

Probabilistic Analysis of Tree Search

C.J.H. McDiarmid

Abstract

Consider the family tree of an age-dependent branching process, where the branches have costs corresponding to birth times. The first-birth problem of Hammersley (1974) then concerns the cost of an optimal (cheapest) node at depth n . Suppose that we must explore the tree so as to *find* an optimal or nearly optimal node at depth n . We now have a suitable model for analysing the behaviour of tree search algorithms, and we may extend the investigations of Karp and Pearl (1983).

1. Introduction

Many algorithms considered in operations research, computer science and artificial intelligence may be represented as a search or partial search through a rooted tree. Such algorithms typically involve backtracking but try to minimise the time spent doing so. This paper extends work of Karp and Pearl (1983), and gives a probabilistic analysis of backtracking and non-backtracking search algorithms in certain random trees. We thus cast some light on the question of when to backtrack: it seems that backtracking is valuable just for problems with ‘dead-ends’.

Let us review briefly the model and results of Karp and Pearl. They consider an infinite rooted tree in which each node has exactly two sons. The branches have independent 0, 1-valued random costs X , with $p = P(X = 0)$. (We have swapped p and $1 - p$ from the original paper.) The problem is to find an optimal (cheapest) or nearly optimal path from the root to a node at depth n .

The problem changes nature depending on whether the expected number $m_0 = 2p$ of zero-cost branches leaving a node is greater than 1, equal to 1 or less than 1 (as was suggested in Hammersley 1974, Note 8). When $m_0 > 1$ a simple ‘uniform cost’ breadth-first search algorithm A1 finds an optimal solution in expected time $O(n)$; and when $m_0 = 1$ this algorithm takes expected time $O(n^2)$. When $m_0 < 1$ any algorithm that is guaranteed to find a solution within a constant factor of optimal must take exponential expected time. However, in this case a ‘bounded-lookahead plus partial backtrack’ algorithm A2 usually finds a solution close to optimal in linear

expected time. This successful performance of the backtracking algorithm A2 for the difficult case when $m_0 < 1$ was taken to suggest that similar heuristics should be of general use for attacking NP-hard problems.

We shall see that with the above search model, a simple non-backtracking bounded-lookahead algorithm A3 performs as successfully as the backtracking algorithm A2. Thus it seems hard to recommend the use of heuristics like A2 on the basis of this search model. Similar comments hold if we allow more general finite random costs on the branches.

However, there is a qualitative difference if we allow nodes to have no sons (or allow branches to have infinite costs) so that there are 'dead-ends'. We extend Karp and Pearl's work by considering search in random trees generated by an age-dependent branching process, in which the mean number of children of an individual is greater than one. The investigation is related to the first-birth (or death) problem, as introduced by J.M. Hammersley (1974) (see also Joffre et al. 1973). This model is discussed further below. Let p_0 be the probability that a node has no sons, and let m_0 be the mean number of zero-cost branches leaving a node (instantaneous births).

Our results concerning algorithms A1 and A2 are natural extensions of Karp and Pearl's results. Thus the breadth-first search algorithm A1 finds an optimal solution in linear expected time if $m_0 > 1$ and in quadratic expected time if $m_0 = 1$. If $m_0 < 1$ then any algorithm with a constant performance guarantee must take exponential expected time, but the backtracking algorithm A2 finds a nearly optimal solution in linear expected time.

However, the performance of the simple non-backtracking bounded lookahead algorithm A3 depends critically on whether $p_0 = 0$ or $p_0 > 0$. Suppose that $m_0 < 1$, so that optimal search is hard. If $p_0 = 0$, so that as in the Karp and Pearl model there are no dead-ends, then algorithm A3 usually finds a nearly optimal solution in linear expected time; that is, it performs as successfully as the backtracking algorithm A2. However, if $p_0 > 0$ then algorithm A3 usually fails to find a solution. Thus our model suggests that backtracking becomes attractive when there is the possibility of dead-ends.

In the next section we give details concerning the search model and the algorithms A1, A2 and A3, then in Section 3 we present our results, and finally Section 4 contains proofs.

2. Model and Algorithms

We suppose that the tree to be searched is the family tree F of an age-dependent branching process of Crump-Mode type (see Crump and Mode 1968). In this model an initial ancestor is born at time $t = 0$ and then produces children at random throughout his lifetime. If $Z_1(t)$ denotes the num-

ber of children born up to time t , then $Z_1(t)$ is an arbitrary counting process, that is a non-negative integer-valued non-decreasing right-continuous random process. We do not insist that $Z_1(0) = 0$. The children of the ancestor form the first generation: from their several birth times they behave like independent copies of their parent. Their children form the second generation, and so on. We let $Z_n(t)$ be the number of individuals born in the n th generation by time t , and let $Z_n(\infty) = \sup_t Z_n(t)$.

We shall always assume that the mean number $m = E[Z_1(\infty)]$ of children of an individual satisfies $m > 1$, so that the extinction probability q satisfies $q < 1$. Let $p_k = P[Z_1(\infty) = k]$. Clearly $q > 0$ if and only if $p_0 > 0$, and these conditions correspond to the existence of ‘dead-ends’.

When searching the family tree F we take the cost of a branch to be the difference between the birth times of the child and the parent, and so we take the cost of a node to be the birth time of the corresponding individual. Thus we seek a first born individual in generation n . We shall denote the corresponding optimal cost by C_n^* (rather than B_n): if $Z_n(\infty) = 0$ then we set $C_n^* = \infty$. Thus

$$P(C_n^* = \infty) = P(Z_n(\infty) = 0) = q_n \rightarrow q \text{ as } n \rightarrow \infty.$$

The interesting case is when the tree to be searched is infinite: we shall often condition on the event S of ultimate survival, and then C_n^* is finite for all n .

Assumptions. Recall that we assume that the mean family size m satisfies $m > 1$: this is essential. For convenience we shall also assume that m is finite and that lifetimes (branch costs) are bounded. We are thus able to show that certain events of interest fail with exponentially small probabilities. (Truncation arguments as in Kingman (1975) may then be used to obtain ‘almost sure’ results under weaker assumptions.) Further, a simple translation allows us to assume that small costs can occur, that is $E[Z_1(\delta)] > 0$ for $\delta > 0$.

The distinction between zero and non-zero costs turns out to be important. Let $m_0 = E[Z_1(0)]$ be the expected number of zero-cost branches from a node (instantaneous births to an individual).

We shall discuss the performance of three algorithms, A1, A2 and A3, the first two of which are taken from Karp and Pearl (1983). Each algorithm maintains a subtree T of the family tree F containing the root; and at each step explores some node of T . Here, exploring a node x means appending to x the next (leftmost) child y of x in F but not yet in T , and observing the cost of the corresponding branch xy ; or observing that node x has no more children.

Algorithm A1 is a ‘uniform cost’ breadth-first search algorithm and will be analysed for the cases $m_0 > 1$ and $m_0 = 1$, when there are many zero-cost branches and search is easy. Algorithm A2 is a hybrid of local

and global depth-first search strategies and will be analysed for $m_0 < 1$. Algorithm A3 consists of repeated local optimal searches, and will be analysed also for $m_0 < 1$. Note that A1 is an exact algorithm, whereas A2 and A3 are approximation algorithms or heuristics.

For each algorithm A_j , we let the random cost of the solution found be $C_n^{A_j}$ ($= \infty$ if no solution is found), and the random time taken be $T_n^{A_j}$. We measure time by the number of nodes of the search tree encountered. The three algorithms are as follows.

Algorithm A1: At each step, explore the leftmost node among those active nodes of minimum cost. Here, a node is active if it is in T and may perhaps have further children. The algorithm halts when it would next explore a node at depth n . That node then corresponds to an optimal solution.

Algorithm A2: This algorithm conducts a staged search with backtracking if a local test is failed. It has three parameters: d , L , and α . By an (α, L) -regular path we mean a path which consists of segments each of length L and cost at most αL (except that the last segment may have length less than L). The algorithm A2 conducts a depth-first search to find an (α, L) -regular path from a depth d node to a depth n node. If it succeeds in reaching depth n , the algorithm returns the corresponding path as a solution: if it fails, the search is repeated from another depth d node. If all the nodes at depth d fail to root an (α, L) -regular path to a depth n node, the algorithm terminates with failure.

Algorithm A3: This simple bounded-lookahead or ‘horizon’ heuristic is a staged-search algorithm which avoids backtracking. It has one parameter L . Starting at the root it finds an optimal path to a node at depth L , makes that node the new starting point and repeats.

3. Results

We summarise our results in six theorems. Theorem 3.1 concerns the region where the mean number m_0 of zero-cost branches leaving a node satisfies $m_0 > 1$, Theorem 3.2 concerns $m_0 = 1$ and Theorems 3.3–3.6 concern $m_0 < 1$. When $m_0 \geq 1$ the main distinction is between zero and non-zero costs. Recall that S denotes the event of ultimate survival.

THEOREM 3.1. *Let $m_0 > 1$.*

- (a) *The random variable $C^* = \lim C_n^*$ is finite almost surely on S , and indeed there exists $\delta < 1$ such that*

$$P(C^* \geq k \mid S) = O(\delta^k) \text{ as } k \rightarrow \infty.$$

- (b) *The time T_n^{A1} taken by algorithm A1 satisfies $E[T_n^{A1}] = O(n)$.*

Thus, if the family tree is infinite, the optimal cost C_n^* remains bounded as $n \rightarrow \infty$, and algorithm A1 finds an optimal path in linear expected time.

Next we consider the critical case $m_0 = 1$. It is convenient here to restrict attention to a Bellman-Harris age-dependent branching process (see for example Harris 1963). Now children are produced according to a simple Galton-Watson branching process, and branch costs are independent and each distributed like some non-negative random variable X .

THEOREM 3.2. *Consider a Bellman-Harris process with $m_0 = mP(X = 0) = 1$ and $E[Z_1(\infty)^2] < \infty$.*

(a) *If further $E[Z_1(\infty)^{2+\delta}] < \infty$ for some $\delta > 0$, $P(0 < X < 1) = 0$ and $P(X = 1) > 0$, then*

$$C_n^*/\log \log n \rightarrow 1 \text{ almost surely on } S \text{ as } n \rightarrow \infty.$$

(b) *The time T_n^{A1} taken by algorithm A1 satisfies $E[T_n^{A1}] = O(n^2)$.*

Part (a) shows roughly that if the optimal cost is finite then it is usually close to $\log \log n$: it is a special case of a result of Bramson (1978). Part (b) states that the algorithm A1 finds an optimal path in quadratic expected time.

Our main interest is in the case $m_0 < 1$. The first result for this case shows that we cannot quickly find guaranteed optimal or near optimal solutions, and so it is of interest to analyse heuristic approximation methods. The next result concerns the optimal cost C_n^* and then we consider the algorithms A2 and A3.

THEOREM 3.3. *Assume that $m_0 < 1$. Let T_n be the least number of nodes explored in any proof that guarantees a certain path of length n to be within a constant factor β of optimal. Then there exists $\eta > 1$ and $\delta < 1$ such that*

$$P(T_n < \eta^n \mid S) = O(\delta^n) \text{ as } n \rightarrow \infty.$$

THEOREM 3.4. *There is a constant $\gamma \geq 0$, defined by equation (4.1) below and satisfying $\gamma > 0$ if and only if $m_0 < 1$, such that for any $\epsilon > 0$ there exists $\delta < 1$ with*

$$P\left(\left|\frac{1}{n}C_n^* - \gamma\right| > \epsilon \mid S\right) = O(\delta^n) \text{ as } n \rightarrow \infty.$$

This result shows roughly that if the optimal cost C_n^* is finite then it is usually close to γn . It is essentially due to Hammersley (1974) and Kingman (1975), see also Kesten (1973), Kingman (1976). We shall find that it follows quite easily from our analysis of the search algorithm A2; see also Biggins (1979).

THEOREM 3.5. *Let $m_0 < 1$, and consider the backtracking algorithm A2. For any $\epsilon > 0$, with appropriate parameters the algorithm runs in linear expected time, and there exists $\delta < 1$ such that*

$$P(C_n^{A2} \leq (1 + \epsilon)C_n^*) = 1 - O(\delta^n) \text{ as } n \rightarrow \infty.$$

THEOREM 3.6. *Let $m_0 < 1$, and consider the non-backtracking algorithm A3.*

(a) *If $p_0 = 0$ then for any $\epsilon > 0$, with appropriate constant lookahead the algorithm runs in linear expected time, and there exists $\delta < 1$ such that*

$$P(C_n^{A3} \leq (1 + \epsilon)C_n^*) = 1 - O(\delta^n) \text{ as } n \rightarrow \infty.$$

(b) *If $p_0 > 0$, then for any constant lookahead there exists $\delta < 1$ such that*

$$P(C_n^{A3} < \infty) = O(\delta^n) \text{ as } n \rightarrow \infty.$$

We thus see that the backtracking algorithm A2 is a good heuristic, and so is the non-backtracking algorithm A3 as long as $p_0 = 0$.

Hammersley (1974, Note 8) asked about the concentration of the random variable C_n^* , in particular in the special case considered by Karp and Pearl (1983) when each individual has exactly two children, both born at time 0 or 1. For this case the bounded differences inequality of Hoeffding (1963), Azuma (1967) in the form given in McDiarmid (1989) shows immediately that for any $t \geq 0$,

$$P(|C_n^* - E(C_n^*)| \geq t) \leq 2e^{-2t^2/n}.$$

4. Proofs

The first lemma below immediately gives part (a) of Theorem 3.1.

LEMMA 4.1. *Suppose that $m_0 > 1$. Let T be the least depth at which an infinite path of zero-cost branches is rooted, where we let $T = \infty$ if there is no such path. Then there exists $\delta < 1$ such that $P(T > n | S) = O(\delta^n)$ as $n \rightarrow \infty$.*

PROOF: The zero-cost branches yield a branching process \tilde{Z} with mean $m_0 > 1$ and thus with extinction probability $\tilde{q} < 1$. Suppose that $p_0 + p_1 > 0$. Then $\alpha = f'(q)$ satisfies $0 < \alpha < 1$. So, by a minor extension of Theorem 8.4 in Chapter I of Harris (1963), $f_n(\tilde{q}) = q + O(\alpha^n)$. Here f_n is the generating function for $Z_n(\infty)$. But

$$P(T > n) = \sum_k P(Z_n(\infty) = k) \tilde{q}^k = f_n(\tilde{q}).$$

Hence

$$q + O(\alpha^n) = P(T > n) = (1 - q)P(T > n | S) + q,$$

and so $P(T > n | S) = O(\alpha^n)$, as required. The case not considered so far is when $p_0 + p_1 = 0$, but then clearly

$$P(T > n | S) = P(T > n) \leq \tilde{q}^{2^n}. \quad \blacksquare$$

Now consider a Galton-Watson branching process \tilde{Z} . Let D_n be the number of nodes encountered in a depth-first search of the family tree which terminates on reaching a node at depth n or on searching the complete tree, and let $d_n = E[D_n]$.

LEMMA 4.2. For each n , $d_n \leq n + 1$.

PROOF: Let $q_n = P(\tilde{Z}_n = 0)$. Of course $d_0 = 1$. Suppose that d_n is finite. Then by conditioning on \tilde{Z}_1 we see that

$$\begin{aligned} d_{n+1} &= 1 + \sum_{k \geq 1} p_k (1 + q_n + \dots + q_n^{k-1}) d_n \\ &= 1 + \frac{d_n}{1 - q_n} \sum_{k \geq 0} p_k (1 - q_n^k) \\ &= 1 + \frac{d_n}{1 - q_n} (1 - f(q_n)) \\ &\leq 1 + d_n \end{aligned}$$

since $f(q_n) = q_{n+1} \geq q_n$. \blacksquare

We may now prove part (b) of Theorem 3.1. Consider the branching process \tilde{Z} corresponding to the zero cost branches. It has mean $m_0 > 1$ and so it has extinction probability $\tilde{q} < 1$. It now follows from Lemma 4.2 (by Wald's equation) that

$$E[T_n^{A1}] \leq \frac{n + 1}{1 - \tilde{q}}.$$

This completes our proof of Theorem 3.1. \blacksquare

PROOF OF THEOREM 3.2: Part (a) has already been discussed, so consider part (b). Consider again the process \tilde{Z} corresponding to the zero-cost branches. This has mean $E[\tilde{Z}_1] = m_0 = 1$ and variance $\tilde{\sigma}^2 = \sigma^2 p^2 + mp(1 - p) < \infty$, where $p = P(X = 0)$. Hence

$$P(\tilde{Z}_n > 0) = \frac{2 + o(1)}{\tilde{\sigma}^2 n}$$

(see for example Athreya and Ney 1972, p. 19). So, arguing as before,

$$E[T_n^{A1}] \leq \frac{n+1}{P(\tilde{Z}_n > 0)} = O(n^2). \quad \blacksquare$$

To consider the case $m_0 < 1$ we must investigate the Crump-Mode model in more detail. The key to the analysis is the function $\phi(\theta)$ introduced by Kingman (1975). For $\theta \geq 0$, let

$$\phi(\theta) = E\left[\sum_r e^{-\theta B_{1r}}\right],$$

where the sum is over the birth times B_{1r} of the children r of the initial ancestor. Note that $\phi(0) = E[Z_1(\infty)] = m < \infty$, and so $\phi(\theta) < \infty$ for all $\theta \geq 0$. Next, for $a \geq 0$ let

$$\mu(a) = \inf\{e^{\theta a}\phi(\theta) : \theta \geq 0\},$$

and define the ‘time constant’ γ by

$$\gamma = \inf\{a \geq 0 : \mu(a) \geq 1\}. \quad (4.1)$$

The next two lemmas may be found (essentially) in Kingman (1975).

LEMMA 4.3. *The function μ on $[0, \infty)$ is continuous; $\mu(0) = m_0$; and for some $b \geq 0$, μ is strictly increasing on $[0, b]$ and $\mu(a) = m$ for each $a \geq b$.*

LEMMA 4.4. *For any $a \geq 0$,*

$$E[Z_n(an)] = (\mu(a) + o(1))^n \quad \text{as } n \rightarrow \infty.$$

PROOF OF THEOREM 3.5: Let $0 < \epsilon < \gamma$ and let $\alpha = \gamma + \epsilon$. By Lemmas 4.3 and 4.4 we may choose a constant L such that $E[Z_L(\alpha L)] > 1$. By considering the (α, L) -sons of a depth d node and their (α, L) -sons and so on we obtain a branching process \tilde{Z} say. This process has mean $\hat{m} > 1$ and thus has extinction probability $\hat{q} < 1$.

We can bound the expected running time of algorithm A2 as follows. By Lemma 4.2 (and Wald’s equation) for each node at depth d , the expected cost of a search to depth n from that node is at most

$$m^{L+1}(\lceil(n-d)/L\rceil + 1) \leq m^{L+1}(n+1).$$

Hence (by Wald’s equation again)

$$E[T_n^{A2}] \leq \frac{d + m^{L+1}(n+1)}{1 - \hat{q}} = O(n).$$

Next consider costs. Let λ be a bound on lifetimes or branch costs, and set $d = d(n) = \lfloor (\epsilon/\lambda)n \rfloor$. If the algorithm A2 succeeds then

$$C_n^{A2} \leq d\lambda + \lceil (n-d)/L \rceil (\alpha L) \leq (\alpha + \epsilon)n \text{ once } d \geq L.$$

But

$$P(C_n^{A2} = \infty) \leq \sum_k P(Z_d(\infty) = k) \hat{q}^k = f_d(\hat{q}).$$

Also, by Lemmas 4.3, 4.4,

$$P(C_n^{A2} \leq (\gamma - \epsilon)n) \leq E[Z_n((\gamma - \epsilon)n)] = O(\delta_1^n)$$

for some suitable $\delta_1 < 1$. Hence

$$\begin{aligned} P\left(C_n^{A2} > \frac{\gamma + 2\epsilon}{\gamma - \epsilon} C_n^*\right) &\leq P(C_n^* \leq (\gamma - \epsilon)n) + P(\{C_n^{A2} = \infty\} \setminus \{C_n^* = \infty\}) \\ &\leq O(\delta_1^n) + f_d(\hat{q}) - f_n(0) \\ &= O(\delta_2^n) \text{ for suitable } \delta_2 < 1. \quad \blacksquare \end{aligned}$$

We do not need to prove Theorem 3.4 here, since it is implicit in Kingman (1975), but note that we have actually done so above. We may adapt the proof idea above to yield a variant of the ‘Chernoff theorem’ of Biggins (1979), which we shall use to prove Theorem 3.3.

LEMMA 4.5. *If $1 < \mu(a)$ then for any $1 < \eta < \mu(a)$ there exists $\delta < 1$ such that*

$$P(Z_n(an) < \eta^n \mid S) = O(\delta^n) \text{ as } n \rightarrow \infty.$$

PROOF: By Lemma 4.3 we may choose $b < a$ such that $\mu(b) > \eta$. Then by Lemma 4.4 we may choose $\epsilon > 0$ sufficiently small and L sufficiently large that $\epsilon\lambda + b < a$ and $(\hat{m} - \epsilon)^{(1-2\epsilon)/L} > \eta$, where $\hat{m} = E[Z_L(bL)]$.

Consider the branching process \hat{Z} , with mean \hat{m} , formed by taking (b, L) -sons and their sons and so on. By a theorem of Seneta and Hyde (see for example Athreya and Ney 1972, Theorem 3, p. 30)

$$P(\hat{Z}_k < (\hat{m} - \epsilon)^k) \rightarrow \hat{q} \text{ as } k \rightarrow \infty.$$

By considering these processes rooted at the nodes at depth $d + i$ where $d = \lfloor \epsilon n \rfloor$ and $0 \leq i < L$, we see that

$$\begin{aligned} P\left(Z_{d+i+kL}((d+L)\lambda + kbL) < (\hat{m} - \epsilon)^k\right) &\leq f_{d+i}(\hat{q} + o(1)) \\ &= q + O(\delta_1^k) \end{aligned}$$

for suitable $\delta_1 < 1$. But, if $n = d + i + kL$ where $0 \leq i < L$, then $an \geq (d + L)\lambda + kbL$ and $(\hat{m} - \epsilon)^k \geq \eta^n$ for n sufficiently large. Hence

$$P(Z_n(an) < \eta^n) \leq q + O(\delta_2^n).$$

Thus

$$P(Z_n(an) < \eta^n \mid S)(1 - q) + q_n \leq q + O(\delta_2^n),$$

and the result follows, since $q_n = q + O(\delta_3^n)$ for some suitable $\delta_3 (< 1)$. ■

PROOF OF THEOREM 3.3: Let $0 < \epsilon < \gamma$, let $\beta' = \beta(\gamma + \epsilon)/(\gamma - \epsilon)$ and let $k = k(n) = \lfloor n/\beta' \rfloor$. The key observation is that if $C_n^* > (\gamma - \epsilon)n$ then

$$T_n \geq Z_k((\gamma - \epsilon)n/\beta) \geq Z_k((\gamma + \epsilon)k);$$

for each node counted by $Z_k((\gamma - \epsilon)n/\beta)$ has cost less than C_n^*/β and so must be explored. Now let $1 < \eta < \mu(\gamma + \epsilon)$. Then

$$\begin{aligned} P(T_n < \eta^k \mid S) &\leq P(C_n^* \leq (\gamma - \epsilon)n \mid S) + P(Z_k((\gamma + \epsilon)k) < \eta^k \mid S) \\ &= O(\delta^k) \end{aligned}$$

for suitable $\delta < 1$, by Lemmas 4.4 and 4.5. ■

PROOF OF THEOREM 3.6:

(a) Let $p_0 = 0$. By Theorem 3.4 we may choose L so that $E[C_L^*/L] < (1 + \epsilon)\gamma$. Now C_n^{A3} is bounded above by $\lceil n/L \rceil$ independent copies of C_L^*/L , and we are done.

(b) Observe that

$$P(C_n^{A3} = \infty) \geq 1 - (1 - p_0)^{n/L}. \quad \blacksquare$$

Acknowledgement

I would like to acknowledge stimulating conversations on tree search with Greg Provan.

REFERENCES

- Athreya, K.B. and Ney, P.E. (1972). *Branching Processes*. Springer-Verlag, Berlin.
- Azuma, K. (1967). Weighted sums of certain dependent random variables. *Tôkoku Mathematical Journal* 19, 357–367.
- Biggins, J.D. (1976). The first- and last-birth problems for a multitype age-dependent branching process. *Advances in Applied Probability* 8, 446–459.
- (1977). Martingale convergence in the branching random walk. *Journal of Applied Probability* 14, 25–37.
- (1979). Chernoff's theorem in the branching random walk. *Journal of Applied Probability* 14, 630–636.
- Bramson, M.D. (1978). Minimal displacement of branching random walk. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 45, 89–108.
- Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics* 23, 493–507.
- Crump, K.S. and Mode, C.J. (1968). A general age-dependent branching process I. *Journal of Mathematical Analysis and Applications* 24, 494–508.
- Grimmett, G.R. (1985). Large deviations in subadditive processes and first-passage percolation. In *Particle Systems, Random Media, and Large Deviations*, ed. R. Durrett, Contemporary Mathematics 41, 175–194.
- Hammersley, J.M. (1974). Postulates for subadditive processes. *Annals of Probability* 2, 652–680.
- Harris, T.E. (1963). *The Theory of Branching Processes*. Springer-Verlag, Berlin.
- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association* 58, 13–30.
- Joffre, A., LeCam, L. and Neveu, J. (1973). Sur la loi des grands nombres pour des variables aléatoires de Bernoulli attachées à un arbre dyadique. *Compte Rendus de l'Académie des Sciences, Paris A* 277, 963–964.
- Karp, R.M. (1976). The probabilistic analysis of some combinatorial search algorithms. In *Algorithms and Complexity*, ed. J.F. Traub, Academic Press, New York, 1–19.
- Karp, R.M. and Pearl J. (1983). Searching for an optimal path in a tree with random costs. *Artificial Intelligence* 21, 99–116.
- Kesten, H. (1973). Contribution to the discussion of Kingman (1973).
- Kingman, J.F.C. (1973). Subadditive ergodic theory. *Annals of Probability* 1, 883–909.
- (1975). The first birth problem for an age-dependent branching process. *Annals of Probability* 3, 790–801.
- (1976). Subadditive processes. In *Ecole d'Eté de Probabilités de Saint Flour V-1975*, ed. P.-L. Hennequin, Lecture Notes in Mathematics no. 539, Springer-Verlag, Berlin.
- McDiarmid, C.J.H. (1989). On the method of bounded differences. Manuscript.

Department of Statistics
Oxford University
Oxford OX1 3TG.