REPORT
Academic Specialist Program
USIA
United States Information Agency
Lecture Tour by
Frank L. Borchardt
in the Netherlands

24 March - 8 April 1993

Durham, 25 April 1993

Observations

The Specialist (Frank L. Borchardt) profoundly relished the time he was allowed to spend in the Netherlands and the work he was allowed to do. The entire tour, from preliminary planning to final execution, was organized with a vanishing minimum of red tape and a maximum of efficiency by USIA Washington, the American Embassy in the Hague, its representative, the cultural attaché, Kate Delaney, and her agent at the Netherlands Institute for Curriculum Development in Enschede, Gé Stocks. The Specialist is much indebted to them all.

The Specialist was received by one conference (OWG [Onderweis Werkgroep], Zeist), two educational service institutions (SLO, Institute for Curriculum Development, Enschede, and CITO, Bureau for Educational Measurement, Arnhem), two Teacher Training Colleges (Utrecht and Tilburg), and three Universities (Nijmegen, Leiden, Twente). In each case without exception, the reception was extravagantly generous, courteous, and kind.

The overall impression was shaped and confirmed repeatedly that the Dutch are very seriously concerned about the successful employment of information technologies in their educational system, that they are worried about their young people having not enough or having the wrong kind of experience with high technology, and that the connection between language learning and high technology is a fertile area for doing something positive about these worries.

The Dutch are far from lacking resources in these areas, either in expertise or materiel. The Specialist encountered exceedingly progressive projects at SLO, the College of Educational Technologies (HIO) in Enschede, and at the Universities of Leiden, Nijmegen, and Twente addressing these questions head on. In the face of funding constraints afflicting education world-wide, others, notably the Teacher Training Colleges at Utrecht and Tilburg, were experimenting with creative low-cost solutions to an array of problems of educational technologies. The OWG Onderweis Werkgroep left a strong impression on the Specialist for its vigor, its grass-roots organizing principles ("bottom-up"), and the broad spectrum of its services.

Recommendations

- If there is a U.S. outfit comparable to the OWG, it would be well for it to seek liaison with this creative and dynamic organization. If there is no such grass-roots, K-12 organization, it would be well to seek to found one. *A recommendation should be made to the Department of Education to contemplate the funding of grass-roots educational technologies initiatives for in-service teachers.*

- The Specialist further recommends that USIA itself issue an RFP for international liaisons in the area of educational technology, especially as applied to modern language learning, between individuals, organizations, and institutions in the United States and specific areas abroad, Western Europe, Central Europe, CIS, Middle East, South Asia, and East Asia, in both directions.

SCHEDULE

Program
for the visit of professor Frank Borchardt to the Netherlands
25 March - 8 April 1993

Thursday 25 March	07.30	arrival at Schiphol airport (Flight DL 38)
	11.00	arrival at Enschede; meeting and lunch with Ms Marita Tholey of SLO/NICL
Friday 26 March	afternoon	Preparation and hardware check for lecture on Saturday with Gerhard Hoogterp at SLO
	19.00	Dinner with G. Stoks
Saturday 27 March	12.00 - 13.00	Lecture "Prophets and Profits" CAL-meeting at Zeist; meeting with US-embassy official
Sunday 28 March	Open	
Monday 29 March	12.30 -13.30	Lunch with Mr Piet de Bruijne, director of SLO and Ms Tineke Zeelenberg, project coordinator for IT at SLO
	14.00 - 16.00	Meeting at SLO, discussion with IT-colleagues
Tuesday 30 March	14.00 - 16.00	Visit to Leiden University: Meeting with professor Gerard Kempen and staff members for neural nets and/or psycholinguistics at Leiden University, Experimental Psychology Contact person: professor Gerard Kempen
Wed'day 31 March	15.00 - 18.00	Visit to Nijmegen University: Lecture "multimedia technologies for second language acquisition" at Nijmegen University, Institute of Applied Linguistics Contact person: Mr Charles van Os
Thursday 1 April	13.00	Meeting with staff of the languages department of Utrecht teacher training college to discuss multimedia technologies for second language acquisition
	14.30 - 16.30	Lecture "Future trends in educational technology" for staff and teachers of the Utrecht teacher training college. Contact person: Mr Ton Koenraad
Friday 2 April	16.00	Colloquy "the future of educational

		technology" with staff and students of the department of educational technology of Twente University, Enschede Contact person: Mr Gerrit Carleer
Saturday 3 April	Open	
Sunday 4 April	16.00 - 18.00	Reception at Gé Stoks' home in Borne
Monday 5 April	4.00 - 16.00	Lecture on "Computer Assisted Assessment" at Cito, Institute for Educational Measurement, Arnhem. Contact person: Mr José Noijons
Tuesday 6 April	14.00 - 16.00	Columbus' lecture "the future of educational technology" at SLO
Wednesday 7 April	14.00 - 16.00	Lecture on "the future of educational technology" at teacher training college in Tilburg Contact person: Mr Harry Jansen
Thursday 8 April	09.00	Breakfast with Ms Kate Delaney, US cultural attaché at the American Embassy in The Hague
	13.55	Departure (flight DL 39) from Schiphol airport

Program coordinator:
Gé Stoks
SLO, Institute for Curriculum Development
P.O. Box 2041
7500 CA Enschede
The Netherlands
Phone: +31.53.840.840

Temporary address in the Netherlands:
Best Western Dish Hotel
Boulevard 1945 2
Enschede
The Netherlands
Phone: +31.53.86.66.66

ITINERARY

The Specialist (Frank L. Borchardt) departed Raleigh/Durham on Wednesday 25 March 1993 for Enschede, the Netherlands, arriving there Thursday 26 March. He was met by Marita Tholey of SLO, the Netherlands Institute for Curriculum Development, with whom he consulted on the disposition of the lecture tour. On Friday 26 March the Specialist met with Gé Stoks, organizer of the tour, and Gerhard Hoogterp to arrange technology for the tour. On Saturday 27 March, the Specialist was driven to Zeist for the meeting of the OWG Onderweis Werkgroep,

at which he lectured on the topic, “Prophets and Profits: On the Future of Educational Technologies” (see attachment 1) and consulted with the Director, Mr. Wim Buwalda. The Specialist met with Tilly de Groot, representative of the U.S. Embassy in the Hague.

> **!** The Specialist was very favorably impressed by the work of the OWG, which is a grass-roots organization, representing programming elementary and secondary school teachers. Their educational technologies are created *ad hoc*, in the trenches, for immediate use in the classroom. OWG accepts their programs, cleans them up, arranges for documentation, and puts them on the open market at cost with modest royalties to the authors.

On Monday 29 March the Specialist was hosted at lunch by Mr. Piet de Bruijne, Director of SLO, and Tineke Zeelenberg, Project Coordinator for Educational Technologies. The Specialist was privileged in the afternoon to observe presentations on the work of students and consultants: Ilse Siteur from the University of Twente (UT) on motivating lower secondary school pupils in the realm of educational technologies; Elly Stienstra and Margré Mollink from the College of Educational Technologies (HIO), Enschede, on in-service, on-the-job training; and Jan Lapeltak on video for SLO.

> **!** The collaborative labors between SLO, UT, and HIO left a very favorable impression, not only for the substance of the projects but also as a model for cooperation of adjacent institutions with comparable missions.

On Tuesday 30 March the Specialist was received at the University of Leiden by Professor Gerard Kempen at his Department of Social & Behavioral Sciences. The Specialist reported briefly on the work of his project at Duke University in the area of Artificial Neural Networks (ANNs) and their application to problems of Natural Language Processing (NLP; see attachment 2). He was thereupon shown the work of Professor Kempen's team in the area of NLP for Dutch and an allied enterprise in ANNs.

! Professor Kempen and his colleagues, Dr. Koenraad deSmedt, Dr. Peter Diesveld, and Alice Dijstra gave evidence of work of the highest quality, fully comparable to their U.S. counterparts in the whole area of Artificial Intelligence (AI) and ANNs. Distinctive and most impressive was their willingness to test their theoretical explorations with actual applications both in AI and ANNs both in software and hardware.

On Wednesday 31 March the Specialist was received at the University of Nijmegen by Dr. Everhard Ditters in the absence of Charles van Os and lectured on the topic: "Multi-Media Technologies and Second Language Acquisition" (see attachment 3).

! The audience in Nijmegen included Dr. Mark Leiblum, who has had a long history of involvement with computer assisted instruction (CAI) in general.

On Thursday 1 April the Specialist was received by Dr. Toen Koenraad at the Faculty of Education at the Hogeschool Midden Nederland in Utrecht, where the Specialist met first with a commission charged to equip a modern multi-media language laboratory in the face of personnel restrictions. He brought along information on instances of multi-media projects in full implementation in the U.S.

The Specialist thereupon lectured on "Trends in Educational Technology" and demonstrated the actual implementation of the Digital Video Interface (DVI), digitized video without special video peripherals on a portable computer from within **CALIS** (**C**omputer **A**ssisted **L**anguage **I**nstructional **S**ystem).

The Specialist spent Friday 2 April at the University of Twente, had consultations with Dr. Betty Collis and Dr. Italo DeDiana, and participated in a small colloquium with Mr. Gerrit Carleer, staff, and student on the future of educational technologies.

On Monday 5 April the Specialist was received by Harriet Robben, José Noijons, and a Vice President of CITO, the Netherlands Bureau for Educational Measurement in Arnhem. After lunch, the Specialist reported on the activities of

CALICO, the CALIS project, and the near-term prospects for educational technologies.

On Tuesday 6 April back in Enschede, the Specialist delivered the Columbus Lecture, "Electronic Media and Educational Revolution," with the help of Gé Stoks (see attachment 4).

This event was repeated on Wednesday 7 April at the English Department of the Teacher Training College, Hogeschool Katholieke Leergangen, Tilburg, with Harry Jansen and Fritz Schulte hosting.

The Specialist departed Tilburg for the Hague, was met by Gé Stoks, and spent the morning of Thursday 8 April in a debriefing with Kate Delaney, U.S. Cultural attaché at the American Embassy in the Hague. The Specialist arrived back at his home institution that evening.

USIA Report
Attachment 1

Prophets & Profits
Forecasting Educational Technologies

Frank L. Borchardt

Abstract

Teaching and learning with *text* is available wherever there is a computer of almost any kind. Countless activities traditionally conducted on the printed and written page can be undertaken as well on the electronic page. The advantage to the teacher can be that the teacher may not have to correct and grade the work (is there a higher advantage to a teacher?). The advantage to the student can be privacy, self-pacing, and an infinitely patient, generally available tutor.

It takes a slightly, but only slightly more up-to-date computers (vintage 1986 would probably suffice) to add *graphics* to the process, cartoons, maps, flags, and alphabets not customarily seen on computers in the West (Arabic, Hebrew, Chinese, Japanese). So far, this is a world very like Mr. Gutenberg's, except for the medium answering back, which Mr. Gutenberg's invention was never intended to do. We start to depart farther from Mr. Gutenberg's invention when we add audio and video to the process. Computers love to govern things. They *like* to turn on your traffic lights and turn off your microwave oven. And the same is true of your audio tape player or videodisc machine. But experience has shown that every new box attached to the computer adds ten-fold complexity, so that only the most enthusiastic techno-teachers will bother to get involved.

The most modern computers (vintage 1991 and later) hold other promises, most importantly, eliminating the new boxes. All the newest computers really want and need is a source of information that will talk to them in their own (binary) language. Ordinary audio tape and videotape or video machines do not yet speak this language. But, the CD player on which most of us hear our recorded music *does* speak the computer's language, as do the encyclopedias, bibliographies, and such available on CD/ROM and probably already in use at your local public or university library. All that is still needs doing is a slight increase in storage capacity and speed and a slight decrease in cost. Then there will be no representational medium that can capture sight and sound that will be outside the capacity of the ordinary desktop computer to transmit, interactively, for the general public, for its entertainment, for its information, for its instruction.

****

There is a saying going around the United States, so common it is becoming a cliché: How can you tell the pioneers? They're the ones with the arrows in their chests.

It means: “be careful if you are the first to reach a new technology, you are engaging in a very dangerous activity; let others take the first risks.” One thinks of the Gutenberg Legend, and his famous bankruptcy.

Computer assisted instruction (CAI) has been around for no less than twenty years, so that the risk of being the very first is much diminished. Indeed, unless someone soon produces a history of CAI, the first generations are likely to be forgotten altogether, with only some vague memory surviving that there was once such a thing as PLATO and plasma terminals to deliver the graphics characters.

Since then, the personal computer revolution has rewritten all the rules. What was once a vague and distant threat, located in its own reserved, climate controlled, sanitized, dust-free, secure temple, overseen by a dedicated priesthood of fanatic technicians, incomprehensible to all others, but influencing their lives nonetheless—that remote main-frame computer has undergone a magical transformation, shrunk in size, been replicated many hundreds of thousands of times, and found its way into virtually every business office on earth, yea, even onto the desks of ordinary citizens, far more powerful, in every sense, than it ever was when it dwellt lonely in its inner sanctum.

The monster was invented in order to count, in order to calculate. Now-a-days, except at tax-time and maybe during a census, most of the machines are used most of the time for other activities altogether. And the great preponderance of those activities has to do with words. It turned out that this machine was not really a calculator to begin with, but rather a logic machine, which could judge true and false very quickly, and thus function, secondarily, as a calculator doing binary arithmetic. It was not long before it dawned on people that the same logic could be employed to represent letters as well as numbers. Since people verbalize more than they count, the machines have themselves become more and more verbal, until now, they have replaced in most markets, the typewriter, except perhaps for the occasional reincarnation of Ernest Hemingway and Paris in the 'Twenties.

Why would anyone want to give up a perfectly good typewriter for a difficult and temperamental word-processor? Some of the reasons are obvious. The word-processor lends itself to revision and correction. One can have both a typewriter and a dictionary in the same machine. That's convenient. And then, one can store the results in a good deal less space, file it, and retrieve it rationally—if one is careful. All of a sudden it starts to become a typewriter and a dictionary and a filing cabinet. It can store mailing lists and "personalize" impersonal letters. It can also, true to its early employment, keep accounts and inventory and connect all of this information in any direction. Well, now it is a typewriter and a dictionary and a filing cabinet and a mailer and a ledger—all without taking up any more room on the desk. Attach a modem and it can answer the phone, route calls, itself initiate phone calls, hook up to another computer with similar connections to the outside world, and provide access to all the information stored on all the connected computers in the world. Well, now it is a typewriter and a dictionary and a filing cabinet and a mailer and a ledger and a telephone, telex, telegraph, and a universal encyclopedia—all without taking up any more room on the desk. Now, that is pretty useful, as any well automated office will tell you.

But have a closer look. Is there anything new there, anything that hasn't been there before? Any service that cannot still be accomplished by prior technologies? The answer is no. We are, even the most advanced among us, still in the "cradle" era of computing technology. The Latin for "cradle" is "incunabulum," and we use the term technically to refer to books printed in the first fifty years of printing with movable types. For those fifty years printers did everything they possibly could to make printed books look like hand-written manuscripts. The fonts imitated the best hand (until Aldus invented italics). Initials were surrounded by white-space to leave room for the miniaturist to illuminate the letters. There were no indexes (maybe a generic table of contents, but then no page or folio designation), no title pages, nothing to disturb the reader of manuscripts, except, of course, that each impression of each type, each copy of each run was pretty much identical, that information found on page so-and-so of one copy of one edition would be found at the same place in another copy of the same edition. And, most importantly, that this marvelously portable, randomly accessible technology, best used in private, perhaps even the creator of privacy, could be reproduced indefinitely at progressively lower cost.

[One wonders when reading aloud gave way to reading in private, and what role reading played in the creation of privacy.]

If the art of printing itself took fifty years to begin to find its own identity, and if the business office itself, constrained by the demands of efficiency, still employs computing power retrospectively, in respect to past technologies, is it any wonder that these technologies, when they reach the classroom, incline to look suspiciously like textbooks? It is no wonder at all.

The whole of the early generations of CAI, including the first of the personal computer era, were basically "drill and practice," what is now disparagingly called "drill-and-kill," in an attempt to force improvement in the deliver of electronic instruction. These exercises took from the textbook, often letter by letter and word-for-word, those dreary, repetitive exercises meant to jog from its slumber the art of memory, which writing and the art of print put to sleep to begin with.

For those kinds of knowledge best acquired by rote, by memorization, these kinds of exercises are still effective. The computer has one immense advantage over its hard-copy forebear, however, in that it can respond immediately, critique, and point in useful directions any time of day or night, in the absence of the teacher, in a private and completely non-threatening environment. It would be safe to generalize: any written assignment given on a Monday, collected on a Wednesday and returned on a Friday—when neither teacher nor student have the vaguest recollection of what was asked or answered, can be better accomplished by a computer. And just about any computer can do this. One doesn't have to have the latest and the best.

The typewriter and the printed page provided the first models for computer assisted instruction. The computer can well imitate both. The question is, should it? Is there not a fundamental difference between paper and ink reflecting and not reflecting light, on the one hand, and on the other, light emanating from an electronic screen? Is it likely, or even possible, that the same aesthetics govern media which are materially so different? This question is far from answered, but its underlying thrust should provide the consumer with some criteria for judging the

courseware offered in the market place. Is the screen covered with words? Does it look like a printed page? Is it hard to read or does it require immense concentration, far more than reading a printed page? If the answer is yes, the consumer is looking at first-generation courseware which has not yet come to terms with the dynamics of the new medium.

In the second generation of the personal computer revolution, with the adoption of graphics standards around 1985-86 and the introduction of the AT 16-bit standard an important step forward was taken. The underlying metaphors remained the typewriter and the book, but this could be a book written in alphabets other than those descended from Rome and adorned with illustrations to one's heart's content. To be sure, the computer itself "understands" only the letters of the Roman alphabet and a few punctuation marks, but it could be fooled into thinking that these Russian or those Greek characters were only Roman characters in disguise. It opened standard office-based computing on the world. About the same time, a competing standard entered the market place, one that abandoned the typewriter metaphor and any pretense that computers were chiefly readers of encoded information. This technology admitted from the very outset, it was a (pseudo-) printed page. The monitors had white screens with black letters and just about everything that appeared on them was presented in graphical mode, that is, drawn or painted, not "typed."

With this generation, the imitation of the book became complete. Graphics could be added to illustrate a point, to reduce the density of print, to present information pointedly, economically. Printed textbooks themselves, probably in response to conditioning by other media, began to present less print-per-page, more white-space, more illustrations, experimenting with the best proportions between print and page, for a market that was beginning to lose the ability to concentrate on dense print pages. Courseware addressing this development is already light-years ahead of "drill and kill." Consumers looking for criteria can be assured, if the program handles graphics intelligently, employing them for legibility as well as for learning strategy, they are looking at no less than second-generation software, courseware gradually trying to come to terms with electronics.

Computers love to govern things. They *like* to turn on your traffic lights and turn off your microwave oven. And the same is true of your audio tape player or videodisc machine. But experience has shown that every new box attached to the computer adds ten-fold complexity, so that only the most enthusiastic techno-teachers will bother to get involved.

The most modern computers (vintage 1991 and later) hold other promises, most importantly, eliminating the new boxes. All the newest computers really want and need is a source of information that will talk to them in their own (binary) language. Ordinary audio tape and videotape or video machines do not yet speak this language. But, the CD player on which most of us hear our recorded music *does* speak the computer's language, as do the encyclopedias, bibliographies, and such available on CD/ROM and probably already in use at your local public or university library. All that is still needs doing is a slight increase in storage capacity and speed and a slight decrease in cost. Then there will be no representational medium that can capture sight and sound that will be outside the capacity of the ordinary desktop computer to transmit, interactively, for the general public, for its entertainment, for its information, for its instruction. Add to this the development of international standards for the representation of all the languages of the world in encoding that the computer can understand—knowing the language is Hebrew, for example, and that the character is aleph (not, as is the case now, pretending to be English and a Roman "a")—and little is left over to make the ideal language teaching machine. We lack only speaker independent, continuous speech recognition; and that, too, may be closer than anyone imagines.

The basic metaphors we have been seeing across this history of applied computing include: the calculator (or adding machine), the typewriter (early computers had no cathode ray tubes and communicated back and forth by means of teletype machines, the unit of communication being the typed line); the typed page, the printed page, the book, the book-shelf, the encyclopedia, and the engines of the business office. We have had to add peripherals, computers plus audio ("Walkman"), computer plus video. And a new metaphor has crept in, one few people dare to name, and it is the infrared remote control.

The principle of "two steps forward, one step back" applies in educational technologies as it does in a great many fields. Just when we think we are making computers as powerful as they can possibly be, what happens? Just as we are introducing multi-media and "interactivity," what happens? Human sloth enters the picture, and, with the "graphical user interface" enter the mouse, with "point-and-click," the computerized version of couch-potato infra-red remote! Not Inter-Activity, "Inter-Passivity"! If one suggests a one-year moratorium on "point-and-click," one is greeted first with the spontaneous laughter of surprise, and then a growing realization of what it would imply and mounting horror at the prospect.

The alternatives, language production, are easily available. There are three basic strategies for answer analysis: simple string-matching, wild-cards, and the so-called "edit distance algorithm." Computers are very good at comparing. They can compare with ease the difference between an anticipated answer and an actual answer. This can be made quite flexible by the same tools one uses to ask a computer to provide a directory listing of all the files names beginning with the letter "s." And then there exists in the public domain a very convenient routine by which a grid is established with an anticipated response on one axis and the actual response on the other; a quick comparison is made, and the algorithm can determine errors in capitalization, wrong letters, wrongly repeated letters, missing letters, and transposed letters, all instantaneously without further author intervention. The tools for most keyboard linguistic input, in sum, exist and are well known. When they are not employed, why are they not employed? Because point-and-click seems easier, does not require typing skills, permits recognition, does not require production. When speaker independent, continuous speech recognition comes on the scene, the mouse will probably go the way of all flesh. But in the meantime, it is almost certainly doing more harm than good to the learning process.

What does the future hold? The fundamental technologies basically exist already. If you can digitize it, a computer can present it. If we had infinitely large storage devices and infinitely fast transmission speeds, we could have it all now. But in their absence we have a problem. One digitized video frame contains about one megabyte of information. Full motion video requires thirty frames per second. Even very fast modern computers have a problem transferring that much information that fast, not to mention what that much storage would leave over on our hard-disks

for other purposes. The industry has responded with “compression.” And now, it is just a matter of time for the standards to slug it out in the market place.

IBM has announced, but not released a product called the *M-wave*. It is a board that can send and receive high resolution faxes, perform modem communications at high speed, compress and present graphics, record (digitized) sound, speed up and slow down without loss of pitch or timbre (there used to have to be a whole box, “InstaVox,” just for that function for interactive video soundtracks), record or play electronic music (MIDI) files, recognize speaker dependent, articulated (space-between-words) speech, display still video, present three-dimensional graphs of sound in real-time, and answer your telephone. If it can be digitized, this board can present it (except for motion video). Whether this particular technology will win out is hard to say. But it already exists and faces now only the test of the market place.

Intel has announced and slowly put into practice a standard called DVI (Digital Video Interface). The advantage of this over others is that it can be realized in software and does not need an extra board or exterior unit. It is the first fulfillment of the promise of universal digitization, including motion video.

What the future holds is quite clearly that: a single mass storage device which will hold all the information the computer could hope to process: data, text, graphics, sound, video. Bringing these all together will be the work of the pioneers. . . to end where we began.

USIA Report
Attachment 2

Position Paper on Artificial Neural Networks

Frank L. Borchardt

The field of Artificial Neural Networks, fallow in North America for almost a generation, felt the plow again in the very few years before early Summer 1987. That June in San Diego, the landmark "First International Conference on Neural Networks" was held before a huge crowd of several thousand participants and observers. (IEEE First International Conference on Neural Networks, San Diego June 21-24, 1987) A colleague who had gone to the ICH conference instead noted wryly that 250 people had shown up at the meeting of an established society, where ten times that many had chosen to hear about a field no one was quite sure existed. What came out, at that landmark meeting, early and often, was a profound hostility to AI as espoused at MIT, even though the guiding light of that lab had early in life designed an elaborate electronic artificial neural network. ANNs were supposed to establish a radically different model, not serial but parallel, not digital but analogue, not symbolic but connectionist, not local but distributed, not deterministic but associative, not *a priori* but inductive. "There are no symbols, only contexts." There seemed to burst forth some kind of new hope) in robotics, in computer vision, in Natural Language Processing) for an actual delivery what AI had been promising for years.

The plausibility of this hope was buttressed by some stunning early experiments. Of these, the most accessible was Terrence Sejnowsky's "NetTalk," a text-to-speech machine. (Sejnowski, Terrence J., and Charles R. Rosenberg."Parallel Networks that Learn to Pronounce English Text," *Complex Systems*, 1 (1987), 145-168; Sejnowski, Terrence J., and Charles R. Rosenberg. "NETtalk: A Parallel Network that Learns to Read Aloud," *JHU/EECS-86/01*, Johns Hopkins University, EE/CS Department, January 1986.) He successfully trained a speech synthesizer to pronounce English without an algorithm, but rather from examples, and more or less as well as a child of three might pronounce English. The final result was not what was so disturbing about this experiment, but rather the process. Part of the way through the learning the result was tested) chiefly to make sure that the network was not memorizing the words but rather generalizing from examples. In the first steps beyond mere gibberish the network had learned the distinction between vowels and

consonants, though, to be sure, it pronounced all consonants as one and all vowels as one. To hear this intermediate and then the final result is enough to make a believer out of anyone) even though Marvin Minski having heard the tape, denied that the final result sounded like speech at all, human or any other kind. But be consoled: word had it that the Sejnowski experiment was successfully replicated in Austria in German, leading one to question the disinterestedness of Professor Minski's hearing.

Closer to home, at the North Carolina State University, a student Sofus Simonsen, (McKee, P.G. 1987: "Learning Genders of German Nouns: A Neural Network Application", Neural Engineering Technical Report 88-10, Center for Communications and Signal Processing, North Carolina State University.), struggling with the gender of German nouns, devised a network similar to Sejnowski's, which he trained to recognize the gender of a random half of the nouns in his textbook. This it did very well, learning over 98% of what it had been taught. He then tested the network for what it had learned on the other half of the nouns, and it got almost 2 out of 3 correct. And the mistakes that it made were rather like those students would make (chiefly, overgeneralizing).

One rather disturbing result of this experiment was the “second-best guess” phenomenon. If the network's highest certainty level was wrong, its second highest level was almost always right, and it had near zero certainty that the third alternative was possible. In all but one or two cases (das “Ende”), the network consistently guessed right about which gender the noun was *not*. This experiment and its variations were shown over the years to many visitors. One of these was the brother of the President of the Czech Republic, Dr. Ivan Havel. He was equally disturbed by the “second best guess phenomenon” and tried to force the network to reveal its reasoning by inputting slight variations of such a word, one vowel off or one consonant off, etc. The network, not surprisingly for people familiar with them, revealed no consistent pattern to these stimuli. In frustration, Dr. Havel input a Czech word. The network responded after some “thought” (processing): “I am near 0% certain this is masculine; I am near 0% certain this is feminine; I am near 0% certain this is neuter.” In sum, the network recognized that the input string was nothing like anything it had been trained on, that is, that the string was not German. We are calling this “the Havel Result.” It basically proves that an Artificial Neural Network can, at the very least, negatively recognize language.

This result was closely parallel to another experiment conducted at Duke, the “bogus input experiment.” Students doing CAI homework had the choice of giving up after one attempt and receiving the correct answer from the program. A few

students, we learned from analyzing input data, routinely typed in garbage as their first attempt, gave up, transcribed the right answer, and repeated the exercise from their notes. It seemed like a worthy exercise for an ANN to see whether one could recognize and intercept garbage and tell the student "Big Brother knows what you are doing and doesn't like it." This network trained on legitimate strings from the original dataset ("right answers") and legitimate but erroneous answer attempts (from student input data). As it turned out, the network almost never classified a right answer or a legitimate attempt as "bogus" but caught bogus attempts only about half the time. We found this result altogether satisfactory for practical purposes. Designating a right answer or a legitimate attempt as bogus would be pedagogically disastrous, whereas missing a few bogus attempts wouldn't hurt, as long as they were caught sooner or later. This kind of "tolerance") one should note in passing) is what distinguishes ANN strategies radically from rule-based systems.

Similar tolerances emerged from yet another Duke project experiment ("Hyphenation and `Bogus' Word detection) An Experiment in the Economic Use of Neural Modelling in Text-Processing." Frank L. Borchardt, Andreas Geyer-Schulz, Wolfgang H. Janko, Nick Staddon, Hongbin Wang. *Diskussionspapiere zum Tätigkeitsfeld Informationsverarbeitung und Informationswirtschaft*. Vienna, April 1991). German hyphenation conventions were employed as a worst case. A network was trained on several corpuses of most frequently used German vocabulary. When novel input tested the results, the network regularly chose correct hyphenation points but often not all of them and rarely chose erroneous hyphenation points. The experiment indicated that an ANN could be used successfully as an alternative to algorithmic hyphenation strategies, and, more importantly, could be used as a strategy for producing acceptable hyphenation mechanisms for underdocumented languages, such as Kurdish, for which no hyphenation rules were known at the time of the experiment.

In collaboration with colleagues in the Slavic Languages Department at Duke, the Project undertook an experiment to compare the performance between human informants and an ANN when it came to nouns of ambiguous gender in Russian. ("Gender Signalling in Russian: a Contrastive Analysis between Native Speakers and Neural Networks.") We are hoping for imminent publication of this experiment (*Language Quarterly*, June 93). The questions were more complex than in other of our experiments and the conclusions are a bit more subtle. In a supervised learning environment, an ANN, once employing purely orthographic representation (spelling alone), once employing morphophonemic representation (phonetic representation of grammatically significant parts of the noun), showed that it could learn 99% of the

words presented to it and replicate correctly even ambiguous gender assignment. It refused to learn the eccentric gender of a “fond diminutive” *da'da*', which form dictates should be feminine, but which meaning dictates should be masculine.

As a control, the identical network was “trained” on a set of nonsense words to see whether it would produce a similar but pseudo-result. Happily, the network found no identifiable patterns in the nonsense dataset, which indicated a high probability that the network was in fact identifying patterns and not memorizing answers. The network, trained on the actual dataset, however, had great difficulty identifying novel input correctly, especially by contrast with the human informants. Once it was supposed that stress, in addition to morphemic information might improve the network's performance. That proved to be inconsequential for the network, but very important for the human informants. We assume that the relatively small size of the dataset (666 words) may be to fault. The most substantial distinctions between ANN and human performance came when semantics played a role. The network was wholly lacking any semantic information whatsoever. In a great number of cases, human decisions on words of morphologically doubtful gender were based on meaning. Perceived foreign origin also seemed to play a part. The experiment seemed to prove that formal marking, while systematic and not random, was by itself insufficient to determine the gender of ambiguous Russian nouns, that phonetics (stress), perceived origin, and semantics were at least equally influential in human beings making gender assignments. This is, apparently, a significant result for the conventions of Russian grammar, which, like German, is generally agreed to assign gender principally on morphological grounds (which is a way of saying, a kind of lost and forgotten history).

Finally, the Duke Project is exploring the possibility of a Chinese “Pseudo-Parser”, a network which would be happy with learning sequences, retain information about previous states of a sentence at a decaying rate, learn what the allowable sequences of word-classes may be, and predict, with enough evidence and training time, the probability of what the next word-class would be for any given phonetic input. This strategy is meant for the disambiguation of phonetic inputting of Chinese.

The network is exposed to the string directly at a certain activation level but also at a variable but lesser decay rate, so that “memory” of one input string will be passing through the network at the same time as but at a lesser rate than the next input string. The network is thus be training on sequences of word-classes which may be as long as two or three while it is learning what the next state of the sentence may be. One can

picture the process as a window of two or thee slots passing across a sentence and learning what the allowable and then more or less probable sequences are in the dataset. If the dataset is truly "typical," and the network functions as envisioned we should be able to take its results and implant them in a Chinese editor at a cost of no more than 6 or 7k. If the feature were turned on, the inputter would, when reaching an ambiguity, be presented only with those alternatives the network learned as allowable at that point of a "typical" sentence, and those alternatives would be presented in order of probability. Alas, the decision to employ a two-byte, UNICODE compliant environment for the Project's inputting of Chinese has delayed final implementation of this scheme, but we trust, not indefinitely.

During the first flush of excitement over the possibilities of strategies alternative to traditional AI approaches, exciting preliminary and experimental results were produced across the field of ANNs. Just as a kind of depression was setting in, waves of new fashions rolled over the same beach, "fuzzy logic" and more recently "Chaos." This may give ANNs the appearance of an old fashion, a kind of hula hoop, bell-bottom trousers, or polyester leisure suit. But actually the History of Science may assign ANNs another role, regardless of their eventual success, and that would be role of the or an opener of the flood-gates for more or less popular understanding of non-linear, not "rule-based," (unless we redefine "rule"), not "symbolic" (unless we redefine "symbol") models of things in general, including language.

As early as 1988 (Hubert L. Dreyfus & Stuart E. Dreyfus, "Making a Mind Versus Modeling the Brain: Artificial Intelligence Back at a Branchpoint." *Dædalus* [Winter 1988], pp. 15-43) the brothers Dreyfus warned of the failures of rule-based AI when confronted with the "combinatorial explosion": even if the presuppositions are so, and that a rule or theory can be articulated for each and every "symbol" (if they exist), then the interplay of the rules will sooner or later exceed the power of even the largest imaginable computer to process. They warned that precisely the same limitation faces connectionist models, perhaps even much earlier than is the case for their symbolic counterparts.

Does this mean then that the laborers in AI fields, particularly in the corners thereof as applications intensive as CALL, are doomed to be laboring in vain? The applications of Loritz, Sanders, DeSmedt, and Bailin, among others, would indicate that this is far from the case. What is perhaps more important, however, even than these successes, is the opening on the horizon provided by ICAI, an opening that always sheds light on the rest of the strenuous day-to-day business of CALL. It may

be possible that the principles behind our practices may some day make a lasting contribution to human understanding of language itself. That possibility alone, and not the utility of an application, justifies stubborn persistence in the work of intelligent CAI.

USIA Report
Attachment 3

Multi-Media Technologies and Second Language Acquisition.

University of Nijmegen; Wednesday 31 March 93

Given the record of Americans in second language acquisition, and the record of the Dutch in the same area, the idea that an American is coming to the Netherlands to lecture the Dutch on second-language acquisition is about as plausible as Russians coming to lecture them on democracy.

Let the stress of my talk, therefore, fall on the first part of the title of this talk and let's leave the second part as subsidiary.

This area has a history but in many respects it is entirely new. There are risks involved in being the first on the field.

Not too many years ago, "multi-media" meant a package that somehow combined text, graphics, audio, and maybe video, maybe computer based exercises, in some kind of coordinated package. It sounds like "multi-media," smells like "multi-media," must be "multi-media." Nowadays that is no longer the case. The term refers explicitly to computer based interaction with sight or sound media.

When used in this sense, multi-media for language instruction is about ten years old. The first such project to reach the stage of marketability was at produced Brigham Young University, observed the metaphor of an "Adventure" game, was meant to teach Spanish, and was called "Montevidisco."

The most recent to reach the market in the United States was developed at the CIA, is called "Exito" and is based on an "InterLex" (Donald Sola, Cornell) or "HyperText" or "Book-Shelf" metaphor, meaning that they provide on-line: lexicons, grammar checkers, idioms, cultural information, etc.

Both have in common that they require a lot of equipment, that that equipment is best presented in a dedicated mode (that is, not used for anything else), and shouldn't be moved around a lot. It is inter-connected with a mass of vermicelli, every strand of which is an invitation to disaster.

The original "Montevidisco" station included a Sony CPM computer, a Videodisc player, an "InstaVox" (a randomly accessible 14" floppy with an analogue re-recording of the video's audio source at different paces and clarity of articulation), separate monitors for text and video, and external speakers. The authoring system was written ad hoc and still exists as a monument to complexity. Montevidisco has been updated for the Macintosh environment and is probably a pretty nice learning environment by now.

Even so, the point needs to be made: modern multi-media, i/av, is not in every classroom or in every copy-shop around the world for a reason -- the hardware environment is still fundamentally too complex. This may not be the case forever or even for very long, but it has been the case up until now.

My personal contributions to the world of IA/V have thus far been limited to the transmission of the following joke which, I am told, has been quoted so often at CALL (Computer Assisted Language Learning) meetings in Europe that people are beginning to refer to it by short-hand, as the "Demo story"--

> A software developer dies and goes to heaven. After a short period there he finds it to be just a bit boring, and he asks St. Peter if he can check out the other place. St. Peter readily agrees. The software developer goes down and finds swaying palm trees and tropical breezes, beautiful women and sumptuous banquets wherever he looks. He returns to St. Peter and gives him his decision, he going to the other place. And St. Peter readily casts him to Hell, where now the software developer finds fire and brimstone, torture and agony, Bosch and Breughel. He complains bitterly to Satan, the devil, "what happened to the swaying palm trees and tropical breezes?" Satan replies, "You're a software developer, you ought to know: that was just the demo."

The point of the story, in this case, is that modern multi-media, IA/V, seen once working properly, will take your breath away and persuade you, at once, it is the solution to any imaginable communication and learning problem. The realities, on the other side of that first impression, are still daunting; and the technological problems are the very least part of the problem.

In fact, "as we speak," piece by piece, these problems are bring solved. IBM has announced, but not released a product called the M-wave. It is a board that can send and receive high resolution faxes, perform modem communications at high speed,

compress and present graphics, record (digitized) sound, speed up and slow down without loss of pitch or timbre (that is what the "InstaVox" was there to do), reproduce electronic music (MIDI) files, recognize speaker dependent, articulated (space-between-words) speech, display still video, present three-dimensional graphs of sound in real-time, and answer your telephone. If it can be digitized, this board can present it (except for motion video, and that, too is being addressed). Whether this particular technology will win out is hard to say. But it already exists and faces now only the test of the market place.

Intel has announced and slowly put into practice a standard called DVI (Digital Video Interface). The advantage of this over others is that it can be realized in software and does not need an extra board or exterior unit. It is the first fulfillment of the promise of universal digitization, including full screen, full motion video.

It is only a matter of time, and that time may be very short, before all of these potentials will be realized on a chip or a couple of chips and be part and parcel of every computer that comes off the assembly line. The interactive video tower will disappear altogether and just be the normal computer everyone has on the desk. This does not, however, mean that everything, all of a sudden, is going to be swaying palm trees and tropical breezes.

The fundamental problem may be signaled, right now, by one easy convention employed just about in every place where modern multi-media are now implemented, and that is: the mouse. The chief inputting method for users is "point-and-click." There is no intrinsic technological or media reason for this convention other than that it is easy and limits the exploration to the avenues intended by the author.

Let us for a moment look more closely at "point-and-click" as an inputting method. What does it resemble most closely in other activities of the day? How about the infra-red remote control of the television set? Just how "interactive" is the activity of lying on the couch and editing one's own television entertainment from the channels available on cable or satellite dish? Is it really interactive? How about "inter-passive" as a name closer to the truth of the matter?

When I suggest to my colleagues a one-year moratorium on "point-and-click" they react first with spontaneous laughter, and then with growing horror at the implications of such a prohibition.

Why should anyone be so sadistic as to suggest a one-year moratorium on

point-and-click? It is short-hand for suggesting that educational technology needs to be less technology-driven and more education-driven, and that the basic educational premises need serious investigation and reflection. One should not use "point-and-click" unless one is truly convinced that the infra-red remote is a good and useful precedent or model for learning.

It is at this point that second-language acquisition theory enters the scene. If your theory holds that manual recognition is sufficient to language learning, point-and-click is okay. If, however, you hold that language production is central to all aspects of language learning, even the demonstration of reception skills, then point-and-click is not enough, and you have every right to demand language production of the technology. To be fair, the two products I have mentioned, "Montevidisco" and "Exito" permit students to record their voices and play them back, but no speech recognition is involved. Otherwise, student so-called "interaction" is point-and-click, point-and-drag, multiple choice, true/false, in other words, technology-driven and not education-driven.

In point of fact, the capacity for better or worse response judgment has been available to the technology for ten years or more. The very techniques by which multiple choice or true/false responses can be judged correct or incorrect can be extended to word-length or line-length responses. The paradigms are Jim Pusack's "Dasher" (University of Iowa), "wild-card" schemes (the same tool as lets you summon a DOS directory of all files beginning with 'b', "dir b*.*), and the so-called "edit distance algorithm" (which builds a grid between the student input and the desired response and recognizes, discretely and separately, errors of capitalization, wrong letters, missing letters, added letters, duplication, and inversion -- instantly and without further author intervention).

Even more sophisticated response analysis has been developed over time, products in the categories of grammar checkers and style checkers. They remain a bit controversial, probably because the trade-offs between completeness and generally available computing power have not yet allowed for a market leader to set the pace of development. There is a great deal of experiment going on in the realm of "intelligent tutoring systems"— many people are concerned about the inadequacy of present-day inputting fashions. These incline, however, still to be slow and imperfect even on the super-powerful engineering work-stations on which they are being developed.

One does not, however, require fast and perfect, "total" response analysis to permit linguistic input. The aforementioned string matching strategies will suffice

nicely, thank you, and are infinitely more satisfactory that the wholly unintelligent alternatives now prevailing. There is a little bit of pure laziness involved. It is easier to employ "point-and-click," both from the programmer's and the course-author's point of view. The question of benefit to the learner seems to get forgotten en route.

Available technology can be made to satisfy almost any second language acquisition theory or strategy. The remaining technological tasks are relatively few and are being addressed. I believe them to be three:

1. encoding of languages not readily representable in the Roman character set. This problem is being addressed by the world-wide UNICODE consortium and will be adopted by the newest generation of 32 bit technology coming out in the next few years. This will permit electronic information exchange in all the languages of the world. A computer will know this character is Hebrew and encoded as aleph and not, as is now generally the case, a Roman 'a' parading as a picture of a letter in another character set.

2. presentation of full screen, full-motion video (30 frames per second) digitized and compressed on the hard disk along with all other data--text, graphics, and audio--and retrievable in real-time.

3. speaker-independent, continuous (not discrete) speech recognition. There is considerable controversy as to whether this will ever be possible. But the advances made in the last few years in combinations that fall short of this ideal provide hope that it is a problem that can be solved in due course by brute force, by sheer overwhelming computing power. There exist already cheap and easy speaker-independent discreet speech (pause between words) recognition devices, as well as speaker dependent continuous speech recognition devices. Now the two have to be brought together.

Personally, I cannot envision any other tool that would be required for the perfect inter-active multi-media delivery station. Does that mean that all out learning and teaching problems are basically solved? In fact, such developments point not to the end of a technology revolution but to the beginning of an educational revolution. That revolution will entail a dynamic and dialectic relationship between educational strategies and human-machine ergonomics or "Human Factors."

Second-language acquisition theory may suggest a possible set of activities as a learning strategy, say, for reading:

1. Preliminary Activities (Prediction and Previewing)
2. Preliminary Processing (Skimming and Scanning)
3. Decoding activities (Semantic chaining, lexicon recovery)
4. Total comprehension activities (Gisting, summarizing, translating)

A computer can be made to provide media for all of these activities, with many advantages over hard-copy, teacher dependent environments (instant response analysis, general, out-of-class availability). But what should the electronic media look like? Should they look like paper? Almost certainly not. What are the "rules" of electronic presentation a) as they differ from paper and b) as they conform to learning strategies? Frankly, no one knows yet.

There will be surprises. Imagine the following statement: "My child has been watching television since he was three. He is completely illiterate." A hearing impaired corresponent of ours employs "closed-captioned" TV, that is, sub-titles. His child wins all the reading prizes in school. "My child has been watching television since he was three and he is supremely literate." Discoveries like these are likely to be among the surprises.

There are ample studies of legibility, aesthetics, and information retention--none of them as yet conclusive and none of them genuinely informed by broad educational theory and still less by learning strategies. It is a brand new field, and those who enter upon it have a great many discoveries to make.

USIA Report
Attachment 4

Electronic Media and Educational Revolution

A Dialogue
followed by
A Lecture-Demonstration
followed by . . .

Idiota: Good Morning, Master Pancognitus. It is a pleasure to see you on such a fine day. May I ask where you are rushing in such a hurry?

Pancognitus: I have no time to stop and chat. Please come along. I don't want to be late for the disputation.

Idiota: What disputation would that be?

Pancognitus: Why, Masters Dollenkopfius and Schlauraffius will be debating a marvelously refined point, haven't you heard?

Idiota: How refined?

Pancognitus: Their disputation concerns the disputation itself!

Idiota: A disputation on disputation!

Pancognitus: Yes, now that you mention it. But actually they will debate many questions: whether poetry, letters, oration, or disputation best demonstrate the learning skills of the learned.

Idiota: That confuses me thoroughly. What, for example, does poetry have to do with learning?

Pancognitus: How critically acute you are! You have always been the best of my students. What, indeed, does poetry have to do with learning? This is a novelty the deceitful Italians have been trying to pass across the Alps for a hundred years, but no worry, we are having none of it.

Idiota: But do not you yourself have us memorize those famous passages from Virgil, Horace, even wicked Ovid, and unmentionable Catullus?

Pancognitus: But that is only to help you memorize. The regularity of verse eases the way for the lazy mind. You may adorn your speech with those famous moments, but it is only the frivolous who would waste their time pursuing poetry any further.

Idiota: It would seem that the private letters of famous men might have more to do with learning than poetry.

Pancognitus: Well, if you insist. The private letters of famous, even of obscure men, are all well and good if they show off the erudition of the writer and his skill as an orator. But realize, private letters are *written* and not spoken. Any scribe can write. But only an orator can speak well.

Idiota: But has not writing become somewhat more important since paper has come on the scene? Our pupils may still memorize the grammar of Donatus, but the children of the prosperous are taking notes on paper and studying them at home, away from the feet of the master.

Pancognitus: And a very unhappy development it is. By using paper, their memories will never be properly developed and without memory, how will they memorize the great models of the ancients? And without the models, how will they speak well?

Idiota: Is it possible that, one day, what students write will be used to judge their skills instead of how they speak?

Pancognitus: I cannot imagine it. Writing is only a craft, like carving wood. If a town needs a wood-carver it will hire a wood carver. If a town needs a writer, it will hire a scribe. There is no more need for everyone to be his own scribe than for everyone to be his own wood-carver. It is the *spoken* word that makes the master.

Idiota: Well then, orations are evidently the final from in which the learned can demonstrate their skills.

Pancognitus: You are almost on the mark. In oration, the speaker can show off all the appropriate skills: memory, invention, erudition, persuasion. In a spoken world, one can ask for little more.

Idiota: Why then bother with disputation?

Pancognitus: The answer to that is quite simple. A disputation is an oration enhanced by argument. Like athletes at their best in competition, the disputation pits orators against one another in a struggle. It brings out the best in the combatants. It combines all the skills of speaking well with those of the struggle. Disputation is a contest of orators. And the best man always wins.

Idiota: As always, Master *Pancognitus*, you win!

Pancognitus: Discretion is the better part of valor.

Idiota: The disputation will therefore always be the premiere form of learned discourse.

Pancognitus: Always.

Idiota: It will never change?

Pancognitus: Never. Why should it?

Idiota: Have you ever heard of Laurentius Janson Coster and the firm of Gutenberg, Fust, Schöffer & Co.?

Pancognitus: The famous bankrupts? Who has not heard of them.

Idiota: What do you think of their invention?

Pancognitus: It is a common failing of the young to be taken in by every new fad or fashion. What is there to think of this (*with contempt*) novelty? It is no more than the written word written many times. Please, pick up any printed book you want. What is it made to look like? Well?

Idiota: I'm not sure I fully understand what you are getting at.

Pancognitus: They look like manuscripts. The printers even leave open spaces around initials in the hopes that one day a miniaturist will illuminate the letters. Printing is only competition for the scribes. And, as I have told you, not everyone needs to be his own scribe.

Idiota: That may be true of most books, but have you seen the new books from the hand of Jenson and the House of Aldus in Venice? The types are crisp and clean and compact, finer than any scribe can write, and readable. You can hold the complete works of Livy in one hand!

Pancognitus: Another Italian novelty, another fad, another fashion! For a thousand years, we have had no need of the complete works of Livy. Memorizing the famous excerpts has always been enough, and I see no reason for that to change.

Idiota: But not only I can hold the works of Livy in one hand. Where they have been perhaps ten or twenty copies hidden on the dusty shelves of monasteries, now a thousand others can hold the works of Livy in one hand, or ten thousand. And we'll all have exactly the same words on exactly the same pages. We can write each other and be sure we are referring to the same text.

Pancognitus: What difference that should make is altogether beyond me.

Idiota: Master Pancognitus, your golden words are now heard by one hundred students a semester and lost upon the air the moment they are spoken. Imagine your wisdom inscribed on a thousand copies, or ten thousand, and broadcast across all the world, from Rotterdam to Cracow, and living beyond your life as long as paper lasts!

Pancognitus: I'm not the least bit sure the vulgar masses of Rotterdam and Cracow have any business dealing with my wisdom. Who knows what nonsense they might say or think or write because they have read or heard my printed word. My audience needs to hear me speak, to recognize my gesture and intonation. All of that is lost in print. It will never last. Do not be deluded by fashion, my simple-minded friend.

Idiota: If some of your form is lost in print, is not your substance, the ideas, the knowledge, the wisdom, preserved in the logic of argument, and is that not wholly preserved in the word? and that not wholly present in print?

Pancognitus: Don't be insolent. The world in which we live and work is spoken. It is not written. It will never be written. It will always be spoken. And now you've made me miss my disputation.

Idiota: I wouldn't worry. The Gerard Leeu is printing the transcripts in Gouda. You can read about it tomorrow in the broadside the students are having printed. Farewell.

****

Having the advantage of hindsight, we now know that *Pancognitus* was not altogether wrong. Good speech retains undiminished importance in the conduct of the affairs of the world. Even the oration survives in a few ceremonial moments, public and academic. And adapted to the modern classroom and seminar, it lives on as the chief means of formal instruction, the lecture-demonstration. But it is no longer the exclusive measure of learning. In most cases the written word, and what is more, the written word in print, has become the measure of learning and the means of dissemination. Some form of disputation survives in academic (and political) life, but almost exclusively in print and only occasionally in dialogue (the question-and-answer session after a typed paper has been read aloud).

How long did it take for the university to adjust to the new technology? Frankly, it seems to have been slow in coming, maybe as long as two hundred and fifty to three hundred years. The students seem to have adjusted first. In the sixteenth century the students demanded and got text-books. This provided them access to the knowledge they needed to acquire at times when they did not have their professor before them. The textbooks were compact, relatively inexpensive, randomly accessible, portable, usable in private. But the students were still memorizing what they read. This is, however, radically different from memorizing what they *heard*. They took the step across the abyss to the new technology first.

It is amazing that the vanity of the professors did not drive them into print immediately. Idiota's argument, that the wisdom of *Pancognitus* would be available to the masses, found no resonance. This probably has to do with the unabating snobbism of the university, which has always been aware that it exists in order to educate the elite. A sure way to guarantee an early end to an academic career is to acquire the designation of “popularizer.” Academic knowledge is produced by experts for the edification of other experts. But the slow pace of academic acceptance had probably also to do with the reward system. What we call “authors” were still “writers” in the sixteenth and seventeenth centuries. That is very hard to grasp in our post-Romantic haze. Writers were craftsmen and not yet “poets.” They were no more or less important to the great publishing houses than engravers or typesetters. Hans Luft the publisher, not Luther, got rich printing the German Bible.

There are a few precedents in the seventeenth century, but it is not until the first modern attempts at university reform in the eighteenth century, under Pietist leadership, that published research became the measure of academic achievement. By

the middle of the eighteenth century, academic publishing is big business. Professors begin routinely to publish their courses, no doubt encouraging their students to buy the "hard copy" version to be assured a complete understanding of the professor's wisdom (and to supplement his meager salary). "Publish or perish" was installed and has remained enthroned from then until now, some two hundred or two hundred and fifty years.

The lecture-demonstration, although it has its roots in an oral culture became "literate," or "typographic," when it became one station in the cycle of academic learning, where all other stations were in print.

The question one needs to ask today is just how permanent, inevitable, and unchanging are the conventions of formal instruction in the time in which we live. The dominant mode of formal instruction world-wide is still the lecture-demonstration. The seminar, in normal circumstances, is little more than a variant of the lecture-demonstration, where the students may be permitted occasionally to play the same role as the teacher. In ideal circumstances, the variant may move over to the disputation, a bit less formal now than the medieval and Renaissance predecessor, but a disputation nonetheless, where new knowledge can be achieved through dialogue, through dialectic. But any new knowledge thus achieved, or achieved any other way, is only "virtual" until it reaches print. Then it turns into a kind of potential reality, until it is read, when it is something that can be talked about, written about, read, talked about, etc.

To be ruthlessly honest with ourselves, how much of our new knowledge is acquired through print nowadays? And if we are incapable of this extreme of self-criticism, how much of the new knowledge of the young is acquired through print, as opposed to electronic media? Whether this is a happy development or not is entirely beside the point. It is wholly possible to argue that manuscript is superior to print, even to be right in that argument. It is still wholly beside the point. The art of printing superseded manuscript. Electronic media appear to be doing something similar to print. [A colleague of mine at my university says: "Well, the Gutenberg Blip appears to be over."]

If the "Gutenberg Blip" is over or at least passing quietly into history, what does that mean for those of us who depend on lecture-demonstrations for a living? The Provost of the most famous language school in the United States, the Defense Language Institute in Monterey, California, Ray Clifford, has often been quoted as saying "Computers will never replace teachers; but teachers who use computers will

replace teachers who don't use computers." It is far from certain that he is right. Those may be in the right who fear new technology as a competitor who threatens their very livelihoods. Examination of one's professional role in society never hurts. It may be that we are near the point of having to do so for our own survival.

For those of you who feel unthreatened, I have brought along a tape:

[The IBM "Ulysses" Tape]