A Game-theoretic Approach for Enumerating Dominant Sets

Samuel Rota Bulò, Andrea Torsello and Marcello Pelillo

Dipartimento di Informatica.
Università "Ca' Foscari" di Venezia.
{srotabul,torsello,pelillo}@dsi.unive.it

Abstract. Recently, a novel graph-theoretic notion of a cluster has been proposed, i.e., the "dominant set", which captures the two basic requirements of a cluster, namely internal coherency and external incoherency. In this paper, we tackle the problem of finding several dominant sets using the replicator dynamics. Specifically, we adopt a game-theoretic perspective to this iterative extraction: Game dynamics are used to locate individual dominant sets, and after each extraction the similarity matrix is transformed in such a way as to make the located cluster unstable under the dynamics, without affecting the remaining groups. This guarantees that once found, a cluster will not be extracted again.

1 Introduction

Recently, a novel graph-theoretic notion of a cluster has been proposed, i.e., the "dominant set", which captures the two basic requirements of a cluster, namely internal coherency and external incoherency [1, 2]. In this framework clusters are in a one-to-one relation with the maxima of a quadratic function, where the non-zero entries of the maxima correspond to elements within a cohesive group. The approach is capable of automatically detect the number of groups present as this is determined by the number of maxima in the objective function. Furthermore, the approach works as a multi figure/ground discrimination algorithm since it only extracts cohesive groups, leaving spurious entries unclustered.

In this paper, following [3], we adopt a game-theoretic perspective to the iterative extraction of the dominant sets: Game dynamics are used to locate individual dominant sets, and after each extraction the similarity matrix is transformed in such a way as to make the located cluster unstable under the dynamics, without affecting the remaining groups. This guarantees that once found, a cluster will not be extracted again.

2 Dominant Sets Framework

In the pairwise clustering framework the data to be clustered are represented (possibly implicitly) as an undirected edge-weighted graph with no self-loops $G = (V, E, w)$, where $V = \{1, \ldots, n\}$ is the vertex set, $E \subseteq V \times V$ is the edge set, and $w : E \rightarrow \mathbb{R}_{+}^{*}$ is the (positive) weight function. Vertices in G correspond to data points, edges represent neighborhood relationships, and edge-weights reflect similarity/dissimilarity between pairs of linked vertices. As customary, we represent the graph G with the corresponding weighted adjacency (or similarity/dissimilarity) matrix, which is the $n \times n$ nonnegative, symmetric matrix $A = (a_{ij})$ defined as:

$$a_{ij} = \begin{cases} w(i,j) \ , & \text{if } (i,j) \in E \\ 0 \ , & \text{otherwise} \ . \end{cases}$$

The dominant set framework has been presented in [1]. Let $S \subseteq V$ be a non-empty subset of vertices and $i \in V$. The *(average) weighted degree* of i w.r.t. S is defined as:

$$\text{awdeg}_S\,(i) = \frac{1}{|S|} \sum_{j \in S} a_{ij} \tag{1}$$

where $|S|$ denotes the cardinality of S. Moreover, if $j \notin S$ we define $\phi_S\,(i,j) = a_{ij} - \text{awdeg}_S\,(i)$ which is a measure of the similarity between nodes j and i, with respect to the average similarity between node i and its neighbors in S.

Let $S \subseteq V$ be a non-empty subset of vertices and $i \in S$. The *weight* of i w.r.t. S is

$$\text{w}_S\,(i) = \begin{cases} 1, & \text{if } |S| = 1 \\ \displaystyle\sum_{j \in S \setminus \{i\}} \phi_{S \setminus \{i\}}\,(j,i)\, \text{w}_{S \setminus \{i\}}\,(j)\,, & \text{otherwise} \end{cases} \tag{2}$$

while the *total weight* of S is defined as:

$$\text{W}(S) = \sum_{i \in S} \text{w}_S(i) \ . \tag{3}$$

Intuitively, $\text{w}_S\,(i)$ gives us a measure of the overall similarity between vertex i and the vertices of $S \setminus \{i\}$ with respect to the overall similarity among the vertices in $S \setminus \{i\}$, with positive values indicating high internal coherency.

A non-empty subset of vertices $S \subseteq V$ such that $W(T) > 0$ for any non-empty $T \subseteq S$, is said to be *dominant* if:

1. $\text{w}_S\,(i) > 0$, for all $i \in S$
2. $\text{w}_{S \cup \{i\}}\,(i) < 0$, for all $i \notin S$.

The two previous conditions correspond to the two main properties of a cluster: the first regards internal homogeneity, whereas the second regards external inhomogeneity. The above definition represents our formalization of the concept of a cluster in an edge-weighted graph.

Now, consider the following quadratic program, which is a generalization of the so-called Motzkin-Straus program [4]:

$$\begin{array}{ll} \text{maximize} & f(\boldsymbol{x}) = \boldsymbol{x}' A \boldsymbol{x} \\ \text{subject to} & \boldsymbol{x} \in \Delta \end{array} \tag{4}$$

where

$$\Delta = \{\boldsymbol{x} \in \mathbb{R}^n : \boldsymbol{e}'\boldsymbol{x} = 1, \, \boldsymbol{x} \geq \mathbf{0}\}$$

is the standard simplex of $\mathbb{R}^n$, and $\boldsymbol{e}$ is a vector of appropriate length consisting of unit entries. The *support* of a vector $\boldsymbol{x} \in \Delta$ is defined as the set of indices corresponding to its positive components, that is $\sigma(\boldsymbol{x}) = \{i \in V \ : \ x_i > 0\}$. In [1], an intriguing connection between dominant sets and local solutions of Program (4) is established. Specifically, it is proven that if S is a dominant subset of vertices, then its (weighted) characteristic vector $\mathbf{x}^S$, which is the vector of Δ defined as

$$x_i^S = \begin{cases} \frac{\mathrm{w}_S(i)}{\mathrm{W}(S)}, & \text{if } i \in S \\ 0, & \text{otherwise} \end{cases} \tag{5}$$

is a strict local solution of Program (4). Conversely, if $\boldsymbol{x}$ is a strict local solution of Program (4) then its support $S = \sigma(\boldsymbol{x})$ is a dominant set, provided that $\mathrm{w}_{S\cup\{i\}}(i) \neq 0$ for all $i \notin S$.

3 A Game-theoretic Perspective

Let $O = \{1, 2, \ldots, n\}$ be the set of *pure strategies* available to the players and $A = (a_{ij})$ the $n \times n$ payoff or utility matrix [5] where a_{ij} is the payoff that a player gains when playing the strategy i against an opponent playing j. In biological contexts, payoff are typically measured in terms of Darwinian fitness or reproductive success whereas in economics applications, they usually represent firms' profits or consumers' utilities.

A *mixed strategy* is a probability distribution $\boldsymbol{x} = (x_1, x_2, \ldots, x_n)'$ over the available strategies in O. Mixed strategies clearly lie in Δ. The expected payoff that a player obtains by playing the element i against an opponent playing a mixed strategy $\boldsymbol{x}$ is $u(\boldsymbol{e}^i, \boldsymbol{x}) = (A\boldsymbol{x})_i = \sum_j a_{ij}x_j$, where $\boldsymbol{e}^i$ is the vector with all components equal zero except for the i^{th}-component which is equal to 1. Hence, the expected payoff received by adopting a mixed strategy $\boldsymbol{y}$ is $u(\boldsymbol{y}, \boldsymbol{x}) = \boldsymbol{y}'A\boldsymbol{x}$.

Evolutionary game theory considers an idealized scenario wherein pairs of individuals are repeatedly drawn from a large population to play a two-player symmetric game. Each player is not supposed to behave rationally or have a complete knowledge of the details of the game, but he acts according to a pre-programmed pure strategy. This dynamic activates some selection process that results in the evolution of the fittest strategies.

A well-known formalization of the selection process is given by the replicator equations [5]:

$$\dot{x}_i = x_i[u(\boldsymbol{e}^i, \boldsymbol{x}) - u(\boldsymbol{x}, \boldsymbol{x})]. \tag{6}$$

If the payoff matrix is symmetric then $\boldsymbol{x}'A\boldsymbol{x}$ is strictly increasing along any non-constant trajectory of any payoff-monotonic dynamics [5]. This result allows us to establish a one-to-one correspondence between the strict local solutions of program (4), namely characteristic vectors of dominant sets of a edge-weighted graph G, and asymptotically stable points of the replicator dynamics with payoff matrix the weighted adjacency matrix of G.

In order to obtain enumeration of dominant sets through a continuous formulation we move from undirected graphs to directed graphs, or, in other words, from symmetric payoff matrices to asymmetric payoff matrices. If we loosen the symmetry constraint, then all the results that bind local solutions to asymptotically stable points and dominant sets do not hold any longer, and $\boldsymbol{x}'A\boldsymbol{x}$ is not a Lyapunov function for the dynamics.

The *best replies* against a mixed strategy $\boldsymbol{x}$ is the set of mixed strategies

$$\beta(\boldsymbol{x}) = \{\boldsymbol{y} \in \Delta : u(\boldsymbol{y}, \boldsymbol{x}) = \max_{\boldsymbol{z}} u(\boldsymbol{z}, \boldsymbol{x})\}\,.$$

A mixed strategy $\boldsymbol{x}$ is a *Nash equilibrium* if it is a best reply to itself, i.e. $\forall \boldsymbol{y} \in \Delta$, $u(\boldsymbol{y}, \boldsymbol{x}) \leq u(\boldsymbol{x}, \boldsymbol{x})$. This implies that for all $i \in \sigma(\boldsymbol{x})$, $u(\boldsymbol{e}^i, \boldsymbol{x}) = u(\boldsymbol{x}, \boldsymbol{x})$, hence the payoff of every strategy in the support of $\boldsymbol{x}$ is constant, while all strategies outside the support of $\boldsymbol{x}$ earn a payoff that is less than or equal $u(\boldsymbol{x}, \boldsymbol{x})$. A strategy $\boldsymbol{x}$ is said to be an *evolutionary stable strategy* (ESS) if it is a Nash equilibrium and for all $\boldsymbol{y} \in \Delta$ such that $u(\boldsymbol{y}, \boldsymbol{x}) = u(\boldsymbol{x}, \boldsymbol{x})$ we have that $u(\boldsymbol{x}, \boldsymbol{y}) > u(\boldsymbol{y}, \boldsymbol{y})$. Intuitively, ESS are strategies such that any small deviation from them will lead to an inferior payoff.

Consider again Program (4), where A is a symmetric matrix. It is well known that the set of ESS's of a two-player game with payoff matrix A is equal to the set of strict local maximizers of (4). But this also implies that ESS's are in one-to-one correspondence with dominant sets of the edge-weighted graph having A as weighted adjacency matrix. In [3], we show that this correspondence between ESS's and dominant sets holds also for the general case of asymmetric affinities.

Through this change in perspective, we can move from a constrained maximization problem, to a game-theoretic setting. Instead of finding local solutions of a quadratic program, we look for ESS of a symmetric two-persons game. An advantage of this new approach is that we can generalize the Motzkin-Straus result also to non symmetric payoff matrices and, hence, edge-weighted directed graphs.

4 Continuous-based Enumeration

We have seen that the replicator dynamics is a simple and elegant tool for finding a dominant set in a edge-weighted graph. If however we reiterate it in order to extract other dominant sets, we have no guarantee that a new one will be found. The work we present overcomes this problem by rendering unstable the stationary points corresponding to dominant sets once they have been extracted using the replicator dynamics. We achieve this goal by repeated asymmetric extensions of the affinity matrix, in a way to preserve all other equilibria.

4.1 Binary Case: Enumeration of Maximal Cliques

In this subsection we will characterize our enumerative approach in presence of binary symmetric affinities, i.e. adjacency matrices of undirected graphs.

Given an undirected graph G with adjacency matrix A_G and $0 < \alpha < 1$, in [6], Bomze provides a one-to-one relation between local maximizers of (4) with $A = A_G + \alpha I$, where I is the identity matrix, and the characteristic vectors of maximal cliques of G. This implies also that in the symmetric binary case, maximal cliques are indeed dominant sets.

Since the notion of dominant set still holds for asymmetric affinities, in [3] we inspected its graph-theoretic sense in the context of directed graphs, yielding an interesting result.

Let $G = (V, E)$ be a directed graph. A *doubly-linked clique* of G is a set $S \subseteq V$ such that for all $u, v \in S$, $(u, v) \in E$ implies $(v, u) \in E$. The clique is *saturated* if there is no $t \in V \setminus S$ such that for all $s \in S$, $(t, s) \in E$.

Theorem 1. *Let $G = (V, E)$ be a directed graph with adjacency matrix A, $S \subseteq V$ is a saturated doubly-linked clique of G if and only if $\boldsymbol{x}^S$ is an ESS for a two-player game with payoff matrix $B = A + \alpha I$, where $0.5 < \alpha < 1$.*

We have already seen that if we consider an undirected graph G and the payoff matrix $A_G + \alpha I$ with $0 < \alpha < 1$, then the ESS's of the related two-player game are in one-to-one correspondence with maximal cliques of G. However if we strengthen the constraint on α to lay between 0.5 and 1, then we can see that the concept of saturated doubly-linked clique is a direct generalization to the asymmetric case of the concept of maximal clique, i.e. ESS's are in one-to-one correspondence with saturated doubly-linked cliques.

In order to render a given ESS $\boldsymbol{x}$ unstable, it is sufficient to drop the Nash condition for $\boldsymbol{x}$. A simple way to do it without affecting other equilibria, is to add a new strategy $\boldsymbol{z}$ that is a best reply to $\boldsymbol{x}$, but to no other ESS. This way, $\boldsymbol{x}$ will be no longer asymptotically stable.

Let $G = (V, E)$ be an undirected graph and $G' = (V, E')$ be its directed version where for all $(u, v) \in E$, $(u, v), (v, u) \in E'$. Given a set Σ of maximal cliques of G, we extend G' obtaining the *Σ-extension* G^Σ of G. The extension is as follows. For each clique $S \in \Sigma$, we create a new vertex v, called *Σ-vertex*, and put edges from v to each vertex in S and from each vertex outside S to v. After this operation, each Σ-vertex v dominates a particular clique S of Σ. Further, each vertex not in S dominates the Σ-vertex v so that it cannot form a new asymptotically stable strategy.

The next theorem guarantees, given a set of maximal cliques Σ of an undirected graph G, that there exists a one-to-one correspondence between the set of maximal cliques of G not in Σ and the set of ESS's of a two-player symmetric game associated to the extended graph G^Σ.

Theorem 2. *Let $G = (V, E)$ be an undirected graph, Σ be a set of maximal cliques of G and A be the adjacency matrix of the Σ-extension G_Σ of G. Let Φ be a two person symmetric game with payoff matrix $A + \alpha I$ with $0.5 < \alpha < 1$. Then $\boldsymbol{x}$ is an ESS equilibrium of Φ if and only if it is the characteristic vector of a maximal clique of G not in Σ.*

Proof. ($\Rightarrow$) By Theorem 1, if $\boldsymbol{x}$ is an ESS of Φ then it is the characteristic vector of a saturated doubly-linked clique S of G_Σ. By construction of G_Σ, the only possible doubly-linked cliques are subsets of V, therefore S is a clique of G. It is also maximal and not in Σ because otherwise it would not be saturated.

($\Leftarrow$) Consider $S \notin \Sigma$ a maximal clique of G. Then by construction of G^{Σ}, it is a saturated doubly-linked clique of G^{Σ} and hence, by Theorem 1, $\boldsymbol{x}^S$ is an ESS equilibrium of $\mathcal{G}$. □

The continuous-based enumerative algorithm uses this result in the following way. We iteratively find an asymptotically stable point through the replicator dynamics. If we have an ESS, then we have found a new maximal clique. After that, we extend the graph by adding the newly extracted clique to Σ, hence rendering its associated strategy unstable, and reiterate the procedure until we have enumerated the selected number of maximal cliques.

4.2 General Case: Enumeration of Dominant Sets

In order to obtain enumeration of dominant sets we iteratively render unstable all previously extracted ESS's by adding new strategies that are best replies to the previous ESS's, but to no other. This way the previous equilibria will no longer be asymptotically stable.

Let Σ be a tuple of ESS's of a game with payoff matrix A. So for example if $\boldsymbol{x}$ and $\boldsymbol{y}$ are ESS's of a doubly symmetric game then $\Sigma = (\boldsymbol{x}, \boldsymbol{y})$ and with Σ_i we select the i-th ESS.

We denote the barycenter of the simplex-face spanned by the vertices in C with $\boldsymbol{b}^C \in \Delta$.

$$b_i^C = \begin{cases} \frac{1}{|C|} & i \in C \\ 0 & \text{otherwise} \end{cases}$$

The Σ-extension $A^{\Sigma} = (a_{ij}^{\Sigma})$ of the payoff matrix A is defined as follows.

$$a_{ij}^{\Sigma} = \begin{cases} a_{ij} & \text{if } i, j \in [1, n] \\ \alpha & \text{if } j > n \text{ and } i \notin \sigma(\Sigma_{j-n}) \\ \beta & \text{if } i, j > n \text{ and } i = j \\ \frac{1}{|\Sigma_{i-n}|} \sum_{k \in \Sigma_{i-n}} a_{kj} & \text{if } i > n \text{ and } j \in \sigma(\Sigma_{i-n}) \\ 0 & \text{otherwise} \end{cases}$$

where $\alpha > \beta$ and $\beta = \max_{i,j} a_{ij}$.

Lemma 1. *Let Φ^{Σ} be a two-player game with Σ-extended payoff matrix A^{Σ} and let $\boldsymbol{x}$ be an ESS of Φ^{Σ}. If i is a Σ-strategy such that $i \in \sigma(\boldsymbol{x})$ then there exists $j \in \Sigma_{i-n} \cap \sigma(\boldsymbol{x})$.*

Proof. Suppose that $\Sigma_{i-n} \cap \sigma(\boldsymbol{x}) = \emptyset$. Then $(A^{\Sigma}\boldsymbol{x})_i = \beta x_i$. But for all $z \in \sigma(\boldsymbol{x}) \cap \{1, \ldots, n\}$ we have that

$$(A^{\Sigma}\boldsymbol{x})_z > \alpha x_i > \beta x_i = (A^{\Sigma}x)_i$$

because by construction of the Σ-extension, $\alpha > \beta$. This contradicts the hypothesis that $\boldsymbol{x}$ is an ESS. □

Lemma 2. *Let Φ^{Σ} be a two-player game with Σ-extended payoff matrix A^{Σ} and let $\boldsymbol{x}$ be an ESS of Φ^{Σ}. Then $\sigma(\boldsymbol{x})$ does not contain Σ-strategies.*

Proof. Suppose i a Σ-strategy such that $i \in \sigma(\boldsymbol{x})$. By lemma 1 there exists $j \in \Sigma_{i-n} \cap \sigma(\boldsymbol{x})$. Take $\boldsymbol{z} = \boldsymbol{x} + \varepsilon(\boldsymbol{e}^i - \boldsymbol{e}^j)$, where $0 < \varepsilon \leq x_j$.

Clearly, we have

$$\boldsymbol{z}' A^{\Sigma} \boldsymbol{x} = \boldsymbol{x}' A^{\Sigma} \boldsymbol{x} \,.$$

The second condition for asymptotic stability is:

$$\begin{aligned} \boldsymbol{x}' A^{\Sigma} \boldsymbol{z} > \boldsymbol{z}' A^{\Sigma} \boldsymbol{z} = \boldsymbol{x}' A^{\Sigma} \boldsymbol{z} + \varepsilon(\boldsymbol{e}^i - \boldsymbol{e}^j)' A^{\Sigma} \boldsymbol{z} \\ (\boldsymbol{e}^i - \boldsymbol{e}^j)' A^{\Sigma} \boldsymbol{z} = (\boldsymbol{e}^i - \boldsymbol{e}^j)' A^{\Sigma} \boldsymbol{x} + \varepsilon(\boldsymbol{e}^i - \boldsymbol{e}^j)' A^{\Sigma} (\boldsymbol{e}^i - \boldsymbol{e}^j) < 0 \\ (\boldsymbol{e}^i - \boldsymbol{e}^j)' A^{\Sigma} (\boldsymbol{e}^i - \boldsymbol{e}^j) = a^{\Sigma}_{ii} - a^{\Sigma}_{ij} - a^{\Sigma}_{ji} + a^{\Sigma}_{jj} = \\ \beta - (\boldsymbol{b}^{\sigma(\Sigma_{i-n})'} A^{\Sigma})_z < 0 \,. \end{aligned}$$

However, $\beta = \max_{i,j} a_{ij} \geq (\boldsymbol{b}^{\sigma(\Sigma_{i-n})'} A^{\Sigma})_z$, hence, we have a contradiction and we conclude the proof. □

From here on we will use the following notation. Let Φ be a two-player doubly symmetric game with payoff matrix A and Φ^{Σ} be a two-player game with payoff matrix A^{Σ}. If $\boldsymbol{x}$ is a mixed strategy of Φ then $\bar{\boldsymbol{x}}$ is a mixed strategy of Φ^{Σ} obtained from $\boldsymbol{x}$ by setting the components relative to Σ-strategies to 0.

Theorem 3. *Let Φ be a two-player doubly symmetric game with payoff matrix A and let Σ be a tuple of ESS's of Φ. Furthermore let Φ^{Σ} be a two-player game with payoff matrix A^{Σ}. Then $\boldsymbol{x}$ is an ESS of Φ not in Σ if and only if $\bar{\boldsymbol{x}}$ is an ESS of Φ^{Σ}.*

Proof. ($\Rightarrow$) Let $\boldsymbol{x}$ be an ESS of Φ not in Σ. We will prove that $\bar{\boldsymbol{x}}$ is an ESS also for Φ^{Σ}.

In first step we prove that $\bar{\boldsymbol{x}}$ is a Nash equilibrium by checking that for all pure strategies i, $(A^{\Sigma}\bar{\boldsymbol{x}})_i \leq \bar{\boldsymbol{x}}' A^{\Sigma} \bar{\boldsymbol{x}}$. It is trivial to verify that $\forall i = 1 \dots n$,

$$(A^{\Sigma}\bar{\boldsymbol{x}})_i = (A\boldsymbol{x})_i \leq \boldsymbol{x}' A \boldsymbol{x} = \bar{\boldsymbol{x}}' A^{\Sigma} \bar{\boldsymbol{x}}'.$$

Consider now the case when i is a Σ strategy, i.e. $n \leq i \leq n + |\Sigma|$ and let $\boldsymbol{z} = \Sigma_{i-n}$

$$(A^{\Sigma}\bar{\boldsymbol{x}})_i = \boldsymbol{b}^{\sigma(\Sigma_{i-n})'} A^{\Sigma} \bar{\boldsymbol{x}} \leq \bar{\boldsymbol{x}}' A^{\Sigma} \bar{\boldsymbol{x}}$$

Here we have the equality if and only if $\boldsymbol{z}$ is a best reply to $\boldsymbol{x}$. This implies that $\boldsymbol{z} = \boldsymbol{x}$ since $\boldsymbol{x}$ and $\boldsymbol{z}$ are ESS. However, this case is not possible because by hypothesis $\boldsymbol{x} \notin \Sigma$.

This implies that any $\boldsymbol{y}$ best reply to $\bar{\boldsymbol{x}}$ must have $\sigma(\bar{\boldsymbol{y}}) \subseteq \{1, \dots, n\}$. Hence, it suffices to check the second condition for evolutionary stability.

$$\bar{\boldsymbol{y}}' A^{\Sigma} \bar{\boldsymbol{x}} = \bar{\boldsymbol{x}}' A^{\Sigma} \bar{\boldsymbol{x}} \Rightarrow \bar{\boldsymbol{x}}' A^{\Sigma} \bar{\boldsymbol{y}} = \boldsymbol{x}' A \boldsymbol{y} > \boldsymbol{y}' A \boldsymbol{y} = \bar{\boldsymbol{y}}' A^{\Sigma} \bar{\boldsymbol{y}} \,.$$

This proves that $\bar{\boldsymbol{x}}$ is an ESS of Φ^{Σ}.

($\Leftarrow$) Let $\boldsymbol{x}$ be an ESS of Φ^{Σ}. By lemma 2 we have that $\sigma(\boldsymbol{x})$ does not contain Σ-strategies. Let $\boldsymbol{y}$ be the mixed strategy obtained removing the Σ-strategies. It is trivial to verify that $\boldsymbol{y}$ is an ESS of Φ. We show that $\boldsymbol{y}$ cannot be in Σ, otherwise there would exist a Σ-strategy i such that

$$(A\boldsymbol{y})_i = (A^{\Sigma}\boldsymbol{x})_i = \boldsymbol{b}^{\sigma(\boldsymbol{x})'} A^{\sigma}\boldsymbol{x} = \boldsymbol{x}' A^{\Sigma}\boldsymbol{x}$$

and by the second condition for evolutionary stability

$$(\boldsymbol{x}' A^{\Sigma})_i = 0 > a_{ii}^{\Sigma} = \beta$$

arriving at a contradiction. Concluding we have an ESS $\boldsymbol{y}$ of Φ^{Σ} only if $\boldsymbol{y} = \bar{\boldsymbol{x}}$ where $\boldsymbol{x}$ is an ESS of Φ not in Σ. □

We use this result to enumerate the dominant sets in the following way: We iteratively find new dominant sets by looking for an asymptotically stable point using the replicator dynamics. After that, we extend the graph by adding the newly extracted set to Σ, hence rendering its associated strategy unstable, and reiterate the procedure until we have enumerated all the groups and hence are unable to find new dominant sets.

5 Conclusions

In this paper we presented a game-theoretic approach for enumerating dominant sets of an edge-weighted graph. Game dynamics are used to locate individual dominant sets and after each extraction, the similarity matrix is transformed in such a way as to make the located cluster unstable under the replicator dynamics, without affecting the remaining groups.

References

1. Pavan, M., Pelillo, M.: Generalizing the motzkin-straus theorem to edge-weighted graphs, with applications to image segmentation. In: Energy Minim. Methods in Computer Vision and Patt. Recogn. Volume 2683 of Lecture Notes in Computer Science., Springer (2003) 485–500
2. Pavan, M., Pelillo, M.: Dominant sets and pairwise clustering. IEEE Trans. Pattern Anal. Machine Intell. **29**(1) (2007) 167–172
3. Torsello, A., Bulò, S.R., Pelillo, M.: Grouping with asymmetric affinities: A game-theoretic perspective. In: IEEE Conf. Computer Vision and Patt. Recogn. (2006) 292–299
4. Motzkin, T.S., Straus, E.G.: Maxima for graphs and a new proof of a theorem of Turán. Canad. J. Math. **17** (1965) 533–540
5. Weibull, J.W.: Evolutionary Game Theory. Cambridge University Press (1995)
6. Bomze, I.: Evolution towards the maximum clique. J. Global Optimiz. **10**(2) (1997) 143–164