Dependence of Adaptability on Environmental Structure in a Simple Evolutionary Model

Jeffrey A. Fletcher

Systems Science Ph.D. Program, Portland State University

P.O. Box 751, Portland, OR 97207-0751, USA

Email: jeff@sysc.pdx.edu

Martin Zwick

Systems Science Ph.D. Program, Portland State University

P.O. Box 751, Portland, OR 97207-0751, USA

Email: zwick@sysc.pdx.edu

Mark A. Bedau

Department of Philosophy, Reed College

3213 SE Woodstock Blvd., Portland, OR 97202, USA

phone: (503) 771-1112, ext. 7337; fax: (503) 777-7769

Email: mab@reed.edu

Abstract

This paper concerns the relationship between the detectable and useful structure in an environment and the degree to which a population can adapt to that environment. We explore the hypothesis that adaptability will depend unimodally on environmental variety, and we measure this component of environmental structure using the information-theoretic uncertainty (Shannon entropy) of detectable environmental conditions. We define adaptability as the degree to which a certain kind of population successfully adapts to a certain kind of environment, and we measure adaptability by comparing a population's size to the size of a non-adapting, but otherwise comparable, population in the same environment. We study the relationship between adaptability and environmental structure in an evolving artificial population of sensorimotor agents that live, reproduce, and die in a variety of environments. We find that adaptability does not show a unimodal dependence on environmental variety alone, although there is justification for preserving our unimodal hypothesis if we consider other aspects of environmental structure. In particular, adaptability depends not just on how much structural information is detectable in the environment, but also on how unambiguous and valuable this information is, i.e., whether the information accurately signals a difference that makes a difference. How best to measure and integrate these other components of environmental structure remains unresolved.

Keywords: adaptation, environment, environmental structure, evolution, sensorimotor function, Shannon entropy.

1 How Does Adaptability Depend on Environmental Structure?

An evolving system consists of a population of agents adapting their behavior to an environment through the process of natural selection. The difficulty of the adaptive challenge obviously depends upon the population, the environment, and the interaction between the two. In this paper, we adopt an environment-centered view, that is, we examine how environments vary in the adaptive challenge which they present. This orientation reflects a kind of figure/ground reversal. One often takes the environment as ground and the adapting population as figure. That is, one treats the adaptive challenge as fixed and examines the resulting dynamics of adaptation, perhaps as a function of different adaptive capabilities of the population. Here, we treat the population as relatively given and study how varying the environment affects the difficulty of the adaptive task to be solved. This reversal of focus is found in some other recent studies (e.g., Wilson, 1991; Littman, 1993; Todd and Wilson, 1993; Todd, Wilson, Somayaji, and Yanco, 1994) and it recalls the earlier work of Emery and Trist (1965) on the causal texture of environments of social organizations.

Recent studies tend to pursue one of two projects: either (i) providing an abstract categorization of environments, or (ii) gathering experimental evidence about how artificial agents actually adapt in different simulated environments. Wilson (1991) and Littman (1993) follow the first project, with Wilson focusing on the degree of non-determinism in an environment and Littman characterizing the simplest agent that could optimally exploit an environment. But neither Wilson nor Littman address the degree to which a given (possibly suboptimal) agent could adapt to a given environment. Experimental investigation of agents adapting in different environments—project (ii) above—is the focus of Todd and Wilson (1993) and of Todd, Wilson, Somayaji, and Yanco (1994). Todd and Wilson introduce an experimental framework for investigating how adaptation varies in response to different kinds of environments, and Todd et al. demonstrate different adaptations in different kinds of environments. In neither case, though, is environmental structure actually classified or measured. Our work pursues both projects (i) and (ii) simultaneously; we experimentally study how the adaptability of given (possibly suboptimal) agents varies in response to environmental structure. Since our characterization of environmental structure is quantitative, we can seek evidence for general laws relating adaptability and environmental structure.

The guiding idea behind our work is the hypothesis that a population's ability to adapt to an environment depends roughly unimodally on the environment's detectable and useful structure. If the environment is too simple because it does not present the population with enough of the kind of information which adaptation can exploit, then adaptation will be difficult. On the other hand, if

the environment is too complex because it swamps the population with too much information, then adaptation will again be difficult. Adaptability would seem to be maximized somewhere between these extremes. This hypothesized dependence of adaptability on environmental structure should apply to both artificial and natural systems. In this paper we explore this hypothesis in a simple artificial evolving system. This makes it comparatively easy to tease apart the relevant issues; it also provides a baseline against which more complex systems can be compared and understood.

The population in our model consists of sensorimotor agents. Each agent responds to limited sensory input from the environment with a single behavioral output specified by the agent's genome. The adaptive task consists of finding an output to associate with each possible input. The difficulty of the adaptive task, therefore, would seem to involve at least the following aspects of the environment:

- the quantity of sensory information, i.e., the variety of sensed environmental conditions with which behaviors must be associated;

- the ambiguity of the information, i.e., the degree to which sensory input accurately represents the objective environment;

- the value of the information, i.e., the benefit of adaptive behaviors over non-adapted behaviors.

In terms of these components, an adaptive task is difficult if the environment sends many messages requiring an adaptive response, if the messages from the environment are ambiguous, or if they have little value.

Our agents' sensory input reflects the environment's local structure, so the first component listed above reflects the agent's perception of the environment's structural variety. This is a salient feature of the environmental structure and the adaptive challenge it presents, and it is the essence of Ashby's (1956) conceptualization of adaptation, according to which environmental variety poses a problem to which behavioral variety is the response. We thus begin our analysis of environmental structure here, seeking initially to ascertain to what extent this factor alone determines the difficulty of the adaptive task. We find that, although the environment's structural variety is indeed an important component of structure relevant to adaptation, it does not characterize it entirely. We speculate that, at least in part, this is because it omits the roles played by ambiguity and value.

2   A Model of Adaptation in Diverse Environments

All of our empirical observations are from computer simulations of adaptation in environments with different kinds of structure. Our model consists of many agents that exist sensing their local

4

environment, moving as a function of what they sense, and ingesting what resources they can find.

2.1 Agent and Environment Interactions

The world is a grid of 128×128 sites with periodic boundary conditions, i.e., a toroidal lattice. All that exists in the world besides the agents is a resource field, which is spread over the lattice of sites and is replenished from an external source. The resource level at a given site is set at a value chosen from the interval $[0-R]$, where R is the maximum resource level (chosen here arbitrarily as 255). These models are a modification of those previously studied by Bedau and Packard (1992), Bedau, Ronneburg and Zwick (1992), Bedau (1994), Bedau and Bahm (1994), Bedau (1995), and Bedau, Giger and Zwick (1995). All of these models are extensions of one originally proposed by Packard (1989). In the framework of Emery and Trist (1965), our model is a type-II ("placid, clustered") rather than type-III ("disturbed, reactive") environment, because the principal consideration is location rather than response to the behaviors or possible behaviors of other agents.

Here we consider only *static* resource fields, i.e., fields in which resources are immediately replenished whenever they are consumed, so that the spatiotemporal resource distribution is constant. In static resource models the population has no effect on the distribution of resources. Nevertheless, since the agents constantly extract resources and expend them by living and reproducing, the agents function as the system's resource sinks and the whole system is dissipative.

Adaptation is resource driven since the agents need a steady supply of resources in order to survive and reproduce. Agents interact with the resource field at each time step by ingesting all of the resources (if any) found at their current location and storing it in their internal resource reservoir. Agents must continually replenish this reservoir to survive for they are assessed a constant resource tax at each time step. If an agent's internal resource supply drops to zero, it dies and disappears from the world. As a practical expedient for speeding up the simulation, each agent also runs a small risk, proportional to population size, of randomly dying.

Each agent moves each time step as dictated by its genetically encoded sensorimotor map: a table of behavior rules of the form IF (environment j sensed) THEN (do behavior k). Only one agent can reside at a given site at a given time, so an agent randomly walks to the first unoccupied site near its destination if its sensorimotor map sends it to a site which is already occupied. (Population sizes range from about 2% to 10% of the number of sites in the world, so at the larger population sizes these collisions will occur with a non-negligible frequency.) An agent receives sensory information about the resources (but not the other agents) in the von Neumann neighborhood of five sites centered on its present location in the lattice. An agent can discriminate only four resource levels (evenly distributed over the $[0-R]$ range of objective resource levels) at

each site in its von Neumann neighborhood. Thus, each sensory state j corresponds to one of 4^5 = 1024 different detectable local environments. Each behavior k is a jump vector between zero and fifteen sites in any one of the eight compass directions (north, northeast, east, etc.). The behavioral repertoire of these agents is finite, consisting of $8 \times 16 = 128$ different possible behaviors. Thus, an agent's genotype, i.e., its sensorimotor map, consist of a movement genetically hardwired for each detectable environmental condition. These genotypes are extremely simple, amounting to nothing more than a lookup table of 1024 sensorimotor rules. On the other hand, the space in which adaptation occurs is vast, consisting of 128^{1024} distinct possible genotypes. (As the next section explains, in some environments some von Neumann neighborhoods do not exit and so the corresponding sensorimotor rules cannot ever be used; this lowers the number of *effectively* different genotypes in these environments.)

An agent reproduces (asexually—without recombination) if its resource reservoir exceeds a certain threshold. The parent produces one child, which starts life with half of its parent's resource supply. The child also inherits its parent's sensorimotor map, except that mutations may replace the behaviors associated with some sensory states with randomly chosen behaviors. The mutation rate parameter determines the probability of a mutation at a single locus, i.e., the probability that the behavior associated with a given sensory state changes. At the extreme case in which the mutation rate is set to one, a child's entire sensorimotor map is chosen at random.

Sensorimotor strategies evolve over generations. A given simulation starts with randomly distributed agents containing randomly chosen sensorimotor strategies. The model contains no *a priori* fitness function (Packard 1989), so the population's size and genetic constitution fluctuates with the contingencies of extracting resources. Agents with maladaptive strategies tend to find few resources and thus to die, taking their sensorimotor genes with them; by contrast, agents with adaptive strategies tend to find sufficient resources to reproduce, spreading their sensorimotor strategies (with some mutations) through the population.

During each time step in the simulation, each agent follows this sequence of events: it senses its present von Neumann neighborhood, moves to the new location dictated by its sensorimotor map, consumes any resources found at its new location, and then goes to a new location chosen at random from the entire lattice of sites. This algorithm constantly scatters the population over the entire environment, exposing it to the entire range of detectable environmental conditions. Since the resource field is static, the set of detectable environmental conditions remains fixed throughout a given simulation. Agents never have the opportunity to put together unbroken sequences of behaviors, since each behavior is followed by a random relocation. And since all agents are taxed equally, rather than being taxed according to distance moved, all that matters to an agent in a given detectable local environment is to jump to the site most likely to contain the most resources. Thus, the adaptive challenge the agents face is to make the best possible

single move given specific sensory information about the local environment. Adaptation occurs through multiple instances of these one-step challenge-and-response trials.

The correctness of an agent's perceived view of its environment, within the limits of its discrimination, is altered by a model parameter specifying a level of "sensory noise." This parameter specifies the probability that an agent's perceived neighborhood will be random, rather than reflecting its actual neighborhood. With maximal sensory noise (probability one), an agent's sensory state is always chosen at random, and so its behavior is always chosen from its sensorimotor map at random. Significant sensory noise will damp a population's ability to adapt, because it undermines the process of inducing optimal connections between sensory inputs and behavioral outputs.

2.2   Varying Environmental Structure

We want to study adaptation in a variety of environments that differ only in their environmental structure. At the same time, to make population size a measure of adaptability that can be meaningfully compared across the different environments, we want all of these environments to have the same total quantity of resources. If we let R be the maximal possible resource level at a site (in the present simulation $R = 255$), we can achieve this goal by engineering the environments so that the average resource level at a site is $\frac{R}{2}$. (Although a site can have any of 256 different objective resource levels, recall that the agents can discriminate only four resource levels.) The following suite of environments meets these desiderata:

1. **Flat.** Each site in this environment has a resource level set to $\frac{R}{2}$.

2. **Random.** Resource levels in this environment are chosen at random with equal probability from the interval $[0-R]$, thus ensuring that the average resource level is $\frac{R}{2}$ (see Figure 1).

<div align="center">

figure 1 about here

</div>

3. **Sinewaves.** Resource levels at each site in these environments are assigned by two sinewaves, one along the x-axis and the other along the y-axis. The amplitude of these waves is scaled in such a way that when both are maximal and overlapping the site has the maximum resource level, when both are minimal the site has no resources, and the average resource level is $\frac{R}{2}$. The frequencies of the two sinewaves can be varied independently and are expressed in the number of sinewave periods which cover the x- or y-axes. Figures 2–7 show top-down views (agent's perspective—limited to four levels of discrimination) of 1×1, 16×16, 64×64, 8×48, 14×63, and 34×42, sinewave environments. Note that, since a given site in

the lattice can have only one resource level, the lattice structure imposes a coarse grain on the sinewaves. When the dimensions of the lattice and the frequencies of sinewaves do not match, the coarse grain structure can look rather unusual, as Figures 5, 6, and 7 reveal.

<div style="text-align:center">

figures 2, 3, 4, 5, 6, 7 about here

</div>

4. **Substituting Flat or Random Levels in Sinewaves.** In these environments the sinewave-generated resource level at randomly chosen sites is substituted with either constant or random values. Since the constant resource level is set equal to $\frac{R}{2}$, and the random resource levels are chosen with equal probability from the interval $[0-R]$, the average resource level per site remains $\frac{R}{2}$ regardless of the density of sites. The density of substituted sites is a model parameter. Figures 8 and 9 show top-down views (agent's perspective) of 16×16 sinewave environments in which fifty percent of the sites have been substituted with flat or random resource levels, respectively. (The random substitutions depicted in Figures 8 and 9 differs from the "sensory noise" described in the previous section. Substitutions, when they exist, are part of the objective and permanent structure in the environment. Sensory noise, by contrast, exists only "in the agents' minds" and obscures the environment's objective and permanent structure.)

<div style="text-align:center">

figures 8, 9 about here

</div>

To develop a feel for the various adaptive challenges posed by our suite of environments, it is useful to apply Wilson's and Littman's classification schemes to them. Wilson's (1991) categorization of environments is sensitive to two independent aspects of environmental structure. The distinction between Class 0 and Class 1 environments depends on whether for every stimulus (Class 0), or only for some stimuli (Class 1), there is an action which results in positive reinforcement. In addition, the distinction between Classes 0 or 1 and Class 2 depends on whether the next stimulus is (Classes 0 or 1), or is not (Class 2), determined by the present stimulus and the action. Wilson further subdivides the non-deterministic Class 2 environments into Classes $2.k$, for each finite k, according to whether the next stimulus is determined by the $k-1$ preceding stimulus/action pairs. The flat environment clearly falls into Class 0 since there is never any question about the next sensory state. At the other extreme, the random environment clearly falls into Class 2, since the present sensory state and action do not come close to determining the

<div style="text-align:center">

8

</div>

next sensory state. When we consider sinewave environments, the classification becomes more subtle. If we ignore the possibility that agents collide and start a random walk, then sinewave environments fall into either Class 0 or Class 2; e.g., the 64×64 environment is Class 0 since in every sensory state there is an action that lands on the top of a food pile, while the 1×1 environment is Class 2 since the outcome of all actions from some sensory states is indeterminate. (This is explained further in section 4.2 below.) Furthermore, recall that the environment randomly relocates each agent each time step. So, if the environment is non-deterministic, i.e., Class 2, it will remain non-deterministic no matter how many past stimulus/action pairs are known, so for no finite k is the environment in Class 2.k. On the other hand, when we take account of the possibility that our agents always run some risk of colliding and being forced to do a random walk to the first unoccupied site, then *all* of our non-flat environments become Class 2, and the environment's random relocation debars them from any Class 2.k. That is, all non-flat environments fall into Wilson's most complex class of environments.

Littman (1993) classifies environments in terms of the simplest agent that can adapt optimally to the environment, where the complexity of agents are characterized as follows (p. 262): "the ideal (h, β)-agent uses the input information provided by the environment and at most h bits of local storage to choose an action that maximizes the discounted sum of the next β reinforcements." It turns out that this categorization of environments is not sensitive to the variety among our environments. Since the environment always randomly relocates each agent after each action, there is no advantage to considering more than the next movement when selecting an action, and, by the same token, there is no evident advantage to storing any information in short term memory. Thus, *all* our environments seem to be $(h = 0, \beta = 1)$-environments—the simplest category in Littman's scheme. (Incidentally, if this is correct, then Littman is wrong to identify Class 2 with $(h > 0, \beta > 1)$-environments (1993, p. 262), although perhaps this identification is true when agents can string together uninterrupted sequences of actions.)

The flat, random, and sinewave environments range from being too simple (flat) to too difficult (random) for adaptation, with a variety in between these extremes (various sinewaves). Substituting resource levels in a sinewave environment with either constant or random values generates even more varied adaptive challenges. Furthermore, these adaptive challenges are controllable so that, as the density of flat or random substituted sites approaches one, the adaptive challenge approaches that posed by the flat or random environments; or, as the density of substituted sites approaches zero, the adaptive challenge approaches that posed by the original sinewave environment. By choosing from this suite of environments, we can vary the kind of adaptive challenge posed by the environment, and then measure the extent to which adaptation is affected.

3 Measures of Adaptability and Environmental Structure

To study how adaptability depends on environmental structure, we define independent measures of environmental structure and adaptability. We then observe how adaptability (our dependent variable) responds when we manipulate environmental structure (our independent variable). The measures we propose are not exhaustive or final. Still, they do illuminate how adaptability and environmental structure interact. In addition, they can promote the search for better measures.

3.1 A Measure of Environmental Structure

Adaptation is sensitive to those aspects of environmental structure that the agents perceive and act upon. One such aspect is the variety of the environmental conditions which the agents can discriminate. A natural way to quantify this is with the information-theoretic uncertainty (Shannon and Weaver 1949) of the distribution of detectable local conditions:

$$H_E = -\sum_i P_E(v_i) \log_2 P_E(v_i), \tag{1}$$

where v_i is the i^{th} detectable environmental condition (in this case, a distinct von Neumann neighborhood), and $P_E(v_i)$ is the probability frequency of occurrence across all sites in environment E of v_i.

H_E measures the information content of the environmental conditions that the agents can detect, i.e., how much information on average does an agent gain by detecting a given local environmental condition. This measure is a particular way of integrating two aspects of the distribution $P_E(v)$: its width (number of different v) and flatness (constancy of $P_E(v_i)$). Everything else being equal, the wider or the flatter $P_E(v)$ is, the more uncertain an agent will be about which neighborhood it will detect, the more information an agent will get when it does detect its neighborhood, and the higher H_E will be. We can equivalently refer to H_E as the detectable environment's uncertainty, Shannon entropy, or information content. H_E measures how much information an agent gains by sensing the environment—how "surprising" the sensory information is—on average. By contrast, Wilson's Class 2.k scheme is related to how much information an agent gains about the *next* sensory state—how surprising it is—*given* the present sensory state and action.

In the environments studied here, the neighborhoods v are different patterns in the detectable resource levels in the five sites that make up the von Neumann neighborhood. Since these environments all have static resource distributions, in every case H_E is constant over time. But H_E will change in environments with dynamic resource distributions. The measure H_E thus applies to a wide variety of environments in addition to those studied here.

To develop a feel for aspects of the detectable environmental structure measured by H_E, consider our suite of environments:

1. If E is the flat environment, all local environmental conditions *are* identical, so they all *look* identical to the agents in the population. Thus, for some specific j, $P_{\text{flat}}(v_j) = 1$, and $P_{\text{flat}}(v_k) = 0$ for all $k \neq j$. Thus, $H_{\text{flat}} = 0$.

2. If E is the random environment, all detectable environments occur with (approximately) equal frequency, which makes H_{random} close to its maximal value, which is $\log_2$ of the number of different v. Since the agents in our model can detect two bits of information about resource levels at each site in their von Neumann neighborhood, there are $4^5 = 2^{10}$ detectable environmental conditions, so $H_{\text{random}} \approx 10$. (In the random environments we generated, typically $H_{\text{random}} = 9.95$.)

3. Sinewave environments vary in the x and y frequency of the sinewaves, and the number and frequency of detectable neighborhoods varies with these frequencies. Thus, $P_E(v)$ can have a variety of shapes, and H_E can take a variety of values. For example, $H_{1 \times 1} = 2.65$, $H_{4 \times 4} = 3.99$, $H_{16 \times 16} = 5.73$, and $H_{34 \times 42} = 7.09$.

4. If some fraction of the sites in a sinewave environment are replaced with flat or random resource levels, H_E values can vary quite a bit. Low density of replaced sites tend to make $P_E(v)$ slightly flatter, which makes H_E slightly higher, regardless of whether the resource levels in the new sites are flat or random. As the density of replaced sites approaches one however, depending on whether the substituted levels are flat or random, $P_E(v)$ approaches the shape of $P_{\text{flat}}(v)$ or $P_{\text{random}}(v)$, so H_E approaches the value of H_{flat} or H_{random}.

Finally, we wish to reiterate that H_E does not simply reflect the objective properties (i.e., the resource field) of the environment; it reflects this field as perceived by agents of the population. In this respect, it is like the ways in which Wilson (1991) and Littmen (1993) characterize environments.

3.2 A Measure of Adaptability

We define adaptability, $A_P(E)$, as the degree of adaptive success achieved by population P in environment E. Adaptability depends upon both properties of the environment and upon the population's internal capacities—such as its sensory capacities, its information processing capacities, its behavioral capacities, and its metabolic capacities. We write E as the argument and P as a parameter to indicate that in this study we are focusing primarily on the effect of different environments, and not different populational capacities, on adaptability.

The model we study here is resource driven, and a population's size reflects its ability to locate the resources found in the environment. Thus, in this context we measure adaptability related directly to population size. (In a different context it might be more useful to measure adaptability in terms of, say, birth rate or performance on some predefined test.) Nevertheless,

we cannot assume that observed population size *by itself* is an accurate measure of adaptability, for a population of agents might be able to sustain a certain size simply due to the inherent capacities of the agents and the nature of the environment. For example, given the quantity of resources available in the environment and given the agents' existence taxes, even if all agents have entirely unreliable sensors and so move entirely at random, some number of agents might still survive just due to the probability of accidently "bumping into" resources. To factor out this possibility, we compare maximal equilibrium population size in a given environment with the equilibrium size of a "reference" population in exactly the same kind of environment. The maximum population size is the largest population size reached as certain parameters are varied. (Here, we varied mutation rate.) The reference population has exactly the same set of internal capacities (sensory capacities, information processing capacities, behavioral capacities, metabolic capacities, etc.) as the observed population, *except* that it is engineered in such a way that it *cannot* adapt to the environment. We denote this reference population size $min(P|E)$, where $P|E$ denotes the equilibrium size of population P given environment E. Thus, the adaptability $A_P(E)$ of a certain population P in environment E is how much larger than the reference population the largest population is, expressed in units of the reference population size.

$$A_P(E) = \frac{max(P|E) - min(P|E)}{min(P|E)}. \tag{2}$$

If $max(P|E) = min(P|E)$, we let $A_P(E) = 0$.

We create the reference population by setting the sensory noise parameter to one, thus ensuring that each agent always acts at random. Reference populations can be created in other ways, as well, for example, by setting the mutation rate to its maximal value. Different reference populations might well have different equilibrium population sizes in a given environment. This should create no confusion, though, provided we bear in mind how the reference populations are defined in each context, and provided these reference populations are appropriate for the purposes at hand.

Let's briefly consider how this measure of adaptability works in two sinewave environments. (We study how adaptability depends on environment in detail in the next section.) The environments are 1×1 and 16×16 sinewaves (recall Figures 2 and 3), and mutation rate, μ, is varied among 0, 0.001, 0.01, 0.1, and 1. Figures 10 and 11 show time series of population size from five simulations at different mutation rates in each of the two environments. Figure 12 shows how equilibrium population size in all ten runs varies as a function of mutation rate.

---

figures 10, 11, 12 about here

We see some variation in equilibrium population size at different mutation rates: a slight

effect in the 1×1 environment and a dramatic effect in the 16×16 environment. When $\mu = 1$ and every IF − THEN behavior in each agent is chosen at random, we observe the lowest population sizes. At the other end of the spectrum, when $\mu = 0$ and no new behaviors ever enter the population, equilibrium population size in the 1×1 environment is slightly higher, and in the 16×16 environment equilibrium population size is much higher. Finally, although we sampled only a few intermediate mutation rates, we see that population size increases away from both of these two extremes. In the 16×16 environment, in particular, population size rises dramatically as μ drops below 1.

The way population size varies with mutation rate has a straightforward explanation. If $\mu = 1$, every agent has a randomly generated sensorimotor strategy, so good sensorimotor strategies cannot be inherited. If $\mu = 0$, selection will favor the best sensorimotor strategies that happen to be present in the initial randomly-produced population, but no innovative behavior rules ever enter the gene pool. Small but positive mutation rates both allow agents to pass on good behaviors and allow new behaviors to be tested by the population. This explanation fits with previous observations in similar models on how adaptation (measured differently) depends on mutation rate (Bedau and Bahm 1994, Bedau and Seymour 1994, Bedau 1995).

Since the largest equilibrium population sizes in the two environments occur when $\mu = 0.001$, populations that evolve at this mutation rate give us the value of $max(P|E)$. The maximum equilibrium populations are about 1430 in the 16×16 environment and 495 in the 1×1 environment. Since the reference population size, $min(P|E)$, was observed to be about 400 in both environments, this yields adaptability values of $A_P(16 \times 16) = 2.57$ and $A_P(1 \times 1) = 0.24$. In other words, the equilibrium population size in the 16×16 environment is 2.57 reference populations larger than the reference population, compared to 0.24 reference populations larger for the 1×1 environment. Evidently, this kind of population of agents can adapt much more successfully to the 16×16 environment than the 1×1 environment. The explanation for this difference is explored in section 4.2.

4 Observations of Adaptability and Environmental Structure

Hundreds of simulations were conducted in various environments. Except for environment, mutation rate and sensory noise, all model parameters were held constant across all simulations. Each simulation lasted for 100,000 time steps (although, as Figures 10 and 11 suggest, in many environments population sizes reached equilibrium levels well before the end of the simulation). Population size data were collected every 1,000 steps, and equilibrium population sizes were calculated by averaging population size data collected during the final 20,000 time steps.

Our experiments occurred in three stages. In the first stage, we concentrated on a simple progression of symmetric sinewave environments (i.e., 1×1, 2×2, etc.). In the second stage, we studied some sinewave environments with similar H_E values, but different spatial properties. The results at this stage led us to take a closer look at the adaptive challenge posed by various environments. As a result, we propose two additional components of environmental structure besides H_E. In the third stage of experiments, we varied the density with which flat or random resource levels were substituted in one sinewave environment.

4.1 Sinewave, Flat, and Random Environments

We conducted two hundred simulations in the first stage of our experiments. In order to get an initial sense of how adaptability depends on uncertainty, H_E, of detectable neighborhoods, we focussed on certain symmetrical sinewave environments (i.e., those in which the x and y frequencies are identical). These environments exhibit a gradual variation in H_E values: $H_{1 \times 1} = 2.65$, $H_{2 \times 2} = 3.18$, $H_{4 \times 4} = 3.99$, $H_{8 \times 8} = 5.05$, $H_{16 \times 16} = 5.73$. In order to study a sinewave environment with a higher H_E value, we also ran simulations in the 34×42 sinewave environment, where $H_{34 \times 42} = 7.09$. Finally, to study the most extreme possible environments, we ran simulations in the flat and random environments, with $H_{\text{flat}} = 0$ and $H_{\text{random}} = 9.95$.

Figure 13 shows that, at least in the selected environments, adaptability $A_P(E)$ depends unimodally with the uncertainty of the detectable environments, H_E. Adaptability is nil in the flat environment, with $A_P(\text{flat}) = 0$. In the series of symmetrical sinewaves, as the H_E value increase, so does the adaptability, reaching a maximum value of $A_{16 \times 16} = 2.57$. When we move beyond the symmetrical sinewaves to the 34×42 environment, $H_{34 \times 42} = 7.09$, adaptability falls to roughly half. This environment was added because it has the highest H_E of all sinewave environments observed. Finally, in the random environment, the most uncertain environment of all, adaptability falls almost to zero.

figure 13 about here

This dependence of adaptability on environment interacts as one would predict with factors that damp adaptability. Figure 13 shows that sensory noise damps adaptability, and this damping increases monotonically with noise level. If we assume that the flat environment has too little environmental structure for adaptation, and that the random environment has too much structure for adaptation, and that H_E measures at least one component of environmental structure, then, so far, our observations in these selected environments are consistent with our suggestion that adaptability depends unimodally on environmental structure.

To further test our tentative result that adaptability depends on H_E, we did simulations with three additional sinewave environments. Two environments, 14×63 and 8×48, were chosen to explore the adaptability curves (see Figure 13) at H_E values between that of the 16×16 environment where adaptability was maximal, and that of the 34×42 environment where adaptability first declined. The H_E values for these two environments are $H_{14\times63} = 6.04$ and $H_{8\times48} = 6.23$; compared to $H_{16\times16} = 5.73$ and $H_{34\times42} = 7.09$. At the other end of the scale we added the high frequency 64×64 environment ($H_{64\times64} = 2.00$) to contrast with the low frequency 1×1 environment ($H_{1\times1} = 2.56$) (see Figures 3, 6, and 7).

Figure 14 adds the adaptability observed in the 64×64, 14×63, and 8×48 sinewave environments to the results presented in Figure 13, with the environments ordered according to their H_E value. Clearly adaptability is not smoothly unimodal in H_E. For one thing, although $H_{14\times63} < H_{8\times48}$ and these environments appear on that part of the H_E scale in Figure 14 in which $A_P(E)$ is falling, $A_P(14 \times 63) < A_P(8 \times 48)$. This indicates that, adaptability depends upon more than just H_E.

---

figure 14 about here

Adaptability in the 64×64 environment dramatically underscores this conclusion. Although its H_E value is relatively low, its adaptability is actually higher than that of *any* other sinewave environment; $A_P(64 \times 64) = 3.02$. Clearly, $A_P(E)$ does not depend on H_E alone.

4.2   Additional Components of Environmental Structure

At least two additional factors, besides H_E, can affect the adaptive significance of the information an agent gains from sensing its neighborhood. Ignoring sensory noise, there is a distinction between the number of *objective* as opposed to *perceived* resource levels, namely 256, as opposed to four. Thus, an agent's information about the resources in its local neighborhood is imperfect. Figures 15 and 16 show side-views of both the objective and perceived resource levels in a cross-section of sites in 1×1 and 4×4 sinewave environments, respectively. Second, recall that an agent can move up to fifteen sites away from its current location, but its sensory information is restricted to its present and four immediately adjacent sites; in other words, the *movement horizon* of the agents in the population greatly exceeds their *sensory horizon*.

---

figures 15, 16 about here

The roles played by these two factors in an environment affect the adaptive challenge set by the environment. Consider an agent in a 1×1 sinewave environment, and imagine the agent is about one third the way up a resource mountain (e.g., located at about site 69 in Figure 15). This agent is located on an objective resource gradient which, if it could be perceived, would unambiguously indicate where to find the most resources. In fact, the agent's movement horizon includes sites about two-thirds the way up the resource mountain. However, given the relatively gentle slope in this sinewave, and given the agent's limited sensory discrimination, the agent cannot detect the resource gradient; instead, all sites in its von Neumann neighborhood will appear to contain the same quantity of resources. That is, the agent will be on a *perceived* resource plateau. The agent's sensory information is useless; it cannot indicate in which direction it is best to move. The resource mountain is in one direction and the resource valley is the other direction, but its sensory information provides no hint of which direction is which. Thus, gaps between perceived and objective resources and between sensory and movement horizons limit the extent to which the population can adapt to the 1×1 sinewave environment. By contrast, consider an agent one-third the way up a resource mountain in a 4×4 sinewave environment (e.g., located at about site 49 in Figure 16). Given the steepness of the resource mountains in this environment, the agent will detect a resource gradient. Furthermore, the very top of the resource mountain is within the agent's movement horizon. Thus, an agent's sensory information is much more useful in the 4×4 sinewave environment, and the population should be better able to adapt to this environment.

These reflections underscore that, over and above H_E, there are at least two additional properties of the environment that are relevant to adaptability:

- **Ambiguity.** An environment's ambiguity reflects how misleading are the environmental indications about the adaptive significance of different behaviors. For example, in each instance of each detectable neighborhood there is some optimal behavior. If the same behavior is optimal in each instance of a given neighborhood, then that neighborhood is unambiguous. On the other hand, if the optimal behavior in some instances of that neighborhood is different from the optimal behavior in other instances, then that neighborhood is ambiguous. The distribution of this ambiguity over all detectable neighborhoods reflects a second aspect of environmental structure. (Ambiguity is related to the non-determinism by which Wilson (1991) demarcates Class 2 from Class 0 or 1 environments, but ambiguity focusses only on the degree to which non-determinism is relevant to adaptation.)

- **Value.** An environment's value indicates how much can be gained by adapting. At a given environment site, different behaviors yield different resource payoffs. For example, one can ask how much better than the average payoff is the optimal payoff; this reflects a site's value. The distribution of values over all sites is a third aspect of an environment's structure. (The value scale is related to the Wilson's (1991) dichotomy between Class 0 and

Class 1 environments.)

Ambiguity and value make opposing contributions to adaptability. Everything else being equal, the adaptability varies directly with value, but inversely with ambiguity. To get a feel for these environmental properties, consider two extreme cases. In a flat environment, ambiguity is nil since there is no variation in the payoff of different behaviors. In addition, value is nil since all behaviors have the same payoff. On the other hand, in a random environment, ambiguity is high since the optimal behavior in a given neighborhood varies greatly across the neighborhood's different instances. The value is also high since at most environment sites a maximum or near maximum resource level is within the jump range of agents.

Ambiguity and value seem promising candidates for explaining the relative adaptability of the 64×64 and 1×1 environments. The 64×64 environment has no ambiguity. Every neighborhood is such that there is a behavior that is optimal in all instances of that neighborhood (see Figure 4). In addition, the value of this environment is maximal since the optimal behavior in each neighborhood yields maximal resources (the top of a resource mountain). By contrast, the 1×1 environment is highly ambiguous. Because of the distinctions between objective and perceived environment and between movement and sensory horizon, the optimal behavior in each of its four predominant neighborhoods varies substantially in different instances of a neighborhood (see Figures 2 and 15). Also, this environment has only moderate value. The likely distance to the optimum site, even were the location of this site unambiguous, typically exceeds the movement horizon. The optimal behavior in many neighborhoods, even were it known, would yield only moderate resources.

It is less obvious how to explain the relative adaptability of the 14×63 and 8×48 environments. Systematic study of ambiguity and value in these environments is required.

H_E measures the *quantity* of information detectable in the environment. By contrast, ambiguity and value reflect the *pragmatic implications* of that information, i.e., how revealing about which behavior is optimal is the environmental information, and how much can be gained by the optimal behavior. Our results with sinewave environments suggests that the properties of an environment relevant to adaptation include ambiguity and value in addition to H_E. We can now express this as follows: The adaptability of a population in an environment depends on both the amount and the pragmatic import of the information the population has about its environment, i.e., on the extent to which the detectable environmental information signals "a difference that makes a difference," to use Bateson's (1972, p. 453) phrase.

4.3 Flat and Random Substitutions in a Sinewave Environment

We have suggested that adaptability depends on the amount and pragmatic import of environmental information—that is, the detectable and useful environmental structure—and that this

quantity reflects both the uncertainty (i.e., H_E) as well as the ambiguity and value of the environment. However, we do not propose to measure ambiguity or value here. Nevertheless, we do think that the extremes of the flat and random environments are good examples of what we mean by "too little" and "too much" detectable and useful environmental structure, respectively. Furthermore, substituting sinewave-generated resource levels at more sites with flat (random) values makes a sinewave environment more like a flat (random) environment. So, without measuring ambiguity or value, much less systematically varying them, we can still get some sense of how adaptability depends on detectable and useful environmental structure by observing how adaptability depends on varying the density with which flat or random sites are substituted in a sinewave environment.

We used the 16×16 sinewave environment as our baseline, since this environment has a structure such that the resource gradient is always sensible from an agent's perspective (minimal ambiguity) and a maximum resource level is always within an agent's jump range (maximum value). We varied the density of substituted sites across the values 0.01, 0.05, 0.10, 0.25, 0.50, 0.75, and 1.00, in a total of 75 simulations. Sensory noise was set to zero in all these simulations.

Figures 17 and 18 show how adaptability depends on density of flat and random substituted sites. In both cases, adaptability falls off monotonically with the degree of substitution. As density of flat or random sites approaches one, adaptability approaches $A_P(\text{flat})$ or $A_P(\text{random})$; that is, it becomes negligible. (Note that the population can *slightly* adapt in random environments, evidently accommodating itself to some aspects of the static environmental structure; see Figure 1.) In addition, as the density of substituted sites approaches zero, adaptability approaches $A_P(16 \times 16)$.

> figures 17, 18 about here

These results provide further support for the suggestion that adaptability depends unimodally on detectable and useful environmental structure, although we cannot yet fully quantify this relationship.

5   Conclusions

Our observations support two kinds of conclusions: methodological conclusions about how to quantify properties like adaptability and environmental structure, and substantive conclusions about how adaptability depends on environmental structure.

Our measures of population adaptability, $A_P(E)$, and information content H_E of detectable environmental conditions have considerable virtues. As we have defined it, $A_P(E)$ can be mea-

sured in many systems; in addition, the idea behind $A_P(E)$ can be implemented in other ways, to capture other kinds of population adaptability. H_E does not capture *all* aspects of environmental structure, but it has very general applicability and does capture *one* important component of environmental structure.

Furthermore, we think that ambiguity and value are additional components of detectable environmental structure. These quantities still need systematic study, and they raise further theoretical issues. For example, each might involve the notion of the optimal behavior in a given neighborhood, but this notion itself needs further clarification. In the presence of ambiguity, a given behavioral rule might have multiple possible consequences. Such a rule could be evaluated based on its average, worst, or best possible consequences, corresponding to the maximum expected value, maximin, and maximax decisions strategies (von Neumann and Morgenstern 1944). These different decision strategies are likely to have different evolutionary consequences, whose merits are unpredictable. One way to cope with this unpredictability in the presence of ambiguity might be to have sensorimotor strategies evolve on a short-term evolutionary time scale, but allow the decision strategy itself to adapt on a long-term evolutionary time scale.

How best to combine H_E, ambiguity, and value into a single measure of environmental structure remains an open question. It is striking how difficult it is to quantify the evolutionary task posed by the environment even in relatively simple static resource models, where sensory and behavioral capacities are limited.

What becomes of our original hypothesis that adaptability depends unimodally on the degree of environmental structure? Now that we view environmental structure as involving at least H_E, ambiguity, and value, it would seem that adaptability does not depend unimodally on any of these components taken singly. One would expect adaptability to fall monotonically with ambiguity and rise monotonically with value. It is still unclear how adaptability depends on H_E, everything else being equal, but it seems possible that this relationship is also monotonic. Consider the 64×64, 8×48, and 34×42 environments (see Figures 4, 5, and 7), all of which would appear to have minimal ambiguity and maximum value. Adaptation falls monotonically as H_E increases for these environments (see Figure 14).

Still, our original hypothesis does receive some support if we distinguish the *quantity* and the *pragmatic import* of the detectable information about the environment. H_E measures the former, while ambiguity and value reflect the latter. Pooling what we have learned prompts us to frame a more precise hypothesis about the unimodal dependence of adaptability on environmental structure: Adaptability is low if the agents have either too little or too much information about the pragmatic import of local environmental conditions. In other words, it is difficult for adaptation to build useful connections between a population's sensory input and behavioral output to the extent that the population is either deprived of, or flooded with, information that makes a dif-

ference. Information that makes a difference might be missing either because the agents' sensory limitations *hide* useful structure in the environment, as in the 1×1 environment or in other highly ambiguous environments, or because the environment simply *lacks* useful structure, as in the flat environment or in other environments which have little or no value.

We conjecture that this unimodal dependence of adaptability on environmental structure, when understood as explained in the paragraph above, will be observable in our static resource models, and in other adapting systems, both artificial and natural. New forms of interaction between adaptability and environmental structure may well be generated in dynamic environments in which the process of adaptation itself changes environmental structure. By extending and developing the methods illustrated here, these conjectures can all be subjected to empirical computational tests.

References

[1] Ashby, W. R. (1956). *An Introduction to Cybernetics*. London: Chapman & Hall.

[2] Bateson, G. (1972). *Steps to an Ecology of Mind*. New York: Ballantine Books.

[3] Bedau, M. A. (1994). The Evolution of Sensorimotor Functionality. In P. Gaussier and J. D. Nicoud (Eds.), *From Perception to Action* (pp. 134–145). New York: IEEE Press.

[4] Bedau, M. A. (1995). Three Illustrations of Artificial Life's Working Hypothesis. In W. Banzhaf and F. Eeckman (Eds.), *Evolution and Biocomputation—Computational Models of Evolution* (pp. 53–68). Berlin: Springer.

[5] Bedau, M. A., Bahm, A. (1994). Bifurcation Structure in Diversity Dynamics. In R. Brooks and P. Maes, (Eds.), *Artificial Life IV* (pp. 258–268). Cambridge, MA: Bradford/MIT Press.

[6] Bedau, M. A., Giger, M., Zwick, M. (1995). Adaptive Diversity Dynamics in Static Resource Models. *Advances in Systems Science and Applications*, 1, 1–6.

[7] Bedau, M. A., Packard, N. H. (1992). Measurement of Evolutionary Activity, Teleology, and Life. In C. G. Langton, C. E. Taylor, J. D. Farmer, S. Rasmussen (Eds.), *Artificial Life II* (pp. 431–461). Redwood City, CA: Addison-Wesley.

[8] Bedau, M. A., Ronneburg, F., Zwick, M. (1992). Dynamics of Diversity in an Evolving Population. In R. Männer and B. Manderick (Eds.), *Parallel Problem Solving from Nature, 2* (pp. 94–104). Amsterdam: North-Holland.

[9] Bedau, M. A., Seymour, R. (1994). Adaptation of Mutation Rates in a Simple Model of Evolution. In R. Stonier and X. Yu (Eds.), *Complex Systems—Mechanisms of Adaptation* (pp. 37–44). Amsterdam: IOS Press.

[10] Emery, F. E., Trist, E. L. (1965). The Causal Texture of Organizational Environments. *Human Relations*, 18, 21–32.

[11] Littman, M. L. (1993). An Optimization-Based Categorization of Reinforcement Learning Environments. In J. A. Meyer, H. L. Roiblat, and S. W. Wilson (Eds.), *From Animals to Animats 2* (pp. 262–270). Cambridge, MA: Bradford/MIT Press.

[12] Packard, N. H. (1989). Intrinsic Adaptation in a Simple Model for Evolution. In C. G. Langton (Ed.), *Artificial Life* (pp. 141–155). Redwood City, CA: Addison-Wesley.

[13] Shannon, C. E., Weaver, W. (1949). *The Mathematical Theory of Communication*. Urbana, IL: University of Illinois Press.

[14] Todd, P. M., Wilson, S. W. (1993). Environment Structure and Adaptive Behavior From the Ground Up. In J. A. Meyer, H. L. Roiblat, and S. W. Wilson (Eds.), *From Animals to Animats 2* (pp. 11–20). Cambridge, MA: Bradford/MIT Press.

[15] Todd, P. M., Wilson, S. W., Somayaji, A. B., Yanco, H. A. (1994). The Blind Breeding the Blind: Adaptive Behavior Without Looking. In D. Cliff, P. Husbands, J. -A. Meyer, and S. W. Wilson (Eds.), *From Animals to Animats 3* (pp. 228–237). Cambridge, MA: Bradford/MIT Press.

[16] Von Neumann, J., and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton, N.J.: Princeton University Press.

[17] Wilson, S. W. (1991). The Animat Path to AI. In J. -A. Meyer and S. W. Wilson (Eds.), *From Animals to Animats* (pp. 15–21). Cambridge, MA: Bradford/MIT Press.

Figure 1: Top-down view of the random environment in a 128 × 128 toroidal lattice of sites. Resource levels (depicted with gray scale) are shown from the agents' perspective: agents can distinguish only 4 resource levels, even though sites objectively can have 256 different resource levels.

Figure 2: Top-down view of the 1 × 1 sinewave resource field (see Figure 1 caption).

Figure 3: Top-down view of the 16 × 16 sinewave resource field (see Figure 1 caption).

Figure 4: Top-down view of the 64×64 sinewave resource field (see Figure 1 caption).

Figure 5: Top-down view of the 8×48 sinewave resource field (see Figure 1 caption). Note that the 128×128 lattice of sites imposes a coarse grain on the sinewaves.

Figure 6: Top-down view of the 14×63 sinewave resource field (see Figure 1 caption). Note that the 128×128 lattice of sites imposes a coarse grain on the sinewaves.

Figure 7: Top-down view of the 34×42 sinewave resource field (see Figure 1 caption). Note that the 128×128 lattice of sites imposes a coarse grain on the sinewaves.

Figure 8: Top-down view of the 16 × 16 sinewave resource field in which the resource levels in fifty percent of the sites have been replaced by *flat* values (see Figure 1 caption).

Figure 9: Top-down view of the 16 × 16 sinewave resource field in which the resource levels in fifty percent of the sites have been replaced by *random* values (see Figure 1 caption).

Figure 10: Population size as a function of time for the 1×1 sinewave environment at 5 mutation rates.

Figure 11: Population size as a function of time for the 16×16 sinewave environment at 5 mutation rates.

Figure 12: Equilibrium population sizes as a function of mutation rate in the 1×1 and 16×16 sinewave environments. Five mutation rates were tested in each environment.

Figure 13: Adaptability, $A_P(E)$, as a function of environment and sensory noise. The environments include flat, random, a number of symmetrical (identical x and y frequency) sinewaves, and one asymmetrical (34×42) sinewave. The environments are ordered by uncertainty of detectable neighborhoods, H_E.

Figure 14: Adaptability, $A_P(E)$, as a function of environment and sensory noise. In addition to those shown in Figure 13, the environments include one more symmetrical (64×64) sinewave environment and two more asymmetrical (8×48 and 14×63) sinewave environments. The environments are ordered by uncertainty of detectable neighborhoods, H_E.

Figure 15: Side view of the 1×1 sinewave environment in a 128×128 toroidal lattice of sites, showing both the objective resource field and the agents' perspective of it. Note that, although the objective resource level at a site can have one of two hundred and fifty six possible values, the agents can distinguish only four resource levels.

Figure 16: Side view of the 4×4 sinewave resource field in a 128×128 toroidal lattice of sites, showing both the objective resource field and the agents' perspective of it. Compare with the previous figure.

Figure 17: Adaptability, $A_P(E)$, as a function of the density with which sites in a 16×16 sinewave environment have been substituted with flat resource levels.

Figure 18: Adaptability, $A_P(E)$, as a function of the density with which sites in a 16×16 sinewave environment have been substituted with random resource levels.
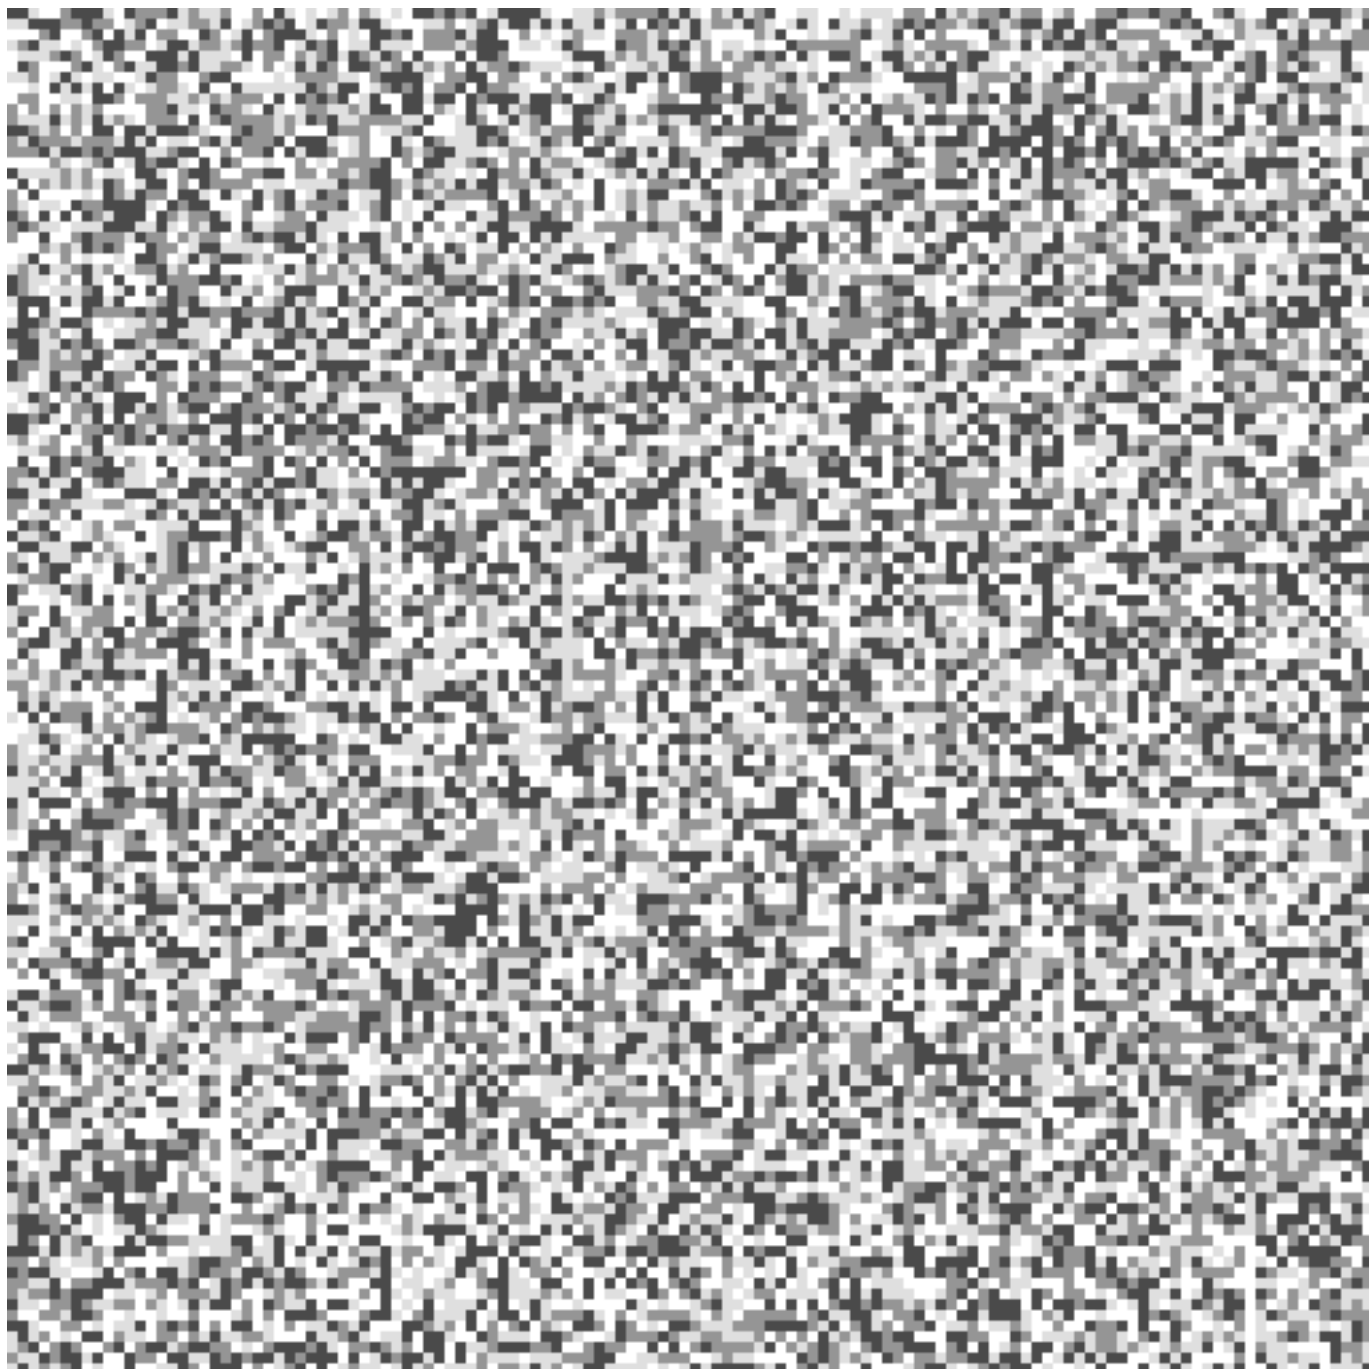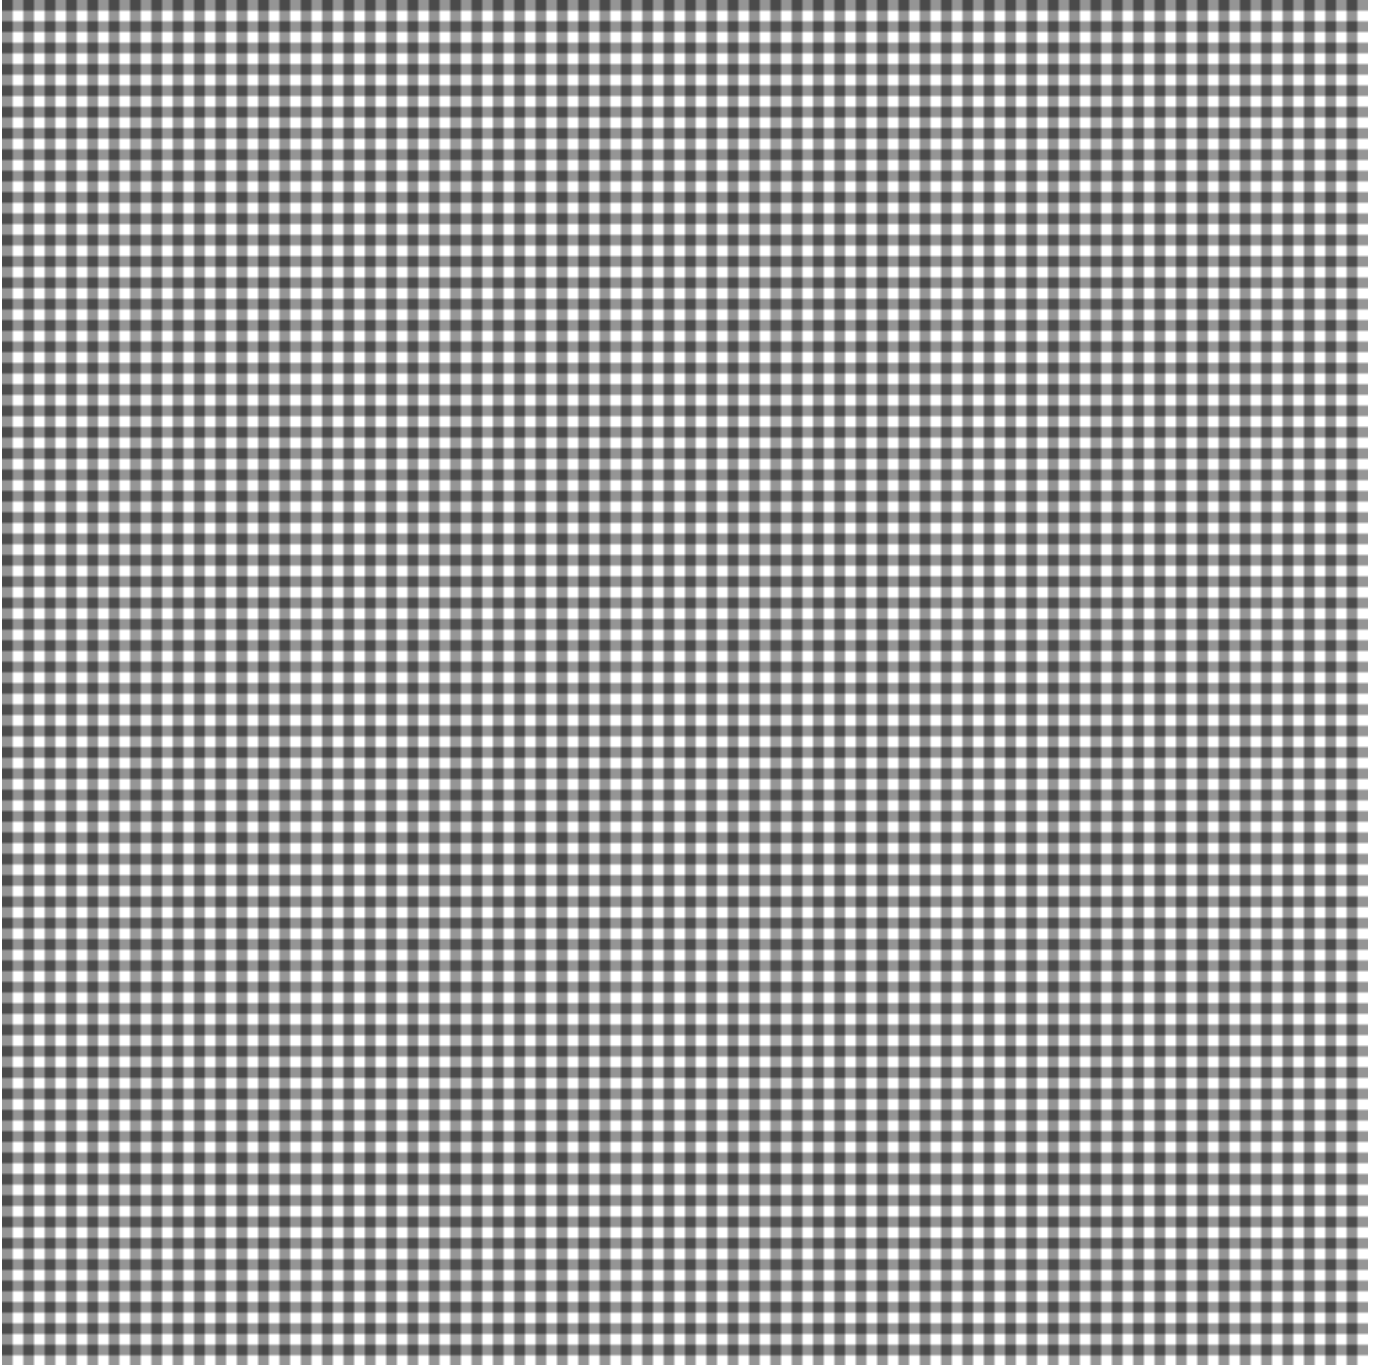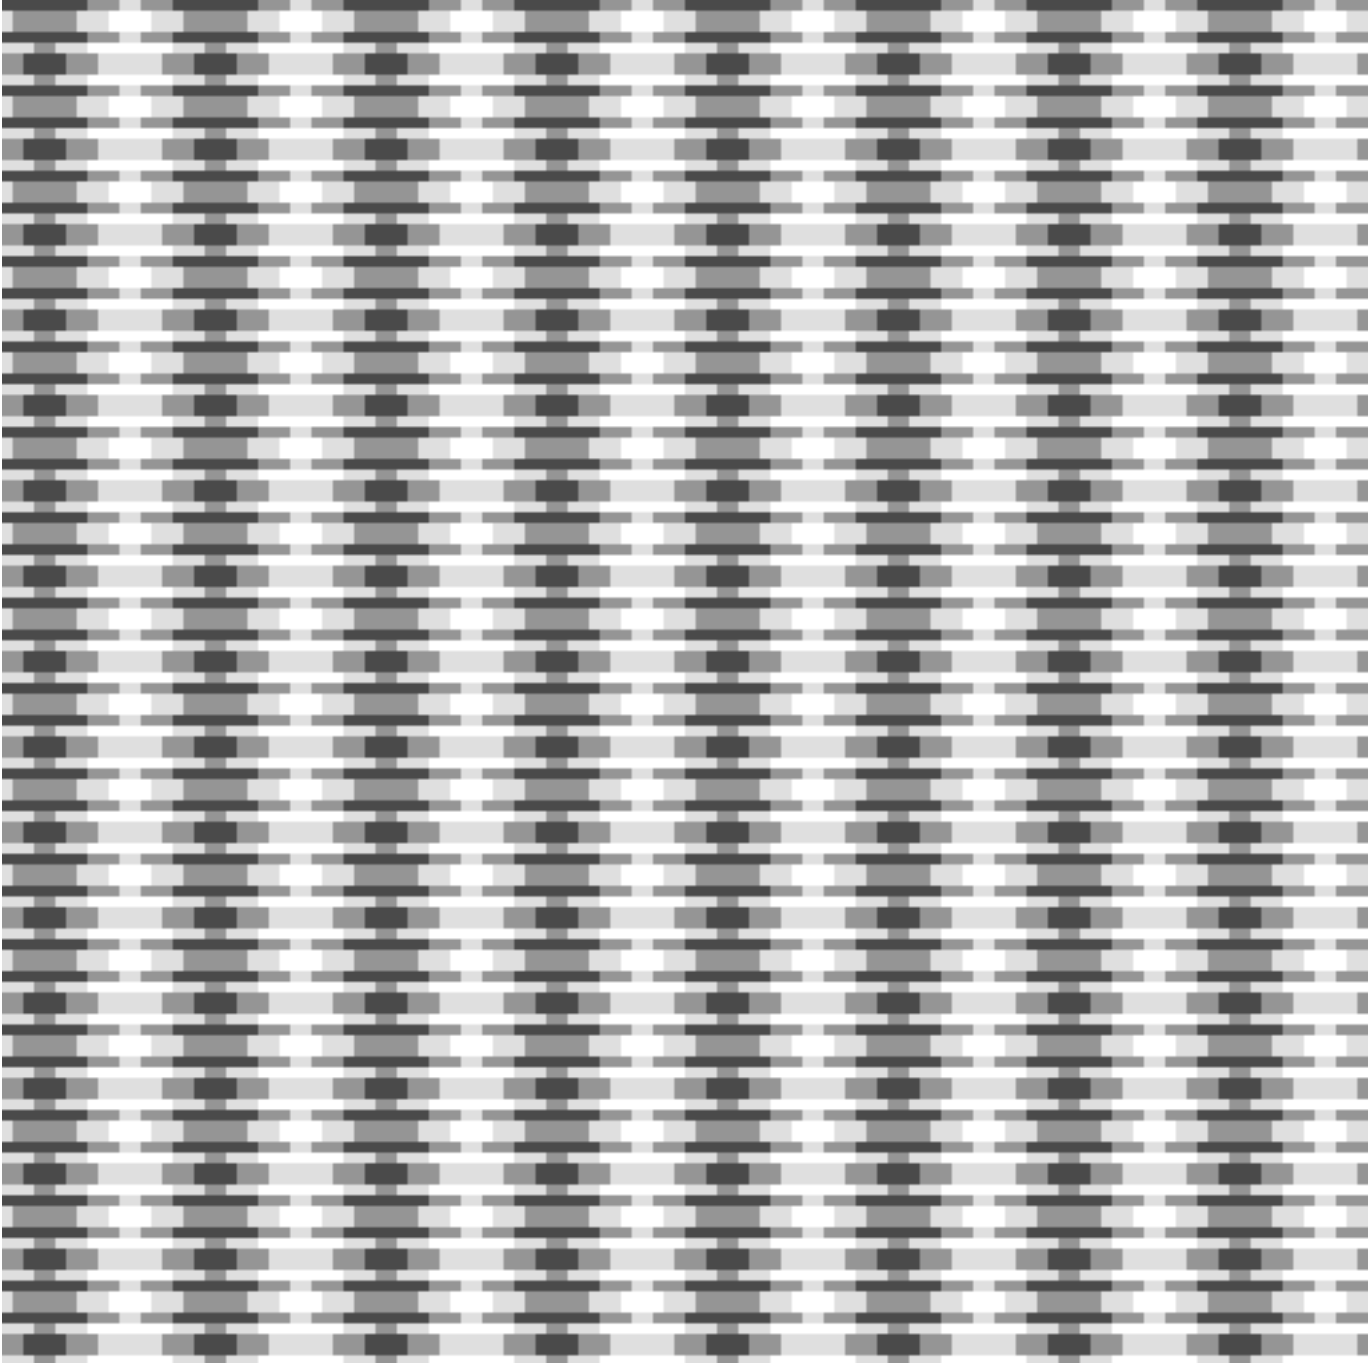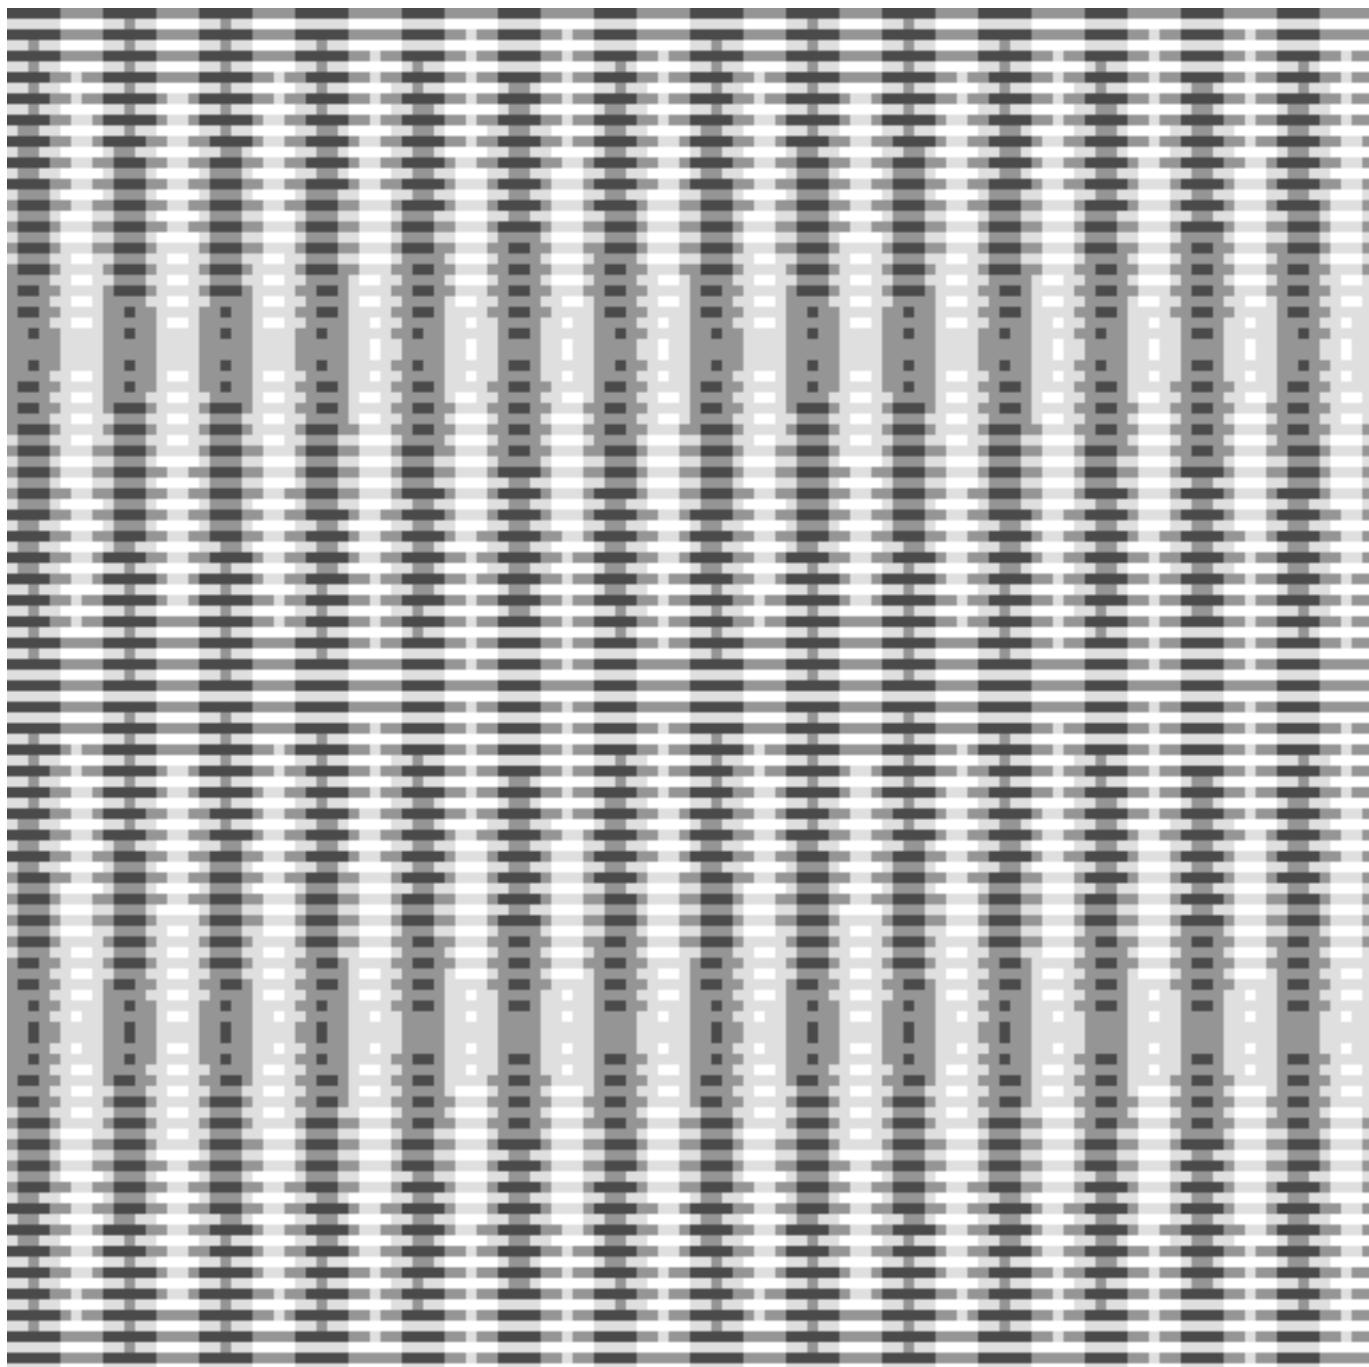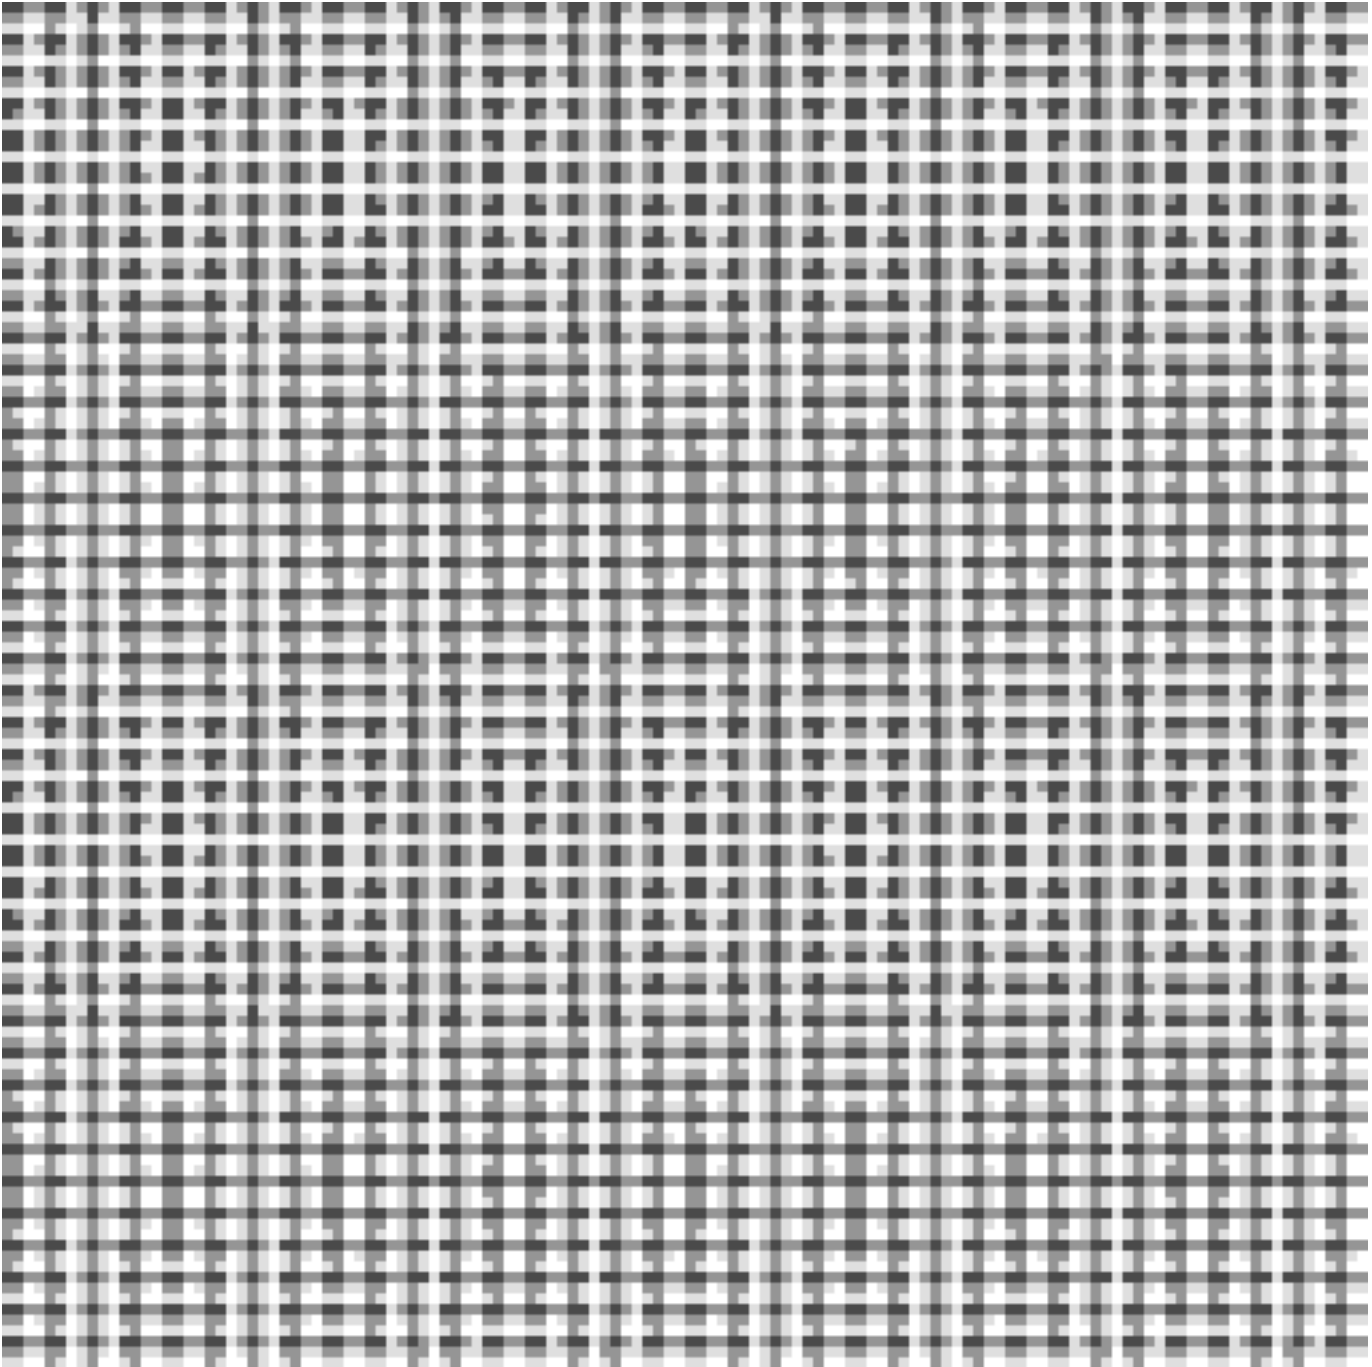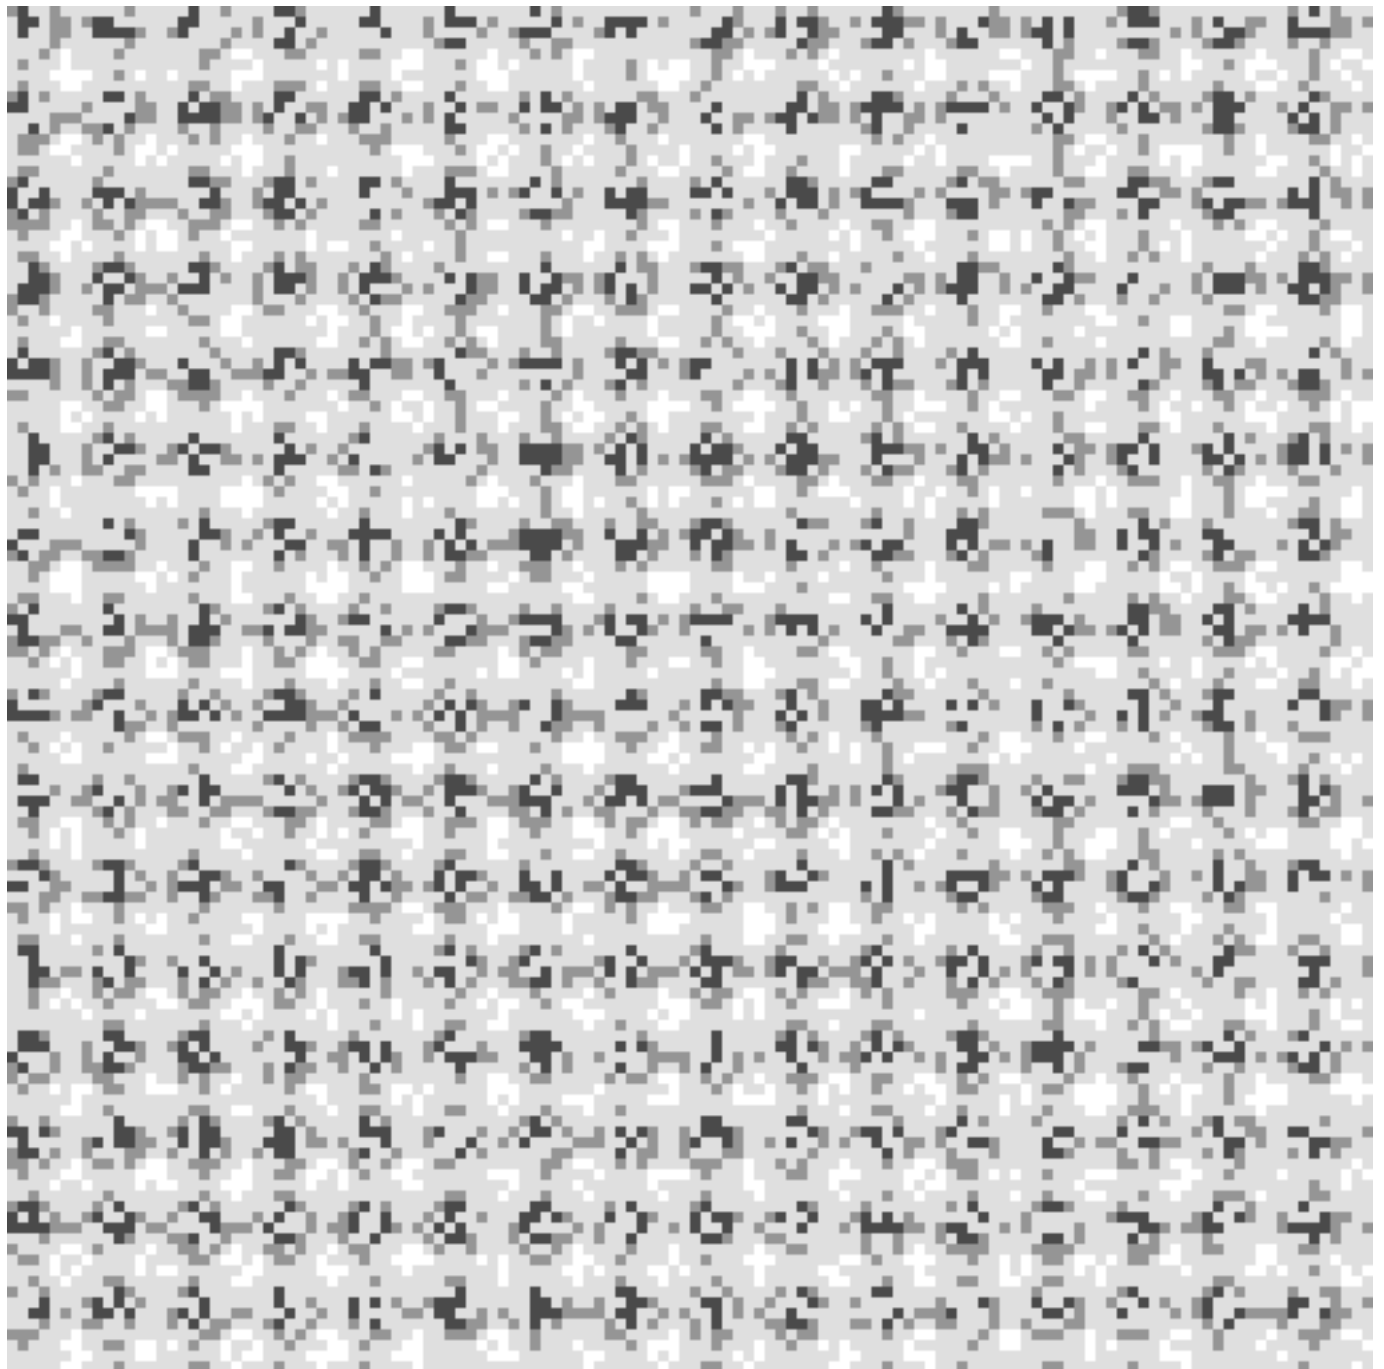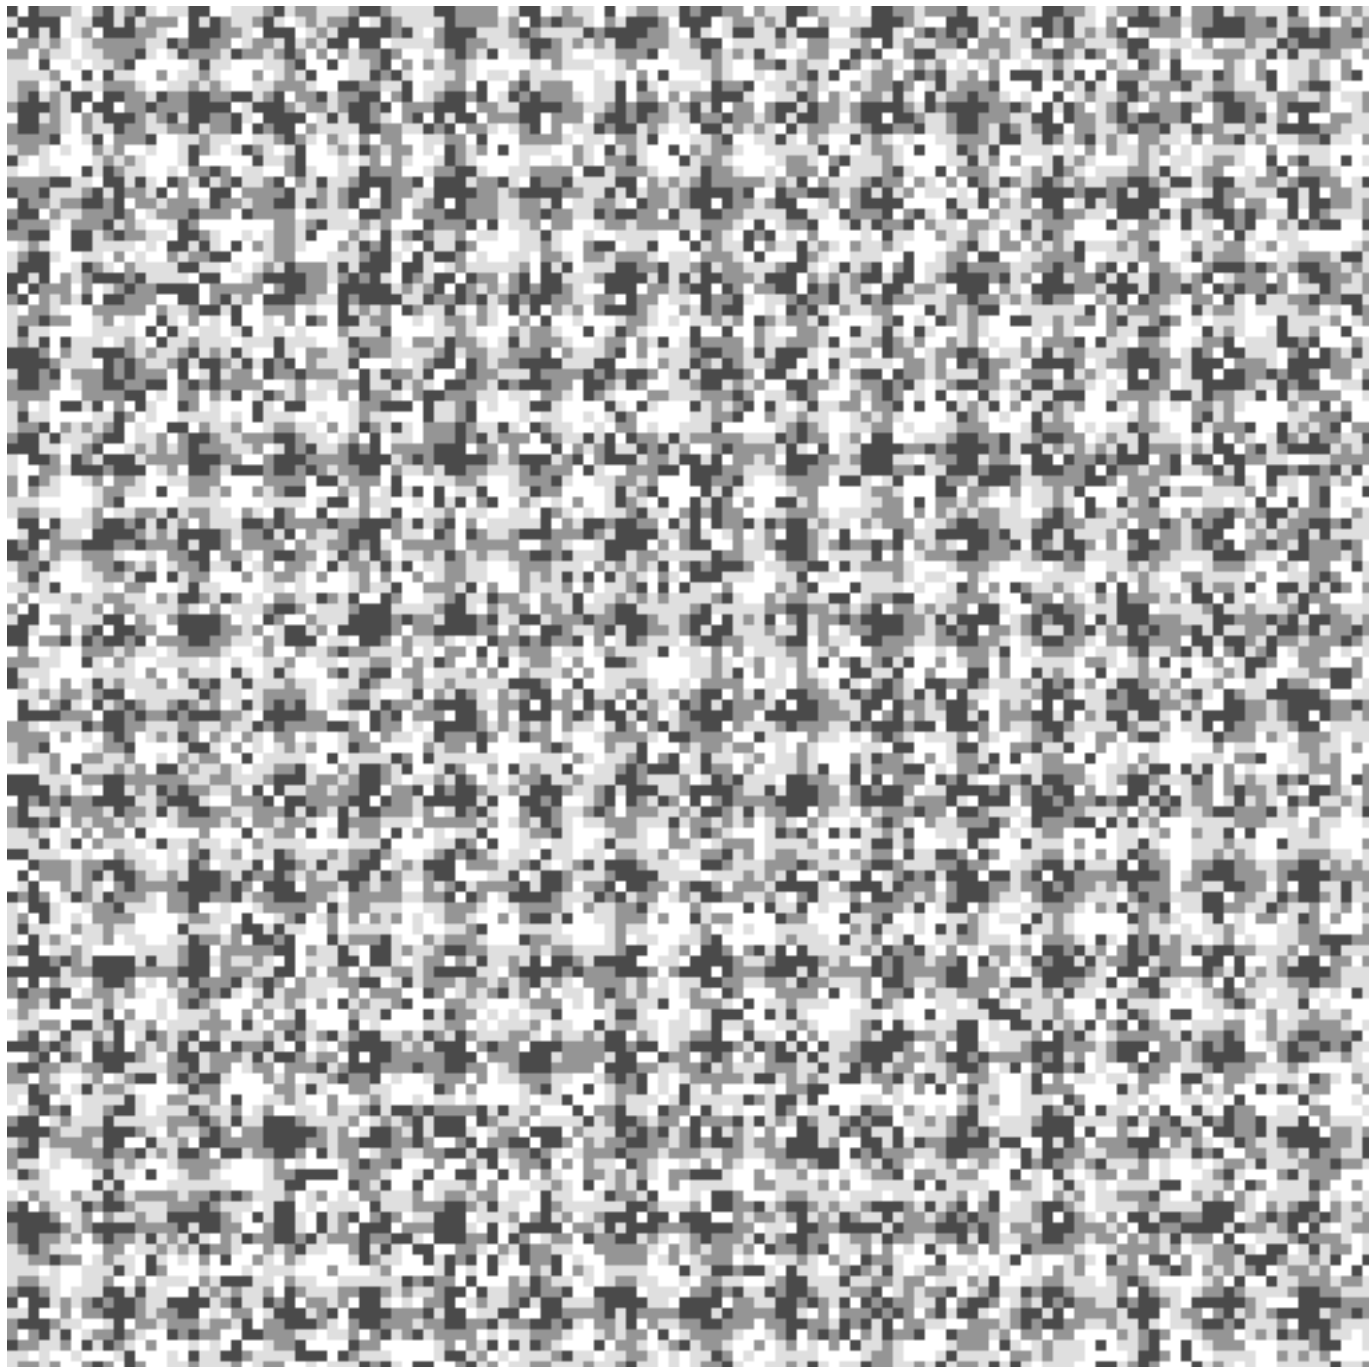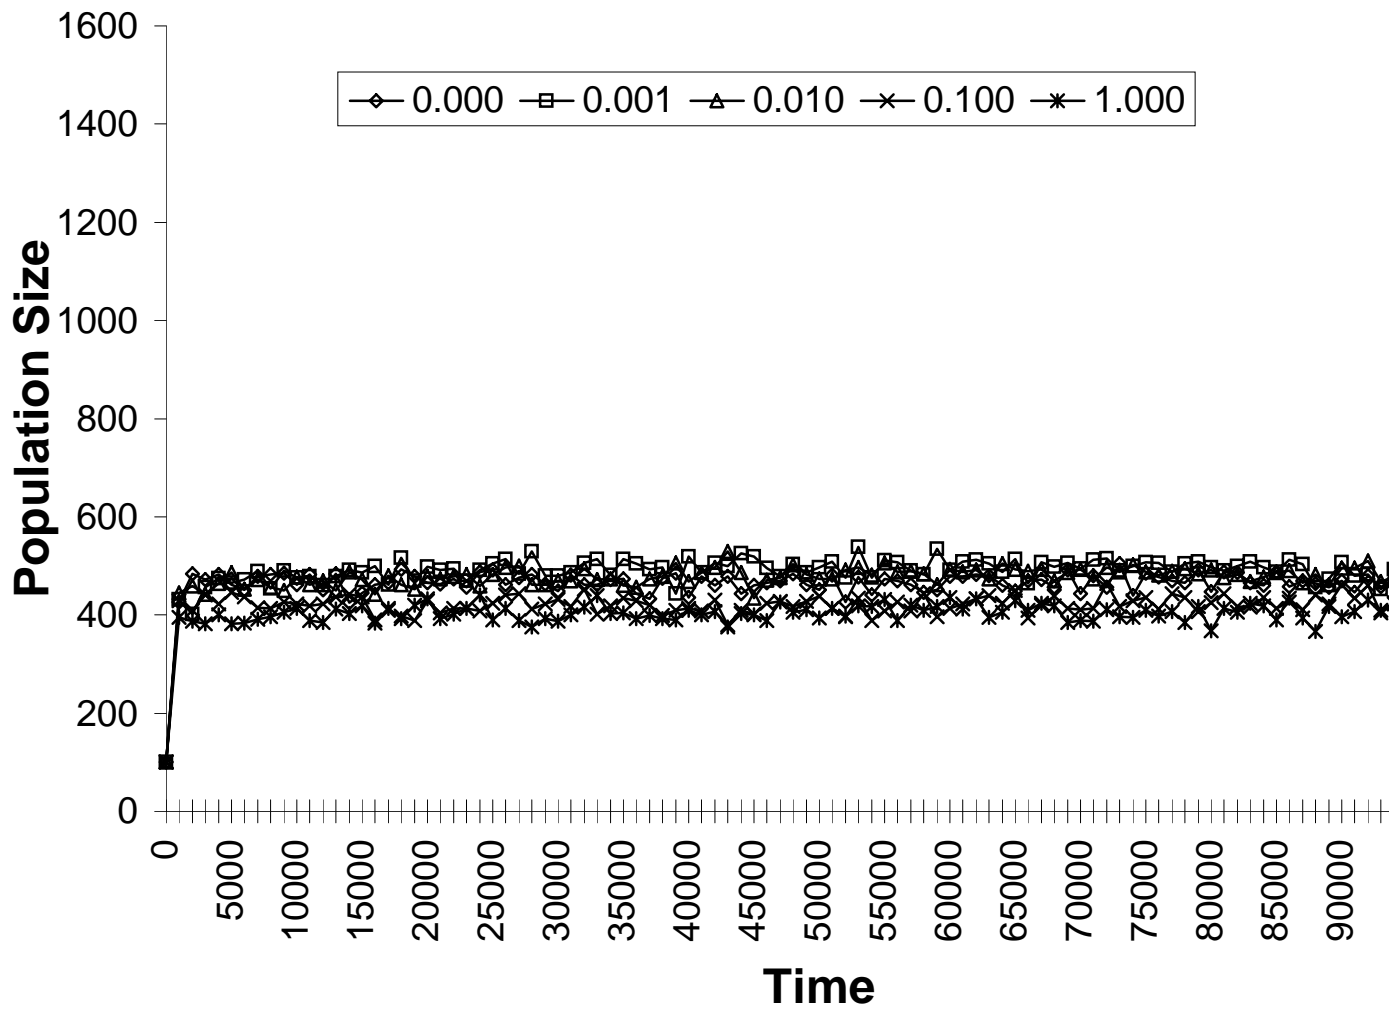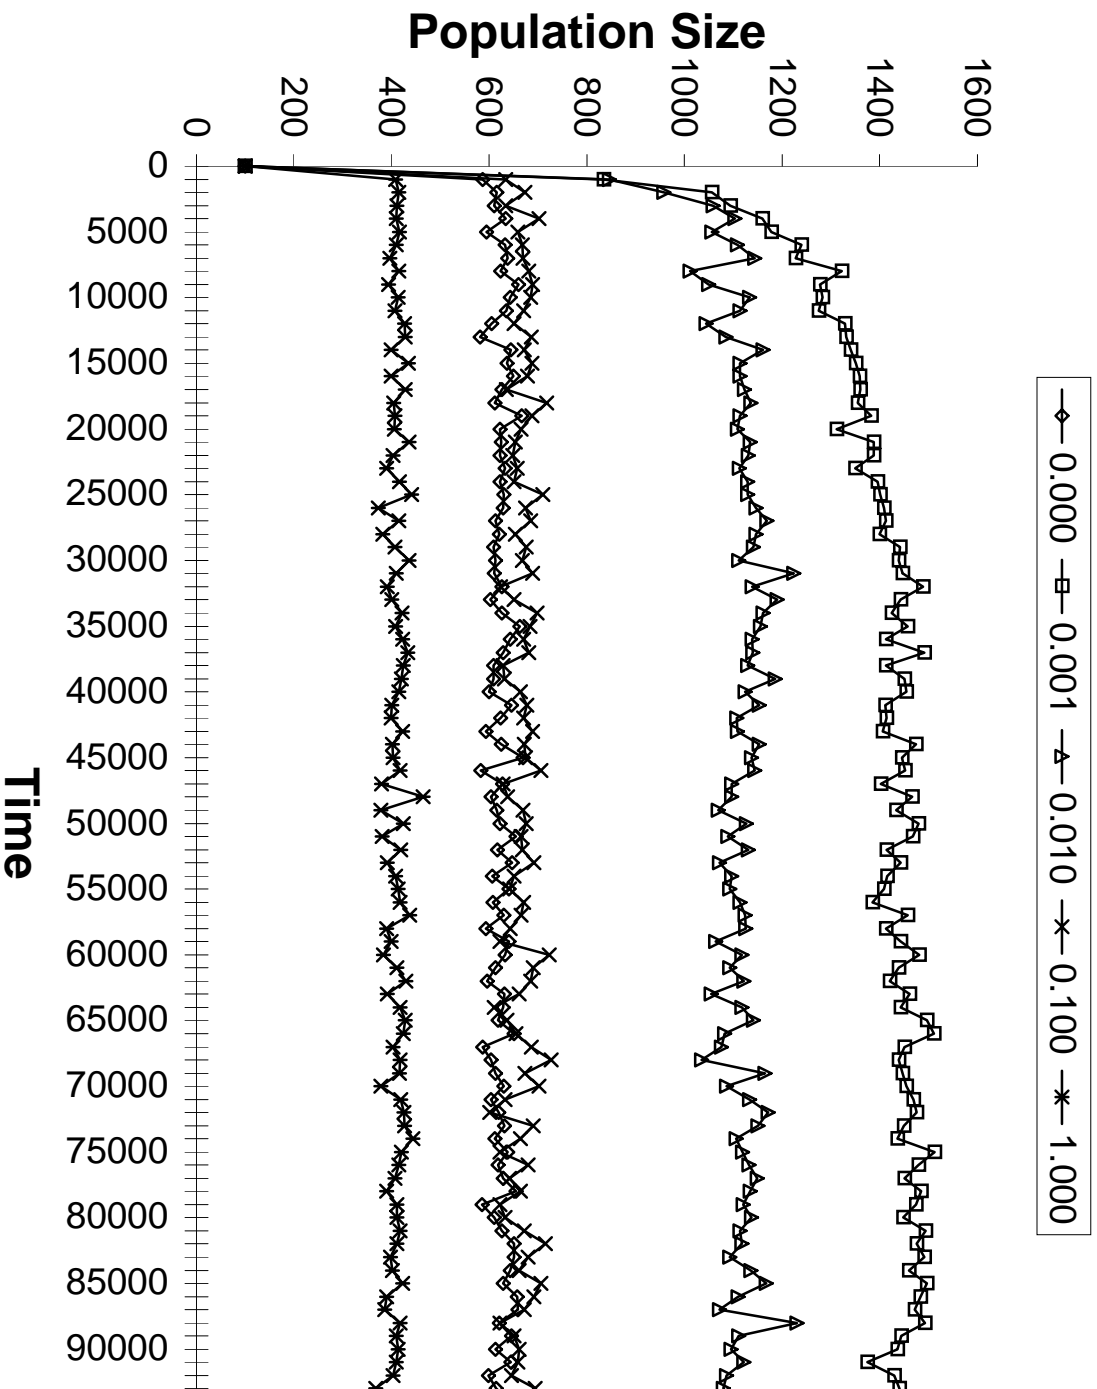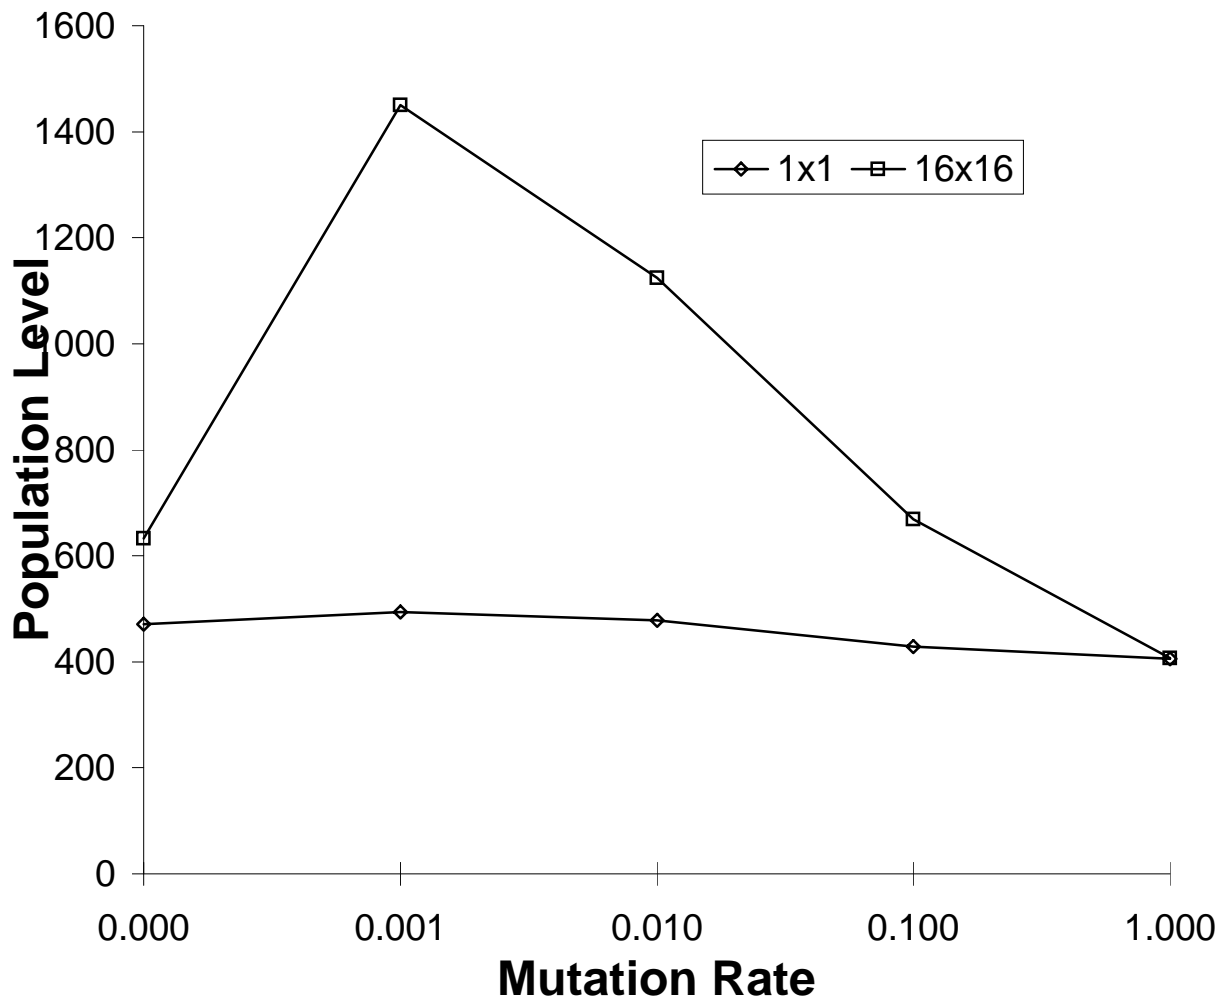
Figure 1
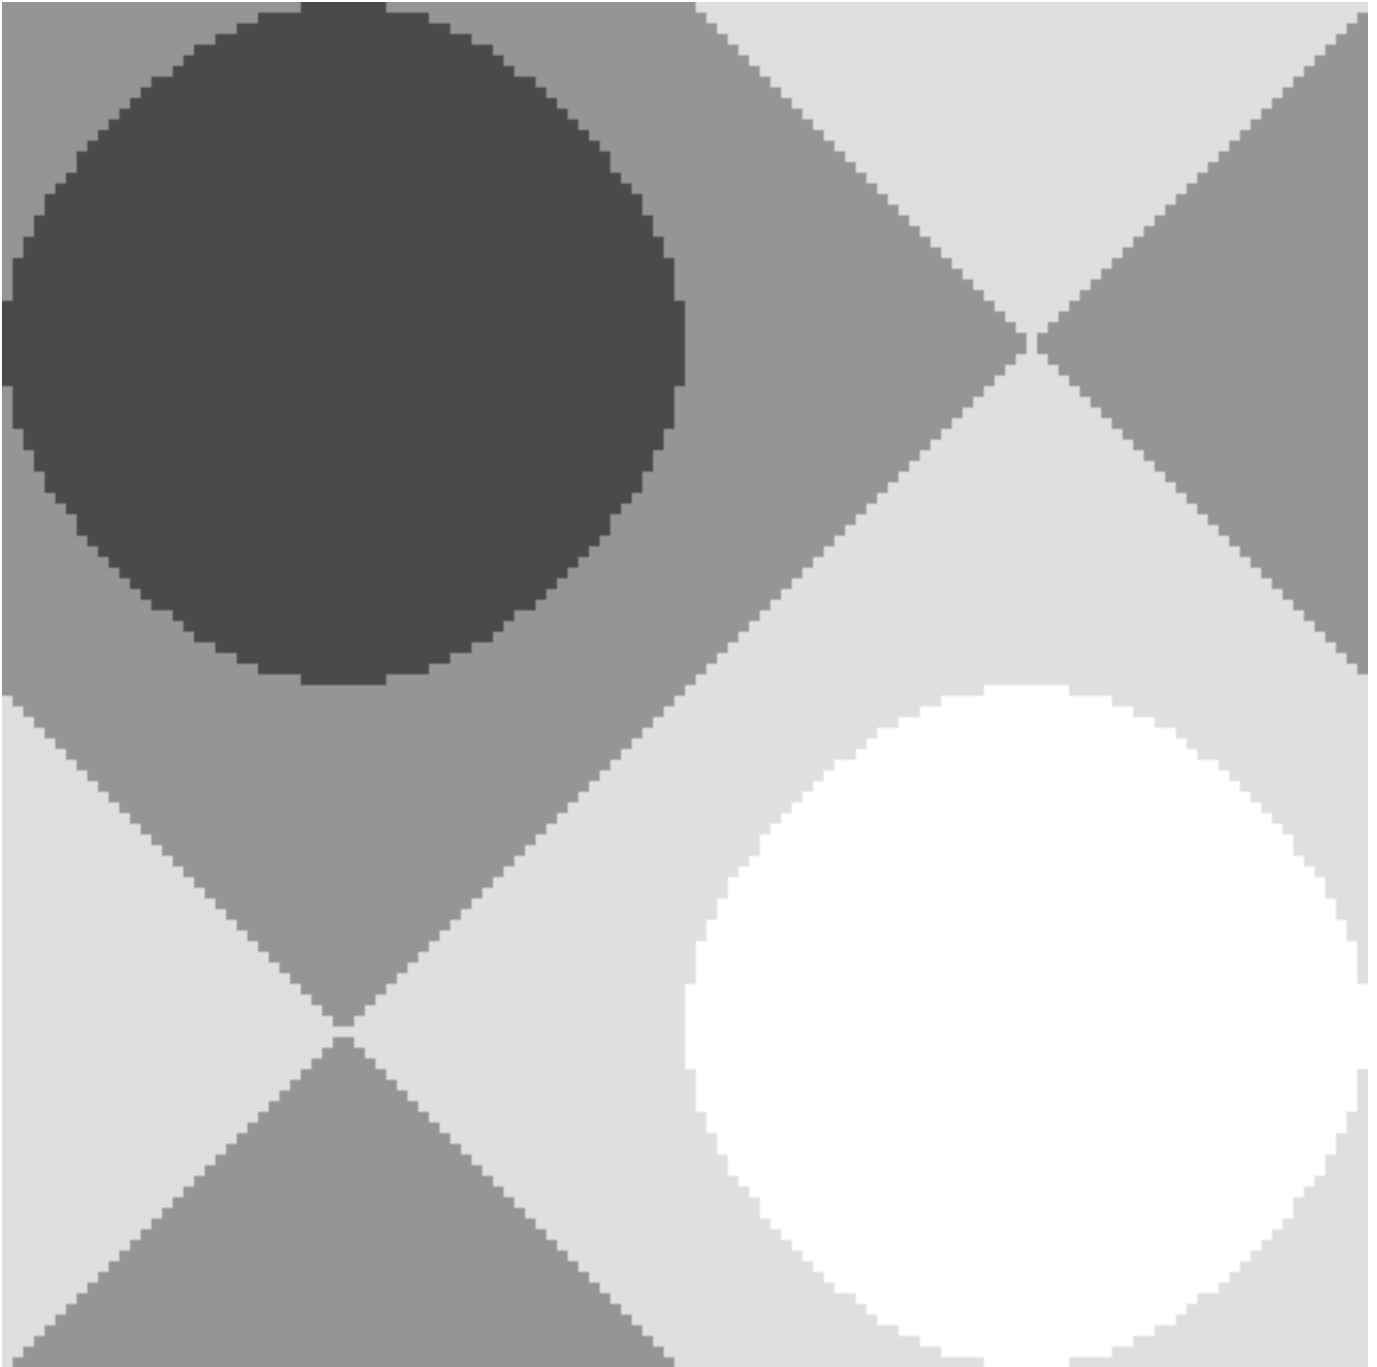
Figure 2
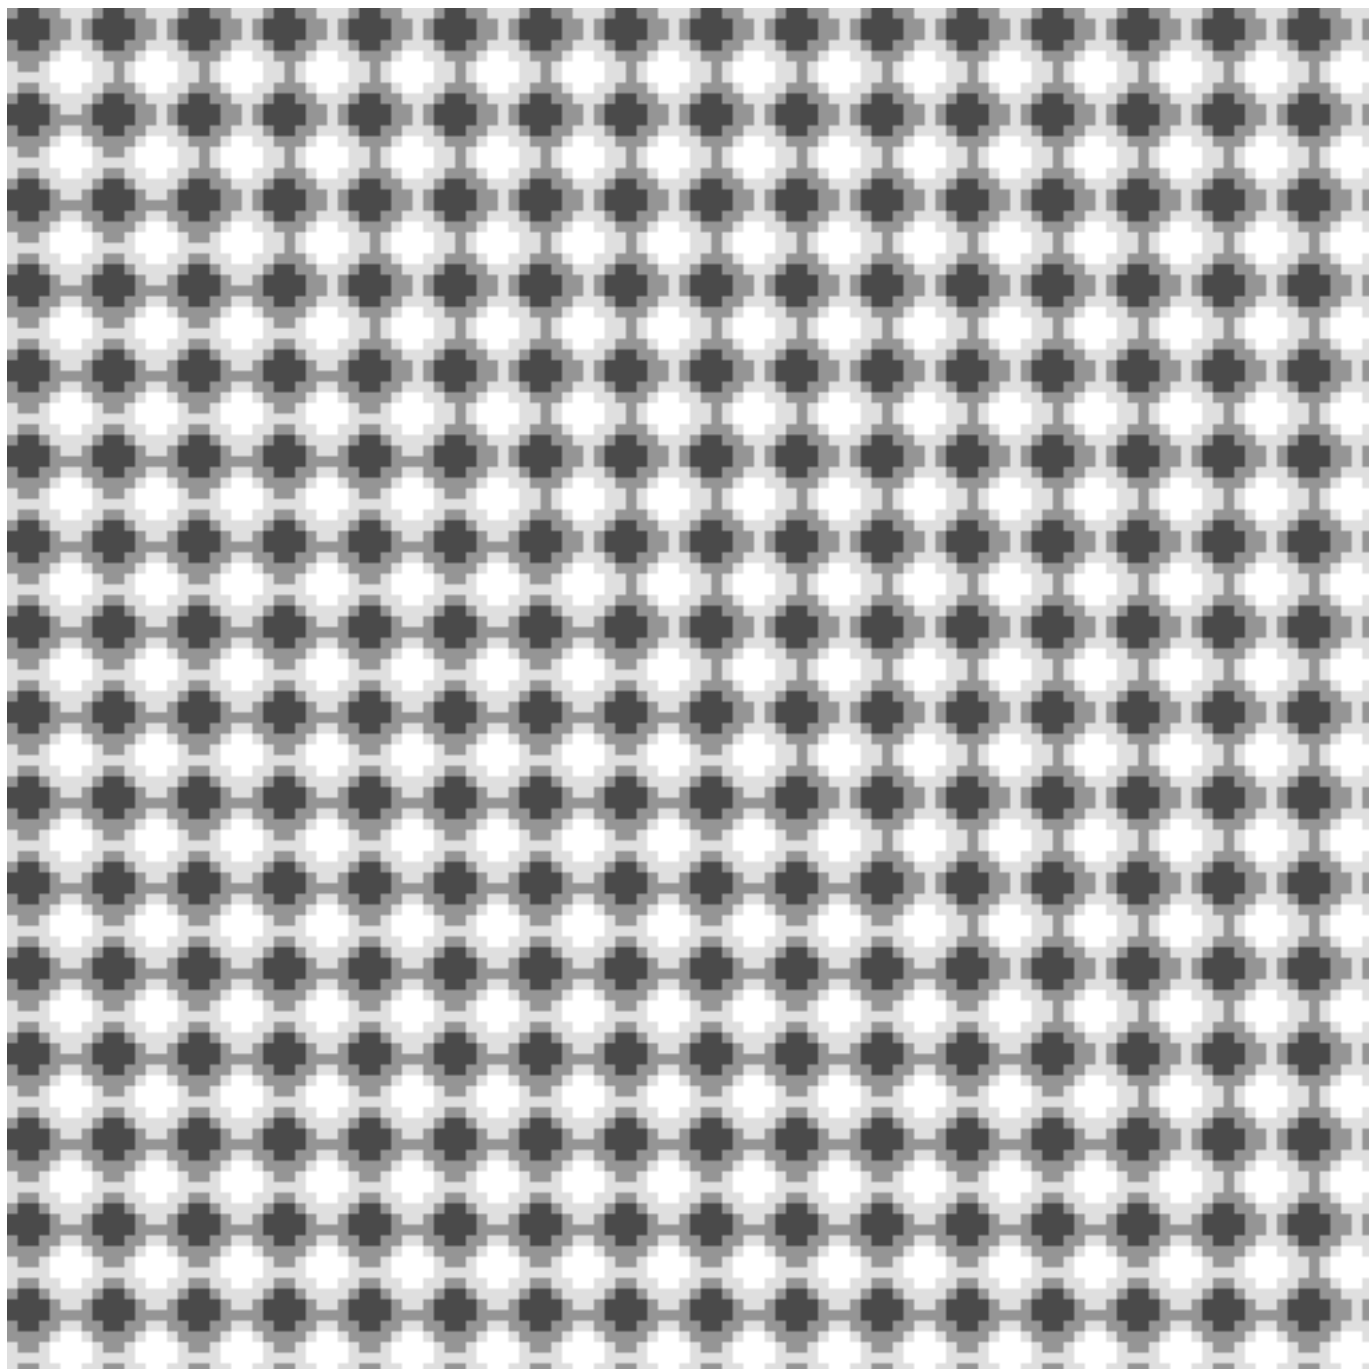
Figure 3
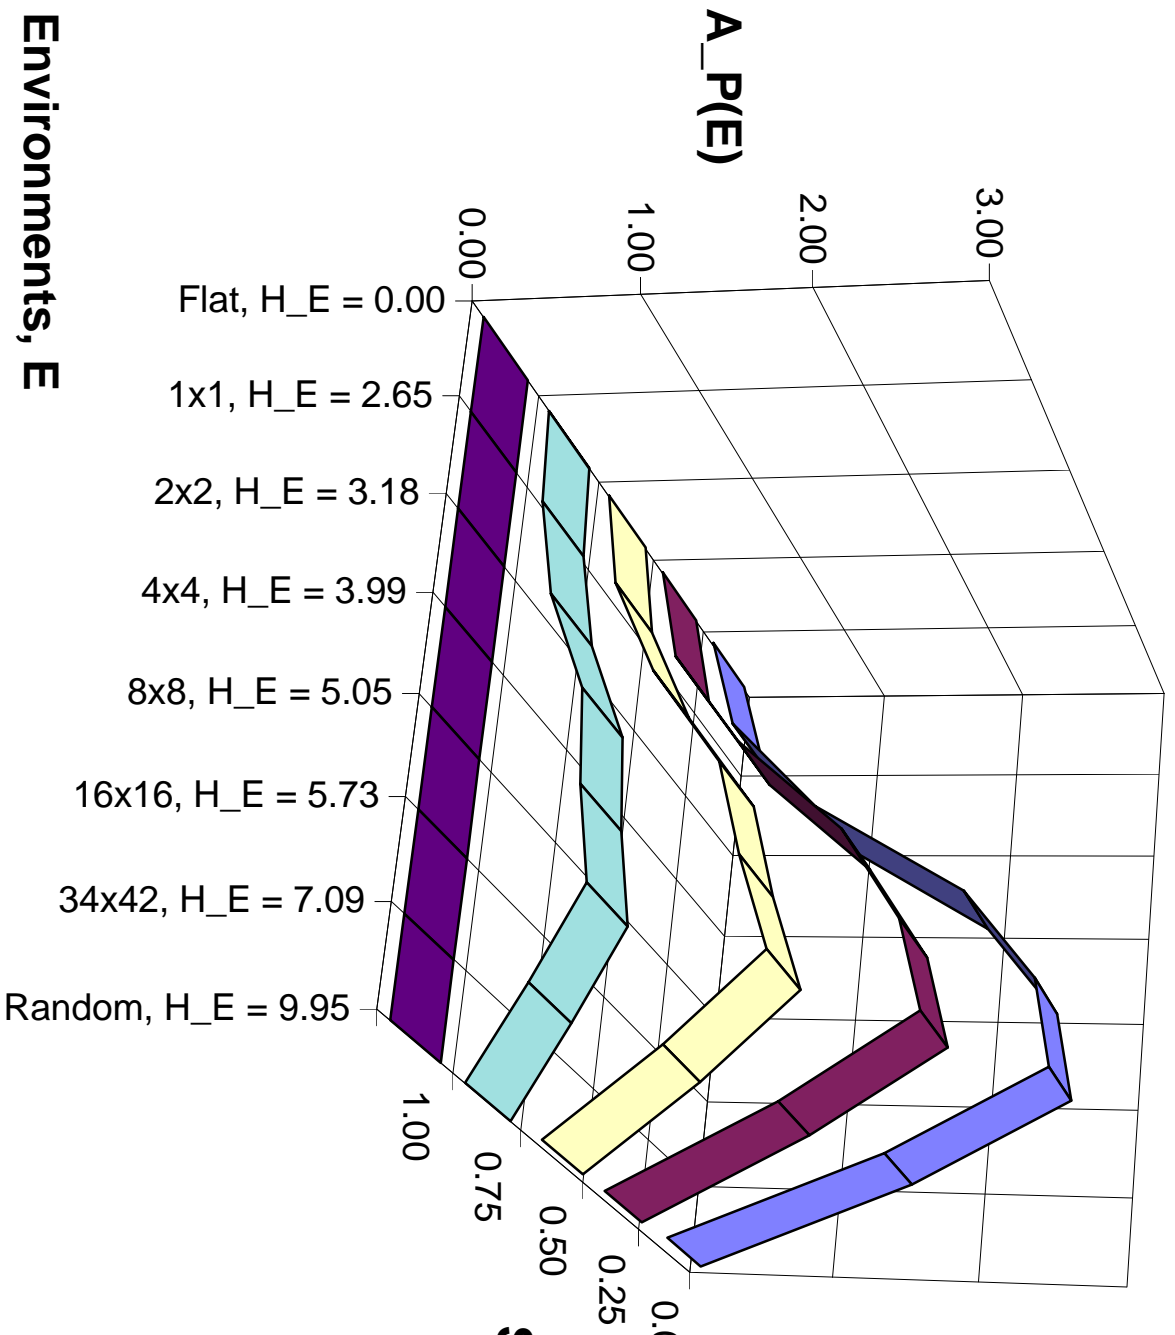
Figure 4

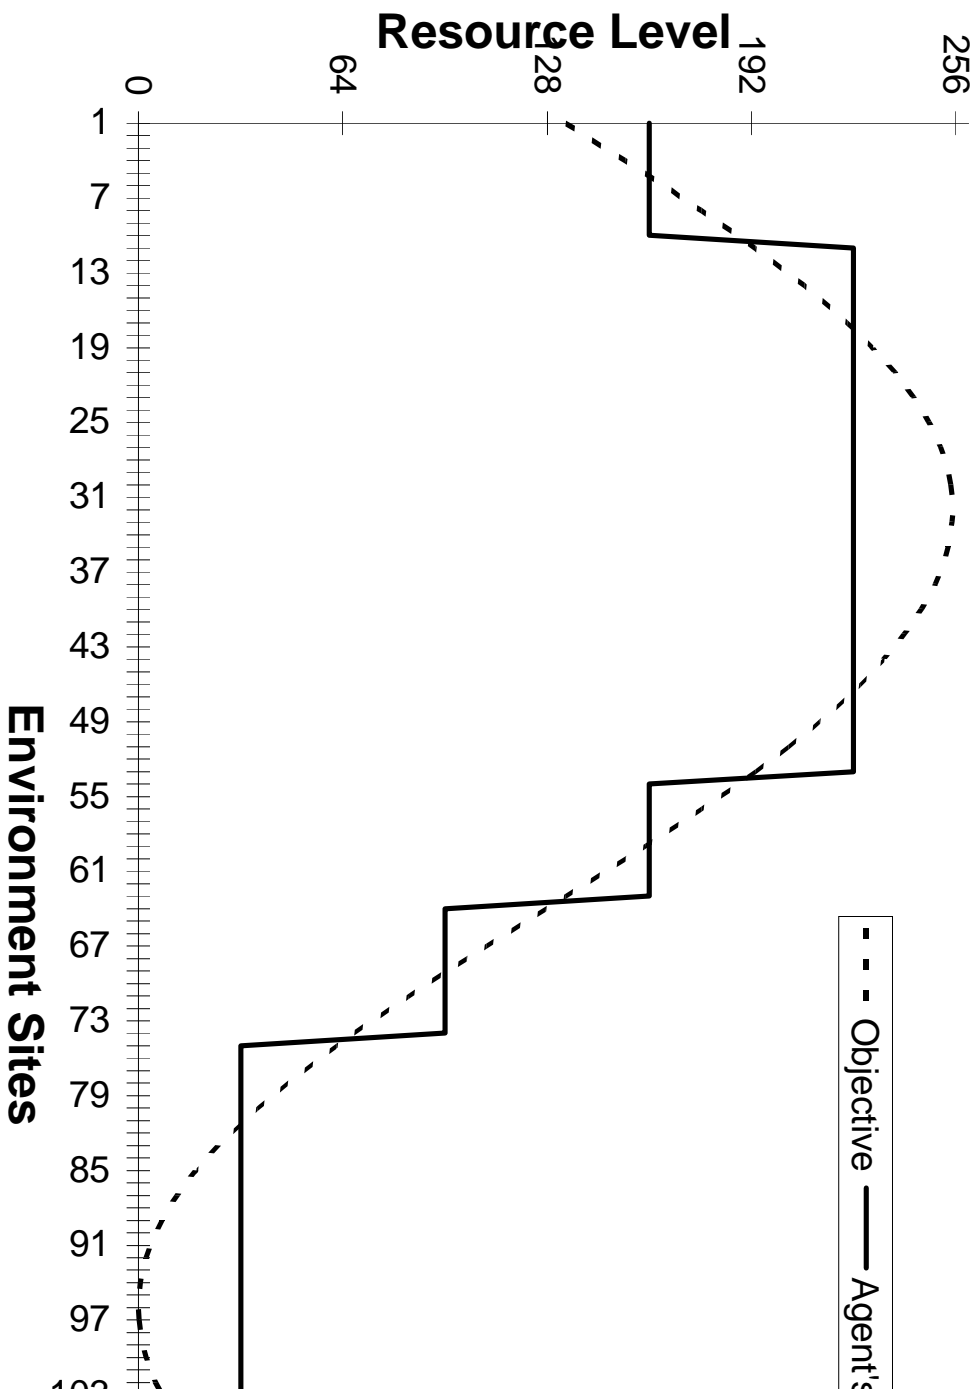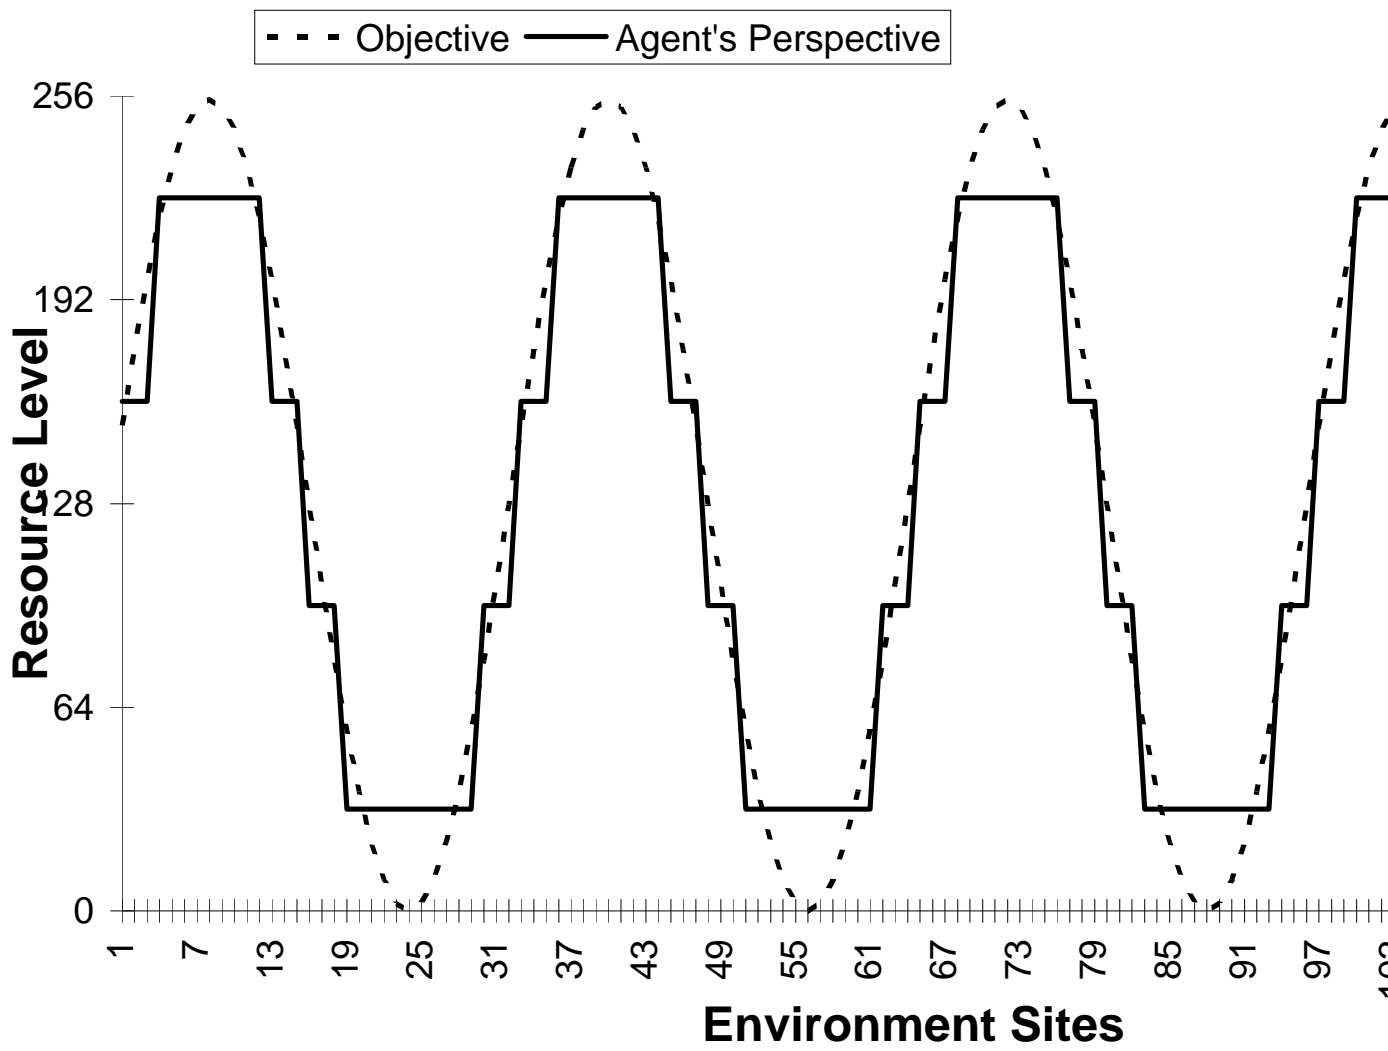Figure 5

Figure 6

Figure 7

Figure 8

Figure 9

Figure 10

Figure 11

Figure 12

A_P(E)

3.00
2.00
1.00
0.00

Environments, E

Flat, H_E = 0.00
1x1, H_E = 2.65
2x2, H_E = 3.18
4x4, H_E = 3.99
8x8, H_E = 5.05
16x16, H_E = 5.73
34x42, H_E = 7.09
Random, H_E = 9.95

Figure 13

1.00
0.75
0.50
0.25

A_P(E)

Environments, E

Flat, H_E = 0.00
64x64, H_E = 2.00
1x1, H_E = 2.65
2x2, H_E = 3.18
4x4, H_E = 3.99
8x8, H_E = 5.05
16x16, H_E = 5.73
14x63, H_E = 6.04
8x48, H_E = 6.23
34x42, H_E = 7.09
Random, H_E = 9.95

Figure 14

Resource Level

256
192
128
64
0

Environment Sites

1
7
13
19
25
31
37
43
49
55
61
67
73
79
85
91
97
103

Figure 15

- - - Objective ——— Agent's
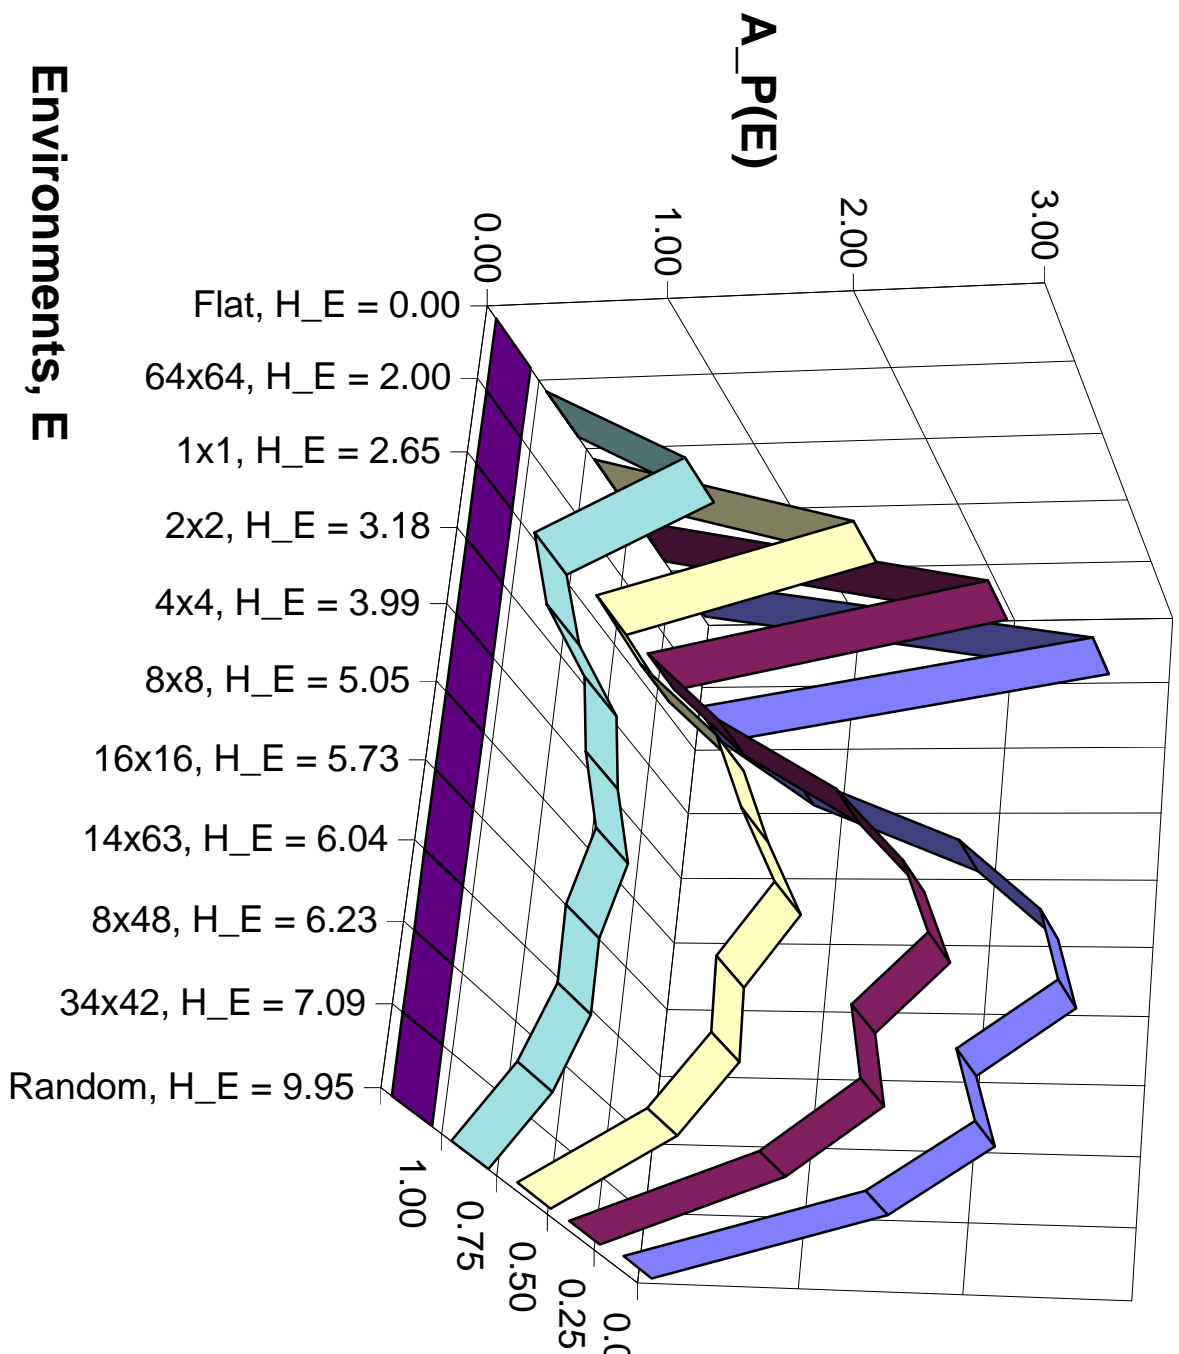
Figure 16

Flat Resource Levels Substituted

Figure 17

Figure 18