

Fast Population Game Dynamics for Dominant Sets and Other Quadratic Optimization Problems

Samuel Rota Bulò¹, Immanuel M. Bomze², and Marcello Pelillo¹

¹ Dipartimento di Informatica - Univ. of Venice - Italy
`{srotabul,pelillo}@dsi.unive.it`

² Department of Statistics and Decision Support Systems - Univ. of Vienna - Austria
`immanuel.bomze@univie.ac.at`

Abstract. We propose a fast population game dynamics, motivated by the analogy with infection and immunization processes within a population of “players,” for finding dominant sets, a powerful graph-theoretical notion of a cluster. Each step of the proposed dynamics is shown to have a linear time/space complexity and we show that, under the assumption of symmetric affinities, the average population payoff is strictly increasing along any non-constant trajectory, thereby allowing us to prove that dominant sets are asymptotically stable (i.e., attractive) points for the proposed dynamics. The approach is general and can be applied to a large class of quadratic optimization problems arising in computer vision. Experimentally, the proposed dynamics is found to be orders of magnitude faster than and as accurate as standard algorithms.

1 Introduction

Dominant sets are a graph-theoretical notion of a cluster [1], which have found application in problems as diverse as the analysis of fMRI data [2], content-based image retrieval [3], detection of anomalous activities in video streams [4], bioinformatics [5], human action recognition [6] and matching problems [7,8].

Computationally, the standard approach to finding dominant sets in an edge-weighted graph is to use *replicator dynamics*, a class of evolutionary game-theoretic algorithms inspired by Darwinian selection processes. However, a typical problem associated with these algorithms is the scaling behavior with the number of data. On a dataset containing N examples, the computational complexity of each replicator dynamics step is $\mathcal{O}(N^2)$, thereby hindering their applicability to problems involving very large data sets, such as high-resolution imagery and spatio-temporal data.

In order to avoid this drawback, in this paper we propose a new population game dynamics for finding dominant sets which turns out to be dramatically faster and even more accurate than standard approaches from evolutionary game theory. Our approach is motivated by the analogy with infection and immunization processes within a population of “players.” The selection mechanism

governing our dynamics iteratively performs an infection step, which consists of spreading (or suppressing) the most successful (unsuccessful) strategies in the population. The infection phase is then protracted as long as the selected “infective” strategy performs better (or worse, if not extinct) than the average population’s payoff. As opposed to standard techniques, such as the replicator dynamics or best-response dynamics, which can be considered interior-point methods, our algorithm resembles a vertex-pivoting method. Each step of the proposed dynamics is shown to have a linear time/space complexity and we show that, under the assumption of symmetric affinities, the average population payoff is strictly increasing along any non-constant trajectory, thereby allowing us to prove that dominant sets (i.e., ESS equilibria of the underlying “grouping game” [9]) are asymptotically stable points for the proposed dynamics.

We provide experimental evidence that the proposed algorithm is orders of magnitude faster than standard dynamics on two computer vision applications, namely image segmentation and region-based hierarchical image matching, while preserving the quality of the solutions found.

Although the main focus in this paper is dominant sets, we note that the proposed approach is general and can be applied to a large class of optimization problem, instances of which abound in computer vision and pattern recognition (e.g., graph matching, stereo matching, image labeling, etc.).

2 Basics of Evolutionary Game Theory

Evolutionary game theory considers an idealized scenario whereby pairs of individuals are repeatedly drawn at random from a large, ideally infinite, population to play a symmetric two-player game. Let $O = \{1, \dots, n\}$ be the set of *pure strategies* available to the players and let A be the $n \times n$ payoff or utility matrix [10], where a_{ij} is the payoff that a player gains when playing the strategy i against an opponent playing strategy j . A *mixed strategy* is a probability distribution $\mathbf{x} = (x_1, x_2, \dots, x_n)^\top$ over the available strategies in O . Mixed strategies lie in the standard simplex Δ of the n -dimensional Euclidean space, which is defined as

$$\Delta = \left\{ \mathbf{x} \in \mathbb{R}^n : \sum_{i=1}^n x_i = 1 \text{ and } x_i \geq 0, i = 1, \dots, n \right\}.$$

We denote by \mathbf{e}^i the i th column of the identity matrix. The *support* of a mixed strategy $\mathbf{x} \in \Delta$, denoted by $\sigma(\mathbf{x})$, defines the set of elements with non-zero probability: $\sigma(\mathbf{x}) = \{i \in O : x_i > 0\}$. The expected payoff that a player obtains by playing the pure strategy i against an opponent playing a mixed strategy \mathbf{x} is $\pi(\mathbf{e}^i | \mathbf{x}) = (A\mathbf{x})_i = \sum_j a_{ij}x_j$. Hence, the expected payoff received by adopting a mixed strategy \mathbf{y} is given by $\pi(\mathbf{y} | \mathbf{x}) = \mathbf{y}^\top A\mathbf{x}$ while the population expected payoff is $\pi(\mathbf{x}) = \pi(\mathbf{x} | \mathbf{x}) = \mathbf{x}^\top A\mathbf{x}$. For notational compactness, in the sequel we will write $\pi(\mathbf{y} - \mathbf{x} | \mathbf{z})$ for the payoff difference $\pi(\mathbf{y} | \mathbf{z}) - \pi(\mathbf{x} | \mathbf{z})$, and $\pi(\mathbf{y} - \mathbf{x})$ for $\pi(\mathbf{y} - \mathbf{x} | \mathbf{y}) - \pi(\mathbf{y} - \mathbf{x} | \mathbf{x})$.

A mixed strategy \mathbf{x} is a (*symmetric Nash equilibrium strategy*) if for all $\mathbf{y} \in \Delta$, we have $\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) \leq 0$. This implies that $\pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x}) \leq 0$ for all $i \in O$, which in turn implies that $\pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x}) = 0$ for all $i \in \sigma(\mathbf{x})$. Hence, the payoff is constant across all (pure) strategies in the support of \mathbf{x} , while all strategies outside the support of \mathbf{x} earn a payoff that is less than or equal $\pi(\mathbf{x})$.

A strategy \mathbf{x} is said to be an *Evolutionary Stable Strategy (ESS)* if it is a Nash strategy (*equilibrium condition*) and for all $\mathbf{y} \in \Delta \setminus \{\mathbf{x}\}$ satisfying $\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) = 0$ we have $\pi(\mathbf{y} - \mathbf{x}|\mathbf{y}) < 0$ (*stability condition*). Intuitively, ESS's are strategies such that any small deviation from them will lead to an inferior payoff. ESS's can be found by *replicator dynamics (RD)*, a classic formalization of a natural selection process [10].

3 Dominant Sets and Their Characterizations

The dominant set framework is a pairwise clustering approach [1] that is based on the notion of a dominant set, which can be seen as an edge-weighted generalization of a clique. The framework is based on a recursive characterization of the weight $W_S(i)$ of element i with respect to a set S of elements, and characterizes a group as a *dominant set*, i.e., a set that satisfies:

1. $W_S(i) > 0$, for all $i \in S$,
2. $W_{S \cup \{i\}}(i) < 0$, for all $i \notin S$.

These conditions correspond to the two main properties of a cluster: the first regards internal homogeneity, whereas the second regards external heterogeneity.

The characteristic vector \mathbf{x}^S of a set $S \subseteq V$ is defined as

$$x_i^S = \begin{cases} \frac{W_S(i)}{W(S)} & \text{if } i \in S, \\ 0 & \text{otherwise.} \end{cases}$$

The following result establishes a one-to-one correspondence between ESS's and dominant sets [9].

Theorem 1. *If $S \subseteq V$ is a dominant set with respect to affinity matrix A , then \mathbf{x}^S is an ESS for a two-player game with payoff matrix A .*

Conversely, if \mathbf{x} is an ESS for a two-person game with payoff matrix A , then $S = \sigma(\mathbf{x})$ is a dominant set with respect to A , provided that $W_{S \cup \{i\}}(i) \neq 0$ for all $i \notin S$.

Under the assumption of a symmetric affinity matrix A there exists a one-to-one correspondence between dominant sets and the (strict) local solutions of the following so-called standard quadratic program (StQP) [1]:

$$\max \{ \mathbf{x}^\top A \mathbf{x} : \mathbf{x} \in \Delta \} . \tag{1}$$

4 A New Class of Evolutionary Dynamics

Let $\mathbf{x} \in \Delta$ be the *incumbent* population state, \mathbf{y} be the *mutant* population invading \mathbf{x} and let $\mathbf{z} = (1 - \varepsilon)\mathbf{x} + \varepsilon\mathbf{y}$ be the population state obtained by injecting into \mathbf{x} a small share of \mathbf{y} -strategists. The *score function* of \mathbf{y} versus \mathbf{x} [11] is given by:

$$h_{\mathbf{x}}(\mathbf{y}, \varepsilon) = \pi(\mathbf{y} - \mathbf{x}|\mathbf{z}) = \varepsilon\pi(\mathbf{y} - \mathbf{x}) + \pi(\mathbf{y} - \mathbf{x}|\mathbf{x}).$$

Following [12], we define the (*neutral*) *invasion barrier* $b_{\mathbf{x}}(\mathbf{y})$ of $\mathbf{x} \in \Delta$ against any mutant strategy \mathbf{y} as the largest population share $\varepsilon_{\mathbf{y}}$ of \mathbf{y} -strategists such that for all smaller positive population shares ε , \mathbf{x} earns a higher or equal payoff than \mathbf{y} in the post-entry population \mathbf{z} . Formally:

$$b_{\mathbf{x}}(\mathbf{y}) = \inf(\{\varepsilon \in (0, 1) : h_{\mathbf{x}}(\mathbf{y}, \varepsilon) > 0\} \cup \{1\}).$$

Given populations $\mathbf{x}, \mathbf{y} \in \Delta$, we say that \mathbf{x} is *immune* against \mathbf{y} if $b_{\mathbf{x}}(\mathbf{y}) > 0$. Trivially, a population is always immune against itself. Note that, \mathbf{x} is immune against \mathbf{y} if and only if either $\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) < 0$ or $\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) = 0$ and $\pi(\mathbf{y} - \mathbf{x}) \leq 0$. If $\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) > 0$ we say that \mathbf{y} is *infective* for \mathbf{x} . We denote the set of infective strategies for \mathbf{x} as

$$\Upsilon(\mathbf{x}) = \{\mathbf{y} \in \Delta : \pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) > 0\}.$$

Consider $\mathbf{y} \in \Upsilon(\mathbf{x})$; clearly, this implies $b_{\mathbf{x}}(\mathbf{y}) = 0$. If we allow for invasion of a share ε of \mathbf{y} -strategists as long as the score function of \mathbf{y} versus \mathbf{x} is positive, at the end we will have a share of $\delta_{\mathbf{y}}(\mathbf{x})$ mutants in the post-entry population, where

$$\delta_{\mathbf{y}}(\mathbf{x}) = \inf(\{\varepsilon \in (0, 1) : h_{\mathbf{x}}(\mathbf{y}, \varepsilon) \leq 0\} \cup \{1\}).$$

Note that if \mathbf{y} is infective for \mathbf{x} , then $\delta_{\mathbf{y}}(\mathbf{x}) > 0$, whereas if \mathbf{x} is immune against \mathbf{y} , then $\delta_{\mathbf{y}}(\mathbf{x}) = 0$. Since score functions are (affine-)linear, there is a simpler expression $\delta_{\mathbf{y}}(\mathbf{x}) = \min\left[\frac{\pi(\mathbf{x} - \mathbf{y}|\mathbf{x})}{\pi(\mathbf{y} - \mathbf{x})}, 1\right]$, if $\pi(\mathbf{y} - \mathbf{x}) < 0$, and $\delta_{\mathbf{y}}(\mathbf{x}) = 1$, otherwise.

Proposition 1. *Let $\mathbf{y} \in \Upsilon(\mathbf{x})$ and $\mathbf{z} = (1 - \delta)\mathbf{x} + \delta\mathbf{y}$, where $\delta = \delta_{\mathbf{y}}(\mathbf{x})$. Then $\mathbf{y} \notin \Upsilon(\mathbf{z})$.*

The proof of this result is straightforward by linearity and can be found, e.g., in [13].

The core idea of our method is based on the fact that $\mathbf{x} \in \Delta$ is a Nash equilibrium if and only if $\Upsilon(\mathbf{x}) = \emptyset$ (we prove this in Theorem 2). Therefore, as long as we find a strategy $\mathbf{y} \in \Upsilon(\mathbf{x})$, we update the population state according to Proposition 1 in order obtain a new population \mathbf{z} such that $\mathbf{y} \notin \Upsilon(\mathbf{z})$ and we reiterate this process until no infective strategy can be found, or in other words, a Nash equilibrium is reached.

The formalization of this process provides us with a class of new dynamics which, for evident reasons, is called *Infection and Immunization Dynamics* (INIMDYN):

$$\mathbf{x}^{(t+1)} = \delta_{\mathcal{S}(\mathbf{x}^{(t)})}(\mathbf{x}^{(t)})[\mathcal{S}(\mathbf{x}^{(t)}) - \mathbf{x}^{(t)}] + \mathbf{x}^{(t)}. \tag{2}$$

Here, $\mathcal{S} : \Delta \rightarrow \Delta$ is a generic *strategy selection* function which returns an infective strategy for \mathbf{x} if it exists, or \mathbf{x} otherwise:

$$\mathcal{S}(\mathbf{x}) = \begin{cases} \mathbf{y} & \text{for some } \mathbf{y} \in \mathcal{Y}(\mathbf{x}) \text{ if } \mathcal{Y}(\mathbf{x}) \neq \emptyset, \\ \mathbf{x} & \text{otherwise.} \end{cases} \tag{3}$$

By running these dynamics we aim at reaching a population state that can not be infected by any other strategy. In fact, if this is the case, then \mathbf{x} is a Nash strategy, which happens if and only if it is fixed (i.e., stationary) under dynamics (2):

Theorem 2. *Let $\mathbf{x} \in \Delta$ be a strategy. Then the following statements are equivalent:*

- (a) $\mathcal{Y}(\mathbf{x}) = \emptyset$: there is no infective strategy for \mathbf{x} ;
- (b) \mathbf{x} is a Nash strategy;
- (c) \mathbf{x} is a fixed point under dynamics (2).

Proof. A strategy \mathbf{x} is a Nash strategy if and only if $\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) \leq 0$ for all $\mathbf{y} \in \Delta$. This is true if and only if $\mathcal{Y}(\mathbf{x}) = \emptyset$. Further, $\delta = 0$ implies $\mathcal{S}(\mathbf{x}) = \mathbf{x}$. Conversely, if $\mathcal{S}(\mathbf{x})$ returns \mathbf{x} , then we are in a fixed point. By construction of $\mathcal{S}(\mathbf{x})$ this happens only if there is no infective strategy for \mathbf{x} .

The following result shows that average payoff is strictly increasing along any non-constant trajectory of the dynamics (2), provided that the payoff matrix is symmetric.

Theorem 3. *Let $\{\mathbf{x}^{(t)}\}_{t \geq 0}$ be a trajectory of (2). Then for all $t \geq 0$,*

$$\pi(\mathbf{x}^{(t+1)}) \geq \pi(\mathbf{x}^{(t)}),$$

with equality if and only if $\mathbf{x}^{(t)} = \mathbf{x}^{(t+1)}$, provided that the payoff matrix is symmetric.

Proof. Again, let us write \mathbf{x} for $\mathbf{x}^{(t)}$ and δ for $\delta_{\mathcal{S}(\mathbf{x})}(\mathbf{x})$. As shown in [13], we have

$$\pi(\mathbf{x}^{(t+1)}) - \pi(\mathbf{x}^{(t)}) = \delta [h_{\mathbf{y}}(\mathbf{x}, \delta) + \pi(\mathbf{y} - \mathbf{x}|\mathbf{x})] .$$

If $\mathbf{x}^{(t+1)} \neq \mathbf{x}^{(t)}$, then \mathbf{x} is no Nash strategy, and $\mathbf{y} = \mathcal{S}(\mathbf{x})$ returns an infective strategy. Hence $\delta > 0$ and

$$h_{\mathbf{y}}(\mathbf{x}, \delta) + \pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) \geq \pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) > 0$$

(in fact, if $\delta < 1$, then even $h_{\mathbf{y}}(\mathbf{x}, \delta) = 0$), so that we obtain a strict increase of the population payoff. On the other hand, if $\pi(\mathbf{x}^{(t+1)}) = \pi(\mathbf{x}^{(t)})$, then the above equation implies $\delta = 0$ or $h_{\mathbf{x}}(\mathbf{x}, \delta) = \pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) = 0$, due to nonnegativity of both quantities above. In particular, we have $\delta = 0$ or $\pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) = 0$. In both cases, $\mathbf{y} = \mathcal{S}(\mathbf{x})$ cannot be infective for \mathbf{x} . Thus $\mathcal{Y}(\mathbf{x}) = \emptyset$ and \mathbf{x} must be a fixed point, according to Theorem 2. This establishes the last assertion of the theorem.

Theorem 3 shows that by running INIMDYN , under symmetric payoff function, we strictly increase the population payoff until we reach a Nash equilibrium at a fixed point. This of course holds for any selection function $\mathcal{S}(\mathbf{x})$ satisfying (3). However, the way we choose $\mathcal{S}(\mathbf{x})$ may affect the efficiency of the dynamics. The next section introduces a particular selection function that leads to a well-performing dynamics for our purposes.

5 A Pure Strategy Selection Function

Depending on how we choose the function $\mathcal{S}(\mathbf{x})$ in (2), we may obtain different dynamics. One in particular, which is simple and leads to nice properties, consists in allowing only infective pure strategies.

Given a population \mathbf{x} , we define the co-strategy of \mathbf{e}^i with respect to \mathbf{x} as

$$\overline{\mathbf{e}}^i_{\mathbf{x}} = \frac{x_i}{x_i - 1}(\mathbf{e}^i - \mathbf{x}) + \mathbf{x}.$$

Note that if $\pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x}) \neq 0$ then either $\mathbf{e}^i \in \Upsilon(\mathbf{x})$ or $\overline{\mathbf{e}}^i_{\mathbf{x}} \in \Upsilon(\mathbf{x})$.

Consider the strategy selection function $\mathcal{S}_{Pure}(\mathbf{x})$, which finds a pure strategy i maximizing $|\pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x})|$, and returns \mathbf{e}^i , $\overline{\mathbf{e}}^i_{\mathbf{x}}$ or \mathbf{x} according to whether $\pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x})$ is positive, negative or zero. Let $\mathcal{M}(\mathbf{x})$ be a pure strategy such that

$$\mathcal{M}(\mathbf{x}) \in \arg \max_{i=1,\dots,n} |\pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x})|.$$

Then $\mathcal{S}_{Pure}(\mathbf{x})$ can be written as

$$\mathcal{S}_{Pure}(\mathbf{x}) = \begin{cases} \mathbf{e}^i & \text{if } \pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x}) > 0 \text{ and } i = \mathcal{M}(\mathbf{x}) \\ \overline{\mathbf{e}}^i_{\mathbf{x}} & \text{if } \pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x}) < 0 \text{ and } i = \mathcal{M}(\mathbf{x}) \\ \mathbf{x} & \text{otherwise.} \end{cases}$$

Note that the search space for an infective strategy is reduced from Δ to a finite set. Therefore, it is not obvious that $\mathcal{S}_{Pure}(\mathbf{x})$ is a well-defined selection function, i.e., it satisfies (3). The next theorem shows that indeed it is.

Proposition 2. *Let $\mathbf{x} \in \Delta$ be a population. There exists an infective strategy for \mathbf{x} , i.e., $\Upsilon(\mathbf{x}) \neq \emptyset$, if and only if $\mathcal{S}_{Pure}(\mathbf{x}) \in \Upsilon(\mathbf{x})$.*

Proof. Let $\mathbf{y} \in \Upsilon(\mathbf{x})$. Then $0 < \pi(\mathbf{y} - \mathbf{x}|\mathbf{x}) = \sum_{i=1}^n y_i \pi(\mathbf{e}^i - \mathbf{x}|\mathbf{x})$. But this implies that there exists at least one infective pure strategy for \mathbf{x} , i.e., $\mathbf{e}^i \in \Upsilon(\mathbf{x})$ for some $i = 1, \dots, n$. The converse trivially holds.

A fixed point of INIMDYN is *asymptotically stable* if any trajectory starting sufficiently close to \mathbf{x} converges to \mathbf{x} .

Theorem 4. *A state \mathbf{x} is asymptotically stable for INIMDYN with \mathcal{S}_{Pure} as strategy selection function if and only if \mathbf{x} is an ESS, provided that the payoff matrix is symmetric.*

Proof. If the payoff matrix is symmetric, every accumulation point of INIM-DYN with \mathcal{S}_{Pure} is a Nash equilibrium [13]. Moreover ESSs are strict local maximizers of $\pi(\mathbf{x})$ over Δ and vice versa [10].

If \mathbf{x} is asymptotically stable, then there exists a neighborhood U of \mathbf{x} in Δ such that any trajectory starting in U converges to \mathbf{x} . By Theorem 3 this implies that $\pi(\mathbf{x}) > \pi(\mathbf{y})$ for all $\mathbf{y} \in U, \mathbf{y} \neq \mathbf{x}$. Hence, \mathbf{x} is a strict local maximizer of $\pi(\mathbf{x})$ and therefore \mathbf{x} is an ESS.

Conversely, if \mathbf{x} is an ESS then \mathbf{x} is a strict local maximizer of $\pi(\mathbf{x})$ and an isolated Nash equilibrium. Hence, there exists a neighborhood U of \mathbf{x} in Δ where $\pi(\mathbf{x})$ is strictly concave and \mathbf{x} is the only accumulation point. This together with Theorem 3 implies that any trajectory starting in U will converge to \mathbf{x} . Hence, \mathbf{x} is asymptotically stable.

This selection function exhibits the nice property of rendering the complexity per iteration of our new dynamics linear in both space and time, as opposed to the replicator dynamics, which have quadratic space/time complexity per iteration.

Theorem 5. *Given the iterate $\mathbf{x}^{(t)}$ and its linear transformation $A\mathbf{x}^{(t)}$, both space and time requirement of one iteration step is linear in n , the number of objects.*

Proof. Again abbreviate $\mathbf{x} = \mathbf{x}^{(t)}$. Now, given $A\mathbf{x}$ we can straightforwardly compute in linear time and space $\pi(\mathbf{x})$ and $\mathcal{S}_{Pure}(\mathbf{x})$. Assume that $\mathcal{S}_{Pure}(\mathbf{x}) = \mathbf{e}^i$, then the computation of $\delta_{\mathbf{e}^i}(\mathbf{x})$ has a linear complexity, since $\pi(\mathbf{x} - \mathbf{e}^i | \mathbf{x}) = (A\mathbf{x})_i - \pi(\mathbf{x})$ and $\pi(\mathbf{e}^i - \mathbf{x}) = a_{ii} - 2A\mathbf{x} + \pi(\mathbf{x})$. Moreover, $A\mathbf{x}^{(t+1)}$ can be also computed in linear time and space since

$$A\mathbf{x}^{(t+1)} = \delta_{\mathbf{e}^i}(\mathbf{x}) [A_i - A\mathbf{x}] + A\mathbf{x},$$

where A_i is the i th column of A . Similar arguments hold if $\mathcal{S}_{Pure}(\mathbf{x}) = \overline{\mathbf{e}}^i_{\mathbf{x}}$. Indeed,

$$\begin{aligned} \pi(\overline{\mathbf{e}}^i_{\mathbf{x}} - \mathbf{x} | \mathbf{x}) &= \frac{x_i}{x_i - 1} \pi(\mathbf{e}^i - \mathbf{x} | \mathbf{x}), \\ \pi(\overline{\mathbf{e}}^i_{\mathbf{x}} - \mathbf{x}) &= \left(\frac{x_i}{x_i - 1} \right)^2 \pi(\mathbf{e}^i - \mathbf{x}), \end{aligned}$$

and finally,

$$A\mathbf{x}^{(t+1)} = \left(\frac{x_i}{x_i - 1} \right) \delta_{\overline{\mathbf{e}}^i_{\mathbf{x}}}(\mathbf{x}) [A_i - A\mathbf{x}] + A\mathbf{x}.$$

Hence the result.

The only step of quadratic complexity is the first one, where we need to compute $A\mathbf{x}^{(0)}$. Even this can be reduced to linear complexity, if we start from a pure strategy \mathbf{e}^i , in which case we have $A\mathbf{x}^{(0)} = A_i$. Note that the latter is impossible, e.g., for the replicator dynamics.

6 Experimental Results

In order to test the effectiveness of our algorithm, we present experiments on some computer vision problems, which have been attacked using the dominant-set framework or related quadratic optimization problems. Our aim is to show the computational gain over the standard algorithm used in the literature, namely the replicator dynamics (RD). Specifically, we present comparisons on image segmentation [1] and region-based hierarchical image matching [8].

The stopping criterion adopted with our dynamics is a measure of the accuracy of the Nash equilibrium, which is given by $\epsilon(\mathbf{x}) = \sum_i \min \{x_i, \pi(\mathbf{x} - \mathbf{e}^i | \mathbf{x})\}^2$. Indeed, $\epsilon(\mathbf{x})$ is 0 if and only if \mathbf{x} is a Nash equilibrium. In the experiments, we stopped the dynamics at accurate solutions, namely when $\epsilon(\mathbf{x}) < 10^{-10}$. As for RD, we stopped the dynamics either when $\epsilon(\mathbf{x}) < 10^{-10}$ or when a maximum number of iterations was exceeded.

6.1 Image Segmentation

We performed image segmentation experiments over the whole Berkeley dataset [14] using the dominant-set framework as published in [1]. The affinity between two pixels i and j was computed based on color and using the standard Gaussian kernel. Our INIMDYN algorithm was compared against standard replicator dynamics (RD) [1] (using the out-of-sample extension described in [15]) as well as the Nyström method [16]. The algorithms were coded in C and run on a AMD Sempron 3 GHz computer with 1GB RAM. To test the behavior of the algorithms under different input sizes we performed experiments at different pixel sampling rates, namely 0.005, 0.015, 0.03 and 0.05, which roughly correspond to affinity matrices of size 200, 600, 1200 and 2000, respectively. Since the Nyström method, as opposed to the dominant set approach, needs as input the desired number of clusters, we selected an optimal one after a careful tuning phase.

In Figure 2(a) we report (in logarithmic scale) the average computational times (in seconds) per image obtained with the three approaches. The computational gain of INIMDYN over the replicator dynamics is remarkable and it clearly increases at larger sampling rates. It is worth mentioning that INIMDYN other than being faster, achieved also better approximations of Nash equilibriums as




Fig. 1. Precision/Recall plots obtained on the Berkeley Image Database (s.r.=sampling rate)




Fig. 2. Average execution times (in logarithmic scale) for the image segmentation and region-based hierarchical image matching applications

opposed to RD. As for the quality of the segmentation results, we report in Figure 1 the average precision/recall obtained in the experiment with the different sampling rates. As can be seen, all the approaches perform equivalently, in particular RD and INIMDYN achieved precisely the same results as expected.

6.2 Region-Based Hierarchical Image Matching

In [8] the authors present an approach to region-based hierarchical image matching, aimed at identifying the most similar regions in two images, according to a similarity measure defined in terms of geometric and photometric properties. To this end, each image is mapped into a tree of recursively embedded regions, obtained by a multiscale segmentation algorithm. In this way the image matching problem is cast into a tree matching problem, that is solved recursively through a set of sub-matching problems, each of which is then attacked using replicator dynamics (see [8] for details). Given that typically hundreds of sub-matching problems are generated by a single image matching instance, it is of primary importance to have at one's disposal a fast matching algorithm. This makes our solution particularly appealing for this application.

We compared the running time of INIMDYN and RD over a set of images taken from the original paper [8]. We run the experiments on a machine equipped with 8 Intel Xeon 2.33 GHz CPUs and 8 GB RAM. Figure 2(b) shows the average computation times (in seconds) needed by RD and INIMDYN to solve the set of sub-matching problems generated from 10 image matching instances. Since each image matching problem generated sub-matching problems of different sizes, we grouped the instances having approximately the same size together. We plotted the average running time within each group (in logarithmic scale) as a function of the instance sizes and reported the standard deviations as error bars. Again, as can be seen, INIMDYN turned out to be orders of magnitude faster than RD.

7 From QPs to StQPs

Although in this paper we focused mainly on dominant sets, which lead to quadratic optimization problems over the standard simplex (StQPs), the

proposed approach is indeed more general and can be applied to a large class of quadratic programming problems (QPs), instances of which frequently arise in computer vision and pattern recognition.

In fact, consider a general QP over a bounded polyhedron

$$\max \left\{ \frac{1}{2} \mathbf{x}^\top Q \mathbf{x} + \mathbf{c}^\top \mathbf{x} : \mathbf{x} \in M \right\}, \quad (4)$$

where $M = \text{conv} \{ \mathbf{v}_1, \dots, \mathbf{v}_k \} \subseteq \mathbb{R}^n$ is the convex hull of the points \mathbf{v}_i , which form the columns of a $n \times k$ -matrix V . Then we can write the QP in (4) as the following StQP:

$$\max \left\{ \mathbf{y}^\top \hat{Q} \mathbf{y} : \mathbf{y} \in \Delta \right\},$$

where $\hat{Q} = \frac{1}{2} (V^\top Q V + \mathbf{e}^\top V^\top \mathbf{c} + \mathbf{c}^\top V \mathbf{e})$.

Thus every QP over a polytope can be expressed as an StQP. This approach is of course only practical if the vertices V are known and k is not too large. This is the case of QPs over the ℓ^1 ball, where $V = [I | -I]$, I the $n \times n$ identity matrix and $\Delta \subset \mathbb{R}^{2n}$ and, more generally, for box-constrained QPs [17]. However, even for general QPs, where the constraints are expressed as $M = \{ \mathbf{x} \in \mathbb{R}_+^n : A \mathbf{x} = b \}$, we can use StQP as a relaxation without using all vertices (see [18] for details).

Acknowledgements

We acknowledge financial support from the FET programme within the EU FP7, under the SIMBAD project (contract 213250).

References

1. Pavan, M., Pelillo, M.: Dominant sets and pairwise clustering. *IEEE Trans. Pattern Anal. Machine Intell.* 29(1), 167–172 (2007)
2. Neumann, J., von Cramon, D.Y., Forstmann, B.U., Zysset, S., Lohmann, G.: The parcellation of cortical areas using replicator dynamics in fMRI. *NeuroImage* 32(1), 208–219 (2006)
3. Wang, M., Ye, Z.L., Wang, Y., Wang, S.X.: Dominant sets clustering for image retrieval. *Signal Process.* 88(11), 2843–2849 (2008)
4. Hamid, R., Johnson, A., Batta, S., Bobick, A., Isbell, C., Coleman, G.: Detection and explanation of anomalous activities: representing activities as bags of event n-grams. In: *CVPR*, vol. 1, pp. 20–25 (2005)
5. Frommlet, F.: Tag SNP selection based on clustering according to dominant sets found using replicator dynamics. *Adv. in Data Analysis* (in press, 2010)
6. Wei, Q.D., Hu, W.M., Zhang, X.Q., Luo, G.: Dominant sets-based action recognition using image sequence matching. In: *ICIP*, vol. 6, pp. 133–136 (2007)
7. Albarelli, A., Torsello, A., Rota Bulò, S., Pelillo, M.: Matching as a non-cooperative game. In: *ICCV* (2009)
8. Todorovic, S., Ahuja, N.: Region-based hierarchical image matching. *Int. J. Comput. Vision* 78(1), 47–66 (2008)

9. Torsello, A., Rota Bulò, S., Pelillo, M.: Grouping with asymmetric affinities: a game-theoretic perspective. In: CVPR, pp. 292–299 (2006)
10. Weibull, J.W.: Evolutionary game theory. Cambridge University Press, Cambridge (1995)
11. Bomze, I.M., Pötscher, B.M.: Game Theoretical Foundations of Evolutionary Stability. Springer, Heidelberg (1989)
12. Bomze, I.M., Weibull, J.W.: Does neutral stability imply Lyapunov stability? Games and Econ. Behaviour 11, 173–192 (1995)
13. Rota Bulò, S.: A game-theoretic framework for similarity-based data clustering. PhD thesis, University of Venice (2009)
14. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In: ICCV, July 2001, vol. 2, pp. 416–423 (2001)
15. Pavan, M., Pelillo, M.: Efficient out-of-sample extension of dominant-set clusters. NIPS 17, 1057–1064 (2005)
16. Fowlkes, C., Belongie, S., Malik, F.C.J.: Spectral grouping using the Nyström method. IEEE Trans. Pattern Anal. Machine Intell. 26(2), 214–225 (2004)
17. Pardalos, P.M.: Quadratic problems defined on a convex hull of points. BIT Num. Math. 28(2), 323–329 (1988)
18. Bomze, I.M., Locatelli, M., Tardella, F.: Efficient and cheap bounds for (standard) quadratic optimization. Technical report, University "La Sapienza" of Rome (2005)