

Remembering Bill Mann

Christian M. I. M. Matthiessen*
Macquarie University

William C. Mann passed away on August 13, 2004, after a long struggle with leukemia. He is known to readers of *Computational Linguistics* as a keen supporter and past president (1987) of ACL, a pioneer in the development of text generation as a field of research in computational linguistics, the originator of rhetorical structure theory (RST) and dialogue game theory (DGT), and the developer (together with David Weber) of dialect adaptation as a technique within machine translation. And he'll be remembered for these and many other contributions, but above all, in this context, he'll be remembered as a unique visionary in computational linguistics (CL) and as a truly wonderful and extraordinarily generous colleague.

Bill Mann made huge contributions to many people's intellectual development, to their careers, and to their lives in general. What follows are my own recollections of him, selected from a crowded jungle of memories. I will leave it to others to fill out the picture, but I shall attempt to give an indication of the intellectual and historical context in which Bill was working and also of the network of researchers that he created; his ideas have radiated throughout this network of scholars around the world.

I met Bill in April 1980 and worked as a research linguist on a succession of projects he directed, or codirected, at the Information Sciences Institute (ISI) at the University of Southern California (USC), until I left the institute to take up a position at the University of Sydney in August 1988. Through that period (which happened to coincide, more or less, with the Reagan years), not only was Bill a very wise and wonderful project leader, but he also became a mentor and guide and a true friend. I followed from afar the further developments in his life: his departure from ISI and move to Kenya with his family to take up a position with the Summer Institute of Linguistics (SIL) as a consultant responsible for work on discourse and computational matters, their return to the United States, and his new phase of research with initiatives in RST and DGT.

Phase I: Electrical Engineering

When I first met Bill, he was already in the third phase of his productive career and intellectual journey of discovery and innovation. The first phase was launched with a degree in electrical engineering, and this was, I think, an important aspect of his enthusiasm for modeling and for problem solving and innovation. After receiving this degree, he did research in electrical engineering and learned never to do classified research again. On starting a project involving classified work, he did a survey of his field of research and learned about another researcher doing interesting work. He tried to identify this person in order to learn from his or her findings, but because this work was also classified, he was never able to—until, on completing his own project, he discovered that this other researcher was Bill himself! He realized the high cost of secrecy in research: the great loss

* Department of Linguistics, Macquarie University, NSW 2109, Australia.
E-mail: cmatthei@ling.mq.edu.au.

of information and also the lack of the benefits of peer review. His own approach came to be just the opposite: He championed the development of reusable, shareable resources.

Phase II: Artificial Intelligence

The second phase started when Bill went back to school to get a Ph.D. in artificial intelligence (AI)/computer science at Carnegie Mellon University. His thesis advisors were two of the founding fathers of AI and contributors to cognitive science, Herbert A. Simon (1916–2001) and Allen Newell (1927–1992). Bill's post-Ph.D. work at ISI did not pursue this AI research directly, but he made good use of AI techniques and tools, including goal pursuit, planning models, and knowledge representation. It was at some time during this second phase, I think, that Bill came across and read the later work by Ludwig Wittgenstein. It had a profound influence on him, as it has had on a number of other scholars coming from a logical and philosophical background rather than a rhetorical and anthropological one. Bill came to give language a much more central place than it had been given in the mainstream of AI and cognitive science, characterizing it as the most powerful human resource for knowledge representation and emphasizing the centrality of communication. Inspired by Wittgenstein, he developed his dialogue game theory, concerning exchanges in dialogue, in the late 1970s (Mann 1979). But this already takes us into Phase III.

Phase III: Computational Linguistics

The third phase started in the mid-1970s with Bill's position as a computational linguistics researcher at ISI, a nonprofit research institute founded by Keith Uncapher (1922–2002) in 1972 (together with a small group of people who'd all been at RAND) in association with the University of Southern California. By the time I joined ISI in 1980, it had developed into a major research center. It was a very stimulating and exciting research environment, with cutting-edge research into VLSI design, network protocols, and so on. It played a central role in the development of ARPANET and its successor, the Internet, with key contributions by Jon Postel (1948–1998), a brilliant, soft-spoken researcher at ISI.

Bill was part of the generation of researchers who transformed ISI into a major research center, contributing a minor but significant strand of computational linguistics research to the major research strands within ISI. By the end of the 1970s, he had, as already mentioned, developed dialogue game theory. He had also, together with Jim Moore, developed an early text generation system, the Knowledge Delivery System (KDS). This experimental system played a significant role in Bill's subsequent major work on text generation. And outside the framework of his ISI research activities, he had started his work together with linguist David Weber on computer-assisted dialect adaptation (e.g., Weber and Mann 1979). This work was to be applied to dialects of Quechua (Kasper and Weber 1986), a language for which Weber was to produce a major reference grammar (Weber 1989). This activity was linked to the Summer Institute of Linguistics, a connection that was to become much more central in Bill's Phase IV work in the 1990s.

Into Text Generation

Around 1980, Bill was preparing to start a major new research effort in text generation with funding from the Air Force Office of Scientific Research. At the time, text

generation was not a developed part of the computational linguistics research spectrum; the main concern in computational linguistics was with parsing and text understanding and forms of representation needed to support these (like transition networks for grammars and knowledge representation for semantics). Machine translation research had had to retreat in the United States and had also been affected negatively in other places after the disastrous ALPAC report in the mid-1960s and was only just beginning to make a comeback.

After various debates in the 1970s, computational linguists had clarified a number of issues fundamental to the modeling of language and to the processing of text. One key issue was the nature of knowledge representation: in the choice between procedural approaches and declarative ones, there was a growing consensus in favor of declarative forms of representation, since such representations could be used by processes of different kinds. This was important to the development of knowledge bases for text generation and also for the move away from representations of grammar with a built-in procedural orientation to parsing (ATNs) toward declarative representations allowing specifications of grammars to serve as resources in both parsing and generation—a move pioneered by Kay (1979) with his functional unification grammar. There were other advances in knowledge representation. After Woods's (1975) critique of the early free-for-all in work with semantic networks, the epistemological status of such networks had been clarified (e.g., Brachman 1979), but the complementarity of frame-based inheritance networks and logic-based representations had not yet been worked out.

At the same time, there were unresolved questions—centrally, the contrast between generalists and particularists (advocates of special-case solutions). This contrast had at least two independent, but often linked, manifestations: a higher-level one and a lower-level one. At the higher level, generalists were oriented toward logical resources, relying on reasoning from first principles (this being seen as the investigation of general intelligence launched by the founders of AI), whereas the particularists were oriented toward experiential resources, relying on repertoires of schemata or scripts representing chunks of experience. At the lower level, generalists favored grammatical approaches to natural language processing, whereas particularists favored lexical approaches (going back to Becker [1975]).

Bill had picked his own way through these various issues: He favored declarative forms of representation, choosing Brachman's KL-ONE as the first knowledge representation system (later to be replaced by NIKL and still later by LOOM) and Halliday's systemic functional grammar, and he tended to favor generalist solutions. He also anticipated certain later developments in NLP: He invested in the creation of large-scale reusable linguistic resources and based the computational design on the investigation of naturally occurring texts.

In the early development of text generation, there were three long-term initiatives that played a significant role: Dave McDonald's MUMBLE system, Kathy McKeown's TEXT system, and Bill's Penman system. Text generation grew significantly as a field of research during the 1980s with a succession of biannual generation workshops (one of which Bill helped organize in a very memorable location: the USC Marine Biology Institute on Catalina Island, in 1988). From his experience with KDS, Bill drew the conclusion that the generation system to be built would need a large-scale, linguistically motivated grammar and consequently also linguists on the research team. This conclusion was supported by the success of Davey's (1978) Proteus text generator and Bill's analysis of the first 20 years or so of CL. He had asked himself why progress was relatively slow in CL and, surveying the field, he concluded that two related problems were (1) that each new project tended to have to start from scratch, because there

was little effort given to developing reusable general-purpose resources, and (2) that linguistically motivated accounts did not inform CL enough.

The challenge was finding such a large-scale, linguistically motivated grammar for the new generation system to be developed. At the time, most of the grammars in CL had been designed for parsing and text understanding. These were not reversible and not well adapted to the meaning-oriented view of grammar needed in generation, so with David Weber's help, Bill surveyed the state of the art in linguistics, looking for an approach to grammar that was meaning-oriented and thus adaptable to text generation and that came with a reasonably comprehensive description of the grammar of English. Comprehensiveness was not part of the agenda of Chomsky's generative grammar, at that time the dominant type of grammar in linguistics. Generative semantics appeared to be of potential interest, but no comprehensive descriptions were available (and it was going out of fashion in any case).

Bill Mann and Michael Halliday: Nigel

The one candidate that Bill judged as holding out great promise for his project was Michael Halliday's systemic functional grammar (SFG): a judgment which was also supported by the research systems using SFG developed by Winograd (1972), who had studied with Halliday in London in the 1960s, and Davey (1978). So with David Weber's help, Bill advertised for a research linguist who had expertise in the area of SFG. I was lucky enough to see a copy of this ad on the notice board in the UCLA Linguistics Department, where I was on a one-year scholarship in 1979–1980 from Lund University in Sweden. Michael Halliday and Ruqaiya Hasan were at the time on sabbatical at Stanford University; by the time I had seen and responded to the ad, they had moved down to UC Irvine, where Halliday had been invited by Nick Colby and was giving a 10-week course on SFG. I had been attending the weekly lectures for a while, spending four hours on the bus in each direction between Los Angeles and Irvine. When I first talked with Bill on the phone, I asked him if he knew that Halliday was "in town." He didn't but was delighted to hear the news and suggested we drive down together to attend the remaining lectures. So my long bus journeys were replaced by wonderful rides with Bill—the first of many memorable trips we undertook together in North America, Europe, and Australia. Halliday's lectures convinced Bill he'd made the right choice in terms of linguistic theory and description.

The end result was that Bill asked Michael Halliday to serve as a consultant on the new project and hired me as a research linguist. Michael was impressed with Bill's conception of language in relation to computing and with linguistics in relation to computer science and accepted Bill's invitation. He wrote later that, against the background of his experience with computing in the 1950s and 1960s, he had asked Bill at the time of this early discussion: "How much would I have to 'simplify' (that is, distort) the grammar in order to make it computable?" Bill answered: "If I can't compute your grammar, I'll have to learn how," later adding, "I don't always understand why linguists describe things the way they do; but I have come to realize they always have a reason." This was a distinctive feature of Bill's scholarship: He didn't let the current versions of computational representations and implementations constrain linguistic theory or his use of it, and he was prepared to work as hard to understand functional linguists as AI researchers developing medical expert systems were to understand doctors.

Nick Colby generously gave us a flying start. As an innovative anthropologist, he had arrived at the conclusion that since text provided a key window on culture (an insight systemic functional linguistics had inherited from the anthropologist Bronislaw

Malinowski [1884–1942]), it was essential to develop a computational text analysis system; so as a step in this direction, he hired a computer expert, Mark James, to work with Michael on the development of a computational representation of a systemic functional grammar of English. James developed the representation, and Michael specified an initial clause grammar of around 80 systems. Colby passed this on to Bill's project. Bill guided the further development of this grammar, transforming it into the well-known Nigel grammar, a key component of the Penman generator he developed at ISI. He organized a mini-workshop on SFG and computation at ISI; one of the participants was Martin Kay, who was working on his functional unification grammar (FUG) at the time. This was just at the beginning of the major expansion of unification-based grammars pioneered by Kay, an expansion that was to include Ivan Sag and Carl Pollard's HPSG and Aravind Joshi's TAG.

By the time I left ISI in 1988, the Nigel grammar had grown to close to 1,000 systems (the documentation of which grew into Matthiessen [1995]). It is now part of the Penman system and John Bateman's freely available KPML system. In this way, Bill's vision of large-scale reusable resources has been realized.

Nigel Wallpaper and Penman Drama

Bill's creative thinking about Penman and the Nigel grammar also included the way they were made accessible and presented. As the Nigel grammar grew in size, he recognized that it would be extremely valuable to have a graphic representation of the full network of systems of which the grammar consisted, thus anticipating the emphasis that was to be placed on visualization in later research. In those days (around 1981), creating such a graphic representation was of course a major challenge; we did not yet have the Xerox and SUN workstations that were to become part of our environment later on. So Bill set to work, guiding one of the team members, Yasutomo Fukumochi (a linguist from UCLA with programming skills), in the development of a layout program for system networks. After several months of hard work, this tool was ready to use. We plotted the full system network using a Techtronix plotter that could handle sheets of about 70 × 100 centimeters, baby-sitting the plotter as it slowly and laboriously drew all the lines that represented the relations making up the system network.

The whole network covered quite a few sheets, and we had to stick them together to create a huge wall chart of the grammar, roughly about 2.5 × 5 meters in size. The problem now was where to display this chart, so Bill managed to convince the person in charge of office space at ISI that we needed a large new office for the chart. The result was truly impressive, and visitors to the project would be fascinated by the "grammar on the wall." As far as I know, this was the first time in the long human history of describing grammars of different languages that the systemic organization underlying grammar had ever been visualized on such a scale, making it possible to see the whole description as far as it had been taken. It became an invaluable map that we consulted in our further development of the Nigel grammar. Bill and I even took the wall chart with us to an international systemic functional congress held at York University in Toronto in 1982. The only place with enough space to mount it was a large hall where it had to cover a fireplace. After that time, the grammar on the wall showed a faint outline of a fireplace!

American Quilt: Long-Term Research Achievements

At ISI, Bill managed to create a long-term research effort out of a succession of projects with different funding sources. The funding sources were military ones, and there were

different views of the ethics of military funding at the time. Noam Chomsky's view was that it didn't matter where the research funding came from; what was important was what researchers did with their discretionary time. Terry Winograd's view was, as I understood it, that it did matter that funding came from the military, because this strengthened the military channels. Bill's view was that since the military was not well enough organized to make use of this kind of research, the research would benefit the public domain long before the military could do anything with it—and I think as far as our research was concerned, he turned out to be absolutely right. The projects had different specific goals, and they involved various domains (an electronic mail and calendar system, the U.S. 6th fleet, and so on), but Bill was able to keep the general foundation in focus, expanding it quite systematically according to a theoretically and empirically motivated model of language, adding components to the general Penman architecture (Mann 1982) with the Nigel grammar (Matthiessen 1981, 1983) that included RST (Mann and Thompson 1988), an RST text planner (Hovy 1988), inquiry semantics (Mann 1983), the "upper model" (Bateman et al. 1990), and an interface to applications using Penman and SPL (Kasper 1989). (Bill had formulated a moral principle about authorship: Contrary to the common practice in many places, his name would only be included among the list of authors of a publication if he had taken part in the writing, not simply by virtue of the fact that he was project leader. A number of publications cited here are thus heavily indebted to Bill, even though this is not indicated by the list of authors.)

The systemic-functional basis of Bill's research into text generation also inspired and influenced other research into text generation, including, around the mid-1980s, Terry Patten's SLANG generation system, John Bateman's research into Japanese text generation at Kyoto University, Robin Fawcett's development of the GENESYS generation system at Cardiff University and, toward the end of the 1980s and into the 1990s, new research projects using the Penman generator (e.g., TechDoc [Rösner and Stede 1994]; Healthdoc [DiMarco et al. 1995]; Drafter [Paris et al. 1995]) or taking it as a model but extending it (Cross 1991).

Bill Mann and Sandy Thompson: RST

The research team that Bill led at ISI in 1980 consisted of Jim Moore, Steve Klein, and me, with Michael Halliday as a consultant (and Jane Robinson as a consultant on another project being completed at the time). Over the years the team would contract (both Jim and Steve left ISI) and expand. At times, it was just Bill and me, but most of the time, the team included several researchers.

One key participant was Sandy Thompson, professor of linguistics first at UCLA and then at UC Santa Barbara. Sandy and Bill developed a wonderful friendship and working relationship, with Sandy as a consultant working on our project. The three of us started developing what was to become rhetorical structure theory around 1982, surveying the research on text organization that had been done at the time in computational linguistics (e.g., J. Hobbs), tagmemic linguistics and related work (e.g., R. Longacre, J. Grimes), systemic functional linguistics (e.g., M. A. K. Halliday, R. Hasan, J. R. Martin) and other approaches, and analyzing texts from different text types. Bill was very keen to develop a general account of text organization that could be used in a goal-pursuit model of text planning, which is reflected in the definitions of RST relations; the actual RST planner was developed only later, by Ed Hovy (1988).

Two other linguists, both of whom were Sandy's Ph.D. students, made substantial contributions at the time, becoming part of Bill's "research family" at ISI: Barbara Fox,

who used RST in her Ph.D. thesis (Fox 1987), and Cece Ford. Barbara analyzed texts as long as articles in *Scientific American*, and she reported that these analyses had to be spread out on very large sheets of paper in her living room. Sandy was thus the guiding linguistic light for us at UCLA; she was also David Weber's, Susanna Cumming's, and my Ph.D. supervisor.

Sandy, Bill, and I had never anticipated the interest that RST would generate in both computational linguistics and linguistics more generally. In the 1990s, Daniel Marcu (1997) met the challenge of developing an approach to automatic RST analysis and also made other significant contributions to the computational modeling of RST, and in the 2000s, in Michio Sugeno's laboratory, Toru Sugimoto has used RST to represent the stages in the "translation" from user requests to language-based programming sequences. RST has certainly become another of Bill's lasting contributions to research on language. Bill even experimented with using RST in the production of his own research articles; he reported to me with satisfaction that this method speeded up the writing process—but he added, with a grin, that RST did not make his articles more entertaining. (But entertaining he certainly was, and very good at anecdotes and other story genres, and his presidential address to ACL was a unique contribution: He produced an anthology of language-based humor, organized in such a way as to illustrate the rich range of linguistic phenomena the humor revealed.)

Green Flashes and Chinese Food

Bill created an extraordinary environment of stimulating research, supportive teamwork, and friendship: It would be impossible to identify all the professional and personal friendships that he was midwife and godfather to. Bill had quite a large office at ISI, on the 11th floor of one of the twin office towers in Marina del Rey where ISI is located. His office faced west, with a panoramic view of the Marina del Rey yacht harbor and the Pacific Ocean. We were always welcome in his office, and he kept a supply of Chinese snacks and dry noodle dishes which we were welcome to sample even when he wasn't there—a welcome source of energy during late-night sessions of Penman development. We would often gather in his office at sunset to take in the spectacular southern California sunsets, enhanced in beauty by a combination of natural dust from the desert and L.A. smog, and Bill would tell us about the green flash that often appears just as the sun disappears behind the horizon, delighting in the common transformation of disbelief into surprised recognition when the flash actually appeared. Bill loved Chinese food, and he would regularly take us on lunch or dinner excursions to Chinese restaurants, locally or in Chinatown. This always happened when Michael Halliday visited from Sydney, and these occasions were a special treat, because Michael would order in fluent Mandarin, and delicacies not available on the English menu would appear. Later Bill himself studied Mandarin.

Family Man

Bill and LaDonna, his wife, would also invite us home to welcoming, warm companionship and delicious food. They had three wonderful children, and guests would be part of this loving family for a while. It was very clear how much strength and support they gave one another, which extended to their involvement in the church and its community. Bill was an only child, and he lost both his parents in the 1980s, but his

faith sustained him and gave him a perspective on the changing phases of life. When his father had passed away, I remember he told me he himself was ready to die—in the sense of being mentally and spiritually prepared, but not in the sense of not wanting to be part of life: Life did seem to be a joy to him, and he made the most of the time he'd been granted.

Population Growth

In the second half of the 1980s, Bill's research team at ISI became quite substantial. During this time, he and Norm Sondheimer (now at the University of Massachusetts) worked together as project leaders, and we were collaborating with a team led by Ralph Weischedel at BBN on a joint system involving both text generation and text understanding. The ISI team now also included John Bateman (a linguist and computational linguist who'd joined us after text generation research at Kyoto University), Ed Hovy (a computational linguist who'd joined after completing his Ph.D. with Roger Schank), Bob Kasper (a computational linguist), Robert Albano and Richard Whitney (computer scientists), Lynn Poulton (a linguist who'd come over from Halliday's department in Sydney to test and improve the Nigel grammar) and Susanna Cumming (a linguist doing her Ph.D. with Sandy Thompson at UCLA). The systems developed at ISI and BBN shared certain resources. While the grammars were distinct, the lexicon was the same; Susanna Cumming masterfully produced lexical specifications that could be used by both the Nigel grammar and the BBN ATN grammar. As part of this bicoastal collaboration, we had memorable workshops in Cape Cod and Palm Springs. By this time, there was also other NLP research at ISI—in particular, research led by Bill Swartout (with a background in expert system explanation) with Cécile Paris (who'd joined ISI after completing a Ph.D. with Kathy McKeown at Columbia) and Johanna Moore (who was working on her Ph.D. at the time); and there was a group developing the LOOM knowledge representation system (led by Robert MacGregor at the time). Cécile and Johanna became key members of our text generation research community at ISI, and drawing on her own previous research, Cécile would contribute, together with John Bateman, to the expansion of the Penman generation capabilities (Bateman and Paris 1991).

Bloomington

Bill's guiding hand was behind these and other developments. At the time when we had started modeling our first domain in the early 1980s, an electronic mail and calendar system, there was no sustained research into ontologies for knowledge bases (or other parts of a system), so there was no general guidance indicating how to model knowledge domains and not enough information of the kind needed in the generation of text. (There was important relevant research to take note of, like Bob Amsler's earlier investigation of implicit organization in the glosses used in *Webster's Dictionary*, revealing implicit taxonomies, and, from another angle, Jerry Hobbs's TACITUS project.)

So Bill telephoned Michael Halliday and me to give us the task of proposing such an ontology. At the time (the summer of 1986), we were in Bloomington, Indiana, Michael to give a course on "grammar in daily life" as part of a semiotics institute and I to attend this and other courses. We set to work on this task, often in cafés, and called the proposal we delivered to Bill "the Bloomington lattice." One key property of this ontology was that it was based on systematic and comprehensive evidence from language, which was in line with Bill's view of natural language as the most powerful knowledge

representation around. The Bloomington lattice turned into the “upper model” used in the Penman system and a number of other systems as well (Bateman et al. 1990; Matthiessen 1987). It is still being actively used today, for example, in a research effort at the University of Bremen conducted by John Bateman and others involving robotics and natural language processing.

Phase IV: Discourse Analysis and Linguistics

My memories of daily interactions with Bill and his team come to an end with my departure from ISI in August 1988, although Bill and I continued our joint work on RST, one result of which was Mann and Matthiessen (1991). As it happens, 1988 was also the year that the founding director of ISI, Keith Uncapher, retired from that position, and Herbert Schorr took over, bringing with him executive and research experience from IBM. Norm Sondheimer also left in 1988, to become the manager of the Information Technology and Artificial Intelligence Laboratories at the GE Corporate Research and Development Centers. The funding climate for computational linguistics research had already begun to change in the United States. In the late 1970s and early 1980s, Bill had made text generation attractive to potential funding agencies as a technique for knowledge delivery. Now, about a decade later, he tried to sell it as the graphics of the 1990s. But he stayed on at ISI for only another couple of years, because new opportunities and challenges opened up for him, and he was ready to embark on Phase IV of his career, one in which the main motif was to be discourse analysis rather than computational linguistics. He took early retirement from ISI in 1990, and Ed Hovy took over as project leader, continuing the research tradition and opening up new areas of research (including collaborative work on machine translation in the context of the Pangloss system).

Bill’s career change in 1990 was a major one: He took up a position as a consultant with SIL and moved with his family to their new base in Nairobi, Kenya, where they stayed until 1996. His time was divided between discourse analysis and translation (where he continued the work on dialect adaptation in African contexts). We had hoped that I would be able to visit, but this never happened, and I learned about this new phase from enthusiastic letters Bill wrote describing the new context and reporting on the new research.

After his return from Africa, Bill focused on linguistic research, continuing work on RST (e.g., Mann and Thompson 2000) and on dialogue games (developing his earlier theory into dialogue macrogame theory), but also pursuing related work on intention and communication, and this could certainly be characterized as Phase V of his intellectual journey. He passed away during a period of full activity, with various projects in the pipeline, including an article with Maite Taboada surveying contributions to RST.

Bill’s Legacy

Bill’s contribution to computational linguistics has been immense. In a way, it’s too early to speak of his legacy—because he was a true visionary, and while we already know that he and others influenced by him have realized key aspects of his vision of what computational linguistics can achieve, other aspects will take a longer time to realize, since they would represent a Kuhnian kind of paradigm shift. His vision was multifaceted, involving many deep insights; but one key motif was his serious engagement with language in his approach to computational linguistics.

Part of Bill's legacy is also the example he gave us of how one can create an academic life guided by a spirit of intellectual curiosity and social responsibility, changing his own career and approach out of conviction rather than because of fashion. He was, as Jim Martin put it to me, in lamenting his death, "a wise fellow, and very kind and considerate." This was, I think, the wisdom of a Renaissance man, equally at home in physics, biblical archaeology, engineering, computer science, computational linguistics, linguistics, and Bible studies—or the game of Go, which he was very good at; but it was also a wisdom he would distill into simple sayings like "Blame's not part of the game" and "Hindsight is wiser than foresight." As a person, he had real presence, radiating a kind of soothing calm, and he was also full of laughter and joy.

References

- Bateman, John A., Robert Kasper, Johanna Moore, and Richard Whitney. 1990. A general organization of knowledge for natural language processing: The penman upper model. Technical report, Information Sciences Institute, University of Southern California.
- Bateman, John A. and Cécile Paris. 1991. *Constraining the deployment of lexicogrammatical resources during text generation: Towards a computational instantiation of register theory*. Mouton de Gruyter, Berlin and New York.
- Becker, J. D. 1975. The phrasal lexicon. In Roger Schank and Bonnie Nash-Webber, editors, *Theoretical issues in natural language processing*. Association for Computational Linguistics, New Brunswick, NJ, pages 38–41.
- Brachman, Ron J. 1979. *On the epistemological status of semantic networks*. Academic Press, New York.
- Cross, Marilyn. 1991. *Choice in text: A systemic approach to computer modelling of variant text production*. Ph.D. thesis, Macquarie University, North Ryde, NSW, Australia.
- Davey, Anthony. 1978. *Discourse production: A computer model of some aspects of a speaker*. Edinburgh University Press, Edinburgh.
- DiMarco, C., G. Hirst, L. Wanner, and J. Wilkinson. 1995. Healthdoc: Customizing patient information and health education by medical condition and personal characteristics. In A. Cawsey, editor, *Proceedings of the Workshop on Patient Education*, Glasgow.
- Fox, Barbara. 1987. *Discourse structure and anaphora: Written and conversational English*. Cambridge University Press, Cambridge.
- Hovy, Eduard. 1988. Planning coherent multisentential text. In *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, pages 163–169, Buffalo.
- Kasper, Robert. 1989. A flexible interface for linking applications to penman's sentence generator. In *Proceedings of the DARPA Workshop on Speech and Natural Language*, Philadelphia, pages 153–158.
- Kasper, Robert and David J. Weber. 1986. User's reference manual for the c quechua adaptation program. Technical Report No. 8, Summer Institute of Linguistics, Dallas, TX.
- Kay, Martin. 1979. Functional grammar. In *Proceedings of the Fifth Annual Meeting of the Berkeley Linguistics Society*, Berkeley, CA, pages 142–158.
- Mann, William C. 1979. Dialogue games. Technical Report ISI/RR-79-77, USC Information Sciences Institute, Marina del Rey, CA.
- Mann, William C. 1982. An overview of the Penman text generation system. Technical Report ISI/RR-83-114, Information Sciences Institute, University of Southern California.
- Mann, William C. 1983. The anatomy of systemic choices. *Discourse Processes*, 8(1):53–74.
- Mann, William C. and Christian M. I. M. Matthiessen. 1991. Functions of language in two frameworks. *Word*, 42(3):231–249.
- Mann, William C., and Thompson, Sandra A. 1988. *Rhetorical structure theory: Toward a functional theory of text organization*. *Text*, 8, 243–281.
- Mann, William C. and Sandra A. Thompson. 2000. Toward a theory of reading between the lines: An exploration in discourse structure and implicit communication. Paper presented at 7th IPra International Pragmatics Conference.
- Marcu, Daniel. 1997. The rhetorical parsing of natural language texts. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics*, pages 96–103, Madrid, July 7–10.
- Matthiessen, Christian M. I. M. 1981. A grammar and a lexicon for a text

- production system. In *Proceedings of the 19th Annual Meeting of the Association for Computational Linguistics*, Stanford University, Stanford, CA.
- Matthiessen, Christian M. I. M. 1983. Systemic grammar in computation: The nigel case. In *Proceedings of the First Conference of the European Chapter of the Association for Computational Linguistics*, Pisa.
- Matthiessen, Christian M. I. M. 1987. *Notes on the Environment of a Text Generation Grammar*, pages 253–278. Martinus Nijhoff, Dordrecht.
- Matthiessen, Christian M. I. M. 1995. *Lexicogrammatical cartography: English systems*. International Language Sciences, Tokyo.
- Paris, C., K. V. Linden, M. Fischer, A. Hartley, L. Pemberton, R. Power, and D. Scott. 1995. A support tool for writing multilingual instructions. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1398–1404, Montréal.
- Rösner, D. and M. Stede. 1994. Generating multilingual documents from a knowledge base: The techdoc project. In *Proceedings of the 15th International Conference on Computational Linguistics*, volume 1, pages 339–346, Kyoto, Japan.
- Weber, David J. 1989. *A grammar of Huallaga (Huánuco) Quechua*. University of California Press, Berkeley and Los Angeles.
- Weber, David J. and William C. Mann. 1979. Prospects for computer-assisted dialect adaptation. Technical Report No. 1, Summer Institute of Linguistics, Dallas, TX.
- Winograd, T. 1972. *Understanding natural language*. Edinburgh University Press, Edinburgh.
- Woods, W. 1975. *What's in a link? Foundations for semantic networks*. Academic Press, New York.

