MASSIVELY PARALLEL AI*

DAVID L. WALTZ

Thinking Machine Corporation
245 First Street
Cambridge, MA 02142-1214, USA

Brandeis University
Computer Science Department
Waltham, MA 02254, USA

ABSTRACT

Most AI researchers would, I believe, agree that truly intelligent machines (i.e., machines on a par with humans) will require at least four orders of magnitude more power and memory than are available on any machine today [20,32]. There is now widespread agreement in the supercomputing community that by the year 2000 all supercomputers (defined as the most powerful machines available at a given time) will be massively parallel [10]. Yet relatively little thought has been given in AI as to how to utilize such machines. With few exceptions, AI's attention has been limited to workstations, minicomputers and PCs. Today's massively parallel machines present AI with a golden opportunity to make an impact, especially in the world of commercial applications. The most striking near-term opportunity is in the marriage of research on very large databases with case-based and memory-based AI. Moreover, such applications are steps on a path that can lead eventually to a class of truly intelligent systems.

1. Economic inevitability of massive parallelism. The performance of serial computers is limited by the "von Neumann bottleneck", (the serial path used to move instructions and data between memory and the CPU) and by I/O limitations. Over the next ten years the use of fast technologies (ECL, GaAs, etc.) and further miniaturization might gain a performance factor of five, cleverer caching and instruction prefetch a factor of two, and the use of multiple functional units yet another factor of four, bringing the fastest uniprocessors (now at less than one GFlops) to perhaps 40 GFlops. Compiler technologies could allow as many as 16 such processors to

be ganged together, yielding perhaps as much as 640 GFlops in total. The price would be high: very dense chips of exotic materials packed close together present daunting cooling and packaging problems. An extrapolation of current trends (supercomputers cost about \$100,000/MFlops in 1977, and about \$8000/MFlops in 1990) suggests that such a machine would cost about \$500 million in the year 2000 (about \$800/MFlops).

In contrast, massively parallel machines (e.g., the Connection Machine[R] CM-2) use PC/workstation technologies, and boast per-computational-unit costs similar to those of these small machines: on the order of \$1000/MFlops in 1990. The largest massively parallel machines already exceed the power of serial supercomputers: the 65,536 processor CM-2 is realistically capable of speeds on the order of 5 GFlops (with a 28 GFlops peak). DARPA has targeted a massively parallel TeraOps (one trillion operations/second) machine by 1995, and the cost is expected to be less than \$100 million (\$100/MFlops). By the year 2000 it should be possible to build massively parallel TeraOps machines for \$10 million (\$10/MFlops — two orders of magnitude cheaper than could be done with serial technologies).

2. A brief history of massively parallel AI. The Connection Machine system was originally designed to handle AI tasks, such as NETL-like marker-passing over semantic networks, and low-level computer vision [9,11]. *LISP, an extension of COMMONLISP, was the first high-level language for the CM and its first front-end computer was a Symbolics 3600. A number of the early CMs were purchased by AI labs following its introduction in 1986, and some early work was done along the lines envisioned by Hillis, for example, CIS, a marker-passing parallel expert system, with one (instantiated) rule per processor [3], and a system for computing stereo disparity from a pair of images, using the Marr-Poggio method [14,6]. (See [35] for a summary of several early AI projects.) However, it is clear in retrospect that there was not much of an AI market for \$50,000 Lisp machines, let alone for \$1 million (and up) Connection Machine systems performing these sorts of applications. There is still not a market for very large expert systems, in part because it is difficult to build a very large expert systems (Blelloch's CIS system contained 100,000 artificially generated rules on a CM-1 with about 8 MBytes of memory; a current CM-2 with 8 GBytes of memory could store 100 million rules!). Likewise, there was not (and still is not) a market for high volume low-level image processing (although the new Landsats may eventually create one).

Fortunately for Thinking Machines, the CM-2, with its floating point option (using up to 2000 Weitek chips), proved to be an excellent match for a wide range of scientific problems: finite element models, operations on large matrices, fluid flow and aerodynamics models, n-body problems (e.g., galactic collisions), interactive scientific visualization, models using cellular

automata, seismic data processing, signal processing, etc. Most of Thinking Machines' sales growth, 50% per year over the last few years, has been fueled by sales for these types of applications.

But the AI dream is still valid. Recent research using novel AI paradigms on real customer problems and real customer databases have identified some great opportunities.

3. The large-scale commercial AI opportunity. Data is being generated faster than it can be digested. Very large (> 100 GByte) databases are becoming common. Such databases contain valuable information on customer credit and buying behaviour, many forms of text: public text, including books, newspapers and news wires, financial journals, technical industry publications, scientific journals and annual reports; and private text, such as studies, memos, manuals, proposals, documentation; visually oriented materials, including maps, schematics, plans and blueprints, as well as fast-expanding video archives; and much more.

Until recently mainframe computers, with their extensive existing software libraries, have been the only choice for those who own and/or wish to mine such databases. But mainframes are pitifully slow and at the same time very expensive, and the software available for database tasks has been quite primitive because answering even primitive questions takes too long. Massively parallel machines offer large storage capacity[1] and high I/O rates, using parallel disk arrays and multiple wide-word I/O channels. The total CM-2 I/O capacity is currently 200 MBytes/second, limited by the disk rates.[2] This means that an entire 100 GByte database can be streamed into the CM-2 in 500 seconds. We estimate based on preliminary but realistic tests, that we can achieve a speedup of at least an order of magnitude over identical data selection operations on a mainframe on a CM-2 costing only a fraction of the price [Stanfill, forthcoming].

This kind of performance opens the possibility for changing the fundamental nature of an organization's use of a database, from a weekly to a daily batch run, or from batch to interactive, and also opens the opportunity for applying dramatically more complex and intelligent processing to each database item than can otherwise be imagined today. In the next several sections, I outline the algorithms and performance for several commercially important applications.

4. Text-based intelligent systems. For most of its history, AI has been concerned with "toy problems." Scaling up presents difficulties: At the

---

[1] A single current Connection Machine CM-2 can be equipped with up to eight I/O channels, each of which can have 16 drops, allowing 15 Data Vault disk arrays of up to 80 GBytes each, for a total storage capacity of $8 \times 15 \times 80$ GBytes = 9,600 GBytes!

[2] The CM-2 itself can handle more than 400 MBytes/second.

two extremes of the spectrum, one can hand-code (as in CYC [14]), or one can use methods to automatically build NLP systems. To date there have been very few practical applications for natural language processing, and this fact has dampened the enthusiasm of funding agencies and companies that support research in this area Fortunately, there are signs that this situation can be improved by strategically merging AI/NLP and Information Retrieval (IR) technologies. This offers novel opportunities both for learning research and for building systems that have immediate practical value.

A series of experiments and discoveries led researchers at Thinking Machines, most notably Craig Stanfill and Brewster Kahle, to devise a document retrieval system that works in parallel on the Connection Machine [28,27]. The resulting system, marketed as DowQuest[R] by Dow Jones has been in commercial use since January 1989. DowQuest provides a high quality search through a clever interface that can be used effectively by a computer-naive person after only about 5 minutes of training. The basic idea is this: A database of documents (e.g., news articles, abstracts, books, etc.) is distributed to each of the 65 536 processors of a CM-2 (if documents are 2K bytes long, each processor equipped with large memory can hold about 256 compressed documents or 16 million total!). The user types a few words (a question, description, or a list of terms will do) and a carriage return; the terms are broadcast to all the processors in parallel along with a numerical "weight" indicating the importance of each term.[3] The search portion of this operation is analogous to the following situation: imagine a stadium with 65 536 people, each with one document and a pocket calculator; an announcer reads each search term followed by a number (representing the importance of a term (the rarer a term, the higher its weight) and each person whose document contains the term adds the score to the calculator. After all terms are read, the persons with the highest scored documents present themselves. (This would be hard in the stadium analogy). (To match the real situation, we would actually need either 16 million people, or 256 documents and 256 calculators per person.)

The headlines for the documents with the highest total scores are then sent to the user. The user can view the text of each of these documents by clicking a mouse while pointing to the headline. When a user sees a document (or paragraph of a document) that answers his/her request, the user can mark the document "good" by pointing and clicking the mouse. The system collects all the terms from all the documents marked "good" along with all the initial words the user typed, and repeats the search process described above, but now using all these terms. Each search requires less than a second, even on databases up to 10 GBytes. This method, called

---

[3] Term weights are assigned automatically by a program that pre-processes the text and updates the database. The number of occurrences of each term is saved, and weights computed proportional to the negative log of the probability of occurrences of each term.

"relevance feedback" [18], generally produces a substantially better search than is possible with Boolean search systems [2].

DowQuest: DowQuest uses the method above, with enough memory so that a 1 GByte database is permanently stored. The 16 documents with the best scores are returned to the user, who can look at the full text (stored on a disk on the front end server) and can mark any document (or paragraph of a document) as "relevant". DowQuest has about six months worth of articles from nearly 300 different sources: *Wall Street Journal, Washington Post, Barrons*, the business sections of about 100 U.S. newspapers, *Fortune, Forbes*, and other magazines, and a number of trade publications (e.g., *Byte*). Ninety-nine percent of all searches take under 1.5 seconds, including front end time. CM time for 100 terms is about 170 ms. The system also parses input text, creates a new surrogate database, and updates the database while the system is operating on-line.

We have recently outlined algorithms that allow interactive access to 1 TeraByte (1 000 000 MBytes) of text [29]. For reference, the Library of Congress probably contains on the order of 40 Terabytes of text.

5. New text opportunities. This system suggests interesting opportunities for exploiting natural language processing results:

Adding NLP: It is highly desirable to add natural language pre- and post-processing to the existing system, to improve its performance, and to extend its capabilities. For example, we are building recognizers that can find, label, and store lists of terms that refer to company names, geographic locations, names of persons, etc. Ultimately, this will help users to ask and obtain answers to questions that would be very difficult to phrase as Boolean queries. For example, "Earnings reports for New England utility companies" would expand to "Earnings reports for Maine, New Hampshire, Vermont, Massachusetts, Connecticut, Rhode Island, utilities, power companies, electric companies, power, light ...". In addition, natural language processing systems will allow us to post-process retrieved documents, to filter out irrelevant articles, and thus improve the performance of the system from the user's point of view.

Adaptive Systems: More intelligent processing could be applied to the users' queries; the system could keep track of user patterns and interests, and adapt itself to be easier to use or even to volunteer information it thinks the user is interested in.

WAIS (Wide Area Information Server): We are building a version of this system to seamlessly search an organization's local data as well as remote (e.g., Dow Jones) databases. This system will allow organizations to locate and reuse proposals, reports and studies, find (by matching biographies) appropriate people for various tasks, and generally allow each person to

locate any text, personal, corporate or public, with a single search.

Hypertext: The same methods that let us locate relevant articles can allow us to automatically build hypertext systems for test distributed on CD-ROMs. These could be generated much more quickly and inexpensively than is possible with hand-building.

Automatic Generation of NL Systems: The retrieval system itself can be adapted to extract phrase, sentence and paragraph "templates" or patterns in order to aid the building of recognizers for particular topics or for types of stories. Such processing can provide empirical data on language usage that would be very difficult to find or invent any other way, leading to "dictionaries" of multi-word and multi-sentence language patterns and to FRUMP-like systems [5] with broad subject coverage.

Other related research, using dictionaries or thesauruses, has become popular in recent years. Some striking successes have been achieved by Ken Church and coworkers at AT&T Bell Laboratories [4,7] using the augmented "Brown Corpus" [13]. The "Brown Corpus" consists of one million words of text, chosen to represent a wide range of text types and styles (newspaper and magazine articles, books on history, economics, etc.). It was "augmented" by Kucera and Francis by assigning each word in the corpus to one of about 450 classes, covering standard grammatical categories (noun, verb, adjective) but also including substantially finer distinctions (e.g., noun-agent of sentence; verb-complements of particular types). Church collected statistics on the probabilities that various words would follow particular other word (or category) combinations. This system has been used to judge the most likely categories for words in novel text taken from-news-wire sources. Success rates for Church's system are in the range of 98–99%, much higher than for the best syntactic parsers (in the range of 33% [19]).

All these current lines of research emphasize *breadth* of coverage rather than *depth* of coverage and are thus complimentary to the goals of traditional AI-NL processing research. All present attractive alternatives to hand-coding [14]. And all can be used to accelerate the research into deep processing. The most attractive part of this effort is that our systems are immediately useful, and thus can pay for the research on their own augmentation. I believe these general approaches will have great importance in the ultimate story of the achieving of truly intelligent systems.

6. Memory-based reasoning. Methods broadly analogous to the text-search algorithms can be used to build "memory-based reasoning" systems to aid in decision-making. These systems perform like artificial neural nets [17,34] or ID3-like learning systems [16]. In memory-based reasoning (MBR) a parallel machine is loaded with a database of the sort that can be used as a training set for learning systems: situations together with actions,

classifications, or desired outputs for each situation. When a new problem is encountered, the MBR system compares it to *all* the known prior cases, and uses the most similar case (or majority vote of several similar cases) to classify the new case. The key to successful MBR operation is the selection of a good similarity metric for matching new problems with known cases [30].

6.1. Advantages of MBR. MBR provides expert system-like behaviour, but does not require extensive hand-coding. MBR provides "explanations", the precedents most similar to the current problem case, in order to justify its actions. MBR provides high performance, superior to artificial neural nets [37] and superior to expert systems (see below). MBR is robust when noise is added to its database; in one experiment [31] on the NETtalk database [22], performance fell off only about 10% when 90% of the database was replaced with noise. MBR is simple to update: new cases can be added and old, obsolete ones removed, and performance will immediately track the changes. This is in sharp contrast to artificial neural nets, which must be totally retained if the world changes, and expert systems, which are notoriously difficult to modify.

There are some disadvantages: MBR generally requires a data parallel computer, which will probably be more expensive than a system used to deliver an expert system or neural net application (though inexpensive data parallel systems without extensive interprocessor communication might suffice); and MBR systems do not operate as rapidly as a trained neural net, though they would generally be much faster than expert systems.

6.2. Classifying census data. We have recently demonstrated that an MBR system can perform impressively on a task to generate one of about 241 industry codes and one of about 509 occupation codes for individual respondents, by comparing their answers (expressed as free text and multiple choice selections) with 132 000 cases that have already been classified by hand. Early results have indicated that by keeping categories where the system has been proven to be correct at least 90% of the time for industry codes and at least 86% of the time for occupation codes, MBR can correctly process at least 70% of the database for industry codes and about 56% of the database for occupation codes. For comparison, an expert system that required more than two years to develop, achieves only 57% and 37% of the database respectively on these two tasks [25]; the MBR system took less than a month to build. For this application, the similarity metric is generated, using statistical operations, on the fly.

6.3. Other applications. Similar methods have been used to build MBR systems for optical character recognition, based on a large number of

examples of handprinted numerals [26]; for medical diagnosis [30]; for controlling a robot to produce near optimal trajectories [1]; for (two-dimensional) object recognition [32]; for automatically generating index terms for news articles or routing articles to appropriate recipients [24]. Recent work [38] on protein structure prediction has shown that a system that combines the results of MBR with neural nets and statistical information dramatically outperforms any previous method. This has relevance to the human genome project, another good target area for AI and IR. Many other applications are clearly possible.

7. Genetically-inspired methods. Market research is a "forest for the trees" problem. One needs to generate insights into the repetitive preference patterns among millions of customers and distill market segment definitions in order to offer consumers the products they are most likely to want. We have developed genetically-inspired algorithms [23] that automatically find trends and categories without being told in advance what the patterns are.

One specific problem that has been addressed by these methods is the following: Suppose that we know purchasing behaviour, demographic and credit information for several millions (or tens of millions of) people, and that we wish to mail catalogs containing items selected from the offerings of hundreds of vendors, such that the greatest possible return (dollar amounts ordered minus the cost of the mailings) is maximized. Clearly, the larger the catalogs, the more the cost for postage and the greater the chance each will be thrown away; the smaller and more tailored the catalogs, the better the return, but the more expensive it will be to print the catalogs and stuff the appropriate envelopes.

In an example run, we started with about 8000 customers on a small CM-2, and first calculated the ideal catalog of five items for each customer. (This step requires a model of consumer behaviour.) Each list of five items is analogous to a piece of genetic material. We also computed an expected return (negative) for sending 8000 tailored catalogs. We then used the communications system of the CM-2 to randomly pair up consumers in parallel. For each pair, we then calculated the change in expected return if consumer 1 took consumer 2's catalog, and vice-versa. Both consumers were grouped into one or the other of the catalogs with a probability based on the change in the expected returns. The scoring scheme also makes it easier to merge a consumer who shares a catalog with a small number of others into a larger group than to pull a customer out of a large group and into a smaller one. Every few steps, random point mutations to the catalog's customers were probabilistically introduced, on the theory that the best catalogs may not have been present for *any* of the original consumers. The process continued until the maximum expected return point was found (in this case, 30 tailored catalogs, and sets of consumers who should receive each).

The overall solution to this problem required on the order of two hours on a 4K processor CM-2; the overall potential search space of solutions is on the order of 10^{87}!

8. A different route to truly intelligent systems. So what does all this have to do with cognition? I want to argue that the basic associative memory operation of selecting relevant precedents in any situation is the essence of what intelligent entities do. ("Precedents" may be actions, options, remindings, etc.) If only a single precedent is found (e.g., when one is operating in a familiar environment on familiar tasks), then there is little involved in acting intelligently. Only when two or more incompatible precedents are found, or when the task space is unfamiliar, is reasoning (in the ordinary sense) required. Combinatorially explosive search can be avoided, since in any given situation only a small number of "operations" (actions) are plausible, making branching factors manageable. Planning can be supported by associative memory retrieval of precedents of the form: [hypothetical situation + goal + operator → new situation] and/or [hypothetical situation → goal]. Even creativity or analogical problem solving might be covered (if the best precedents match structurally but are not literal matches).

I am not imagining that a monolithic flat database could model memory. First, generalizations over memory and other structures need to be matched in addition to episodic items. (See [12,8] for descriptions of massively parallel frame systems.) Second, there ought to be situation-specific priming that changes the overall searchable space (or relevance judgements) for precedents. Overall, I am persuaded by society of mind arguments and examples, and feel that the structure of memory also contains many agents responsible for recognizing special situations and either priming or censoring memories.

9. Summary. If AI is to succeed, it is important to find ways to justify ongoing research costs, to substitute profits for promises. Commercial massively parallel applications already offer opportunities for changing the ways business is done, because existing limits on database size and speed of access can be transcended. Moreover, the excess processing capacity of massively parallel systems makes it possible to add greater intelligence to applications. And there are great potential payoffs for this kind of AI: even modest ideas, if they are sufficiently general to apply to an entire large database, can produce results that seem wonderfully magical. Most successes of this sort to date have used data parallel methods, especially memory-based reasoning. MBR applications can often be generated automatically from existing databases. MBR and case-based reasoning may also form the basis of new paradigms for cognition that can scale to human levels as massively parallel machines develop.

REFERENCES

[1] C. ATKESON, *Roles of knowledge in motor learning*, MIT AI Lab. Technical Report 942 (Sept. 1986).

[2] D. BLAIR AND M. MARON, *An evaluation of retrieval effectiveness for a full-text document retrieval system*, Communications of the ACM, 28, (1985), pp. 289–299.

[3] G. E. BLELLOCH, *AFS-1: A programming language for massively concurrent computers*, Tech. Report 918 Cambridge, MA, MIT AI Lab (1986).

[4] K. CHURCH, *A stochastic parts program and noun phrase parser for unrestricted text*, Unpublished manuscript, AT&T Bell Labs, Murray Hill NJ (1988).

[5] G. DEJONG, *An overview of the FRUMP system*, in Strategies for natural language processing, ed. by W. Lehnert and M. Ringle, Hillsdale, NJ, Lawrence Erlbaum Associates (1982).

[6] M. DRUMHELLER, *Connection machine stereomatching*, Proc. of the 5th National Conf. on AI, Philadelphia (1986), pp. 748–753.

[7] E. EJERHED, *Finding clauses in unrestricted text by stochastic and finitary methods*, Unpublished manuscript, AT&T Bell Labs, Murray Hill NJ (1988).

[8] M. EVETT, J. HENDLER AND L. SPECTOR, *PARKA: Parallel knowledge representation on the Connection Machine*, Tech Report CS-TR-2409, Univ. of Maryland (Feb. 1990).

[9] S. E. FAHLMAN, *NETL: A system for representing and using real-world knowledge*, Cambridge, MA, MIT Press (1979).

[10] G. FOX, *The Teraflop Computer*, Computers in Physics 4, No. 1 (Jan/Feb 1990), p. 112.

[11] W. D. HILLIS, *The connection machine*, Cambridge, MA, MIT Press (1985).

[12] J. KOLODNER AND R. THAU, *Design and implementation of a case memory*, Research report. GIT-ICS-88/34, School of Information and Computer Science, Georgia Tech, October (1988).

[13] H. KUCERA AND W. FRANCES, *Frequency analysis of English usage*, Boston, Houghton Mifflin Company (1982).

[14] G. LENAT, M. PRAKASH AND M. SHEPHERD, *CYC: Using common sense knowledge to overcome brittleness and the knowledge acquisition bottleneck*, AI Magazine, 4 (1986), pp. 65–85.

[15] D. C. MARR AND T. POGGIO, *Cooperative computation of stereo disparity*, Science 194, (1976), pp. 283–287.

[16] R. QUINLAN, *Learning efficient classification procedures and their application to chess end games*, in Machine Learning: An Artificial Intelligence Approach, ed. by R. S. Michalski, J. Carbonell and T. Mitchell, Los Altos, CA, Tioga Publishing (1983), pp. 463–482.

[17] D. RUMELHART AND J. MCCLELLAND et al., *Parallel distributed processing*, Cambridge, MA, MIT Press (1986).

[18] G. SALTON, *The smart retrieval system — experiment in automatic document classification*, Englewood Cliffs, NJ, Prentice-Hall (1971).

[19] G. SALTON, *Personal communication*.

[20] R. C. SCHANK, *Dynamic memory*, Cambridge, UK, Cambridge Univ. Press (1982).

[21] J. T. SCHWARTZ, *The new connectionism*, Daedalus 117, No. 1, pp. 123–142 (Winter 1988).

[22] T. J. SEJNOWSKI AND C. R. ROSENBERG, *NETtalk: A parallel network that learns to read aloud*, The John Hopkins Univ. Electrical Engineering and Computer Science Technical Report JHU/EECS-86/01 (1986).

[23] A. SINGER, *Implementations of artificial neural networks on the connection machine*, Tech Report RL90-2, Thinking Machines Corp. (Jan. 1990) (to appear in Parallel Computing).

[24] S. SMITH, *Current limits of text-based intelligent systems*, AAAI Spring Symp. on Text-Based Intelligent Systems, Stanford (Mar. 1990) 130–133.

[25] S. SMITH, B. MASAND AND D. WALTZ, in preparation.

[26] S. SMITH AND H. L. VOORHEES, in preparation.

[27] C. STANFILL, *Parallel computing for information retrieval: recent developments*, Thinking Machines Corporation Technical Report #DR88-1 (1988).

[28] C. STANFILL AND B. KAHLE, *Parallel free text search on the connection machine system*, Communications of the ACM, 29, 12 (1986).

[29] C. STANFILL, R. THAU AND D. WALTZ, *A parallel indexed algorithm for information retrieval*, SIGIR '89 (Proc. 12th Ann. Int. ACMSIGIR Conf. on R&D in Information Retrieval), Cambridge, MA, (June 1989), pp. 88–97.

[30] C. STANFILL AND D. WALTZ, *Toward memory-based reasoning*, Communications of the ACM 29, 12 (1986), pp. 1213–1228.

[31] C. STANFILL AND D. WALTZ, *The memory-based reasoning paradigm*, Proc. Case-Based Reasoning Workshop, Clearwater Beach, Fl. (May 1988), pp. 414–424.

[32] L. TUCKER, C. FEYNMAN AND D. FRITZSCHE, *Object recognition using the connection machine*, Proc. Computer Vision and Pattern Recognition, Ann Arbor, MI (June 1988), pp. 871–878.

[33] D. L. WALTZ, *The prospects for truly intelligent machines*, Daedalus 117, 1 (Winter 1988), pp. 191–212.

[34] D. WALTZ AND J. FELDMAN, eds. *Connectionist models and their implementations*, Norwood, NJ, Ablex (1988).

[35] D. L. WALTZ AND C. STANFILL, *Artificial intelligence related research on the connection machine*, Proc. Int. Conf. on Fifth Generation Computer Systems, Tokyo (Dec. 1988), pp. 1010–1024.

[36] D. L. WALTZ, *Is indexing used for retrieval?* Proc. Case-Based Reasoning Workshop, Pensacola Beach FL. (May 1989), pp. 41–44.

[37] D. WOLPERT, *Generalization theory, surface fitting, and network structures*, Ph.D Thesis, Physics Dept., Univ. of California, Santa Barbara (Winter 1989).

[38] X. ZHANG AND D. WALTZ, *Developing hierarchical representations for protein structures: an incremental approach*, AAAI Spring Symp. on Molecular Biology, Stanford (Mar. 1990), pp. 149–152.