Using Directed Acyclic Graphs in Epidemiological Research in Psychosis: An Analysis of the Role of Bullying in Psychosis

Giusi Moffa[1,2,11], Gennaro Catone[3,4,11], Jack Kuipers[5], Elizabeth Kuipers[6,7], Daniel Freeman[8], Steven Marwaha[9], Belinda R. Lennox[8], Matthew R. Broome[8,10], and Paul Bebbington[*,1]

[1]Division of Psychiatry, University College London, London, UK; [2]Institute for Clinical Epidemiology and Biostatistics, University Hospital Basel and University of Basel, Basel, Switzerland; [3]Department of Mental and Physical Health and Preventive Medicine, University of Naples SUN, Naples, Italy; [4]Faculty of Educational Science, Suor Orsola Benicasa University, Naples, Italy; [5]D-BSSE, ETH Zurich, Basel, Switzerland; [6]Department of Psychology, Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, UK; [7]Biomedical Research Centre, South London and Maudsley NHS Foundation Trust, Beckenham, UK; [8]Department of Psychiatry, Warneford Hospital, University of Oxford, Oxford, UK; [9]Division of Mental Health and Wellbeing, Warwick Medical School, University of Warwick, Coventry, UK; [10]Warneford Hospital, Oxford Health NHS Foundation Trust, Oxford, UK

[11]These authors contributed equally to the article.

*To whom correspondence should be addressed; Division of Psychiatry, University College London, 67-73 Riding House St, London W1T 7NF, UK; tel: +44-(0)20-7679-9465, e-mail: p.bebbington@ucl.ac.uk

Modern psychiatric epidemiology researches complex interactions between multiple variables in large datasets. This creates difficulties for causal inference. We argue for the use of probabilistic models represented by directed acyclic graphs (DAGs). These capture the dependence structure of multiple variables and, used appropriately, allow more robust conclusions about the direction of causation. We analyzed British national survey data to assess putative mediators of the association between bullying victimization and persecutory ideation. We compared results using DAGs and the Karlson–Holm–Breen (KHB) logistic regression commands in STATA. We analyzed data from the 2007 English National Survey of Psychiatric Morbidity, using the equivalent 2000 survey in an instant replication. Additional details of methods and results are provided in the supplementary material. DAG analysis revealed a richer structure of relationships than could be inferred using the KHB logistic regression commands. Thus, bullying had direct effects on worry, persecutory ideation, mood instability, and drug use. Depression, sleep and anxiety lay downstream, and therefore did not mediate the link between bullying and persecutory ideation. Mediation by worry and mood instability could not be definitively ascertained. Bullying led to hallucinations indirectly, via persecutory ideation and depression. DAG analysis of the 2000 dataset suggested the technique generates stable results. While causality cannot be fully determined from cross-sectional data, DAGs indicate the relationships providing the best fit. They thereby advance investigation of the complex interactions seen in psychiatry, including the mechanisms underpinning psychiatric symptoms. It may consequently be used to optimize the choice of intervention targets.

Key words: probabilistic graphical models/directed acyclic graphs/mediation/bullying/persecutory ideation/psychosis/worry/depression/anxiety

Introduction

In recent years, considerable advances have been made in our understanding of psychotic disorders. This has come about partly because of a productive divergence in research strategies: the disease model has been complemented by the emergence of an *interactional model*, which treats both the symptoms associated with psychosis and their hypothesized environmental antecedents as a causal system of individual interacting variables.[1–8] This second model has been applied in predominantly psychosocial contexts. Particular adverse contexts increase the likelihood of a range of psychological symptoms, including psychotic experiences. These symptoms may then have causal effects on each other without reference to an underlying condition. In particular, specific nonpsychotic symptoms modulate the relationship between adversity and psychotic experiences as moderators and mediators. The interactional model thus has the advantage of identifying candidate targets for psychological treatments.

This article concerns the interactional model and its reliance on statistical analysis to disentangle potentially

complex causal chains. We argue that such causal links can be more robustly inferred by the introduction of probabilistic graphical models based on directed acyclic graphs (DAGs). We know of only 2 prior examples of their use in a psychological context.[3,4] In contrast to these authors, we follow a Bayesian approach to learn the DAG and estimate putative causal effects in a consistent fashion.

The gold standard for causal inference in a putative causal system remains direct intervention. However, even where feasible, intervention studies are expensive. It is therefore sensible to prioritize interventions in relation to their likely effectiveness. Accordingly, attempts have been made to optimize causal inference in datasets that do not involve interventions. While causality cannot be determined with certainty from observational data, we can nevertheless ascertain which causal patterns fit the data best.

Much of the adversity literature in psychosis involves testing for putative mediational effects. The concept of mediation postulates that the causal effect between 2 variables is transmitted by a third.[9] So conceived, causation has 2 possible components: one operating directly, and one operating indirectly through the purported intermediary variable.[10] These direct and indirect effects can be quantified, enabling statements about their relative importance. Several recent studies have assessed mediation using logistic regression (in particular through an increasing reliance on the Karlson–Holm–Breen [KHB] commands and their congeners in STATA[11–13]). However, as extensions of regression, these commands cannot of themselves underwrite causal direction. They do give a provisional idea of the plausibility of candidate mediators (if a given variable does not drive an indirect effect, the hypothesis of mediation is not sustainable). However, any further causal inference has to be made using arguments external to the statistical procedure.

The Advantages of Probabilistic Models Based on DAGs

The inadequacies of statistical methods based on logistic regression have encouraged the development and use of techniques better able to compare alternative causal pathways. Bayesian networks, a widely used class of probabilistic graphical models, have been recruited for this purpose.[3,6,14,15] They model the overall dependence structure of multiple variables, visualized in DAGs. DAGs incorporate "nodes" (the specific variables being analyzed), joined by "edges," which are lines representing identified directions of effect. These conventions are illustrated in figure 2. A direct edge in the graph implies a direct cause. A directed *pathway* from one variable to another also implies a causal relationship, albeit one effected through the intermediate variables on the path. Nodes are described as the "parents" of those immediately below them in the causal chain, which are conversely referred to as "children."

The logic behind the use of DAGs to depict hypothetical causal structures is set out in detail by Pearl.[16] While causes can never be proved from observational data only, the use of DAGs and kindred analyses can still provide insights about admissible causal relationships. Take a simple example in which A and B both have to precede C. In cross-sectional data we might then observe: only A; only B; A and B; A, B and C; or none. In contrast if there is no relationship between A, B and C, then any possible combination could be observed. Cross sectional data capture participants at different stages of this causal progression. We would therefore expect to observe different patterns under different models, with differential support from the observed data. In other words, *the time dependence is to some extent encoded in the data, and is not completely lost*. Our task is to decode it. For each model encoded by a DAG, we would expect different patterns in the data. In our inferential procedure we effectively invert the problem, in order to identify the models best supported by the patterns in the data. Thus the structural properties of a set of variables are derived by *learning* the underlying graph from the data.[17]

DAGs encode conditional independence relationships, and characterize the joint probability distribution of the chosen variables. This can then be decomposed as a product of the conditional distribution of each node, given its parent nodes in the graph. DAG analysis is capable of pointing to the most likely directional links between multiple variables, thereby locating them in a putative causal cascade, in which upstream variables constitute the causes of downstream variables. There will nevertheless be circumstances where DAG analysis is unable to specify the causal links between given variables, which then require to be placed at the same level of the graph (ie, neither upstream nor downstream) (supplementary material).

The DAG analysis presented here has 2 special features. First, our statistical method allows us to quantify both the *strength* and *direction* of causal effects. Secondly, since several networks may explain the data reasonably well, we account for the uncertainty in the inference by sampling from the range of possible DAGs in proportion to their posterior distribution.

Our use of separate databases using the same methods of assessment provided instant replication.

DAG Analysis of Bullying Victimization and Persecutory Ideation

We present an example of DAG analysis based on a novel Bayesian method for learning the graphical structure[17] and predicting the range of causal effects for binary variables. There is increasing evidence of an association between bullying victimization and the onset of psychotic symptoms in both clinical and nonclinical populations.[18,19] There are a number of possible mechanisms, both direct and indirect. We analyze data from 2 separate British National Surveys

to examine further the link we previously established between a history of bullying and psychotic symptoms.[18]

Note that the role of theory in the current analysis lies only in our choice of the candidate variables. We do not test a specific DAG, but seek to identify the most plausible mechanisms, using the data to learn potential causal structures.[20] Our core assumption is thus that the variables are causally linked (and can be represented by a DAG structure).

We here present analyses based on the 2007 British Adult Survey of Psychiatric Morbidity,[21] which provides data from a large general population sample ($N = 7403$). Equivalent analyses using the 2000 survey ($N = 8580$) are reported in the supplementary material. The surveys provided information on bullying victimization and paranoid ideation. In line with the model proposed by Freeman and colleagues,[22] we predicted specifically that the link between bullying victimization and persecutory ideation would operate through a range of affective symptoms. We had access to measures of worry, depression, anxiety, mood instability, and sleep disturbance. These symptoms are important correlates of psychotic phenomena, not least because they may offer viable targets for psychological treatments.[1,23–25] The derivation of a causal structure would thus add to the plausibility of such interventions in persecutory ideation associated with bullying. We also examined drug use and the experience of hallucinations, both of which have been suggested as plausible mediators. Following Kalisch et al,[26] we dichotomized the data.

In order to illustrate its advantages, we compare the results of DAG analysis with those based on logistic regression as embodied in the commonly used KHB commands.

Further detail on methodology is provided in the supplementary material.

KHB Analysis

We carried out KHB mediation analysis in 2 ways. In the first, each variable was considered by itself as a potential mediator between bullying and paranoia; in the second, all the potential mediators were entered in the same analysis to estimate their combined mediation effect.

When variables were analyzed individually, worry had the largest effect (26.6%), followed by depressed mood (20.4%), sleep disturbance (17.3%), mood instability (16.9%), and anxiety (15.8%). There was almost no mediation by drug use (4.1%) or by auditory hallucinosis (3.0%). When all the variables are entered into the model together they account for 43.2% of the link between bullying and psychosis. The sum of the percentages of indirect effect for the individual variables considerably exceeds the value of the total indirect effect when the variables are entered together. This indicates that the individual indirect effects are partly due to associations between mediators or to their artificially constrained antecedence to persecutory ideation. The DAG analysis on the other hand accounts for the associations between all the putative mediators and their relation to persecutory ideation.

DAG Analysis

In figures 1 and 2 we summarize our DAG analysis of the 2007 dataset. We made the single stipulation that bullying was antecedent (the equivalent to its status as the independent variable in the KHB analysis), although the relationship is unlikely to be so simple.[27] In contrast to the KHB analysis, no restrictions were imposed on the other variables. We learnt multiple DAGs using our partition MCMC method.[17] In total, we sampled 50 000 DAGs in proportion to their posterior probability, so that better fitting graphs were chosen more often (there are over 200 trillion ways of combining our 9 variables). Each sampled DAG gives a model of the data, thereby providing a particular estimate of the probability of each variable being 0 or 1 depending on the state of its parents in the network. Moving downstream from bullying, we were able to sample each variable in turn and hence to generate new data from the graph.

From each individual DAG we obtained a value for the effect of one variable on another. For example we might set bullying to 0 and calculate the probability that each remaining variable is 0 or 1, then set bullying to 1 and repeat the process. The difference between the 2 values is what Pearl would term the causal effect of bullying on the other variables.[16] We were then able to work out the overall effect of changing each variable on all of the others by averaging the effects over the whole sample of DAGs. In figure 1 we show the distribution of causal effects of the row labels on the column labels. The variables are arranged in order of the number of their downstream effects. Where there were no effects between 2 variables, a zero was entered in the relevant box. In the interest of clarity we truncate the effect distributions to the range −0.1 to +0.5 in all other boxes. Zero causal effect is represented in each box by the red vertical line. If the 95% credible interval (the Bayesian counterpart of confidence limits) does not straddle the zero causal effect line, the whole box is colored to indicate that the effect is significant. Each of these boxes also contains a number quantifying the relevant *average causal effect*. As we stipulated that bullying was antecedent (and therefore not caused by any of the other variables under consideration), it occupies the top row. It is notable nonetheless that bullying does indeed have positive causal effects on every one of the other variables.

We repeated the analysis without stipulating the prior position of bullying, and checked to see whether it fitted better elsewhere in the DAG. We found that, without the restriction on its position, it joined the group of variables whose directions weren't determined. Thus the data did not refute the possibility that bullying was antecedent, and our prior belief sanctioned its placement on the top row.

Figure 2 shows the DAG related to the findings above. It reveals a complex structure of relationships beyond

**Fig. 1.** Distributions of downstream causal effects: 2007 dataset.

Fig. 2. Directed acyclic graph of relationships between variables relating to bullying: 2007 dataset.

standard logistic regression, and enables the direction of several effects to be inferred more securely. The strength of the links is represented by the color intensity of the arrows joining the variables. The variables are either joined by single arrows, indicating a causal effect, or by arrows in both directions, implying that the causal influence cannot be uniquely identified from the data. This corresponds to the presence of bimodal peaks in the plots in figure 1.

For example, in figure 1 there is a 4-by-4 block comprising persecutory ideation, worry, mood instability and drug use: each variable sometimes has a causal effect on the others and sometimes none (ie, when the distribution of effects straddles the red line). Note that there is no lack of power: the size of the dataset is sufficient to learn the graph. Additional data are unlikely to resolve the issue since these directions cannot be learnt from the observations.

In figure 2, paranoia, worry, mood instability and drug use lie immediately downstream of bullying. Depression is further down, followed by another group made up of sleep disturbance and anxiety. Finally, hallucination appears on its own. Worry and mood instability have strong downstream effects on depression, sleep and anxiety.

In figure 1, the numbers in the boxes reflecting mean causal effect can be interpreted in terms of a thought experiment about intervention. Consider the worry row. Imagine 2 interventions, one that turns worry on, and another that eliminates it. Applying these interventions to 2 new groups of individuals, the probability of having depression in the forced worry group would be predicted to be 0.33 higher on average than in the second group, while the probability of paranoia would be 0.05 higher (a small effect, consistent with the inability of the DAG to determine causal direction between worry and paranoia).

The parallel analysis of the data from the 2000 National Survey is provided in the supplementary material. It is largely consistent with the results from the 2007 survey presented here, and confirms the stability of the analytic procedure.

Discussion

Our intentions in this article were to demonstrate the potential of DAGs, and to argue for their use in psychiatric epidemiology. DAGs have the advantage over other network analyses based on Markov random fields of suggesting explicit directions for the causal relationships. Particular advantages of our Bayesian method of DAG analysis are that it (a) provides effect estimates of potential interventions and (b) accounts for the uncertainty in the relationship between variables. It thus serves the current interest in the psychological phenomena that underlie the development and maintenance of psychiatric disorders,[1,28] and identifies the most efficient candidates for targeting in the development of psychological interventions.

KHB logistic regression analysis strongly suggested mediation of the link between bullying and persecutory ideation was unlikely to involve either drug use or auditory hallucinations. The variables whose potential role in mediation was not refuted remain as candidates, but we cannot say more than that. The effects provide an estimate of the *maximum* mediation effect. However, due to the existence of potential confounders (such as the other candidate mediators excluded from each individual model), there is no guarantee that such an effect exists. This is highlighted in the DAG analysis, in which the variables are modeled jointly: no effects were found for most of the potential mediators. The discrepant results from the KHB analysis are due to the fact that it obliges paranoia to be placed downstream of the mediator. This is not the case in the DAG analysis.

Our DAG analysis allowed more robust inferences of causality, and in the process revealed a much more complex picture. It refuted the mediation of the link between bullying and persecutory ideation by several of the putative KHB mediators. Thus DAG analysis provided no support for depression, anxiety, sleep disturbance, and hallucinations as mediators, since these variables were all found to be downstream of persecutory ideation. The DAG analysis failed to disambiguate the relationships between worry, mood instability and persecutory ideation. This issue cannot be resolved using these datasets: both directions are plausible. As this is not due to inadequate power, an intervention trial is the ideal way of settling the issue. Indeed our type of DAG analysis not only identifies plausible relationships between potential mediators, but also produces corresponding distributions of potential intervention effects. A recent clinical trial has shown that an intervention to reduce worry does indeed diminish paranoia and therefore supports the role of worry as a mediator.[25] Similarly, persecutory ideation may turn out to be alleviated by reducing sleep disturbance and modifying depressive cognitions.[29,30]

Limitations

DAGs represent a considerable advance over standard logistic regression techniques. They are more informative, by virtue of being capable of inferring both the *strength* and *direction* of the possible connections between the entire set of variables under analysis, and will indicate improbable causal relationships. Even so, caution is still required, as causal inference is only fully defensible under the assumptions of *faithfulness* and *causal sufficiency*.[31] Broadly speaking, faithfulness requires that the joint distribution of the variables satisfies *all* the conditional independence relationships encoded by the DAG, and *only* those relationships. Causal sufficiency refers to the absence of unmeasured confounders or selection variables.[14,15,32,33] Only under these assumptions can the direction of the edges in a Bayesian network learnt from observational data be interpreted causally.

DAGs should be distinguished from recent applications of network approaches to complex interactions based on partial correlations between variables.[6–8,14,15] While the lack of an edge in the latter would exclude a direct causal link between 2 variables, the converse is not true: the presence of an edge does not of itself imply a direct causal link. Moreover, partial correlations do not produce information on the direction of effects.

Like any form of analysis, the use of DAGs is hostage to the quality of the data. The data in the British National Surveys were obtained through carefully conducted interviews, but may not fully capture the underlying constructs. Moreover, we dichotomized the variables, which inevitably led to some information loss. Uncertainty deriving from the data quality would

however be captured to some extent by our Bayesian approach, whereby we obtain flatter posterior distributions with sparse or noisy data. Bayesian networks in the form of DAGs also have the limitation of being unable to model the feedback loops likely in social/psychological systems (though such loops might be identifiable by subjecting time course data to analysis in terms of dynamic Bayesian networks[34]).

Despite these reservations, DAGs can provide important insights into possible causal relationships between the observed variables. These in turn may guide the design of experimental interventions to validate the observational findings.

Supplementary Material

Supplementary material is available at *Schizophrenia Bulletin* online.

Acknowledgments

References

1. Bebbington P. Unravelling psychosis: psychosocial epidemiology, mechanism, and meaning. *Shanghai Arch Psychiatry*. 2015;27:70–81.
2. Bebbington P. Causal narratives and psychotic phenomena. *World Psychiatry*. 2016;15:127–128.
3. McNally RJ. Can network analysis transform psychopathology? *Behav Res Ther*. 2016;86:95–104.
4. Borsboom D, Cramer AO. Network analysis: an integrative approach to the structure of psychopathology. *Annu Rev Clin Psychol*. 2013;9:91–121.
5. Isvoranu AM, Borsboom D, van Os J, Guloksuz S. A network approach to environmental impact in psychotic disorder: brief theoretical framework. *Schizophr Bull*. 2016;42:870–873.
6. Isvoranu AM, van Borkulo CD, Boyette LL, Wigman JT, Vinkers CH, Borsboom D; Group Investigators. A network approach to psychosis: pathways between childhood trauma and psychotic symptoms. *Schizophr Bull*. 2017;43:187–196.
7. Wigman JT, de Vos S, Wichers M, van Os J, Bartels-Velthuis AA. A Transdiagnostic Network Approach to Psychosis. *Schizophr Bull*. 2017;43:122–132.
8. Fried EI, van Borkulo CD, Cramer AO, Boschloo L, Schoevers RA, Borsboom D. Mental disorders as networks of problems: a review of recent insights. *Soc Psychiatry Psychiatr Epidemiol*. 2017;52:1–10.
9. MacKinnon DP. *Introduction to Statistical Mediation Analysis*. New York, NY: Lawrence Erlbaum Associates; 2008.
10. Pearl J. Direct and indirect effects. Paper presented at: Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence; 2001; San Francisco, CA.
11. Karlson KB, Holm A, Breen R. Comparing regression coefficients between same-sample nested models using logit and probit. A new method. *Sociol Methodol*. 2011;42:286–313.
12. Marwaha S, Broome MR, Bebbington PE, Kuipers E, Freeman D. Mood instability and psychosis: analyses of British national survey data. *Schizophr Bull*. 2014;40:269–277.
13. Marwaha S, Thompson A, Bebbington P, et al. Adult attention deficit hyperactivity symptoms and psychosis: epidemiological evidence from a population survey in England. *Psychiatry Res*. 2015;229:49–56.
14. Isvoranu AM, Borsboom D, van Os J, Guloksuz S. A network approach to environmental impact in psychotic disorder: brief theoretical framework. *Schizophr Bull*. 2016;42:870–873.
15. McNally RJ, Robinaugh DJ, Wu GWY, Wang L, Deserno M, Borsboom D. Mental disorders as causal systems: a network approach to posttraumatic stress disorder. *Clin Psychol Sci*. 2015;3:836e849.
16. Pearl J. *Causality: Models, Reasoning, and Inference*. 2nd ed. Cambridge, UK: Cambridge University Press; 2009.
17. Kuipers J, Moffa G. Partition MCMC for inference on acyclic digraphs [published online ahead of print April 2015]. *J Am Stat Assoc*. http://arxiv.org/pdf/1504.05006.
18. Catone G, Marwaha S, Kuipers E, et al. Bullying victimisation and risk of psychotic phenomena: analyses of British national survey data. *Lancet Psychiatry*. 2015;2:618–624.
19. van Dam DS, van der Ven E, Velthorst E, Selten JP, Morgan C, de Haan L. Childhood bullying and the association with psychosis in non-clinical and clinical samples: a review and meta-analysis. *Psychol Med*. 2012;42:2463–2474.
20. Maathuis MH, Kalisch M, Buhlmann P. Estimating high-dimensional intervention effects from observational data. *Ann Statist*. 2009;37:3133–3164.
21. McManus S, Meltzer H, Brugha T, Bebbington P, Jenkins R. *Adult Psychiatric Morbidity in England, 2007: Results of a Household Survey*. Leeds, UK: National Centre for Social Research/NHS Information Centre; 2009.
22. Freeman D, Garety PA, Kuipers E, Fowler D, Bebbington PE. A cognitive model of persecutory delusions. *Br J Clin Psychol*. 2002;41:331–347.
23. Freeman D, Stahl D, McManus S, et al. Insomnia, worry, anxiety and depression as predictors of the occurrence and persistence of paranoid thinking. *Soc Psychiatry Psychiatr Epidemiol*. 2012;47:1195–1203.
24. Freeman D, McManus S, Brugha T, Meltzer H, Jenkins R, Bebbington P. Concomitants of paranoia in the general population. *Psychol Med*. 2011;41:923–936.
25. Freeman D, Dunn G, Startup H, et al. Effects of cognitive behaviour therapy for worry on persecutory delusions in patients with psychosis (WIT): a parallel, single-blind, randomised controlled trial with a mediation analysis. *Lancet Psychiatry*. 2015;2:305–313.
26. Kalisch M, Fellinghauer BA, Grill E, et al. Understanding human functioning using graphical models. *BMC Med Res Methodol*. 2010;10:14.
27. Shakoor S, McGuire P, Cardno AG, Freeman D, Plomin R, Ronald A. A shared genetic propensity underlies experiences of bullying victimization in late childhood and self-rated paranoid thinking in adolescence. *Schizophr Bull*. 2015;41:754–763.
28. Freeman D, Garety P. Advances in understanding and treating persecutory delusions: a review. *Soc Psychiatry Psychiatr Epidemiol*. 2014;49:1179–1189.
29. Freeman D, Waite F, Startup H, et al. Efficacy of cognitive behavioural therapy for sleep improvement in patients with persistent delusions and hallucinations (BEST): a prospective,

assessor-blind, randomised controlled pilot trial. *Lancet Psychiatry*. 2015;2:975–983.

30. Freeman D, Pugh K, Dunn G, et al. An early Phase II randomised controlled trial testing the effect on persecutory delusions of using CBT to reduce negative cognitions about the self: the potential benefits of enhancing self confidence. *Schizophr Res*. 2014;160:186–192.

31. Verma T, Pearl J. Equivalence and synthesis of causal models. In Proceedings of the Sixth Conference on Uncertainty in Artificial Intelligence, pages 220–227, Cambridge, MA, 27–29 July, 1990.

32. Friedman N, Murphy K, Russell, S. Learning the structure of dynamic probabilistic networks. In Proceedings of the Fourteenth conference on Uncertainty in Artificial Intelligence, pages 139-147, Madison, WI, 24–26 July, 1998. Morgan Kaufmann Publishers; 1998.

33. Dawid AP. Beware of the DAG! *Journal of Machine Learning Research Workshop and Conference Proceedings*. 2010;26:99–157.

34. Spirtes P, Glymour CN, Scheines R. *Causation, Prediction, and Search*. 2nd ed. Cambridge, MA: MIT Press; 2000.