus of possibility, eventuality and probability*, 1975, *The logic of automata*, 1975). I also learned from my Czech graduate student, Ladislav Kohout, that Polish and Czech logicians had already developed multivalued logics and notions of 'linguistic' that encompassed Lotfi's fuzzy logic and linguistic approximation (*Precision and Reasonableness in Linguistic Communication: A Fuzzy Model*, 1975, *Multivalued logics and fuzzy reasoning*, 1975, *Possible automata*, 1975, *Foundations of fuzzy reasoning*, 1976, *Fuzzy reasoning and the logics of uncertainty*, 1976).

In 1978 I published a paper that as far as I was concerned wrapped up the topic, showing that probability theory and fuzzy logic could both be modelled as valuations on a modular lattice of propositions, with fuzzy logic deriving from a requirement of truth-functionality, and probability theory from the requirement of the excluded middle (*Fuzzy and probability uncertainty logics*, 1978). This was also the year I published the culmination of my studies of a mathematically sound derivation of the structure of a stochastic automaton from its behavior (*General system identification - fundamentals and results*, 1978), an exposition of the role of the computer as a vehicle to travel in Popper's world 3 (*Computers in world three*, 1978), and forecasts of the development of natural language communication and human-comparable intelligence (*Man-computer communication - what next?,* 1978).

Professor of Computer Engineering and Head of Department

In 1975 I was appointed Professor of Computer Engineering and Head of Department. The Department had grown rapidly and yet everything still reported to the Head leaving some very talented members under-utilised and not developing the skills of leadership essential to its future. I appointed 4 carefully selected members to be deputy heads, one for each year of the program. That worked out very well and by the time I was in my third year I used to joke there was no job left for me. After my time as head, I took a sabbatical year and was concerned about my own future. Essex had been through some very troubled times with many disruptive student strikes, and the British government had introduced three new pro vice chancellor positions to strengthen the leadership team. It was rumoured I was in line for one whose portfolio included student discipline and labour relations, probably because I was the first person to break a student strike (by persuading our engineering students to attend student union meetings and vote them down) and by being good on the Trade Union negotiation meetings (by being facile in mental arithmetic and seeing through tactics for gradually inching up score counts enough to re-rank jobs, merited or not). However, I was only 37 and wanted to pursue research for many more years, but there were no attractive research positions emerging elsewhere.

Monotype & the Lasercomp High-Resolution Laser Phototypesetter

One did come up from an unexpected source. Peter White, my Manager at STL, had been head-hunted to be Managing Director of Linotype, a British printing equipment company, for 5 years and, after that time, approached me to ask if I knew of any interesting business opportunities. I had told him of what I was doing with Laurie Milbank and suggested there was scope for a company to sell such systems in collaboration with them but not one that interested me. Peter approached them and set up City Computer Systems, asking me to be Technical Director and Consultant part time.

Then Peter was asked by the National Enterprise Board to consult on a proposal from the Monotype company, another British printing equipment manufacturer, for funding to sustain the development of a new laser typesetter after its finances had been stripped by a property developer when Monotype had been depreciating its properties in London at a time when they were going up strongly in value. Peter asked me to join him in the assessment of the value of the new technology which was heavily computer-based, and together we put together a business plan for revitalized Monotype.

After a while the NEB came back to us with the proposal that they would put in the funding provided Peter came in as Managing Director with me as Technical Director to carry out the plan. It was an intriguing proposal given my circumstances, a short-term appointment at four times my university salary with a company car and plenty of opportunity to travel since Monotype had 40 subsidiaries around the world. I was also intrigued by the new possibilities offered by the Lasercomp and by the computer support that it would need. We accepted subject to Peter being Chairman and Chief Executive and me being Deputy Chief Executive and Technical Director.

Mildred and Myself

There were also developments in my personal life. Mildred Shaw had submitted a paper to IJMMS in 1977 on her doctoral research on eliciting and analysing Kelly's conceptual grids through conversational computing techniques. I was very impressed by her work and let her know that when I accepted her paper. Six months later I was invited to be her doctoral examiner and met her for the first time in June 1978 for the exam at Brunel University. We met again at Gordon Pask's Decision conference in Richmond and in January at my first SGSR Conference as President in Houston. It was there that I noticed her as a person and started meeting with her socially when we returned to England. In July 1979 we moved in together in my apartment in the City and have been together until her death in December 2025.

We also began publishing together in areas of mutual interest and I became increasingly attracted to Kelly's constructivist epistemology, which I did not find inconsistent with positivism. Mildred's computer implementation of Kelly's elicitation of personal construct systems as conceptual grids could be seen as a scientific instrument for modelling the 'subjective' (*Decision: foundation and practice II*, 1979, Externalizing the personal world: computer aids to epistemology, 1979). We also began researching new techniques for analyzing such grids (*New directions in the analysis and interactive elicitation of personal construct systems,* 1980, *Fuzzy semantics for personal construing*, 1980, *Fuzzy entailment analysis*, 1980). Our later engagement with the *expert system* development community and later with the *neural nets* development community stems from being prepared to value knowledge generated experientially not only through rigorous scientific procedures.

Systems Research

I have already mentioned the early influence of the *cyberneticians* like Wiener and Ashby on my thinking and recognized the similarity of *system theorists* such as Lotfi Zadeh and George Klir whose works had influenced my doctoral studies and behaviour/structure modeling. I realized that a system-theoretic-approach was common to much of my research such as my adoption of Popper's classification of phenomena into three abstract 'worlds' to clarify the role of the computer in the world of ideas, my reduction of the inductive problem of sequence prediction to a deductive one of determining the Galois connection[10] over a family of models between complexity of model and goodness of fit (relative to the definitions of these measures). There was a beauty and a generality in these simple solutions to what appeared to be complex problems, ones that solved far more problems than one had originally considered. I recognized the phenomena in Euclid's geometry and in Newton's laws of motion. A systemic perspective is a powerful problem-solving tool.

So, when in the middle of the Monotype negotiations, George Klir phoned me from a Society for General Systems Research (SGSR) meeting in the USA and invited me to become its next President, I was ready to become more actively involved in the systems community. George was an old friend whose work I had drawn on but I was not a member of the Society and nor had I attended any of its meetings. George explained the society had been restructured and there was no acceptable member for President. I had been put forward as an acceptable neutral candidate. It was only after I accepted that I found that the incoming President was responsible for organising the next annual meeting—it was not entirely an honorary position!

[10] Order-reversing relation

George had started a new Journal, the *International Journal of General Systems* in 1974, and Ladislav and I had contributed to the early issues (*The logic of automata*, 1975, *Protection as a general systems problem*, 1976) and in January 1976 I had submitted the culmination of my sequence prediction studies (*System identification, approximation and complexity*, 1977). It was very comprehensive, covering the philosophical issues as well as the technicalities with 240 citations—George swore never to publish one as complex again. He invited me to give a keynote address at a major Systems Conference at Binghamton in 1977 and I noted that "it is the systems approach emphasizing lack of disciplinary boundaries, the freedom to apply knowledge and techniques gathered in one field to problems in another, or to suggest that two distinct fields are in fact one, the disciplined freedom of the unconstrained intellect—the approach rather than the subject area, that has been the source of dynamism and progress. Perhaps the most telling progress of all is that we can so confidently speak of a common field of interest knowing that we could not, and would not wish to, agree on a definition of what a system is." (*Progress in general systems research*, 1978).

In my President's address in the 1979 SGSR Yearbook, I propose what appears to be a vacuous definition, "**Systems are whatever we like to distinguish as systems**. Has anything been said? Is there any possible foundation here for a systems science? I want to answer both these questions affirmatively and show that this definition is full of content and rich in its interpretation.... to distinguish some entity as being a system is a necessary and sufficient criterion for its being a system, and this is uniquely true for systems." (*General systems research: quo vadis?,* 1980).

In a succession of papers to SGSR meetings we made these notions more explicit in terms of seminal systemic metatheories. The first paper models the systemic processes whereby communities having common interests form using Popper's notion of '3 worlds', Pask's of 'P-Individuals', Kelly's of 'personal scientists', Woolf's notion of 'surrender' in Gray's 'relevant, nurturing environment' (in World 2) to constructs (in World 3) of the community (in World 2) in order to 'catch' their essence (in World 3) which characterizes that community (*The personal scientist in the community of science*, 1981). The second paper analyses the system-forming in terms of Brown's Laws of Form, Popper's 3 worlds, and Klir's notion of a hierarchy of systems epistemological levels (A *programme for the development of a systems methodology of knowledge and action*, 1981).

In our final two papers at SGSR we enrich the systemic foundations with notions from Bertalanffy, Checkland, Maturana and Varela (*Hierarchies of distinctions as generators of system theories*, 1984), and, finally, present an example of how to use computer tools deriving from personal construct psychology to undertake system building in the way

expounded in the papers (*Some experience in eliciting knowledge structures for expert systems*, 1984).

This part of our story has now overlapped the next time segment in which we move to Canada, and then to Calgary, but it seemed sensible to keep the systems research saga together. It now segues directly to 1986 and the knowledge-acquisition for knowledge-based-system research, and the following section returns to 1979.

Chinese Laser Printing System

My next three years in Monotype were intensely busy and very interesting. I realized from the start that Monotype was losing its status in the print industry through the obsolescence of what it had to offer apart from the laser typesetter. It needed to put that new technology at the forefront of public awareness, and a perfect opportunity came up because China was now post-Mao and the western world was offering capital for it to industrialize. A delegation from China came to visit Monotype to discuss modernizing its print industry which was still using hot metal typesetting equipment supplied by Monotype 60 years ago. I borrowed an ex-graduate student of mine who was an excellent programmer from City Computer Systems to develop a video editor for Chinese. I set up a project in the typography department to repurpose a Monotype 400 key mathematics keyboard with new key caps for the Chinese keyboard designed by Professor S. C. Loh from the Chinese University of Hong Kong that used a keyboard of 256 characters representing the radicals that make up the Chinese characters. I went out to California and purchased Z80 microcomputers having S-100 buses that could provide 16 memory boards to hold the 16 x16 bitmap images for 16,000 Chinese characters. By the end of the year, we were able to take a complete Chinese typesetting system to Hong Kong to display to the Chinese (*Monotype cracks the Chinese puzzle*, 1979), even though we had few Chinese characters digitized for the laser typesetter yet—we proposed to China Printing Corporation (CPC) that they do the digitization and own the intellectual property rights.

We signed an MOU[11] with CPC for a 6-month trial in Beijing and Shanghai and sent 40 engineers over to install the system and train their Chinese colleagues. I went over for 6 weeks to ensure there were no problems and for discussions with senior CPC officials. At the end of the training period the Monotype engineers went back to the UK and the Chinese ones staged an exhibition for some 5,000 Chinese visitors from government and industry. CPC purchased the system and brought it into routine use (*Unravelling the Chinese typesetting puzzle*, 1981).

[11] Memorandum of understanding

I undertook various other projects for Monotype, notably: setting up a research laboratory in Trinity College Science Park for a small R&D team around David Hedgeland, the inventor of the laser typesetter; arranging for the digitised Monotype typeface library to be converted to a form suitable to those needing fonts for laser printers, notably IBM and Adobe; and writing a series of articles for the print industry on the impact of digital typography on what was now possible (The digital revolution II: the bounds of possibility, 1981, Some experience in text processing in the Chinese language, 1981, Opportunity or disaster?, 1982, Digital know-how, 1982).

When my time at Monotype was completed I had the same dilemma as before, I did not wish to take up another industrial or university appointment in the short term but accepted a number of consultancy projects that were offered to me, primarily in Japan and North America. I continued to fit in academic activities, and in May 1982 was consulting in Japan on my way to give the banquet address at NAFIP-1, the First North American Fuzzy Information Processing Group Workshop, at Utah State University. The President of the Japanese Corporation asked me to stay an extra day to go over a speech he had to give in English. As a result, I missed my flight from Salt Lake City to Logan and had to drive, arriving just in time to give the banquet speech (Precise past—fuzzy future, 1983). I went on to San Francisco for more consulting and then home, opened the door, said hi to Mildred and then fell over. We went to a doctor in the Barbican who listened to me, then turned his back on me, harangued me, saying I see lots of men like you, generally lying dead on the floor from a heart attack—go to Dillons and buy a handbook on family medicine and look after your body as well as you do your computer.

Move to Canada

Mildred said we should move nearer my work, and I said somewhere within an hour of Toronto as I was consulting with Moore Corporation there every 2 weeks. She took that as flying time and interviewed for a Computer Science position from Chicago through Canada to Maine, accepting an offer from York University as Associate Professor of Computer Science. I decided to take up an offer from Heward Grafftey, who had been Minister of Science in the Canadian Government, to fund a research company focusing on innovative technologies of interest to major Canadian companies, such as Moore Business Forms and handwriting recognition, if I would move to Canada to run it. I was also offered an adjunct Professorship in Industrial Engineering at the University of Toronto. We both were granted significant NSERC[12] funding. We moved in August 1982, bought a house in North York and embarked on our new life in North America.

[12] Natural Sciences and Engineering Research Council of Canada

John Boose and Knowledge Acquisition

Mildred and I had written a paper for a special issue of IJMMS on *Recent Advances in Personal Construct Technology* (*New directions in the analysis and interactive elicitation of personal construct systems*, 1980) in which we discussed the utility of her grid tools in eliciting knowledge for expert systems. John Boose, who had recently been hired to lead a team investigating the utility of expert systems in the aerospace industry, visited us in Toronto, discussed that with us and arranged for Boeing to purchase the source code of our programs. We became friends and collaborators and later initiated the series of world-wide Knowledge Acquisition Workshops that are still running 40 years later. Mildred and I also began a series of empirical studies of whether well-established knowledge structures, such as Carlsson's analysis of business record structures could be elicited from businesspeople with varying degrees of expertise (*A computer aid to knowledge engineering*, 1983).

Expert Systems and Manufacturing

My experience in a wide range of manufacturing industries led me to explore the utility of expert systems in manufacturing as part of my contribution to the Industrial Engineering Department. I received a 3-year Strategic Grant of $165K from NSERC for 'Management of new industrial technology' as *Principal Investigator*, joint with Harvey Kolodny, Neville Moray and Burhan Turksen (*Expert systems and simulation in the design of an FMS advisory system*, 1986, *Knowledge engineering for an FMS advisory system*, 1986).

Mildred as Associate Dean of Education, Brian's Options

In 1984 Mildred was promoted to Full Professor of Computer Science at York University and Andy Efrat, the Dean of Education, offered her the half time position of Associate Dean of Education which she was pleased to accept. I was offered a Full Professorship of Computer Science at York to strengthen their application for an MSc in Computer Science, and made that my base for a year, but was also discussing becoming Head of the Department of Industrial Engineering at University of Toronto and becoming Dean of a new Faculty of Information Technology at George Mason University (GMU).

Japanese 5th and 6th Generation Plans

In 1982 MITI in Japan had published detailed national plans for what they termed a *fifth generation* of computers based on major advances in artificial intelligence technologies. I saw the significance of what they proposed, studied it in depth and published an extensive book chapter situating it in the current state of the art (*Perspectives on fifth generation computing*, 1984), as well as papers at major conferences (*A framework for the fifth*

generation, 1984, From ergonomics to the Fifth Generation: 30 years of human-computer interaction studies, 1984).

In Japan the program was criticized for being too hardware-oriented and a second proposal was prepared for the Science and Technology Agency and issued in March 1985 with a greater emphasis on artificial intelligence and knowledge science. I again analysed it in detail (*Sixth generation computing: a conspectus of the Japanese proposals*, 1986), and I very much liked the notion of 'knowledge science' for a program with even broader scope that studied the role of knowledge in society from an economic as well as a technological perspective, and drew up a plan for a *Knowledge Science Institute* (KSI) that addressed that need. It was initially directed at GMU, but a new opportunity arose that resulted in it being presented to the University of Calgary (UofC).

Killam Chair, University of Calgary, Knowledge Science Institute

The UofC had inherited an Izaak Walton Killam Research Chair when it split from the University of Alberta. The holder for 25 years had died in 1984 leaving it vacant, and Departments in the Faculties of Science, Engineering and Medicine were entitled to nominate a candidate. Two old friends from England, David Hill and Ian Witten, were now Profs in the Department of Computer Science and phoned to ask if I would like to be nominated. I said no thank you as Mildred and I knew that Calgary got the first snow in the winter and was very cold. They said you and Mildred can come out for a holiday trip, and we can see you and show you around. All you have to do is give a public lecture and attend the selection committee. I edited the Knowledge Science Institute proposal to apply to Calgary and Canada (*The Formation of a Knowledge Science Institute in Canada—The 1985 Proposal*) and presented that to the selection committee, saying that would be my main objective. We went back and forgot it. Some two months later we got a phone call from the Vice President Research who said that the selection committee wished to offer the Killam Chair to me. It was when he said "the duties are, well, there aren't any duties specified" that the post became attractive—I agreed to accept. So, in 1985 we relocated once again, Mildred took a year's sabbatical leave from York to give her a much-needed rest, and I commenced founding the Knowledge Science Institute.

In October I gave the Killam Chair inaugural address saluting the life and times of the Killam family and introducing Knowledge Science in Canada (*The Entrepreneur in the Digital Age: Killam Chair Inaugural Address*, 1985). Two colleagues came up to see me afterwards and said they would like to work with me. We became close friends and coworkers. Joan Vickers from Kinesiology and I pursued joint research into knowledge acquisition by elite athletes and into hypermedia. Doug Norrie from Mechanical Engineering and I pursued joint research into advanced manufacturing and knowledge management. They both offered me

adjunct appointments, but I only accepted mechanical engineering since it would have seemed foolish for a person who avoided physical education all their life to pretend to be a professor of kinesiology.

Joan and I wrote several papers on hypermedia and Mildred and my work on knowledge acquisition had a significant influence on Joan's book, *Instructional Design for Teaching Physical Activities: A Knowledge Structures Approach*. We also learned from her innovative usage of our knowledge acquisition tools for research paradigms we had not envisioned but the mathematical completeness supported; as often happens, a researcher found a valid use for a dual paradigm.

The Knowledge Acquisition Workshops Era

We discovered a beautiful conference centre, the Banff Centre, nestled in the Rocky Mountains at Banff, an hour's drive west of Calgary and suggested to John Boose that we organize a workshop there on knowledge acquisition for knowledge-based systems (KAW). John said he had contacts in AAAI who would provide $5,000 in funds to support student attendance if we acknowledged their sponsorship—which was a prestigious mark of approval anyway. We set the criteria for attendance high, a 20-page journal-quality paper which would be refereed by 7 people from a prestigious program committee. We set a 40-attendance limit to maintain a participatory workshop ethos, but had 500 applications to attend from 30 countries, 120 submissions, 60 accepted—which the wrap-up session decided was too many and limited to 45 in future. The organization was democratic with the attendees from one year setting the guidelines for the next—it had been intended that the workshop circulate in North America but every year they voted to come back to Banff!

The Europeans at the workshop decided they wanted their own KAW series (EKAW) and the Australians and Japanese decided to have a shared series (AKAW, JKAW, PKAW). John and I were invited to be co-chairs for the early ones, and attended them, but they soon developed their own informal organizations that ensured their continuity. All the workshops made their proceedings available to participants in advance facilitating brief presentations and lengthy discussions. The Banff ones had them printed locally, sometimes as 3 volumes, the European ones moved to Springer series, the others varied.

Most of the Banff papers appeared in journals, often the *International Journal of Man-Machine Studies* (IJMMS) that I edited. In 1989 Academic Press agreed to publish a new journal, *Knowledge Acquisition* (KA), edited by myself and John Boose that took up that role. The volume of KA papers rose to be too great for its page allocation so in 1995 Elsevier agreed to merge KA with the *International Journal of Human-Computer Studies* (revised name for IJMMS) which solved the problem.

The KAW meetings allowed the KSI to play a supporting role for a community of significant researchers of human knowledge processes and their computer support and provided a forum for critical comment on our research. The need to address these fora on a regular basis was a stimulus to continue research when the going was tough and the hoped-for breakthroughs were not easy to come by. One was also calibrated by how well the results deriving from the KAW community were valued by colleagues when presented in other relevant communities of practice—which were many given the breadth of issues relevant to 'knowledge acquisition'.

It was also satisfying because the KAW community valued results as much as principles and I could build systems and extend representations to make them more human-comprehensible (e.g. rules with exceptions) and present them at the same forum as deep philosophical questions about whether the knowledge we elicited from experts *pre-existed* or was *created* in the elicitation process (or what that question means!).

The KAW community provided one important component of the KSI, giving it an important role in developing technologies to support the knowledge economy. I had already initiated research in another important component, to model and understand the knowledge economy. I gave a keynote address at Interact'84 (From ergonomics to the Fifth Generation: 30 years of human-computer interaction studies, 1984) and remarked on the empty 3rd ring in the Japanese diagram of 5th generation systems, that of 'human application systems': "Only when we know as much about the natural phenomena of people as we do about the artificial phenomena of hardware and software will human factors achieve its desired maturity. Perhaps the natural/artificial distinction will then have disappeared as we merge with our artefacts into true man-machine symbiosis. There is a Chinese curse, 'May you live in exciting times'—we surely do." That remark is particularly apposite today, over 40 years later!

BRETAM Model of Information Technology Infrastructure

As I sat listening to the paper after my keynote, it struck me that the underlying model of the exponential growth of information technology lay before me, that developments in electronics had enabled those in computers, which had enabled those in software, which had enabled those in computer interaction, ..., and each new layer had facilitated progress in the preceding layers as well as expedited that in succeeding ones. We had a hierarchy of positive feedback loops resulting in super-exponential growth beyond our control. The genie was out of the box and all we could do was to track it and try to predict the consequences. That became the KSI's second task, with the model being updated every few years and the predictions looking more and more like science fiction and yet turning out to be frighteningly realistic. In our era that fright has entered public consciousness.

My first publication of the tiered BRETAM (breakthrough, replication, empirical, theory, automation, maturity phases) model of the growth of the information technology infrastructure was in 1985 (*Human-computer interaction in fifth generation computing: problems and opportunities*, 1985), followed by a systemic model that Mildred and I developed in 1986 (*A learning model for forecasting the future of information technology,* 1986). Since then, there have been many updates tracking impacts on particular industries (*From teaching machines to knowledge-based systems: changing paradigms for CAL,* 1987), explaining convergence phenomena (*Convergence to the information highway,* 1996), and one which earned the best paper prize in the journal, *Technological Forecasting and Social Change* (*The learning curves underlying convergence*, 1998). The last was in 2019 and covers 88 years (*From facilitating interactivity to managing hyperconnectivity: 50 years of human–computer studies*, 2019).

I grouped the 12 layers into 3 epochs: *Computer Science*; *Knowledge Science*; *Symbiosis Science*. The first 2 have been my stomping grounds. The last is where we are now, where our constructs become our partners, maybe our superiors. They need to be involved in this analysis because it goes beyond our experience.

An anecdote here might help to lighten an over-serious discussion. At one time there was a proposal at KAW that we organize a wisdom engineering conference. John was enthused but there was a backlash from those who felt we had no competency to do so and would undermine our credibility. That view prevailed. I see many people now using LLMs to address dilemmas in their lives, seeing them as sources not only of knowledge but of wisdom, much as they would have gone to their elders or other respected persons in the past. We know that this can have tragic consequences, particularly for children. Knowledge in itself is not wisdom and dilemmas may have many resolutions, some of which are most unwise. Is that another dimension of knowledge or does it require a different mode of thought, and can it be inculcated in an LLM? I do not think this lightens the discussion but maybe will increasingly direct it.

Synopsis of Our Research During the KAW Years

The story of the KAW meetings, all of which continue to this day, 40 years after their inception, would take a book to describe—it is rich with anecdotes, with the growth of international collaborations and collegial friendships, and with steady growth towards the AI world of today. I cannot do it justice here but can summarize a few of our contributions during our decade of major involvement. Not only were we participating in 2 KAWs each

year but also in AAAI/IJCAI[13] conferences and associated workshops and specialist meetings like *Knowledge Discovery in Databases*. We were also invited as guest speakers to conferences targeted on AI in particular professions, such as accountancy and manufacturing, to keep them up to date with the latest developments in knowledge acquisition techniques and learn of their specialist needs.

Personal Construct Psychology—Kelly's Conceptual Grid as a KA Tool

Our initial contribution to knowledge acquisition for knowledge-based systems was based on Mildred's doctoral studies on computer-based elicitation of George Kelly's *conceptual grids* designed to make explicit and model their conceptual structures. I am deliberately using Kelly's terminology in his 1955 exposition of personal construct psychology rather than propagate the attempt by his initial expositors to avoid his term 'concept' by using his name for one type of conceptual grid, his 'role repertory grid.' This had the effect of detaching PCP from mainstream mathematical psychology or cognitive science where it belongs (*Kelly's psychological geometry as an intensional logic of knowledge representation*, 1992, *Personal construct psychology and the cognitive revolution*, 2003, Computer aided constructivism, 2012).

We, and the research team at Boeing under John Boose, started from Mildred's tools and extended them to accommodate the needs of practical knowledge-based system development (*Knowledge acquisition tools based on personal construct psychology*, 1993, *Eliciting knowledge and transferring it effectively to a knowledge-based systems*, 1993).

Conventional Knowledge Media, Electronic Journals

In studying knowledge acquisition for knowledge-based systems it seemed to us that the continuing major role of conventional knowledge media, such as books and journals should not be neglected. They needed to be enhanced and integrated with developing technologies so the best of both were available in next generation systems. We were already experimenting with active documents (*A minimalist approach to the development of a word processor supporting group writing activities,* 1991), so, when AAAI in 1992 announced a workshop on *Communicating Scientific and Technical Knowledge,* we decided to participate with a paper that went back to fundamentals as to the origins and purpose of scholarly publication, and delineated reasonable objectives for future developments. We concluded that "The major impediment to significant advances in the application of knowledge support technologies arising out of such research is the difficulty of integration with the obsolescent functionality of the paper journal. The development of

[13] Association for the Advancement of Artificial Intelligence/International Joint Conferences on Artificial Intelligence

digital journals is a prerequisite to the development of knowledge support systems in which human and artificial intelligence can be effectively combined" (*Social dynamics of scholarly publication,*1992).

We were approached at the meeting by the editor of a special issue on electronic journals and asked whether we could extend our workshop paper into a journal one (*An agenda for digital journals: the socio-technical infrastructure of knowledge dissemination*, 1993). Then by Alan MacDonald, the University of Calgary Librarian, to give a keynote address at a medical librarians' conference (Social and technical dimensions of electronic journals, 1993). Then for a chapter in a book on *Computer Networking and Scholarly Communication in the Twenty-First Century University* (*Dimensions of electronic journals,*1995).

The transition from the printed journal to an electronic successor is still ongoing but we have made substantial progress towards open-access, and a cultural expectation that should be the norm. Claerbout's vision of reproducibility is gradually coming about, the publication where not only can one reproduce the results but experiment with changes in the data, theory and parameters. Wolfram's *notebooks* offer that for a wide range of disciplines, but 35 years ago we were too optimistic and thought the transition would be rather sooner.

We continued our research on collaborative writing (Supporting collaboration in digital journal production, 1993). In 1990 we published a paper in a book that demonstrated the expert system described in the book operating from within the electronic version: "This paper is written in a document production tool that appears to a user as a word processor but also acts as an expert system shell with frame and rule representations supporting deductive inference. The electronic version of the document is active, providing typographic text and page layout facilities, versioning, hypermedia, sound and movies, hypertext links, and knowledge structures represented in a visual language" (Documents as expert systems, 1992).

MultiMedia Active Documents

We came across many instances of knowledge-based systems where images, photos and diagrams were valuable ways of presenting knowledge and began to make that affordance available in all our tools, images rather than names representing elements in WebGrid, images representing concepts in semantic networks or less formal concept maps (*Hyper-KSE: The integration of acquisition, interpretation and documentation*, 1990, *Concept maps as hypermedia components*,1995, *Embedded interactive concept maps in web documents,* 1996). We increasingly saw the active document as our general-purpose

medium for acquisition, representation and application of knowledge (*Open architecture multimedia documents,*1993, *Active documents combining multimedia and expert systems*, 1994, *Embedding formal knowledge models in active documents*, 1999).

We continued to study scholarly communities' use of technological support (*Knowledge acquisition and representation techniques in scholarly communication*, 1994, *Modeling the human factors of scholarly communities supported through the Internet and World Wide Web*, 1997). I took the opportunity of an invited keynote address at the annual conference *Conceptual Graphs* community, which collaborated largely through the web, to review the major large-scale knowledge representation services that I felt should play an increasing role in scholarly discourse, as well as the evolution of the more conventional ones discussed above (Representation, discourse, logic and truth: situating knowledge technology, 1993).

Differences in the Usage of the Technical Vocabulary of a Community

A major problem across all media was how to manage the issue identified by Wittgenstein that the 'meaning' of words in a language was not a matter of formal definitions but of use. We found that different experts in the same community of practice used the same technical term to make different distinctions, and different technical terms to make the same distinction (*Comparing conceptual structures: consensus, conflict, correspondence and contras*t, 1989). Early expert system developers noticed the phenomenon and some recommended using only one 'expert', but in practice it made for a more robust and usable system to identify the discrepancies and support them in the system design.

The Need for a Principled Knowledge Representation Server

At the beginning of our knowledge acquisition research, we were content to accept the received wisdom that the KA tools needed to deliver the encoded knowledge to an expert system *shell* and worked closely with Alain Rappaport of Neuron Data to use their *Nexpert* system (*Integration of acquisition and performance systems*, 1988, *The automatic generation of classes, objects and rules at the interface between knowledge acquisition tools and expert system shells*, 1989, *Combining paradigms in knowledge engineering*, 1992), and with Marc Linster at GMD[14] to use their *Babylon* system (*Development of second generation knowledge acquisition systems*, 1990, *Supporting the knowledge engineering life cycle*, 1994).

These were very effective collaborations, but the inference systems used by the shells were not well-defined in logical terms having had pragmatic constraints imposed on them to

[14] German National Research Centre for Information Technology

satisfy industry needs. We decided to develop our own inference engine based on the Aristotelian logic used by Kelly supplemented by ideal elements representing conjunction, disjunction and negation if needed, and by Reiter-style default reasoning if needed (*Knowledge representation servers: a generic technology for knowledge acquisition and knowledge-based systems*, 1990, *Empirical investigations of knowledge representation servers: Design issues and applications experience with KRS*, 1991). We also implemented a visual language for specifying logical structures (*The Sisyphus problem solving example through a visual language with KL-ONE-like knowledge representation. Sisyphus Working Papers Part 2*, 1991) and later showed it was as formal a specification as a symbolic one (*Designing visual languages for description logics*, 2009).

In retirement I became fascinated by the notion that the visual language could provide rigorous foundations for logic as a science of patterns (*Universal logic as a science of patterns*, 2015) and have since shown that the same patterns underlie order relations and mereology—so, Aristotle's mereological and Kelly's geometric approach to reasoning are naturally explained.

The Role of Machine Learning in Knowledge Acquisition

The machine learning research community and the KAW community were largely separate at the time of the first KAW, but I saw them as two approaches to developing knowledge-based systems and benefiting from collaboration. One of our early studies was to compare knowledge acquisition from 'experts', assumed to know the relevant attributes and correct outcomes, with induction from cases, possibly contaminated by irrelevant attributes and incorrect outcomes. The induction algorithms coped well with noise or some irrelevant attributes but the combination of the two caused a massive jump in the amount of data required. However, the overall conclusion was that a good induction algorithm could produce the same intelligible knowledge structure from a large collection of lower quality data as it did from a small collection of carefully curated data from 'experts' (*The trade-off between knowledge and data in knowledge acquisition*, 1991).

Making Induced Knowledge Understandable

One of our primary objectives was to ensure that all knowledge structures were understandable so that they could be assessed by people. This was one attraction of the grids because the vocabulary came from those in the community of practice and the graphical analyses were suggestive of structures or created logical structures. However, we recognized that there are many diverse sources of knowledge such as text, diagrams, tables of data, and so on. We had designed rule induction algorithms and studied the work

of others, such as Ross Quinlan but were not happy with the comprehensibility of the results even if they *performed* well in terms of accuracy with new cases.

Our first studies were of how imperative rules might be incorporated effectively into descriptive term subsumption inference engines (*Integrating rules in term subsumption knowledge representation servers*, 1991) and the realization that the ways the rules were factored into different structures made a big difference in their size and comprehensibility (*Extracting knowledge from data*, 1989). It also became clear that describing with rules that had exceptions resulted in smaller more natural knowledge structures that were also simple computationally (*Exception DAGs as knowledge structures*,1994, *Structured and unstructured induction with EDAGs,* 1995, *Inducing knowledge*,1995, *Transforming rules and trees into comprehensible knowledge structures*, 1996).

In 1986 Mildred was at a conference on expert systems and simulation in Tokyo and met Paul Compton, becoming aware of his research with Bob Jansen on *ripple-down rules* (RDR) with two types of exception that could be developed incrementally with no rule ever being discarded. She invited him to EKAW 1988 in France, and we met at a KAW every year thereafter. I added RDR to my induction programs as they were easy to induce but the size of the directly induced set was about half that of the incremental one. They did not satisfy our requirement for understandability but had become a well-proven technique for developing an expert system. Paul and I were able to carry out some interesting studies of the long-term dynamics of expert systems because of the longevity of the Garvan project (*Induction of meta-knowledge about knowledge discovery*,1993, *Induction of ripple-down rules applied to modeling large databases*,1995).

WebGrid II incorporated induction into the elicitation process to encourage the introduction of new elements that would not be predicted by the existing ones. The client could partially specify an element, then run a chosen induction algorithm on the existing grid to predict the element's missing attribute(s). If the induction was correct, the user could decide whether to add the element to the grid (and likely get a request to add a construct to distinguish it from a similar one) or, if it was wrong, fill in the missing value(s) and add it (WebGrid II: Developing hierarchical knowledge structures from flat grids,1998). That version was used on several basic AI courses because it gave a very good feel for the process of induction, and that one learnt from failures more than successes.

Knowledge Creation Processes in Society

Another stream at the KAW meeting was the bigger scene of knowledge processes in society, what role did knowledge perform, how was it formed, disseminated, applied, stored, passed on, how was expertise inculcated, and so on. I studied this as a separate

exercise from developing tools but tried to present the theories in a way relevant to the support of the processes (*Foundations of expert systems*,1985, *An overview of knowledge-acquisition and transfer*,1987).

One focus was the processes whereby human agents become recognized as 'experts' in some domain, where the positive feedback loop between having a reputation for a skill leads to more opportunities to exercise that skill and learn from more cases (How do experts acquire expertise?, 1987, Positive feedback processes underlying the formation of expertise, 1988, Modeling and extending expertise, 1993). These models highlighted the social aspect of expertise formation and led to the ultimate abstraction, that one should regard the knowledge acquiring entity as a collective rather than an individual (The collective stance in modeling expertise in individuals and organizations,1994, Adaptive interactions in societies of agents, 1996, Knowledge management in societies of intelligent adaptive agents, 1997). We were also asked to explain to the PCP community how PCP has become significant to the KAW community (Expertise and expert systems: emulating psychological processes, 2005).

Manufacturing as Knowledge Science

My research in manufacturing continued and I presented it from a knowledge science perspective in a keynote address at an *IFIP Conference on Advances in Production Management Systems* (*A new science of manufacturing*, 1987). Doug Norrie in the Department of Mechanical Engineering invited me to work with him and his graduate students (*Object-oriented management planning systems for advanced manufacturing*, 1989, *A knowledge-based decision support system for flexible manufacturing*, 1989, *Object-oriented expert systems in management planning*, 1989, *An integrated object-oriented architecture for manufacturing knowledge systems*, 1989).

In 1993 Doug and I were invited to participate as Canadian representatives in the international Intelligent Manufacturing Systems (IMS) research program (*Knowledge systematization in the international IMS research program*, 1995), with Industry Canada providing $100,000 to support our activities with the GNOSIS project, and with me as Canadian Coordinator. We had been working on Mediator, a knowledge support system for all those involved in the manufacturing process from requirements through design, engineering, production, to maintenance and recycling. It was designed to facilitate communication, compliance with constraints including physical restrictions and legal obligations, and to generally represent knowledge about any activity or sub-system relevant to the manufacturing process (Mediator: an intelligent information system supporting the virtual manufacturing enterprise, 1995).

We foresaw many problems in coordinating a heterogeneous group of international projects and decided to make our contribution to GNOSIS the application of Mediator to the coordination of the GNOSIS project itself. For this purpose, we implemented a version based on the newly emerging World Wide Web (WWW). Mediator became the primary communication and record-keeping service for GNOSIS, subject to the limitations of the spread of technology in 1994 such as one project not having access to the Internet (communicated by fax). We supplied the final GNOSIS report on a CD-ROM containing all the project documents indexed by concept-maps together with videos of meetings and other associated materials (*Coordinating societies of research agents - IMS experience*, 1997). The development of Mediator for its original objectives in manufacturing also continued (*Concurrent manufacturing on the web*, 1996, *Distributed agent systems for intelligent manufacturing*, 1996).

The Current State of Our Conceptual Modeling Software

Rep Plus, the latest version of our conceptual modelling software that commenced as a by-product of Mildred's doctoral research has been freely available for many years, and the source code is also now freely available on our web site in the hope that others may maintain it when we are gone. In retirement we made major enhancements that were requested by colleagues or had been on our agenda waiting until we had time.

One enhancement was a major re-write of the manuals to make them far more tutorial, not just how to use the tools but what was appropriate usage and experimental method.

We have added *content analysis* based on Honey's method where the meaning of the construct pole names is compared, as Devi Jankovicz has advocated. This is also very useful in reducing the number of commonly named constructs in Mildred's *Mode* analysis.

We have added methods for the analysis of grids with all elements and constructs in common, often used in management studies of large cohorts from a community of practice—see Lina Bellman's publications, and examples in the Grids Manual.

We also implemented the *theory of anticipation* that Kelly describes in his book and found that it serves as an inference engine for the kinds of data we had investigated in our knowledge-based system research. We added an additional primitive to Kelly's element and construct, that he termed an 'intersect' but we called a 'class.' It may be regarded as an abstract skeletal element. This has not been written up for a journal, but is fully documented, with examples, in the RepGrid Manual.

Industrial Software Engineering Chair, Requirements Engineering

Mildred's and my story do not end with the formation of the KSI. That dominated our lives for some 13 years but left room for many other activities. Mildred was allocated a course in software engineering largely because no one else wanted to teach it and she saw the importance of it to our students and wanted to address their needs. She read the literature widely and realized that software engineering involved many issues beyond coding, such as understanding client requirements and, if necessary, helping them make the requirements explicit. This was so similar to knowledge elicitation for expert systems that she started up a line of research on requirements elicitation and presenting it at the first UK conference in 1994 on the topic (*Requirements acquisition,* 1996). On the undergraduate side, she designed a course where students played a role as customers in one group of 12 and as suppliers in a different one (*Learning Environments for Requirements Engineering*, 1996). The customers had to specify the system they needed, and the suppliers had to design one that addressed that need.

The course was very attractive to industry as the students learnt to work in teams, understand customer requirements, make presentations to clients, skills that were not normally taught in computer science departments yet vital to the computer industry. The Department applied to NSERC for a Chair in Software Engineering to strengthen its offerings in that area. The applicant selected by the Faculty Committee was somewhat mathematical and not particularly attractive to the industry sponsors. He was not accepted by the NSERC evaluation panel and when the Dean asked the industry sponsors what kind of person they would like they answered 'someone like Mildred Shaw' and she was appointed as Industrial Software Engineering Chair in 1996. She then went on to establish a research master's program in software engineering and a doctoral program in requirements engineering.

Dean of Graduate Studies, Associate VPs Research and Academic

My Killam Chair was coming to an end after 13 years and I would have liked to retire but did not want to do so until Mildred was ready, so I accepted a nomination to be Dean of Graduate Studies, was selected and also appointed Associate VP Research. In addition I became the research and graduate studies representative on a university-wide 'Change Program' that was examining all University processes with a view to improvement. I shortly became responsible for the entire program when the Associate VP Academic who had been responsible resigned and the Vice President (VPA) phoned from Ottawa and asked me to take on the position for 6 months as they searched for a replacement. So, I finally became an administrator (or three!), as payback for a life-long happy working environment, but kept my research and publication going.

It was a very heavy workload. I would get to work at 7.30 in the morning and rarely leave until 7.30 at night, have supper, go to bed at 10.30, and when Mildred fell asleep get up and work on my computer until 4am. I used to joke that I was on committees all day and that is where I got my sleep. However, at the retirement dinner the President gave us when we left, in his speech he remarked, "at my Executive Committee as each item comes up on the agenda, I look at Brian and if he is glaring at me, I know it will be contentious, but if he is asleep then I know it is not a problem." So maybe the joke had substance!

Retirement on Vancouver Island

After 3 years Mildred and I wanted more time with one another and decided to take early retirement. Everyone was shocked—I had colleagues earnestly advising me that I should not retire because I was a workaholic and would be dead within six months. Little did they know that I had a huge library on the history of science, the Oldenburg correspondence, Newton and Boyle's collected works and correspondence, all the histories of the Royal Society, historiographies of science, the collected works of obscure scientists such as Cadwalader Colden, and so on. I had studied the philosophy of science, methodologies of science and science itself, but not the raw history of science and those who had created it. I had a thirst for the raw data. So, I spent the first decade of retirement studying my library and finding a very different picture from any previously presented to me. I also studied the historiography of science to ascertain what historical portrayals were methodologically sound—the scientific method of historians.

Retirement Research and Writing

I continued to study various topics that interested me and write academic papers, largely when they were requested for special issues or book chapters.

In a book chapter for the *Handbook of Knowledge Management,* I added a *World 4* to Popper's 3, that of the *a priori*, innate or assumed implicit presuppositions and ideology that are an essential, but tacit, precursor to any of our activities (Fig. 1) (*Organizational knowledge acquisition*, 2003).

In 2009, Jean-Yves Béziau, the editor of Logica Universalis issued a call for papers for a special issue addressing the questions:

1. Do all human beings have the same capacity of reasoning?
2. Does reasoning evolve?
3. Do we reason in different ways depending on the situation?
4. Do the different systems of logic reflect the diversity of reasonings?
5. Is there any absolute true way of reasoning?

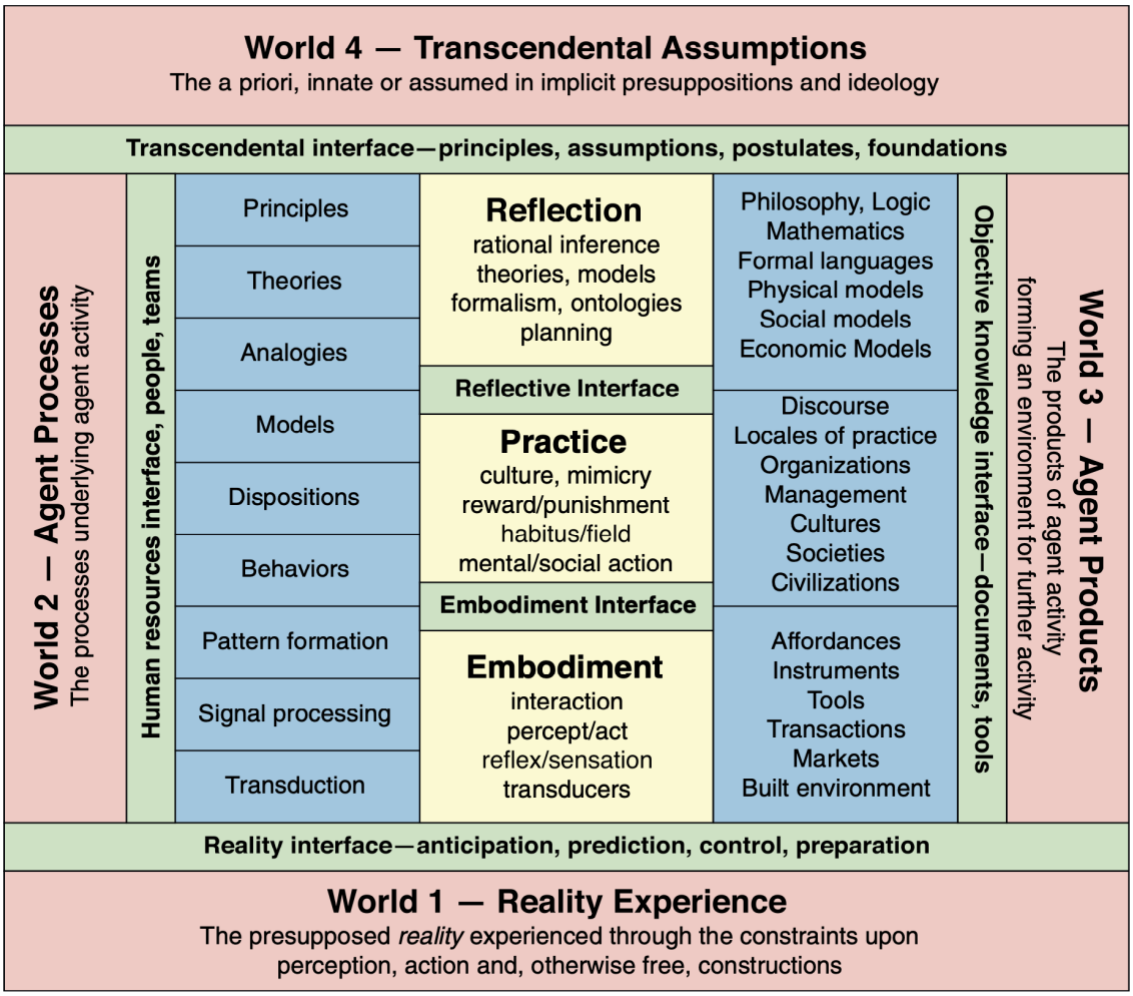

Figure 1: Popper's 3+ Worlds and Klir's epistemological hierarchy

I found the questions intriguing and wrote a paper that addressed each in turn (Human rationality challenges universal logic, 2010). The conclusions are so long and complex it is not appropriate to reproduce them here. More importantly, they predate the rise of LLMs and their partnerships with people which makes for reasoning entities the like of which we have never seen before. The questions need to be asked again in this very different world.

The 25th anniversary of the first KAW occurred in 2011and we were invited out of retirement when KCAP[15], its successor returned to Banff and I was asked to give a keynote address (Knowledge capture through the millennia: from cuneiform to the semantic web, 2011, Presentation visuals for "Knowledge capture through the millennia", KCAP2011, 2011), followed by a paper celebrating 50 KAWs world-wide in a special issue of IJHCS (Knowledge acquisition: Past, present and future, 2013).

[15] International Conference on Knowledge Capture

In a Festschrift for Lotfi Zadeh, I reflected on my experience of the era of fuzzy sets research and on his development of computational algorithms for *computing with words* that have inspired so many theoretical innovations and industrial applications (*Living in an uncertain universe,* 2013).

I wrote an article surveying five decades of research on human-computer interaction for the 50th anniversary of IJHCS and came to the depressing conclusion that by our idealistic pursuit of universal access to computers we had opened the doors of our homes to criminals world-wide. I also extended the BRETAM model of the information technology infrastructure, its growth and interdependencies, to include a new era of *Symbiosis Science* where artificial intelligence is partnering with natural intelligence in ways that are making us uncomfortable (Fig. 2) (*From facilitating interactivity to managing hyperconnectivity: 50 years of human–computer studies*, 2019).

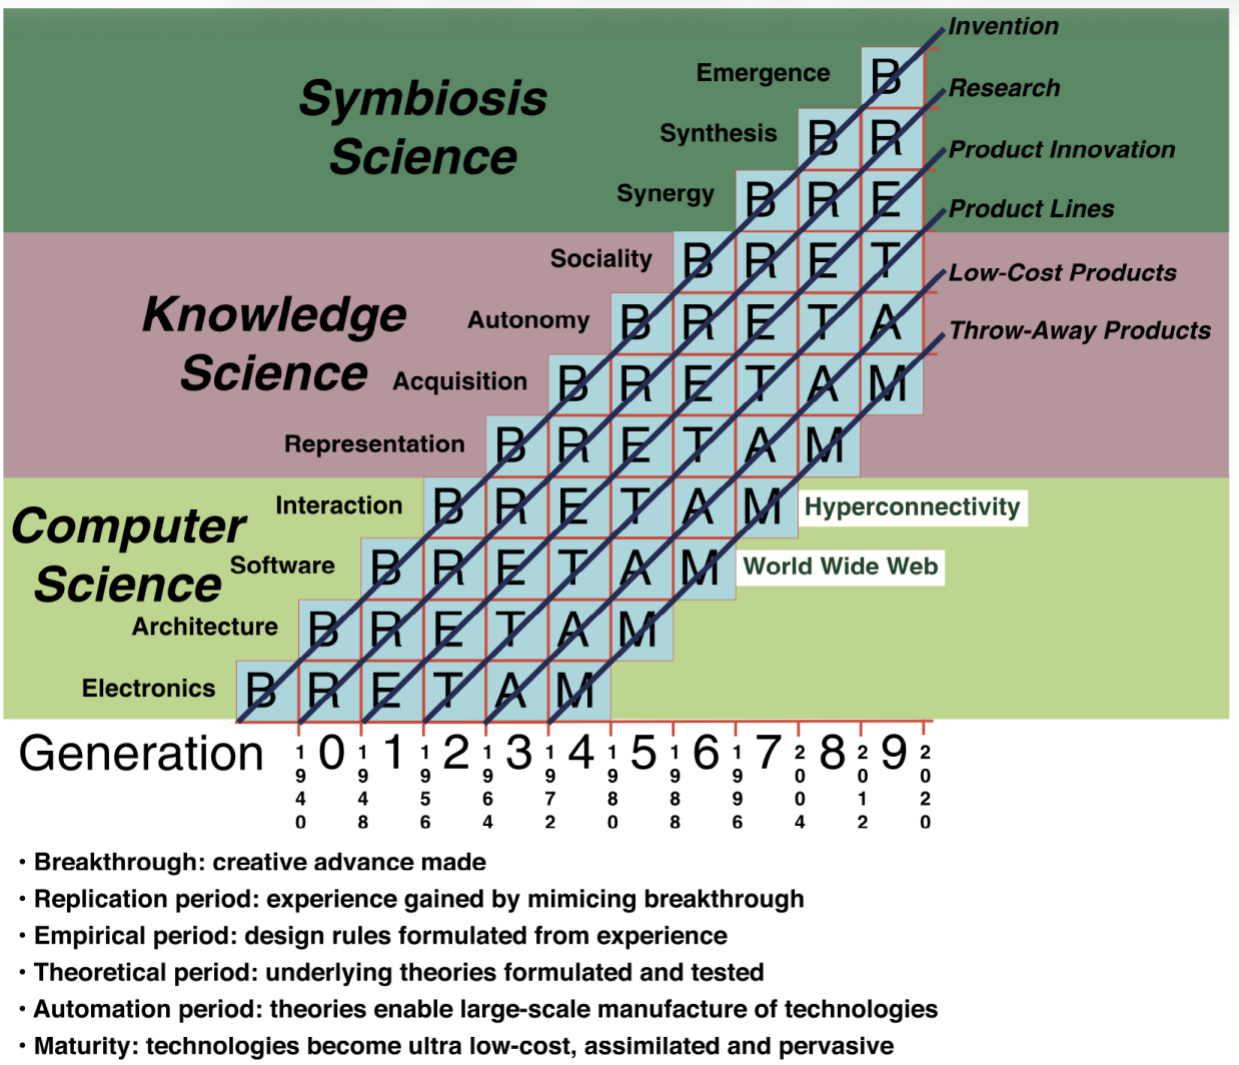

Figure 2: The latest BRETAM model of the infrastructure of information technology

One project that took up much of my time in retirement was reformulating logic as a science of patterns that avoided as much as possible the circularities of most attempts at foundations. It was implemented in the last generation of our knowledge representation servers and includes Reiter's default reasoning in a way that enables it to replicate the phenomena of natural language conditionals. An initial paper was included in Festschrift for Jean Yves Béziau (*Universal logic as a science of patterns*, 2015).

I was asked to contribute to a book on stochastic computing and wrote a chapter on its origins. I was delighted to find that digital libraries enabled me to supplement the known history with related work that was not known at the time, and also to identify an additional inventor, William Thistle, a scientist at the Canadian Armament Research and Development Establishment in Québec, whose only publications are a patent and a report that was not circulated externally until my colleague, David Hill, discovered it in the 1970s. Copies of them are available on this website (*Origins of stochastic computing*, 2019).

I also contributed an article on *stochastic neuromorphic computing* to a special issue of *IEEE Design and Test* that notes that deep learning has its own BRETAM hierarchy developing as LLM[16]s are being used to design the next generation of LLMs (*A conceptual framework for stochastic neuromorphic computing*, 2021).

Ballroom Dancing, Adjunct Professor at University of Victoria

Mildred and I took up ballroom dancing, something neither of us had ever done before, with a rigorous and well-documented method and criteria of success, something that needed us to be supportive of one another, accept our mutual dependence, and trust one another completely. We took it very seriously, had 4 private teachers, attended many workshops from a diversity of teachers, went to 2 or 3 professionally run dances each week, and generally treated it with the same seriousness as we had our scholarship. We were commended in our first exam and went through all levels of the syllabus to gold+. Then came COVID, dancing stopped for 3 years and we never restarted. I accumulated an extensive ballroom dancing library, having all the editions of the major ballroom dancing handbooks intending to write it up as a perceptual motor skill that evolved from the 1920s to today but never got around to it, one of many incomplete projects.

So I was probably busier in retirement than when employed, but at least the agenda was of our own choosing. Mildred's top doctoral student in requirements engineering, Dana Damian, had accepted a position in the Department of Computer Science at the University of Victoria, our local University where I had also been offered an Adjunct position. I served on two of her students' doctoral committees and mentored them in the nuances of

[16] Large Language Machines

presenting their results which were very interesting. I also gained some knowledge of the university and was impressed by the drive of the President, David Turpin, to raise its research status. So, when the tax law changed in Canada such that we had to realize a substantial capital gain on securities we held in the UK, I suggested to Mildred we experiment with donating it to the University to see if they used it wisely. We had already decided to leave a substantial part of our accumulated capital to charities but not which particular ones.

Donors to the University of Victoria

Mildred wanted to donate to the Faculty of Education as she felt that Medicine, Engineering and Science got the lion's share of donations, but school education was foundational to a healthy society. So, we set up an endowment for the Faculty that turned out to be the largest that they had ever received and enabled 3 successive Deans to initiate major projects, keeping us in touch with the outcomes. The money was used well and encouraged us to invest more in helping the University achieve the excellence to which it aspired.

President David Turpin used to invite Mildred and me to lunch once a year to keep us in touch with progress at the University and the lunch on 20th June 2013 included Jonathan Bengtson who had joined as Librarian the previous year. He and I discussed libraries, their roles and future at great length and I was impressed by his progressive views and plans for the Library. We discussed our book collection, which was left to the University and contains some very rare books on printing and Aristotle that should be digitized. We also discussed the formation of something akin to the KSI as an institute hosted by the Library for the study of scholarship and its role in society. I had already discussed with Mildred my interest in donating to the Library and indicated to Jonathan our interest in doing so. We had a long discussion through 2015 about possible projects. Jonathan indicated he was not yet ready to found an Institute but would like to use our funds to found a journal in the area of knowledge science which was launched in 2017 as *KULA: Knowledge Creation, Dissemination, and Preservation Studies* and rapidly established an excellent reputation.

The Kula Academy

Jonathan emailed me in November 2022 that he had been considering a 'Kula Institute' that would subsume some existing initiatives and sent a draft of a proposal. We worked with him to adjust the terms of our wills to leave the residue of our estates to support the Kula Institute. In May 2024 he reported great enthusiasm in the senior admin for the proposal, which had become: *KULA: Library Futures Academy,* with a mission of *advancing the research role of academic libraries*. We set up a *Project Fund for Kula* so that the whole of

our annual donation was immediately available rather than just endowment revenue and by October 2024 Matt Huculak had been appointed Director and was initiating projects. The Kula Academy rapidly attracted attention and initiated partnerships with other libraries in Canada and internationally, appointing Fellows whose activities have become widely recognized.

To complete the circle, Dan Freeman, one of my grad students from the 1980s, who has had a stellar career in the infotech industry and founded Composium.ai which specializes in AI systems whose responses are fully grounded and not subject to 'hallucinations' has now been appointed UVic Libraries' first Kula: Library Futures Academy Industry and Community Fellow. He contributed the query system for our publications on this website, the podcast based on them, and the AI-generated festschrift, which probably establishes a new literary genre, first exemplified in this Library!

We intended this website to be a project of the Kula Academy as an early role model for what might be done to disseminate the contents of fonds in Libraries' special collections more effectively in Libraries of the Future. We only barely made it in time, and it is less than we hoped, but perhaps enough to challenge others to do better. It is particularly fitting that the content of the collection is some small part of the story of how the information age came into being.

That is where we are now and where Mildred's and my story ends.

"In my end is my beginning." T. S. Eliot, Four Quartets

But, as Eliot reminds us in the last sentence of *East Coker*, it is also where so many new stories begin.

I have been fortunate to live long enough to see Jonathan take my vague idea and turn it into an asset for this era that builds on *this* library's strengths to create a new future-looking institution that embodies the culture, values and ethics that have made libraries respected over the millennia.

It has been my pleasure to see Matt bring that Institution into being by facilitating those who have their own dreams of how a library of the future might enhance the services available using new technologies to bring those to life.

It has been inspiring to see that facilitation being available to a researcher like Chloë Farr who while working in the Library recognized that there were difficult problems in OCR that could now be solved by capabilities in AI that she had analyzed in her linguistics research. Kula provided the resources to let her demonstrate that very successfully, making the

results available to a growing network of partnering libraries, and launching Chloë on a new career.

It is continuing to increase its activities and widen them to support more people and projects and is itself attracting more support from those who recognize its value.

That to Mildred and me is what the good life is all about, a win-win-win situation where cooperation, collaboration and sharing let many people realize their dreams, and enhance the lives of many others.

We thank all those who over many years have supported us in achieving our dreams and those who have made it possible for us to pass it on to others.