

# Speciation Dynamics: Generating Selective Pressure Towards Diversity

Heiko Hamann

Department of Computer Science,  
University of Paderborn, Paderborn, Germany  
heiko.hamann@uni-paderborn.de

July 4, 2013

## Abstract

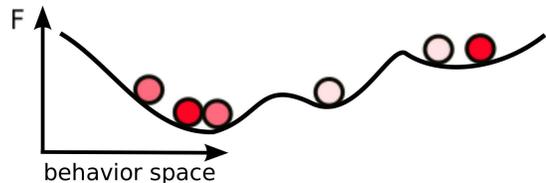
Recent approaches in evolutionary robotics (ER) propose to generate behavioral diversity in order to evolve desired behaviors more easily. These approaches require the definition of a behavioral distance which often includes task-specific features and hence a priori knowledge. Alternative methods, that do not explicitly force selective pressure towards diversity (SPTD) but still generate it, are known from the field of artificial life such as artificial ecologies (AE). In this study, we investigate how SPTD is generated without task-specific behavioral features or other forms of a priori knowledge and detect how methods of generating SPTD can be transferred from the domain of AE to ER. A promising finding is that in both types of systems, in systems from ER that generate behavioral diversity and also in the investigated speciation model, selective pressure is generated towards unpopulated regions of search space. We conclude by hypothesizing how knowledge about self-organizing SPTD in AE could be transferred to the domain of ER.

## 1 Introduction

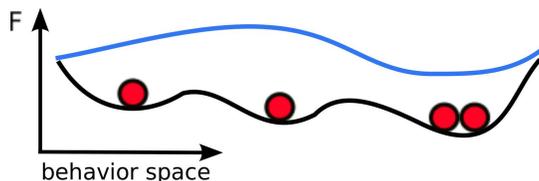
Methods of evolutionary computation have been successful as optimization technique for many years. Also the optimization of behaviors, which can justifiably be called ‘generation of behaviors’, in the field of ER (Nolfi and Floreano, 2000) has proven to be effective. However, the next step in this research towards more complex behaviors and tasks seems to be particularly difficult. Such a complex task could involve, for example, several successive sub-tasks whereas learned earlier sub-tasks have no utility before later sub-tasks are learned as well. The relative simplicity of investigated tasks in ER, especially when compared to natural systems, is, for example, discussed by Nelson et al. (2009).

Evolving robot behaviors becomes even more challenging if the necessary a priori knowledge is minimized which is necessary to achieve generally applicable approaches. Notably this concerns the fitness function and how elaborated it is. Nelson et al. (Nelson et al., 2009) define several fitness function classes such as the behavioral fitness functions which incorporate a lot of a priori knowledge (fitness function ‘selects for behavioral features of a presupposed solution to a given task’, i.e. *how* the task is accomplished) and the aggregate fitness functions which incorporate a very low degree of a priori knowledge (fitness function measures *what* the robot has accomplished and not *how* it was accomplished). Behavioral fitness functions are applied in ER because otherwise behaviors of certain complexities cannot be evolved with a reasonable commitment of resources.

An option is to increase the diversity in the population, e.g. by fitness sharing, see (Sareni and Krähenbühl, 1998). However, these methods include the measurement of a distance between



(a) novelty search, black line is fitness function based on novelty in the behaviors, behaviors of all generations are taken into account



(b) multi-objective behavioral diversity, black line is the fitness component based on behavioral diversity, blue line is the fitness component based on the actual objective function, behaviors of only the current generation are taken into account

Figure 1: Schematic representation of the fitness functions generated by novelty search (Lehman and Stanley, 2011) and MOBD (Mouret and Doncieux, 2009); circles represent known behaviors, selective pressure is towards bigger values of  $F$ .

genotypes which is computationally intractable for common encodings in ER, such as artificial neural networks (ANN) (Mouret and Doncieux, 2012). Instead promising recent results suggest to increase the behavioral diversity during the search or within the current population and measure the distances between behaviors which can be done efficiently. Examples are novelty search (Lehman and Stanley, 2011) and the approach by Mouret and Doncieux (Mouret and Doncieux, 2009, 2012). Novelty search (Lehman and Stanley, 2011) operates without an actual objective function. Instead selective pressure is generated towards behaviors that have not been seen before in the evolutionary run. A desired behavior is that that maximizes the behavioral distance to all known behaviors. For illustration and later use this is shown schematically in Fig. 1a. The circles represent behaviors that were found during the evolutionary run (graded colors represent different generations) and their position in search space. Behaviors that are close to known behaviors are undesirable which is represented by low values of the fitness function  $F$  around circles. Consequently selective pressure towards unpopulated regions in search space is generated which is represented accordingly by steep, upward slopes around circles. Multi-objective behavior diversity (MOBD, Mouret and Doncieux, 2009) includes the behavioral distance only as a component in the multi-objective fitness function. In contrast to novelty search it only accounts for the behaviors in the current population. This is sketched in Fig. 1b.

For approaches of behavioral diversity a distance between behaviors needs to be defined, which typically will be task-specific. Hence, this is a process similar to defining an appropriate fitness function for a given task in standard ER. It is argued that even naive behavioral distance definitions can improve the evolutionary process, such as definitions based on the final position of the robot or of movable objects at the end of the evaluation (Mouret and Doncieux, 2009; Lehman and Stanley, 2008). However, if we follow the analogy to fitness functions, it seems likely that for more complex tasks the measure of behavioral distance will also need to be more complex. This factual connection gets even more plausible if we reconsider the above mentioned comment by Nelson et al. (2009) about behavioral fitness functions (which select for behavioral features). In the design

of the behavioral distance measure not only *what* is accomplished is relevant but also *how* it is accomplished. In a simple exploration task this is not really apparent, if the behavioral distance is defined on the robot’s final position (Mouret and Doncieux, 2009), but it addresses at least the order of how the maze is explored. In a more complex task such as the locomotion of a biped robot (Lehman and Stanley, 2011) it gets more apparent as the behavioral distance is defined by the trajectory of the center of mass (how) instead of just the total offset (what). Hence we state the hypothesis that also methods based on behavioral distances will run into similar problems as seen in fitness function design such as having to define ‘task-specific hand-formulated functions that contain various types of selection metrics’ (Nelson et al., 2009) and consequently having to involve a high degree of a priori knowledge.

Based on our hypothesis we assume that neither genotypic distances nor behavioral distances are able to generate sustainable diversity in ER. What could be candidate solutions? Natural evolution represents a perfect standard for the generation of diversity. Particularly we are interested in evolutionary radiation which is an increase in taxonomic diversity. A typical example of a radiation is the Cambrian explosion. It generated a diversity that is comprehensible when looking at the corresponding phylogenetic tree which gets literally bushy within a comparatively short period of time. Each branching corresponds to an event of speciation. Hence, to understand diversity we want to understand speciation and the underlying process that generates SPTD. Our objectives are 1) to detect how SPTD is generated in a self-organizing system (e.g., speciation) and 2) how to transfer this knowledge to ER.

Next we summarize the knowledge on how speciation operates. Coyne and Orr (2004) ask: ‘Why are there species?’ and they comment: ‘we regard it as one of the most important unanswered questions in evolutionary biology’. Accordingly they do not answer but only discuss the question. They point to Maynard Smith and Szathmary (1998) who consider three explanations. One, species are discrete ‘stable states’ formed by a self-organizing system. However, this option lacks a mechanism that would explain the origin of species. Two, species fill discrete ecological niches. Third, reproductive isolation is an inevitable result of evolutionary divergence. The latter two are dependent on reproductive isolation and are not mutually exclusive (Coyne and Orr, 2004). A conclusive concept that connects both is that of ‘adaptive peaks’ (Dobzhansky, 1951). Species are adaptive peaks that are separated by ‘adaptive valleys’ which are genotypes that are unfit for survival. Whether these adaptive valleys are due to ecological or environmental effects is kept open. Another aspect is whether asexual or sexual reproduction is considered. While in the case of asexual reproduction ecological niches seem to be indispensable to create a reproductive barrier, in the case of sexual reproduction the reproductive barrier might be generated without explicit ecological niches especially because sexual selection could be effective. To summarize we note that evolutionary radiation and speciation are yet difficult to understand as also pointed out by Venditti et al. (2010): ‘Attempts to understand species-radiations [...] should look to the size of the catalogue of potential causes of speciation shared by a group of closely related organisms rather than to how those causes combine.’ The conclusion concerning inspirations for new methods to generate diversity in artificial evolution, especially ER, regrettably has to stay inconclusive for now due to the limited knowledge about the natural system itself. However, it is the starting point of the following investigations.

Instead of detecting methods to generate speciation, what are prerequisites for evolutionary radiation, that is, for active speciation? A prerequisite could be the existence of a complex environment that favors the formation of adaptive peaks and that initiates specialization in the organisms. Another prerequisite could be the existence of an ecology which creates complexity in the interaction of different organisms and species. The suitability of providing a complex environment within applications of ER is limited, notably if a complex environment is not part of the desired task. On the contrary, a considerable amount of research on the creation of AE in order to evolve behaviors has been reported. For example, the minimal ecology, that of just two species, is investigated in studies of coevolution (Floreano and Nolfi, 1997) and AE, possibly with many species, are popular in the field of artificial life (Ray, 1991). The studies of coevolution have an

emphasis on the actually evolved behaviors combined with considerations about their utility and complexity. It turns out that ‘the co-evolutionary process tends to fail into dynamical attractors in which the same solutions are adopted by both populations over and over’ (Nolfi and Floreano, 1998b). In AE studies the actual behaviors are of less interest, instead the evolutionary process as a whole is usually investigated in more detail. In addition it is also unclear how these AE would have to be designed to generate desired behaviors for a given task. Common to both is their high sensitivity to parameters and the challenge of creating actively progressing evolutionary processes. However, they are interesting examples of how SPTD is generated in a self-organizing process without incorporating a priori knowledge about how a certain behavior is accomplished. Summing up, the prerequisite to generate diversity is a minimum of complexity that provokes the emergence of adaptive peaks. The trigger could be ecological or environmental features but also effects of sexual selection. The goal of this study is 1) to detect and measure SPTD which is generated by a self-organizing process in AE, and 2) to determine how methods from the domain of AE could be transferred to ER. To the knowledge of the author, published studies on behavioral diversity tend to focus either on self-organizing diversity without aiming for the solution of a given task (e.g., Ray (1991)) or on explicitly imposed SPTD while searching for the solution of a particular task (e.g., Lehman and Stanley (2011)). In the following we investigate a speciation model of Woehrer et al. (2012) and an extension to it in order to investigate how SPTD is generated in this self-organizing system. The results are then compared to approaches that explicitly impose SPTD.

## 2 A model of speciation

One distinguishes several types of speciation, such as allopatric speciation and sympatric speciation. Previously it was thought that speciation happens mostly allopatrically, that is, by spatial separation of populations which then develop reproductive isolation. More recent results (Coyne and Orr, 2004) support that sympatric speciation, which is speciation within the same geographic region, might be more common than expected. Here, sympatric speciation is of more interest because it is a self-organizing, evolutionary process while allopatric speciation occurs due to external forces (arguably except for migration). In terms of the application of speciation to increase the behavioral diversity in ER, sympatric speciation is preferred as it does not need a priori knowledge while allopatric speciation would need an implementation of a cause.

Woehrer et al. (2012) report an artificial life model of sympatric speciation based on sexual selection, in particular assortative mating which is a mating pattern where mating between individuals with similar genotypes or phenotypes is more likely. The model is inspired by the natural system of finches on the Galápagos Islands. Although an island setting might let allopatric speciation appear as a good and exclusive explanation this does not seem to be the case on the Galápagos Islands (Woehrer et al., 2012). Woehrer et al. (2012) point to the specialty of the proposed system that combines natural selection and sexual selection acting on the same trait which is directly related to so-called ‘magic traits’ (Servedio et al., 2011). Here we reproduce their results, report an extension of the model, and perform additional measurements in simulations.

The artificial system models a bird population of dynamic size. A bird is modeled by age, beak size, energy level, and gender. To survive the birds have to forage (implemented as random search) for seeds. These are modeled by energy, location, uniformly distributed size, and distributed in discrete space of size 100 units  $\times$  100 units (see table 1 for used parameter settings). Selective pressure is imposed by the limited resource of seeds. Initially of the dry season a number of seeds is placed in the world which decreases consecutively over a period of 30 or 43 simulated days (depending on which setting is used) as the birds forage from it. The birds’ search for seeds is limited by their beak size  $s$  because they can only feed on seeds of size  $[s - 1, s + 1]$ . The search costs energy on each day and fed seeds add energy to the bird’s energy. If a bird runs out of energy during the dry season it dies. Those that survive may attempt to reproduce. Reproduction is based on sexual selection and assortative mating (with random mating no speciation was observed (Woehrer

parameter	FAM	EAM
max. age	4 years	4 years
max. energy	2.0	2.0
max. energy per seed	2.0	2.0
search energy cost	0.1	0.1
max. num. male mating	5	5
max. num. generations	1000	1000
initial number of birds	400	150
initial beak size mean	5.5	5.5
initial beak size variance	0.5	3.5
beak size interval	[1, 10]	[1, 10]
dry season length	43 days	30 days
initial number of seeds	5000	6300
feeding square size	10	10
world size	100	100
variance of offspring prop.	0.03	0.01
assortative range $\Delta$	0.5	[0.01, 10]
seed width $W$	n.a.	[0.01, 10]

Table 1: Parameter settings for fixed assortative mating (FAM) and for evolved assortative mating (EAM); some parameters differ from (Woehrer et al., 2012).

et al., 2012)). Females with beak size  $s_f$  select only mates with beak size  $s_m$  within the female’s beak size interval  $s_m \in [s_f - \Delta, s_f + \Delta]$  for assortative range fixed to  $\Delta = 0.5$ . The offspring has a beak size averaged over its parents plus Gaussian noise and random gender. For all remaining details please see (Woehrer et al., 2012) and table 1.

A typical run is shown in Fig. 2a which is a plot of all beak sizes that occur in the population over 375 generations. The resemblance to a phylogenetic tree is obvious. Also the drift of species, branching into two species, and the extinction of species can be noticed. Hence the model of Woehrer et al. (2012) is a simple model of self-organized speciation and can be used as an easy-to-handle analogy to the studies on behavioral diversity in ER. The interval of allowed beak sizes [1, 10] is the equivalence to the behavior space and a bird’s beak size would be the 1-d equivalence to the behavior defined by ANN. The extreme difference between high-dimensional ANN and the simplistic 1-d beak size interval is not a limiting factor of this analogy because the speciation model possesses the one qualitative feature that is relevant for this study, namely self-organized generation of diversity. The assortative mating corresponds to allowing recombination only for ANN that share a ‘magic trait’ which is defined by sexual selection and could be a behavioral feature. In addition, this way the emergence of speciation relies crucially on a pre-defined parameter (assortative range  $\Delta$ ). In order to avoid such a pre-defined measure for the above mentioned reasons we can allow the evolutionary algorithm to vary features of sexual selection. The perfect solution would be to evolve the full process of sexual selection which is, however, beyond the scope of this paper. Instead we restrict the following investigations to the evolution of the allowed difference between females’ and males’ beak size defined by the assortative range  $\Delta$  in this particular case of assortative mating. While this solves the problem of having pre-defined parameters, this would still correspond to a pre-defined process of sexual selection. That way we are able to investigate whether the evolutionary process can self-organize towards a higher degree of diversity without being forced to do so by a parameter setting.

It turns out that the model is very sensitive to settings of the assortative range  $\Delta$  which is typical for such systems as pointed out above (Sec. 1). If  $\Delta$  is set too low, species extend over a narrow interval of beak sizes, feed from a small set of seeds that they are able to eat, and

become extinct often (data not shown). If  $\Delta$  is set too big, no speciation is observed because one big connected component of birds in ‘beak size space’ emerges. Still we proceed and allow the assortative range  $\Delta$  to be evolved. The extension of the above model is described in the following. The parameter  $\Delta$  that defines sexual selection by setting the beak size range is defined now as an individual bird property. It is passed on by an average over the parents plus Gaussian noise (see table 1 for parameters). Furthermore, we introduce an evolved parameter of each individual bird called ‘seed width’  $W$  to increase the attractiveness of being a specialist. In addition to the beak size it also determines the interval of seed sizes a bird is able to feed on ( $[s - W, s + W]$ , priority is with the more restrictive interval) and the energy of a seed is scaled by  $1/W^2$  (for  $W < 1$  a seed’s energy is increased quadratically in  $W$ ). The system shows also speciation without this additional feature of seed width  $W$  but the inclusion of seed width stimulates speciation (data not shown).

For the following presentation and analysis of our results we define a technical concept of species in this simple model. In order to do so we interpret the distribution of a population’s beak sizes as a graph whereas each bird’s beak size represents a node and two such nodes are connected to each other if they are within each others individual assortative range of sexual selection defined by  $\Delta$ . A species is defined by the graph-theoretic concept of connected components: a subset of nodes and for each pair of such nodes there is a path connecting both. With this definition we are able to implement an automatic classifier that determines where and how many species exist in a given configuration.

### 3 Results

In the following we investigate whether speciation is observed, we investigate the distribution of beak sizes over time to find what is beneficial for speciation in this extended model, the distribution of branch lengths, and the dynamics of species. With this extended model we obtain for different random initializations results (for parameters see Tab. 1) that are characterized by three classes: showing speciation (see Fig. 2b), intermediate (see Fig. 2c), and not showing speciation (see Fig. 2d). Following our connected-component definition, species within limited beak-size intervals and clear separation are noticed in Fig. 2b. In Fig. 2c species are not clearly separated at all times. For example, at  $t \approx 600$  one species spans almost over the whole beak-size interval. In Fig. 2d, a single species covers the whole interval. Hence this system does not reliably self-organize towards diversity, at least for the tested parameters.

For the following investigations we classify occurrent configurations. It turns out that four configurations are frequent: only species in the left part of the beak size interval ( $s < 4.5$ , called *left*, frequency: 3.24%), only species in the right part ( $s > 6.5$ , called *right*, frequency: 3.14%), species distributed over the whole interval (called *all-over*, frequency: 88.37%), and species in the two outer parts but not in the middle (called *symmetrical*, frequency: 5.24%). Concerning the evolved assortative range  $\Delta$ , the all-over configuration is distinguishable from the three other configurations. In Fig. 3a the distributions of all occurring assortative ranges over a number of evolutionary runs are compared. The mean for all-over configurations is 1.9 and bigger than those of the others (about 1.3). With bigger assortative range a species spans over big intervals more easily. Consequently big assortative ranges counter diversity in terms of speciation. The populations that spread over the whole interval are bigger than those showing diversity (mean of about 200 birds compared to about 100) because they exploit the energy provided by seeds fully. Consequently they are less prone to fluctuations and have a smaller risk of extinction. At the same time they make sure to exploit seeds of all sizes. Hence the low-diversity solution actually seems to be the evolutionary more robust approach which raises the question of how speciation could be stimulated additionally (an optional target of an investigation which is beyond this paper would be the tradeoff between a generalist’s advantages and costs due to seed width  $W$ ). In turn it is possible to force the system into speciation by forcing bi- or multi-modal distributions of seed sizes (Woehrer et al., 2012). However, this is tweaking the environment which is not a good option

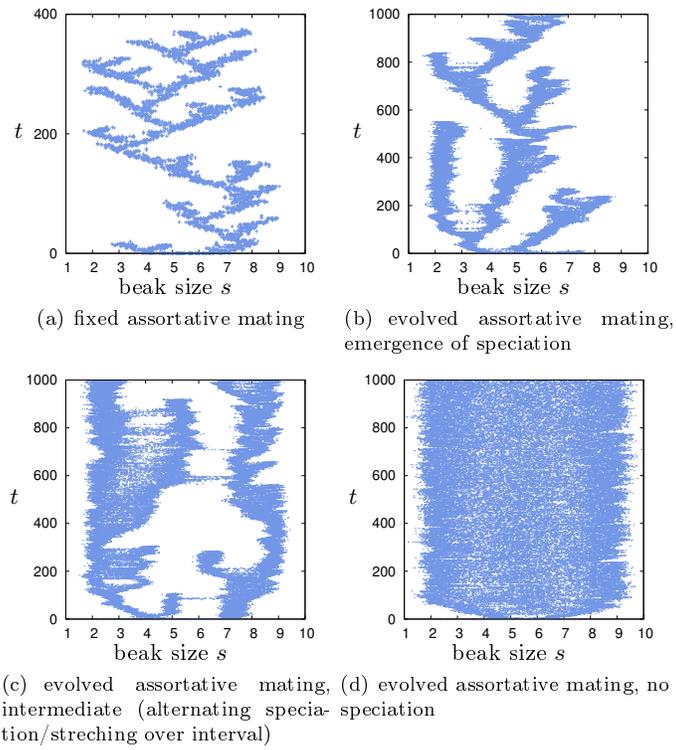


Figure 2: Beak sizes  $s$  over generations  $t$  for fixed and evolved assortative mating; for parameters see Tab. 1; sub-figures b-d differ only in random initialization.

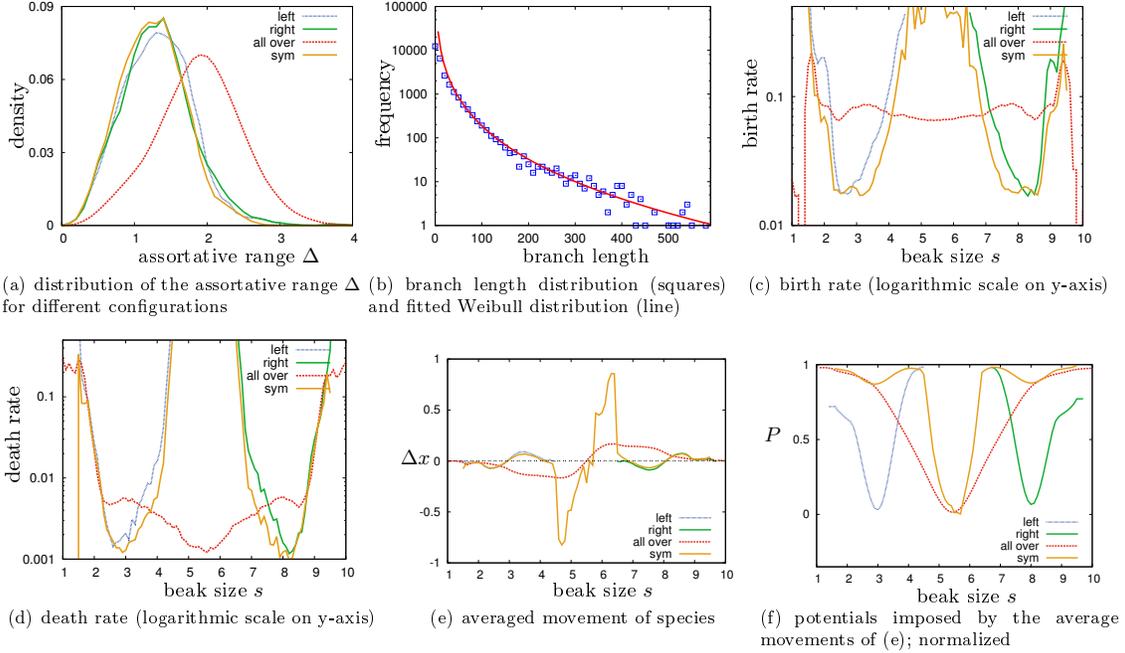


Figure 3: Distribution of the assortative range  $\Delta$  for different configurations and the distribution of branch lengths in the phylogenetic tree with fitted Weibull distribution; birth and death rate of species over beak size for different configurations; averaged movement of species and implied potentials over beak size for different configurations.

for our application in ER.

Generally, populations with big assortative ranges seem to be more stable and there is also a tendency of spreading over the whole interval with increasing time. We are able to support this claim by investigating the distribution of branch lengths in the phylogenetic trees. The branch length is the time period between split-ups of species and/or their extinction. Based on an automatic check for connected components, species and the corresponding branch lengths are determined in the implementation of the model. The distribution of branch lengths based on independent evolutionary runs ( $4.4 \times 10^5$  samples) is shown in Fig. 3b (squares). In the analysis of this data we follow Venditti et al. (Venditti et al., 2010) who discuss and interpret different branch length distributions of phylogenetic trees. While they apply their methods to natural systems we found that they can also be applied to this artificial system. The best fit we found is the Weibull distribution ( $ca/bx^{(a-1)} \exp(-(x/b)^a)$ ,  $a = 0.32$ ,  $b = 0.065$ ,  $c = 29 \times 10^3$ ) which Venditti et al. (2010) interpret in the following way: “the Weibull density can accommodate the probability of speciation changing according to the amount of divergence from the ancestral species. This model will fit the data if, for example, species are either more or less likely to speciate the older they get.” This supports our finding that old species tend to speciate less as they tend to span the whole interval.

Despite these shortcomings in the sustainability of the evolution of species we are able to investigate the dynamics of speciation in this system. Our aim is to detect and measure how SPTD is generated without explicitly forcing it. The evolution of species in space and time based on the above connected-component definition can be interpreted as a discrete birth–death process combined with drifting motion. For the above defined four classes of configurations the birth and death rate depending on the beak size were measured, see Fig. 3c and d (qualitative, no errors shown). The birth rate is bigger by about an order of magnitude because in our implementation the merger of species was not classified as death; only actual extinction was classified as death.

Interestingly the birth rate for the all-over configuration is almost homogeneous while the other configurations have a dip at  $s \approx 3$  and  $s \approx 8$ . This is most likely because species in the all-over configuration are in average almost evenly distributed while in the other configurations species are more likely to be positioned at  $s \approx 3$  and/or  $s \approx 8$ . Once this ‘niche’ is covered, a birth at the same position is not possible. The death rates of all configurations have peaks at the bounds which is explained by the smaller number of available seeds once birds cover also seed sizes that do not occur ( $< 1$  or  $> 10$ ). Interestingly beak sizes keeping a distance to the bounds support survival for all configurations.

In the following we want to measure the dynamics of species, that is, the average movement of species in beak size space in certain configurations. A measurement of how the species drift for the four configurations is shown in Fig. 3e.  $\Delta x$  gives the average displacement of a species from one to the next generation for  $1.9 \times 10^7$  samples. Positive values  $\Delta x > 0$  describe the drift of a species towards bigger beak sizes and negative values  $\Delta x < 0$  describe drift towards smaller beak sizes. For the configurations *left*, *right*, and *all-over* the average motion of species indicates spreading and species keep moving towards the bounds even when already approaching them which means they move into regions of high death rates (Fig. 3d). The configuration *symmetrical* is different because it has a stabilizing effect at  $s \approx 4$  and  $s \approx 6.5$  which are maxima. This is, however, a direct effect of the classification and due to situations when an *all-over* configuration turns into a *symmetrical* configuration. Still it is valid data within our classification scheme. To draw a direct connection to applications in ER we determine the potentials that are implied by the average movement of species (for example, similar to gravitational potentials) which is here merely an integration over the average movement. Fig. 3f shows the 4 potentials of the 4 configurations (normalized to similar scales). These potentials are important findings for this study because they are emergent fitness functions with selective pressure towards bigger values in the same way as in the above schematic representations of the fitness functions for novelty search and MOBD (see Fig. 1). Similarly also here currently populated regions in beak-size space are less desirable and there is pressure towards unpopulated regions. For example, the potential for configuration *left* has a minimum at  $s = 3$  which corresponds to the typically populated position in this configuration as determined by low birth and death rates around  $s = 3$  (Fig. 3c and d) and an average motion of  $\Delta x = 0$  (Fig. 3e). These potentials are the confirmation that this self-organizing system certainly generates SPTD. While, for example, in novelty search that pressure is explicitly enforced by pushing towards behavioral diversity, here the selective pressure is a feature of the ecological system. Unpopulated regions in beak-size space correspond to big resources of energy in the form of seeds no one forages for. Birds that manage to push into these regions compete with few fellows, gather plenty of energy, and increase their fitness for survival. This analogy shows on the one hand that approaches based on behavioral distance should not just be considered as engineered abstractions but rather as bio-inspired approaches that have a direct connection to ecological features of speciation. On the other hand it shows an option of how to generate a self-organizing SPTD as discussed in the following.

## 4 Discussion and Conclusion

In this paper we have detected self-organizing generation of SPTD in an AE and we have motivated the need for methods that generate diversity, particularly behavioral diversity in the field of ER. Standard methods to increase the diversity in a population but also recent methods based on behavioral distances include a considerable amount of a priori knowledge because they (tend to) include behavioral features in the definition of the behavioral distance that measure *how* the robot accomplishes a task. Mouret and Doncieux (2012) discuss the analogy between fitness function design and behavioral distance measure design: “More importantly, novelty search critically depends on a good behavior characterization to create a gradient. Researchers in ER used to craft fitness function to create a perfect fitness gradient; novelty search users have to craft the behav-

ior distance to create a similar gradient. This last option may be easier for some problems but eventually some distances will be hard to define.” Hence, even with methods not relying on fitness functions the former problem of including a priori knowledge about the task persists. Mouret and Doncieux (Mouret and Doncieux, 2012) discuss how the ideal setting for ER might look like: “In an ideal ER setup, ER researchers would only define a high-level fitness function and let the generic evolutionary process do the rest. This goal could be achieved with a generic behavioral distance function that could be used with most ER tasks while still improving the evolutionary process.” Whether such a generic measure of behavioral distances can be found is unknown. Instead we study diversity in natural evolution that is generated in a self-organizing process and our focus is on diversity by speciation. Due to limited biological knowledge about why species exist it is also difficult to determine prerequisites for the generation of diversity in artificial systems. Therefore we investigated the simple speciation model by Woehrer et al. (2012) and extended it to allow the evolution of features that define the process of sexual selection based on assortative mating. Our findings indicate that self-organizing speciation is achievable but the system is sensitive to parameter settings as known from other AE.

It seems difficult to provoke evolutionary dynamics that favor diversity over uniform solutions without imposing diversity by predefined environmental features (e.g., multi-modal distributions of seeds, see (Woehrer et al., 2012)). In fact it seems questionable whether sustainable generation of diversity is possible without stimulating influences from a dynamic environment or a complex ecology with intensive inter-species interaction. A beneficial finding is the analogy between methods based on behavioral distances in ER and the investigated self-organizing ecology in terms of the selective pressure that is generated. In both systems selective pressure is generated towards unpopulated regions of the search space (cf. Figs. 1 and 3f). Hence we follow that both systems relate to each other and that there might be a way of transferring the methods of generating selective pressure in AE to ER.

In the speciation model this selective pressure is an ecological feature because unpopulated regions of search space hold plenty of seeds which increase the fitness for survival once they are foraged. While this is easily implemented in this speciation model by defining a beak-size search space, it is unknown how distributing seeds in search space transfers as analogy to the behavior space of ER. This triggers an important research question: How to define AE in the context of ER that generate SPTD without explicitly addressing particular task-specific behavioral features and without including a priori knowledge. This could be done by a data structure that covers the full behavior space and keeps track of which regions of this abstract space have been visited (following the analogy: regions where most of the seeds have been eaten up). However, this would be most likely intractable and probably also task-specific. Solutions for this dimensionality reduction problem based on heuristics would possibly have similar difficulties as, for example, function approximation in reinforcement learning (Sutton and Barto, 1998). Presumably any method that maps behaviors from the actual behavior space into a smaller feature space will suffer from being either task-specific or of limited benefit although very simple mappings were shown to be beneficial in simple tasks (Mouret and Doncieux, 2009; Lehman and Stanley, 2008). Seemingly the statistics about the frequencies of behaviors is embedded into the environment in natural systems similarly to the concept of stigmergy in swarm intelligence. A candidate solution would be to evolve behaviors in an embodied system which allows for embedding behavioral statistics in the environment. Fortunately, an embodied approach is feasible in ER (Bredeche et al., 2012; Stradner et al., 2012).

Furthermore we have reported measurements (Figs. 3c-f) that allow for modeling speciation as a discrete birth–death process combined with a spatial feature determined by drifting motion towards unpopulated regions. That way speciation truly is linked “to rare stochastic events that cause reproductive isolation” (Venditti et al., 2010). In addition we have shown that the analysis of branch length distributions by Venditti et al. (2010) is also applicable in artificial systems and might prove to be instrumental in classifying artificial systems of speciation. In future work we plan to continue investigations of how to integrate AE into behavior space that is not task-specific

and that generates SPTD in complex tasks for ER. In addition it might be desirable to investigate also models of speciation that evolve themselves features of sexual selection that generate diversity.

## References

- Bredeche, N., Montanier, J.-M., Liu, W., and Winfield, A. F. (2012). Environment-driven distributed evolutionary adaptation in a population of autonomous robotic agents. *Mathematical and Computer Modelling of Dynamical Systems*, 18(1):101–129.
- Coyne, J. A. and Orr, H. A. (2004). *Speciation*. Sinauer Associates.
- Dobzhansky, T. (1951). *Genetics and the origin of species*. Columbia University Press.
- Floreano, D. and Nolfi, S. (1997). Adaptive behavior in competing co-evolving species. In *Proceedings of the fourth European Conference on Artificial Life*, pages 378–387. MIT Press.
- Lehman, J. and Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In Bullock, S., Noble, J., Watson, R., and Bedau, M. A., editors, *Artificial Life XI: Proc. of the 11th Intern. Conf. on the Simulation and Synthesis of Living Systems*, pages 329–336. MIT Press.
- Lehman, J. and Stanley, K. O. (2011). Improving evolvability through novelty search and self-adaptation. In *Proceedings of the 2011 IEEE Congress on Evolutionary Computation (CEC'11)*, pages 2693–2700. IEEE.
- Mouret, J.-B. and Doncieux, S. (2009). Using behavioral exploration objectives to solve deceptive problems in neuro-evolution. In *Proceedings of the 11th annual conference on Genetic and Evolutionary Computation (GECCO'09)*, pages 627–634. ACM.
- Mouret, J.-B. and Doncieux, S. (2012). Encouraging behavioral diversity in evolutionary robotics: an empirical study. *Evolutionary Computation*, 20(1):91–133.
- Nelson, A. L., Barlow, G. J., and Doitsidis, L. (2009). Fitness functions in evolutionary robotics: A survey and analysis. *Robotics and Autonomous Syst.*, 57:345–370.
- Nolfi, S. and Floreano, D. (1998b). How co-evolution can enhance the adaptive power of artificial evolution: Implications for evolutionary robotics. In Husbands, P. and Meyer, J.-A., editors, *Evolutionary Robotics: 1st European Workshop*, vol. 1468 of *LNCIS*, pages 22–38.
- Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press.
- Ray, T. S. (1991). Evolution and optimization of digital organisms. In *Proc. of the 1990 IBM Supercomputing Competition*, pages 489–531. The Baldwin Press.
- Sareni, B. and Krähenbühl, L. (1998). Fitness sharing and niching methods revisited. *IEEE Transactions on Evolutionary Computation*, 2(3):97–106.
- Servedio, M. R., Doorn, G. S. V., Kopp, M., Frame, A. M., and Nosil, P. (2011). Magic traits in speciation: ‘magic’ but not rare? *Trends in Ecology and Evolution*, 26(8):389–397.
- Maynard Smith, J. and Szathmáry, E. (1998). *The Major Transitions in Evolution*. Oxford University Press.
- Stradner, J., Hamann, H., Zahadat, P., Schmickl, T., and Crailsheim, K. (2012). On-line, on-board evolution of reaction-diffusion control for self-adaptation. In Adami, C., Bryson, D. M., Ofria, C., and Pennock, R. T., editors, *Alife XIII*, pages 597–598. MIT Press.
- Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.
- Venditti, C., Meade, A., and Pagel, M. (2010). Phylogenies reveal new interpretation of speciation and the Red Queen. *Nature*, 463:349–352.
- Woehrer, M., Hougen, D., and Schlupp, I. (2012). Sexual selection, resource distribution, and population size in synthetic sympatric speciation. In Adami, C., Bryson, D. M., Ofria, C., and Pennock, R. T., editors, *Proc. of the 13th Intern. Conf. on the Simulation and Synthesis of Living Syst. (Alife'13)*, pages 137–144. MIT Press.