

Finding optimal behaviors with genetic algorithms

Thomas S. Hoffmeister and Éric Wajnberg

Abstract

Optimality models used in behavioral ecology, and especially on insect parasitoids, have taken a variety of approaches, from classical analytical tools to individual-based simulations. The increasing awareness that much of the observable behavior in parasitoids depends on the state of the insect (be it the physiological or informational state) has led to the increasing use of stochastic dynamic programming models. However, optimal behaviors of one individual often depend upon the behavior of conspecifics, further complicating the issue. While classical game theory may be applied when behavior is not state dependent, genetic algorithms (GA) provide a powerful way in finding optimal behaviors for situations where such optimal behavior depends upon an animal's state and on the frequency of alternative behaviors of conspecifics. More generally, GAs can be used when there is a need to find the optimal strategy among a number of different alternative behaviors whose number is far too great to be exhaustively checked. GAs are search algorithms that proceed in a fashion analogous to natural selection. They are individual-based simulations that identify optimal solutions by searching the enormous space of potential solutions mimicking the process of evolution in biological systems.

In this chapter, we discuss research questions in parasitoid behavioral ecology that might benefit from applying GA as a research tool for optimality models. Examples will be developed and explained. Among others, using the foraging problem of estimating habitat quality, we show how fast GAs find optimal solutions for parameter spaces that are impossible to solve numerically otherwise.

17.1 Introduction

Behavioral ecology studies the ecological and evolutionary basis of animal behavior and the roles of behavior in enabling an animal to adapt to its environment. Often, the aim is to try to understand what sort of optimal behavioral strategies animals should adopt,

in a particular environmental situation, in order to maximize their production of progeny (i.e. their so-called 'fitness'). Along with other areas of evolutionary biology, behavioral ecology has incorporated a number of techniques, which have been borrowed from optimization theory. An optimized strategy is a strategy that offers the highest fitness return to an animal, given all the different factors and constraints it is facing. Such optimal behavioral rules are then tested experimentally on real animals, and this has especially been done several times on insect parasitoids, producing important and interesting results (see Godfray 1994, for several significant examples).

Several different methods are available to find optimal theoretical predictions, depending on the behavioral ecological problem to be solved. Most of the time, and especially for early work, models were sufficiently simple to be solved using standard maximization techniques. This is the case for most static, optimal foraging theoretical models in which the state of the animal remains unchanged during its foraging time and when there are no competitors trying simultaneously to maximize their own fitness in the environment (e.g. the marginal value theorem (MVT) of Charnov 1976). However, behavioral ecologists, and especially those working on insect parasitoids, are progressively addressing questions that are closer to real situations and that are correspondingly becoming increasingly complex and difficult to solve analytically. For example, if changes in the state of the animal during the foraging process are able to change the optimal behavioral decisions it should adopt, then stochastic dynamic programming models (SDPs) should be used to find the optimal trajectory of the foraging animal within the state-space, leading to maximal lifetime fitness gain (Clark & Mangel 2000, and see also Chapter 15 by Roitberg and Bernhard). If animals need to learn the different features of their habitat to behave in an optimal way and if they can do that by updating some prior information with current foraging experience, then so-called Bayesian approaches can be used (see also Chapter 16 by Pierre and Green). Also, if foraging animals are trying to maximize their fitness gain in the presence of competitors trying to reach the same goal, then models from game theory (Maynard-Smith 1982, Giraldeau & Caraco 2000) can be used to find the best behavioral strategy they should adopt, eventually leading to Evolutionarily Stable Strategies. Finally, these different approaches can be combined when the problem to be solved involves several of the corresponding situations they can address.

Finding optimal behavioral strategies with one or several of these methods is certainly the most elegant way to solve the questions being addressed but, in an increasing number of cases, the problems to be solved appear to be too complicated to find optimal solutions that way. This will be the case, for example, if the number of parameters (and/or their combination) is leading to an enormously large state-space of possible solutions, making it infeasible to search the state-space of these so-called NP-complete problems entirely (i.e. without using heuristics, see Garey and Johnson (1979)). In these cases, the optimal behavioral strategies animals should adopt can still be found by means of numerical methods. Among these methods, Genetic algorithms (GA) appears to be both the simplest one to be implemented on a computer and one of the most efficient to find a rapid and supposedly optimal solution. The aim of this chapter is to present what this numerical method is and how it works, but also, by means of several working examples, how it can be used to solve a variety of research questions that can be raised on the behavioral ecology of parasitoids.

GAs were invented in the early 1970s to mimic some of the processes observed in natural evolution. GAs, based on the observation that life with such a level of complexity was able to evolve in a relatively short time, use the power of evolution to solve optimization

problems. The idea was invented and then initially explored by Holland (1975), with the idea of using mathematically-based artificial evolution as a method to conduct a structured search for solutions to complex problems. As stated by another pioneer in GA optimization, Goldberg (1989): 'Three billion years of evolution can't be wrong. It's the most powerful algorithm there is.' Hence, the design of the GA is inspired by Darwinian natural selection and biological genetic evolution. Thus, GAs include common ingredients present in real organisms that allow them to evolve, such as chromosomes, genetic recombination through crossing-over, and mutation, as well as forces acting on them such as fitness and selection.

Biologists willing to study in a formalized way the biological systems they are working on have used, since Leonardo da Pisa in the 12th century – and are still using – methods and concepts than have been designed by mathematicians. GAs are one of the few examples (with neural networks, Haykin 1999) where, in turn, biological concepts were used as a template for designing methods used by optimization theoreticians. However, despite being biologically-inspired, applications of GAs are mainly outside the field of biology, and it is only recently that such optimization methods started to be used to solve ecological research problems.

In the following, we will outline how a GA works, illustrating this with the example of the traveling salesman problem (TSP), a well-known example for a simple problem with a huge search space. Further, we will provide three examples from parasitoid behavioral ecology, in order to illustrate various technical approaches on different research topics.

17.2 Outline of a GA

Although a GA of the TSP contains some special features that differ from many other forms of a GA, we use this example since it represents a widely known optimization problem with a relatively huge state space of solutions that is almost impossible to solve by enumeration. In the TSP, there are N cities with a given distance between each of them. A traveling salesman starting at one of the cities, has to visit all of the N cities and then return to the first city. However, his task is to find a sequence of cities to minimize the distance traveled in total, without visiting a city more than once. To search the state space of possible solutions, we need to define some solution and calculate the traveled distance. In order to compare solutions, we define a given set of solutions and evaluate the performance of each of them. We will drop poor solutions and create new solutions. Since we create new solutions by altering existing ones, the new solutions only differ in small parts and we search the space around good solutions only. This process is called hill-climbing and is extremely efficient in finding optimal solutions.

This is essentially how a GA works. As you will see, the outline of the basic GA is very general. There are many parameters and settings that can be implemented differently in various problems. A GA typically comprises a population of individuals (or solutions to a problem) that exhibit heritable characters. These characters are numerically coded on a 'string' (Goldberg 1989), which is often called a 'chromosome' because of its intuitive conformity with the evolutionary coding of characters in organisms. Different 'traits' of an individual are coded on 'genes', and different 'character states' are coded as alternative 'alleles' of a gene. The numerical coding of character states may either be in binary code, or represented by integer or real values (Fig. 17.1).




Fig. 17.1 Coding of 'genes' on a 'chromosome' for a GA, in binary (upper row) and integer (lower row) coding. Each 'gene' codes for a character state of the organism.

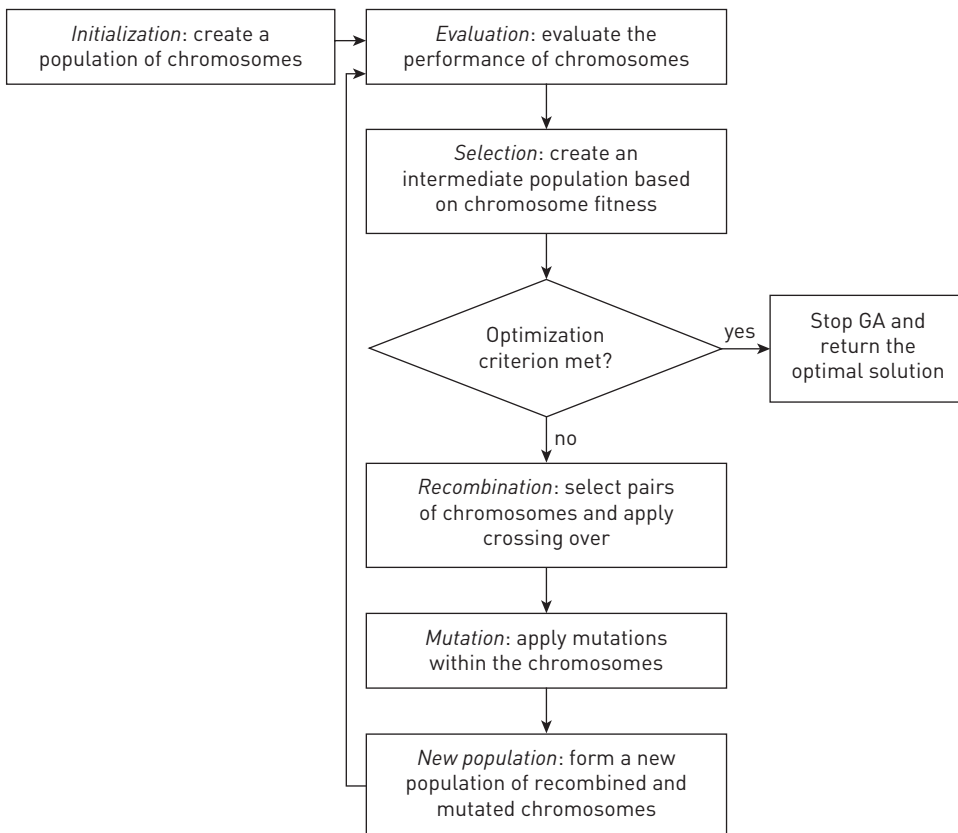


Fig. 17.2 Flow chart of a GA. After initialization, the process is iterated until the optimization criterion is met.

The first step in implementing a GA is the initialization of a population by generating a random population of n chromosomes (Fig. 17.2). Obviously, n should be suitable for the research question to be tackled. Obitko and Slavík (1999) suggest that a population of 20 to 30 chromosomes often performs best (Kumsawat et al. 2005), but some GA problems might require populations of 50 or 100 chromosomes. Generally, the number

of chromosomes in the population should be related to the length of the chromosome, i.e. longer chromosomes require larger populations of chromosomes. Note that each chromosome represents an individual that can be tested against a fitness function or, more generally, against the problem defined in the GA. In our example of the TSP, the chromosome holds N genes with the number-code of one of N cities (each city occurring a single time) and the sequence of visits to cities is given by the sequence of genes on the chromosome. This type of coding is called permutation encoding.

The next step in a GA is the evaluation of each chromosome j using a fitness function $f(j)$, or – as is often the case in game theoretical or other frequency and density dependent approaches – a simulation where chromosome j plays against all other chromosomes in the population and the performance of each chromosome j is evaluated at the end of the simulation (Fig. 17.2). If a stochastic simulation (or so-called Monte Carlo, Rubinstein 1981) is used, the values reached in the evaluation of chromosome j will depend upon stochastic events and might vary noticeably between simulation runs. In order to achieve reliable values, the evaluation process may be replicated sufficiently often (Barta et al. 1997, Hoffmeister & Roitberg 1998, Reinhold et al. 2002) to receive a stable average performance value for a chromosome. In the case of the TSP, the performance of a chromosome does not depend on the coded behavior of alternative chromosomes and thus the total distance traveled on the journey between cities only depends upon the sequence of cities that are visited and the distances between cities. Typically, the position of cities would be coded as the x/y -position on a grid, the sequence of cities that are visited is read from the chromosome, and the distances traveled are summed up across all N cities visited.

While the performance values of chromosomes in the evaluation process are absolute values, in the next step of the GA, a selection process among chromosomes is performed that is based upon relative fitness values, rather than the absolute performance of a chromosome. In this step an intermediate population is formed from the original population of chromosomes, where the number of copies at which a given chromosome is represented depends upon the fitness of the chromosome.

The next question is how to select parents for the next generation. The general idea of the selection process is that chromosomes with low fitness will be eliminated and chromosomes with high fitness will spread in the population by bestowing their code to multiple copies. This can be done in many ways. The most basic approach is called ‘roulette wheel selection’. Imagine all genotypes (represented by chromosomes) are arranged in a pie chart in slices of a size proportional to their relative fitness. Around that pie chart, a roulette wheel is placed. For population size n , n spins of the roulette wheel are performed and the chromosome at which the pointer stops is selected with replacement and introduced to the intermediate population (Fig. 17.3a). That way, chromosomes with high fitness values will be represented with numerous copies, while poorly performing chromosomes will be represented with only a few copies or no copy at all.

However, due to the stochastic nature of roulette wheel selection, the representation of chromosomes might not match exactly the relative fitness ($= (r/\text{sum of fitnesses}) \times 100$, with r being the fitness of chromosome i) of each chromosome. An improvement to this selection process is the stochastic remainder selection (Goldberg 1989). Imagine a population of $n = 100$. When the total sum for the fitness is set to 100, a chromosome j_1 with relative fitness 1.4 would receive 1 copy in the intermediate population, a chromosome j_2 with relative fitness 5.9 would receive 5 copies, and so on. In a first step of selection, these

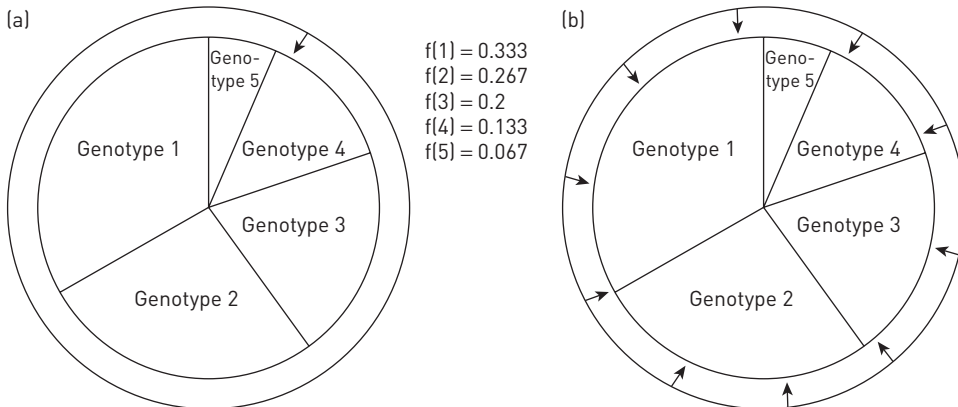


Fig. 17.3 Selection process of a GA by (a) roulette wheel selection or (b) stochastic remainder selection. Values in the center of the graph represent relative fitness values of the five genotypes present in the parental generation. The pointers show selected genotypes. The wheel is spun n times in (a) (with n = population size), whereas n equally spaced pointers are spun once in (b) [Aarsen 1983].

integer parts of the fitness are allocated. Only the non-integer parts of the relative fitness (0.4 for j_1 , 0.9 for j_2 etc.) would enter the roulette wheel selection in a second step of selection as described above. The same gain in precision is reached with the so-called universal selection. Universal selection can be implemented by placing n equally spaced pointers at the roulette wheel on which the genotypes are arranged according to their proportional fitness. A single spin of the wheel determines all selected chromosomes. For each pointer, a single chromosome is created from the genotype the pointer aims at (Fig. 17.3b).

If there are huge differences between fitness values, few other than the best chromosome will have a chance to be selected in the above-mentioned process. Under such conditions, it is advisable to use rank selection rather than a selection based on proportional fitness values. Otherwise, the optimization process might converge to a local optimum. In rank selection, the worst chromosome receives fitness 1, the second worst receives fitness 2, and so on, and the best chromosome receives fitness n . This improves the chances of low-fitness chromosomes being selected at the cost of the speed at which the GA converges to the optimal solution. However, the risk that the GA is caught on a local optimum should be decreased. A further alternative method is steady state selection where a large number of chromosomes with high fitness values are retained and a number of low fitness values are deleted, and the intermediate population is filled up with copies of the retained chromosomes.

One method that distinctly improves the performance of GAs is the introduction of elitism into the selection process. The following step of recombination and mutation in a GA, discussed below, includes the risk that the best chromosomes might be destroyed in the process. Elitism is the fact that the best or a few best chromosomes are retained unaltered in the next generation. If j elite chromosomes are selected, the intermediate population then has the size $n-j$ chromosomes and is produced by one of the above-mentioned selection schemes from the entire set of chromosomes in the population, including those that have been selected as elite chromosomes. In the case of our TSP example,

we might use elitism combined with roulette wheel selection. From our population of 30 chromosomes we retain the two chromosomes with the highest fitness values and form an intermediate population of 28 chromosomes, based upon the relative fitness of all 30 chromosomes. This intermediate population of 28 chromosomes now undergoes recombination and mutation, which is the next step in a GA.

Recombination by genetic operators, such as crossing over and mutations, form the core of GA and are applied in a probabilistic fashion. Crossing over combines elements of two chromosomes in a filial chromosome by exchanging substrings and can, for instance, combine adaptive character states from different genes. Mutations alter individual genes on a chromosome and both crossing over and mutation help to search the previously unsearched state space. While crossing over in a population that converges to the optimum will often lead to no changes (two similar chromosomes are paired and exchange sequences), mutations always change the code of a chromosome. Usually, crossing over rates chosen in GAs are generally high, for example $p_{\text{cross}} = 0.8$, while the probability that a mutation occurs at a given position of a chromosome should be low, e.g. $p_{\text{mut}} = 0.01$. However, there is no fixed value for either crossing over or mutation rate that performs best and optimal rates are a matter of selection by trial and error for each GA. There is a number of different ways of applying crossing over and mutation, and the following description presents only some possibilities.

To employ crossing over, two chromosomes from the intermediate population are randomly drawn. With probability p_{cross} , crossing over occurs at a randomly chosen position on the chromosome. In binary coded chromosomes, the crossing over point may be at any position, i.e. also within genes, while in chromosomes coded in integer or real values, the crossing over point will always be between genes. In crossing over, the sequence of the first chromosome, up to the position of crossing over, is written on the filial chromosome and combined with the sequence of the second chromosome from the position of crossing over until the end of the chromosome, and vice versa (Fig. 17.4).

An exception to this mechanism can be found in chromosomes with permutation encoding, as in our example of the TSP. Since each city can only occur once, the sequence of cities up to the point of crossing over is taken from the first chromosome. Following, the remaining slots on the filial chromosome are filled by selecting cities from the second chromosome in the sequence they occur in this chromosome, if they have not yet been occurring in the sequence on the filial chromosome. For example, let the first and second chromosomes having a length of 10 and the sequences 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 and 1, 3, 5, 7, 9, 2, 4, 6, 8, 10. With crossing over after position 6 we get for the filial chromosome A: 1, 2, 3, 4, 5, 6, 7, 9, 6, 8, 10 (sequence from second parental chromosome underlined). The next step in a GA is to incorporate mutations which are applied with p_{mut} at each position of a chromosome. In binary encoded chromosomes, a mutation changes the allele from 0 to 1 or 1 to 0. In value encoded chromosomes, a new value is either chosen randomly within the range of possible values or the current value is changed by a small increment or decrement. Whether the random value should be drawn from a uniform or non-uniform distribution depends on the phenotypic behavior the gene is coding for. For example, if residence times of a parasitoid (that are usually skewed to the right) are considered, it makes intuitive sense to draw values from an exponential distribution rather than a uniform distribution. Our example of the TSP again requires a different scheme. Since the number of a given city cannot be changed to another value and since all cities are to be represented a single time on the chromosome, for each mutation, two positions

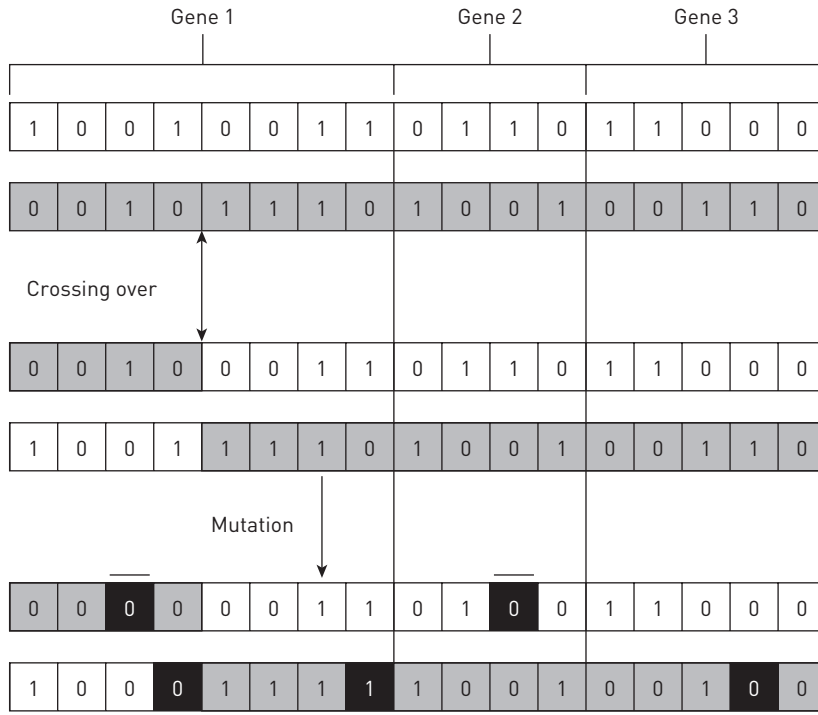


Fig. 17.4 Recombination and mutation in a GA: daughter chromosomes are created from parent chromosomes by crossing over and mutation.

on the chromosome are randomly drawn and the numbers are swapped, thereby altering the sequence of cities to be visited.

After applying crossing over and mutations, a new filial population of chromosomes is created that contains those chromosomes of the parent generation that had high fitness values, but many of these will be slightly altered through crossing over and mutation. By applying selection and the genetic operators of crossing over and mutation, we generate a new set of individuals that have a better than average chance of performing well. By closing the generational loop in a GA (Fig. 17.2) and iterating the process of evaluation, selection, and recombination across many generations, the overall fitness of the population improves as the GA converges to the optimal solution. The obvious next question is when should a GA be stopped? Again, a number of possibilities exist. GAs may either be run for a fixed number of generations (Barta et al. 1997, Bouskila et al. 1998, Reinhold et al. 2002, Perry & Roitberg 2005) or for a fixed computing time. An alternative to these arbitrary stop criteria is to analyze the performance variation of the GA. The GA might be stopped, if a given number of generations has passed with no improvement in the fitness of the best chromosome, or if the improvement in the fitness of the best chromosome has not exceeded a certain increment within a specified number of generations. In our example of the TSP, we stop the GA after 400 generations. As can be seen in Fig. 17.5, the distance traveled decreases most rapidly within the first few generations of the GA runs and quickly converges on a minimum value for the traveled distance that is stable from




Fig. 17.5 Evolution of the length of the best tour across generations of the GA for the TSP. The most drastic reduction in tour length is achieved in the first couple of generations by the optimizing algorithm.

generation 353 onward. Likewise, the number of genotypes, i.e. chromosomes with different encoded sequences, which is close to the population size in the beginning of the GA, also decreases across generations. Note that always some variability in chromosomes will be retained due to crossing over and mutation events.

17.3 A first example of parasitoid behavioral ecology: Learning and optimal patch visitation

Most of insect parasitoids are exploiting hosts that are patchily distributed in the environment (Godfray 1994). There is now a huge amount of both theoretical and experimental work that have been developed to understand what should be the optimal time female parasitoids should invest in each host patch they are exploiting before leaving to find another patch in the environment (Wajnberg 2006, see also Chapter 8 by van Alphen and Bernstein). For this reason, several optimal foraging theoretical models were built in order to understand what behavioral strategy should be adopted in order to maximize that rate at which hosts are encountered and exploited per time unit.

In this respect, one of the most important rate maximization models is the MVT (Charnov 1976) that predicts the optimal time a foraging animal should remain on a host patch (Wajnberg 2006, see also Chapter 8 by van Alphen and Bernstein). This model is based on several assumptions that are probably unrealistic, the most important one being that animals are omniscient, having complete knowledge of the quality of all available patches in the environment and of the mean time needed to reach them (Stephens & Krebs 1986).

This assumption was addressed by several authors, which considered that animals sample patches in the habitat and learn progressively from both their profitability and inter-patch travel times. This is the case for the learning process proposed by McNamara and

Houston (1985), which appears to be both simple and efficient. In their model, the foraging animal initially does not know the quality of the available patches in its habitat and the travel times needed to reach them, but is able to collect such information progressively through the following learning process: patches of different quality are visited sequentially, t_n is the time spent on the n^{th} patch, R_n is the fitness acquired on that patch, and τ_n is the time spent traveling between the n^{th} and $n+1^{\text{th}}$ patch. Let $T_n = t_n + \tau_n$, so $T_1 + T_2 + \dots + T_n$ is the total time between arrival on the first patch and arrival on the $n+1^{\text{st}}$. It is supposed that the animal has some initial constant values R_0 and T_0 ($R_0, T_0 > 0$) and thus has an initial estimate $\gamma_0 = R_0/T_0$ of the average profitability of its environment. According to the MVT, the animal should leave the first patch encountered when its rate of fitness acquisition drops to γ_0 . When it arrives on the next patch, T_1 and R_1 will be known and an updated estimate γ_1 can be computed that will be used to determine when this other patch should be abandoned, and so on. Then, the whole learning process works iteratively according to the two following processes: (i) On arriving on the i^{th} patch, and only at that time, the animal computes an estimate of average environment profitability γ_i where

$$\gamma_i = \frac{R_0 + R_1 + \dots + R_i}{T_0 + T_1 + \dots + T_i} \quad (17.1)$$

and (ii) the animal leaves this i^{th} patch when its rate of fitness acquired on that patch reaches γ_i .

McNamara and Houston (1985) demonstrated that animals using such a simple learning rule asymptotically (i.e. when $n \rightarrow \infty$) converge to the optimal patch allocation strategy predicted by the MVT (i.e. γ_i converges to the true average profitability of the environment γ^*), even if there are different patch qualities in the habitat. The rate of convergence is lower if there is an important variation in patch quality. However, in practical situations, female parasitoids are not sampling an infinite number of patches in their lifetime and might not reach an optimal behavioral strategy when adopting such a learning rule. In this case, as noted by McNamara and Houston (1985), the order in which the different patch types are visited influences the estimates γ_i and thus influences both the time animals will spend on each visited patch and their total fitness gain.

Insect parasitoids are short-lived animals and this is known to influence their patch time allocation foraging strategy (Wajnberg et al. 2006). Thus, they are obviously not visiting an infinite number of patches in their lifetime and, if we assume that they are using the learning rule proposed by McNamara and Houston (1985) in an environment in which several patch qualities can be found, we can ask the question of what should be the optimal order in which all patches should be visited in order for them to maximize their fitness gain. When just a low number of patches are visited, a stochastic dynamic programming model (SDP) can be used to find the optimal visiting order (see also Chapter 15 by Roitberg and Bernhard). However, when the number of patches the animal can visit increases, and especially if there is a large number of patch qualities in the habitat, the number of possible combinations will be too large and a GA should be used instead to find the optimal solution. As an example, suppose that a parasitoid female is able to visit 20 patches in her lifetime, and that the 20 patches are all of different quality. There are thus $20!$ (i.e. about 2.43×10^{18}) possible orders in which the patches can be visited. If a computer is able to evaluate the fitness gained for 100 of these orders per second, it will need more than 770 millions of years to perform all of the computations!

Thus, we developed a GA to find the optimal visiting sequence of 20 patches of different quality. The GA we used for this is similar to the one used to solve the classical TSP presented above. The differences are that: (i) it is the optimal order of patches, and not cities, that we look for: and (ii) the fitness of each genotype does not correspond to the total distance traveled between cities, but to the final γ_i after the 20th patch visited is left. In our example, the quality (in terms of number of hosts to be attacked) of the i^{th} patch is $10 \times i$, and the female is traveling exactly 100 time steps between each patch before reaching each of them. We designed a GA with 50 chromosomes, each bearing 20 genes indicating in what order the different patches should be visited (i.e. so-called permutation encoding). In order to compute the fitness of each chromosome, we used the learning process proposed by McNamara and Houston (1985) with $\gamma_0 = 0.1$, and the fitness gain on each patch was computed as in Wajnberg et al. (2006), using iteratively a Holling (1959) Type II functional response, with attack rate $a = 0.01$ and handling time $Th = 0.1$. At each generation, the two best chromosomes were kept and used for the next generation (i.e. elitism), and each pair of chromosomes experienced a crossing over event at a rate of 80%. In order to always have all patches visited, and after a location is drawn randomly, a crossing over consisted of copying a part of the chromosome of the first parent (i.e. part of the patch order) and the rest of the patches was taken in the same order as in the second parent. Finally, in order to increase the convergence rate of the GA, mutations were allowed using a systematic process: patches were systematically chosen and exchanged if it led to an increase in the resulting fitness. After 300 generations, the best, although not necessary optimal order in which to visit the 20 patches in order to obtain the highest fitness gain rate, is shown in Fig. 17.6.




Fig. 17.6 Sequence of 20 visited patches of different quality leading to the higher fitness rate gain, produced by a GA, for female parasitoids learning progressively the quality of their habitat, using the McNamara and Houston's (1985) learning process.

The best visiting sequence thus seems to be to visit first medium-quality patches. Doing this, females are most likely getting an estimate of the average quality of the habitat that is as close as possible to the true average profitability of their environment. Then, females should progressively and alternatively sample patches of higher and lower quality, leading them to refine their first estimate and then progressively approach the true value of average habitat profitability.

17.4 A second example of parasitoid behavioral ecology: optimal patch time allocation for competing parasitoids with a limited time horizon

As already explained above, the question regarding optimal patch time allocation for female parasitoids foraging for hosts has arguably been the most studied problem in behavioral ecology over the past few decades (van Alphen et al. 2003, Wajnberg 2006). We have seen that the optimal time females should spend on each patch of hosts is given by the MVT (Charnov 1976), but we also saw that such a theoretical maximization model suffers from being based on several assumptions that are most likely unrealistic. Besides considering that female parasitoids are omniscient and should thus know the quality of all available patches in the environment, and on the mean time needed to reach them (see the GA we developed above to see how such an assumption can be addressed), this model also supposes: (i) that females never reach the end of their life (i.e. infinite time horizon); and (ii) that there is no competition because females are supposed to forage alone in each patch (Yamamura & Tsuji 1987).

The effect of relaxing the assumption that foraging parasitoid females are living during a never-ending period has been addressed theoretically using an SDP model by Wajnberg et al. (2006). They demonstrated that time-limited females should spend a longer time on each visited patch than what is predicted by Charnov's (1976) MVT and the deviation should increase when females are approaching the end of their life. The effect of competition was addressed using game theoretical models by Sjerps and Haccou (1982), Haccou et al. (1999), Haccou et al. (2003), and Hamelin et al. (2007a,b). These authors demonstrated that, in most cases, females are engaged in a war of attrition type when foraging simultaneously on a host patch (see also Chapter 9 by Haccou and van Alphen).

Building a theoretical deterministic model, while relaxing simultaneously both the assumptions of no time horizon and no competition, will obviously be a difficult task, if even feasible. Thus, we will consider here, as a simplified example, how such a problem could be approached by using a GA. As usual, the main difficulty in developing a GA is to find a way to encode the question to be solved in terms of a number of genes along chromosomes. For the sake of simplicity, we will consider that:

- 1 all patches in the environment are of the same quality;
- 2 the travel time between patches is fixed;
- 3 all foraging females live for the same finite time; and
- 4 each female parasitoid is foraging in the environment either alone or in the presence of a single competitor.

The genes on chromosomes should code for the behavior of the female in these two different situations. A possible, simplified way to code such a situation can be to have six

genes, three coding the phenotype of the female when it is alone, the other three coding the same phenotype when it is with a competitor. In each of these two cases, one gene is coding for the time the female will invest on each patch she has just encountered. A second gene is coding for a switching time after which the (additional) effect of a third gene, in terms of patch residence time, is taken into account. Doing so, the optimal increase in patch residence time, when a time-limited female is approaching the end of her life (Wajnberg et al. 2006), is explicitly considered. The six genes are coding for phenotypes on a continuous scale, with the only restriction being that all values should be positive (or null). So, the mutation of a gene will consist of replacing its value by a value randomly drawn from an exponential or a Weibull distribution (for patch times) or from a uniform distribution (for switching times). Moreover, the foraging female does not know what behavioral strategy its competitor (if any) – who is also a time-limited forager trying to maximize its own fitness – is adopting. She is also unaware of the probability that a competitor will be foraging simultaneously on the same patch. Thus, we are in the above-mentioned situation in which the fitness of each chromosome in the population cannot explicitly be computed, so Monte Carlo simulation should be used instead. For this, at each generation, the fitness of each chromosome will be estimated by its average fitness computed in several (say a hundred) simulated situations. For each of them, another chromosome is drawn randomly in the population and used to code for the behavioral strategy played by the competitor. Also, for each sequentially visited patch, the presence of the competitor is drawn randomly with a given, fixed probability, which should ideally be around 0.5 in order to have both parts of the chromosome evolving at the same speed.

Running such a GA will provide interesting information about both the optimal time females should invest on a patch, with or without the presence of a competitor, and when and what should be the optimal changes in patch residence upon approaching the end of their life, comparatively in the two situations.

17.5 A third example of parasitoid behavioral ecology: optimal mark persistence in a parasitoid threatened by hyperparasitoids

Many animals leave chemical information in the environment, either in order to communicate with conspecifics or to mark sites for future recognition. The latter occurs frequently in herbivorous and parasitoid insects that use small and discrete breeding sites for the production of offspring. While early investigators focused on the deterring effects that such substances might have on conspecifics searching for oviposition sites (Prokopy et al. 1977, Boller 1981, Prokopy 1981, McNeil & Quiring 1983, Lou & Wang 1985), an evolutionary approach determined that the evolutionary benefit of marking depends upon the fact that a female may encounter her own marking and might use it to avoid wasting eggs and time on resources that have already been exploited (Roitberg & Mangel 1988). Conspecifics encountering such a mark are informed about the current quality of the resource and may adaptively decide to superparasitize or to refrain from oviposition (Roitberg & Mangel 1988, van Alphen & Visser 1990, Nufio & Papaj 2001, Hoffmeister & Roitberg 2002). As long as the behavior of an animal in response to encountering a mark is in the interest of the marker, the marking behavior is of mutual benefit. However, marks, and especially external marks that are placed on rather than into the resource, are an example of public information (information derived from the behavior of other individuals,

see Danchin et al. 2004, Dall et al. 2005) that can be exploited by lower (Hoffmeister & Roitberg 1997) and higher trophic levels (Prokopy & Webster 1978, Roitberg & Lalonde 1991, Baur & Yeagan 1994, Hoffmeister & Gienapp 1999).

Thus, marking behavior does not necessarily lead only to fitness increments but may in turn incur costs, for example, through an increased mortality risk of the offspring developing in the marked substrate. Given that the marking pheromone can evolve toward higher or lower persistence, we may predict that the detectability of a marking pheromone will be related to the costs and benefits that a mark has for the marking individual. Unfortunately, the persistence of marking pheromones has been analyzed only in a small number of cases, yet variation in persistence may suggest that the residual activity of a marking pheromone is subject to evolutionary selection (Quiring & McNeil 1984, Averill & Prokopy 1987, Höller et al. 1991). This sets the stage for a game between trophic levels, since the cost of marking may depend on the frequency of marking individuals in the population and thus, the density of marked and non-marked resources an organism from a higher trophic level might exploit. Hoffmeister and Roitberg (1998) use such logic for an evolutionary game between a marking herbivore host and a mark-exploiting parasitoid of the herbivore offspring that they modeled using a GA. We can easily translate this problem into a GA for a game between a parasitoid marking the oviposition site externally and its hyperparasitoid exploiting the marks to find the primary parasitoid's offspring.

Let us assume that the offspring of a primary parasitoid is vulnerable against competition with conspecifics. The degree of vulnerability may change over time with older larvae being competitively superior to younger larvae. Thus, the mortality risk through a second oviposition into the host, resulting in competition, might be equal for both contestants (50%) right after the oviposition and show an exponential decay over time for the first larva afterwards. Consequently, the benefit of a marking pheromone to a female's offspring is highest directly after oviposition and decays over time. For the female herself, the benefit of the mark stems not only from the fact that she does not kill her own first offspring in sibling competition, but also from the fact that she can avoid ovipositing into previously used hosts, thereby exposing a second offspring to sibling competition. Thus, the female also benefits from the mark if her first offspring is no longer vulnerable to competition.

However, a female parasitoid might not stay at the same site for her entire life and, over time, the probability of revisiting a host she has parasitized some time ago will also follow an exponentially decreasing function. Consequently, the fitness return from a mark will be highest for a female directly after oviposition and decay over time. How steep the decay function will be depends on the fate of her offspring in competition and her spatial searching behavior. The costs of a marking pheromone might be twofold. First of all, they impose costs related to the production of the chemical, and chemical substances less vulnerable to oxidation might be more costly to synthesize than others. Second, costs do accrue from mark-exploiting hyperparasitoids, i.e. natural enemies of the parasitoid that search more intensely in areas where they detect marking pheromone of the parasitoid. Again, these costs depend on the window of vulnerability of the primary parasitoid's offspring. For a chemical substance, we may assume that the decay of the marking pheromone follows first-order kinetics. Thus, the detectability of the marking pheromone M_t at time t is a function of the initial detectability β at t_0 and the decay rate τ , thus:

$$M_t = \beta e^{-t/\tau} \quad (17.2)$$

To allow the analysis of the full state space of possible optimal decay rates of the marking pheromone, we define a gene with a binary character state for marking or non-marking and a second gene for the decay rate τ , with 30 discrete character states in value encoding.

Because the decision whether or not to accept a previously parasitized host for oviposition is a function of a parasitoid's life expectancy and egg load (van Alphen & Visser 1990, Minkenberg et al. 1992, Fletcher et al. 1994), we need two genes for a reaction norm of a parasitoid's tendency to accept previously parasitized and marked hosts as a function of the above-mentioned parameters. While the acceptance rate of marked hosts will generally increase with increased time and egg limitation, the gene may hold a variety of reaction norms that differ in the form (concave or convex) and the steepness of the function (Hoffmeister & Roitberg 1998). Since it is unknown to a parasitoid whether or not she or a conspecific will return to a marked site and at which physiological state, the number of possible combinations, and thus the search space for the benefit of marking is enormous. Further, the probability that a marked site is detected and exploited by hyperparasitoids will be dependent on the rate of marking and the detectability of marks in the parasitoid population. If few marks exist and the hyperparasitoid invests more searching time on marked compared to unmarked sites (Roitberg & Lalonde 1991, Hoffmeister & Gienapp 1999), hyperparasitoid attack will concentrate on marked sites, because more time will be invested in searching for rather than on sites. In contrast, if most genotypes in a population of parasitoids apply marks of high detectability, each marked site will benefit from a dilution effect. Consequently, the parasitoids in our scenario evolve in the frequency-dependent context of an evolutionary game between trophic levels and the search space for optimal solutions is too large to employ a deterministic fitness function. Since we need to use a stochastic fitness function, it will be necessary to employ a repeated evaluation GA, where averaging across 100 Monte Carlo repeats of the same generation leads to reliable fitness estimates for coded strategies.

In the simulation we may use a population of 100 parasitoids that, in the starting condition, represents a random draw from the population of possible genotypes. Each parasitoid has an egg load of 50 eggs and searches and oviposits in a habitat containing 4000, 8000, or 12,000 hosts, thus different scenarios cover a range of conditions from high to low competition for hosts. To evaluate the effect of marking costs we employ scenarios with no, intermediate, and high attack rate of hyperparasitoids, each run for 1000 generations and replicated with 5 different starting conditions, to see whether the GA converges repeatedly on the same strategy. Running such a GA elucidates how environmental characteristics shape the use of a chemical signal in a frequency and density dependent way and points toward important sources that select for the variability of marking systems across a range of ecological scenarios.

17.6 Conclusion

GAs are powerful tools to find optimal behavioral strategies for research questions, whose answers lie within a huge search space. While deterministic models are certainly the most elegant way to elucidate optimal solutions for theoretical approaches in behavioral ecology, their scope is limited to research questions tackling a small-sized search space of solutions. We feel that there is a large array of research questions in parasitoid behavioral

ecology where approaches would benefit from including state-dependence of optimal behavioral decisions, as well as the fact that many such decisions are made in a frequency-dependent game context within and between trophic levels. Naturally, both these mechanisms increase the search space of solutions considerably. When investigating optimal life history strategies in herbivores, Bouskila et al. (1998) found importantly different solutions when modeling a dynamic life history problem in a game context compared to a non-game context. However, few papers exist to date that used GAs to solve a research question in parasitoid behavioral ecology (Hoffmeister & Roitberg 1998, McGregor & Roitberg 2000). GAs are certainly not the only approach that can be used to find optimal behavioral solutions in a dynamic game context, but for research questions with a huge search space of possible solutions they represent an efficient means of finding good solutions. As has been described above and can be seen from the examples given here, GAs cannot only be applied to a variety of research questions, but there is also a variety of ways to employ a GA. Such variability in available methods and the fact that optimizing a GA is subject to fine tuning, often by trial and error, may account for the fact that GAs have been rarely used by parasitoid behavioral ecologists. Moreover, GA is simulation-based, and simulations tend to be viewed as less elegant compared to deterministic approaches. Yet, GAs have been shown to be a very powerful method to search for optimal solutions within enormous and often complicated search spaces. By running them repeatedly with different starting conditions, it is relatively easy to demonstrate whether or not the algorithm found the optimum of the search space. Thus, GAs offer a great potential to research questions in parasitoid behavioral ecology and, with an increased complexity of approaches used in experiment and theory, it is likely that they will become a widely used research tool for parasitoid behavioral ecologists.

References

- Aarsen, L.W. (1983) Ecological combining ability and competitive combining ability in plants: Toward a general theory of coexistence in systems of competition. *American Naturalist* **122**: 707–31.
- Averill, A.L. and Prokopy, R.J. (1987) Residual activity of oviposition-detering pheromone in *Rhagoletis pomonella* (Diptera: Tephritidae) and female response to infested fruit. *Journal of Chemical Ecology* **13**: 167–77.
- Barta, Z., Flynn, R. and Giraldeau, L.A. (1997) Geometry for a selfish foraging group: A genetic algorithm approach. *Proceedings of the Royal Society of London Series B Biological Science* **264**: 1233–8.
- Baur, M.E. and Yeorgan, K.V. (1994) Developmental stages and kairomones from the primary parasitoid *Cotesia marginiventris* (Hymenoptera: Braconidae) affect the response of the hyperparasitoid *Mesochorus discitergus* (Hymenoptera: Ichneumonidae) to parasitized caterpillars. *Annals of the Entomological Society of America* **87**: 954–61.
- Boller, E.F. (1981) Oviposition-detering pheromone of the European cherry fruit fly: Status of research and potential applications. In: Mitchell, E.R. (eds.) *Management of Insect Pests with Semiochemicals: Concepts and Practice*. Plenum Press, New York, pp. 457–62.
- Bouskila, A., Robinson, M.E., Roitberg, B.D. and Tenhumberg, B. (1998) Life-history decisions under predation risk: Importance of a game perspective. *Evolutionary Ecology* **12**: 701–15.
- Charnov, E.L. (1976) Optimal foraging: the marginal value theorem. *Theoretical Population Biology* **9**: 129–36.
- Clark, C.W. and Mangel, M. (2000) *Dynamic State Variable Models in Ecology – Methods and Applications*. Oxford University Press, New York.

- Dall, S.R.X., Lotem, A., Winkler, D.W. et al. (2005) Defining the concept of public information. *Science* **308**: 353–6.
- Danchin, E., Giraldeau, L.A., Valone, T.J. and Wagner, R.H. (2004) Public information: From noisy neighbors to cultural evolution. *Science* **305**: 487–91.
- Fletcher, J.P., Hughes, J.P. and Harvey, I.F. (1994) Life expectancy and egg load affect oviposition decisions of a solitary parasitoid. *Proceedings of the Royal Society of London Series B Biological Science* **258**: 163–7.
- Garey, M.R. and Johnson, D.S. (1979) *Computers and Intractability: a Guide to the Theory of NP-completeness*. W.H. Freeman, San Francisco.
- Giraldeau, L.A. and Caraco, T. (2000) *Social Foraging Theory*. Princeton University Press, Princeton.
- Godfray, H.C.J. (1994) *Parasitoids. Behavioral and Evolutionary Ecology*. Princeton University Press, Princeton.
- Goldberg, D.E. (1989) *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley, Reading.
- Haccou, P., Sjerps, M. and van der Meijden, E. (1999) To leave or to stay, that is the question: Predictions from models of patch leaving strategies. In: Olff, H., Brown, V.K. and Drent, R.H. (eds.) *Herbivores: Between Plants and Predators*. Blackwell Science, Oxford, pp. 85–108.
- Haccou, P., Glaizot, O. and Cannings, C. (2003) Patch leaving strategies and superparasitism: An asymmetric generalized war of attrition. *Journal of Theoretical Biology* **225**: 77–89.
- Hamelin, F., Bernard, P., Nain, P. and Wajnberg, É. (2007a) Foraging under competition: Evolutionarily stable patch-leaving strategies with random arrival times. 1. Scramble competition. In: Jørgensen, S., Quincampoix, M. and Vincent, T. (eds.) *Advances in Dynamic Games and Applications*, vol. 9, Annals of the International Society of Dynamic Games. Birkhäuser, Basel, pp. 327–48.
- Hamelin, F., Bernard, P., Shaiju, A.J. and Wajnberg, É. (2007b) Foraging under competition: Evolutionarily stable patch-leaving strategies with random arrival times. 1. Interference competition. In: Jørgensen, S., Quincampoix, M. and Vincent, T. (eds.) *Advances in Dynamic Games and Applications*, vol. 9, Annals of the International Society of Dynamic Games. Birkhäuser, Basel, pp. 349–65.
- Haykin, S. (1999) *Neural Networks: a Comprehensive Foundation*, 2nd edn. Prentice-Hall, Englewood Cliffs.
- Hoffmeister, T.S. and Roitberg, B.D. (1997) Counterespionage in an insect herbivore – parasitoid system. *Naturwissenschaften* **84**: 117–19.
- Hoffmeister, T.S. and Roitberg, B.D. (1998) Evolution of signal persistence under predator exploitation. *Ecoscience* **5**: 312–20.
- Hoffmeister, T.S. and Gienapp, P. (1999) Exploitation of the host's chemical communication in a parasitoid searching for concealed host larvae. *Ethology* **105**: 223–32.
- Hoffmeister, T.S. and Roitberg, B.D. (2002) Evolutionary ecology of oviposition marking pheromons. In: Hilker, M. and Meiners, T. (eds.) *Chemoecology of Insect Eggs and Egg Deposition*. Blackwell Wissenschaftsverlag, Berlin, pp. 319–47.
- Holland, J.H. (1975) *Adaptation in Natural and Artificial Systems*. The University of Michigan Press, Ann Arbor.
- Höller, C., Williams, H.J. and Vinson, S.B. (1991) Evidence for a two-component external marking pheromone system in an aphid hyperparasitoid. *Journal of Chemical Ecology* **17**: 1021–35.
- Holling, C.S. (1959) Some characteristics of simple types of predation and parasitism. *Canadian Entomologist* **91**: 385–98.
- Kumsawat, P., Attakitmongkol, K. and Srikaew, A. (2005) A new approach for optimization in image watermarking by using genetic algorithms. *IEEE Transactions on Signal Processing* **53**: 4707–19.
- Lou, H.Z. and Wang, X.S. (1985) An oviposition deterrent extracted from the larval frass of *Ostrinia furnacalis*. *Chinese Journal of Biological Control* **1**: 53.
- Maynard-Smith, J. (1982) *Evolution and the Theory of Games*. Cambridge University Press, Cambridge.

- McGregor, R.R. and Roitberg, B.D. (2000) Size-selective oviposition by parasitoids and the evolution of life-history timing in hosts: Fixed preferences vs frequency-dependent host selection. *Oikos* **89**: 305–12.
- McNamara, J.M. and Houston, A.I. (1985) Optimal foraging and learning. *Journal of Theoretical Biology* **117**: 231–49.
- McNeil, J.N. and Quiring, D.T. (1983) Evidence of an oviposition-detering pheromone in the alfalfa blotch leafminer, *Agromyza frontella* (Rondani) (Diptera: Agromyzidae). *Environmental Entomology* **12**: 990–2.
- Minkenberg, O.P.J.M., Tatar, M. and Rosenheim, J.A. (1992) Egg load as a major source of variability in insect foraging and oviposition behavior. *Oikos* **65**: 134–42.
- Nufio, C.R. and Papaj, D.R. (2001) Host marking behavior in phytophagous insects and parasitoids. *Entomologia Experimentalis et Applicata* **99**: 273–93.
- Obitko, M. and Slavík, P. (1999) Visualization of genetic algorithms in a learning environment. In: *Spring Conference on Computer Graphics, SCCG'99*. Comenius University, Bratislava, pp. 101–6.
- Perry, J.C. and Roitberg, B.D. (2005) Games among cannibals: Competition to cannibalize and parent-offspring conflict lead to increased sibling cannibalism. *Journal of Evolutionary Biology* **18**: 1523–33.
- Prokopy, R.J. (1981) Oviposition-detering pheromone system of apple maggot flies. In: Mitchell, E.R. (eds.) *Management of Insect Pests with Semiochemicals*. Plenum Publishing Corporation, New York, pp. 477–94.
- Prokopy, R.J. and Webster, R.P. (1978) Oviposition-detering pheromone in *Rhagoletis pomonella*: a kairomone for its parasitoid *Opius lectus*. *Journal of Chemical Ecology* **4**: 481–94.
- Prokopy, R.J., Greany, P.D. and Chambers, D.L. (1977) Oviposition-detering pheromone in *Anastrepha suspensa*. *Environmental Entomology* **6**: 463–5.
- Quiring, D.T. and McNeil, J.N. (1984). Intraspecific competition between different aged larvae of *Agromyza frontella* (Rondani) (Diptera: Agromyzidae): Advantages of an oviposition-detering pheromone. *Canadian Journal of Zoology* **62**: 2192–6.
- Reinhold, K., Kurtz, J. and Engvist, L. (2002) Cryptic male choice: Sperm allocation strategies when female quality varies. *Journal of Evolutionary Biology* **15**: 201–9.
- Roitberg, B.D. and Mangel, M. (1988) On the evolutionary ecology of marking pheromones. *Evolutionary Ecology* **2**: 289–315.
- Roitberg, B.D. and Lalonde, R.G. (1991) Host marking enhances parasitism risk for a fruit-infesting fly *Rhagoletis basiola*. *Oikos* **61**: 389–93.
- Rubinstein, R.Y. (1981) *Simulation and the Monte Carlo Method*. John Wiley & Sons, New York.
- Sjerps, M. and Haccou, P. (1982) Effect of competition on optimal patch leaving: A war of attrition. *Theoretical Population Biology* **46**: 300–18.
- Stephens, D.W. and Krebs, J.R. (1986) *Foraging Theory*. Princeton University Press, Princeton.
- van Alphen, J.J.M. and Visser, M.E. (1990) Superparasitism as an adaptive strategy for insect parasitoids. *Annual Review of Entomology* **35**: 59–79.
- van Alphen, J.J.M., Bernstein, C. and Driessen, G. (2003) Information acquisition and time allocation in insect parasitoids. *Trends in Ecology & Evolution* **18**: 81–7.
- Wajnberg, É. (2006) Time allocation strategies in insect parasitoids: From ultimate predictions to proximate behavioral mechanisms. *Behavioral Ecology & Sociobiology* **60**: 589–611.
- Wajnberg, É., Bernhard, P., Hamelin, F. and Boivin, G. (2006) Optimal patch time allocation for time-limited foragers. *Behavioral Ecology & Sociobiology* **60**: 1–10.
- Yamamura, N. and Tsuji, N. (1987) Optimal patch time allocation under exploitative competition. *American Naturalist* **129**: 553–67.