Machine Learning, 18, 115–119 (1995)

Book Review: *Neural Network Perception for Mobile Robot Guidance* by Dean A. Pomerleau. Kluwer Academic Publishers, 1993.

GEOFFREY TOWELL towell@scr.siemens.com

Siemens Corporate Research, Inc., 755 College Road East, Princeton, NJ 08540

Editor: Alberto Segre

1. Overview

More than almost any other body of work in the neural network literature, Dean Pomerleau's work on autonomous vehicle navigation stands out because it describes significant, original, interesting research. *Neural Network Perception for Mobile Robot Guidance* brings together Pomerleau's work in a coherent, easily understood package. As a result, the book is both accessible to people with little background in the field and useful to those with greater experience. Overall, it is a book worthy of being read.

The result of Pomerleau's work, vehicles that are able to drive without human control, is interesting to anyone who has gotten bored while driving. Perhaps it is interesting simply because it allows the anthromorphization of computers. Because of this widespread base of interest, the results have been reported in the popular press. Indeed, it is difficult not to have seen one of Pomerleau's computer-controlled trucks either on television or in a magazine. Such reports make it easy to dismiss Pomerleau's work as "mere engineering." Certainly engineering is important to the success of the project. Original research is, however, an equally important contributor to this success. Without both, ALVINN (an acronym for the project) would have remained a lab curiosity; Pomerleau's ability to combine the requisite engineering with original research make ALVINN a success.

The combination of science and engineering makes it difficult to review this book on a single level. So instead, I will review the book three times. First, I review it as a chronicle of the development of ALVINN in which the book details progress from a vehicle that was unsafe at any speed to one capable of cruising the interstate highway system. Second, I review the book as a report about an achievement and the engineering requisite for that achievement. Third, I review the book as a report on a series of innovations which together represent a significant piece of research. That I found this book successful on all three levels is a testament to the quality of the work described therein.

2. As a chronicle of ALVINN

The easiest way to read this book is as a story about work on a complex, difficult problem. In fact, the story line only lasts through Chapter 5. This makes the story one that is both easily read and fairly brief. Read as a chronicle of research, the book is reminiscent of books about the development of computers such as Tracy Kidder's *The Soul of a New Machine* (Kidder, 1981). Unlike these other mass-market books, Pomerleau suppresses his role and that of all other people who worked on ALVINN. As a result, there is a tendency for the story to fall into a rut which can be characterized as 'faceless, tireless scientists who advance despite obstacles.' In these ruts, the story bogs down and is not particularly interesting.

Fortunately, the story breaks out of these ruts with anecdotes about the difficulty of seemingly trivial aspects of the research process. My favorite anecdote involves the measurement of the accuracy with which ALVINN maintains its position in the middle of a road. To make this measurement, Pomerleau attached a tank of water and a siphon hose to the vehicle. As ALVINN drove, it left a trail of water to mark where it had been. Pomerleau then measured the deviation of the trail from a preset baseline. Phrased in this way, the story indicates the actual difficulty of a seemingly simple task; otherwise, the story is quite pedestrian. To me this anecdote evoked an image of Pomerleau as a hunter trying to follow his quarry based on a trail of its spoor, or of ALVINN as a modern-day Hansel and Gretel, trying to remember its way home by dropping high-tech breadcrumbs.

As a chronicle, the book is also one man's story of the research process. It clearly points out that the process involves solving both mundane and involute problems, and, furthermore, that problems which appear mundane are often involute. For example, converting a high resolution video image to a low resolution image suitable for training ALVINN is much harder than I imagined. However, it is probably easier than finding a low resolution video camera.

Finally, because it is a chronicle, the book challenges the reader to solve problems along with the author. It is a tribute both to Pomerleau's writing and his scientific ability that I rarely found myself wondering why Pomerleau's solutions were different (or more complex) than mine. More often than not, Pomerleau's solutions are less complex – and actually work – whereas I have doubts about many of my own 'Rube Goldberg' solutions.

3. As an engineering achievement

The ultimate criteria for evaluating a piece of engineering is "does it work?" To this question, ALVINN receives an unequivocal "yes." However, the mere fact that ALVINN works does not make it interesting. After all, ALVINN is built from fairly well understood components: a neural network with one layer of hidden units makes decisions and a controller implements those decisions. What makes ALVINN interesting is that these two systems are put together to form an effective mechanism.

ALVINN can drive forward or backward on multiple road types using multiple type of sensors. This is accomplished by training multiple networks to become experts for each different task, where a task is defined by the triple [direction, road type, sensor type]. By itself, this is an achievement. An equal achievement is the development of techniques for integrating these expert networks that make it possible for ALVINN to drive without continually needing retraining (or at least a pause for manual selection of the appropriate expert). A pleasant benefit of this work is that the integrated system is able to recognize situations in which it cannot reliably control the vehicle. As a result, ALVINN is unlikely to confidently drive off of a cliff.

I found the sections detailing the efforts to integrate multiple expert networks less compelling than they could have been because they lack a convincing evaluation. These sections are certainly important; the idea of using an autoencoding network to gauge network reliability is an excellent – and possibly generally applicable – idea. However, the lack of sufficient supporting evidence makes it is impossible to be sure which of the three proposed methods is the most effective, although the chronological nature of the book invites the supposition that the final method is also the most effective. For example, simple experimental tests could have been run to answer the following questions: 1) how likely is each method to choose the wrong network, 2) how long does each method require to recognize and respond appropriately to transitions between road types, and 3) how likely is each method to correctly recognize that the road type is unknown? Had these questions been addressed experimentally, I would have felt confident that this work represents the principle contribution of the book.

To be fair, additional tests would have required a lot more driving around with ALVINN. Still, questions such as these deserve to be answered. As a more general criticism, thorough testing and evaluation of the ideas presented in the book are lacking. For instance, the table on page 30 shows that the Gaussian output representation is best able to track a road. Unfortunately, the table is presented without adequate context, so I have no feeling for minimum acceptable standards. It could be that all of the output encoding methods are as good or better than the average driver.

In addition to missing tests, the testing methodology seems somewhat sloppy in some instances. For example, Pomerleau asserts that trained networks are insensitive to variations in lighting. The evidence offered to support this statement (on page 72) is that networks trained on one day worked acceptably on the next day. This is little more than anecdotal evidence: substantiating this claim would have required an experiment in which at least three networks were trained on different days at different times (e.g., morning, midday, and afternoon). Those three networks would then be tested at least three times a day for several days of varying climatic character. One way to do this testing (without reattaching the water tank and siphon hose) would be to use a panel of judges who subjectively rate ALVINN's driving. Such a panel might rate each of four passes over the same section of road; one for each network and one using a human driver. Admittedly, this test requires considerably more time and effort than the one performed. Given that the point is fairly trivial, the lack of a thorough test is certainly forgivable. Nonetheless, more aspects of ALVINN deserve to be thoroughly tested. Missing tests and

answers aside, I found the description of the engineering aspects of the book consistently well written, well described, and interesting.

4. As a research contribution

At the start of the book, Pomerleau makes five specific claims outlining the work's research contribution; because those claims cover much of the space of the research described in the book, I base my evaluation solely in terms of these claims. From my perspective, four of the five claims are well supported. I am less certain about the fifth claim although I have no doubt that others will find my objections less than compelling.

The first claim is that artificial neural networks can quickly learn mobile robot navigation. Other than quibbles about "quickly" this claim is almost undeniable. My only reservation is that in some cases, it is not so much the network which learns, but a combination of Pomerleau and the network. For instance, Pomerleau developed the method for normalizing color images to grey scale. I would have been more comfortable if the network had learned its own normalization scheme because using the scheme developed by Pomerleau requires human intervention when changing sensors. Pomerleau claims that the network is capable of learning normalization, but that it takes more time. To me, this would have been time well spent.

The second claim is that the system robustly handles a wide variety of situations. Again, this claim is well supported. Not only is ALVINN able to drive on multiple types of roads, with equal facility it can drive backwards using a different set of eyes (i.e., replacing the video camera with a laser range finder.) In addition, the same techniques have been applied to a walking robot. Although the chapter on the walking robot was poorly linked to the rest of the book, it did support the flexibility and robustness claims made for the system.

The third claim regards the development of techniques for the analysis of the internal representations of neural networks. The methods developed here represent an impressive amount of work. However, they may not generalize outside of fields which, like vision, have a two-dimensional structure in their inputs. For instance, neural networks have been widely applied in molecular biology, particularly to the analysis of DNA sequences. It is not clear that the methods Pomerleau describes have any application to this domain. Still, given a problem with the appropriate structure, the techniques in this book are well described and informative.

The fourth claim is that it is possible to accurately assess the reliability of neural networks. Claims of graceful degradation have been made for years with respect to neural networks. This is among the first pieces of work to back up these claims with a method of estimating when the degradation has proceeded to the point that answers are unreliable. Aside from my earlier desire for additional quantitative testing, I found this claim more than adequately supported. Moreover, the methods for measuring degradation should be directly applicable to a wide variety of domains.

The final claim is that the a combination of neural networks with symbolic reasoning can yield both accurate control and apparently intelligent behavior. "Intelligent behavior" is defined as the ability to navigate from point A to point B where the two points are not

connected by a single, intersectionless road. Given this definition, the system is adequate to satisfy this claim. However, the level of detail required of the map is daunting. The map must contain not only distances and intersections, but also information about appropriate speed, intersection control, and so on. Requiring such detailed information makes it hard for me to consider the behavior intelligent. Furthermore, the behavior may give the appearance of intelligence, but only in situations considered by the mapmakers. For instance, appropriate speed in a school zone depends on whether school is in session. If school is in session, then driving quickly is clearly unintelligent. Conversely, if school is not in session, then driving slowly is quite stupid. Thus, I found support for this claim weak. The appearance of intelligence is an illusion easily dispelled.

5. Conclusions

Pomerleau succeeds where others have failed. His system can reliably drive without human intervention in a plethora of situations. Equally impressive (at least to me) is his book about that system. *Neural Network Perception for Mobile Robot Guidance* tells a story, describes an engineering enterprise, and details an impressive research achievement all at the same time. Aside from some relatively minor complaints about testing methodology, everything is done well. As a result, I confidently recommend this book to people with no background in neural networks; the concepts are sufficiently well described that this background is unnecessary. While background is unnecessary, I also recommend this book to people with a thorough background in neural networks. The ideas and their application are sufficiently novel that even the most experienced researcher is likely to find something useful.

References

Kidder, T. (1981). *The Soul of a New Machine*, Boston, MA: Little, Brown.