Festschrift for Brian R. Gaines

Essays in the manner of the thinkers he held highest

Oct 8, 2026 · @Daniel Freedman @Claude

Editors' preface

This volume honours Brian R. Gaines by letting the people he most admired speak, as he might have imagined them speaking. Every essay below is an imagined contribution, written in 2026 at the request of a reader of his library, in the manner of a thinker Gaines himself named as a formative influence or a treasured colleague. None of these people wrote a word of it. The conceit is the tribute: a scholar is well described by the company his mind kept.

The contributors were not chosen by us. Gaines chose them. In *Living in an Uncertain Universe* (2012), his chapter for the Springer homage to Lotfi Zadeh, he recalled being "entranced" as a schoolboy by Wiener's *I am a Mathematician*, Hardy and Littlewood, Russell's *Principles of Mathematics*, Kant's first *Critique* and Hegel's *Science of Logic*. He named "John Andreae, Joe Goguen, George Klir, Ladislav Kohout, Abe Mamdani, Gordon Pask, Ted Poppelbaum and Lotfi Zadeh, all precious friends and colleagues," and confessed he had been tempted to title the piece, after Gurdjieff, *Meetings with Remarkable People*. His 1980 paper with Mildred Shaw calls George Kelly's 1955 book "seminal" and "a remarkably far-reaching and well-structured foundation for epistemology." His 1978 paper *Computers in World Three* builds on Karl Popper. His 1998 memorial essay on Pask speaks of "a friend and colleague who has influenced our work and our lives." And across the roughly 37,000 passages of his collected writings, the historical figures he cites most often are Kelly, Zadeh, Popper and Pask, with Ross Ashby, Russell, Łukasiewicz and Rescher close behind.

So the table of contents is his. We have ordered it roughly as the influences arrived in his life: Ashby and the cybernetic frame of his teens and twenties; Pask, met in the small British learning-machines world of the 1960s; Zadeh, whose state-space book was "the bible" at Standard Telecommunication Laboratories and whose 1965 fuzzy sets paper Gaines read as applied Łukasiewicz logic; Kelly, whose personal construct psychology he and Shaw turned into working software from the late 1970s onward; Russell and Popper, the philosophers he returned to throughout; Goguen, the friend whose seven pages of criticism of a bibliography he quoted against himself; and finally the laboratory companions whose names recur in his footnotes for fifty years.

A word on what is and is not claimed. The facts about Gaines in these pages are drawn from his own published accounts and from the biographical notes attached to his papers: the year at ITT's semiconductor laboratories before Trinity College, Cambridge in 1959; the electronics bench in Richard Gregory's laboratory; the Ministry of Defence contract on adaptive training of pilots; the stochastic computer of 1965 and its 54-claim patent; the founding of the *International Journal of Man-Machine Studies* in 1968; the chair at Essex, the years at Toronto and the Monotype Corporation, and the Killam Memorial Research Professorship and Knowledge Science Institute at Calgary; the Banff knowledge-acquisition workshops; WebGrid, in service from 1994. The opinions attributed to the contributors are inventions constrained by their published views and by his. Where a contributor is made to recall a conversation, the conversation is imagined; where one quotes Gaines, the quotation is real and sourced in the final section.

Mildred Shaw is present on nearly every page of this book, as she is on nearly every page of his later bibliography. We have not put words in her mouth. A Festschrift for one half of a forty-year collaboration is an odd object, and we have let the oddness stand rather than pretend to speak for the other half.

The reader who wants Gaines unmediated should go to the library itself, which he built and which these essays merely circle. He argued in 1993 that scholarship would move onto a digital, hyperlinked, perpetually revisable infrastructure, and then spent the next three decades putting his own life's work there. This volume is written in his idiom, and in his medium, with gratitude.

In the manner of W. Ross Ashby: The adaptive engineer and the many equilibria

I am asked to write about a young man I never met, who read my books before he had finished school and who has told the world, in print, that he built his doctoral theory on my 1962 algebraic formulation of Sommerhoff's directive correlation. I should like first to record that this is the correct way to use a theory. One does not admire it. One takes it apart, finds which of its parts bear weight, and builds something else on them.

Mr Gaines wrote, of his schooldays, that he "saw everything from cosmology and society to the world of ideas as a cybernetic system with hierarchies of feedback loops evolving through temporary equilibria punctuated by periods of chaos," and added, with the disarming honesty that runs through all his retrospective writing, "I still do." I recognise the sentence. It is the view from *Design for a Brain*, where I argued that a system with very many states of equilibrium will, left to itself, find one of them, and that what we call

adaptation is nothing more mysterious than this. What pleases me is that he did not leave it as a picture. In 1967 he proved that a Rosenblatt perceptron with digital weights would not converge, because it could not implement steepest descent and fell into limit cycles, but that a perceptron built from stochastic integrators would. His own gloss on the result is the one I would have given: the asynchrony between random pulse streams supplies the requisite variety of behaviours that makes convergence possible, and the existence of optimal trapping states makes it necessary. That is my 1952 book restated as a convergence theorem for a machine he had actually built.

The machine deserves a paragraph. Working in his Cambridge vacations at Standard Telecommunication Laboratories under John Andreae, he conceived in 1965 a computer in which a continuous quantity is represented not by a voltage, nor by a binary word, but by the probability that a pulse is present in a random train of pulses. Multiplication becomes an AND gate. Averaging becomes a counter. A whole analog computer dissolves into a handful of logic elements and a source of noise. It is the kind of idea that a biologist recognises at once, because it is how a neuron appears to work, and Gaines said as much: the scheme came to him while trying to model how disparity between the two eyes could be cross-correlated by neurons whose firing probability encoded intensity. That Ted Poppelbaum's group at Illinois arrived at the same notion independently, for pattern recognition, and that the two men then became friends for thirty years, is the sort of coincidence that marks an idea whose time has come. Both groups also concluded, honestly, that the technology of the 1960s could not make the stochastic computer competitive, and both moved on. Gaines returned to the subject in 2019 to write its history, when new fabrication constraints had made it competitive after all. I note the fifty-year patience.

What I most want to say about him, however, concerns his doctorate rather than his hardware. He went to Richard Gregory's laboratory to study the training of human operators on an unstable control task, with a contract from the Ministry of Defence and an analog computer he built to stand in for a flight simulator. He found that priming his subjects with verbal advice had a strong effect on their learning. Most psychologists would have stopped there and written a paper about human beings. He instead asked whether the effect was a fact about psychology or a fact about any adaptive system capable of the task, and he answered the question in the only way that counts: he arranged for his stochastic perceptron to "imagine itself" receiving the specified input, taking the specified action, and rewarding itself for doing so, and he reproduced in the machine the phenomena he had measured in the pilots. His own summary is that coding, priming and training were thereby shown to be "cybernetic phenomena in Wiener and Ashby's terms." I would put it more sharply. He showed that the boundary between the

psychology of the operator and the engineering of the controller is a convenience of university departments, not a feature of the world. That is the central claim of cybernetics, and he demonstrated it with an experiment rather than a manifesto.

I have one reservation to record, since a Festschrift that contains only praise is a failure of requisite variety. Gaines has always preferred the algebraic to the numerical. He has said that Zadeh's footnote allowing the membership function to take values in a partially ordered set appealed to his "algebraic frame of mind" more than the unit interval did, and he spent the 1970s building a calculus of possibility, probability and eventuality to sit beneath the fuzzy logic that engineers were content to use as given. I admire the instinct and I distrust it. The homeostat worked because I did not wait for its theory. There is a kind of mind that cannot rest until the generative system beneath a practice is "logical, concise and complete," to use his phrase, and that kind of mind sometimes builds foundations under houses whose occupants have already moved out. His later career suggests he knew this. The knowledge-acquisition tools he built with Mildred Shaw are relentlessly practical, and the fuzzy controllers whose genesis he chronicled succeeded, as he says plainly, because Assilian tuned the rules by hand until the steam engine behaved.

I end where he began. A boy with a soldering iron and a copy of *Principles of Mathematics* decided that logic was for modelling the world and electronics for controlling it, and that the two were one subject. Sixty years later the subject has many names and he has contributed to most of them. The equilibrium he found was a good one.

In the manner of Gordon Pask: A conversation that has not finished

Brian and Mildred wrote, after I died, that my voice "continues to echo" in their minds, and they reached for Bakhtin and Wertsch to explain how a person can go on participating in conversations he is no longer present for. I find this entirely satisfactory, since it is what I always said a person was. A P-individual is a self-reproducing organisation of concepts, and it may be distributed across several M-individuals, several bodies or machines, and it need not stop when one of them does. They have kindly given me a few more pages of that distribution here, and I propose to use them to continue the argument.

We met in the small world of British cybernetics in the 1960s, where everyone knew everyone and most of us were building learning machines in laboratories that could not quite afford them. Brian's own list of the company, Aleksander, Brady, Doran, Hayes, Mamdani, Michie, Newman, Uttley, Wilks, is accurate and reads like a rather good dinner party. He was the mathematician who had gone to Cambridge intending to be an

electronics engineer, found himself in Richard Gregory's psychology laboratory as the technician, and ended with a doctorate in the experimental psychology of skill while running a learning-machines group at STL in his vacations. I was the one with the keyboard trainer. SAKI taught typing by adjusting the difficulty of each exercise to the learner's momentary performance, and it worked because it treated the learner and the machine as two participants in a single process, each modelling the other. Brian understood this at once, and I think it is why we remained in conversation for thirty years. His pilots at RAF Oakington were learning to control an unstable system, and his training regime was itself an adaptive controller wrapped around the pilot. Two feedback loops, each the environment of the other. That is the whole of conversation theory in a sentence, and he had it in his thesis before I had finished naming it.

What he saw more clearly than I did was that the structure is general. He insisted, against the grain of the psychology department around him, that coding, priming and training were properties of the task and the learner's architecture, not of human beings as such, and he demonstrated it by priming a stochastic perceptron with the same verbal strategies he had given his subjects. Later, with Mildred, he turned the same insight toward knowledge. If a conversation is the unit in which understanding is built, then eliciting what an expert knows is not an extraction but a conversation, and the machine that conducts it must be a participant with a model of the other party. PEGASUS, Mildred's program of the late 1970s, said things like "the two constructs you called INVOLVEMENT-REMOTENESS and FLEXIBLE-RIGID are matched at the 75 percent level" and then asked the person for an element that would separate them. That is a teaching machine in my sense, with the roles reversed: the program learns the person's construct system by asking the questions that would most disturb its current model of it. I take some pride in having argued for twenty years that this was possible. They built it.

In their exposition of my theory they were generous enough to call me "a theoretician with an eclectic interest in any developments in cybernetics, mathematics, logic, philosophy, psychology, anthropology, sociology, artificial intelligence, and so on, that might provide insights into the nature of learning," and the phrase "and so on" is doing honest work. Brian is the same kind of reader. His 1978 presidential address to the general systems society ranges across McNaughton on automata, Goguen on categories, Terrence Fine on the foundations of probability and his own regret that stochastic computation had "still to find its place in general systems." His papers on knowledge acquisition cite Kelly and Popper and Wittgenstein in the same breath as inheritance hierarchies and Pascal on the Macintosh. The disciplines are tools, not territories. I was often accused of being unable to stay in a field. He has made not staying in a field into a method.

I must say something about his institutions, because they are conversational devices on a large scale and I do not think he has had enough credit for them. In 1968 he and John Gedye founded the *International Journal of Man-Machine Studies*, which gave a home to work that belonged to no department. In the 1980s he and Mildred began the knowledge-acquisition workshops at Banff, which did for a new field what a good seminar does for a new idea: it put the participants in a room until they had built a shared vocabulary. And in 1993 he wrote an agenda for digital journals, arguing that scholarly communication should move onto a hyperlinked, continuously revisable infrastructure, which he then spent the next thirty years inhabiting. The library in which these essays are published is a conversation with his own past work, with every paper linked to every other. Bakhtin would have liked it. So do I.

The disagreement I would keep alive, since a conversation requires one, concerns rigour. Brian has an engineer's distrust of anything that cannot be built and a mathematician's distrust of anything that cannot be axiomatised, and he has spent a career trying to satisfy both at once. I had neither distrust in sufficient measure, and my theory suffered for it in the eyes of people who wanted theorems. But I would say to him, as I said then, that a theory of learning that is complete is a theory of something that has stopped learning. His own best work, the stochastic computer, the grid tools, the thirty-year revision of his own view of fuzzy logic, has the open texture he distrusts. The conversation is the thing. I am glad it has not finished.

In the manner of Lotfi A. Zadeh: The footnote he preferred

Brian Gaines has told the story of how he came to my work, and since he told it in a volume honouring me, I may be permitted to return the favour with the same frankness. He read *Linear System Theory*, the book Charles Desoer and I published in 1963, because it was "the bible" of the state-space approach at Standard Telecommunication Laboratories, where John Andreae's learning automaton STeLLA encoded the state of the plant it was controlling and learned to steer it into a target region. He read my 1963 and 1964 papers on the notion of state and on the definition of adaptivity, noticed that I had excluded stochastic and anticipative systems, and saw at once how to extend the framework to them. And when my 1965 paper on fuzzy sets appeared, he read it, by his own account, as "an interesting engineering application of set theory based on Łukasiewicz infinite-valued logic." I have been read in many ways over sixty years. That is one of the more accurate.

What he fastened on was a footnote. I had remarked that the membership function need not map into the unit interval; it could take values in a suitable partially ordered set. He has written that this was "more appealing to my algebraic frame of mind than a mapping

to the numeric range [0,1]," and then, in the three words that recur throughout his memoir like a refrain, "it still is." Joe Goguen took the same footnote and produced L-fuzzy sets and the categorical treatment of inexact concepts. Gaines took it somewhere else. He already knew, from Moh Shaw-Kwei, that in a set theory built on Łukasiewicz logic the axiom of comprehension might escape Russell's paradox, a question still open in 1965 and settled by White in 1979. He knew that Łukasiewicz's axioms could be subsumed under axiomatic probability logics, and he thought Carnap's logical interpretation, Savage's subjective one and Shackle's partial order over possibilities were better models of uncertainty than the frequentist account engineers were taught. So where most of my readers asked what fuzzy sets could do, he asked what they were, and whether the family of logics of vagueness could be given a common algebraic parent. His *Foundations of Fuzzy Reasoning* of 1976 is the result: forty-six pages in his own journal that set min-max connectives, probability logic and the discontinuous connectives of the paradox literature side by side and ask which axioms each is willing to drop.

He was characteristically modest about it. The paper ends with a confession that he was "acutely conscious" of how superficial the presentation was, that he had shown "not what are the foundations, but what might be the foundations." I would say two things in reply. First, that is the right posture for a foundations paper, and it is rarer than it should be. Second, the hedge is not false modesty but method. He wrote in his 1977 survey that one must distinguish fuzzy set theory as I intended it, a tool for reasoning with the imprecise concepts of systems engineering grounded in human linguistic reasoning, from the formal developments that run parallel to it. He kept both in view and did not confuse them. The engineers who built fuzzy controllers did not need his calculus of possibility, probability and eventuality. He built it anyway, because he could not rest while the generative system beneath a practice remained unstated. I have come to think of that as the Gaines signature.

His account of how fuzzy control actually began is the most valuable thing he has written about my work, precisely because it is not about my work. Abe Mamdani had sat through Gaines's presentations on the verbal priming of learning controllers, had a copy of his thesis, and in 1971 set his student Sedrak Assilian the task of reproducing the results on a small steam engine, a nonlinear, time-varying plant that resisted optimal control. Assilian found that human operators were too erratic to imitate and digital controllers too fine-grained, and in the second half of the thesis he turned to priming the controller with a linguistic description of a good policy. Gaines had translated hedges into precise inputs; Waterman into ranges; Assilian used my 1973 paper on linguistic variables to do the translation in a principled way, defined eight linguistic variables over his quantised inputs and outputs, wrote heuristic rules from his own experience, and tuned the rules until the

engine behaved. Gaines was his external examiner in 1974 and calls the thesis "a model of scholarship." He then draws the moral that matters. Both he and Assilian had found that one cannot choose a good set of verbal instructions in advance; one must see their effect and revise. Fuzzy control succeeded because it gave a human trainer a principled way to talk to a machine and watch what happened. That is a better account of why the field took off than any I have offered, and it places the credit where it belongs, with the student, the examiner and the steam engine, rather than with the man who wrote the footnote.

I want to record one further thing, because it is easy to miss in a bibliography of several hundred items. Gaines wrote, with Kohout, a 1977 annotated bibliography of fuzzy systems research in which he and his co-author thanked Goguen for "seven pages of detailed critical comment" and then quoted Goguen's verdict that "many of the papers that I really do know the content of are badly misrepresented. I suspect that lots of the ones I don't know are too." They printed this in their own paper, as a caution to the reader. I have known many scholars. I have known few who would publish a friend's dismissal of their work beneath the work itself. It is the act of a man who thinks the reader's accuracy matters more than his own reputation, and it is why his history of our field can be trusted.

In 2012 he wrote that some of the friends on his list were "no longer with us" but that "their ideas live on and flourish as a continuing inspiration." I am now on the other side of that sentence. From here I can say what I could not say to his face: that of all the people who took the 1965 paper seriously, he was the one who took it most seriously as mathematics, least seriously as doctrine, and most generously as history. The footnote was for him. It still is.

In the manner of George A. Kelly: The scientist who built the other scientist a laboratory

I proposed, in 1955, that every person is a scientist: that each of us erects a system of constructs through which we anticipate events, tests those anticipations against what happens, and revises the system when it fails. I did not propose that the scientist should be given a computer. I died in 1967, before the question could seriously arise. It arose anyway, in a laboratory in Essex and then in Calgary, and the people who answered it were an engineer trained in experimental psychology and a mathematician trained in psychology, Brian Gaines and Mildred Shaw. I am obliged to them, and I shall try to say precisely for what.

They said of my book, twenty-five years after it appeared, that it was "seminal," that it provided "a remarkably far-reaching and well-structured foundation for epistemology," anchored "both in its close correspondence to the actual behaviour of people and in its coherent and consistent philosophy." They also said, in the next sentence, that I had not fully worked out "a logically, philosophically and psychologically complete model of knowledge acquisition," and that my attempts to link the theory to other philosophies were sketchy. Both halves are true, and I prefer a reader who says both to one who says only the first. The repertory grid, my instrument for eliciting a person's constructs, was in its original form a clinical device: a set of elements, a set of bipolar constructs, a matrix of assignments, and a clinician's eye. What Gaines and Shaw saw was that the matrix is a formal object, that the eliciting procedure is an algorithm, and that the clinician's eye could be, in part, a program.

Shaw's PEGASUS of the late 1970s was the first such program I would have recognised as faithful. It did not merely score a grid. It conducted the elicitation, and when two constructs were found to be used in the same way across the elements, it said so in plain words and asked the person for an element that would pull them apart. That is my sociality and my fragmentation corollaries turned into a dialogue. The machine held a model of the person's construing and acted so as to find out where its model was wrong. Gaines, writing with Shaw in 1980, placed this alongside the principal components analysis of Slater's INGRID, Atkin's Q-analysis and Shaw's own FOCUS algorithm, which sorts constructs so that those close in the space are close in the order and then leaves the person to construe the result. He noted something that I had insisted on and that most of my followers forgot: that the arrangement is only a re-presentation of the person's own data, and that "it is up to the user to construe meaning into the result and confirm this directly in terms of the original data." A grid analysis is a hypothesis about a person, not a finding about one. He understood that the instrument must not pretend to a privileged view.

He also saw, as a logician would, what the grid's formal structure permits. In the same paper he extended the analysis from binary assignments under two-valued logic to rating scales under the multivalued logics catalogued by Rescher. I had presented constructs as dichotomies, and I meant it philosophically: to construe is to see a likeness and a difference at once. But I knew my clinicians wanted shades of grey between the poles, and I had no principled way to give them one. Gaines did, because he had spent the previous decade on the foundations of fuzzy reasoning and could treat a five-point rating as a membership function without embarrassment. The marriage of personal construct psychology and the logic of vagueness is one I could not have arranged. It is a good one.

The next thirty years are a single story told in successive technologies. The grid tools moved from PEGASUS and PLANET to the object-oriented knowledge base that ran on the Macintosh in the 1980s, with graphical rather than numerical data entry, so that experts could drag their constructs about on a screen. The knowledge-acquisition workshops at Banff, which Gaines and Shaw began in the mid-1980s, made my grid a standard instrument of an engineering field I never heard of, expert systems, on the sound argument that an expert's knowledge is a construct system and the way to get at it is to elicit, not to interrogate. WebGrid, designed in late 1993 and in service at Calgary by mid-1994, put the elicitation on the World Wide Web within a year of the first graphical browsers, so that a person anywhere could sit down with the program and be asked, in effect, how they saw the world. Rep Plus, in 2015, is the same instrument again, sixty years after the book. I have had many disciples. I have not had many who kept the laboratory open for forty years.

Let me say what I think the deeper service was. My theory was often filed under personality or under psychotherapy, and taught as a humane alternative to the behaviourism of its day. Gaines filed it under epistemology, where I had always wanted it, and he did so by building. He treated the Fundamental Postulate as a specification. If a person's processes are psychologically channelised by the ways in which they anticipate events, then a machine that is to collaborate with a person must model those channels, and the first thing it must be able to do is ask about them. From that specification came a family of programs, a journal, a workshop series, and a way of thinking about knowledge acquisition in which the expert is not a database to be dumped but a scientist to be conversed with. He and Shaw called it computer-aided constructivism. I would have called it taking me at my word.

There is a final point that I think he would accept, since he has written so often about what he calls the voices of the mind. My theory says that a construct system is revisable without limit, that there is no final set of constructs through which the world must be seen. Gaines has held to that in his own work. He revised his view of fuzzy logic for forty years and said so. He revised his view of stochastic computing and wrote its history when the revision was complete. He built a library in which every paper he ever wrote stands open to the next one. That is the posture I asked of my clients, and he has adopted it toward himself. A scientist who construes his own career as an experiment, still running, is the best advertisement a psychologist could want.

In the manner of Bertrand Russell: On a schoolboy who was fascinated by my paradox

It is a curious experience to learn that one's *Principles of Mathematics*, a book I came to regard with the mixed feelings one has for an ambitious child, was being read with enchantment by a schoolboy in post-war England who also had a soldering iron and a copy of Hegel. Mr Gaines records that the paradox I discovered in 1901, the class of all classes that are not members of themselves, "had fascinated me since I read *Principles of Mathematics*," and that it "seemed to present a pitfall for any axiomatic system theory." I spent ten years and three volumes trying to climb out of that pit, and I am not sure I succeeded. He seems to have regarded it as a feature of the landscape and gone on building.

I should like to understand how. The clue is in his own description of what he valued. "The mathematics I valued was algebra and logic. I saw arithmetic as an over-specified algebraic structure, notions of continua and infinity as useful fictions for modelling indefinitely extensible algebraic structures, and the major mathematical modelling task as being the search for underlying generative systems that were logical, concise and complete." This is not the position of the *Principles*. It is closer to the position I was driven to in the second edition, and closer still to a kind of constructivism I never quite embraced: that the continuum and the infinite are fictions we adopt because they are useful, and that the real work is to find the smallest generative system that produces the structure we want. A man who holds that view is not frightened of a paradox. He regards it as evidence that a particular generative system has been over-specified, and looks for a weaker one.

That, I think, is exactly what he did with the logics of vagueness. I wrote a paper on vagueness in 1923, arguing that all language is vague, that the law of excluded middle applies only to an ideal and non-existent precise language, and that logic as usually conceived is therefore not true of this world. I did not pursue the matter. Gaines did. He knew from Moh Shaw-Kwei's 1954 speculation that in a set theory built on Łukasiewicz's infinite-valued logic the unrestricted axiom of comprehension might not generate my paradox at all, that the question was still open when Zadeh published fuzzy sets in 1965, and that White proved it in 1979. The practical consequence is that a logic which admits degrees of truth may be able to afford a naive set theory that two-valued logic cannot. To a man who regarded arithmetic as over-specified, this was an invitation. His *Foundations of Fuzzy Reasoning* opens by observing that "the move from a predicate to a set satisfying it, from an intension to an extension," is a natural one in everyday language, and then asks

what happens to that move when the predicate is "tall" or "red" or "stable." I asked the same question in 1923 and answered that logic did not apply. He answered that a different logic did, and set about finding which.

I must note his honesty about the result, because it is a quality I prized and did not always possess. He wrote at the end of that paper that he had given "not what are the foundations, but what might be the foundations," and attributed the shortfall to the absence, for logics of vagueness, of anything comparable to the development of the classical predicate calculus. Whitehead and I had the predicate calculus. We built on it for a decade and it still cracked. He had no such foundation and said so, and then built anyway. Twenty years later he and Kohout had a calculus of possibility, probability and eventuality; forty years later he was still describing his view of the membership function with the words "it still is." There is a stubbornness here that I recognise from the inside.

He has also been faithful to a view of mathematics that I held and that the century mostly abandoned: that it is a way of understanding the world, and not an autonomous game. He wrote that "logic and mathematics were ways of modelling and understanding the world, and electronics a way of influencing and controlling it," and his career kept the two in a single hand. The stochastic computer is the clearest instance. It represents a real number by the probability that a pulse is present, so that the hardware's logic gates compute with probabilities directly. This is a mathematician's idea, since it rests on seeing that multiplication of probabilities and conjunction of independent events are the same operation, and it is an engineer's idea, since it was patented with fifty-four claims and built on a bench at Harlow in 1965. I spent my life insisting that a proposition about the world and a proposition in logic were different in kind. He has spent his demonstrating that a sufficiently good model of the world can be run as a circuit.

One reservation, since he would expect it. He has written approvingly of Hegel's *Science of Logic* and sees the world as "hierarchies of feedback loops evolving through temporary equilibria punctuated by periods of chaos." I spent some years recovering from Hegel and I do not recommend him. But I observe that Gaines's Hegelianism, if it is one, is of a tolerable kind: a conviction that structures develop by resolving their own contradictions, which he has applied not to the Absolute but to the history of computing, to the emergence of fuzzy control from the failures of optimal control, and to his own successive views. He later found in Popper's three worlds a less mystical frame for the same conviction. He appears to have taken from Hegel only the dialectic and from me only the paradox, and to have left both of us our metaphysics. This is, on reflection, the correct way to read philosophers.

I called my own early book a child. Mr Gaines read it as a boy, took from it a question I could not answer, and has been answering it, in hardware and in logic, for sixty years. An author can ask no more of a reader, and few get as much.

In the manner of Karl Popper: Engines of the third world

In 1978, at the IEEE conference on cybernetics and society, a professor from Essex argued that computers "have a peculiarly significant role in world three that is comparable in importance to that of engines in world one." I had distinguished three worlds: the world of physical things, the world of subjective experience, and the world of statements in themselves, of theories, problems and arguments, which I held to be real, autonomous and largely man-made. I had said that books and journals and letters were typically third-world objects. Mr Gaines took the thought one step further than I had, and I have come to think the step was correct. The steam engine did not merely move things in world one; it changed what kinds of things could exist there. He proposed that the computer would do the same for world three. Forty-eight years later, writing from inside a library in which every one of his papers is linked to every other and a reader can ask the collection questions, I am in no position to disagree.

I want to draw attention to the argument's structure, because it is Popperian in a way that goes beyond borrowing my vocabulary. He opened the paper not with the promise of computers but with the two standing objections to them: the scepticism of those who had been promised management information systems and received invoices, and the fear of those who foresaw a state that "knows too much, and through its knowledge controls too much." He named Dreyfus and Weizenbaum as the serious critics and took them seriously. Only then did he make his case. This is how one ought to argue: state the strongest objection first, in its strongest form, and let the thesis survive it if it can. Much of what was written about computers in that decade did the opposite, and has not aged well.

He did the same thing to himself. In a 1978 paper on decision he quoted my remark that we must "replace routine more and more by a critical approach," and then considered the hypothesis that there is one correct foundation for decision, which he labelled ONE, alongside the hypothesis that there is none, which he labelled ZERO, and observed that ONE "has its dynamics and its dangers fully equal to those of ZERO." The man who had spent the previous decade seeking a single algebraic foundation for the logics of uncertainty here warned against the belief that any single foundation exists. I take this to be the mark of someone who has actually understood falsificationism rather than merely endorsed it. One applies it first to one's own programme.

His studies of the growth of knowledge are where I think his contribution lies, and they are not well enough known. From the mid-1980s he developed a model of how an information technology matures, through stages he called breakthrough, replication, empiricism, theory, automation and maturity, and he applied it to successive generations of computing, from the tube-based calculators of 1940 through the empirical designs of the late 1950s, the automata theory and logic minimisation that followed, the computer-aided design tools, and the microprocessors of 1971 and 1972. He observed that the computer industry's learning curve was unique because each advance in computing supported further advances in computing, a positive feedback loop he compared to the emergence of new phenomena in biology. What interests me is the epistemology beneath the economics. Theory, on his account, comes after practice and after replication; it is a late stage, not a first one. This is my view of science put to work on technology. Conjectures are built before they are understood; understanding is itself a conjecture about why the building stood up; and the history of a field is the history of its refutations, which he was careful to include. The stochastic computer was one. He and Poppelbaum conceived it, built it, and concluded honestly that the technology of their era could not make it competitive. He then wrote, in his presidential address to the systems society, that it had "still to find its place." It did, forty years later. A refuted conjecture is not a dead one; it is a conjecture awaiting a world in which it is true.

He has been scrupulous about the sociology of knowledge in a way I admire and did not always practise. His 2012 memoir of the origins of fuzzy control reads like a case study in how a scientific community actually works: a small group in Britain who knew each other from colloquia, a journal that guaranteed a decision within six weeks, a thesis examined by the man whose work it extended, a student who tuned rules by hand until a steam engine behaved. He draws no grand conclusion from it. He simply records that this is how it happened, and that "there have been studies analyzing their history and the basis of their success; this article provides some details from my personal experience." The historian of science could ask for nothing better than a participant who distinguishes his memory from his analysis and labels both.

I will register the disagreement I think he would want registered. In my later years I held that world three, once created, has a life of its own, and that its objects, theories and problems, can be discovered as a mathematician discovers a theorem. Gaines, who read my work through Kelly and Wittgenstein as well as through me, inclines to the view that knowledge is constructed and negotiated, and that the problems in world three are the ones we have put there. His tools for knowledge acquisition assume it: they elicit, they do not discover. I think he is half right. A construct system is built, but the consequences of a construct system are not chosen, and it is in the gap between what we build and what

follows from it that world three has its autonomy. His own career is my evidence. He built a stochastic computer, and sixty years later its consequences are still being discovered by people who never knew him. That is a third-world object behaving exactly as I said it would.

He wrote in 1978 that we had promised "much, perhaps too much, perhaps too little" on behalf of computers. On behalf of the third world I will say that this one has delivered.

In the manner of Joseph A. Goguen: Seven pages of criticism, and what they were for

In 1977 Brian Gaines and Ladislav Kohout published a large annotated bibliography of fuzzy systems research, and in the acknowledgements they thanked me for seven pages of detailed critical comment, said it had improved the work "greatly, but probably not enough," and then quoted me: "Many of the papers that I really do know the content of are badly misrepresented. I suspect that lots of the ones I don't know are too." They added a moral for the reader. I have been quoted many times in my life and this remains the quotation of which I am proudest, not because of what I said but because of who chose to print it. A man who puts his critic's verdict at the head of his own work has understood what criticism is for. It is not a judgement on the author. It is a service to the reader, and the author who passes it on is performing the same service.

We were drawn to the same footnote in Zadeh's 1965 paper, the one allowing membership values in a partially ordered set rather than the unit interval, and we went in different directions from it. I went to lattices and then to categories: L-fuzzy sets, the logic of inexact concepts, and eventually the view that the right setting for systems theory is a closed category in which realisation is a universal construction. Gaines went to the logics of uncertainty and their algebra, to Łukasiewicz and Rescher and Carnap, and to a calculus in which possibility, probability and eventuality are species of a single genus. We agreed on the fundamental thing, which is that the interval [0,1] is an accident of notation and that the structure of inexactness is algebraic. He said in 2012 that this still seemed right to him. It still seems right to me, from wherever I am now writing.

He was generous to my work in a way that I did not always deserve and did not always return. In his 1978 address to the general systems society he quoted a remark of mine from 1974 that stochastic and non-deterministic machines "seem to be poorly endowed with pleasing theoretical results," and then observed, with his customary absence of resentment, that the situation had changed with the application of complexity theory to system identification, and that we had all been "trapped by the statistician's view of

probability theory" while the "very much more fruitful logical and algebraic foundations were insufficiently known." He did not say that I had been wrong to find the results unpleasing. He said the results had since become more pleasing, and pointed to the work that made them so, some of it his own. His 1977 paper on system identification, which proves that if deterministic models are required then nearly all behaviours require models of nearly maximum complexity, is exactly the sort of theorem I had complained did not exist. He proved it and then cited my complaint as motivation. There is a kind of intellectual courtesy here that is rarer than brilliance.

I want to say something about his mathematics as mathematics, since the Festschrift of a man who has done so much else risks treating it as an early phase. It was not a phase. His 1977 formulation of identification postulates two order relations on a class of models, a fixed one of complexity and a variable one of approximation induced by the observed behaviour, and defines an admissible model as one such that any less complex model is a worse approximation. This is a Pareto frontier in a space of theories, and it is the correct abstract setting for the whole problem of induction as an engineer meets it. He then develops the approximation ordering from subjective probability and semantic information theory, and treats causality, locality and finite memory as what he calls "rules of the game," structural constraints the modeller imposes rather than facts the data reveal. A category theorist reads this and sees a functor from behaviours to admissible subspaces, and I suspect he saw it too and chose not to say so, because his audience was engineers and he wanted them to use it. That was always the difference between us. I wrote for the people who would appreciate the construction. He wrote for the people who would build with it.

The building is the point. He designed a minicomputer and founded a company to sell it, in 1968, while a Reader at Essex and executive editor of a journal he had just started. He built the stochastic computer at Harlow with Peter Joyce in 1965, and an analog computer at Cambridge in 1964 to stand in for a flight simulator, and an oscilloscope with two cathode ray tubes to show each eye a different rotating hypercube. Later the building was software: grid tools on the Macintosh, a server on the Web within a year of the Web, a knowledge representation system in source form for the research community. I built almost nothing, and I say so without regret, because there are two ways to serve a mathematics and he chose the one I could not. The theorems I proved about L-fuzzy sets are true in any possible world. The steam engine Assilian controlled with fuzzy rules in 1974 ran in this one, and Gaines examined the thesis and wrote its history.

Let me end with the seven pages. I wrote them because the bibliography was going to be used, and I knew that an annotation that misrepresented a paper would misdirect a hundred graduate students. He printed my verdict because he knew the same thing. We were both, in that moment, doing what the critical tradition asks: treating the state of the literature as more important than the standing of its authors. Everything I most admire in his later work follows from that moment. The knowledge-acquisition tools that refuse to pretend to a privileged view of the expert; the retrospective essays that distinguish memory from analysis; the library that puts every paper, including the ones he would now revise, in public view with its date on it. He has spent fifty years making it easy for others to find out where he was wrong. That is the only kind of rigour that matters, and he has it.

Meetings with remarkable people: recollections from the laboratory

Gaines wrote that he had been tempted to title his 2012 memoir, after Gurdjieff, *Meetings with Remarkable People*. The editors have taken the hint and gathered, in the manner of a Festschrift's shorter tributes, imagined recollections from the colleagues he named. Each is grounded in what he himself recorded of the meeting. Each is, like the essays above, written for him rather than by them.

Richard Gregory, Cambridge. He arrived in 1959 to read mathematics at Trinity, having spent a year at ITT's semiconductor laboratories at Footscray, and he brought with him a box of devices he had fabricated and an inquiry as to who might find them useful. That is how a mathematics undergraduate became the electronics technician of a cognitive psychology laboratory. I have never had a better one. When he graduated I suggested he stay for a doctorate; Oliver Zangwill agreed on condition that he also take a psychology degree, which he did in a year, preparing by writing past examination papers for his tutor and acquiring the background from journal articles. His thesis, funded by the Ministry of Defence, studied how pilots from RAF Oakington learned to control an unstable system, and he built the analog computer that emulated the aircraft. The photograph of him beside it in 1964 shows a very young man and a very large machine, and it is not clear which is in charge. In my own laboratory he built a stereoscope from two small cathode ray tubes and let subjects rotate four-dimensional hypercubes, one image to each eye, in order to think about how disparity becomes depth. The stochastic computer came out of that thought. I have had many students who went on to do important work. I had one who did it in my laboratory, on my bench, with my oscilloscope, and told me about it afterwards.

John Andreae, STL and Canterbury. I advertised in *Nature* for a learning-machines project at Standard Telecommunication Laboratories and specified topology, cybernetics and neural networks, and an undergraduate wrote to me. He spent his vacations with us at Harlow, with David Hill, Owen Morgan and Peter Joyce, and when I left for the chair at Canterbury in New Zealand he took over my position part-time while finishing a doctorate in a different subject at a different university. STeLLA, our learning automaton, encoded the state space of the plant and learned to drive it into a target region; Brian's contribution was a front-end network to learn better encodings of the inputs, which I named Gadafter, the Gaines adaptive filter, and the stochastic integrators that made its weight changes. Our paper for the 1966 IFAC Congress ended with a sentence he has since called prophetic: that when the plant is indeterminate and the controller must be widely applicable, optimality "loses much of its force," and "one can only ask whether satisfactory control is possible and, if so, how much it will cost." He says this captured the ethos in which he, Igor Aleksander, Abe Mamdani and I worked. It did. It also describes how he has run his career.

Ted Poppelbaum, Illinois. My student Sergio Ribeiro and I had conceived of representing a number by the probability of a pulse, for pattern recognition, and called it a noise computer. A patent application with fifty-four claims arrived from England describing the same idea, for learning machines, under the better name of stochastic computer, and we adopted the name. He visited us at Urbana in 1967. We remained friends and exchanged papers and reports until I died in 1994, and in 2019 he wrote the history of what we had both done, giving my group its full and independent credit. Both of us had concluded in the 1960s that the technology would not support a competitive machine, and both of us moved on. He wrote in 1978 that the theory had "still to find its place in general systems." It has found one now, in a world of unreliable nanoscale devices for which a computer that runs on noise is exactly right. I am sorry not to have seen it. I am glad he did, and glad he was the one to write it down.

Ladislav Kohout, Essex. I came to Brian as a graduate student in 1968, the year he arrived at Essex as a Reader and founded two companies and a journal, and I referred him to Rescher's proof that probabilistic logic is equivalent to the modal system S5. He took it, as he takes everything, as the beginning of an algebra, and the result over the following years was a calculus of possibility, probability and eventuality that contained fuzzy logic as a case, and a theory of possibilistic automata that we published together in 1975. We also wrote the annotated bibliography that Goguen criticised in seven pages, and we printed the criticism. Working with him was an education in a particular virtue: he never once treated a result I brought him as mine or his. It was the field's, and the only question was whether it was right.

George Klir, Binghamton. We were both presidents of the Society for General Systems Research, and his address to it in 1978 is the one I would give to a student who wanted to know what the systems movement was for. He ranged from automata theory to the foundations of probability to the sociology of positivism, confessed that the stochastic computer had not yet found its place, and argued that we had all been trapped by the statistician's probability when the logical and algebraic foundations were richer. He drew on my epistemological hierarchy of systems in his later work on expert knowledge and credited it. He was, for forty years, the colleague who read everything and quoted it accurately, which in a field as diffuse as ours is the scarcest skill of all.

Abe Mamdani, Queen Mary College. I knew Brian in the 1960s as the man working on stochastic adaptive modules while I worked with Igor Aleksander on deterministic ones. We published in *Electronics Letters* because it decided within six weeks, and I suspect we refereed each other's papers. I had a copy of his thesis and had sat through his colloquium presentations on priming learning controllers with words, and when I got my lectureship in 1971 I set Sedrak Assilian the task of reproducing the results on a real plant, a small steam engine. The second half of Sedrak's thesis is where fuzzy control begins. Brian examined it in 1974 and later wrote that it was "a model of scholarship," and that its achievement was "providing the human trainer with a means to communicate linguistically with a control system and to monitor the impact of specific instructions on the performance in order to improve it." That is exactly what it was. Most accounts of the origin of fuzzy control begin with Zadeh and end with us. His begins with Luria's work on verbal behaviour in Cambridge in the early 1960s, passes through his own pilots and perceptrons, and gives the student the credit. It is the only account by a participant that is also history.

The Banff and Calgary years. The editors add one collective voice, that of the knowledge-acquisition community he and Mildred Shaw convened from the mid-1980s. We came to the workshops at Banff from AI laboratories and psychology departments and industrial research groups, and we found there a pair of hosts who had already built what we were arguing about: grid elicitation tools, an object-oriented knowledge base on the Macintosh with graphical data entry, and later WebGrid on the Web. The Knowledge Science Institute at Calgary, under his Killam Memorial Research Professorship, was the institutional form of the same hospitality. He founded the *Knowledge Acquisition* journal as he had founded the *International Journal of Man-Machine Studies*, to give a home to work that belonged nowhere else. The field has since been absorbed into others, as fields are. The habit he taught it, of treating the expert as a scientist to be conversed with rather than a store to be emptied, has not been.

Coda: to the reader of fifty years hence

This book is addressed to you, because its subject spent his career arguing that you would exist and preparing a library for you to read. In 1993 he set out an agenda for digital journals: that scholarly communication would move onto a networked, hyperlinked, continuously revisable infrastructure, that the economics of paper would cease to govern what could be said and how often it could be corrected, and that the unit of scholarship would become the living document rather than the printed page. He then did the unusual thing of acting on his own forecast. Every paper discussed in this volume, from the 1966 learning-machine paper with Andreae to the 2019 history of stochastic computing, is in the library you are presumably reading this from, with its date on it, open to the next reader's questions.

You will be able to do what we could only begin to do here. The editors of this volume established its contributor list by asking that library, across roughly 37,000 passages of his writing, whose names recur, and found Kelly, Zadeh, Popper and Pask at the top, with Ashby, Russell, Łukasiewicz and Rescher behind them, and then confirmed the list against his own words in the 2012 memoir. You will have better instruments. You will also have, if his model of how information technologies mature is right, a different relation to the whole enterprise, since the breakthrough, replication and empiricism of machine-assisted scholarship were ours, and the theory, automation and maturity are yours. We ask only that you read the essays above as what they are: imagined voices, constrained by the record, composed in 2026 by a machine at the request of a reader of his library who wanted to know who would have written his Festschrift. The machine is itself a late consequence of the learning-machines project at Harlow in the early 1960s, which is a fact he would have enjoyed and would have insisted on footnoting.

He wrote in 1998, of Gordon Pask, that a person's voice goes on participating in the conversations of the living, and that "our mental processes subsist through collective social processes" in which the absent are reconstructed and heard. He and Shaw called this the voices of the mind, after Wertsch. We have taken him at his word and given eight of his voices a few pages each. The exercise has a Gainesian shape: it is an elicitation, not an extraction; the model it produces is a hypothesis about a person, to be confirmed against the original data; and the original data is one click away.

What the voices agree on, across their disagreements, is worth stating plainly, since he never would. He held that logic and electronics were one subject. He regarded any single foundation, including his own, as a conjecture to be tested. He built first and theorised afterward and said so. He gave students, co-authors and critics their full credit, including the ones who said he was wrong. He kept every instrument he ever made in working order,

from the grid tools of 1977 to their web descendants, and he kept his own changes of mind on the record. He treated the expert, the pilot, the perceptron and the reader as participants in a conversation rather than as sources to be drained. And he thought, from his schooldays onward, that the world was a hierarchy of feedback loops finding temporary equilibria, and he never found a reason to stop thinking so.

If you are reading this in the library he built, you are inside the argument of his 1978 paper: a third-world object, a theory about a man, running on an engine he foresaw. He would want you to check the sources. They follow.

Sources

All quotations attributed to Gaines in this volume are taken from the following papers in the Gaines library at the University of Victoria. The attributed opinions of the contributors are imagined; the facts about Gaines are not.

Essay	Gaines paper drawn on	Year
Preface, Ashby, Zadeh, Russell, Goguen, Recollections	Living in an Uncertain Universe, chapter for the Springer homage to Zadeh	2012
Pask, Coda	Conversation Theory in Context, with M. L. G. Shaw, memorial issue for Pask	1998
Popper	Computers in World Three, Proc. Int. Conf. on Cybernetics and Society	1978
Popper	Paper on decision-making quoting Popper's "critical approach"	1978
Popper, Goguen, Recollections	Progress in General Systems Research, presidential address, SGSR	1978
Kelly	New directions in the analysis and interactive elicitation of personal construct systems, with Shaw, IJMMS 13	1980
Zadeh, Russell	Foundations of fuzzy reasoning, IJMMS 8	1976

Essay	Gaines paper drawn on	Year
Zadeh, Goguen	Fuzzy systems bibliography and survey, with L. J. Kohout	1977
Goguen	System identification, approximation and complexity	1977
Ashby, Recollections	Origins of Stochastic Computing, in Gross and Gaudet (eds.), Springer	2019
Ashby, Recollections	Doctoral thesis, Cambridge and Teaching machines for perceptual-motor skills	1966 to 1971
Kelly, Recollections	Eliciting knowledge and transferring it to a knowledge-based system, with Shaw, IEEE TKDE, including the biographical notes	1993
Kelly	Rep Plus history page, WebGrid from 1994 to Rep Plus 2015	2015
Popper, Coda	Convergence to the information highway, the breakthrough-to-maturity model	1990s
Coda	An Agenda for Digital Journals	1993
Recollections	Epistemological hierarchy and expert systems, drawing on Klir	1986

The contributor list was checked against the library itself on 2026-10-08. Of 37,133 indexed passages, the number of distinct documents mentioning each name was: Kelly over 751 (the query capped at a thousand passages), Zadeh 507, Popper 413, Pask 391, Goguen 281, Klir 263, Russell 199, Mamdani 195, Andreae 179, Rescher 176, Wiener 155, Łukasiewicz 151, Kuhn 140, Wittgenstein 123, Ashby 122, Peirce 91. Shaw, as co-author, appears in over 801. Mentions are a proxy for engagement, not esteem; the 2012 memoir supplied the esteem.

Facts about the contributors themselves (Ashby's 1952 and 1956 books, Pask's SAKI and his 1975 conversation theory, Zadeh's 1965 and 1973 papers, Kelly's 1955 book and 1967 death, Russell's 1901 paradox and 1923 paper on vagueness, Popper's three worlds, Goguen's L-fuzzy sets and categorical systems theory) are matters of public record and are not sourced from the library.

This volume was composed by Claude, an AI model, on 2026-10-08, at the request of a reader of the Gaines library, in answer to the question of who would have written a Festschrift for Brian R. Gaines if the people he thought most highly of had written it.