Canadian Artificial Intelligence

Intelligence Artificielle au Canada

Canada's National AI Publication | *La Publication Nationale en IA au Canada*

Spring 1999 | *No. 43* | *printemps 1999*

Contents / Contenu

News / Nouvelles

March 1999: Space robotics division of Spar Aerospace purchased by MacDonald Dettwiler

February 1999: Boost to research funding in federal budget

December 1998: European Community launches four-year AI funding program

December 1998: CAnet 3 operational by mid-1999, connected to Internet2 by end-1999

... continued on page 6

Web Version Web

This issue is on the web via http://cscsi.sfu.ca/cai.html – give your userID and password at the login window. For information about your userID and password

(1) **don't throw away** the **envelope** in which CAI/IAC arrived and

(2) please read the about access to CSCSI/SCEIO Website on page 2 (over page).

Canadian Artificial Intelligence is published by the Canadian Society for Computational Studies of Intelligence (CSCSI). ***Intelligence Artificielle an Canada*** est publiée par la Société canadienne pour l'étude de l'intelligence par ordinateur (SCEIO). Canadian Publications Mail Product Sales Agreement No. 1497200.

ISSN 0823-9339

Editor's Message

Dan Fass

Dear reader:

After many difficulties, issues 42 and 43 of *Canadian Artificial Intelligence / Intelligence Artificielle au Canada* (hereafter *CAI/IAC*) have been produced. Issue 42 should arrive within a few days of this issue. We apologize for the interruptions in producing the magazine and hope to return to a regular publishing schedule.

This issue (#43) and some past issues of *CAI/IAC* are available on the members-only area of the CSCSI/SCEIO website (http://cscsi.sfu.ca/cai.html). To access the area, you need to type your userID and password at the login window. Your userID is the first letter of your first name plus up to seven letters of your last name. For example, the userID for Pierre Elliott Trudeau is ptrudeau. Your password is based on your CSCSI/SCEIO membership number which is **printed on the envelope** in which this issue arrived. Take that number and prepend to it the first letter of your first and last name. For example, if Pierre Elliot Trudeau's membership number was `765432`, then his password would be `pt765432`.

Because of the delays in producing issues 42 and 43, we are making copies available to former members of CSCSI/SCEIO whose memberships expired in 1997, 1998, or early 1999. Unfortunately, it was not possible to synchronize mailing lists for issues 42 and 43, so some people who are no longer members may receive a hard copy of #43 but not #42. To cover this possibility, #42 has been made freely available on the CSCSI/SCEIO website at http://cscsi.sfu.ca/cai.html. If you do not receive a copy of #42 for any reason, please collect one from this URL.

The CSCSI/SCEIO Executive has decided to roll into one the jobs of editing and production of *CAI/IAC, and* administering the CSCSI/SCEIO website – jobs that for some years were done by different people. The Executive feels that one way to ease the workload of doing all three jobs is to somewhat simplify *CAI/IAC*. As a result, this issue of *CAI/IAC* looks more newsletter-like. As the first two jobs consume a lot of time, plans for improving the website have had to, unfortunately, be shelved for now.

I would like to thank the past editor, Suhayya Abu-Hakima, for the issues of *CAI/IAC* she edited alone and, earlier, with former co-editor Peter Turney. I would also like to thank the production team of Arlene Merling, Carol Tubman, and Greg Klymchuk for their contributions to *CAI/IAC* over the years.

After taking over as editor, I sought editorial direction of *CAI/IAC*, aware that CSCSI/SCEIO marked its 25-year anniversary in 1998, I turned to past issues of *CAI/IAC* for a historical perspective and editorial ideas. I was fortunate enough to find in Vancouver an almost complete collection of newsletters and magazines put out by CSCSI/SCEIO since its inception in 1973. Alan Mackworth was kind enough to lend me his almost complete collection of early newsletters from 1973 until 1984. All that was missing were two issues from the late 1970s or early 1980s. I would be very interested in getting hold of copies of these issues. They are:

CSCSI/SCEIO Newsletter, Vol. 2, No. 2 (produced after December 1978, but before December 1980).

Newsletter of CSCSI/SCEIO, CM-CCS, and CIPPRS, vol. 1, No. 1 (also produced after December 1978, but before December 1980).

Fred Popowich has a complete collection from September 1984 of what was then known as *Canadian Artificial Intelligence Newsletter*, when Graeme Hirst took over as editor. Reading through all these back issues, I was struck by the amount of work that had gone into them and the terrific historical resource that they represented.

That reading of back issues led to two articles. The first of these articles is "The Infrastructure of Artificial Intelligence R&D in Canada." The first part of this article appears in this issue; the second part will appear in the next. It is intended for two audiences. One audience is "old-timers" who were involved in Canadian AI in the 70s and early 80s, for whom it may provide some reminiscences. The second audience is those involved more recently, for whom the article will hopefully be a primer. I'm part of the second group and, during the writing of the article, I came to understand better many acronyms that I had heard over the years, acronyms such as CIAR, STEAR, and IRAP.

... continued on page 17

President's Message

Robert Mercer

... sitting in my office ... Middlesex College ... The University of Western Ontario ... trying to collect my thoughts, thoughts that will go into my upcoming President's Message. Pencil ... paper ... here, let's put my feet up, I do my best thought collecting with my feet up ... there, what a great day outside ... wish I were outside ... oh, yeah, President's message ... oh, my! last year, 1998, was the 25th anniversary of CSCSI/SCEIO. Twenty-five years ... more than half my life ... let's sit back ... relax ... collect more thoughts ...

... yes, twenty five years ago a small group of Canadian researchers got together on a holiday weekend at Western and formed a Steering Committee which in time became the executive of CSCSI/SCEIO. I was but a first year graduate student in 1973. Western's Ted Elcock became the first chair of this Steering Committee. Now, me, another Western faculty member is in that (since renamed) position. Hmmmm ... 25 years ... the circle has closed. Of that small group some have retired, several hold senior positions at Canadian and American universities. Just think! Some of my former fellow grad students are department chairs, faculty deans, ...

... what a great day outside ... at least I have a nice view out my office window. Trees are still leafless. I can see across to Delaware Hall ... why that was the residence for CSCSI'84 ... that was the first time that I visited this campus ... gave a paper ... met so many people from across Canada, the US, ... BONG, BONG, ...

Reality check! The Middlesex College clock tower ... 11 ... 12 ... it's noon. Better get my president's message into draft form before going to lunch. Now where was I, yes, the challenges, the opportunities ...

What was I thinking about ... oh, yes, the people. That's what is important about this society, the opportunity to meet people, face-to-face and through its publications. It was this organization which introduced me to so many of you, and it all started when I was a graduate student ... yes, this is the idea that I want to discuss ... the future generation becomes the current generation ... This is the biggest challenge that faces any society — how does it cultivate the next generation, the future of the society. Good. But there was something else ... ah, yes ... I've been reminded on occasion that the society was named CSCSI/SCEIO because the original view was to be broad in the aims of the society. The original group was comprised of computer scientists, psychologists, engineers, and mathematicians. Well, AI has redefined itself a number of times. Some of what was AI research a decade ago has renamed itself. What was AI research a short time ago is now just a piece of a business application that no one considers

AI any longer. Hmmmm ... an opportunity exists to reexamine who we are, rediscover our roots, and remake ourselves. We have an opportunity to create stronger links between academia and business.

AI/GI/VI '2000 to be held in Montreal in June 2000

Challenges and opportunities ... this is the future. I had better say something about the present ... We are currently organizing for our next biennial conference. It will be held in Montreal in June 2000 in the now usual AI/GI/VI format, but in conjunction with ISR'2000 an international robotics conference and the IRIS/Precarn conference ... Much of the efforts of the executive over the past few months has been to improve communication with the membership. Although the past year has proven a bit bumpy, we have finally produced what you have in your hands. The herculean efforts of our (new and) most capable editor, Dan Fass, have produced this edition of the CAI/IAC magazine ... he has others in the pipeline ... oh yes, I should introduce the rest of the executive ...Fred Popowich (the-never-did-he-think-that-he-was-going-to-be-so-busy-as Past President), Russ Greiner (Vice-president), Howard Hamilton (Treasurer), and Guy Mineau (Secretary) ... these three people have already indicated that they will play an important role in addressing the challenges and the opportunities that face the society.

Challenges and opportunities ... I'd better get some emails off while these ideas are fresh ... I'll have to tidy up this draft later ...

Treasurer's Report

Howard Hamilton

This report is for fiscal year 1998. CSCSI/SCEIO had three major financial activities in fiscal year 1997: holding our biennial conference (AI'98), member servicing and the production of one issue of the Canadian AI magazine. As is normal during a conference year, income exceeded expenses. The balance will be used to cover expenses during the following non-conference year.

Opening Balance Sheet, 1 January 1998

Assets	
Bank balance	40,246.61
Sept. membership cheque	140.24
Dec. membership cheque	328.46
Uncollected GST refund	27.38
Prepaid deposit on AI'98	4,000.00
Total assets	$44,742.69
Liabilities	
Uncashed cheques	35.00
Unpaid scholarship	500.00
Unpaid GST	51.42
Total liabilities	$586.42
Capital (assets - liabilities)	$44,156.27

Income

Memberships, including some multiyear	4,497.15
Interest income	718.58
AI'98 conference	10,426.70
Net GST	308.62
Total	$15,951.05

Expenses

CAI magazine, Issue #42	3,090.20
AI'98 conference	7,256.85
Website development	2,083.33
Total	$12,430.38

Income - Expenses $3,520.67

Closing Balance Sheet, 31 December 1998

Assets	
Bank balance	21,299.30
GIC	30,000.00
Dec. membership cheque	202.37
Amount owed by CIPS	22.95
Total assets	$51,524.62
Liabilities	
Unpaid expenses: GST, website, printing issue #42	3,812.68
Uncashed cheque	35.00
Total liabilities	$3,847.68
Capital (assets - liabilities)	$47,676.94

Change in capital = New capital - old capital $3,520.67

Membership in CSCSI increased from 133 to 146 in 1997, partially as a result of a successful drive to increase memberships at AI'98. The society agreed to spend more than it earned in the two year period starting July 1, 1998 to further increase the number of members and the services they receive. Funds were earmarked for updating the website as well as providing more frequent magazine issues.

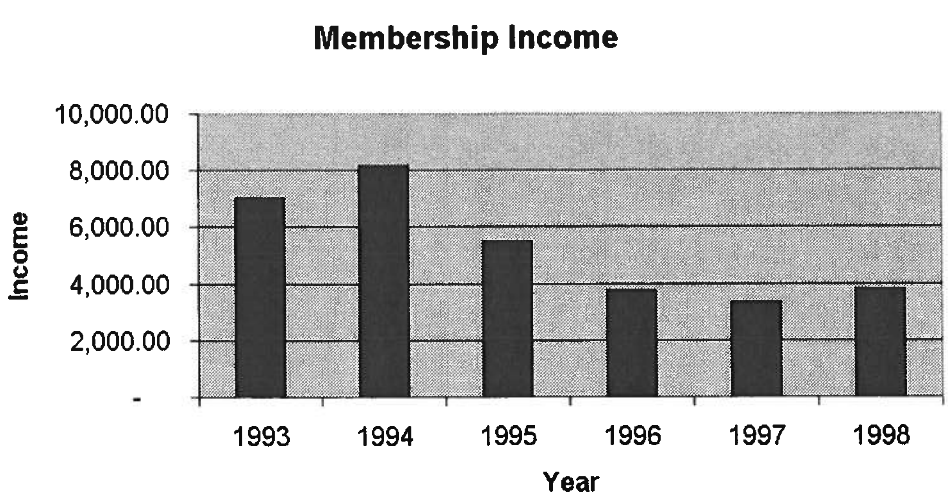

Secretary's Report

Guy Mineau

Minutes of the 1998 CSCSI/SCEIO AGM

Held 18 June 1998 at Simon Fraser University Harbour Center in Vancouver, the meeting started at 17h30 and ended at 18h30, during the Canadian AI'98 conference. About 20 people were present throughout the meeting, including all executive members of CSCSI/SCEIO and the chairs of AI'98: Fred Popowich (president), Stan Matwin (past-president), Renée Elio (vice-president), Guy Mineau (secretary), Howard Hamilton (treasurer), Sue Abu-Hakima (editor of CAI/IAC), Robert Mercer and Eric Neufeld (co-chairs of AI'98). The eight agenda items adopted in the meeting are described below.

1. CAI/IAC Magazine

Due to the actual production costs, the time constraints of the editor Sue Abu-Hakima and of Arlene Merling, the producer of the CAI/IAC magazine, it was decided that the magazine would become a newsletter. That this newsletter would be sent regularly to all CSCSI/SCEIO members (4 times a year), and that an electronic version of the newsletter would be made available on our web site. This newsletter will be produced at SFU under the supervision of Fred Popowich. It will include fewer articles but many AI related news. The idea is to strengthen the link between CSCSI/SCEIO and its members.

2. Membership

After a constant decrease in membership, it is now stable at 133 members. We will eventually have to deal with this issue, and come up with ways of improving it. For one, we should look into additional discounts for scientific journals to which CSCSI/SCEIO members subscribe.

3. CSCSI/SCEIO Executive Nominations

Nominations needed to be submitted to the president of CSCSI/SCEIO, Fred Popowich, through electronic mail at popowich@cs.sfu.ca, by 30 June 1998.

4. Distinguished Service Award

It was decided to maintain the award even though it was not awarded this year. A call for nomination will be sent to all members through the newsletter for the year 2000 award.

5. Conference Report

There were 71 submissions, from which 35 papers were selected for the conference. Springer published the proceedings and seem to be interested in repeating the experience in year 2000. A consensus among CSCSI/SCEIO members exists to the effect that the AI conference should remain a biennial event.

The joint venue with the IRIS conference did not produce additional registrants for AI'98, as was originally anticipated. Fewer people attended the IRIS conference to begin with, and many of the IRIS attendees now tend to focus their work more toward robotics.

Keeping that in mind, we plan on repeating the experience with a bigger crowd, having the next AI conference in Montreal jointly with GI and VI as usual, but also jointly with (overlapping and on the same premises of) the ISR-2000 conference (the 31st International Symposium on Robotics conference), from 14–17 May 2000. We hope to provide this conference with some insights on fundamental and pragmatic aspects of artificial intelligence, particularly with regard to robotics; it may be of some interest to them. Discussions with the organizing committee of the ISR-2000 conference will take place early next fall so that we know early on what kind of financial arrangement can be negotiated with them with regard to co-locating with them.

6. Web Page

The *cscsi.org* URL will be sought; it is a small fee to ensure continuity and stability of the link to our web site. A link to our web site will also be made available from the CIPS web site, if possible. Fred Popowich will look into these two matters.

7. Financial Report

A financial report report was presented. In brief, this year's conference will probably make a small profit (still to be confirmed); while the society will make a small deficit.

8. Other Business

8.1 Travel Grants

Guy Mineau has volunteered to follow up on the students who were awarded a travel grant to attend and present a paper to IJCAI-97 in Nagoya last year, to see if the money they received from IJCAI did mention that $500.00 came from CSCSI/SCEIO. It was decided to keep the travel grant program and to advertise it in the first issue of the newsletter.

8.2 Other

Alan Mackworth moves to congratulate the chairs and the executive on the excellent work done in organizing the AI'98 conference. Unanimously carried.

News
Nouvelles

19 March 1999

Robotics division of Spar sold. The Space Robotics Division of Spar Aerospace (http://www.spar.ca), maker of the Canadarm, was purchased by MacDonald Dettwiler and Associated Ltd. (http://www.mda.ca) for $63 million. MacDonald Dettwiler is owned by Orbital Sciences Corporation (http://www.orbital.com) of Virginia.

12 March 1999

New Canadarm ready. The robotic arm has been developed and tested at Spar's labs in Brampton, Ontario and will shortly be moved to NASA's Kennedy Space Centre.

16 February 1999

Boost to research funding in Canadian Budget 1999. The Natural Sciences and Engineering Research Council (NSERC), which funds computer science research, as well as the Social Sciences and Humanities Research Council (SSHRC) and Medical Research Council (MRC), will receive $405 million in extra funding over the next three years.

The Networks of Centres of Excellence (NCE) program will receive an extra $90 million over three years. This should support as many as eight new networks.

Technology Partnerships Canada will receive an additional $150 million over three years, beginning 1999-2000.

The1999 budget booklet *Building a Stronger Economy Through Knowledge and Innovation* (http://www.fin.gc.ca/budget99/ecoe/ecoe.html) contains more information on these funding increases.

1 February 1999

New BC law relaxes high-tech work restrictions. British Columbia's Technology Minister Andrew Petter announced that restrictions have been removed on hours of work, overtime and statutory holidays for high-technology "professionals" in, for example, software development, engineering, and animation. The aim is to make BC's high-tech industries more competitive.

22 December 1998

European Community launches Fifth Framework. The Fifth Framework program, or FP5 for short (http://www.cordis.lu/fp5/) is a Europe-wide R&D program in high technology for 1998-2002. It succeeds FP4, which ran 1994-1998. Its budget is 13,700 million euro. AI research falls largely under its theme 2 (Creating a User-Friendly Information Society), also known as the Information Society Technologies (IST) thematic programme, which is to receive 3,600 million euro. (As of going to press, 1 euro = 1.617 Canadian dollars.)

15 December 1998

Progress with CAnet 3. Officials at CANARIE (the Canadian Network for the Advancement of Research, Industry and Education) announced an agreement to connect CAnet 3 with the US Internet2 (also known as Abilene) at Chicago by the end of 1999. CAnet 3, which will be operational by mid-1999, will be the world's first fibre-optical network of its kind. It will operate at 40 gigabits (GB) per second. Abilene is aiming for 2 GB per second.

(In 1993, the original CAnet operated at 56,000 bits per second. CAnet 2, which began in 1997, was 3,000 times faster at 155 Megabits per second. That capacity has since been doubled.)

5 November 1998

SFU economists to study Canadian AI brain drain. Three economics professors at Simon Fraser University have been awarded $135,000 for a two-year investigation of the brain drain from Canada's high-tech industries to the US. The study will include those working in AI, robotics, telecommunications, and high-tech medicine. See http://www.sfu.ca/mediapr/sfnews/1998/nov5/braindrain.html

8 March 1998

NASA renews Spar's Canadarm contract. NASA has given the Space Robotics Division of Spar Aerospace a five-year, $91.5-million contract to continue development of its Canadarm. Four Canadarms are currently in use, all built by Spar.

24 February 1998

Increased research funding in Canadian Budget 1998. The NSERC, SSHRC, and MRC will receive a 14% rise in funding. NSERC's then-current budget was $434 million. It was to receive $494m instead of $423m in 1999-2000 and $501m instead of $416m in 2000-2001.

Learning to Uncover the Black Soup Pot

Ken Barker

Résumé

Cet article décrit un système semi-automatique pour regrouper les modificateurs (noms et adjectifs) d'un nom en anglais. Quelques problèmes de regroupement de ces modificateurs sont présentés, l'algorithme de regroupement est décrit et son fonctionnement est évalué.

Abstract

This article describes a semi-automatic system that can bracket noun phrases containing an unlimited number of premodifying nouns and adjectives. Problems bracketing noun phrases are presented, the bracketing algorithm is described and its performance is evaluated.

Introduction

__Black soup pot cover__: is that a black cover on a soup pot or a cover on a black soup pot?
Does a __large clothing store__ sell a lot of clothes, or only those of a certain size?
Is an __advanced calculus professor__ really that much more evolved than the rest of us?

The examples in bold are not just a source of wordplay, they pose a serious problem for natural language processing. Before an NLP system can attempt to interpret the meaning of complex noun phrases, it must first untangle the modification pattern of the head noun's premodifiers. The modification pattern is often called a "bracketing" of the noun phrase; and determining the pattern is often called "bracketing." (!)

This article describes a semi-automatic system for bracketing noun phrases containing an unlimited number of premodifying adjectives or nouns. Since the system is intended to start processing with no prior knowledge, it gets trained as it brackets. That is, it starts from scratch and accumulates bracketing evidence while processing a text under user supervision (see also Barker 1998a, Barker & Szpakowicz 1998).

Experiments show that generalizations of the structure of complex modifier sequences allow the system to bracket previously unseen compounds correctly. Furthermore, as more compounds are bracketed, the number of bracketing decisions required of the user decreases.

Noun Compounds, Complex Nominals and Bigger Things

Many papers that talk about bracketing deal with a head noun along with one or more premodifying nouns only. Such a sequence of nouns is often called a *noun compound.* Others allow certain kinds of adjectives as premodifiers and name the sequence a *complex nominal.* The bracketer I will describe has to deal with any sequence of an unlimited number of any kind of premodifying nouns or adjectives. I sometimes refer to the sequence of premodifiers and head noun as the *modifier sequence.*

Bracketing

Bracketing premodifiers consists of identifying all of the nested modifier-head pairs in the modifier sequence. Identifying these pairs requires decisions about whether a given sequence of three elements is *left-branching* or *right-branching* (where an element is a word or a nested modifier-head pair). The sequence in (1) is *left-branching* and has the bracketing shown in (2); (3) is *right-branching* and has the bracketing shown in (4).

(1) *laser printer manual*
(2) *((laser printer) manual)*
(3) *desktop laser printer*
(4) *(desktop (laser printer))*

Several people have proposed empirical solutions to the bracketing problem. Liberman & Sproat (1992), Pustejovsky *et al.* (1993) and Resnik (1993) take a similar tack: for a given sequence X-Y-Z, compare the number of occurrences of X-Y in isolation in a corpus (a large body of text) with the number of occurrences of Y-Z. If X-Y occurs more frequently in the corpus than Y-Z, X-Y-Z is taken to be left-branching. Lauer (1995) calls the model that compares X-Y to Y-Z the *adjacency* *model* and offers a different model, the *dependency model*. The dependency model compares the number of occurrences of X-Y to the number of occurrences of X-Z (instead of Y-Z). In Lauer's bracketer, the dependency model outperforms the adjacency model.

What's Wrong with Adjacency?

In fact, most experiments comparing the adjacency and dependency models find that dependency is more dependable than adjacency (see also ter Stal 1996, Jun & Changning 1998). Is it mere coincidence, or is the adjacency model more fundamentally flawed?

Consider phrase (5): both right and left bracketings are possible. Previous occurrences of *large store* would be evidence for right-branching. Occurrences of *large clothing* would be evidence for left-branching. Occurrences of *clothing store* could be evidence for either. (6) restricts the store; (7) restricts the clothing. But both refer to clothing stores. So occurrences of Y-Z in a text support *either* interpretation and tell us nothing about the bracketing of X-Y-Z.

(5) *large clothing store*
(6) *(large (clothing store))*
(7) *((large clothing) store)*

Reduced Modifier Subbracketings

Unfortunately, when trying to bracket the sequence X-Y-Z, it is often the case that the pairs X-Y and X-Z occur infrequently in isolation in a given text. In order to increase the chances of finding branching evidence, it would be useful to generalize the pairs and look for occurrences of the generalizations. Lauer (1995) generalizes the nouns X, Y and Z to the Roget's Thesaurus categories that contain them: R_X, R_Y and R_Z. Instead of looking for other occurrences of X-Y and X-Z in the text, Lauer's bracketer looks for occurrences of U-V and U-W such that $R_U = R_X$, $R_V = R_Y$ and $R_W = R_Z$. The technique limits generalization to nouns that occur in Roget.

An alternative to using a *semantic* generalization of the words is to try to find a *structural* generalization. Consider phrase (8) and a reasonable bracketing for it in (9).

(8) *dynamic high impedance vocal microphone*
(9) *(dynamic ((high impedance) (vocal microphone)))*

Each non-atomic element of each bracketed pair can be considered a *subphrase* of the original phrase. Given the bracketing in (9) the subphrases for phrase (8) are phrase (8) itself as well as the subphrases in (10).

(10) *high impedance vocal microphone*
high impedance
vocal microphone

Each subphrase consists of one or more modifiers and a head local to the subphrase. Local heads in (9) are *microphone* (the head of three subphrases) and *impedance*. The subphrases are generalized by reducing modifiers and modificands to their local heads. If this concept of reduction is applied to a bracketed phrase, the result is a set of reduced subbracketings of the original phrase. The reduced subbracketings of (9) appear in (11).

(11) *(dynamic microphone)*
(impedance microphone)
(high impedance)
(vocal microphone)

The reduced subbracketings together are a structural generalization of the original modifier sequence. Instead of simply memorizing complete modifier sequences and their bracketings, the bracketer I present stores the subbracketings. This allows it to analyze different modifier sequences that have only subbracketings in common with previous sequences.

The Bracketing Algorithm

The algorithm for noun premodifier bracketing handles modifier sequences of any length by dealing with a window of three elements at a time, where an element is a word or a bracketed pair of elements. It is shown in Table 1.

Confidence in Branching Decisions

The sequence noun-adjective-noun is confidently right-branching since adjectives precede the nouns they modify. The exception is postpositive adjectives, which occur relatively infrequently within premodifier sequences (see Barker 1998b for a more detailed discussion of the problem of postpositive adjectives).

For any other sequence of three elements X-Y-Z, the bracketer reduces X, Y and Z to their local heads X_h, Y_h and Z_h. The sequence X-Y-Z is considered confidently right-branching if the frequency of previous occurrences of the reduced subbracketing (X_hZ_h) is greater than the frequency of previous occurrences of the reduced subbracketing (X_hY_h). If (X_hY_h) has occurred more frequently than (X_hZ_h), X-Y-Z is confidently left-branching.

1	Start with the rightmost three elements, X-Y-Z.	...	V	W	X	Y	Z
2a	If X-Y-Z is confidently right-branching (see below), bracket it X-(YZ) and restart the algorithm with the rightmost three elements W-X-(YZ).	...	V	W	X	(YZ)	
2b	If X-Y-Z is confidently left-branching, move the window one element to the left and repeat the algorithm with W-X-Y. Note that X-Y-Z being confidently left-branching does not necessarily mean that it can be bracketed ...(XY)-Z, since left-branching may also be bracketed ...X)Y)-Z.	...	V	W	X	Y	Z
3	When the leftmost element in the whole sequence appears in the window, a left-branching triple U-V-W can be left-bracketed (UV)-W; restart with the three-element window expanded back to the right of the sequence.	(UV)		W	X	Y	...

Table 1: Bracketing algorithm.

User Interaction

When the system cannot find sufficient evidence in favour of right-branching or left-branching, it turns to the user to supply the decision. Such decisions may be difficult and unintuitive. Here are several ways to lessen the burden. By using these techniques, a 'yes' answer to any question will provide confident left-branching; a 'no' answer will mean confident right-branching.

- ask only yes-no questions about right-branching — don't ask the user to supply bracketing information directly.

good: in the context of *tomato soup pot*, does *tomato soup* make sense?

bad: does *tomato soup pot* bracket left or right?

- phrase questions in the context of three individual words by using subphrase reductions.

good: in the context of *steel soup pot*, does *steel soup* make sense?

bad: in the context of *cotton soup pot cover holder*, does *cotton soup pot cover* make sense?

- ask the user only about the acceptability of X-Y; do not ask the user to compare X-Y and X-Z — it is possible for *both* X-Y and X-Z to be unacceptable if X-Y-Z is in the middle of a modifier sequence, which would make the user's decision much more difficult.

good: in the context of *steel tomato soup*, does *steel tomato* make sense?

bad: in the context of *steel tomato soup*, which makes more sense: *steel tomato* or *steel soup*?

An Example

Assume that phrases (12) and (13) have already been bracketed. This section traces through the bracketing of (14).

(12) *(soup bowl)*
(13) *(wooden (pot handle))*
(14) *wooden French onion soup bowl handle*

Start with the rightmost three elements, *soup-bowl-handle*:

wooden	French	onion	**soup**	**bowl**	**handle**

soup-bowl-handle is confidently left-branching, since *(soup bowl)* has occurred and *(soup handle)* has not. Move the window one element to the left and restart the algorithm:

wooden	French	**onion**	**soup**	**bowl**	handle

Neither *(onion soup)* nor *(onion bowl)* have occurred previously and *soup* is not an adjective, so there is no confidence in right-branching or left-branching. Ask the user if *onion soup* makes sense in the context of *onion soup bowl*. The user answers 'yes', providing confidence in left-branching. Move the window one element to the left and restart the algorithm:

wooden	**French**	**onion**	**soup**	bowl	handle

Neither *(French onion)* nor *(French soup)* have occurred previously. Ask the user if *French onion* makes sense in the context of *French onion soup*. The user answers 'no', so the sequence is confidently right-branching. Bracket *(onion soup)* and expand the window one element to the left:

wooden	French	(onion soup)	bowl	handle

French is an adjective, so *wooden-French-(onion soup)* is confidently right-branching. Bracket *(French (onion soup))*. Since there are no more elements to the left of *wooden*, expand the window back to the right:

wooden	(French (onion soup))	bowl	handle

Neither *(wooden bowl)* nor *(wooden soup)*, which is the reduction of *wooden-(French (onion soup))*, have occurred previously. *(French (onion soup))* is obviously not an adjective, so there is no confidence in either right-branching or left-branching. Ask the user if *wooden soup* makes sense in the context of *wooden soup bowl*, which is the reduction of *wooden-(French (onion soup))-bowl*. The user answers 'no', providing confidence in right-branching. Bracket *wooden-(French (onion soup))-bowl* as *wooden-((French (onion soup)) bowl)*. Since there are no more elements to the left of *wooden*, expand the window to the right:

wooden	((French (onion soup)) bowl)	handle

(wooden bowl), which is the reduction of *wooden-((French (onion soup)) bowl)* has not occurred previously; *(wooden handle)* has occurred previously as a reduction of *(wooden (pot handle))*. *wooden-((French (onion soup)) bowl)-handle* is therefore confidently right-branching. Bracket *(((French (onion soup)) bowl) handle)*:

wooden	(((French (onion soup)) bowl) handle)

Since there are only two remaining elements in the sequence, bracketing is trivial:

(wooden (((French (onion soup)) bowl) handle))

Finally, the system stores all of the reduced subbracketings (15) for future processing. It also stores the complete bracketing (16): if the phrase is ever encountered in its entirety, the bracketer can simply look up the complete bracketing instead of going through the steps of the algorithm.

(15) *(onion soup)*
(French soup)
(soup bowl)
(bowl handle)
(wooden handle)

(16) *(wooden (((French (onion soup)) bowl) handle))*

Evaluation

This section gives results of using the bracketer in the context of two experiments. The *sparc* experiment applied the bracketer to the first 500 non-trivial noun phrases in a computer installation guide. In this context a non-trivial noun phrase has at least one premodifier (adjective or noun) or postmodifying prepositional phrase. Bracketing the 500 noun phrases in the test resulted in 645 modifier-head pairs.

The second experiment was in the context of a complete knowledge acquisition experiment (Barker *et al.* 1998) of a book on the mechanics of small engines. Bracketing resulted in 733 modifier-head pairs.

Performance

In both experiments, most of the modifier-head pairs occurred in noun phrases with a single premodifier and head. These simple compounds required no bracketing decisions. In the *sparc* experiment, the 645 pairs required 188 bracketing decisions. The system made 122 (65%) of these decisions correctly, with the rest made by the user. Of the 66 user decisions, 47 (71%) were required during the first half of the experiment with only 19 in the second half. The running totals of user and system bracketing decisions appear in Figure 1.

The *small engines* experiment required 164 bracketing decisions. The system made 101 (62%) decisions correctly, with the rest made by the user. Due to the consistent terminology in the *small engines* text the cumulative number of decisions made automatically by the bracketer was always greater than the number required from the user (after the first decision by the user). The running totals for user and system bracketing decisions appear in Figure 2.

The Effect of the Threshold

The bracketer determines whether a given triple X-Y-Z is confidently right-branching if the subbracketing $(X_h Z_h)$ has previously occurred N times more frequently than

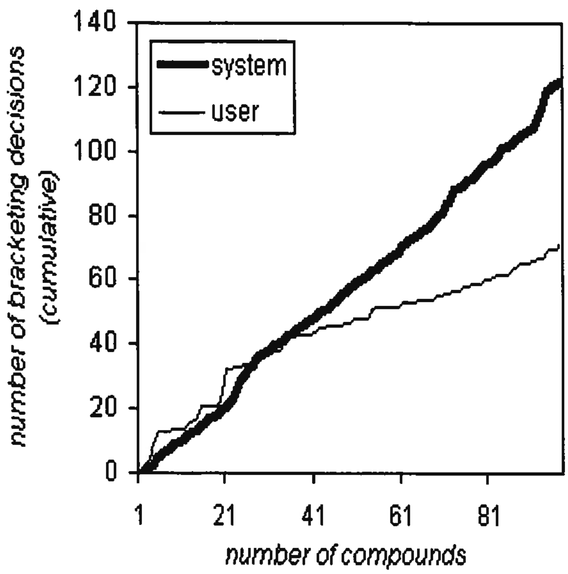

Figure 1. Bracketing decisions in the *sparc* experiment

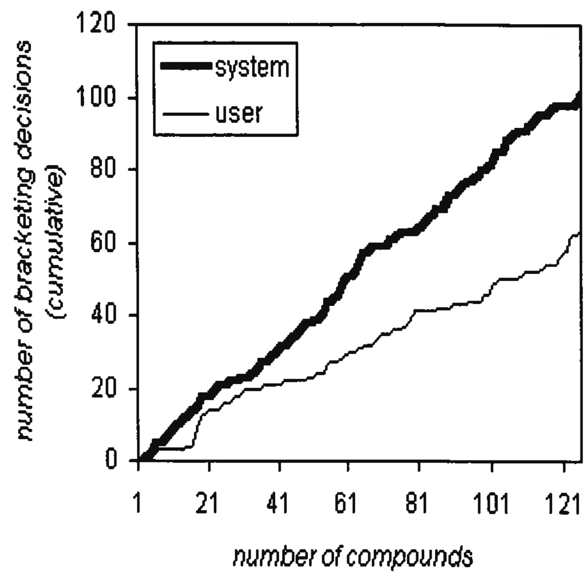

Figure 2. Bracketing decisions in the *small engines* experiment

the subbracketing (X_hY_h), where N is a threshold that can be set by the user. In the absence of sufficient evidence, the system asks the user to supply right-branching information.

If the value of the threshold is set high, the number of previous occurrences of (X_hZ_h) must greatly outweigh the number of occurrences of (X_hY_h) for the system to assume right-branching. High values of the threshold cause the system to be more conservative. Low values of the threshold (close to 1.0), make it more aggressive: the system requires less evidence to commit to a branching decision.

The *sparc* experiment was run twelve times with different threshold values. As expected, the number of system decisions, both correct and incorrect, was highest for low threshold values. For higher threshold values, the number of incorrect system decisions decreased, but so did the number of correct decisions, and the user made more decisions.

For the *small engines* text, changing the threshold had no effect. This result suggests that for any triple X-Y-Z in that text, if (X_hZ_h) appears as a reduced pair, (X_hY_h) does not. In general, however, both may appear in any given text.

Branching Frequencies

Ter Stal (1996) confirms earlier results of Resnik (1993) and Lauer & Dras (1994) that between 60% and 70% of noun-noun-noun compounds in text are left-branching. A bracketer *could* guess left-branching when there is no confidence in right-branching. Results from the *small engines* experiment confirm the bias for left-branching. For the *sparc* experiment, however, the data in Table 2 show that guessing left-branching would have produced poor results. The predominance of left-branching compounds is apparently not universal. If the system were modified to guess left in the absence of other evidence, there are texts (like the *sparc* text) for which the bracketer would perform poorly.

Conclusions

In this article I have presented a semi-automatic system for bracketing noun modifiers. The system can start processing noun phrases with no prior knowledge and, with the help of a user, learn to bracket the majority of noun modifier sequences automatically. Experiments have shown that the system gets better at bracketing as more noun phrases are encountered, that conflicting bracketing evidence is rare within a single, coherent technical text, and that general modifier sequences are not universally predominantly left-branching.

Acknowledgments

This research was supported by the Natural Sciences and Engineering Research Council of Canada. I would like to thank Stan Szpakowicz for valuable input to my work on noun phrases. I am also grateful to Robert Mer-

		left-branching	*right-branching*
small engines	*noun-noun-noun*	47 (96%)	2 (4%)
	adjective-noun-noun	31 (84%)	6 (16%)
	total	78 (91%)	8 (9%)
sparc	*noun-noun-noun*	41 (55%)	33 (45%)
	adjective-noun-noun	11 (26%)	31 (74%)
	total	52 (45%)	64 (55%)

Table 2: Branching frequencies for *small engines* and *sparc.*

cer, Eric Neufeld and the AI'98 program committee for providing fora for the ideas in this article.

References

[1] Barker, Ken (1998a). A Trainable Bracketer for Noun Modifiers. *Proceedings of the Twelfth Canadian Conference on Artificial Intelligence.* Vancouver, 196-210.

[2] Barker, Ken (1998b). Semi-Automatic Recognition of Semantic Relationships in English Technical Texts. Ph.D. dissertation, School of Information Technology and Engineering, University of Ottawa.

[3] Barker, Ken & Stan Szpakowicz (1998). Semi-Automatic Recognition of Noun Modifier Relationships. *Proceedings of COLING-ACL'98*, Montréal, 96-102.

[4] Barker, Ken, Sylvain Delisle & Stan Szpakowicz (1998). Test-Driving TANKA: Evaluating a Semi-Automatic System of Text Analysis for Knowledge Acquisition. *Proceedings of the Twelfth Canadian Conference on Artificial Intelligence.* Vancouver, 60-71.

[5] Jun, Zhao & Huang Changning (1998). A Quasi-Dependency Model for Structural Analysis of Chinese Base NPs. *Proceedings of COLING-ACL'98*, Montréal, 1-7.

[6] Lauer, Mark (1995). Corpus Statistics Meet the Noun Compound: Some Empirical Results. *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics.* Cambridge. 47-54.

[7] Lauer, Mark & Mark Dras (1994). A Probabilistic Model of Compound Nouns. *Proceedings of the 7th Australian Joint Conference on Artificial Intelligence.* Armidale. 474-481.

[8] Liberman, Mark & Richard Sproat (1992). Stress and Structure of Modified Noun Phrases. *Lexical Matters* (*CSLI Lecture Notes*, 24). Stanford: Center for the Study of Language and Information.

[9] Pustejovsky, James, S. Bergler & P. Anick (1993). Lexical Semantic Techniques for Corpus Analysis. *Computational Linguistics* 19(2). 331-358.

[10] Resnik, Philip Stuart (1993). Selection and Information: A Class-Based Approach to Lexical Relationships. Ph.D. thesis, IRCS Report 93-42, University of Pennsylvania.

[11] ter Stal, Wilco (1996). Automated Interpretation of Nominal Compounds in a Technical Domain. Ph.D. thesis, University of Twente, The Netherlands.

About the Author

Ken Barker obtained his Ph.D. in Computer Science from the University of Ottawa in June 1998 and was hired by the School of Information Technology and Engineering in July. His research focuses on knowledge-scant tools for natural language engineering. His projects include systems for semi-automatic extraction of knowledge from technical texts, fully automatic text summarization and word sense disambiguation. Dr. Barker's recent awards include the Canadian Advanced Technology Association prize for graduate research in Computer Science and Software Engineering and the Best Paper Award at the Twelfth Canadian Conference on Artificial Intelligence.

Ken Barker can be contacted at
School of Information Technology and Engineering
University of Ottawa
161 Louis Pasteur, P.O. Box 450 Stn A,
Ottawa Ontario Canada K1N 6N5
Tel: (613) 562-5800 x6190; fax: (613) 562-5175
Email: kbarker@site.uottawa.ca
WWW: http://www.site.uottawa.ca/~kbarker

Learning Visual Landmarks for Pose Estimation

Robert Sim and Gregory Dudek

Résumé

Traduction non disponible.

Abstract

We present an approach to vision-based mobile robot localization which learns a set of visual features we refer to as landmarks. The landmark learning mechanism is designed to be applicable to a wide range of environments. Each landmark is detected as a local extremum of a measure of uniqueness and is represented by an appearance-based encoding. Our formulation of uniqueness is inspired by models of biological visual attention. Localization is performed using a method that matches observed landmarks to learned prototypes and generates independent position estimates for each match. The independent estimates are then combined to obtain a final position estimate, with an associated uncertainty. Quantitative experimental evidence is presented that demonstrates that accurate pose estimates can be obtained.

Introduction

This paper addresses the question of position estimation for a robot located in a previously explored region of the environment. The robot is equipped with a single camera, and does not require an *a priori* pose estimate. An accurate position estimate is desired without any motion on the part of the robot. One might imagine that the robot must consistently re-localize itself after periodic shutdowns for maintenance. We build on previous work by Sim and Dudek which demonstrated that position estimation could be accurately performed in a more constrained environment using a similar technique (Sim & Dudek, 1998; Sim, 1998). In contrast to alternative localization methods, our approach provides accurate pose estimates at low computational cost without making domain-dependent assumptions (Sim, 1998).

Our approach to the problem employs image-domain features we refer to as *landmarks* to perform position estimation. We extract these landmarks from a preliminary traversal of the environment. Candidate landmark selection is based on a local distinctiveness criterion-that is, local maxima of edge density, a measure which is motivated by supporting evidence in biological models of visual attention. Candidate selection is later validated by verifying the appearance of the candidate landmarks against a set of landmark prototypes. The method consists of an off-line "mapping" phase and on-line "localization" phase. The off-line phase is performed once, upon initial exploration of the environment, and consists of learning a set of tracked landmarks considered useful for position estimation The on-line phase is performed as often as a position estimate is required, and consists of matching candidate landmarks in the input image to the learned tracked landmarks, followed by position estimation using an appearance-based linear combination of views.

Method

Off-line "Map" construction: (Figure 1)

1. Training images are collected by sampling a range of poses in the environment (represented as the set of nine rectangles or "images" in Fig 1).

2. Candidate Landmarks (represented as the smaller squares in each "image" in Fig 1) are extracted from each image as local maxima of edge density.

3. Tracked Landmarks, each of which is represented by a characteristic prototype, are extracted by tracking sets of candidate landmarks over the configuration space of the robot. For each image, a local search is performed in the neighbourhood of the candidate landmarks in the image in order to locate optimal matches to the prototypes.

4. A measure of confidence is computed for each tracked landmark, based on the quality of the pose estimates obtained from the tracked landmark as it is used to estimate the (known) position from which each of its member candidate landmarks was observed.

5. The set of tracked landmarks and corresponding measures of confidence is stored for future retrieval.

On-line localization: (Figure 2)

1. When a position estimate is required, a single image is acquired from the camera.

2. Candidate landmarks are extracted from the input image as local maxima of edge density.

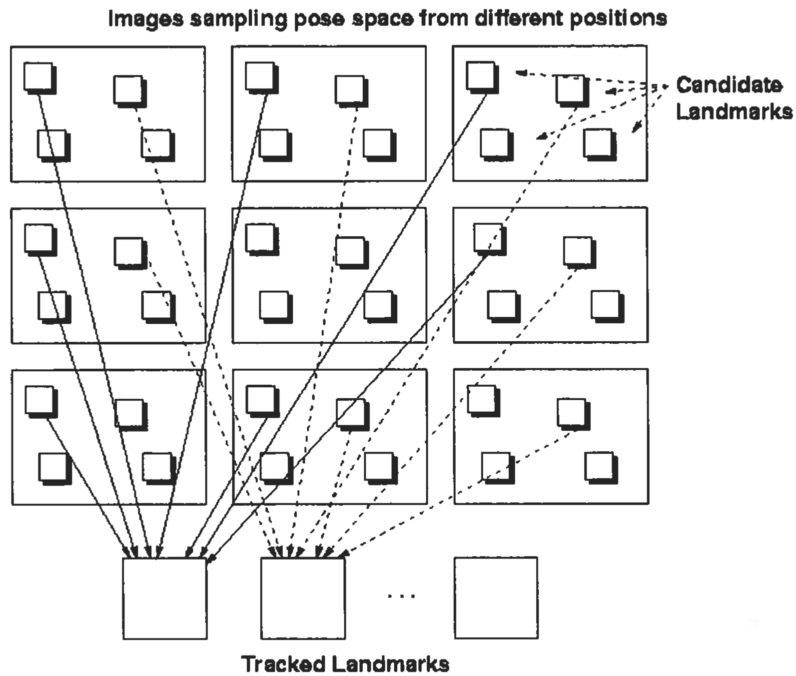

Figure 1: The offline (training) phase.

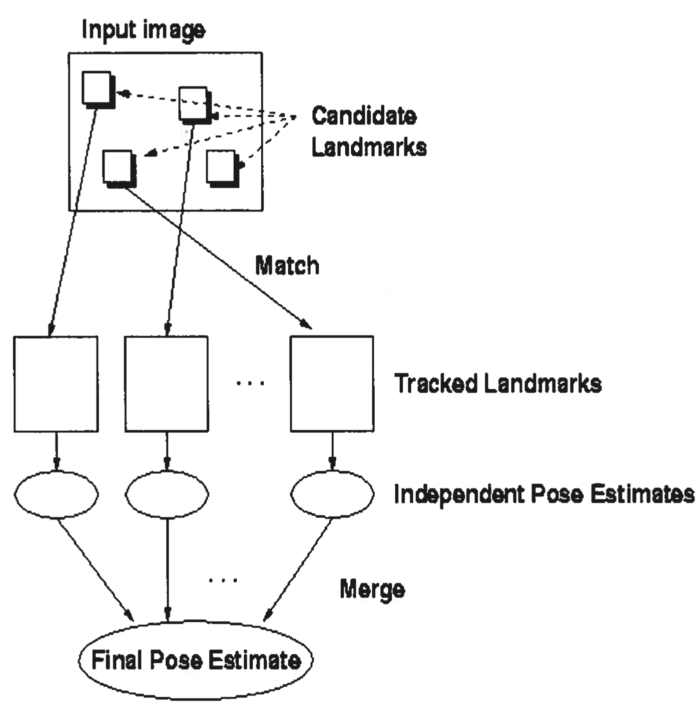

Figure 2: The online (pose-estimation) phase.

3. The candidate landmarks are matched to tracked landmarks by matching a principal components description of the candidate to the learned prototypes,

4. A position estimate is obtained for each matched candidate landmark. This is achieved by computing a reconstruction of the candidate based on a principal components decomposition of the intensity distributions and image positions of the (previously learned) candidates and their known poses in the tracked landmark. The result is a position estimate obtained as a linear combination of the views of the candidates in the tracked landmarks. Our confidence in the pose estimate obtained for a particular candidate landmark is assigned as the precomputed measure of confidence in the tracked landmark to which it matched.

5. A final position estimate is computed as the robust average of the individual estimates of the observed candidates, taking into account the measure of confidence assigned to each estimate.

Experimental Results

An indoor scene is depicted in Figure 3. In this scene, a camera was mounted on an RWI B-12 mobile robot. Training images were taken at 20.0cm intervals over a 2.0m by 2.0m grid. In this experiment the orientation is fixed to face straight ahead. Ground truth pose measurements were obtained by hand, accurate to approximately 0.5cm, and 0.5° of rotation about a vertical axis. Despite reasonably good dead reckoning, the unevenness of the floor led to some variation in image alignment (i.e. rotation about an arbitrary horizontal axis). Once training images were collected, a series of 30 test images were taken from random positions in order to test the method. Figure 4 presents the set of estimates obtained using the method, plotted against their ground-truth. Each 'x' represents the pose estimate obtained for the image taken at the corresponding true position 'o'. Grid crossings denote the locations from which training images were obtained. The mean error in position is 6.3cm or 31% of the spacing between training samples.

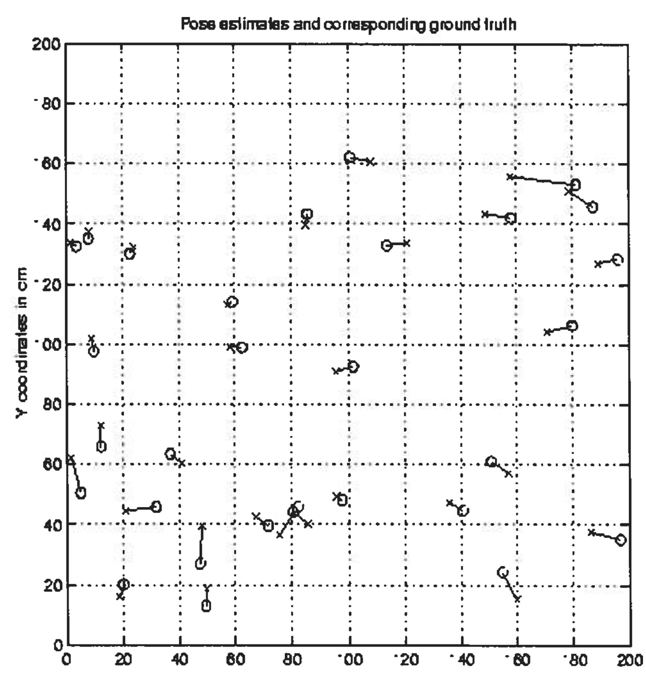

Figure 3: The explored environment.

Conclusions

We have presented a method for *learning* a set of image-domain landmarks which are suitable for robotic pose estimation. The method consists of an offline "training" phase, performed once, and an online "estimation" phase, performed when a pose estimate is required. In the offline phase, a model of visual attention is employed to select suitable candidates, which are then tracked over the configuration space. The online method exploits the interpolating properties of principal components analysis methods in order to obtain pose estimates. When the orientation of the robot is constrained, a single image is sufficient to produce a pose estimate accurate to a fraction of the space between training samples. In other work, we have shown that incorrect orientation of the camera can be detected and corrected, and that in sufficiently constrained environments, highly accurate pose estimates can be obtained.

References

[1] Sim, R and Dudek, G., Mobile Robot Localization from Learned Landmarks, *Proceedings of the IEEE/RSJ Conference on Intelligent Robots and Systems (IROS)*, Victoria, BC, 1998.

[2] Sim, R., Mobile Robot Localization from Learned Landmarks, Master's Thesis, McGill University, Montreal, QC, 1998 (also available as a Centre for Intelligent Machines technical report CIM-98-03).

Figure 4: The set of pose estimates obtained using the method. Each 'x' represents the estimated position for the corresponding 'true' position, marked by an 'o'. The mean error in estimation is 6.3cm.

About the Authors

Robert Sim (simra@cim.mcgill.ca) obtained a B.Eng (1996) in Computer Engineering and M.Sc. (1998) in Computer Science at McGill University. He is currently a Ph.D. candidate in Computer Science at McGill's Centre for Intelligent Machines, where he is considering the problems of robot exploration and pose estimation.

Gregory Dudek (dudek@cim.mcgill.ca) is an Assistant Professor with the School of Computer Science and a member of the McGill Research Centre for Intelligent Machines (CIM) and an Associate member of the Dept. of Electrical Engineering at McGill University. He directs the McGill Mobile Robotics Laboratory. He has published over 100 research papers on subjects including visual object description and recognition, robotic navigation and map construction, distributed system design and biological perception. His research interests include perception for mobile robotics, navigation and position estimation, environment and shape modelling and computational vision.

Editor's Message (contd from p. 2)

The second of these articles, "Twenty Five Years of CSCSI/SCEIO and CAI/IAC," will appear in the next issue of CAI/IAC. This article is written for the same two audiences.

Also appearing in the current issue are papers by Ken Barker and by Robert Sim and Gregory Dudek. These papers are based on work that was selected as having particular merit at Canadian conferences last year. Ken's paper "A Trainable Bracketer for Noun Modifiers" was selected as best paper at last year's AI '98 conference. Robert was one of three graduate students to win best poster prizes at IRIS-PRECARN '98. The other prize winners were Richard Dearden and Terence Gilhuly. A paper by Richard will appear in the next issue of *CAI/IAC*. (Terence's poster seems to be more robotics than AI, so it does not seem appropriate to publish his work in *CAI/IAC*.)

While I have ideas for content for future issues of *CAI/IAC*, ideas and especially contributions would be very welcome. Feel free to send in Letters to the Editor. I would like to continue the News / Nouvelles section on page 6 and would welcome tips, leads, tidbits, news items, and the like for the next issue. Please contact me at fass@cs.sfu.ca

Advances in Board and Card Game Software

Scrabble, Backgammon Software Rival Human Best

At a computer games exhibition at AAAI-98, Scrabble and backgammon software each played three-day matches against the human world champions in those games (see http://www.aaai.org/Conferences/National/1998/aaai98-exhibition.html#champions).

PC-Based Chess Program Beats World #2

The Dutch-made computer chess program Rebel played eight games against the world number-two chess player in the world, Vishy Anand, during 21-23 July 1998. Rebel won 5-3. Rebel ran on a boosted 450-MHz AMD K6-2 chip cooled with a KryoTech Cooler (see http://www.rebel.nl/anand.htm).

GIB Bridge Program Increasingly Competitive

GIB (Goren in a Box), designed by US AI Professor Matthew Ginsburg, finished 12th in the 1998 world championships. Of the approximately 13,000 points lost by GIB, 1000 were in time penalties, 6000 were because GIB failed to understand the bidding (see http://www.cirl.uoregon.edu/ginsberg/gibnews.html).

The Infrastructure of Artificial Intelligence R&D in Canada, Part 1

Dan Fass

Résumé

Cet article présente une brève histoire du domain de l'intelligence artificielle, se concentrant particulièrement sur la situation au Canada. Il décrit ensuite les sociétés et organismes canadiens qui s'occuppent de promouvoir l'IA (comme par exemple la SCEIO/CSCSI), le soutient pour l'IA au Canada venant du secteur privé et du gouvernement fédéral (par exemple, CIAR [Institut canadien pour la recherche avancée], PRECARN [Réseau de recherche appliqué préconcurrentielle] et IRIS [l'Institut de robotique et d'intelligence des systèmes]), diverses politiques initiées durant les années 80 et le tout début des années 90 (par exemple, l'atelier sur l'IA organisé par le Conseil de science en 1983), les institutions fédérales s'occuppant de la recherche en IA (comme Industrie Canada) et certaines initiatives provinciales (par exemple, le Conseil de recherches de l'Alberta).

Abstract

This article presents a brief history of AI with special reference to Canada. It then describes Canadian societies and organizations involved in promoting AI (such as CSCSI/SCEIO), private-sector and federal-level public support for Canadian AI (e.g., CIAR, PRECARN, and IRIS), various policy initiatives in Canadian AI during the 1980s and very early 1990s (for example, the 1983 Science Council workshop on AI), federal institutions that are involved in AI research in some way (such as Industry Canada), and provincial initiatives in AI (e.g., the Alberta Research Council).

1. Introduction

This article looks at the "infrastructure" of Canadian R&D (research and development) in AI (artificial intelligence). By "infrastructure" I mean institutions that support or have supported AI in Canada by either funding, promoting, or facilitating it.

Section 2 provides some context to the description of infrastructure by outlining the establishment and gradual consolidation of AI through the late 1950s, 1960s, and 1970s. A look is then taken at the hyping of AI in the 1970s and 1980s, and the industrialization of AI in the 1990s.

In Section 3, Canadian societies and organizations are listed that are involved in some degree in the promotion of Canadian AI. These include CIPS, CSCSI/SCEIO, CHCCS, CIPPRS, ITAC, and CATA.

Private-sector and federal-level public support for Canadian AI is reviewed in Section 4. Descriptions are given of NSERC, the formative role of CIAR in the early-mid 1980s, PRECARN, IRIS, STEAR and CSA, the 1990 Artificial Intelligence Research and Development Fund, and a number of other sources.

Sections 5 and 6 describe various policy initiatives in Canadian AI during the 1980s and very early 1990s. These initiatives varied greatly in their format and source. Formats included committees, workshops, and surveys. Sources included federal government, provincial government, corporate leadership, and the AI R&D community. These initiatives can be divided into two main groups. Section 5 describes the first group of initiatives, which took place during 1983-1991 and consisted largely of AI policy ideas produced by the Canadian AI R&D community. Several R&D community initiatives stemmed from a January 1983 workshop organized by the Science Council: a policy discussion paper from the R&D community was circulated at the workshop, the workshop proceedings contained some comments about AI R&D policy that were controversial and seemingly influential, and a survey of AI R&D in Canada (conducted in 1983-1984) was initiated at the same workshop. There was a second policy discussion paper in 1985, produced by the Canadian Society for Fifth Generation Research. Finally, there was a 1986-1991 Associate Committee on AI (ACAI). This committee had elements of involvement from the AI R&D community and federal government. It was initiated by people from the former but was sponsored by the NRC, which is a branch of federal government.

(Some criticisms of NSERC funding of AI were made in the proceedings of the 1983 and elsewhere. The criticisms led to some debate, which was covered in Canadian Artificial Intelligence/Intelligence Artificielle au Canada (often referred to hereafter as Canadian AI/IA au Canada or CAI/IAC). This debate is reviewed in the description of NSERC in Section 4.)

Section 6 describes the second group of initiatives, which occurred during 1983-1991 and consisted of studies of AI initiated by federal or provincial governments. The federal government commissioned from private consulting firms a 1983-1984 study of AI in Canada (the five-volume Cognos Report), a 1985 study of natural language processing in Canada, a 1986 study on the application of expert systems to the Canadian transportation sector, a 1986 study of potential AI applications in federal government, and a large 1989 study of AI R&D in Canada, overseen and published by ACAI in 1991 (ACAI, 1991). The Québec government produced a 1984 paper on its view of AI and a 1986 AI bibliography of useful references.

In Section 7, descriptions are given of federal institutions that are involved in AI research in some way. These include Industry Canada, its Council of Science and Technology Advisors, its Strategis online presence, its Technology Partnerships program, and its two research centres concerned with AI: the National Research Council (NRC) and Communications Research Centre. Particular attention is paid to the NRC, its Institute for Information Technology which does AI research, and its Canadian Technology Network (CTN) and Industrial Research Assistance Program (IRAP), both of which provide technology assistance to small and medium-sized Canadian businesses.

Many provinces have begun their own high-technology initiatives. Some of these initiatives relevant to AI R&D are reviewed in Section 8.

In Section 9, some analysis and a few summary observations are made. First, comparisons are made between this article and the three main sources on which it draws: Canadian AI/IA au Canada, and the 1983-1984 and 1989 surveys of AI R&D in Canada. Second, certain 1980s documents are identified that seem to have had a notable influence on subsequent AI policy of the federal government. Third, an attempt is made to highlight the main elements of the infrastructure that's been described in the article.

Apologies are made in advance for the heavy use of acronyms in this article. Their use is made all the more confusing because a high percentage contain the letter 'C' and a considerable number also contain an 'A', 'T' or 'I', or some combination of them (e.g., CAI, CATA, ITAC, ITSC, CSA). I would defend the heavy use of acronyms because many of them were and are favoured over their longer, proper names. For example, people generally speak of PRECARN and not the Pre-Competitive Applied Research Network.

The world-wide web was used heavily in researching this article. Many URLs appear, as do quotes from web pages, and/or the results of searches with search engines. Please be aware that all such information derived from the web may be out-of-date or otherwise incorrect.

2. The Ups and Downs of AI

The term "artificial intelligence" was coined by John McCarthy for a proposal to obtain funding for the now-famous, seminal AI conference at Dartmouth College in 1956. In the years following the conference, AI research centres were established at the Massachusetts Institute of Technology; Carnegie Mellon University; Stanford University; Bolt, Beranek & Newman (BBN); Stanford Research Institute (now SRI International); and Edinburgh University in Scotland.

The computational foundations of AI were established by, among other things, the development of LISP in the late 1950s and early 1960s. AI soon divided into subfields such as vision, problem solving, and natural language processing (NLP). Techniques were developed in the areas of representation (e.g., logic, semantic nets, frames), inference (e.g., theorem-proving, constraint propagation), and control (e.g., demons), and problem-solving architectures (e.g., rule-based, constraint-based). The first AI compendium, ``Computers and Thought," appeared in 1963 (Feigenbaum & Feldman, 1963).

The first faculty with AI backgrounds appeared in Canadian universities in the early 60s and faculty began to be appointed in AI (as a distinct subject) in the mid-late 60s. Also in the late 60s, forums for AI publications began to appear. For example, the first IJCAI (International Joint Conferences on Artificial Intelligence) was held in 1969 and Artificial Intelligence journal began in 1970.

In the 1970s, logic programming and Prolog were created. New techniques were developed in representation (e.g., production rules, semantic networks, scripts, frames, knowledge representation languages), inference (e.g., unification), and problem-solving architectures (e.g., blackboard systems). A major application was expert systems. In the US, DARPA (Defense Advanced Research Projects Agency) provided considerable funding to AI research, notably in speech recognition.

In the 1980s, dedicated LISP machines started to appear. LISP Machines Incorporated (LMI) introduced the first commercial LISP machine in 1981. Symbolics produced their own LISP machine the same year. Texas Instruments produced their own machine, the TI Explorer, in 1985.

The first Canadian AI companies started up in the early-mid 1980s. Three that continue to this day are Applied AI Systems (AAI) (http://www.aai.ca), Acquired Intelligence (http://vvv.com/ai), and Comdale Technologies (http://www.comdale.com).

Applied AI Systems of Kanata, ON, was set up in January 1983 and is the oldest AI company in Canada (http://www.aai.ca/who.htm). Applied AI Systems has worked in AI, robotics, artificial life, speech recognition, neural networks, and real-time knowledge-based systems.

Acquired Intelligence was established in Victoria, BC, in 1987. (These days, Acquired Intelligence is a specialist in knowledge acquisition and knowledge-based development. It markets an expert systems shell called Acquire (http://vvv.com/ai/acquire/pcm.html).)

Comdale Technologies of Toronto, ON, was incorporated in 1987. In its early days, Comdale developed AI software for industrial process control and marketed a rule-based expert system development tool. (These days, Comdale's products incorporate techniques from fuzzy logic, expert systems, neural networks and genetic algorithms (cf. http://www.comdale.com/general/about.htm).)

Other companies starting AI R&D in the mid-1980s included CompEngServ which was using AI techniques from 1983; Bell-Northern Research, Ottawa, ON, which formed a core AI group in 1984; Canadian Pacific, which started using AI technology in 1984; CAE Electronics Ltd., which began its activities in AI in 1985; and Shell Canada, which began doing AI work in 1986.

The 80s also saw the establishment of AI consulting companies. One of the first such companies was the Canadian Cognicom Inc. which comprised seven Canadian academic researchers: Richard Kittredge, John Mylopoulos, Zenon Pylyshyn, Ray Reiter, John Tsotsos, Robert Woodham, and Steve Zucker. The company was set up in 1983, initially in response to a request for a report on AI in Canada (later to become the Cognos study — see Section 6). The company provided advice to industry and governments throughout the 80s and early 90s before being dissolved in 1997.

Numerous workshops, conferences and journals appeared in AI's subfields. More techniques were developed in inference (e.g., inheritance, case-based reasoning), control (e.g., metaplans), and problem-solving architectures (e.g., object-oriented, neural network).

In the mid-late 1970s and 1980s, in particular, AI was heavily hyped. This was due in part to the announcement of the Japanese Fifth Generation project in October 1981, which caused something of a panic in North America and Europe. As a response, the UK began the Alvey Programme (later renewed as Alvey-2), West Germany started a program, the EEC began its ESPRIT program (renewed as ESPRIT-2), and the US initiated the Microelectronics and Computer Technology Corporation (MCC) and its Strategic Computing Initiative. Even Finland began its FINPRIT (Finnish Program for R&D in Information Technology). The media were accused of playing a role in the hype, as were some AI practitioners (cf. McDermott et al., 1985).

The hype led to problems. According to Moravec (1994, p. 86), "[t]he first AI companies rushed to market academically interesting but underdeveloped techniques with few applications or customers." Symbolics was a major player. Its and LMI's LISP machines were sold to secondary AI companies who "sold application-oriented or generic expert systems at exorbitant prices" (Ibid.).

Eventually, there was a downturn, referred to by some as "AI winter" (see McDermott et al., 1985). The bubble began to burst when cheap Unix workstations (produced by, e.g., Sun Microsystems and Hewlett-Packard) and then PCs appeared that could run LISP code as well as LISP machines could. Also, expert systems and other AI software began to be written in C and other non-AI languages, with the result that companies found they could develop their own applications in-house at a lower price.

Companies like Symbolics, Palladin, Intellicorp, Teknowledge, Gold Hill and others either went under or survived in much reduced form. Symbolics became Macsyma, Inc.; Teknowledge became a small division (Moravec, 1994, p. 86). In Canada, Nexa Corporation, formed in 1982, went under in 1988. Its subsidiaries included Canadian Artificial Intelligence Products Corporation (CAIP, which had begun in 1984), Symbolics

Canada, International Artificial Intelligence Inc. (IAI), Inference Canada (makers of ART (Automated Reasoning Tool)), and Logicware Inc. (who marketed MPROLOG and the TWAICE expert-system builder).

Adding to the downturn was the termination of the Japanese Fifth Generation project in 1992, which caused major government initiatives in other countries to also end (e.g., the Alvey-2 program in the UK).

In the 1990s, considerable attention has been paid to large-scale, engineering-type statistical approaches. There are numerous applications of AI techniques such as fuzzy logic (Munakata & Jani, 1994), neural networks (Widrow et al., 1994), and distributed AI (Chaib-draa, 1995). Machine learning (ML) techniques have been very influential. Data mining, which uses ML techniques, has become a major application.

In 1993, the global market for AI systems was estimated to be worth about US$900 million (Charles, 1995, p. 70), with 70-80% of Fortune 500 companies using AI, and the US Department of Defense being "by far the largest single user of AI in the world" (Ibid., p. 71). There are now thousands of domain-specific expert systems (Durkin, 1996) and case-based assistants in use worldwide; speech recognition and generation applications are common (Hayes-Roth, 1997, p. 105). Some game programs operate at the level of highly skilled human players, e.g., in chess, checkers, and bridge (Hedberg, 1997). An example is the Chinook drafts program, whose principal developer was Jonathan Schaeffer of the University of Alberta, was the first to play and win a World Championship (Schaeffer, 1997; Schaeffer et al., 1996; http://www.cs.ualberta.ca/~chinook). Another example is IBM's Deep Blue chess program which has beaten world chess champion Garry Kasparov.

3. Canadian Societies and Organizations Involved in Promoting AI

In this section, CIPS, CSCSI/SCEIO, CHCCS, CIPPRS, ITAC, CATA, and the Canadian Society for Fifth Generation Research are described. All are or have been involved to some degree in the promotion of AI in Canada, though the involvement of ITAC and CATA has been less than the others.

The Canadian Information Processing Society (CIPS) (http://www.cips.ca) was founded in 1958 with the mission of defining and fostering the IT profession, encouraging and supporting IT practitioners, and advancing the theory and practice of IT, while safeguarding the public interest (http://www.cips.ca/organization.htm). CIPS has more than 6,000 Canadian members (http://www.cips.ca/toronto).

CIPS has four national-level SIGs (special-interest groups). These are the Canadian Society for Computational Study of Intelligence / Societé Canadienne des Études d'Intelligence par Ordinateur (CSCSI/SCEIO) (http://www.cscsi.sfu.ca), Canadian Human-Computer Communications Society (CHCCS) (http://zeppo.cs.ubc.ca:2001/home.html), Canadian Image Processing and Pattern Recognition Society (CIPPRS) http://www.gel.ulaval.ca/~cipprs/), and Canadian Information Processing Security SIG (SEC) (http://cipsnet2.cips.ca/sigs/security/default.htm). For more on CSCSI/SCEIO, see the forthcoming article "Twenty Five Years of CSCSI/SCEIO and CAI/IAC."

The Toronto Chapter of CIPS formed a SIGKBS (Special Interest Group on Knowledge-Based Systems), but that no longer seems to exist.

The Information Technology Association of Canada (ITAC) (http://www.itac.ca), together with various partner organizations, represents over 1200 IT companies in the computing and telecommunications fields (http://www.itac.ca/frame.htm). (PRECARN is a member of ITAC.) ITAC conducts research and has produced publications on, among other subjects, "A Knowledge-Based Canada: The New National Dream" (January 1993).

There was an article by Roy Woodbridge of the Canadian Advanced Technology Association (CATA). CATA's views on Canadian high tech R&D were often reported in CAI/IAC. CATA presently has 450 corporate members. Its mandate is:

> to ensure that Canada continues to be a good place from which advanced technology companies can do business. To do this, CATA encourages a public policy environment favourable to investment in scientific endeavour, the conduct of research and development, and the production of strategic and emerging technologies, products, and services in Canada (http://vvv.com/VIATeC/canadtec.htm).

There is also a CATA Alliance (http://www.cata.ca/cata/catainfo/index.html), a national trade association of more than 1000 "new economy" enterprises.

The Fifth Generation Society, also known as the Canadian Society for Fifth Generation Research (CSFGR/ SCRSCG) was formed in March 1984 and was very active for several years. Its members overlapped considerably with those of CSCSI/SCEIO, and its activities, e.g., its annual meetings, were reported in Canadian AI/ IA au Canada. The Fifth Generation Society was very active for many years, producing a draft Canadian Fifth Generation Plan in 1985 (see March 1985 supplement of CAI/IAC). It is now apparently defunct.

4. Private-Sector and Federal-Level Public Funding Support for Canadian AI

This section describes the main sources of funding for Canadian AI R&D in the last 25 years. These are NSERC, the formative role of CIAR in the early-mid 1980s, PRECARN, IRIS, STEAR and CSA, and the 1990 Artificial Intelligence Research and Development Fund. Several of these sources have involved substantial private sector input and funding, notably CIAR and PRECARN. Some other, smaller funding sources are also listed.

4.1. Natural Sciences and Engineering Research Council (NSERC)

Most Canadian AI funding comes from NSERC (Natural Sciences and Engineering Research Council) (cf. October 1989, p. 9). NSERC (http:/www.nserc.ca), which began its work in May 1978, is one of Canada's three federally-funded granting councils. The other two are the MRC (Medical Research Council) and SSHRC (Social Sciences and Humanities Research Council).

NSERC has three goals (http:/www.nserc.ca/fact.htm): (1) to support high-quality research by means of research grants, (2) to promote industrial R&D through its Research Partnerships Program, and (3) to train the next generation of researchers through its scholarships and fellowship programs. NSERC manages jointly with other organizations various Networks of Centres of Excellence (NCEs) (http://www.nserc.ca/programs/resguide/nce.htm), including IRIS, of which more later.

"Support is provided for a wide variety of disciplines, including artificial intelligence, through a number of NSERC programs including the Operating Grants Program, the Strategic Grants Program, Research Partnerships, Scholarships, and International Programs. NSERC invested in excess of $3.6 million for artificial intelligence research in 1989-90" (ACAI, 1991, p. 5). There was a controversy about NSERC's funding policy in 1994 (CAI/IAC, autumn 1994, pp. 4-9). It was alleged in Canadian AI/IA au Canada that NSERC's Research Grants Program was being phased out, but the program still exists in 1998. A listing of 1998 awards in computing and information science (including AI) can be found at http:/www.nserc.ca/programs/result/1998/rg/07.htm

After the publication of the proceedings of the AI workshop organized by the Science Council of Canada (see Section 5), there was some disagreement between Nick Cercone, then president of CSCSI/SCEIO, and Gordon MacNabb, president of NSERC, about some things said about NSERC in the proceedings. Some of their correspondence on this was published in the March 1984 issue of the Newsletter of CSCSI/SCEIO, CMCCS/AACHO, and CIPPRS (Ibid., pp. 107-118). In a letter to Stuart Smith, President of the Science Council of Canada, Gordon MacNabb expressed concern that on page 60 of the Science Council workshop report (1983), university researchers declared that they were "poorly served by NSERC" and argued that "improvement in funding is needed rapidly" (MacNabb, 1984, p. 110). MacNabb attached an annex to his letter showing that total funds granted to computing and information science had increased 243% over five years.

Nick Cercone took issue with some of the contents of Gordon MacNabb's letter to Stuart Smith. Cercone said that he did not believe there was an expert on artificial intelligence at the time on the NSERC Grant Selection Committee for Computer and Information Science (Cercone, 1984, p. 116). He agreed that computing science support from NSERC had increased markedly in recent years, but said that it still lagged significantly behind the level of support given to other disciplines (Ibid.).

4.2. Networks of Centres of Excellence (NCE)

The federally-funded Networks of Centres of Excellence (NCE) program was launched in 1989. The NCE Steering Committee includes the presidents of NSERC, MRC and SSHRC, and the Deputy Minister of Industry Canada.

There are eleven currently funded NCEs (http://www.nce.gc.ca). These are grouped into five areas. The three areas most relevant to AI are Information Technology, Infrastructure, and Human Resources. The Information Technology grouping includes the Institute for Robotics and Intelligent Systems (IRIS), which is funded for 1989-2005 (subject to a positive mid-term progress review) and Micronet — Microelectronic Devices, Circuits and Systems, funded from 1989-2005 (subject to a mid-term positive review). Infrastructure includes Intelligent Sensing for Innovative Structures (ISIS), funded for 1995-2002 (subject to a positive mid-term review). Human Resources includes the Tele-Learning-NCE.

IRIS is administered by PRECARN Associates Inc. (PRECARN is short for the PRE-Competitive Applied Research Network). PRECARN and IRIS are described at some length in Sections 4.4. and 4.5. The TeleLearning-NCE is described briefly in Section 4.6.

4.3. Canadian Institute for Advanced Research (CIAR)

To understand the origins of PRECARN and IRIS, it is helpful to know about the Canadian Institute for Advanced Research (CIAR), which was founded in 1982 by Fraser Mustard and a group of Canadian academics and business people as a private sector initiative, with support from individual corporations, foundations, and the Ontario government. "The federal government recently agreed to match private sector support for CIAR up to $7 million over the next four years" (CAI/IAC, January 1989, p. 11).

Fraser Mustard was CIAR's first president from 1982-1996. The Artificial Intelligence and Robotics (AIRS) program, approved in 1983 and founded 1 July 1984, was CIAR's first program and its largest in 1991, eventually supporting 14 fellows and 11 associates. In its five-year review, CIAR fellowships were viewed as key to recruiting to Canada talented people in AI. The fellowships also helped keep in Canada a significant pool of talented AI people at a time when such people were being aggressively recruited around the world.

Presently, the CIAR is based in Toronto. Although its AI program has been terminated, most of CIAR's members continue as participants of IRIS (see Section 4.5).

4.4. Pre-Competitive Applied Research Network (PRECARN)

PRECARN Associates Inc. (PRE-Competitive Applied Research Network) was set up by Gordon MacNabb (the first president of NSERC) and Fraser Mustard (first president of CIAR) and was incorporated as a non-profit organization in May 1987 (PRECARN, 1997, p. 34). The mandate of PRECARN is "To improve the capacity of Canadian industry to understand, receive and employ advances in intelligent systems technologies" (PRECARN, 1997, p. 3). The first president and CEO was Gordon MacNabb, who retired in the fall of 1993. The second president was Mac Evans, who served from fall 1993 to June 1995, and left to become president of the Canadian Space Agency. The third since 1 July 1995 has been Harry Rogers.

PRECARN is an industry-led consortium and is based in Nepean, ON. Its members include software companies such as Corel Corporation; defence contractors such as Spar Aerospace and MacDonald Dettwiler; and resource-based industry and energy sector companies such as INCO, BC Hydro, Hydro Québec, and Atomic Energy of Canada Ltd. PRECARN's members also include federally-funded research institutions involved in AI R&D such as the National Research Council and Department of National Defence (see Section 7). Also members are some provincially-funded research institutions concerned with AI R&D, for example, the Alberta Research Council, BC's Advanced Systems Institute, and Québec's centre de recherche informatique de Montréal (CRIM).

In all, the first phase of PRECARN funded about eight projects varying in length from two to five years. The first project to get under way was APACS (Advanced Process Analysis and Control System), led by Ontario Hydro. The second PRECARN project was the 5-year IGI/Visiwall (Intelligent Graphics Interface) project, begun in 1991 and led by MPR Teltech. Other projects were led by Spar Aerospace (a project known by the acronym KAD), Defence Research Establishment Valcartier (PASSPORT), and Atomic Energy of Canada Ltd. (TCHE). A listing of all past and present PRECARN projects can be found at http://www.precarn.ca/precproj.htm

PRECARN's funding comes from industry and government. On average, every dollar invested by industry in PRECARN research programs is combined with $5 from other industry partners, and $11 from government (http://www.precarn.ca/moreprcn.htm). As far as government support is concerned, PRECARN was given $10 million by Industry Canada in June 1989, and a further $6 million in December 1992 (PRECARN, 1997, p. 34).

PRECARN-II was announced 6 January 1995. It is backed by a contribution from Industry Canada of $19.4 million over five years through to 31 March 2000 (PRECARN, 1997, p. 34). On PRECARN's web pages, it is noted that through leveraging, "funding for IRIS projects will enable investments in excess of $65 million between 1996 and 2000" (http://www.precarn.ca/moreprcn.htm). Phase 2 projects were "shorter and more commercially driven than those conducted during Phase 1" (PRECARN, 1997, p. 12). The Phase 2 projects were in six sectors: manufacturing and engineering, mining, forestry, transportation and infrastructure (e.g., vehicle maintenance and marine navigation), health care, and process industries (power distribution, power generation, and food inspection) (PRECARN, 1997, p. 12).

A business plan for PRECARN-III was submitted to Industry Canada in 1998. When this article went to press, there was no response from Industry Canada.

4.5. Institute for Robotics and Intelligent Systems (IRIS)

PRECARN helped create IRIS (Institute for Robotics and Intelligent Systems) in the summer of 1990, a 4-year program, one of the (then) 14 research networks in the Federal Networks of Centres of Excellence. PRECARN administers IRIS, reviewing and approving its projects. This administrative arrangement was formalized with NSERC in a July 1990 agreement.

The mission of IRIS is "to promote high-quality collaborative research in intelligent systems which is of strategic importance to Canadian industry and to strengthen the R&D interaction between universities and industry, thereby improving the competitiveness of Canadian firms" (PRECARN, 1997, p. 3).

IRIS has been through two phases and has entered a third. IRIS-I was a $24.8 million research program that was completed in June 1994. IRIS-II, a $25 million program, ran from June 1994 to June 1998. IRIS-I conducted research in three major areas: computational perception, knowledge-based systems, and intelligent robotics systems. IRIS-II conducted research in the areas of intelligent computation, human-machine interfaces, machine sensing and actuation, and integrated systems. Funding for IRIS-III is through to 2005, subject to a satisfactory mid-term review. The first four years, from 1 April 1998 to 31 March 2002, has received $17.5 million in funding (PRECARN, 1997, p. 26).

The first IRIS/PRECARN conference was held 12-14 June 1991, attended by approximately 290 people. There have been annual IRIS-PRECARN conferences ever since.

4.6. TeleLearning-NCE

Human Resources includes the TeleLearning-NCE (http://www.telelearn.ca), which is funded for 1995-2002. Having been given a positive mid-term review, the NCE has been granted $12.8 million for 1998-2002. The original TeleLearning mandate explicitly mentioned the use of AI techniques, and many current projects do, especially those in Theme 6 (coordinated by Gilbert Paquette and Tom Carey), e.g., those of Gordon McCalla and Jim Greer, and of Claude Frasson.

4.7. Canadian Space Agency (CSA)

The Canadian Space Agency (CSA) has conducted AI research. Some of this work has been in collaboration with PRECARN.

The Canadian Space Agency (http://www.cta-otc.gc.ca) was announced on 1 March 1989. The agency sponsored the Strategic Technologies for Automation and Robotics (STEAR) program as part of Canada's International Space Station Program. STEAR was established in 1987 and helped develop the Shuttle Remote Manipulator System (SRMS or Canadarm) and Mobile Servicing System (MSS), "a multi-armed robotic unit which will be used to assemble and maintain the Space Station" (CAI/IAC, July 1990, p. 7), due to be launched in late 1997. Among the STEAR AI projects was an expert system for training astronauts and technical personnel (Ibid.).

On 17 August 1994, the Canadian Space Agency (CSA) and PRECARN signed a cost-sharing agreement by which the CSA would contribute $1.25 million to STEAR, industry would match the $1.25 million, and PRECARN would contribute $2.25 million. This money was to go into developing intelligent systems technologies for the Mobile Servicing System robotic unit mentioned previously, was successfully completed in 1998 and is to be integrated with the remainder of the space station.

4.8. Artificial Intelligence Research and Development Fund

On 5 April 1990, the federal government announced the

Artificial Intelligence Research and Development Fund, a $10 million five-year program, part of the Strategic Technologies Program of Industry, Science and Technology Canada (ISTC). The Strategic Technologies Program "assists Canadian industry in responding to the challenge of rapid technological change in information technology, biotechnology and advanced industrial materials" (CAI/IAC, July 1990, p. 8).

The Artificial Intelligence Research and Development Fund is "a procurement-based program which will use the federal government as a test bed to assist in the development of private sector capabilities in artificial intelligence" (Ibid.). Together with the announcement of the fund, it was reported that eight projects, costing a total of $2 million, would be funded immediately. These projects were being undertaken by "the Department of Communications, Energy, Mines and Resources Canada, Environment Canada, Health and Welfare Canada, Transport Canada, and the Department of Defence, on a cost-shared basis" (Ibid.). "The Department of Supply and Services manages the process of contracting out the work on these systems according to its standard procedures. Each sponsoring department is responsible for the management of its own project" (Ibid.).

On 16 January 1991, a further $3.9 million was allocated, meaning that $6.3 million had been allocated from the original $10 million to a total of 20 active projects (CAI/IAC, February 1991, p. 5). The twenty projects are listed in Canadian AI/IA au Canada (Ibid., pp. 6-8).

It is not clear what happened to this fund after 1995 (no results found from a search for the fund at Strategis). The idea for this fund seems to resemble the recommendations on procurement made in the study conducted by the Nordicity Group and CAIP Corporation in 1986 (see section 6).

4.9. Microelectronics and Systems Development Program (MSDP)

In a section of the 1989 survey of AI in Canada on "Government Support Programs for AI R&D in Canadian Industry," there was mention of the Artificial Intelligence Research and Development Fund and the Strategic Technologies Program. Also mentioned was the Microelectronics and Systems Development Program (MSDP), which was funded by Industry, Science and Technology Canada (ISTC). "For AI projects to qualify for support [in the MSDP], the software development activities must be an integral part of the development of a technology or system qualifying for support under MSDP" (ACAI, 1991, p. 53).

4.10. Technology Outreach Program (TOP)

Also mentioned in the 1989 survey is the Technology Outreach Program (TOP), funded by the ISTC. "[The program] improves productivity and competitiveness of Canadian industry by supporting technology centers whose national activities and services accelerate the acquisition, development, and diffusion of technology, especially in small- and medium-sized enterprises" (ACAI, 1991, p. 53).

4.11. Federally-Funded Regional Support Programs

A number of regional programs are also mentioned in the survey that "may complement the research and development activities of Canadian AI firms" (ACAI, 1991, p. 53). These include the Western Diversification Program and the Atlantic Canada Opportunities Agency (ACAI, 1991, p. 53).)

4.12. CANARIE

Some AI work has been conducted as part of CANARIE (http://www.canarie.ca), though not much. CANARIE is a non-profit, industry-led and managed consortium that has the goal of stimulating the development of the Information Highway in Canada (http://www.canarie.ca/eng/org/about_e.html). (The connections between AI and the Internet were discussed in a special issue of Canadian AI/IA au Canada (summer 1995)). However, CANARIE's work has had little direct bearing on Canadian AI. Only one of the 43 projects approved in CANARIE Phase-I involved artificial intelligence. This was InfoScan, e-mail filtering software developed by Machina Sapiens, a company with experience in AI and NLP (results from Strategis search). Also, CANARIE's Healthcare division, Technology and Applications Development (TAD) program, includes an On-Line Interactive Health Care System that uses AI techniques based on medical modelling (results of search at http://www.canarie.ca/frames/startsearch_e.html).

End of Part 1.

... Part 2 in next issue of *CAI/IAC*

Margaret Dalziel

Director, Strategy and Marketing, PRECARN

http://www.precarn.ca/

Proposal for PRECARN-III Submitted

In October 1998, PRECARN submitted to Industry Canada its proposal for a new, third phase of collaborative research extending through 2005. At the time of going to press, there was no news from Industry Canada about the proposal.

PRECARN Phase III: Smart Systems for Smart Industries

PRECARN Phase III is premised on the need for greater scale and stronger links with 10 important sectors of the Canadian economy (see Figure 1). Greater scale will help ensure that Canada keeps up with the rapid development of intelligent systems technologies taking place in other advanced nations and will permit the development of high impact research projects in all 10 targeted sectors. Intelligent systems are key to productivity increases and the creation of new value-added products in a number of these sectors ranging from forestry to agri-food to manufacturing. New intelligent systems-based products are the basis of emerging sectors such as medical imaging and virtual reality.

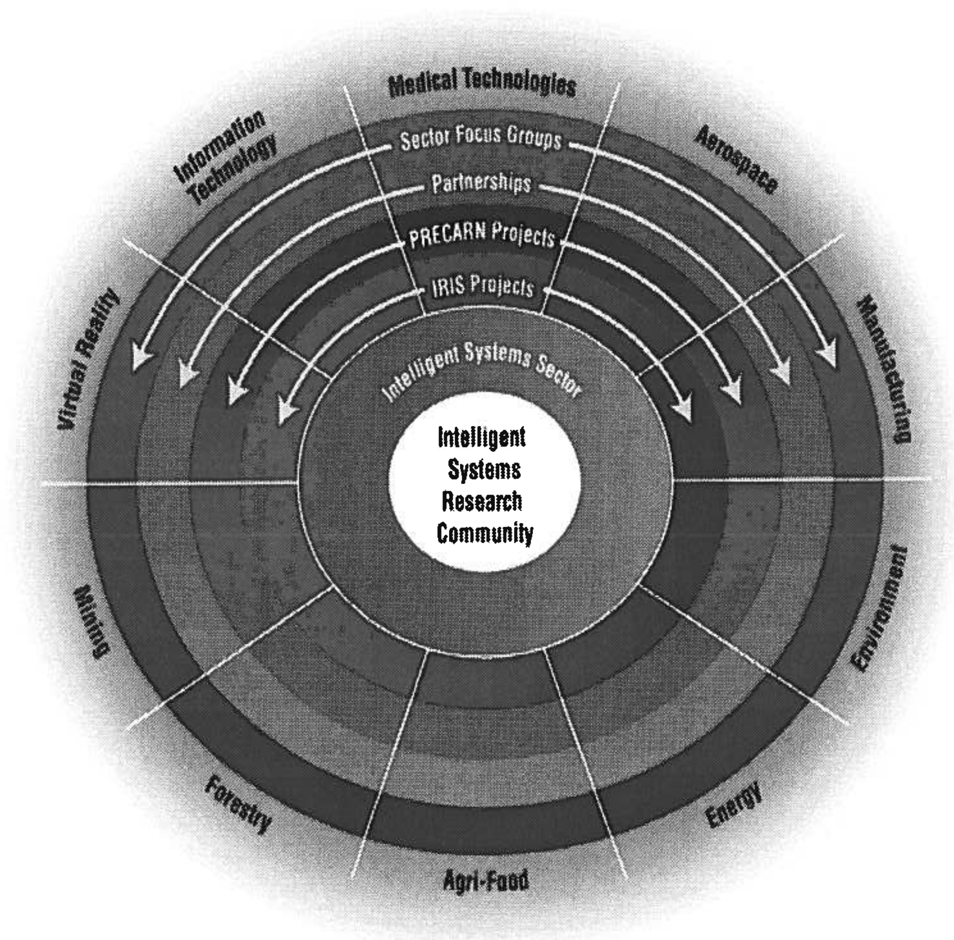

Figure 1: Intelligent systems research, IRIS, PRECARN and 10 targeted sectors.

PRECARN's core constituency, the intelligent systems sector, is undergoing a period of tremendous growth. Advances in low cost, high performance computers, software, and sensors are enabling the creation of systems which can truly deliver on the early promises of artificial intelligence research of the '60's and '70's. All industrialized nations are investing in intelligent systems research—global government investments in intelligent systems research conducted by industry are estimated at between $2 and $3 billion.

Achievements of PRECARN-I and II

PRECARN, which stands for PRE-Competitive Applied Research Network, was created in 1987 to encourage Canadian companies, government research organizations, and universities to work together to create and deploy intelligent systems technologies. Achievements include:

- **A network which extends from coast-to-coast**
  - Almost 50 members, half of whom are small and medium-sized enterprises
  - Over 1200 people and 100 companies and organizations involved in 32 major collaborative projects
  - 300+ companies and organizations active in a nation-wide network
  - 10 important sectors of the economy engaged

- **A leadership role in university-industry collaboration**
  - PRECARN created IRIS (Institute for Robotics and Intelligent Systems), a leading Network of Centres of Excellence, and is now guiding the start-up of IRIS Phase III
  - PRECARN-IRIS wins NSERC-Conference Board Synergy Award in 1998 as an innovative model of university-industry interaction

- **Over 100 new technologies created**
  - 60% already embedded in new products and processes

- **Between $270 and $450 million in future profits forecast over 10 years**
  - Estimates developed by independent consultants based on a sample of five PRECARN projects
  - Industry Canada investment in the projects was $10.7 million

- **A significant and growing intelligent systems sector**
  - 250 firms, 40% founded since 1990
  - $3.8 billion in revenues, majority exports
  - Annual growth exceeds 10% for most firms
  - 23,000 jobs in all regions of the country

Vision Interface '99

Delta Hotel, Trois Rivières
Québec, Canada

19-21 May 1999

Computer Vision for Industrial Applications

VI '99 — the twelfth Canadian conference devoted to computer vision, signal and image processing, and pattern recognition — is sponsored by the Canadian Image Processing and Pattern Recognition Society and the International Association for Pattern Recognition. VI '99 will have a theme: "Computer Vision for Industrial Applications." The 1999 conference will be held in conjunction with Quality Control by Artificial Vision.

WWW: http://www.DMI.USherb.CA/conferences/VI-99/

Les linguistiques de corpus

Benoît Habert, Adeline Nazarenko, André Salem

Paris: Editions Armand Colin, 1997, ISBN 2-200-01775-8, FF 125.-

Comte rendu par
Caroline Barrière
Université Ottawa

Le livre *Les linguistiques de corpus*, des auteurs Benoît Habert, Adeline Nazarenko et André Salem, paru aux éditions Armand Colin, 1997, contient 218 pages qui se lisent très aisément. Les auteurs présentent un survol des techniques et problèmes associés à la constitution de corpus, à l'analyse manuelle et automatique de corpus et à l'interprétation des résultats statistiques obtenus sur ces corpus.

Le livre ne présente pas une analyse poussée, un travail original, ou une critique, mais plutôt offre au lecteur une organisation par thèmes des travaux de plusieurs chercheurs qui touchent de près ou de loin la linguistique de corpus. Je le recommenderais comme bouquin de référence pour quiconque s'intéresse aux corpus, du chercheur établi au nouvel étudiant gradué débutant sa thèse. Même si le livre est écrit en français, par des auteurs travaillant en France, leur survol relate surtout de travaux se basant sur des corpus en anglais et américain (les auteurs distinguent les deux "langues"), pas par choix, mais plutôt par la force des choses, étant donné un nombre beaucoup plus vaste de travaux effectués.

De par sa nature "survey," ce livre sera bientôt désuet, malheureusement. C'est une réalité à laquelle font face tous les chercheurs en linguistique computationnelle; la recherche et les ressources disponibles dans ce domaine évoluent très rapidement. D'autre part, de par sa nature didactique, ce livre pourra garder sa place comme livre de référence au fils des ans. "Les linguistiques de corpus" donne au lecteur beaucoup d'information technique, et ce dans un langage près de la vulgarisation.

Le livre donne une introduction et par la suite se divise en 3 parties: (i) les corpus annotés et leurs utilisations, (ii) les dimensions transversales, (iii) méthodes et techniques.

Dans l'introduction, tout d'abord, les auteurs présentent différents corpus, commentent sur leur accès de plus en plus facile et leur taille de plus en plus grande. Les corpus présentés tout au long de l'ouvrage sont: Brown, LOB (Lancaster-Oslo-Bergen), Susanne, London-Lund, Lancaster/IBM Treebank, Helsinki, Archer, BNC (British National Corpus), Penn Treebank. Pour constituer un corpus, les auteurs soutiennent que des textes doivent être sélectionnés et organisés selon des critères linguistiques pour servir d'échantillon du langage. Le débat sur ce qui constitue un "corpus représentatif" ne peut pas être plus actuel, tel qu'ont pu voir récemment les lecteurs du groupe électronique de nouvelles "corpora."

La première partie du livre présente des corpus avec leur type d'étiquetage. Deux aspects intéressants discutés par les auteurs sont: (1) la dépendence du choix d'étiquetage sur les buts subséquents envisagés pour un corpus, (2) l'assujettement des corpus étiquettés aux erreurs humaines; lorsqu'un chercheur travaille en analyse de corpus, il doit s'attendre à des résultats indicatifs, mais non pas parfait.

Les auteurs différencient les corpus étiquettés, et les corpus arborés. Ces derniers sont beaucoup moins fréquents et servent de patrons à des analyses syntaxiques. L'étiquetage sémantique est aussi mentionné, mais *"ces corpus porteurs d'annotations sémantiques n'existent qu'à l'état embryonnaire."* Les auteurs discutent des énormes problèmes que présentent l'annotation sémantique. Tout d'abord il n'y a pas de standards établis; ainsi, un étiquetage utilisant une ressource sémantique (un dictionnaire particulier) donnera des résultats tout à fait différents d'un étiquetage fait à l'aide d'une différente ressource (un autre dictionnaire, un "thesaurus" ou Wordnet par exemple). Après la parution de ce livre, en 1998, une première tentative de comparaison d'étiquetage sémantique automatique, mise sous forme du concours SENSEVAL, s'est effectuée.

La deuxième partie présente l'analyse de corpus vers des applications particulières. (1) Analyse de textes et assignation de sens aux mots. Cela retouche les problèmes de compatibilité de ressources présentés à l'étape précédente et mène les auteurs à présenter l'idée de construire ses propres ressources à l'aide de textes. Ils présentent les travaux de Grefenstette sur la construction automatique d'entrées de dictionnaire. (2)

Analyse du langage au fil des ans. (3) Comparaison des langues. Ils présentent les corpus alignés et différents types d'analyses pouvant être effectués sur ces corpus.

La troisième partie reprend plusieurs thèmes mentionnés dans la partie I lors de la présentation des corpus, mais maintenant présente plus en détails les méthodes et techniques. On reprend les thèmes de la constitution d'un corpus, des normes possibles, des critères pour le choix des textes. Les auteurs reviennent sur l'annotation d'un corpus, mais cette fois présentent les méthode d'annotation automatique contrairement aux méthodes et problèmes présentés au chapitre 1 sur l'étiquetage manuel. Une dernière section parle des analyses statistiques, de la "quantification des phénomènes langagiers."

Les linguistiques de corpus présentent bien les défis fascinants qui s'ouvrent devant nous par le fait que l'informatique et la linguistique se voient enfin "forcés" à former des associations, à joindre leurs forces pour permettre de continuer des recherches empiriques, pour faire des analyses de groupe de textes, ceux-ci qui mèneront à diverses découvertes et permettront à la linguistique computationnelle de pousser plus loin les limites.

Caroline Barrière s'intéresse au traitement du langage naturel, à l'acquisition de connaissances à partir de textes et dictionnaires, et à la représentation de ces connaissances.

The Cambridge Quintet: A Work of Scientific Speculation.

John L. Casti

Reading, MA: Little, Brown/Addison-Wesley, 1998, ISBN: 0-201-32828-3 181 pp.

This novel describes an imaginary dinner party. The host is C. P. Snow, who has been asked to advise government on the feasibility of building a machine that "thinks like a human being." Invited dinner guests are Alan Turing, Erwin Schrödinger, J. B. Haldane, and Ludwig Wittgenstein.

The Twelfth International Conference on

Industrial & Engineering Applications of

Artificial Intelligence & Expert Systems IEA/AIE-99

Le Meridien Cairo, Cairo, Egypt

May 31- June 3, 1999

Sponsored by:

International Society of Applied Intelligence

Organized in Cooperation with:

AAAI, ACM/SIGART, CSCSI/SCEIO, IEE, INNS, JSAI, SWT

WWW: http://mason.gmu.edu/~iimam/ieaaie99/ieaaie99.html

CALL FOR PAPERS

The Thirteenth International Conference on

Industrial & Engineering Applications of

Artificial Intelligence & Expert Systems IEA/AIE-2000

New Orleans, Louisiana, USA

June 19-22, 2000

Sponsored by:

International Society of Applied Intelligence

Organized in Cooperation with: (confirmations are pending)

AAAI, ACM/SIGART, ECCAI, CSCSI/SCEIO, INNS, JSAI, IEE, SWT, USL

Authors are invited to submit one copy of their paper, written in English, of up to 10 single spaced pages, presenting the results of original research or innovative practical applications relevant to the conference. Practical experiences with state-of-the-art AI methodologies are also acceptable when they reflect lessons of unique value to the conference attendees. Shorter works, up to 6 pages, to be presented in 10 minutes, may be submitted as SHORT PAPERS representing work in progress or suggesting possible research directions. **Submissions** must be received by **17 November 1999**.

WWW: http://www.cacs.usl.edu/~ieaaie2000/

CSCSI/SCEIO Membership Application

I wish to join CSCSI/SCEIO and receive *Canadian Artificial Intelligence/Intelligence Artificielle au Canada*

Yes ___ / No ___ Web Access Only ($30.00* Cdn./yr.)

Yes ___ / No ___ Printed Copy and Web Access ($40.00* Cdn./yr.)

I am a student

Yes ___ / No ___ Web Access Only ($15.00* Cdn./yr.)

Yes ___ / No ___ Printed Copy and Web Access ($15.00* Cdn./yr.)

I am a member of CIPS

Yes ___ / No ___ Web Access Only ($25.00* Cdn./yr.)

Yes ___ / No ___ Printed Copy and Web Access ($30.00* Cdn./yr.)

*Includes applicable G.S.T.

Name ....

Job Title ....

Company ....

Address ....

....

....

....

E-mail ....

Phone ....

Fax ....

Visa or Amex Number ....

Expiry Date ....

Signature ....

Please send your membership application to:

Mail: CIPS National Office, One Yonge St., Suite 2401, Toronto, Ontario M5E 1E5

Phone: (416) 861-2477

Fax: (416) 368-9972

E-mail: dawn@cips.ca

For more information, contact CIPS or a member of the CSCSI/SCEIO executive.

Canadian Artificial Intelligence
Intelligence Artificielle au Canada

Planned for Issue 44

Bayesian Q-Learning
Richard Dearden

The Infrastructure of Artificial Intelligence R&D in Canada, Part 2
Dan Fass

Twenty Five Years of CSCSI/SCEIO and CAI/IAC
Dan Fass

Projeté pour la 44^{e} édition

Apprentissage-Q bayesien
Richard Dearden

L'infrastructure de la R-D en intelligence artificielle au Canada, 2^{e} partie
Dan Fass

Vingt-cinq ans de SCEIO/CSCSI et IAC/CAI
Dan Fass

Submission information • Guide de soumission

Canadian Artificial Intelligence / Intelligence Artificielle au Canada welcomes submissions on any matter related to artificial intelligence. Please send your contribution, preferably in electronic form, with an abstract to:

Dan Fass
Editor, *Canadian AI / IA au Canada*
School of Computing Science
Simon Fraser University
Burnaby, British Columbia, Canada
V5A 1S6 or — fass@cs.sfu.ca
Phone (604) 291-3208
Fax (604) 291-3045

Advertising rates are available upon request from the editor.

Book reviews and candidate books to review should be sent to:

Graeme Hirst
Canadian Artificial Intelligence / Intelligence Artificielle au Canada
Department of Computer Science
University of Toronto
Toronto, Ontario, Canada
M5S 1A4 or — gh@ai.toronto.edu

Recycled / Recyclable