

Graph Transduction as a Noncooperative Game

Aykut Erdem

aykut.erdem@hacettepe.edu.tr

Faculty of Engineering, Hacettepe University, 06800 Beytepe, Ankara, Turkey

Marcello Pelillo

pelillo@dsi.unive.it

DAIS, Università Ca' Foscari Venezia, 30172 Venezia Mestre, Italy

Graph transduction is a popular class of semisupervised learning techniques that aims to estimate a classification function defined over a graph of labeled and unlabeled data points. The general idea is to propagate the provided label information to unlabeled nodes in a consistent way. In contrast to the traditional view, in which the process of label propagation is defined as a graph Laplacian regularization, this article proposes a radically different perspective, based on game-theoretic notions. Within the proposed framework, the transduction problem is formulated in terms of a noncooperative multiplayer game whereby equilibria correspond to consistent labelings of the data. An attractive feature of this formulation is that it is inherently a multiclass approach and imposes no constraint whatsoever on the structure of the pairwise similarity matrix, being able to naturally deal with asymmetric and negative similarities alike. Experiments on a number of real-world problems demonstrate that the proposed approach performs well compared with state-of-the-art algorithms, and it can deal effectively with various types of similarity relations.

1 Introduction ---

In the machine learning community, semisupervised learning (SSL) has gained considerable popularity over the past decade (Chapelle, Schölkopf, & Zien, 2006; Zhu, 2005), and within the existing paradigms, graph-based approaches to SSL, namely, the graph transduction methods, constitute an important class of algorithms. These methods model the geometry of the data as a graph with nodes corresponding to the labeled and unlabeled points and edges being weighted by the similarity between the points; they then try to estimate the labels of unlabeled points by propagating the coarse information available at the labeled nodes to the unlabeled ones. Performing this propagation in a consistent way relies on a common a priori assumption known as the cluster assumption (Zhou, Bousquet, Lal, Weston, & Schölkopf, 2004; Chapelle et al., 2006), which is reminiscent of the

homophily principle used in social network analysis (Easley & Kleinberg, 2010). The assumption simply states that (1) points that are close to each other are expected to have the same label and (2) points in the same cluster (or on the same manifold) are expected to have the same label. Building on this assumption, traditional graph-based approaches formalize graph transduction as a regularized function estimation problem on an undirected graph (Joachims, 2003; Zhu, Ghahramani, & Lafferty, 2003; Zhou et al., 2004).

This letter presents a novel game-theoretic perspective to the problem of graph transduction. Specifically, graph transduction is formulated in terms of a multiplayer noncooperative game where the players are the data points that take part in the game to decide their class memberships. In this setting, while the strategies played by the labeled points are already decided at the outset, because each of them knows which class it belongs to, the possible strategies available to unlabeled points are the whole set of hypotheses of being a member of one of the available classes. Within this formulation, the well-known Nash equilibrium concept for noncooperative games provides a well-founded way of consistent labeling for the unlabeled points.

The game-theoretic interpretation presented here is appealing for a number of reasons. To begin with, it is intrinsically a multiclass approach. But more importantly, it can cope with both asymmetric and negative similarities. Since the seminal work of Tversky (1977), a large body of psychological studies suggests that human similarity judgments are nonmetric (in particular, asymmetric). Nonmetric similarities or dissimilarities also arise naturally in many practical applications, like comparing shapes (Jacobs, Weinsall, & Gdalyahu, 2000) and protein sequences (Meila & Pentney, 2007). Common examples include the directed Hausdorff distance between sets, the Kullback-Leiber divergence between probability distributions, and Tversky's contrast model (Tversky, 1977; Santini & Jain, 1999). Nonmetricity has been largely regarded as an artifact of poor choice of features or algorithms. Note, however, that it has been shown that nonmetricity is sometimes essential to the nature of the problem, and rendering the similarities metric may destroy relevant information (Laub, Roth, Buhmann, & Müller, 2006).

The organization of the letter is as follows. Section 2 shows that transduction on an unweighted undirected graph can be formulated as a binary constraint satisfaction problem, which provides the motivation for the proposed approach. Section 3 briefly reviews some basic notions of noncooperative game theory. Next, section 4 extends our analysis to the general case of transduction on weighted (directed) graphs, where graph transduction is formalized as a noncooperative game using the connection between so-called relaxation labeling processes and game theory (Miller & Zucker, 1991). Section 5 investigates the loose relation between the proposed game-theoretic formulation and the traditional energy-based formulations. Section 6 reports experimental results on a number of real-world classification problems. Section 7 concludes.

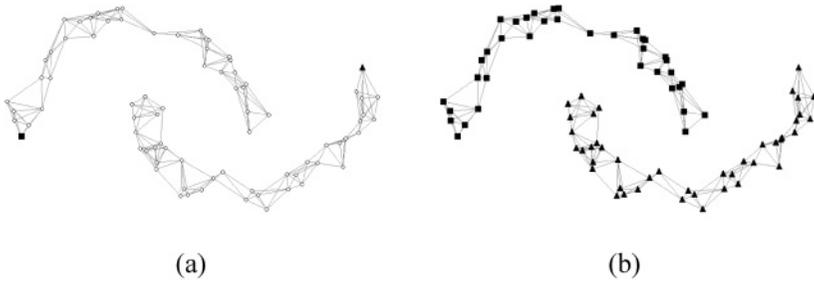


Figure 1: Transductive learning on an unweighted undirected graph. (a) An unweighted undirected graph describing the similarity relationships among the points. Two nodes of the graph are marked with different labels, respectively represented by a square and a triangle. (b) The consistent labeling of points satisfying the constraints based on the cluster assumption of SSL.

2 Transductive Learning on Unweighted Undirected Graphs ---

The theoretical motivation for the proposed approach stems from analysis of the simplest case of graph transduction where the graph expressing the similarity relationships among the data points is an unweighted undirected graph. To give an example, the graph can be seen as a k -nearest neighbor (k -NN) graph with 0/1 weights over points in which the presence of an edge simply denotes the perfect similarity between a pair of two data points; otherwise, the points are completely dissimilar. To illustrate this toy problem, consider the graph shown in Figure 1 in which edges are unweighted. The classification task is to estimate the labels of the unlabeled points based exclusively on the information available at the two labeled points, each marked with a different label. Recall the cluster assumption of semisupervised learning that neighboring objects and objects in the same cluster (or on the same manifold structure) tend to belong to the same class. Clearly, in the unweighted graph setting, the cluster assumption is also valid and can be expressed as the hypothesis that every node in a connected component of a binary similarity graph has the same class label as each connected component describes a manifold.

Following this observation, we can formulate this toy version of graph transduction as a (binary) constraint satisfaction problem (CSP) (Tsang, 1993; Marriott & Stuckey, 1998). CSPs are widely used to solve combinatorial problems in a variety of application domains, such as artificial intelligence and computer vision. In the computer vision literature, the problem is often known as the consistent labeling problem (Waltz, 1975; Haralick & Shapiro, 1979).

A binary CSP is defined by a set of variables representing the elements of the problem being modeled and a set of binary constraints representing

the relationships among variables. A solution of the problem is simply an assignment of values to the variables, which satisfies all the constraints. If there is no such assignment, the problem is unsatisfiable. When each variable can take a value from a finite domain, a binary CSP can be described in a formal manner as a triple (V, D, R) , where $V = \{v_1, \dots, v_n\}$ is a set of variables; $D = \{D_{v_1}, \dots, D_{v_n}\}$ is a set of domains of the variables, each D_{v_i} denoting a finite set of possible values for variable v_i ; and $R = \{R_{ij} \mid R_{ij} \subseteq D_{v_i} \times D_{v_j}\}$ is a set of binary constraints, each R_{ij} describing compatible pairs of values for the variables v_i and v_j . If the cardinalities of the domains of variables are p and q , respectively, then R_{ij} can be expressed by a 0/1 matrix of size $p \times q$, where $R_{ij}(\lambda, \lambda') = 1$ if the assignment $v_i = \lambda$ is compatible with the assignment $v_j = \lambda'$. For a general CSP on a finite domain, the problem of finding a solution is known to be NP-complete (Haralick, Davis, Rosenfeld, & Milgram, 1978). The simplest way to obtain an assignment satisfying all the given constraints or to report nonexistence of such a solution is to perform backtracking. However, it is time-consuming, so in practice, either a constraint propagation technique or a local search method is used to solve the problem (Tsang, 1993; Marriott & Stuckey, 1998).

Returning back to the motivating problem of transductive learning on an unweighted undirected graph, suppose that we are given a data set $\mathcal{D} = \{\mathcal{D}_\ell, \mathcal{D}_u\}$ consisting of labeled points $\mathcal{D}_\ell = \{d_1, \dots, d_\ell\}$ and unlabeled points $\mathcal{D}_u = \{d_{\ell+1}, \dots, d_n\}$ and a set of labels $\Phi = \{1, \dots, c\}$ such that the labels provided for the first ℓ labeled points are given by $\{\phi_1, \dots, \phi_\ell\} \in \Phi$. The task of transductive learning is to estimate the unknown labels $\{\phi_{\ell+1}, \dots, \phi_n\}$ of unlabeled points $\{d_{\ell+1}, \dots, d_n\}$. Now further suppose that the relationships among the data points are given by an unweighted, undirected graph $\mathcal{G} = (\mathcal{D}, \mathcal{E})$, where \mathcal{D} is the set of nodes and \mathcal{E} is the set of edges such that an edge $e_{ij} \in \mathcal{E}$ shows that points i and j are perfectly similar to each other. Let $A = (a_{ij})$ denote the 0/1 adjacency matrix of \mathcal{G} . Reflecting the constraints imposed by the cluster assumption of SSL, the problem of graph transduction on an unweighted graph can be formalized as a binary CSP as follows:

- The set of variables: $V = \{v_1, \dots, v_n\}$
- Domains: $D_{v_i} = \begin{cases} \{\phi_i\}, & \text{for all } 1 \leq i \leq \ell \\ \Phi & \text{for all } \ell + 1 \leq i \leq n \end{cases}$
- Binary constraints: $\forall ij : \text{if } a_{ij} = 1, \text{ then } v_i = v_j$ —for a two-class problem, for example $R_{ij} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$.

Each assignment of values to the variables satisfying all the constraints is a solution of the CSP and provides a consistent labeling for the unlabeled points.

Table 1: Payoff Matrix of the Rock-Paper-Scissors Game.

	<i>rock</i>	<i>paper</i>	<i>scissors</i>
<i>rock</i>	(0, 0)	(-1, 1)	(1, -1)
<i>paper</i>	(1, -1)	(0, 0)	(-1, 1)
<i>scissors</i>	(-1, 1)	(1, -1)	(0, 0)

Classical CSPs such as the one given in this section assume crisp constraints, in the sense that constraints are either completely satisfied or completely violated. However, for many real-world applications, such a formulation is too restrictive to be practical. A classical generalization to deal with soft constraints is described in Hummel and Zucker (1983), in which each constraint is assigned a weight representing a level of confidence. Later, it was shown that the notion of consistency they proposed is related to the Nash equilibrium concept in noncooperative game theory (Miller & Zucker, 1991). In this study, we build on this connection to devise a graph transduction game, which serves as a generalization of the binary CSP for the motivating problem. For completeness, we next provide an overview of some basic notions from noncooperative game theory (Nash, 1951).

3 Noncooperative Games and Nash Equilibria

Noncooperative game theory deals with models of strategic interactions (games) among anonymous agents (players), where the goal of each player is to maximize its own utility or payoff. Each player has a set of possible actions (pure strategies) to play, called the pure strategy set, and receives a payoff based on its own choice and those of the other players. In the definitions below, we restrict ourselves to multiplayer games expressed in normal form and follow the notations used in Weibull (1995).

In normal form, a game with many players can be expressed as a triple $G = (\mathcal{I}, S, \pi)$, where $\mathcal{I} = \{1, \dots, n\}$, with $n \geq 2$, is the set of players, $S = \times_{i \in \mathcal{I}} S_i$ is the joint strategy space defined as the Cartesian product of the individual pure strategy sets $S_i = \{1, \dots, m_i\}$, and $\pi : S \rightarrow \mathbb{R}^n$ is the combined payoff function, which assigns a real valued payoff $\pi_i(s) \in \mathbb{R}$ to each pure strategy profile $s \in S$ and player $i \in \mathcal{I}$. In the case of two players, payoff functions can be represented as two $m_1 \times m_2$ matrices. To illustrate, consider the well-known rock-paper-scissors game, a two-player game in which each player has three possible strategies, $S_i = \{\textit{rock}, \textit{paper}, \textit{scissors}\}$, for each $i = 1, 2$. Both players simultaneously choose a strategy and consequently receive a payoff based on their actions with respect to the rules “rock beats scissors,” “scissors beats paper,” and “paper beats rock.” The payoff matrix of the game is given in Table 1 in which player 1 is the row

player and player 2 is the column player. For example, if player 1 plays *rock* and player 2 plays *paper*, player 1 loses the game, that is, receives a payoff of -1 , and player 2 wins the game, that is, receives a payoff of 1 .

A mixed strategy of player $i \in \mathcal{I}$ is a probability distribution over its pure strategy set S_i , which can be described as the vector $x_i = (x_{i1}, \dots, x_{im_i})^T$ such that each component x_{ih} denotes the probability that the player chooses to play its h th pure strategy among all the available strategies. Mixed strategies for each player $i \in \mathcal{I}$ are constrained to lie in the standard simplex of the m_i -dimensional Euclidean space \mathbb{R}^{m_i} :

$$\Delta_i = \left\{ x_i \in \mathbb{R}^{m_i} : \sum_{h=1}^{m_i} x_{ih} = 1, \text{ and } x_{ih} \geq 0 \text{ for all } h \right\}.$$

Accordingly, a mixed strategy profile $x = (x_1, \dots, x_n)$ is defined as a vector of mixed strategies, each $x_i \in \Delta_i$ representing the mixed strategy assigned to player $i \in \mathcal{I}$, and each mixed strategy profile lives in the mixed strategy space of the game, given by the Cartesian product

$$\Theta = \times_{i \in \mathcal{I}} \Delta_i. \tag{3.1}$$

For simplicity, let $z = (x_i, y_{-i}) \in \Theta$ denote the strategy profile where player i plays strategy $x_i \in \Delta_i$ whereas the other players $j \in \mathcal{I} \setminus \{i\}$ play based on the strategy profile $y \in \Theta$, that is, $z_i = x_i$ and $z_j = y_j$ for all $j \neq i$. The expected value of the payoff that player i obtains can be determined by a weighted sum for any $i, j \in \mathcal{I}$ (when all the other players' strategies are kept fixed) as

$$u_i(x) = \sum_{s \in S} x(s) \pi_i(s) = \sum_{k=1}^{m_j} u_i(e_j^k, x_{-j}) x_{jk}, \tag{3.2}$$

where $u_i(e_j^k, x_{-j})$ denotes the payoff that player i receives when player j adopts its k th pure strategy, and $e_j^k \in \Delta_j$ stands for the extreme mixed strategy corresponding to the vector of length m_j whose components are all zero except the k th one, which is equal to one. Note that for player j , playing its k th pure strategy is probabilistically equivalent to playing the extreme mixed strategy e_j^k .

The mixed best replies for player i against a mixed strategy $y \in \Theta$, denoted by $\beta_i(y)$, are the set of mixed strategies such that no mixed strategy other than the ones included in this set gives a higher payoff to player i against strategy y :

$$\beta_i(y) = \{ x_i \in \Delta_i : u_i(x_i, y_{-i}) \geq u_i(z_i, y_{-i}) \quad \forall z_i \in \Delta_i \}.$$

Subsequently, the combined mixed best replies is defined as the Cartesian product of the best replies of all the players $\beta(y) = \times_{i \in \mathcal{I}} \beta_i(y) \subset \Theta$.

Definition 1. A mixed strategy $x^* = (x_1^*, \dots, x_n^*)$ is said to be a Nash equilibrium if it is the best reply to itself, $x^* \in \beta(x^*)$, that is

$$u_i(x_i^*, x_{-i}^*) \geq u_i(x_i, x_{-i}^*) \quad (3.3)$$

for all $i \in \mathcal{I}$, $x_i \in \Delta_i$, and $x_i \neq x_i^*$. Furthermore, a Nash equilibrium x^* is called strict if each x_i^* is the unique best reply to x^* , $\beta(x^*) = \{x^*\}$.

Nash equilibrium constitutes the key concept of game theory. It is motivated by the idea that a theory of rational decision making should not be a self-destroying prophecy that creates an incentive to deviate for those who believe it. Indeed, the notion itself is a stability condition, which states that no player can obtain a higher payoff by changing unilaterally its own strategy once such an equilibrium state is reached. Note that not all games have a Nash equilibrium in pure strategies (e.g., the rock-paper-scissors game), but a fundamental result of game theory states that any normal-form game has at least one mixed Nash equilibrium (Nash, 1951). For example, the rock-paper-scissors game has a (unique) mixed Nash equilibrium that corresponds to the case where each player picks a strategy uniformly at random: $x_i = (1/3, 1/3, 1/3)$ for each $i = 1, 2$. The algorithmic issue of computing a Nash equilibrium is discussed in section 4.2.

4 The Graph Transduction Game (GTG)

Consider the following graph transduction game (GTC). Assume each player $i \in \mathcal{I}$ participating in the game corresponds to a particular point in a data set $\mathcal{D} = \{d_1, \dots, d_n\}$ and can choose a strategy among the set of strategies $S_i = \{1, \dots, c\}$, each expressing a certain hypothesis about its membership in a class and c being the total number of classes. Hence, the mixed strategy profile of each player $i \in \mathcal{I}$ lies in the c -dimensional simplex Δ_i . By problem definition, the players of the game can be categorized into two disjoint groups: those that already have knowledge of their membership, referred to as labeled players and denoted with the symbol \mathcal{I}_ℓ , and those that do not have any idea about this at the beginning of the game, which are hence called unlabeled players and correspondingly denoted with \mathcal{I}_u .

The so-called labeled players of the game can be distinguished further based on the strategies they follow without hesitation, coming from their membership information. In formal terms, $\mathcal{I}_\ell = \{\mathcal{I}_{\ell|1}, \dots, \mathcal{I}_{\ell|c}\}$, where each disjoint subset $\mathcal{I}_{\ell|k}$ stands for the set of players always playing their k th pure strategies. It thus follows from this statement that each player $i \in \mathcal{I}_{\ell|k}$

plays its extreme mixed strategy $e_i^k \in \Delta_i$. In other words, x_i is constrained to belong to the minimal face of the simplex Δ_i spanned by $\{e_i^k\}$. In this regard, it can be argued that the labeled players do not play the game to maximize their payoffs since they have already chosen their strategies. In fact, the transduction game can be easily reduced to a game with only unlabeled players \mathcal{I}_u , where the definite strategies of labeled players \mathcal{I}_ℓ act as bias over the choices of unlabeled players.

It is important to note that any instance of the proposed transduction game will always have a Nash equilibrium in mixed strategies (Nash, 1951). Recall that for the players, such an equilibrium corresponds to a steady state such that each player plays a strategy that could yield the highest payoff when the strategies of the remaining players are kept fixed, and it provides a globally consistent labeling of the data set. Once an equilibrium is reached, the label of a data point (player) i is simply given by the strategy with the highest probability in the equilibrium mixed strategy of player i as

$$\phi_i = \arg \max_{h=1\dots c} x_{ih}, \quad (4.1)$$

thereby yielding a crisp classification.

4.1 Payoff Functions. We assume in this letter that the proposed graph transduction game is an instance of a special subclass of multiplayer games, known as polymatrix games (Janovskaya, 1968; Howson, 1972), in which players are nodes of a graph and every edge denotes a two-player game between corresponding pair of players. In other words, we suppose that only pairwise interactions are allowed in the game, and the payoffs associated with each player are additively separable so that the payoff of each player is given by the sum of the payoffs gained from each game played with one of its neighbor. Formally, for a pure strategy profile $s = (s_1, \dots, s_n) \in S$, the payoff function of every player $i \in \mathcal{I}$ is in the form

$$\pi_i(s) = \sum_{j=1}^n A_{ij}(s_i, s_j), \quad (4.2)$$

where A_{ij} is the partial payoff matrix between players i and j . It follows that in terms of a mixed strategy profile $x = (x_1, \dots, x_n)$, the payoffs are computed as $u_i(e_i^h) = \sum_{j=1}^n (A_{ij}x_j)_h$ and $u_i(x) = \sum_{j=1}^n x_i^T A_{ij}x_j$.

In an instance of the transduction game, since each labeled player is restricted to play a definite strategy of its own, all of these fixed choices can be reflected directly in the payoff function of an unlabeled player $i \in \mathcal{I}_u$ as

follows:

$$u_i(e_i^h) = \sum_{j \in \mathcal{I}_U} (A_{ij} x_j)_h + \sum_{k=1}^c \sum_{j \in \mathcal{I}_{D|k}} A_{ij}(h, k), \quad (4.3)$$

$$u_i(x) = \sum_{j \in \mathcal{I}_U} x_i^T A_{ij} x_j + \sum_{k=1}^c \sum_{j \in \mathcal{I}_{D|k}} x_i^T (A_{ij})_k. \quad (4.4)$$

What is left is how to specify partial payoff matrices between each pair of players. For doing so, the binary logical constraints for the toy transduction problem can be simply replaced with their weighted versions. Let the geometry of the data be modeled with a weighted graph $\mathcal{G} = (\mathcal{D}, \mathcal{E}, w)$ in which \mathcal{D} is the set of nodes representing both labeled and unlabeled points and $w : \mathcal{E} \rightarrow \mathbb{R}$ is a weight function assigning a similarity value to each edge $e \in \mathcal{E}$. Representing the graph with its weighted adjacency matrix $W = (w_{ij})$, the partial payoff matrix between two players i and j is set as $A_{ij} = w_{ij} \times I_c$, where I_c is the identity matrix of size c . Note that when partial payoff matrices are represented in block form as $A = (A_{ij})$, the matrix A is given by the Kronecker product $A = I_c \otimes W$. Notice that in the case of binary-valued similarity relations, the above partial payoff matrices coincide with the compatibility matrices defined for the binary CSP given in section 2. Further, if only pure strategies are allowed, the transduction game reduces to the presented CSP. In a pure Nash equilibrium of such a game, the neighboring players all play the same pure strategy to get the maximum support from their neighbors. However, when the data contain noise, then class manifolds may be connected; hence, a pure Nash equilibrium may not exist, or, in other words, the CSP may be unsatisfiable.

Empirically we observed that specifying payoffs in terms of the normalized similarity matrix $\widehat{W} = D^{-1/2} W D^{-1/2}$, with $D = (d_{ii})$ being the diagonal degree matrix of W whose elements are given by $d_{ii} = \sum_j w_{ij}$, performs better than the case with the original similarities. In that regard, the use of normalization is a common practice in graph-based approaches because it can typically achieve a better performance. To give an example, while the GFHF method (Zhu et al., 2003) uses original (unnormalized) similarities, the LGC method (Zhou et al., 2004) employs the normalized similarity matrix in its formulation. Moreover, while not directly related to graph transduction, it has been shown that the use of normalization has nice convergence properties in spectral clustering (von Luxburg, Bousquet & Belkin, 2004). In terms of game theory, however, both versions of the transduction game (with and without normalizing input similarities) belong to the so-called class of normalized games—those with payoffs in a unit-length interval (Daskalakis, 2011)—but the gap in their classification performance requires further investigation.

4.2 Computing Nash Equilibria. There has been growing interest in the computational aspects of Nash equilibria. The general problem of computing a Nash equilibrium is shown to belong to the complexity class PPAD-complete, a newly defined subclass of NP (Daskalakis, Goldberg, & Papadimitriou, 2009; Daskalakis, 2011). Nevertheless, there are many refinements and extensions of Nash equilibria that can be computed efficiently, and, moreover, the former result does not apply to certain classes of games (Nisan, Roughgarden, Tardos, & Vazirani, 2007). This letter uses the well-established *evolutionary approach* (Weibull, 1995; Hofbauer & Sigmund, 1998), initiated by, Maynard Smith (1982) to find a Nash equilibrium in a multiplayer game, but we mention that other options do exist, such as the simplicial subdivision method (van der Laan, Talman, & van der Heyden, 1987), continuation methods (Govindan & Wilson, 2003), and enumeration-of-support methods (Mangasarian, 1964; Dickhaut & Kaplan, 1991; Porter, Nudelman, & Shoham, 2008).

The dynamic interpretation of Nash equilibria through the evolutionary approach imagines that the game is played repeatedly, generation after generation, during which a selection process acts on the multipopulation of strategies, thereby resulting in the evolution of the fittest strategies. The selection dynamics is commonly modeled by the following set of ordinary differential equations:

$$\dot{x}_{ih} = g_{ih}(x)x_{ih}, \quad (4.5)$$

where a dot signifies derivative with respect to time and $g(x) = (g_1(x), \dots, g_n(x))$ is the growth rate function with open domain containing $\Theta = \times_{i \in \mathcal{I}} \Delta_i$, each component $g_i(x)$ being a vector-valued growth rate function for player i . Hence, g_{ih} specifies the growth rate at which player i 's pure strategy h replicates. It is generally required that the function g be regular (Weibull, 1995), that is, g is Lipschitz continuous, and $g_i(x) \cdot x_i = 0$ for all $x \in \Theta$ and players $i \in \mathcal{I}$. While the first condition guarantees that system 4.5 has a unique solution through every initial state, condition $g_i(x) \cdot x_i = 0$ ensures that the simplex Δ_i is invariant under equation 4.5.

The class of regular selection dynamics includes a wide subclass known as payoff monotonic dynamics, in which the ratio of strategies with a higher payoff increases at a higher rate. Formally, a regular selection dynamics, equation 4.5, is said to be payoff monotonic if

$$u_i(e_i^h, x_{-i}) > u_i(e_i^k, x_{-i}) \Leftrightarrow g_{ih}(x) > g_{ik}(x) \quad (4.6)$$

for all $x \in \Theta$, $i \in \mathcal{I}$ and pure strategies $h, k \in S_i$.

A particular subclass of payoff monotonic dynamics, which is used to model the evolution of behavior by imitation processes, is given by

$$\dot{x}_{ih} = x_{ih} \left[\sum_{l \in S_i} x_{il} (\phi_i[u_i(e_i^h - e_i^l, x_{-i})] - \phi_i[u_i(e_i^l - e_i^h, x_{-i})]) \right], \tag{4.7}$$

where $\phi_i(u_i)$ is a strictly increasing function of u_i . The multipopulation version of the replicator dynamics is obtained when ϕ_i is taken as the identity function, that is, $\phi_i(u_i) = u_i$, as

$$\dot{x}_{ih} = x_{ih} (u_i(e_i^h, x_{-i}) - u_i(x)). \tag{4.8}$$

The following theorem states that the fixed points of equation 4.8 are Nash equilibria:

Theorem 1. *A point $x \in \Theta$ is the limit of a trajectory of equation 4.8 starting from the interior of Θ if and only if x is a Nash equilibrium. Further, if point $x \in \Theta$ is a strict Nash equilibrium, then it is asymptotically stable, additionally implying that the trajectories starting from all nearby states converge to x .*

Proof See Weibull (1995).

In the experiments, the discrete-time counterpart of equation 4.8 given below is utilized, where the mixed strategies of each unlabeled player are initialized to uniform probabilities, the barycenter of the simplex Δ_i :

$$x_{ih}(t + 1) = x_{ih}(t) \frac{u_i(e_i^h)}{u_i(x(t))}. \tag{4.9}$$

The discrete-time replicator dynamics, equation 4.9, has essentially the same dynamical properties as the continuous version (see, e.g., Pellilo, 1997, for a detailed analysis).

4.3 Computational Complexity. The computational complexity of finding a Nash equilibrium of a transduction game using equation 4.9 can be given by $\mathcal{O}(kcn^2)$, where n is the number of players (data points), c is the number of pure strategies (classes), and k is the number of iterations needed to converge. In theory, it is difficult to predict the number of required iterations, but experimentally, we noticed that it typically grows linearly on the number of data points.¹ Note that the complexity of popular graph

¹In the case of asymmetric similarities, we have no Lyapunov function for the dynamics, so convergence is not guaranteed. However, if the dynamics converges to a fixed point, it will definitely be a Nash equilibrium (see theorem 1).

transduction methods is also close to $\mathcal{O}(n^3)$. It is of interest that for two-player games, a fast evolutionary game dynamics has been proposed recently, which exhibits linear space and time complexity per iteration (Rota Bulò & Bomze, 2011). One can use this kind of dynamics to increase the efficiency of proposed approach.

5 Connection to Graph-Based Approaches

Contrary to our derivation, the vast majority of graph-based SSL studies start with writing down an objective function that casts the transductive learning problem as an energy minimization problem and most of the focus is on how to compute the optima of the corresponding objective function. In general, all of these methods attempt to estimate an optimal classification function, which is defined on the nodes of the graph by minimizing an objective function with two terms. One term penalizes the mismatch between the initial label assignments and the labels estimated by the classifier. The second term is a regularization term that enforces the smoothness of the classification function. Although the game-theoretic perspective shifts the focus from optima of objective functions to equilibria of the noncooperative games, in this section, we shed some light on the connection between the proposed transduction game to the energy-based graph transduction methods. It is important to note that this analysis investigates only a special case in which the pairwise similarities are assumed to be symmetric and follows from the following property of polymatrix games.

Consider a polymatrix game with $A = (A_{ij})$ being the block matrix representation of partial payoff matrices between players; the average payoff for the whole population can be defined as

$$E(x) = \sum_{i=1}^n x_i^T \left(\sum_{j=1}^n A_{ij} x_j \right) = x^T A x. \quad (5.1)$$

The following proposition establishes a link between local maximizers of $x^T A x$ in Θ and Nash equilibria of polymatrix games with symmetric payoff matrices:

Proposition 1. *Suppose A is symmetric, that is, $A_{ij} = A_{ji}$ for all $i, j \in \mathcal{I}$. Then any local maximum $x^* \in \Theta$ of equation 5.1 is a Nash equilibrium point of the polymatrix game (Hummel & Zucker, 1983; Miller & Zucker, 1991).²*

²This does not hold for the asymmetric case because $x^T A x$ is not a Lyapunov function for the dynamics.

Now consider a special instance of the proposed transduction game where the given pairwise similarities are symmetric: $w_{ij} = w_{ji}, \forall i, j \in \mathcal{D}$. Recall that graph transduction games belong to the family of polymatrix games, and its partial payoff matrices are all in the form $A_{ij} = w_{ij} \times I_c$. To compute Nash equilibria of a transduction game with symmetric similarities, one can use proposition 1 and come up with the following constrained quadratic optimization problem by considering the special form of the partial payoff matrices:

$$\begin{aligned} & \text{maximize } E(X) = \text{tr}\{X^T W X\} \\ & \text{subject to } x_i \in \Delta_i \forall i \in \mathcal{I}_{\mathcal{U}} \\ & \quad x_i = e_i^k \forall i \in \mathcal{I}_{\mathcal{D}|k}, \end{aligned} \tag{5.2}$$

where $X = [x_1 \cdots x_n]^T$ is the $n \times c$ matrix of mixed strategies. Clearly, there is no guarantee that the solution found by the replicator dynamics equation 4.9 will be a global maximizer of equation 5.2. However, it was experimentally shown that on a related problem, the basins of attraction of optimal or near-optimal solutions are quite large and the dynamics often converges to one of them (Pelillo, 1999; Pelillo, Siddiqi, & Zucker, 1999). Moreover, it should be noted that although globally optimal solutions are favorable because they yield the largest consistency, unlike standard approaches, in our game-theoretic interpretation, local optima have a meaningful interpretation and do indeed correspond to solutions of our problem, that is, Nash equilibria of the transduction game.

For the special case above, we can now investigate its relation to energy-based formulations. To begin, the functional in the above problem resembles the continuous relaxation of the k -way normalized cut criterion (Yu & Shi, 2003). However, note that there is a key difference in the game-theoretic formulation in that each mixed strategy x_i is constrained to lie in the c -dimensional standard simplex Δ_i . This subtle difference is very important since it provides robustness against noise and outliers (Pavan & Pelillo, 2007). Moreover, unlike the proposed approach, hard labeling constraints cannot be embedded into the normalized cuts framework in an explicit way, such that partial grouping constraints could be enforced by introducing extra linear equality constraints (Eriksson, Olsson, & Kahl, 2007; Xu, Li, & Schuurmans, 2009; Yu & Shi, 2004). The framework suggested recently in Ghanem and Ahuja (2010) is an exception, but it is inherently a two-class clustering approach and requires a recursive strategy to solve multiclass problems.

In section 4.1, we defined partial payoff matrices for the graph transduction game as $A = I_c \otimes \widehat{W}$, where $\widehat{W} = D^{-1/2} W D^{-1/2}$ is the normalized similarity matrix. Suppose instead that they were specified as $A = I_c \otimes -L$, where $L = D - W$ is the unnormalized graph Laplacian. Then the resulting

optimization problem becomes.

$$\begin{aligned}
 & \text{minimize } E(X) = \text{tr}\{X^T L X\} \\
 & \text{subject to } x_i \in \Delta_i \quad \forall i \in \mathcal{I}_{\mathcal{U}} \\
 & \quad \quad \quad x_i = e_i^k \quad \forall i \in \mathcal{I}_{\mathcal{D}|k}.
 \end{aligned} \tag{5.3}$$

The above optimization problem is in fact equivalent to that of the graph Laplacian regularization used in Zhu et al. (2003), which is known to be a special case of the regularization in Zhou et al. (2004) with the parameter $\mu = \infty$ and the graph Laplacian being unnormalized. In this way, one can argue that the method in Zhu et al. (2003) solves a special case of transduction games in which the pairwise similarities are symmetric and the partial payoffs are specified in terms of negative graph Laplacian.

6 Experimental Results

In this section, we analyze the performance and effectiveness of the proposed approach on some real-world classification problems. To highlight the property that the game-theoretic formulation can naturally deal with symmetric, asymmetric, and negative similarity relations alike, three groups of experiments are carried out. In each group of experiments, we compare our results against state-of-the-art graph-based semisupervised learning algorithms and have obtained reasonably good results. It is important to note that none of the considered graph transduction methods can cope with all three types of similarities.

6.1 Experiments with Symmetric Similarities. In this section, experiments are conducted on four real-world data sets: USPS,³ YaleB (Georghiadis, Belhumeur, & Kriegman, 2001), Scene (Oliva & Torralba, 2001), and 20-news.⁴ Here are some details:

- USPS contains images of handwritten digits 0 to 9 downsampled to 16×16 pixels, and it has 7291 training and 2007 test examples. As in Zhou et al. (2004), only the digits 1 to 4 from the training and test sets are selected, which gives a total of 3874 data points.
- YaleB is composed of face images of 10 subjects captured under varying poses and illumination conditions. As in Breitenbach and Grudic (2005), each image is downsampled to 30×40 pixels, and a subset of 1755 images is considered, which corresponds to the images of individuals 2, 5, and 8.

³<http://www-stat.stanford.edu/~tibs/ElemStatLearn/>.

⁴<http://people.csail.mit.edu/jrennie/20newsgroups/>.

Table 2: Data Sets Used in the Experiments with Symmetric Similarities.

	USPS	YaleB	Scene	20-news
Number of <i>objects</i>	3874	1755	2688	3970
Number of <i>dimensions</i>	256	1200	512	8014
Number of <i>classes</i>	4	3	8	4

- Scene is a scene classification data set consisting of 2688 natural scene images classified into one of eight classes. Each image is represented with a 512-dimensional GIST descriptor (Oliva & Torralba, 2001), which combines the outputs of Gabor-like filters specifically designed to capture the structural properties of a scene.
- 20-news is the text classification data set used in Zhou et al. (2004), which contains 3970 news group articles selected from the 20-news groups data set, all belonging to the topic *rec*, which is composed of the subjects *autos*, *motorcycles*, *sport.baseball*, and *sport.hockey*. As Zhou et al. (2004) describes, each article is represented in 8014-dimensional space based on the TFIDF representation scheme.

Table 2 shows a summary of the data sets. For USPS and YaleB, each image pixel is treated as a single feature; thus, each example is represented in 256- and 1200-dimensional space, respectively. The similarity between two examples d_i and d_j is computed using the gaussian kernel as $w_{ij} = \exp(-\frac{\text{dist}(d_i, d_j)^2}{2\sigma^2})$, where $\text{dist}(d_i, d_j)$ is the distance between d_i and d_j and σ is the kernel width parameter. Among several choices for the distance measure $\text{dist}(\cdot)$, the Euclidean distance $\|d_i - d_j\|$ is evaluated for USPS, YaleB, and Scene, and the cosine distance $\text{dist}(d_i, d_j) = 1 - \frac{\langle d_i, d_j \rangle}{\|d_i\| \|d_j\|}$ is evaluated for 20-news.

In the experiments, the proposed approach (denoted here with GTG) is compared against four well-known graph-based SSL algorithms: the spectral graph transducer (SGT; Joachims, 2003),⁵ the gaussian fields and harmonic functions-based method (GFHF; Zhu et al., 2003),⁶ the local and global consistency method (LGC; Zhou et al., 2004),⁷ and Laplacian regularized least squares (LapRLS; Belkin, Niyogi, & Sindhwani, 2006).⁸

⁵The optimal value of the parameter c is selected as the one with the best mean performance from the set {400, 800, 1600, 3200, 6400, 12800}.

⁶In obtaining the hard labels, the class mass normalization step is employed as suggested in Zhu et al. (2003).

⁷As in Zhou et al. (2004), the parameter α is set as 0.99.

⁸The optimal values of the extrinsic and intrinsic regularization parameters γ_A and γ_I are selected from the set $\{10^{-6}, 10^{-4}, 10^{-2}, 1\}$ for the best mean performance.

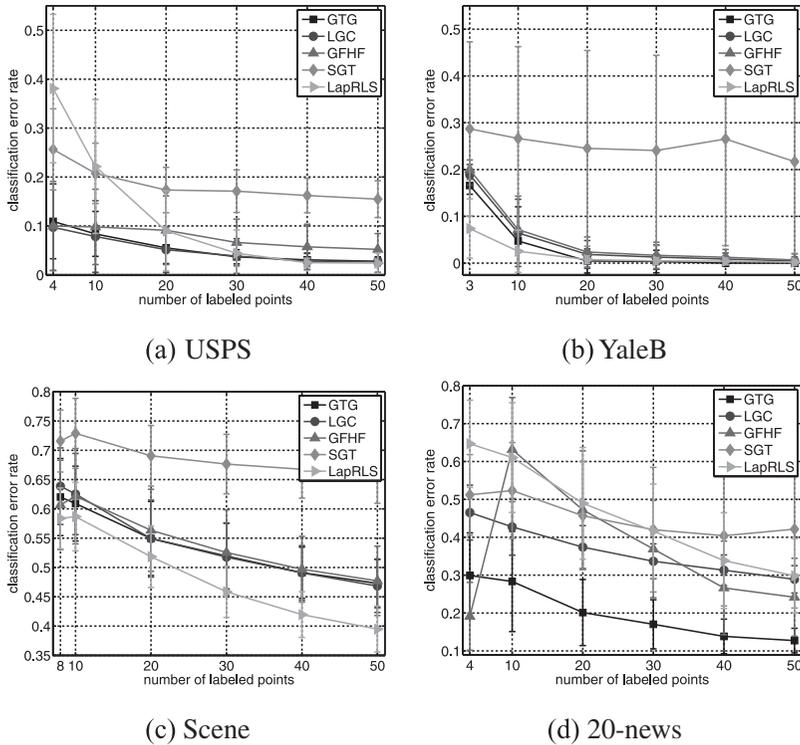


Figure 2: Performance comparisons on classification problems with symmetric similarities.

A crucial factor in the success of graph-based algorithms is the construction of the input graph as it represents the data manifold. As a result, graph construction for classification has become a subject of interest in recent years (Jebara et al., 2009), but the problem is still open. To be fair in the evaluation, for all the methods, a fixed set of kernel widths is used, and nine 20-NN graphs are generated by setting $w_{ij} = 0$ if x_j is not among the 20 nearest neighbors of x_i . In particular, the kernel width σ ranges over the set $\text{linspace}(0.1r, r, 5) \cup \text{linspace}(r, 10r, 5)$ with r being the average distance from each example to its twentieth nearest neighbor and $\text{linspace}(a, b, n)$ denoting the set of n linearly spaced numbers between and including a and b .

Figure 2 shows the mean test error rates and one-standard-deviation error bars over 100 trials with different sizes of labeled data.⁹ The LapRLS

⁹We randomly select labeled samples such that each set contains at least one sample from each class.

Table 3: Data Sets Used in the Classification Experiments with Asymmetric Similarities.

	SCOP	Cora	Citeseer	Cornell	Texas	Washington	Wisconsin
Number of <i>objects</i>	451	2708	3312	827	814	1166	1210
Number of <i>classes</i>	5	7	6	2	2	2	2

method gives the best results for the relatively small data sets, YaleB and Scene. However, for the other two, its performance is poor. In general, the proposed GTG algorithm is either the best or the second best-algorithm. Although its success is almost identical to that of the LGC method in USPS, Yale-B, and Scene, it gives superior results for 20-news.

6.2 Experiments with Asymmetric Similarities. In this section, several experiments are carried out on a real-world protein data set derived from SCOP (structural classification of proteins; Murzin, Brenner, Hubbard, & Chothia, 1995),¹⁰ and three document data sets: Cora, Citeseer (Sen et al., 2008),¹¹ and WebKB.¹² The details of these data sets are given below:

- SCOP contains 7329 proteins, which are hierarchically divided into seven classes based on structural and evolutionary relationships. Each class is divided into folds, and each fold is further divided into superfamilies. Similar to the setup in Meila and Pentney (2007), only the proteins from the five largest folds of the all alpha class are selected, which give a total of 451 protein sequences to be classified by fold.
- Cora contains 2708 machine learning publications classified into seven classes, for a total 5429 citations.
- Citeseer consists of 3312 scientific publications, each of which belongs to one of six classes, for a total of 4732 links.
- WebKB contains Web pages collected from the computer science departments of four universities (Cornell, Texas, Washington, and Wisconsin), and each is classified into seven categories. Following the setup in Zhou, Huang, and Schölkopf, (2005), here we concentrate on classifying student pages from the others. Each subset respectively contains 827, 814, 1166, and 1210 Web pages and 1626, 1480, 2218, and 3200 links.

The data sets are summarized in Table 3. For SCOP, *E*-values of the PSI-BLAST search calculated by Weston, Elisseeff, Zhou, Leslie, and Noble

¹⁰Version 1.59 with less than 95% identity.

¹¹Cora and Citeseer are available at <http://www.cs.umd.edu/projects/linqs/projects/lbc/>.

¹²Available at <http://www.nec-labs.com/~zsh/files/link-fact-data.zip>.

(2004) are considered as the dissimilarity scores.¹³ These dissimilarity values are not symmetric. Instead of constructing NN graphs, this time we use the full similarity matrices where the kernel width σ was optimized for each method with respect to the set $\text{linspace}(0.5, 2.5, 5) \cup \text{linspace}(5, 20, 4) \cup 40$. For Cora, Citeseer, and WebKB, as in Zhou et al. (2005), only the citation/link structure is considered, even though one can also assign some weights by utilizing the textual content of the documents. Specifically, the experiments are performed on the link matrix $W = (w_{ij})$, where $w_{ij} = 1$ if document i cites document j and $w_{ij} = 0$ otherwise.

Unlike the proposed game-theoretic approach, the standard methods of SGT, GFHF, LGC, and LapRLS are subject to symmetric similarities. Hence, in this context, they can be applied only after rendering the similarities symmetric, but this could result in loss of relevant information in some cases. In the evaluation, only the graph-based methods that can directly deal with asymmetric similarities are considered. Specifically, the proposed game-theoretic approach is compared against our implementation of the method in Zhou et al. (2005), denoted here with LLUD. This algorithm is based on the notion of random walks on directed graphs, and it is equivalent to LGC in the case of symmetric similarities. However, it assumes the input similarity graph to be strongly connected, so Zhou et al. (2005) consider the teleporting random walk (trw) transition matrix as input, which is given by $P^\eta = \eta P + (1 - \eta)P^u$, where $P = D^{-1}W$ and P^u is the uniform transition matrix. This suggests a second variant for our framework, denoted with GTGtrw, where payoffs are defined in terms of this transition matrix. In the experiments, we fix $\eta = 0.99$ for both LLUD and GTGtrw. To provide a baseline, we also report the results of our approach that works on the symmetrized similarity matrices, denoted with GTGsym. For that case, we used the transformation $\tilde{W} = 0.5 \times (W + W^T)$ for SCOP and the symmetrized link matrix $\tilde{W} = (\tilde{w}_{ij})$ for the others, where $\tilde{w}_{ij} = 1$ if either document i cites document j or vice versa, and $\tilde{w}_{ij} = 0$ otherwise.

The test errors averaged over 100 trials are shown in Figure 3. Recall that the replicator dynamics used to find a Nash equilibrium is not guaranteed to converge in the case of asymmetric similarities. In fact, GTG succeeded in finding a solution only on SCOP and failed on the others. However, we did not face any convergence problem with GTGtrw. Hence, we suspect that this might be related to the high sparseness of the data. Notice that the performances of GTGtrw and LLUD are quite similar on the classification problems in the WebKB data sets. On the other hand, GTGtrw is superior in the multiclass problems in SCOP, Cora, and Citeseer. It is important to note that symmetrization sometimes can provide good results. As shown

¹³Available online at <http://www.kyb.tuebingen.mpg.de/bs/people/weston/rankprot/supplement.html>.

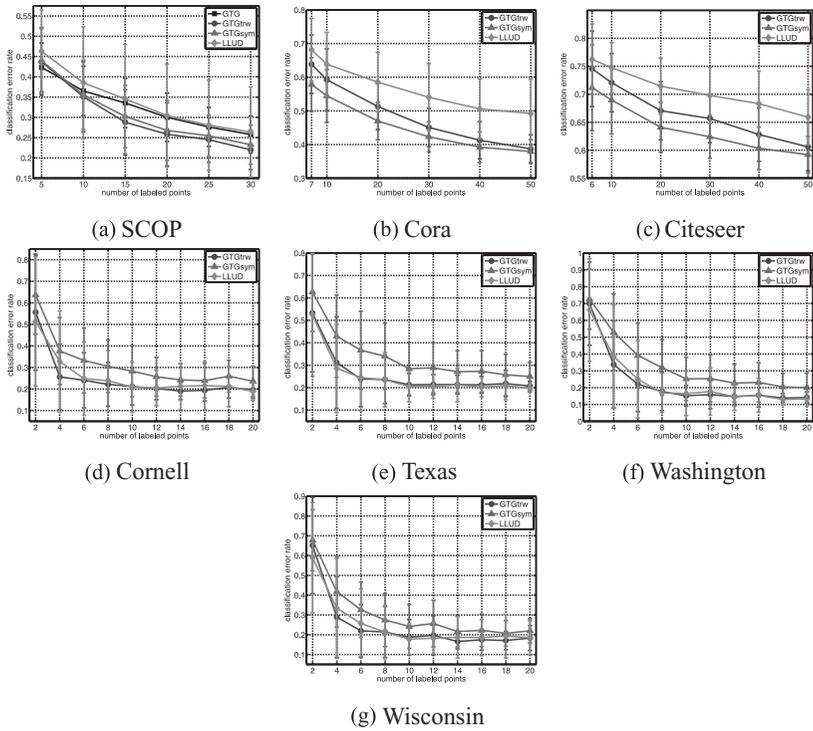


Figure 3: Performance comparisons on classification problems with asymmetric similarities.

in Figures 3(b) and (c), in Cora and Citeseer, GTGsym performs better than the other two methods.

6.3 Experiments with Negative Similarities. The game-theoretic formulation can also handle negative similarity or dissimilarity information in two-class semisupervised classification problems in a natural way. In practice, such dissimilarity relationships could arise in the computations, or they can be provided explicitly in terms of a set of cannot-link pairs where the objects in each pair are expected to be assigned to different classes. The standard methods cannot accept negative similarities as well, since negative similarities could make their energy functions negatively unbounded. However, there are some graph-based SSL methods specifically designed for dealing with dissimilarity information, that is, the mixed label propagation method (MLP; Tong & Jin, 2007) and manifold regularization with dissimilarity method (MRWD) (Goldberg, Zhu, & Wright, 2007), which is

Table 4: Data Sets Used in the Classification Experiments with Negative Similarities.

	Ionosphere	Diabetes
Number of <i>objects</i>	351	768
Number of <i>dimensions</i>	34	8
Number of <i>classes</i>	2	2

adopted from Belkin et al. (2006).¹⁴ In this section, the proposed method is compared against these two algorithms.

The experiments are performed on two data sets from the UCI repository:¹⁵ Ionosphere and Diabetes. These data sets are described in Table 4. Notice that the similarities derived from the data sets did not originally contain any negative values, but we adopted the procedure in Goldberg et al. (2007) and introduced oracle dissimilarity relations (cannot links) by randomly sampling pairs of examples having different labels. These pairs do not contain any labeled samples, and to enforce the maximum degree of dissimilarity, the edge weights were set to the maximum similarity value existing in the data. In the experiments, the size of the labeled data is fixed as 50, and the dissimilarity edges are varied between 3 and 12,800. For all data sets, full similarity matrices are used, and the kernel width σ is optimized with respect to the set $\{0.01\} \cup \text{linspace}(0.05, 0.25, 5) \cup \text{linspace}(0.25, 2.5, 10) \cup \{5, 10, 20, 25\}$.

The average test errors over 10 trials with randomly selected labeled examples and dissimilarity edges are given in Figure 4. Because the methods under consideration explore both similarity and dissimilarity information, their accuracy improves as the size of the dissimilarity edges increases. Although there is no considerable difference in the performances of the methods on Diabetes, GTG is clearly more successful on Ionosphere.

Conclusion

This article introduces a novel game-theoretic interpretation to graph transduction in which the problem is formulated in terms of a polymatrix game whereby any equilibrium coincides with the notion of a consistent labeling of the data. As compared to existing approaches, the main advantage of the proposed framework is that there is no restriction on the pairwise relationships among data points; similarities, and thus the payoffs, can be

¹⁴In the absence of negative similarities, MRWD reduces to LapRLS. In the experiments, we used the regularized least squares (RLS) classifier and selected the optimal values of the extrinsic and intrinsic regularization parameters as described for LapRLS.

¹⁵<http://archive.ics.uci.edu/ml/>.

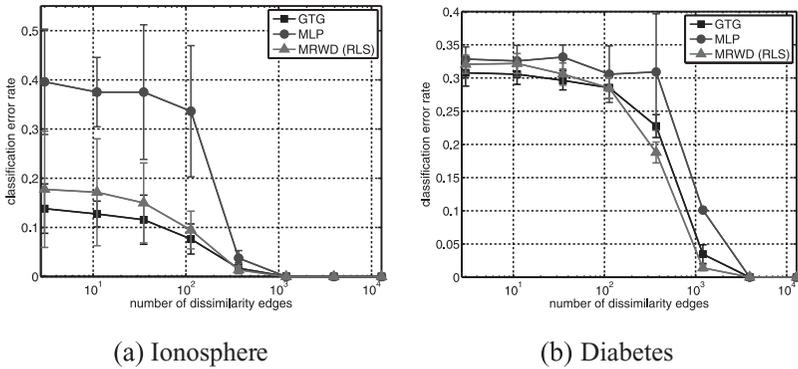


Figure 4: Performance comparisons on classification problems with negative similarities.

negative or asymmetric. Apart from that, our approach is easy to implement and can be applied to multiclass problems. The experimental results show that the game-theoretic approach is not only more general but also competitive with standard approaches. As future work, we plan to focus on improving efficiency. In our current implementation, we use the standard replicator dynamics to reach an equilibrium, but we can study other selection dynamics that are much faster (Porter et al., 2008; Rota Bulò & Bomze, 2011). Another possible direction for future research is to generalize the presented approach to transductive learning in hypergraphs (Agarwal, Branson, & Belongie, 2006; Zhou, Huang, & Schölkopf, 2007). This will require replacing the pairwise interactions with higher-order interactions in defining payoffs, along the lines proposed in Rota Bulò and Pelillo (2009) for unsupervised learning. In this context, it would be especially interesting to explore whether different classes of games such as action-graph games (Jiang, Leyton-Brown, & Bhat, 2011) will be more suitable.

Acknowledgments

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

References

- Agarwal, S., Branson, K., & Belongie, S. (2006). Higher order learning with graphs. In W. W. Cohen & A. Moore (Eds.), *Proc. 23rd Intl. Conf. Machine Learning* (pp. 17–24). New York: ACM.
- Belkin, M., Niyogi, P., & Sindhvani, V. (2006). Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *J. Mach. Learn. Res.*, 7, 2399–2434.

- Breitenbach, M., & Grudic, G. Z. (2005). Clustering through ranking on manifolds. In L. De Raedt & S. Wrobel (Eds.), *Proc 22nd Intl. Conf. on Machine Learning* (pp. 73–80). New York: ACM.
- Chapelle, O., Schölkopf, B., & Zien, A. (Eds.) (2006). *Semi-supervised learning*. Cambridge, MA: MIT Press.
- Daskalakis, C. (2011). On the complexity of approximating a Nash equilibrium. In L. De Raedt & S. Wrobel (Eds.), *Proc. 22nd Intl. Conf. on Machine Learning* (pp. 73–80). Philadelphia: SIAM.
- Daskalakis, C., Goldberg, P. W., & Papadimitriou, C. H. (2009). The complexity of computing a Nash equilibrium. *Commun. ACM*, 52(2), 89–97.
- Dickhaut, J., & Kaplan, T. (1991). A program for finding Nash equilibria. *Math. J.*, 1(4), 87–93.
- Easley, D., & Kleinberg, J. (2010). *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge: Cambridge University Press.
- Eriksson, A. P., Olsson, C., & Kahl, F. (2007). Normalized cuts revisited: A reformulation for segmentation with linear grouping constraints. In *Proc. IEEE 11th Intl. Conf. on Computer Vision*. Piscataway, NJ: IEEE.
- Georgiades, A. S., Belhumeur, P. N., & Kriegman, D. J. (2001). From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23, 643–660.
- Ghanem, B., & Ahuja, N. (2010). Dinkelbach NCUT: An efficient framework for solving normalized cuts problems with priors and convex constraints. *Int. J. Comput. Vis.*, 89, 40–55.
- Goldberg, A., Zhu, X., & Wright, S. (2007). Dissimilarity in graph-based semisupervised classification. In M. Melia & X. Shen (Eds.), *Proc. 11th Intl. Conf. on Artificial Intelligence and Statistics*. N.p.: Society for Artificial Intelligence and Statistics.
- Govindan, S., & Wilson, R. (2003). A global Newton method to compute Nash equilibria. *J. Econ. Theory*, 110, 65–86.
- Haralick, R. M., Davis, L. S., Rosenfeld, A., & Milgram, D. L. (1978). Reduction operations for constraint satisfaction. *Inf. Sci.*, 199–219.
- Haralick, R. M., & Shapiro, L. G. (1979). The consistent labeling problem: Part I. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1(2), 173–184.
- Hofbauer, J., & Sigmund, K. (1998). *Evolutionary games and population dynamics*. Cambridge: Cambridge University Press.
- Howson, J. T. (1972). Equilibria of polymatrix games. *Management Science*, 18, 312–318.
- Hummel, R. A., & Zucker, S. W. (1983). On the foundations of relaxation labeling processes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 5, 267–287.
- Jacobs, D. W., Weinshall, D., & Gdalyahu, Y. (2000). Classification with nonmetric distances: Image retrieval and class representation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22, 583–600.
- Janovskaya, E. B. (1968). Equilibrium points in polymatrix games. *Litovskii Matematicheskii Sbornik*, 8, 381–384 (in Russian).
- Jebara, T., Wang, J., & Chang, S. (2009). Graph construction and b -matching for semisupervised learning. In A. P. Danyluk, L. Bottou, & M. L. Littman (Eds.), *Proc. 26th Intl. Conf. on Machine Learning* (pp. 441–448). New York: ACM.

- Jiang, A. X., Leyton-Brown, K., & Bhat, N.A.R. (2011). Action-graph games. *Games and Economic Behavior*, 71, 141–173.
- Joachims, T. (2003). Transductive learning via spectral graph partitioning. In T. Fawcett & N. Mishra (Eds.), *Proc. 20th Intl. Conf. on Machine Learning* (pp. 290–297). Cambridge, MA: AAAI Press.
- Laub, J., Roth, V., Buhmann, J. M., & Müller, K-R. (2006). On the information and representation of non-Euclidean pairwise data *Pattern Recognition*, 39, 1815–1826.
- Mangasarian, O. (1964). Equilibrium points of bimatrix games. *J. Soc. Ind. Appl. Math.*, 12, 778–780.
- Marriott, K., & Stuckey, P. J. (1998). *Programming with constraints: An introduction*. Cambridge MA: MIT Press.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.
- Meila, M., & Pentney, W. (2007). Clustering by weighted cuts in directed graphs. In *SIAM Conference on Data Mining*. Philadelphia: SIAM.
- Miller, D. A., & Zucker, S. W. (1991). Copositive-plus Lemke algorithm solves poly-matrix games. *Operations Research Letters*, 10, 285–290.
- Murzin, A. G., Brenner, S. E., Hubbard, T., & Chothia, C. (1995). *SCOP: A structural classification of proteins database for the investigation of sequences and structures*. <http://scop.berkeley.edu/>.
- Nash, J. (1951). Non-cooperative games. *Annals of Mathematics*, 54, 286–295.
- Nisan, N., Roughgarden, T., Tardos, E., & Vazirani, V. V. (2007). *Algorithmic game theory*. Cambridge: Cambridge University Press.
- Oliva, A., & Torralba, A. (2001). Modeling the shape of the scene: A holistic representation of the spatial envelope. *Int. J. Comput. Vis.*, 42, 145–175.
- Pavan, M., & Pelillo, M. (2007). Dominant sets and pairwise clustering. *IEEE Trans. on Pattern Anal. and Mach. Intell.*, 29, 167–172.
- Pelillo, M. (1997). The dynamics of nonlinear relaxation labeling processes. *J. Math. Imaging Vis.*, 7, 309–323.
- Pelillo, M. (1999). Replicator equations, maximal cliques, and graph isomorphism. *Neural Computation*, 11, 1933–1955.
- Pelillo, M., Siddiqi, K., & Zucker, S. W. (1999). Matching hierarchical structures using association graphs. *IEEE Trans. on Pattern Anal. and Mach. Intell.*, 21, 1105–1120.
- Porter, R. W., Nudelman, E., & Shoham, Y. (2008). Simple search methods for finding a Nash equilibrium. *Games and Economic Behavior*, 63, 642–662.
- Rota Bulò, S., & Bomze, I. M. (2011). Infection and immunization: A new class of evolutionary game dynamics. *Games and Economic Behaviour*, 71, 193–211.
- Rota Bulò, S., & Pelillo, M. (2009). A game-theoretic approach to hypergraph clustering. In Y. Bengio, D. Schuurmans, J. Lafferty, C.K.I. Williams, & A. Culotta (Eds.), *Advances in neural information processing systems*, 22 (pp. 1571–1579). Cambridge, MA: MIT Press.
- Santini, S., & Jain, R. (1999). Similarity measures. *IEEE Trans. on Pattern Anal. and Mach. Intell.*, 21, 871–883.
- Sen, P., Namata, G. M., Bilgic, M., Getoor, L., Gallagher, B., & Eliassi-Rad, T. (2008). Collective classification in network data. *AI Magazine*, 29, 93–106.

- Tong, W., & Jin, R. (2007). semisupervised learning by mixed label propagation. In *Proc. 22nd AAAI Conf. on Artificial Intelligence* (pp. 651–656). Cambridge, MA: MIT Press.
- Tsang, E. (1993). *Foundations of constraint satisfaction*. San Diego, CA: Academic Press.
- Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*, 327–352.
- van der Laan, G., Talman, A., & van der Heyden, L. (1987). Simplicial variable dimension 1987 for solving the nonlinear complementarity problem on a product of unit simplices using a general labelling. *Math. Operations Res.*, *12*, 377–397.
- von Luxburg, U., Bousquet, O., & Belkin, M. (2004). On the convergence of spectral clustering on random samples: The normalized case. In J. Shawe-Taylor & Y. Singer (Eds.), *Proc. 17th Annual Conf. on Learning Theory* (pp. 457–471). Berlin: Springer.
- Waltz, D. L. (1975). Understanding line drawings of scenes with shadows. In P. H. Winston (Ed.), *The psychology of computer vision* (pp. 19–92). New York: McGraw-Hill.
- Weibull, J. W. (1995). *Evolutionary Game Theory*. Cambridge, MA: MIT Press.
- Weston, J., Elisseeff, A., Zhou, D., Leslie, C., & Noble, W. S. (2004). Protein ranking: From local to global structure in the protein similarity network. *Proc. National Academy of Science*, *101*, 6559–6563.
- Xu, L., Li, W., & Schuurmans, D. (2009). Fast normalized cut with linear constraints. In *CVPR*, 2866–2873.
- Yu, S. X., & Shi, J. (2003). Multiclass spectral clustering. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition* (pp. 313–319). Piscataway, NJ: IEEE.
- Yu, S. X., & Shi, J. (2004). Segmentation given partial grouping constraints. *IEEE Trans. Pattern Anal. Mach. Intell.*, *26*, 173–183.
- Zhou, D., Bousquet, O., Lal, T. N., Weston, J., & Schölkopf, B. (2004). Learning with local and global consistency. In L. Saul, Y. Weiss, & L. Bottou (Eds.), *Advances in neural information processing systems*, *17*. Cambridge, MA: MIT Press.
- Zhou, D., Huang, J., & Schölkopf, B. (2005). Learning from labeled and unlabeled data on a directed graph. In L. De Raedt & S. Wrobel (Eds.), *Proc. 22nd Intl. Conf. on Machine Learning* (pp. 1036–1043). New York: ACM.
- Zhou, D., Huang, J., & Schölkopf, B. (2007). Learning with hypergraphs: Clustering, classification, and embedding. In J. C. Platt, D. Koller, Y. Singer, & S. Roweis (Eds.), *Advances in neural information processing systems*, *20* (pp. 1601–1608). Cambridge, MA: MIT Press.
- Zhu, X. (2005). *Semi-supervised learning literature survey* (Tech. Rep. No. 1530). Madison: Computer Sciences, University of Wisconsin-Madison.
- Zhu, X., Ghahramani, Z., & Lafferty, J. (2003). semisupervised learning using gaussian fields and harmonic functions. In T. Fawcett & N. Mishra (Eds.), *Proc. 20th Intl. Conf. on Machine Learning* (pp. 912–919). Cambridge, MA: MIT Press.