VU Research Portal

Increasing Reward in Biased Natural Selection Decreases Task Performance

Haasdijk, E.W.; Eigenhuis, F

published in

The 15th International Conference on the Synthesis and Simulation of Living Systems (ALife XV)
2016

document version

Early version, also known as pre-print

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Haasdijk, E. W., & Eigenhuis, F. (2016). Increasing Reward in Biased Natural Selection Decreases Task Performance. In *The 15th International Conference on the Synthesis and Simulation of Living Systems (ALife XV)* MIT Press.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Increasing Reward in Biased Natural Selection Decreases Task Performance

Evert Haasdijk and Floor Eigenhuis
Vrije Universiteit Amsterdam, The Netherlands
e.haasdijk@vu.nl, f.m.eigenhuis@student.vu.nl

Abstract

This paper presents an investigation into a population of robots that evolves through embodied evolution — an evolutionary process that is not centrally controlled, but emerges from robot interactions just as natural evolution does. The robots select their partners randomly, without reference to any assessment of task performance, but the environment is biased to promote task behaviour by awarding additional lifetime to robots that pick up pucks. The experiments show that the robots do learn to pick up pucks in such a setting. Contrary to what one might expect, increasing the amount of additional lifetime awarded decreases task performance for all settings considered. Closer analysis shows that this decrease is in part due to the fact that the increased lifespan decreases the number of opportunities to spread a robot's genome, but that increasing the award level also negatively affects selection pressure when there is opportunity for robots to spread their genome. We conclude that higher rewards overly emphasise one aspect of robot behaviour and in doing so prevent evolution from exploring the behaviour space.

Introduction

Embodied evolution, or more generally on-line evolutionary robotics, subscribes to a view of collectives of robots that are released in uncharted, possibly changing, environments. The robots learn to operate in that environment, of which the particulars are unknown at design time, through evolution (Watson et al., 2002; Haasdijk et al., 2014). Just as in natural evolution, the process is not centrally orchestrated (in contrast to most evolutionary computing research), but evolution emerges from local interactions between the individuals: they survive to meet other individuals and between them decide to procreate, or not. For such systems to be of practical relevance, evolution must serve two purposes. Firstly, the robots must adapt to their environment so that they get the opportunity to procreate. The robots must, for instance, learn to move about to spread their genomes, or they must maintain their energy levels by regularly visiting charging stations. Secondly, the robots must perform some user-defined task: monitoring, patrolling, surveying, mining or harvesting are often considered in these kinds of scenarios (Bellingham and Rajan, 2007).

Although the environment in which robots operate does not specify any crisply defined objective function, it does indirectly circumscribe goals for the population of robots to survive and evolve. This implies environmental selection pressure that drives adaptation to the environment without reference to any user-defined task. There are two ways to augment this selection to also drive towards task performance: firstly, the robots can explicitly assess their task performance and add a second tier of selection (*mate selection*). Secondly, and more commonly in artificial life than in evolutionary robotics research, the environment can be modified so that individuals with appropriate behaviour receive an environmental advantage that increases their chances of reproductive success, e.g. by increasing their lifespan. In the latter case, task performance is not explicitly selected for, but the natural selection process is biased to promote task performance.

A well-known example of an evolutionary system which explicitly selects for task performance –and origin of the term *embodied evolution*– is that by Watson et al. (2002). Watson et al. added selection on the basis of an individual's prowess in a resource gathering task by varying the frequency at which robots would attempt to broadcast their genomes: better task performance increased this frequency. Thus, the task is explicitly defined and the robots assess their own task performance to drive selection. Haasdijk (2015) investigated the interplay of natural, environmental selection and explicit task-based selection, finding that explicit selection for task performance results in substantially higher selection pressure than that imposed indirectly by the environment, causing robots to prefer environmentally detrimental behaviour if that improves their odds in explicit selection.

This paper considers the alternative method of promoting task behaviour: the rules of the environment are modified so that robots that act appropriately benefit through the environment improving their chance of reproductive success. A good example of such *biased natural selection* (Bredeche and Montanier, 2012) can be seen in the Avida system by Adami and Brown (1994). Here, the individuals are programs in a virtual machine that can procreate by making

copies of themselves, just as in Ray’s seminal Tierra system (1991). Selection for task performance is then added by increasing the clock rate for individuals that perform particular calculations well: this increases the speed at which they can generate copies and so bestows a reproductive advantage. The famous poisonous food experiment by Todd and Miller (1990) provides another example. In this experiment, the task is to collect resources, and virtual agents that gather (‘eat’) the right type of resource increase their lifespan, while eating the wrong type of resource decreases lifetime. The difference in lifetime implies a difference in reproductive success: agents that live longer (i.e., agents that eat ‘healthy’ plants) have more opportunity to procreate. Similar approaches use ‘virtual energy’ as an indicator of task performance to then increase lifespan and so give reproductive advantage, e.g., Elfving et al. (2005) and Weel et al. (2012).

The experiments in this paper take a similar approach in a resource gathering scenario: simulated robots can collect pucks that extend their lifetime by certain fraction of their initial lifespan (the *reward level*). The resultant reproductive benefits promote task behaviour and so induce selection pressure towards puck gathering behaviour. The research question we consider in this paper is how the size of the reward influences the course of evolution.

Experimental Set Up

The experiments in this paper are based on the MONEE experiments by Haasdijk (2015) (from which a substantial part of this section is taken), which in turn extend Bredeche et al.’s mEDEA system (2012). In mEDEA, there is no objective to optimise: robots can exchange genetic material that encodes their controllers whenever they come within a certain maximum distance of each other (e.g., in range for infrared communication). A robot’s controller is active for a fixed amount of time and when this time expires, it randomly selects one of the received genomes and activates a mutated copy of this genome as the new active controller. Thus, robot controllers procreate by transmitting their genome to eggs, and the more eggs a robot inseminates, the more chances it has for procreation. Because the transmission of genomes is continuous and at close range, the more a robot moves about the arena, the better its chances of producing offspring.

The robot controller lifecycle in our experiments consists of two phases: life and rebirth. The robot controllers have a limited, fixed, lifetime during which they perform their actions; moving about, foraging, et cetera (this lifetime may be extended by picking up pucks as described below). When their lifetime ends, they enter a rebirth phase and become ‘eggs’: stationary receptacles for genomes that are transmitted by passing live robots. The rebirth phase also lasts a fixed amount of time, and once this has passed, the egg randomly selects parents from the received genomes to create a new controller. The robot then reverts to the ‘life’ role with

Figure 1: Experiment screenshot. Robots are shown as small circles with sensor beams indicated. Pucks are shown as small green squares (the blue squares show a second puck type that is disregarded in the experiments in this paper). The shaded orange rectangles indicate arena walls and obstacles.

this new controller. The resulting evolutionary process is essentially the same as that in mEDEA: the more ambulant a robot, the higher its reproductive success rate. In contrast to MONEE, there is no further selection criterion because the eggs select the parent genome randomly. In these experiments, the active phase lasted 2,000 time steps, and the egg phase 200 time steps.

Environment and Control

The experiments were conducted in a simple 2D simulator called RoboRobo (Bredeche et al., 2013), simulating 100 e-puck robots in an environment that contains obstacles and pucks.¹ The sides of the square arena are roughly 330 robot body lengths long (1,024 pixels in the simulator), and it contains a number of obstacles (see figure 1) and pucks. The pucks are spread throughout the arena, and they are immediately replaced in a random location when picked up. The robots move around the arena, spreading their genome as they encounter eggs and dying when their time has passed.

Robots can collect pucks simply by driving over them; picking up a puck extends the robot’s lifetime by a certain amount. To detect pucks, the robots have 8 special sensors, laid out in the same manner as the standard e-puck infrared

¹Code for the experiments and analysis scripts is available from <https://github.com/ci-group/monee.git>.

sensors: 6 face forward, 2 face to the rear. Each robot is controlled by a single-layer feed forward neural network which controls its left and right wheels. The inputs for the neural network are the robot’s puck and obstacle sensors as well as two bias nodes (18 inputs in total). The robot’s genome directly encodes the neural network’s weights as an array of reals. The robots select a single parent from the received genomes and their current controller is discarded, so there is no crossover. Variation is applied by adding small gaussian perturbations ($\mathcal{N} = 0, \sigma = 0.1$) to the connection weights.

As mentioned, the robots alternate between periods of active puck gathering (life phase) and motionless genome reception (egg phase). The egg phase lasts 200 time steps, the life phase is initialised at 2,000 time steps, but to prevent synchronised cycles among the robots, we add a small random number to each robot’s initial lifetime. This desynchronises switching between life and rebirth even though the runs start with all robots in sync at the first time-step of their lifetime.

Biased Natural Selection

MONEE extends mEDEA by having the robots select explicitly for task behaviour instead of selecting randomly from the received genomes. In the current set of experiments, however, this explicit selection is disabled. Instead, we provide reproductive advantage and so promote puck-gathering behaviour by rewarding robots for pucks they pick up: each puck yields an increase in lifetime. The amount of added lifetime is defined as the percentage of lifetime added. A reward level of 0.1, for example, means that a robot’s lifetime is extended by 10%, i.e., by 200 time steps. We ran experiments with varying reward levels ranging from 0.05 to 0.8, with 50 replicate runs of for each setting.

The behaviour of collecting pucks and consequently living longer improves the reproductive chances of the robot: the more pucks a robot collects, the longer its lifespan; robots that are skilled at picking up pucks thus live longer and consequently have more opportunities to disperse their genomes by inseminating robots in egg state. Robots with less effective behaviour return to the dormant egg stage sooner, accelerating the distribution of genes that lead to puck collecting behaviour. Through repetition of this process of selection and gene dispersion the entire population’s aptitude increases.

Quantifying Selection Pressure

Haasdijk et al. (2014) introduced a measure to quantify selection pressure that calculates the likelihood of random associations between behaviour and number of offspring in a population. This measure is based on the premise that an increasing level of certainty that the relation between behaviour and fecundity is not random indicates a higher selection pressure. If there were no selection pressure, the relationship between behaviour and fecundity would be random,

and contrariwise, if an individual’s chances of generating offspring depend on its behaviour, the relationship is systemic. Fisher’s exact test (Fisher, 1925) determines the certainty of nonrandom associations between the categories in a contingency table. We construct contingency tables by considering the distance covered, number of pucks collected and offspring count over the lifetime of the robots in the population. We split these individuals into classes with and without offspring and we split them along the median distance travelled or the median number of pucks collected during their lifetime to create two 2×2 contingency tables: one relating offspring and distance travelled and one relating offspring to number of pucks collected. The cells of the contingency tables contain the count of individuals for that cell (e.g., the number of individuals with offspring and below median distance travelled).

Fisher’s exact test estimates the likelihood that the two classes in each contingency table (having offspring and above/below median distance travelled or pucks collected, respectively) are associated. The p-values resulting from these tests indicate the probability that there is no relationship between having offspring and having above- or below-median distance travelled or pucks collected. Thus, low p-values indicate high selection pressure and vice versa. Because the p-values are very small, we ease interpretation and comparison by reporting the log-likelihood multiplied by -1 .

Results and Analysis

Figure 2 shows the amount of pucks collected over time for a range of reward levels. The reproductive benefit of picking up pucks clearly leads to puck-collecting behaviour: the robots pick up more pucks as evolution runs its course. However, the plot also clearly shows that the number of pucks collected is substantially lower for higher reward levels, and that the decrease in performance is systemic. This seems counter-intuitive: not only could one expect increased reward to provide a stronger incentive, but a longer, possibly even infinite, lifetime also implies that a robot can spend more time collecting pucks, because the egg stage in which a robot remains passive is deferred or eliminated. Figure 3 indicates that larger rewards do lead to longer lifetimes: it shows the number of deaths over time for different reward levels. Increasing the reward level decreases the number of deaths in the population per time interval, and with it the amount of time spent in the passive egg state (remember, each death initiates a 200 time steps egg phase). Thus, the robots spend more time in the active state where they can collect pucks.

To analyse the mechanisms that lead to this surprising effect of increasing reward, consider that when the reward is small, robots need to collect multiple pucks to increase their lifespan substantially. Robots that are not so skilled return to the egg stage quickly, providing receptacles for the genomes

Figure 2: Number of pucks collected per 1,000 time steps vs time for different reward levels. The number of pucks collected initially increases more rapidly for higher reward levels, but lower reward levels ultimately lead to better task performance. The line plots were smoothed to emphasise the trends; the shaded area indicates the 95% confidence interval for the median.

of more skilled robots. However, the reproductive advantage from puck gathering for skilled robots decreases as the reward level increases. There are two possible causes: firstly, higher rewards allow relatively unskilled but lucky individuals to gain substantial lifetime extensions, allowing them to maintain and spread their genomes for much longer than would be the case for low reward levels. Secondly, fewer robots die at higher reward levels, leaving less opportunity for robots with relatively effective behaviour to pass on their genomes because there are fewer eggs available for insemination. Figure 3 shows that for reward levels 0.4 and higher, the number of deaths decreases very rapidly, and this stalls evolution because no robots become available for insemination.

Figures 4 and 5 provide more detail about the distribution of puck gathering behaviour and longevity over the populations. The plots show a positive correlation between reward and skewness: high rewards lead to populations where a few individuals collect large numbers of pucks and have very long lifetimes, while the majority of individuals perform at a much lower level. The majority gathers fewer pucks for high reward levels so that the median number of pucks collected per individual decreases as the reward increases. High rewards thus do lead to excellent task behaviour in a select

Figure 3: Number of deaths per 1,000 time steps vs time for different reward levels. The number of deaths decreases as the population evolves to pick up pucks. The line plots were smoothed to emphasise the trends; the shaded area indicates the 95% confidence interval for the median.

few individuals (high best fitness, in evolutionary algorithm terms). This, however, is not a relevant measure in an on-line setting such as this: here, the performance of each individual in the population counts, and with high rewards there is

Figure 4: Distribution of age for different reward levels for individuals that died in the last 10,000 time steps of the runs. Each violin plot shows the probability density of age values for a reward level (note that the x axis is not to scale). Next to the violin plots are boxplots showing median and interquartile range. The plots show combined data over all repeats for each setting (individual runs show much the same pattern).

Figure 5: Distribution of number of pucks collected by individuals that died in the last 10,000 time steps of the runs for different reward levels. Each violin plot shows the probability density of values for a reward level (note that the x axis is not to scale). Next to the violin plots are boxplots showing median and inter-quartile range. The plots show combined data over all repeats for each setting (individual runs show much the same pattern).

such a preponderance of very poorly performing individuals that the overall performance of the population becomes ineffectual.

We perform a second series of experiments where more robots die also for higher reward settings by introducing ‘random death’. In these experiments, robots can die at every time step with a probability of $1.5 \cdot 10^{-4}$, regardless of the amount of lifetime they actually should have left. Thus, eggs are more readily available for all reward levels. The results of these experiments are displayed in figure 6. The graph shows the number of pucks collected during the last 1,000 time steps at the end of each run for the normal condition (in blue) and the random death (in red) conditions. It shows the value for each run as a small circle and the median for each condition as a $\times$ symbol. The plot also shows the result of linear regression model of the relationship between the median number of pucks collected and the log of the reward level.

For lower reward levels the random death condition is not beneficial: the robots actually collect fewer pucks. At these lower reward levels, sufficient eggs are available all the time and more deaths are only counterproductive. Around reward level 0.2, the lines cross, and from this point onward the robots collect more pucks on average when robots die randomly. Note that the clear pattern of robots collecting fewer pucks for higher reward levels persists. This leads us to deduce that the decreasing task performance is also caused by large rewards decreasing selection pressure because even relatively poorly behaving robots receive substantial rewards.

The graphs in figure 7 show how selection pressure develops over time for a number of conditions. There are

Figure 6: The number of pucks collected in the last 1,000 time steps vs the reward level with and without random death. Note that the horizontal axis is logarithmic. The circles represent individual runs, the $\times$ symbols indicate the medians for each reward level setting and the red and blue lines show the result of log linear regression through the medians. For reward levels lower than ca. 0.2, adding random death decreases task performance, for higher values this increases the number of pucks collected.

two components of behaviour that the environment selects for (albeit indirectly). Firstly, robots that move about the environment more have a greater chance of encountering eggs where they can leave copies of their genome. Secondly, picking up pucks increases lifetime and so increase the number of opportunities to spread the genome. Figs 7a and 7b show how the selection pressure deriving from these two components develops over time. For reward levels up to 0.2, selection pressure increases initially as the relevant behaviour spreads through the population. Then, as this behaviour becomes prevalent, the consequent relative reproductive benefit and with it the selection pressure decreases slowly to an intermediate level. Similar trends were also reported by Haasdijk et al. (2014). The selection pressure decreases as the reward level increases, supporting our deduction. For reward levels higher than 0.2, the picture is different: after an initial rise, selection pressure rapidly decreases, then fluctuates to settle at a low value. Selection pressure settles much sooner than it does for low reward levels.

Figure 7c compares selection pressures with and without the random death condition for a reward level of 0.2.

Figure 7: Development of selection pressure over time for different reward levels (7a and 7b) and comparing runs with and without random death. The selection pressure decreases with increasing reward level; for a reward level of 0.8, the selection pressure becomes minimal. Adding random death increases the selection pressure substantially, in particular the selection pressure resulting from puck collecting. Selection pressure is quantified as -1 times the log-likelihood of a random association between number of offspring and speed or number of pucks collected. The plots were smoothed to emphasise the trends.

It shows that increasing the number of deaths in this way substantially increases the selection pressure, in particular regarding the number of pucks selected. Interestingly, the selection pressure becomes much greater and stays much higher than it does without random death for lower reward levels, even though the number of pucks collected remains substantially lower.

The fluctuations for higher reward levels after the initial rise in selection pressure correspond with the fluctuations in Figs 2, 3 and 8. The fluctuations become increasingly pronounced and persistent as the reward level increases. In all cases, these fluctuations occur for reward levels higher than 0.2, which is also the tipping point at which adding random death starts improving puck gathering behaviour. This seems to indicate that this fluctuating trend is related to the lack of available eggs. Because of the high reward levels, there are few available eggs until robots that have poor behaviour but are lucky eventually do start dying off. At that point, there is a period where robots can spread their genome and robots with more appropriate behaviour enjoy some reproductive advantage, increasing selection pressure until there are few eggs available again. The cycle then repeats until the behaviour stabilises.

Higher reward levels do imply a faster increase in the number of pucks collected. The limited increase in lifespan for lower reward levels implies that moving at speed to be able to impregnate many eggs is an important component of reproductive success. This is also borne out by figure 7: initially, the selection pressure from movement is higher than from puck collection. As the median speed of the robots

Figure 8: Median robot speed vs time for different reward levels. The development of speed shows the same trend as that of the number of pucks collected in 2: speed increases more rapidly for higher reward levels, but lower levels ultimately lead to higher speeds. The plots were smoothed to emphasise the trends.

risers (figure 8), collecting pucks becomes more advantageous as now it is possible to live substantially longer even with low reward levels. At higher reward levels, speed is not necessary to live for a very long time, and consequently still be able to spread one's genome. Figure 7 shows that speed matters only initially, and the benefits of collecting pucks appear sooner than for lower reward levels. Thus, it seems that high reward levels emphasise the puck collecting task too much to the detriment of movement, causing evolution to focus on the quick win of collecting a couple of pucks. More modest rewards cause evolution to explore the benefits of movement equally, and in the end this is beneficial for collecting pucks as well.

Conclusion

In this paper, we investigated an evolutionary system where robots randomly exchange genetic material in an environment that bestows benefits on robots that collect pucks. This benefit takes the form of an increase of the robot's lifespan. We showed that the population of robots does learn to pick up pucks as a result of the natural selection that is biased by the reward of additional lifetime without any further explicit selection.

When considering different reward levels, we saw that, unexpectedly, increasing the reward actually leads to poor performance, although the number of pucks collected initially rises more rapidly. High rewards emphasise one component of behaviour (in this case, collecting pucks) to the detriment of other components of behaviour (in this case, movement). Consequently, evolution focusses too much on the quick win of collecting pucks and neglects movement, becoming stuck in sub-optimal behaviour. If the reward is too big, there is no gradient for evolution to exploit: the benefit of mediocre or even poor performance is so big that there is little incentive to improve behaviour and evolution bogs down.

The most immediate conclusion, then, is: if the goal is for a population of robots to collect as many pucks as possible in a setting with biased natural selection, the reward for puck collecting should be minimal. However, it is tenuous to generalise this precise conclusion to other systems, e.g., where the benefit is awarded by increasing the speed of movement.

More generally put we showed that unduly rewarding behaviour in one aspect limits evolution's capability to explore the behaviour space, and the mechanisms that generate selection pressure must be balanced with care. This resonates with findings of research into explicitly selecting for diverse behaviour, for instance by Lehman and Stanley (2011).

It is a truism that evolution requires death, and in our experiments, the effects of overly focussing on puck collection are exacerbated by the lack of opportunities to spread a robot's genome because the increased lifespan reduces the number of available receptacles. This lack seems to become

particularly pressing when the reward level exceeds the tipping point of increasing lifespan by 20%.

The research presented here is part of an ongoing effort to research the interacting selection processes in embodied evolution, and further investigations, in particular to relate the findings here to experiments with multiple tasks and with explicit selection for task behaviour, are underway.

References

- Adami, C. and Brown, C. (1994). Evolutionary learning in the 2d artificial life system *avida*. *Proceedings of Artificial Life IV*, pages 377–381.
- Bellingham, J. G. and Rajan, K. (2007). Robotics in remote and hostile environments. *Science*, 318(5853):1098–1102.
- Bredecche, N. and Montanier, J.-M. (2012). Environment-driven open-ended evolution with a population of autonomous robots. In *Evolution of Physical Systems, ALIFE XIII Workshop, online proceedings*, pages 7–14.
- Bredecche, N., Montanier, J.-m., Liu, W., and Winfield, A. F. T. (2012). Environment-driven distributed evolutionary adaptation in a population of autonomous robotic agents. *Mathematical and Computer Modelling of Dynamical Systems*, 18(1):101–129.
- Bredecche, N., Montanier, J.-M., Weel, B., and Haasdijk, E. (2013). *Roborobo! a Fast Robot Simulator for Swarm and Collective Robotics*. Technical report, arXiv.org.
- Elfwing, S., Uchibe, E., Doya, K., and Christensen, H. I. (2005). Biologically Inspired Embodied Evolution of Survival. In *2005 IEEE Congress on Evolutionary Computation*, volume 3, pages 2210–2216. IEEE Press, Piscataway, NJ.
- Fisher, R. A. (1925). Theory of Statistical Estimation. *Mathematical Proceedings of the Cambridge Philosophical Society*, 22(05):700–725.
- Haasdijk, E. (2015). Combining Conflicting Environmental and Task Requirements in Evolutionary Robotics. In *Ninth IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, pages 131–137, Boston, MA, USA. IEEE Press.
- Haasdijk, E., Bredecche, N., and Eiben, A. E. (2014). Combining Environment-Driven Adaptation and Task-Driven Optimisation in Evolutionary Robotics. *PLoS ONE*, 9(6):e98466.
- Lehman, J. and Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evol. Comput.*, 19(2):189–223.

- Ray, T. (1991). Is it alive or is it GA? *Proceedings of International Conference on Genetic Algorithms*.
- Todd, P. M. and Miller, G. F. (1990). Exploring adaptive agency ii: simulating the evolution of associative learning. In *Proceedings of the first international conference on simulation of adaptive behavior on From animals to animats*, pages 306–315, Cambridge, MA. The MIT Press.
- Watson, R. A., Ficici, S. G., and Pollack, J. B. (2002). Embodied Evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, 39(1):1–18.
- Weel, B., Hoogendoorn, M., and Eiben, A. E. (2012). On-line evolution of controllers for aggregating swarm robots in changing environments. In Coello, C. A. C., Cutello, V., Deb, K., Forrest, S., Nicosia, G., and Pavone, M., editors, *Parallel Problem Solving from Nature - PPSN XII*, pages 245–254, Berlin, Heidelberg. Springer.