

More approaches to the travelling salesman guide

SIR—In David G. Bounds's interesting exposition¹ of several intriguing new approaches to combinatorial optimization, he is unfortunately and seriously misleading when he suggests that these new approaches "seem likely to lead to much better algorithms for finding acceptably low-cost minima" for the travelling salesman problem (TSP). The Lin–Kernighan algorithm², well-known to researchers in computer science and operations research but apparently not to Bounds or many of the authors of his references, substantially outperforms all these approaches, both in speed and the quality of solution found.

The Lin–Kernighan algorithm is a sophisticated successor to the naive 3-OPT algorithm mentioned by Bounds. It finds better tours, and is orders of magnitude faster. It has been widely implemented, and is well-known to be practical for instances substantially larger than those discussed by Bounds. For instances consisting of points in the plane under the euclidean metric, our own implementation of Lin–Kernighan finds tours that are within 2% of optimal for instances with as many as 50,000 cities. For 10,000 cities the running time is little more than an hour on a VAX 8550. For 100 cities, it is only 3 s.

The new approaches discussed by Bounds offer Lin–Kernighan little competition. Consider simulated annealing, and the type of instance studied by Kirkpatrick and Toulouse³ and cited by Bounds. For 400-city instances of this type, Kirkpatrick and Toulouse found tours of length about 2.4. Lin–Kernighan, on the other hand, finds tours with length 2 (and for each instance comes within 2–3% of the Held and Karp⁴ lower bound on optimal tour length²). Moreover, whereas Lin–Kernighan takes about 20 s to find these tours, we find simulated annealing requires more than 6 h.

In addition to simulated annealing, Bounds covers two biologically motivated approaches: the neural-net approach of Hopfield and Tank⁵ and the genetic approach of Brady⁶, and mentions a fourth approach based on an analogy with rubber bands⁷.

The biggest instance Hopfield and Tank considered was 30 cities. I provided them with this instance, and when I did so I also told them the length of the optimal tour, which can be generated in less than a second using Lin–Kernighan, and whose optimality can be verified in 3 s. In all their many neural-net simulations on this instance; Hopfield and Tank never came within 17% of the optimal tour length⁵.

Brady also worked on only a tiny instance (64 cities, where the best of 10 runs of Lin–Kernighan is likely to be optimal), and reports that taking the best of a

collection of random-start 2-OPT runs was essentially as good as any of his techniques involving mating⁶. Multiple-start 2-OPT is a well-known procedure, and is not competitive with Lin–Kernighan.

The best of these last three approaches appears to be that of Durbin and Willshaw⁷, who found an optimal solution to the 30-city instance mentioned above. For instances with $n = 100$, however, they apparently performed no better than 3-OPT, worse than Lin–Kernighan.

Note that for instances as small as those on which the last three approaches were tested, there now exist programs that routinely find verified optimal solutions in reasonable amounts of time. The best such optimization algorithm is that of Padberg and Rinaldi⁸, which regularly finds verified optima for 100-city instances in what would be about 5 min on a VAX 8550. (The largest TSP they have solved

was a real-world instance with 2,392 cities, requiring 27 h on a supercomputer.)

The TSP is an intriguing problem, and those of us who have been working on it for years welcome its new-found popularity and the consequent influx of new ideas and approaches. I hope that in future, however, the evaluation of these new approaches will more fully take into account the strong competition provided by the techniques already in use.

DAVID JOHNSON

AT&T Bell Laboratories,
600 Mountain Avenue,
Murray Hill, New Jersey 07984-2070, USA

1. Bounds, D.G. *Nature* **239**, 215–219 (1987).
2. Lin, S. & Kernighan, B.W. *Oper. Res.* **21**, 498–516 (1973).
3. Kirkpatrick, S. & Toulouse, G. *J. Phys., Paris* **46**, 1277–1292 (1985).
4. Held, R.M. & Karp, R.M. *Oper. Res.* **18**, 1138–1162 (1970).
5. Hopfield, J.J. & Tank, D.W. *Biol. Cybern.* **52**, 141–152 (1985).
6. Brady, R.M. *Nature* **317**, 804–806 (1985).
7. Durbin, R. & Willshaw, D. *Nature* **336**, 689–691 (1987).
8. Padberg, M. & Rinaldi, G. *Oper. Res. Lett.* **6**, 1–8 (1987).

Returns on boomerangs

SIR—When a boomerang is thrown, the far end has a positive velocity relative to the chord but the near end has a negative velocity and therefore a negative Reynolds number. If the blade is twisted so that when thrown right-handed it has wash-in or positive incidence at the far end it will return to the thrower. If it has wash-out or is thrown left-handed the aerodynamic forces will tend to balance out because the nearward travelling blade will have positive incidence. Under these circumstances, if thrown in the horizontal plane the boomerang will travel on a curve having a large radius, approximating to a straight line.

The determination of whether the boomerang described by Valde-Nowak, Nadachowski and Wolsan (*Nature* **329**, 436–438; 1987) would return could easily be decided by measurement of the incidence variation across the span. Determining whether the palaeolithic thrower was left- or right-handed would be more difficult.

BRENNIG JAMES

Cherry Orchard,
Marlow Common,
Buckinghamshire, UK

The burden of proof in linking AIDS to scrapie

SIR—We have questioned^{1,2} the significance of similarities noted by Haseltine and Patarca³ between the scrapie-associated protein, PrP 27–30, and the human immunodeficiency virus, HIV-1. Patarca *et al.*⁴ defend the significance of this sequence resemblance. This supposed relationship could affect directions of AIDS research; also, it calls into question the value of sequence similarity measures as evidence for ancestral or functional

relationships between molecules. We must therefore point out weaknesses in the arguments of Haseltine and colleagues.

The central question in sequence comparison is biological relevance, not statistical significance. Statistical significance is typically used as a clue to biological relevance, which is more difficult to assess. Patarca *et al.*⁴ maintain that probability values of 10^{-3} and 10^{-1} for the HIV–PrP comparison are significant. However, there are about 12,500 entries in the GenBank(R) sequence database. Thus, for the HIV–PrP alignments with a probability value of 10^{-3} , we can expect to find 12–13 unrelated sequences with similar comparison scores in a search of the database. Actual searches by one of us (J.F.B.) yielded 7 and 13 equally good or better alignments for HIV and PrP, respectively. Yet, it is unlikely that any of the aligned sequences are truly related to HIV or PrP; they come from genes as diverse as chloroplast ribosomal RNA and rat cytochrome P-450c.

Patarca *et al.*⁴ assert that Bazan *et al.*² used statistical scores from two algorithms that were not comparable. In fact, the program FASTN was only used to search GenBank for loci with sequence similarity to HIV or PrP. The statistical significance of each similarity was then evaluated using ALIGN.

PrP shows no similarity to visna virus, equine infectious anaemia virus or HIV-2 (ALIGN scores of 0.2 s.d., –2.1 s.d. and 0.5 s.d.), yet these viruses are related to HIV-1 (refs 5, 6). Patarca *et al.*⁴ dismiss this evolutionary comparison on the grounds that the viruses may have undergone different rates of sequence divergence which could obscure their relationship to PrP. We note that Haseltine and Patarca³ were the first to invoke an evolutionary argument when they used a visna virus–PrP alignment to corroborate their