

Bifurcation Angles in Ant Foraging Networks: A Trade-Off between Exploration and Exploitation?

Luc Berthouze and Alexander Lorenzi

Centre for Computational Neuroscience and Robotics,
University of Sussex,
Brighton BN1 9QH, UK

L.Berthouze@sussex.ac.uk

<http://www.informatics.sussex.ac.uk/users/lb203/>

Abstract. The distribution of bifurcation angles found in ant foraging networks has been shown to give polarity to the networks so that nest-bound ants reaching a bifurcation can choose the appropriate direction. In this paper, we use an individual-based model to test the hypothesis that this distribution is an emergent property of a population of foraging ants optimising the trade-off between exploitation of the existing network to maximise food intake and exploration of the environment to maximise the population's ability to rapidly adapt to novel or changing environments. We identify a parameter regulating an ant's drives to forage existing trails and explore uncovered areas of the environment as a collective variable controlling the distribution of bifurcation angles in the foraging network and we show that when the exploration-exploitation trade-off is realised, the resulting distribution exhibits the same informational characteristics as that found in the original study.

Keywords: Exploration-exploitation trade-off, ant foraging, network polarisation, bifurcation angle, stigmergy, self-organisation.

1 Introduction

Whilst many ant species use a wide variety of cues to orientate themselves (see [1,2,3,4], for some examples), some ants rely primarily on chemical trails. For those species, reorientation when displaced may be harder unless trails are polarised. In a recent study on Pharaoh's ants, Jackson et al. [6] showed the distribution of bifurcation angles in their foraging networks to follow a normal distribution with mean 53° and standard deviation 15° . They argued that this value (which is shared by other trail-laying species such as the leaf-cutters and seed harvester ants studied by Acosta et al. [5]) is particularly information rich since a returning forager would have a choice of two paths, one only of which deviating greatly from its current heading. In a series of neat experiments, they validated their geometry hypothesis by showing that whilst in straight trails reorientation was as likely to be incorrect as correct, with a bifurcation angle

of 55° , the ratio of correct reorientations to incorrect reorientations was a significant 5.63. Forty three percent of fed ants heading away from the nest (the ‘wrong’ direction) made a U-turn on meeting a bifurcation, and only 8% when heading in the correct direction. Similarly, 47% of unfed ants made a course change when heading in the wrong direction, and again, only 8% made incorrect changes when heading away from the nest. Testing their hypothesis further, the authors showed that the ability of ants to make correct course changes degraded as the bifurcation angle was increased to 120 degrees, at which point the trail lost its polarisation. These results clearly suggest that ants can use the trail geometry to orientate themselves. As Collett and Waxman [7] pointed out, however, the study does not answer the question of how these angles are formed. In particular, it does not explain whether they occur as a result of directional decisions made by individuals when they first branch from a trail, or whether they are formed when the trails become well-travelled.

Central to this question are the concepts of stigmergy [8] and self-organisation, that is, the processes by which large numbers of agents interacting on a local level can coordinate these interactions to produce the complex global behaviours that are apparent in social insects colonies. Ants leave chemical trails in their environment so as to recruit more ants to forage newly discovered food sources, this in turn recruits more ants to the trail in a positive feedback loop. This method of communication through the environment was first described by Grassé [8] and is known as stigmergy. Importantly it does not need any form of centralised control, and individuals need only to react to local changes left by others within their environment. Holland and Melhuish [9] described two minimal qualities needed by an agent and its environment to support stigmergic interactions: (i) the agent must be able to move within the environment, and modify parts of it; (ii) the environment must be able modifiable, and the modifications should persist long enough for other agents acting within it to be affected by them. With such a small set of key features, it is no surprise that we should see many cases of stigmergic interactions with nature. Many of the amazing displays by social insects can in part be explained by stigmergy, however, on its own, it is not enough. Stigmergy only describes how individual colony members are able to communicate indirectly, but it does not explain how these interactions are coordinated to produce the complex colony level behaviours. For this, we need to include the ideas of self-organisation, that is, the “set of dynamical mechanisms whereby structures appear at the global level of a system from interaction among its lower-level components” [10]. Bonabeau et al. [10] described four basic ingredients needed for self-organising patterns to emerge (positive and negative feedback, amplification of fluctuations, and the presence of multiple interactions), the signatures of which are the creation of spatiotemporal structures in an initially homogeneous medium, the possibility of multistability and the existence of parametrically determined bifurcations. Applied to ant foraging, the positive feedback relates to trail recruitment through the leaving of chemical markers. Negative feedback can occur when food sources become depleted, or through saturation when ants are physically unable to access a food source or to enter

the nest owing to the number of other ants. Amplification of fluctuations occurs in situations where ants lose the trail they are following and stumble upon a new food source, the resulting trails will be amplified as more ants are recruited to the new source. The presence of multiple interactions naturally results from the sheer size of ant colonies. Bifurcation angles then become the signatures of the resulting self-organised foraging network. These ideas are central to the concept of quantitative stigmergy [11] which gives us a process by which we can coordinate and regulate stigmergic interactions.

In this paper, we put forth the hypothesis that the distribution of bifurcation angle observed by Jackson et al. is an emerging property of a population of foraging ants optimising the trade-off between exploitation of the existing network to maximise food intake and exploration of the environment to maximise the population's ability to rapidly adapt to novel or changing environments. To test this hypothesis, we constructed a simple individual-based model that incorporates the qualities necessary for self-organisation and stigmergy to take place.

2 Model

The methods used to test our hypothesis are based on common techniques used to examine ant behaviour and pattern formation in natural sciences. The simulation used is a cellular automata (CA), of which the two major components of our model are the ants and the pheromone concentration map. In the context of CA, the ants are modeled as a Lattice Gas class of CA, in which particles operate in a bi-directional grid, and their movement contains a random component. The pheromone concentration map on the other hand falls in the deterministic class of CA, in which states are associated with sites on the grid, and a state's change is a deterministic function of its surrounding sites [12,18]. The environment in which the ants operate is a 500x400 lattice (the unit of distance used throughout this study is the length of one ant). The single nest is located at the centre of the grid, and has a radius of 7.5 units. Ants are released from the nest at a rate of 1 ant every 5 updates. The food sources are randomly placed in the environment, and have a finite amount of food. Once a food source is depleted, it is removed from the environment, and replaced by a new food source in a different location.

Each ant is described by its position on the lattice, its heading and velocity and two coefficients that represent how fatigued the ant is (c_f) and the food load it is carrying (c_{fl}). At each update, the ant's fatigue coefficient increases by 0.0005 until its maximum value of 1 when the ant is considered dead from exhaustion and is removed from the environment. Only when the ant is either at a food source, or in the nest, does the fatigue coefficient decrease (multiplied by a factor 0.9 at each update). The food load coefficient increases when an ant is at a food site (+0.1 unit per update), and decreases when the ant is in the nest (-0.1 unit per update).

Environmental stimuli and internal variables modulate five drives that control the ant's heading and velocity. For simplicity, each of these drives will be described in a polar coordinate system where the pole is the position of the ant

on the lattice, and the polar axis is the ant's current heading. Radial coordinates will denote the velocity component of each drive.

Exploratory drive: This drive implements an ant's 'urge' to change to a new heading deviating randomly from its current heading by $\pm 15^\circ$. The trajectory is smoothed by spreading the change over α_r updates.

$$d_1 = \left(0, \frac{-15 + \text{rnd} * 30}{\alpha_r}\right) \quad (1)$$

where rnd is a random value between 0 and 1 drawn from a uniform distribution.

Exploitation drive: The movement of the ant is affected by the pheromone concentrations ϕ of its eight Moore neighbours

$$d_2 = \left(\sum_{i=1}^8 \phi_i (1 - c_{fl}) \left(1 - \frac{|\theta_i|}{180}\right), \theta_i\right) \quad (2)$$

where θ_i is the angular deviation to the i^{th} Moore neighbour. Note that this expression favours pheromone sites ahead of the ant. This is to prevent the ant from turning around to follow its own trail.

'Move to food' drive: Each food source has a draw area (of diameter $f_d = 3$) in which ants become 'aware' of the food source, either visually or, in the case of the blind army ants, through olfactory perception. This drive is only activated when the ant gets within this area.

$$d_3 = \left((1 - c_{fl}) \left(1 - \frac{d_f}{f_d}\right), \theta_f\right) \quad (3)$$

where θ_f is the angular deviation to the food source and d_f is the distance to the food source. The attraction to the food source is inversely proportional to the distance to the food source.

'Return to nest' drive: An ant will be drawn back to the nest either when it has collected food, or when its fatigue coefficient increases above a threshold. For simplicity, ants maintain a memory of the exact location of the nest relative to their current position. This isn't biologically plausible, of course, but since ants can use path integration and other cues to locate their nest (see [19], for example), it is an acceptable compromise.

$$d_4 = \left(\sum_{i=1}^8 \phi_i c_{fl} \left(1 - \frac{|\theta_i|}{180}\right)^r, \theta_n\right) + \left(e^{-k(1-c_f)}, \theta_n\right) \quad (4)$$

where θ_n is the angular deviation to the nest, θ_i is the angular deviation to the i^{th} Moore neighbour, r regulates the ant's directional selectivity, and k determines at which fatigue level the drive will kick-in. The bias toward stronger pheromone concentrations in front of the ant corresponds to observations that ants move

quicker on trails than on unmarked areas [11]. Since the fatigue coefficient is updated at each step, this is a plausible mechanism from an energetic viewpoint. A value of $r = 4$ was used that implements a good compromise between accuracy and energy saving by way of trail following. Parameter k was set to 10 throughout the study, which corresponds to a fatigue coefficient of 0.5. The drive to return to the nest increases in a non-linear fashion with the fatigue coefficient.

'Move forward' drive: This drive simply 'motivates' the ant to go forward.

$$d_5 = (\min_v, 0) \quad (5)$$

where \min_v is a constant ant velocity (set to 0.3 units per update in all runs). It guarantees that all ants will move at a minimum speed, and also puts a limit to the ant's angular velocity (a similar constraint was used in [20]).

The ant's actual movement is given by the weighted sum

$$d = \gamma d_1 + (1 - \gamma)d_2 + d_3 + d_4 + d_5 \quad (6)$$

where γ is in the range $[0, 1]$ and is used to regulate how exploratory/exploitative the ant is. At $\gamma = 1.0$, the ant is fully exploratory and ignores pheromone concentrations. Conversely, at $\gamma = 0.0$, the ant will fully exploit existing chemical trails. Consistent with animal observations [14,15] and other models [16,17,18], ants lay pheromones both when leaving the nest and returning with food, although returning ants do so at a higher concentration ($0.1c_{fl}$ unit of food per location visited for ants returning to the nest compared with a fixed 0.01 for ants leaving the nest). In addition, the level at which pheromones are no longer deposited differs whether ants are returning to the nest (maximum concentration of 1) or leaving the nest (maximum concentration of 0.2). A similar constraint was used in [16,17].

Finally, two constraints complete the model: (i) as in [16,17], an ant's maximum velocity is limited to 1 unit per update; (ii) a single lattice location can hold a maximum of 15 ants. If an ant's move takes it into a location that contains 15 ants, it attempts to move to the next grid location 1 unit away from its current position while retaining its desired heading. This minimises the occurrence of unwanted branching that might confound our results. If no site is available, the ant remains in place.

3 Results

To examine the relationship between our putative controlling variable γ and the distribution of bifurcation angle in the resulting foraging trail network, we run simulations varying γ in the interval $[0, 1]$ in steps of 0.1. The environments in each simulation had two food sources. To avoid the random bias of differing food source locations, the locations of food sources stayed constant between simulations. Simulations were run for 10000 iterations, with a colony size of 1500 and food source size of 2000.




Fig. 1. Left: Cumulated food intake (vertical axis, per run of 10,000 iterations) as a function of γ (horizontal axis) over 16 runs. Food intake is given in unit of food where 1 is the maximum food load of an ant entering the nest at a given time. Solid line: logistic regression ($y=1274.00/(1+\exp(-(0.93-x)/0.038))$, d.f.=157, $p < 0.0001$). Right: Recovery times (per run of 10,000 iterations) as a function of γ (horizontal axis) over 16 runs. No data is available at $\gamma = 1.0$ since no food sources got depleted. Times are given in number of iterations. Solid line: cubic fit by nonlinear regression ($y=1788.08x^3-3190.93x^2+1367.34x+49.88$, d.f.=6, $p < 0.01$).

The trade-off between exploitation of the network and exploration of the environment was assessed through two 'fitness' measures based on the amount of food entering the nest at each time step. Cumulated food intake over the 10000 iterations of each run provided a measure of the ability of the network to exploit the existing network. Robustness of the population to environmental changes was assessed by the average time (recovery time) it took for the population to find a new food source once a food source had been depleted. As shown by Figure 1, simulations confirmed the role of γ in controlling the overall behaviour of the population. Lower γ settings (highly exploitative) resulted in higher cumulated food intake, whilst higher γ settings (highly explorative) resulted in faster recovery times. The fact that recovery times were faster at $\gamma = (0.1, 0.2)$ than at $\gamma = 0.3$ was unexpected but finds its origin in a flaw of the model owing to an absence of published data on the behaviour of the first ant off the nest in a novel environment. The model being an individual-based model, such data is critical to determine what the behaviour of a highly exploitative ant should be when no established trail network exists. To circumvent this problem, ants were sent off the nest at an angle drawn from a uniform distribution, with a minimal amount of exploratory drive. In a small environment, when food sources are not too far from the nest, this actually amounts to a fairly effective systematic exploration of the environment, thus the fast recovery times. For a given colony size, increasing the size of the environment or placing the food source further from the nest should suppress this artefact.




Fig. 2. Left: Snapshot of a foraging network. The nest is found at the centre of the figure. In the upper right quadrant, a new food source has just been discovered. The brightness of a trail is proportional to its concentration in pheromone. The purple lines denote bifurcations as used to construct Figure 3. Right: Close-up on a bifurcation. Only Y junctions were considered, and a bifurcation angle θ was determined as the angle between branches away from the nest.

To determine the optimal trade-off between exploration and exploitation, we derived a 'fitness' function $y=I(1-1/RT)$ maximising food intake I for low recovery time RT with I and RT approximated by the nonlinear regressions given in Figure 1, after normalisation to $[0, 1]$. Theoretically, the maximum of this function is obtained when $I = 1/RT$. Experimentally, the maximum was obtained at $\gamma = 0.8$, a highly explorative setting.

As in Jackson et al. [6], bifurcation angles were measured from well-structured foraging networks over three of the 16 simulation runs collected¹. Only networks that showed well-established Y junctions were considered (see Figure 2, left). A bifurcation angle was defined as the angle between the two branches away from the nest (see Figure 2, right). The distribution of bifurcation angle for each γ was determined using normal mixture modelling [21]. At $\gamma = 0.8$, the putative optimal trade-off between exploration and exploitation, the density function showed a mean angle of 40.84° and a standard deviation of 12.69° ($n = 67$). In comparison, Jackson et al. [6] found a mean of 53.48° , and a standard deviation of 14.88° . The model used in this study being an abstract model with parameters having little or no biological relevance, any outright comparison of the numerical values of the means and standard deviations would have little significance. Nevertheless, the data shown in Figure 3 (top) display three important characteristics. First, the probability of bifurcation angles of more than 80° is very small, which means that the resulting network is polarised. It is not the case in more exploitative settings. Second, the number of clusters (or components) as identified by Bayesian Information Criterion (BIC) decreases as γ increases. As in Jackson et al.'s study, the distribution of the bifurcation angles at high γ (highly explorative population) is a normal distribution. Third, as shown by

¹ This is a very laborious activity which we are trying to automatise.

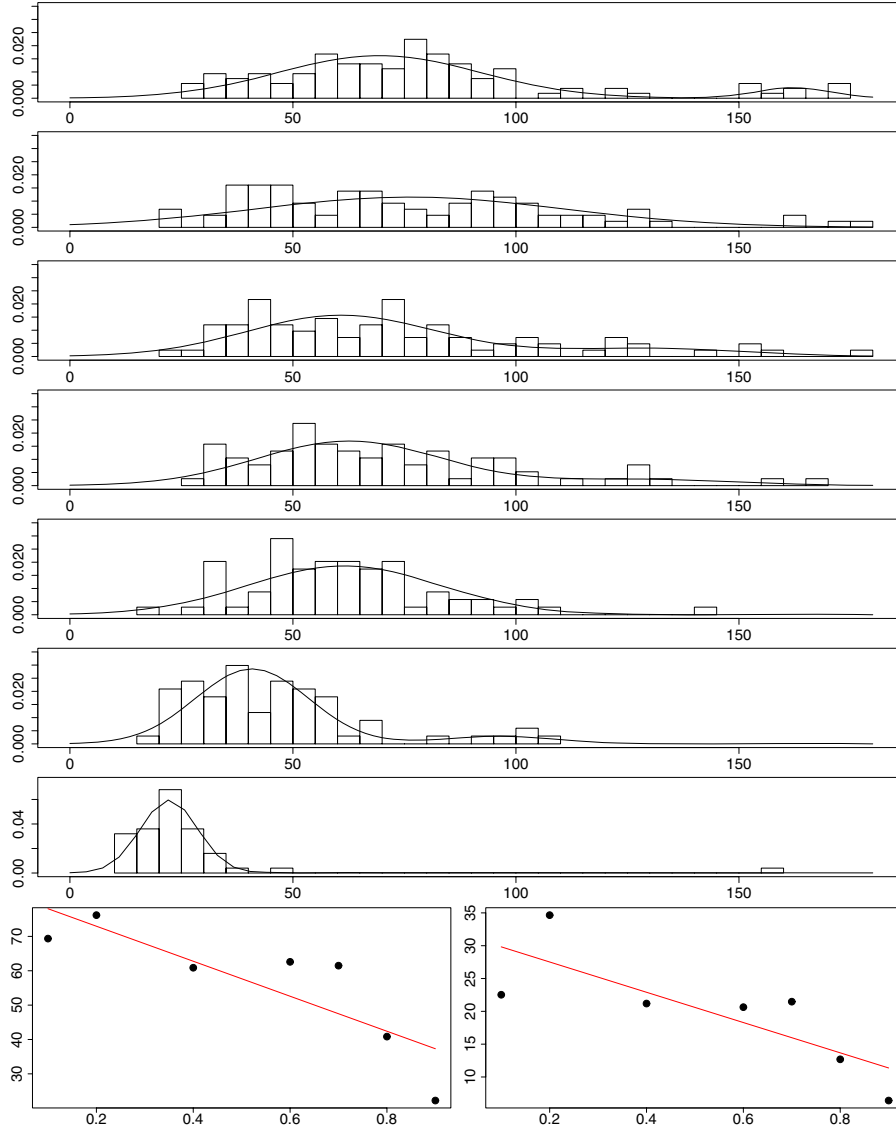


Fig. 3. Top: Densities of bifurcation angles for $\gamma = (0.1, 0.2, 0.4, 0.6, 0.7, 0.8, 0.9)$ (from top to bottom) from three separate runs ($n > 69$ for all γ setting except for $\gamma = 0.9$ when $n = 50$). Note the different scale in vertical axis for $\gamma = 0.9$). Horizontal axis denotes angular values in degrees. Solid line, fit from normal mixture modelling [21]. Bottom left: Mean (in degrees) of the largest cluster as a function of γ . Fit by linear regression ($R^2 = 0.700$; $d.f. = 5$, $F = 11.69$, $p = 0.019$). Bottom right: Standard deviation (in degrees) of the largest cluster as a function γ . Fit by linear regression ($R^2 = 0.637$; $d.f. = 5$, $F = 8.76$, $p = 0.032$).

Figure 3 (bottom), both mean and standard deviation of the main cluster (component) show a statistically significant correlation with γ . This suggests that γ , an individual-based control parameter, is indeed a control variable for the distribution of bifurcation angle in the foraging network.

4 Conclusion

In providing evidence that an individual-based parameter can turn into a collective variable controlling the geometry of the foraging network, this paper aims to answer an important question raised by Jackson et al.'s study, namely, whether bifurcation angles occur through directional decisions made by individual ants when they first branch from a trail, or whether they form when travelled by many ants [7]. Our results support the latter view, that is, bifurcation angles are an emergent property of the stigmergic system. Our results also support our hypothesis that the distribution of bifurcation angles observed by Jackson et al. results from an optimisation of the trade-off between exploitation of the existing network to maximise food intake and exploration of the environment to maximise the population's ability to rapidly adapt to novel or changing environments. Whilst such an hypothesis is not unreasonable from an evolutionary perspective, the limitations of using an abstract model with parameters have little or no biological relevance are obvious. Yet, and as often when using the animat approach, constructing a model raises interesting questions that warrant further study: (i) what is the behaviour of the first ant off the nest? could individual ants' trajectories be modelled by correlated random walk? (ii) how could the γ parameter be implemented: switching individual behaviour or changing ratio of scouts to foragers?

Acknowledgments. We are grateful to Pierre Nouvellet, David Waxman, and Paul Graham (all from the Department of Biology and Environmental Science at University of Sussex) for useful discussions and advices.

References

1. Carthy, J.D.: The orientation of two allied species of british ant. I. Visual direction finding in *Acanthomyops (Lasius) niger*. *Behaviour* 3, 275–303 (1951)
2. Banks, A.N., Syrgley, R.B.: Orientation by magnetic field in leaf-cutter ants *Atta colombica* (hymenoptera forcidae). *Ethology* 109, 835–846 (2003)
3. Carthy, J.D.: The orientation of two allied species of british ant. II. Odour trail laying and following in *Acanthomyops (Lasius) niger*. *Behaviour* 3, 304–318 (1951)
4. Hölldobler, B.: Recruitment behaviour, home range orientation and territoriality in harvester ants, *Pogonomyrmex*. *Behavioral Ecology and Sociobiology* 1, 3–44 (1976)
5. Acosta, F.J., López, F., Serrano, J.M.: Branching angles of ant trunk trails as an optimization cue. *Journal of Theoretical Biology* 160, 297–310 (1993)
6. Jackson, D.E., Holcombe, M., Ratnieks, F.L.W.: Trail geometry gives polarity to ant foraging networks. *Nature* 432, 907–909 (2004)

7. Collett, T.S., Waxman, D.: Ant navigation: Reading geometrical signposts. *Current Biology* 15(5), 171–173 (2005)
8. Grassé, P.P.: La reconstruction du nid et les coordinations interindividuelles chez *bellicositermes natalensis* et *cubitermes* sp. La théorie de la stigmergie: Essai d'interprétation des termites constructeurs. *Insectes Sociaux* 6(1), 41–80 (1959)
9. Holland, O., Melhuish, C.: Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life* 5(2), 173–202 (1999)
10. Bonabeau, E., Theraulaz, G., Fourcassié, V., Deneubourg, J.: The phase-ordering kinetics of cemetery organization in ants. *Physical Review E* 57(4), 4568–4571 (1998)
11. Bonabeau, E., Dorigo, M., Theraulaz, G.: *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York (1999)
12. Ermentrout, B., Edelstein-Keshet, L.: Cellular automata approaches to biological modeling. *Journal of Theoretical Biology* 160, 97–133 (1993)
13. Resnick, M.: *Turtles, Termites, and Traffic Jams. Explorations in massively parallel microworlds*. MIT Press, Cambridge (1997)
14. Aron, S., Pastells, J.M., Deneubourg, J.L.: Trail-laying behaviour during exploratory recruitment in the argentine ant, *Iridomyrmex humilis*. *Biology of Behavior* 14(3), 207–217 (1989)
15. Fourcassié, V., Deneubourg, J.L.: The dynamics of collective exploration and trail-formation in *Monomorium pharaonis*: experiments and model. *Physiological Entomology* 19, 291–300 (1994)
16. Deneubourg, J.L., Gross, S., Franks, N., Pasteels, J.M.: The blind leading the blind: Modeling chemically mediated army ant raid patterns. *Journal of Insect Behavior* 2, 719–725 (1989)
17. Solé, R.V., Bonabeau, E., Delgado, J., Fernández, P., Marn, J.: Pattern formation and optimization in army ant raids. *Artificial Life* 6(3), 219–227 (2000)
18. Watmough, J., Edelstein-Keshet, L.: Modelling the formation of trail networks by foraging ants. *Journal of Theoretical Biology* 176(3), 357–371 (1995)
19. Santschi, F.: Observations et remarques critiques sur le mécanisme de l'orientation chez les fourmis. *Revue Suisse de Zoologie* 19, 305–338 (1911)
20. Couzin, I.D., Frank, N.R.: Self-organized lane formation and optimized traffic flow in army ants. *Proceedings of the Royal Society B: Biological Sciences* 270(1511), 139–146 (2003)
21. Fraley, C., Raftery, A.E.: *MCLUST Version 3 for R: Normal Mixture Modeling and Model-Based Clustering*. Technical Report, No. 504, Department of Statistics, University of Washington (2006)