Intransitivity Revisited
Coevolutionary Dynamics of Numbers Games

Pablo Funes
Icosystem Corporation
10 Fawcett. St.
Cambridge MA 02138 USA

pablo@icosystem.com

Enrique Pujals
Instituto de Matematica Pura e Aplicada (IMPA)
Rua Dona Castorina 110
Jardim Botanico CEP 22460
Rio de Janeiro RJ Brasil

enrique@impa.br

ABSTRACT

Relative fitness, or "evaluation by tests" is one of the building blocks of coevolution: the only fitness information available is a comparison with other individuals in a population, therefore they evolve in response to each other, without a global fitness to provide a reference. This can lead to failure, in the form of Red Queen Effect, or cycling. *Numbers Games* have been studied by several authors as minimal models of intransitivities which could lead to cycling. Here we carry out an analytical study of the dynamics of minimalistic coevolutionary algorithms in the presence of intransitivities, focusing on two-dimensional real-valued numbers games. We show that depending on the characteristics of the problem, the coevolutionary (1+1) hill-climber either makes good progress with constant average speed, or fails, behaving as a random walk. Larger populations fail to bring qualitative changes into those pathological problems, but teacher-learner separation does. Thus this exercise ends up revealing a fundamental difference between single- and separate-population coevolutionary dynamics.

Categories and Subject Descriptors

I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search—*Heuristic Methods*
; F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*Sorting and searching*

General Terms

Algorithms, Theory

Keywords

Coevolution, intransitivity, numbers games, teacher-learner.

1. INTRODUCTION

The meaning of the word *coevolution* (for computer science) is still being debated. Some suggest that coevolution should be used

as a synonym for *genotypical segregation* between two or more subpopulations. We prefer the definition given by de Jong and Pollack [8] that coevolution "evolves individuals on an evolving set of tests". This notion focuses in the feedback loop inherent to the process, where adaptations lose their benefits as the environment, composed of a co-evolving set of tests, reacts to them.

Elitism can be used, in evolutionary algorithms, to guarantee that the sequence of best individuals from each generation is a non-decreasing one. But coevolution brings an uninvited guest, namely the Red Queen effect: even though individuals in the next generation are better than those in the current one, they could actually be worse than those from a few generations ago. In spite of an evolutionary ladder that pushed for progress at each successive step, we can end up back in the beginning.

This issue, of relative versus absolute notions of progress, brings up the notion of *intransitivity*: what seems to be happening is that generation n is better than generation $n-1$ which is better than generation $n-2$, but this does not necessarily imply that n is better than $n-2$.

Watson and Pollack [11] proposed the Intransitive Numbers Game (ING), an abstract model to study intransitivity in evolutionary algorithms. Others followed, studying variants of the problem and the behavior of different algorithms [2]. De Jong and Pollack [8] demonstrated a two population Pareto-ranking algorithm capable of enforcing monotonic progress in at least one variant of ING.

Here we study Watson and Pollack's original ING, together with an infinite family of related problems. We analyze the behavior of different classes of algorithms under those landscapes. Interestingly, we found a fundamental difference between one- and two-population algorithms: with a single population (individuals are both learners and teachers simultaneously) one cannot solve what is in fact solvable by two populations (teachers and learners have separate goals).

2. INTRANSITIVE GAMES

Two-player games are a natural domain for coevolution: there is no obvious fitness function that could be used to rank game strategies; instead, only paired comparisons are available. Therefore, coevolving game players by having them adapt to each other is relatively straightforward, as opposed to trying to come up with some kind of external evaluator.

A coevolutionary algorithm, working on a two-player game, evolves a set of strategies. The individuals in each generation are used to create a local fitness to select the parents of the next one. This can be done either using a single population ("single popu-

lation coevolution"), or multiple ones. A common setup involves co-evolving two populations, "teachers" and "learners".

In the single population case, the word *coevolution* refers to the fact that individual and landscape are changing together. Two or more populations can be thought of as virtual species changing in response to each other.

Definition 1. Let $\mathcal{Y}$ be an optimization or "genotype" space for evolution, with a suitable mutation operator M. A *test function* (or "coevolutionary game") over $\mathcal{Y}$ is a function $f : \mathcal{Y} \times \mathcal{Y} \to \{-1, 0, 1\}$.

As a fitness evaluator, the role of f is to tell us if a given strategy $y \in \mathcal{Y}$ is heuristically better than $y' \in \mathcal{Y}$ ($f(y, y') < 0$) or not ($f(y, y') \geq 0$). In game terminology, $f(y, y') = -1, 0, 1$ correspond to y winning, drawing and losing against y'.

Definition 2. A test function $f : \mathcal{Y} \times \mathcal{Y} \to -1, 0, 1$ is *symmetric* if for every pair $(y, y') \in \mathcal{Y}$,

$$f(y, y') = -f(y', y) \tag{1}$$

NOTATION 1. *A test function defines a relation $\leq$ in the genotype space $\mathcal{Y}$ by*

$$y \leq y' \Leftrightarrow f(y, y') \geq 0$$

and conversely, any relation $\leq$ on $\mathcal{Y}$ defines a test function by

$$f(y, y') = \begin{cases} 1 & \text{if} & y \leq y' & \text{and not} & y' \leq y \\ -1 & \text{if} & y \geq y' & \text{and not} & y' \geq y \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

In the rest of this paper we use the notations $f(y, y') \geq 0$ (resp. $f(y, y') \leq 0$, $f(y, y') > 0$, $f(y, y') < 0$) and $y \leq y'$ (resp. $y \geq y'$, $y < y'$, $y > y'$) interchangeably.

2.1 Intransitivity

A consequence of having a two-player fitness function is the possibility of internal cycles or "intransitivities", as in the well-known *rock, paper, scissors* game:

f(x,y)	rock	paper	scissors
rock	0	1	-1
paper	-1	0	1
scissors	1	-1	0

Definition 3. A test function is *transitive* if and only if for all $A, B, C \in \mathcal{Y}$

$$f(A, B) = 1 \wedge f(B, C) = 1 \Rightarrow f(A, C) = 1 \tag{3}$$

and *intransitive* if it is not transitive. In other words, f is transitive when the induced relation satisfies

$$A < B \wedge B < C \Rightarrow A < C$$

($<$ is a preorder in this case) so if B beats A and C beats B then C beats A.

2.2 Number Games

As mentioned in the introduction, the first papers in coevolution had to deal with a kind of intransitivity across generations: even though generation $n+1$ is always *better* in some sense than generation n (with respect to the fitness of generation n), it is possible to have later generations that are in fact *worse* than earlier ones.

This issue gave rise to the concept of *Red Queen effect:* the landscape is changing, so it is conceivable that the coevolutionary algorithm, like the Red Queen of Lewis Carrol, is going nowhere fast [5].

"Number games" for the study of coevolutionary dynamics were proposed first by Watson and Pollack and later analyzed by others [2, 3, 7, 9–11]. The idea is that intransitivity in the fitness evaluation, as with an intransitive game for example, can lead to a sequence of generations that do not show overall improvement. A carefully designed algorithm should able to avoid this trap, creating new generations that are in fact better than all previous ones.

The first example was the intransitive numbers game [11, eq. 3]:

$$(x_1, y_1) \leq (x_2, y_2) \iff \begin{cases} |x_1 - x_2| \leq |y_1 - y_2| \text{ and } x_1 \leq x_2 \\ \text{or} \\ |x_1 - x_2| \geq |y_1 - y_2| \text{ and } y_1 \leq y_2 \end{cases} \tag{4}$$

ING was designed to illustrate the fact that, even when there is a straightforward way to "improve" — namely, increasing the values of both x and y — there can be an twist. Two points are compared only on the basis of their most similar dimension. If they are closer on x they will be compared with respect to x and if they are closer in y they will be compared in y. There is no confusion if one player beats the other in both dimensions, but if one pair is better than the other only in one dimension, then the most similar dimensions prevails. This twist leads to intransitivities, for example: $(0,0) \geq (-1,2)$ and $(-1,2) \geq (2,1)$ but $(2,1) \geq (0,0)$.

Bucci [2] studies the intransitive numbers game as well as the *focusing game* he defines, which is not symmetric:

$$f((x_1, y_1), (x_2, y_2)) = 1 \iff \{ x_2 > x_1 \text{ or } y_2 > y1 \tag{5}$$

For contrast we can define the *sum* game, a game that is both transitive and symmetric[1],

$$(x_1, y_1) < (x_2, y_2) \iff x_1 + y_1 < x_2 + y_2 \tag{6}$$

2.3 Positive Regions

Numbers games like the ones above are fully characterized by their *positive regions*.

Definition 4. The positive region of a fixed point $y_0 \in \mathcal{Y}$ is the set of all elements in the genotype space that are "better" than y_0,

$$C^+(y_0) = \{y \in \mathcal{Y} : y_0 < y\} \tag{7}$$

and conversely, the negative region is $C^-(y_0) = \{y \in \mathcal{Y} : y_0 > y\}$

Positive regions for the three games above (and two others) are shown in figure 1.

All the numbers games mentioned so far are uniform, in the sense that the shape of the positive regions is the same in any point in space,

Definition 5. A test function f over $\mathbb{R}^n$ is *uniform* if and only if

$$\forall y, y', z \in \mathcal{Y} : f(y, y') = f(y + z, y' + z) \tag{8}$$

REMARK 2. *A uniform test function f is fully characterized by the positive region of zero. Given a black and white coloring of $\mathcal{Y}$, f is reconstructed by*

$$f(y, y') = \begin{cases} 1 & \text{if } y - y' \text{ is white and } y' - y \text{ is black} \\ -1 & \text{if } y - y' \text{ is black and } y' - y \text{ is white} \\ 0 & \text{otherwise} \end{cases}$$

In fact, all that is required is a coloring of the region M_0 that can be reached from $0 \in \mathcal{Y}$ in a single mutation.
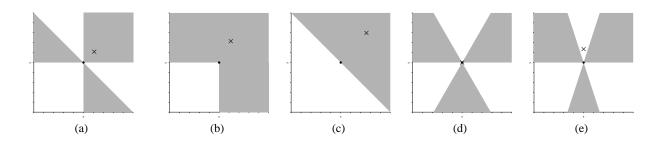
---

[1]Also in [11, eq. 2]

Figure 1: Positive regions: (a) Watson and Pollack's intransitive numbers game (eq. 4); (b) Bucci and Pollack's *focusing game* (eq. 5); (c) s*um game* (eq. 6); (d) *balanced game* and (e) *deceiving game*. A pair (x, y) at the center, is defeated by all strategies in the gray area (positive region), and in turn defeats the strategies in the white area (negative region). Game (b) is asymmetric – A beats B does not imply that B loses to A — because opposite areas do not have opposite colors. (c) is symmetric and transitive. (a, d and e) are symmetric but not transitive. The x marks the center of mass (see theorem 7).

NOTE 3. *In this paper we use uniform random mutations in the unitary ball, so $M_0 = B_0(1) = \{(x,y) : ||(x,y)|| < 1\}$.*

Definition 6. A test function f is a *pie game* if and only if f is defined, as in the symmetric examples above, by a pie-shaped positive region. Formally, there are angles $0 = \alpha_0 < \alpha_1 < \ldots < \alpha_{2k} = 2\pi$ such that $\alpha_{i+k} = \pi + \alpha_i$ and $\forall y_0, y \in \mathbb{R}^2$

$$f(y_0, y) = 1 \Leftrightarrow \exists i : \alpha_{2i} < \theta < \alpha_{2i+1}$$

where $\theta = \angle\{y_0 + (0,1), y_0, y\}$.

3. EVOLUTION OR RANDOM WALK?

In order to do an analytical study of the dynamics of coevolution in the presence of intransitivities, we started with simplest kind of coevolutionary setup possible.

Definition 7. The *coevolutionary hillclimber* $C(1+1)$ is the following algorithm:

Given a fitness test $f : \mathcal{Y} \times \mathcal{Y} \to \{-1, 0, 1\}$

1. Take an initial point $y \in \mathcal{Y}$

2. Repeat forever:

    (a) Generate a random mutation y' of y

    (b) Replace y with y' whenever $y \le y'$

We only need to consider steps in the positive regions of the game. That is, even though all mutations are equally likely to occur, only those which improve the fitness are kept. Therefore it is useful to think only in terms of the steps taken,

Definition 8. $C^*(1+1)$ is the following algorithm: Replace step 2(a) in def. 7 above with

2 (a)' Generate a random mutation y' of y such that y' beats y

In C^* we only increment the generation counter when the mutation has been successful.

We want to analyze if a population is evolving in some sense. With these geometrical problems, there are two notions we can use,

- Victories: newer generations should be better than previous ones.

- Exploration: the newer generations should move within the problem space, as opposed to wandering aimlessly.

The first notion defines progress in terms of the fitness function exclusively. The second notion is a geometric one that is specific to these problems. ING and related problems are interesting precisely because the geometrical aspect gives us a bird's eye view of evolutionary progress.

Definition 9. **Red Queen Property** (RQ) Let $y_0, y_1, \ldots$ be the sequence of the (best) individuals from each generation in a coevolutionary run. The sequence has the *Red Queen Property* if early generations defeat an infinite number of future generations. Formally, for all $k \in \mathbb{N}_0$,

$$\forall n > k : \exists m > n : y_k \text{ defeats } y_m \tag{9}$$

This definition agrees with the notion first proposed by Cliff and Miller [5], that coevolutionary progress can be measured by the increased frequency by which the best individual of the last generation beats the best individuals of earlier generations. We propose to define RQ as the opposite property: progress stagnates and, no matter the number of generations, the evolutionary process keeps generating individuals that are no better than the previous ones.

Definition 10. **Random Walk** (RW) $y_0, y_1 \ldots, y_n, \ldots$ is a *random walk* if it wanders around the starting point without going in any particular direction. Formally,

$$\lim_{n \to \infty} \frac{1}{n} \sum_{k=0}^{n} y_k = 0$$

Random walks have the curious property of *recurrence*: they go back to the starting point an infinite number of times. In fact,

Definition 11. **Recurrent Random Walk** (RW') $y_0, y_1 \ldots$ is a *recurrent random walk* when it visits all open regions U of the space:

$$\forall n \in \mathbb{N}_0, U \subseteq \mathcal{Y} \text{ open}, \exists m \ge n : y_n \in U \tag{10}$$

REMARK 4. *For uniform test functions in $\mathbb{R}^n$, each point starts an identical stochastic process, thus the RW and RQ properties above are valid for all $k \ge 0$ if and only if they are valid for a single k.*

Number games have been defined geometrically. We intuitively feel that progress must mean movement in some direction (as opposed to just wandering around the starting point).

THEOREM 5. *In pie games, RW $\Rightarrow$ RQ (with probability one)*

Proof. Let Y_n be a random variable representing the n-th generation element of $C^*(1+1)$ and let $X_n = Y_{n+1} - Y_n$. Then $\{X_n\}_{n \in N}$ is a sequence of independent, identically distributed random variables. The strong law of large numbers says that $P(\frac{1}{n}Y_n \to E(X)) = 1$ (where $X = X_0$). Therefore, if RW is observed, with probability 1 we know that $E(X) = 0$. Furthermore, by the central limit theorem, Y_n converges in distribution to $N(0, \sigma\sqrt{n})$, where σ is the standard deviation of X.

Consider the negative region $C^- = C^-(0)$. The probability $P(Y_n \in C^-)$ of Y_n falling in the negative region converges to $P(N \in \sqrt{n}C^-)$ for N normal with mean 0 and variance σ^2. Therefore, if there is an $\varepsilon > 0$ such that $P(N \in \sqrt{n}C^-) > \varepsilon$ for all n, then $P(Y_n \in C^-) > \frac{\varepsilon}{2}$ for all n greater than some n_0. Thus the probability that Y_n lies within C^- for some $n \geq n_0$ is 1.

In pie games, negative (and also positive) regions are circular sectors, and thus $\sqrt{n}C^- = C^-$, concluding the proof.

PROPOSITION 6. *For test functions with open negative regions, RW' $\implies$ RQ (with probability one).*

Let y_k be such that $y_k > y_0$. Since y_0 is in the negative region of y_k, which is open, there is an open neighborhood of y_0 contained in the negative region of y_k. A strong random walk visits such neighborhood an infinite number of times, therefore y_k defeats an infinite number of future generations.

THEOREM 7. *Let f be a pie-shaped numbers game in $\mathbb{R}^2$. Then the algorithm C(1+1):*

1. *Is RW if and only if the center of mass is zero*

2. *Moves in the direction of the center of mass with speed proportional to its norm.*

3. *If f is a pie game, it is RQ if and only if the center of mass lies outside C^+*

Definition 12. We define the *center of mass* as the expected value $E(X)$ of the mutation operator

$$E(X) = \int_{(x,y) \in C^+ \cap B_0(1)} (y,x) dx dy \qquad (11)$$

Proof of theorem 7. As shown in the proof to theorem 5, a random walk happens (with probability one) if and only if $E(X) = 0$.

For part 2, the central limit theorem tells us that Y_n converges in probability to $N(\mu, \sigma\sqrt{n})$, with $\mu = E(X)$. Therefore, the expected location Y_n of the algorithm after n steps is $n\mu$ with a variance proportional to $\sqrt{n}$.

Finally, if the center of mass is in the interior of the winning region for Y_0, the distance between $nE(X)$ and the winning region grows linearly, but the variance only with the square root of n, therefore the probability of the process visiting the losing region approaches zero. Conversely, if the center of mass is inside the losing region of zero, then the probability that Y_n is beaten by zero approaches 0.

REMARK 8. *For a general class of uniform test functions in $\mathbb{R}$ and $\mathbb{R}^2$ (with open negative regions), the convergence to $N(0, \sigma\sqrt{n})$, implies that the steps taken by the algorithm vibrate uniformly in all directions, generating a stochastic process with the recurrence property, visiting every open set in the plane (see [1,4]). Therefore, part 1 of theorem 7 is in fact true for RW' and for uniform test functions with open negative regions.*

3.1 Comments

Theorem 7 shows that the basic coevolutionary algorithm moves in a reasonable direction, with linear speed, avoiding the Red Queen Effect, for most problems. That is, problems where the center of mass is located inside a positive region.

Well-behaved problems include the original ING and Bucci's focusing game and of course, all transitive problems. Looking back at figure 1, we can characterize the behavior of a pie game by looking at the position of x that indicates the center of mass. "Good problems" are those for which the x is inside a dark region.

This supports the conclusion of [6] in the sense that intransitivity in the problem domain does not pose such a big problem. In reality, generic problems avoid the RW/RQ effects and simply progress towards the most promising direction.

However, pathological problems with center of mass at zero do lead to hopeless random walks. Problems with the center of mass in the wrong place (e.g. the "deceiving" problem of fig. 1e), confused by opposing intransitivities, go deep into the losing region. This is the worst kind of intransitivity, with the probability of losing vs. earlier generations approaching 1.

Can we use larger populations to escape from RW/RQ behavior? Below we analyze two simple algorithms: $C(1+N)$, where instead of a random opponent it is the winner of a tournament of mutants that is chosen, and also $C_<(N+1)$ which tries to escape RQ by means of a finite archive.

4. LARGER POPULATIONS

4.1 C(1+N)

A straightforward improvement to $C(1+1)$ is to generate several mutants, instead of one, and choose the winner of a tournament as the next generation. This will have the effect of concentrating the probabilities into the largest winning sector, provided there is one, and the periphery of the mutation region (figure 2).

This could turn a deceptive problem (fig 1e) into a non-deceptive one (provided one of the regions is larger than the other) but it cannot break the symmetry when the center of mass is at the center.

We have left the details of the proof that the probability measure concentrates in the largest sector out for reasons of space and clarity. However, it should be obvious that, because of symmetry, no matter how many mutants are chosen, an initially symmetrical situation will remain symmetric, as in figure 2 (bottom row), and therefore, not much can be gained from $C(1+N)$ when the problem is balanced.

4.2 C_<(N+1)

Keeping a finite archive seems like a good idea: perhaps we can break RQ by remembering some of the previous steps, and then moving forward if and only if we know that we can beat all the previous cases. The question is: is there a memory size large enough that it can guarantee monotonic progress?

Definition 13. $C_<(N+1)$ is the following algorithm: given a set of initial points $\{y_1, ..., y_N\}$

1. Generate a mutant y_{N+1} of y_N such that y_{N+1} beats y_i for all $1 \leq i \leq N$

2. Replace $\{y_1, ..., y_N\}$ with $\{y_2, ..., y_{N+1}\}$.

THEOREM 9. *If f is a pie-shaped, symmetric game such that $C(1+1)$ is a random walk, then $C_<(1+N)$ is also RW for all N.*
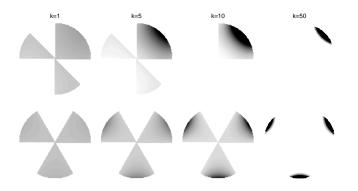
Figure 2: C(1+N): Probability density of beating (0,0) plus *k* random opponents for increasing *k* in ING (top row) and balanced game (bottom row). The probabilities of movement become more concentrated for larger populations, but remain equilibrated in the balanced case, leading to a random walk.
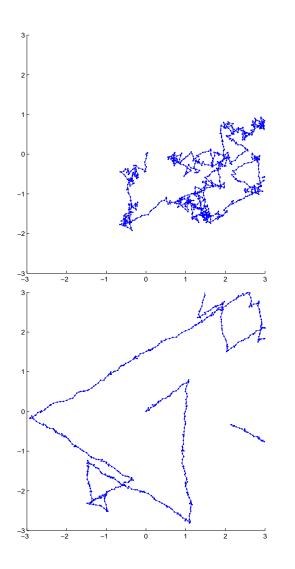

Figure 3: Sample runs of $C_<(N+1)$ for a balanced problem. $N=2$(top) and $N=5$ (bottom). Direction changes become less frequent with higher N but do not disappear.

Proof. If the problem is symmetric, the center of mass E_N of $C_<(1+N)$ is always 0 because the algorithm is unbiased. In other words, if $C_i, i=1,\ldots m$ are the disjoint sectors or "cones" that make up C^+, the algorithm must choose any of them with equal probability.

It is conceivable that the memory mechanism could, in spite of initial symmetry, lead to monotonic improvement. Once the first few moves are made, the remainder of the algorithm could manage to "choose" a single direction and stick to it.

If, however, notwithstanding how many steps have been taken in a given direction, there remains a non-vanishing probability of switching to a different direction, then such direction switches will continue to occur forever, and the algorithm is a random walk.

LEMMA 10. *In the conditions of theorem 9, the probability of changing directions does not approach zero.*

Proof. Let y_k be the k-th step in a $C_<(N+1)$ run. We can assume without loss of generality that the current direction of movement is C_1. y_{k+N} must beat y_k, so $y_{k+N} \in C^+ + y_k$. The algorithm continues in the "good" direction if and only if $y_{k+N} \in C_1 + y_k$, but it switches directions if y_{k+N} falls in a different subsector of $C^+ + y_k$. Without memory, we would have the probability of X_{k+1} lying in "bad" sector C_2 proportional to the relative area of C_2. We can write this as: $P_1(X_{k+1} \in C_2) = p_2 > 0$, (with X_i defined as in the proof to theorem 5), where P_i denotes the probability measure associated with the algorithm $C_<(1+i)$.

Assuming all previous steps also went in the correct direction C_1, the region $C_1 + y_k$ is contained in $C_1 + y_{k-j}$ for all $1 \leq j \leq N$, because they are a succession of pie slices contained in each other.

Let $p = P_1\{X \in B_{\varepsilon/N}(0) \cap C_1\}$ for a given ε. The probability that by chance, all steps $k+1,\ldots,k+N$ are all very short, advancing ε or less, is p^N or more, and therefore, the probability that N steps together advance a distance of no more than $N\varepsilon$ is at least p^N: $p^N \leq P_N\{X_{k+1} + \ldots + K_{k+N-1} \in B_{N\varepsilon}(y_k)\}$. p^N is a lower bound for the probability of taking a succession of small steps, so that we remain very close to y_k after N more steps.

Take ε small enough so that $P_1(X + u \in C_2) \geq \frac{1}{2}p_2 \forall u \in B_{N\varepsilon}(0)$. Therefore, $P(X_{k+N} \in C_2) \geq \frac{1}{2}p^N p_2$. This lower bound does not depend on either k nor y_k and we conclude that a change of direction occurs with a non-vanishing frequency.

Therefore, $C_<(N+1)$ is RW for all N (see fig 3).

5. TEACHER-LEARNER DYNAMICS

De Jong [8] described the *Delphi* algorithm, based in two populations called *learners* and *evaluators*. These populations coevolve against each other using Pareto dominance as the fitness criterion. Delphi, as opposed to the algorithms described above, is able to break the symmetry and make steady progress by focusing in a single positive sector, thus avoiding the Red Queen effect.

We introduce a minimalistic teacher-learner setup in order to show the essence of the symmetry-breaking mechanism involved.

Definition 14. TL(2,1) is the following algorithm:

1. Initialize $T_0 = T_1 = L$ to the same starting point in $\mathcal{Y}$

2. Generate random mutants T_0', T_1', L' of each member

3. Replace (T_0, T_1, L) with (T_0', T_1', L') If and only if:

    (a) L' beats both T_0' and T_1', and
    (b) T_0' and T_1' both defeat L

The difference between teachers and learners here is that learners are generalists (they must defeat both teachers) but teachers are allowed to specialize, beating the learner but not the other teachers.

The requirement that a new learner, better than the two new teachers, must be found simultaneously is given because, otherwise, teachers could move too fast, creating an impossible situation for the learner.

The TL algorithm enforces learnability by moving only when it can find two new challenges and a solution at the same time, keeping the teachers in the "proximal zone".

THEOREM 11. *TL(2,1) is not RQ for any pie problems.*

Proof. We offer only the sketch of a proof.

The algorithm works because at some point the teachers end up delimiting a single region for the learner, from far away.

Assuming the teachers and learners are away from each other, a *disengaged* situation occurs when the movements of the three participants can only take place within a single cone of their counterparts' positive region (fig. 4).

If such disengaged situation is reached, the directions of movement for the three members are limited to one particular, transitive subregion for each one of them. Without reachable intransitivities, the RQ property is avoided as long as disengagement is maintained.

Given that T_1, T_2 are defeated by L, but move to defeat it in the next step, we can assume without loss of generality the position of T_1 and look at where T_2 can be located with respect to T_1 and L. There are five cases, labelled A through E in the figure. The corresponding moves for L are labelled a through e.

Observe that in all cases, L necessarily moves within a single sector of the positive region, and cannot change while the same configuration remains, because the alternative sectors are farther than one step away.

At least one of the configurations *(E)* is stable in the sense that the average direction of movement of T_1, T_2 and L takes them move away from each other, reducing with each step the probability that they will meet again in the future.

The remaining configurations are unstable, because the agents do not tend to separate, and will eventually produced new engaged situations, where agents are close to each other.

All the "engaged" configurations are unstable, because when the characters are close to each other their positive regions overlap in several subsectors, leading to opportunities for moving in all directions.

The stable configuration thus is attractive, and the eventually the system must find it and remain there.
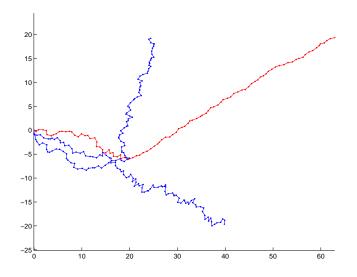

Figure 5: Sample run of TL(2,1) in a uniform intransitive problem. Teachers' trajectories are shown in black, and learner in gray. The direction of movement changes a few times in the beginning but, once the system falls into a stable configuration (see text), they continue making improvements along a single positive sector. Delphi [8] has a similar specialization property.

6. CONCLUSIONS

The notion that intransitivity between generations, due to coevolution between population and fitness landscape (Red Queen effect), can originate from a problem domain that is intransitive, is not always true, as is shown in [8]: there can be intransitive generations without intransitive problems, and intransitive problems do not necessarily lead to intransitive generations.

Here we studied a family of intransitive problems in the euclidean plane, inspired Watson and Pollack's Intransitive Numbers Game. We analyzed the simplest algorithm first, which is the coevolutionary hillclimber, showing that in general, it is not fooled by the intransitivities in the domain.

We characterized the pathological problems of the "Watson family", namely, balanced and deceiving cases (*d* and *e* in fig. 1) investigating the behavior of other coevolutionary algorithms in those cases.

We found the interesting property that finite memory algorithms (section 4.2) are equally incapable of dealing with a balanced situation. They also generate random walks in spite of their (finite) archive of previous states. It is therefore not sufficient to have $A_1 < A_2 < \ldots < A_k$, for any k, to guarantee that $A_1 < A_{k+1}$. Even more interesting is the fact that the simplest teacher-learner coevolutionary algorithm, with just one teacher and two learners, is sufficient to break the Red Queen effect, producing sustainable improvement.

This led us to an important conclusion: there is a qualitative difference between one- and two-population coevolution. The teacher-learner setup is fundamentally different from the learner-learner case because teachers can focus in subareas, subdimensions of the problem, in ways that learners cannot, because they are required to be generalists.

Pie-shaped games do have a claim to generality. If we assume that fitness comparisons originate in a continually differentiable test function $f(y_1, y_2) : \mathbb{R}^k \times \mathbb{R}^k \to \mathbb{R}$, then the implicit function theorem implies that, in a neighborhood around y_1, there is a $k-1$-
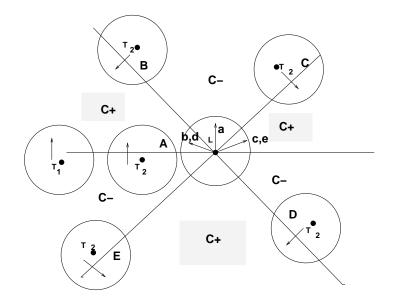
Figure 4: Disengaged situations for TL(2,1). The learner L is at the center, with the radial lines delimiting the different positive (C+) and negative (C-) regions. Both teachers T_1, T_2 are currently inside C-, but their mutants jump to C+ in the direction of the arrows. T_2 can be located in 5 different regions A, B, C, D, E with respect to T_1. The corresponding directions of movement for L are labeled a, b, c, d, e. The only stable configuration is E, where all distances tend to increase.

dimensional submanifold that divides the space between winners ($f(y_1, y) > 0$) and losers ($f(y_1, y) < 0$).

Therefore, if f is differentiable at y_1, then the problem is locally transitive and shaped as in figure 1c. Local intransitivities correspond to saddle points of f, and are shaped like the pies we have studied here. Uniform intransitive problems are extreme in this respect, because the test function is saddle-shaped in a large area, not just a in a single-point singularity.

It remains to be seen whether it is possible to design an algorithm that can discover the alternative directions of progress, and explore them all. Both our TL(2,1) algorithm and de Jong's Delphi [8] succeed because their are able to choose a single direction and maintain it. They are incapable, however, of discovering and exploring the other alternatives. A direct application of Pareto techniques does not help in this case because intransitive number games are non-dominated in all points in the plane, thus the Pareto front can grow indefinitely without ever leaving a small area around the starting point — creating a "Brownian dust", a cloud of particles performing random walks.

7. ACKNOWLEDGMENTS

8. REFERENCES

[1] L. Breiman. *Probability*. Society for Industrial and Applied Mathematics, 1992.

[2] A. Bucci and J. Pollack. Focusing versus intransitivity: geometrical aspects of coevolution. In e. a. Cantú-Paz, E., editor, *Genetic and Evolutionary Computation - GECCO 2003*, number 2724 in LNCS, Berlin, 2003. Springer-Verlag.

[3] A. Bucci and J. B. Pollack. A mathematical framework for the study of coevolution. In *Foundations of Genetic Algorithms (FOGA VII)*, 2003.

[4] Chung and Fuchs. On the distribution of values of sums of random variables. *Mem. of the American Mathematical Society*, 6, 1951.

[5] D. Cliff and G. Miller. Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In F. Morán, A. Moreno, J. J. Merelo, and P. Chacón, editors, *Advances in Artificial Life: Third European Conference on Artificial Life*, number 929 in Lecture Notes in Computer Science, pages 200–218, Berlin, New York, 1995. Springer.

[6] E. De Jong. Intransitivity in coevolution. In *Proceedings of the 8th International Conference on Parallel Problem Solving from Nature PPSN-04*, pages 843–851, 2004.

[7] E. D. De Jong and J. B. Pollack. Learning the ideal evaluation function. In *Proceedings of GECCO 2003*, pages 277–288, 2003.

[8] E. D. de Jong and J. B. Pollack. Ideal evaluation from coevolution. *Evolutionary Computation*, 12(2), 2004.

[9] S. G. Ficici and J. B. Pollack. Game theory and the simple coevolutionary algorithm: Some preliminary results on fitness sharing. In *GECCO 2001 Workshop on Coevolution: Turning Adaptive Algorithms upon Themselves*, 2001.

[10] L. Panait and S. Luke. Comparison of two competitive fitness functions. In W. B. Langdon, editor, *GECCO-2002: Proceedings of the Genetic and Evolutionary Computation Conference*, pages 503–511. Morgan Kauffman, 2002.

[11] R. Watson and J. B. Pollack. Coevolutionary dynamics in a minimal substrate. In . e. a. Spector, L, editor, *Proceedings of the 2001 Genetic and Evolutionary Computation Conference*. Morgan Kaufmann, 2001.