Harris V. Georgiou (MSc, PhD)

Games People Play

Conflicts, mechanisms and collective decision-making in expert committees

A short collection of introductory papers

Athens, Greece – 2015

Harris V. Georgiou (MSc, PhD)

Games People Play

Conflicts, mechanisms and collective decision-making in expert committees

A short collection of introductory papers

Athens, Greece – 2015

Harris V. Georgiou
(MSc, PhD)

Signal & Image Processing Lab (SIPL),
Department of Informatics and Telecommunications (DIT),
National & Kapodistrian University of Athens, Greece (NKUA/UoA).

Email: harris@xgeorgio.info – URL: http://xgeorgio.info

Games People Play
Conflicts, mechanisms and collective decision-making in expert committees
(A short collection of introductory papers)

Athens, Greece – 2015

Preface

Game Theory is one of the most challenging and controversial fields of applied Mathematics. Based on a robust theoretical framework, its applications range from analyzing simple board games and conflict situations to modeling complex systems and evolutionary dynamics.

This book is a short collection of introductory papers in the field, aimed primarily as reading material for graduate- and postgraduate-level lectures in Game Theory and/or Machine Learning. The four papers included here are all original works already published as open-access or conference publications, spanning a timeframe of several years apart and a wide range of topics. Hence, each paper is self-contained and can be studied on its own, without any prerequisite knowledge from the previous ones. However, their presentation order is consistent with going from the most elementary issues to the more advanced and experiment-rigorous topics.

The first paper presents an overview of Game Theory in general, its core issues and building blocks, game analysis and methods for identifying Minimax solutions and Nash equilibria, as well as a brief introduction to coalitional gaming and collective efficiency. There is also a short summary of other important elements like signaling, credibility, threats/promises, etc. The second paper extends some of the topics from coalitional gaming, focusing more on collective efficiency, optimal voting mechanisms and weighted voting, as well as a brief proposal for applying this game-theoretic framework to optimal combination of experts. The third paper builds upon this proposed framework and employs it in Pattern Recognition (Machine Learning) within the context of combining pattern classifiers. A “static” model for weighted majority voting with an analytical model for the voting weights is experimentally tested against other similar models. Finally, the forth paper presents an extension of this game-theoretic approach for classifier combination, employing “adaptive” voting weights via local accuracy estimates; in other words, the ensemble of classifiers is adapted to local efficiency priors (instead of static globals) but keeping the same analytical model for the voting weights, i.e., without the need to acquire them via training. This new approach is experimentally validated against state-of-the-art combination methods for pattern classifiers and it is proven highly competitive with much lower complexity overhead.

These papers are all part of the author’s PhD work, conducted at the Department of Informatics and Telecommunications (DIT), National & Kapodistrian University of Athens, Greece (NKUA/UoA). The author wishes to give special thanks to prof. Sergios Theodoridis (supervisor) and prof. Michael Mavroforakis, colleagues and friends, for their valuable collaboration in parts of these works.

Table of Contents

Section 1: Elements of Game Theory

Summary:

In this paper, a gentle introduction to Game Theory is presented in the form of basic concepts and examples. Minimax and Nash's theorem are introduced as the formal definitions for optimal strategies and equilibria in zero-sum and nonzero-sum games. Several elements of cooperative gaming, coalitions, voting ensembles, voting power and collective e-ciency are described in brief. Analytical (matrix) and extended (tree-graph) forms of game representation is illustrated as the basic tools for identifying optimal strategies and “solutions” in games of any kind. Next, a typology of four standard nonzero-sum games is investigated, analyzing the Nash equilibria and the optimal strategies in each case. Signaling, stance and third-party intermediates are described as very important properties when analyzing strategic moves, while credibility and reputation is described as crucial factors when signaling promises or threats. Utility is introduced as a generalization of typical cost/gain functions and it is used to explain the incentives of irrational players under the scope of “rational irrationality”. Finally, a brief reference is presented for several other more advanced concepts of gaming, including emergence of cooperation, evolutionary stable strategies, two-level games, metagames, hypergames and the Harsanyi transformation.

Citation:

"Elements of Game Theory – Part I: Foundations, acts and mechanisms",
H. Georgiou, ArXiv.org: 16-Jun-2015 (arXiv: 1506.05148v1 [cs.GT]).

"Games people play: An overview of strategic decision-making theory in conflict situations", H. Georgiou, viXra.org: 15-Jun-2015 (viXra: 1506.0114 [GenMath]).

Elements of Game Theory

Part I: Foundations, acts and mechanisms.

*Harris V. Georgiou (MSc, PhD)**

Department of Informatics and Telecommunications,
National & Kapodistrian University of Athens, Greece.

Abstract

In this paper, a gentle introduction to Game Theory is presented in the form of basic concepts and examples. Minimax and Nash's theorem are introduced as the formal definitions for optimal strategies and equilibria in zero-sum and nonzero-sum games. Several elements of cooperative gaming, coalitions, voting ensembles, voting power and collective efficiency are described in brief. Analytical (matrix) and extended (tree-graph) forms of game representation is illustrated as the basic tools for identifying optimal strategies and "solutions" in games of any kind. Next, a typology of four standard nonzero-sum games is investigated, analyzing the Nash equilibria and the optimal strategies in each case. Signaling, stance and third-party intermediates are described as very important properties when analyzing strategic moves, while credibility and reputation is described as crucial factors when signaling promises or threats. Utility is introduced as a generalization of typical cost/gain functions and it is used to explain the incentives of irrational players under the scope of "rational irrationality". Finally, a brief reference is presented for several other more advanced concepts of gaming, including emergence of cooperation, evolutionary stable strategies, two-level games, metagames, hypergames and the Harsanyi transformation.

Keywords: Game Theory, Minimax theorem, Nash equilibrium, coalitional gaming, indices of power, voting ensembles, signaling, bluff, credibility, promises, threats, utility function, two-level games, hypergames, evolutionary stable strategies, Harsanyi transformation, metagames.

GAME THEORY is a vast scientific and research area, based almost entirely on Mathematics and some experimental methods, with applications that vary from simple board games to Evolutionary Psychology and Sociology-Biology in group behavior of humans and animals. Conflict situations are presented everywhere in the real world, every day, for thousands of years - not only in human societies but also in animals. The seller and the buyer have to come up with a mutually acceptable price for the grocery. The employer and the employee have to bargain in order to reach a mutually satisfying value for the salary. A buyer in an auction

*Email: harris@xgeorgio.info — URL: http://xgeorgio.info

has to continuously estimate the cost/gain value of making (or not) the next higher bid for some object. The primary adversaries in a wolf pack have to decide when it is beneficial to fight over the leadership and when to stop before they are severely wounded. A swarm of fish has to collectively "decide" what is the optimal number and distance of the piket members or "scouts" that serve as the early warning for the group, perhaps even self-sacrificing if required. All these cases are typical examples, simpler or more complex, of conflict situations that depend on bargaining, coordination and evolutionary optimization. Game Theory provides a unified framework with robust mathematical foundations for the proper formulation and analysis of such systems.

1 The building blocks

In principle, the mathematical theory of games and gaming was first developed as a model for situations of conflict. *Game Theory* is the area of research that provides mathematical formulations and a proper framework for studying adversarial situations. Although E. Borel looked at similar problems in the 1920s, John Von Neumann and Oskar Morgenstern provided two breakthrough papers (1928, 1937) as a kick-start of the field. Since the early 1940's, with the end of World War II and stepping into the era of the Cold War that followed, the work of Von Neumann and Morgenstern has provided a solid foundation for the most simple types of games, as well as analytical forms for their solutions, with many applications to Economics, Operations Research and Logistics. However, there are several limitations that fail to explain various aspects of real-world conflicts [25], especially when the human factor is a major factor. The application of game-theoretic formulations in designing experiments in Psychology and Sociology is usually referred to as *gaming* [46, 6].

1.1 Games, strategies and solutions

The term *game* is the mathematical formulation of adversarial situations, where two or more *players* are involved in competitive or cooperative acts. The *zero-sum* games are able to model situations of conflict between two or more players, where one's gain is the other's loss and vice versa. Most military problems can be modeled as some form of two-player zero-sum game. When the structure of the game and the rationale of the players is known to all, then the game is one of *complete information*, while if some of these information is somehow hidden or unknown to some players, it is one of *incomplete information*. Furthermore, if all players are fully informed about their opponents' decisions, the game is one of *perfect information*. In contrast, if some of the information about the other players' moves, the game is one of *partial* or *imperfect information*. Such games of both complete and perfect information are all board games, like Chess, Go and Checkers, and they are all zero-sum by nature.

Von Neumann and Morgenstern [48] proved that there is at least one optimal plan of decisions or *strategy* for each player in all zero-sum games, as well as a *solution* to the game that comes naturally as a result of all players following their optimal strategies. At the game's solution, each player can guarantee that the maximum gain an opponent can gain is kept under a specific minimal limit, defined only by this player's own strategy. This assertion was formulated as a theorem called *Minimax* and in the simple case of two opposing players with only two strategies each the Minimax solution of the game can be calculated analytically as a solution of a 2x2 set of linear equations, which determine the stable solution or *saddle-point*.

The consequences of the Minimax theorem have been thoroughly studied for many years after its proof. As an example, it mathematically proves the assertion that all board games, including the most complex ones like Chess, have at least one solution, i.e., an optimal (pure) strategy for both players that can be analytically calculated, at least in theory [44, 46, 37]. Of course, in the case of Chess the game space is so huge that it is still unfeasible today to calculate this theoretically optimal strategy, even with the help of parallel processing in supercomputers. In contrast, Checkers is a much smaller (3x3) and simpler game, making it possible to create the complete game space in any typical desktop computer[1] and calculate the exact optimal strategy - in fact, it is the same strategy that every child soon learns by trial-and-error, playing in a way that always leads to a win or a draw (never loose).

In general, if the chosen strategy of one player is known to its opponent, then an optimal counter-strategy is always available. Hence, in *simultaneous* games where the opposing moves are conducted at the same time, each player would normally try not to employ a deterministic way of choosing its strategy and conceal this choice until the very last moment. However, the Minimax theorem provides a mathematically solid way of nullifying any stochastic aspect in determining the opponent's choice and, in essence, make its exact choice irrelevant: no matter what the opponent does, the Minimax solution ensures the *minimum losses* to each player, given a specific game setup. In other words, it provides an analytic way to determine the best *defensive* strategy, instead of a preference to offensive strategies. In some zero-sum games this leads to one stable outcome or *equilibrium*, where each player would have no incentive not to choose its Minimax strategy; however, if this choice leads to a negative handicap for this player if it is known with complete certainty by the others, then this choice should not be manifested as certain. In practice this means that the Minimax solution would not be any single one of the player's *pure* strategies but rather a weighted combination of them in a *mixed* strategy scheme, where each weight corresponds to the probability of choosing one of the available pure strategies via a random mechanism. This notion of using mixtures of pure strategies for randomly choosing between them leads to a false sense of security in single-turn games, since the optimality of the expected outcome of the mixed strategy scheme refers to the asymptotic (long-term) and not the "spot" (one-shot) payoff. Moreover, a game may involve an infinite number of strategies for the players, in a discrete or continuous set; in this case the game is labeled as *continuous* or *infinite*, while a *finite* game is one with a limited number of

[1] In Checkers, the board size is 3x3 and each position can be either empty or host the mark of of one of the two players, "X" or "O". Hence, if the two players are treated as interchangeable (i.e., who plays first) and no other symmetries are considered, the total number of all possible distinct board setups is: $9 \cdot 8 \cdot \ldots \cdot 2 \cdot 1 = 9! = 362,880$. After applying the game rules and pruning the game tree for early stops (with incomplete boards), the true number of game states is about 2/3 of that set. Using simple tree-node representation for each board setup, e.g. a 3-value 9-positions vector dictionary ($= 3^9 \simeq 2^{14.265} \leq 2^{15} < 2^{16} = 2$ bytes), such a program would only require about 484 KB or less than 0.5 MB. This is roughly the size of a small-sized photo taken by the camera of a low-end smart-phone today, while in the '80s this was almost the total size of RAM in a typical PC.

(discrete) strategies [14, 46].

When the game is inherently repetitive or *iterative*, i.e., includes multiple turns and not just one, even the pure strategy suggested by Minimax should not be chosen deterministically in every turn if according to the game setup this information might provide a handicap to the opponent. This is a topic of enthusiastic discussion about the optimality of the Minimax solution and its inherent defensive nature, as it is not clear in general when information about an opponent's *next* move is available and trustworthy enough to justify any deviation from this Minimax strategy.

Summary:

- In *zero-sum* games, one player's gains is another's losses (and vice versa).
- Information about the game structure and the opponents' moves may be *complete* or not, *perfect* or not.
- All board games are inherently zero-sum, of complete and perfect information.
- The *Minimax theorem* assures that all board games have at least one theoretically optimal way to play them, although its exact calculation may be unfeasible in practice for some games (e.g. Chess, Go).
- The *Minimax solution* of a game is the combination of players' strategies that lead to an *equilibrium* or *saddle-point.*

1.2 Nonzero-sum games and Nash equilibria

Although the Minimax theorem provided a solid base for solving many types of games, it is only applicable in practice for the zero-sum type of games. In reality, it is common that in a conflict not all players receive their opponents' looses as their own gain and vice versa. In other words, it is very common a specific combination of decisions between the players to result in a certain amount of "loss" to one and a corresponding "gain", not of equal magnitude, to another. In this case, the game is called *nonzero-sum* and it requires a new set of rules for estimating optimal strategies and solutions. As each player's gains and losses are not directly related to the opponents', the optimal solution is only based on the assertion that it should be the one that ensures that the player has "no regrets" when choosing between possible decision options. This essentially means that, since each player is now interested in his/her own gains and losses, the optimal solution should only focus on maximizing each player's own *expectations* [33, 28, 13]. The Minimax property can still be applied in principle when the single most "secure" option must be identified, but now the solution of the game gains a new meaning.

During the early 1950's, John Nash has focused primarily on the problem of finding a set of *equilibrium points* in nonzero-sum games, where the players eventually settle after a series of competitive rounds of the game [29, 30]. The failure of the Minimax approach to predict real-world outcomes in nonzero-sum games comes from the fact that the players are assumed to act independently and simultaneously, while in reality they usually are not. Experience shows that *possibly* better payoffs with what a player *might* choose, after observing the opponent's moves, is a very strong motivator when choosing its actual strategy [27]. In strict mathematical terms, these equilibrium points would not be the same in essence with the Minimax solutions, as they would come as a result of the players' competitive behavior over several "turns" of moves and not as an algebraic solution of the mathematical formulation in a single-turn game.

In 1957 Nash has successfully proved that indeed such equilibrium points exist in all nonzero-sum games, in a way that is analogous to the Minimax theorem assertion. This new type of stable outcome is referred to as *Nash equilibrium* after his name and can be considered a generalization of the corresponding Minimax equilibrium in zero-sum games. In essence, they are the manifestation of the *no regrets* principle for all players, i.e., not regretting their final choice after observing their opponents' behavior [44, 46]. However, although the Nash theorem ensures that at least one such Nash equilibrium exists in all nonzero-sum games, there is no clear indication on how the game's solution can be analytically calculated at this point. In other words, although a solution is known to exist, there is no closed form for nonzero-sum games until today. Seminal works by C. Daskalakis & Ch. Papadimitriou in 2006-2007 and on have proved that, while Nash equilibria exist, they may be unattainable and/or practically impossible to calculate due to the inherent algorithmic complexity of this problem, e.g. see: [12, 34].

It should be noted that players participating in a nonzero-sum game may or may not have the same options available as alternative course of action, or the same set of options may lead to different gains or *payoffs* between the players. When players are fully interchangeable and their ordering in the game makes not difference to the game setup and its solutions, the game is called *symmetrical*. Otherwise, if exchanging players' position does not yield a proportional exchange of their payoffs, then the game is called *asymmetrical*. Naturally, symmetrical games lead to Nash equilibrium points that appear in pairs, as an exchange between players creates its symmetrical counterpart.

Summary:

- In *nonzero-sum* games, the payoffs of the players are separated (although may be correlated).
- If players are allowed to observe their opponents moves over several iterations, then the "no regrets" principle is a strong incentive to revise their own strategies, even though their payoffs are separated.
- The *Nash equilibrium theorem* ensures that, under these conditions, there are indeed stable solutions in nonzero-sum games, similarly to the Minimax theorem for zero-sum games.
- However, calculating the optimal strategies and the game solution for these Nash equilibria is a vastly more complex and generally unfeasible task.

2 Cooperation instead of competitiveness

The seminal work of Nash and others in nonzero-sum games was a breakthrough in understanding the outcome in real-world adversarial situations. However, the Nash equilibrium points are not always the globally optimal option for the players. In fact, the Nash equilibrium is optimal only when players are strictly competitive, i.e., when there is no chance for a mutually agreed solution that benefits them more. These strictly competitive forms of games are called *noncooperative* games. The alternative option, the one that allows communication and prior arrangements between the players, is called a *cooperative* game and it is generally a much more complicated form of nonzero-sum gaming. Naturally, there is no option of having cooperative zero-sum games, since the game structure itself prohibits any other settlement between the players other than the Minimax solution.

2.1 The cooperative option

The problem of cooperative or possibly cooperative gaming is the most common form of conflict in real life situations. Since nonzero-sum games have at least one equilibrium point when studied under the strictly competitive form, Nash has extensively studied the cooperative option as an extension to it. However, the possibility of finding and mutually adopting a solution that is better for both players than the one suggested by the Nash equilibrium, essentially involves a set of behavioral rules regarding the players' stance and "mental" state, rather than strict optimality procedures [27]. Nash named this process a *bargain* between the players, trying to mutually agree on one solution between multiple candidates within a *bargaining set* or *negotiation set*. In practice, each player should enter a bargaining procedure if and only if there is a chance that a cooperative solution exists and it provides at least the same gain as the best strictly competitive solution, i.e., the best Nash equilibrium. In this case, if such a solution is agreed between the players, it is called *bargaining solution* of the game [28, 33].

As mentioned earlier, each player acts upon the property of no regrets, i.e., follow the decisions that maximize their own expectations. Nevertheless, the game setup itself provides means of improving the final gain in an agreed solution. In some cases, the bargaining process may involve the option of *threats*, that is a player may express the intention to follow a strategy that is particularly costly for the opponent. Of course, the opponent can do the same, focusing on a similar threat. This procedure is still a cooperative bargaining process, with the threshold of expectations raised for both players. The result of such a process may be a mutually *deterring* solution, which in this case is called a *threatening solution* or *threat equilibrium*. There is also evidence that, while cooperative strategies do exist, in some cases "cooperation" may be the result of *extortion* between players with unbalanced power and choices [36].

In his work, Nash has formulated a general and fairly logical set of six axioms, the *Nash's bargaining axioms*, regarding the behavior of *rational* players, in order to establish a non-empty bargaining set, i.e., to have at least one stable solution (equilibrium) [28, 33, 29]. In non-strict form, these axioms can be summarized in the following propositions:

- Any of the cooperative options under consideration must be feasible and yield at least the same payoff as the best strictly non-cooperative option for all players, i.e., cooperation must be mutually beneficial.
- Strict (mathematical) constraints: Pareto optimality, independence of irrelevant alternatives, invariance under linear transformations, symmetry [46, 33, 28].

The first proposition essentially defines the term "rationality" for a player: he/she always acts with the goal of maximizing own gains and minimizing losses, regardless if this means strictly competitive or possibly cooperative behavior. The second proposition names a set of strict mathematical preconditions (not always satisfied in practice), in order for such a bargaining set to exist. Having settled on these axioms, Nash was able to prove the corresponding *bargaining theorem*: under these axioms, there exists such a *bargaining process*, it is unique and it leads to a bargaining solution, i.e., equilibrium. However, as in the general case of strictly competitive games, Nash's bargaining theorem does not provide analytical means of finding such solutions.

The notion of bargaining sets and threat equilibrium is often extended in special forms of games that include iterative or recursive steps in gaming, either in the form of multi-step analysis (*meta-games*) or focusing on the transitional aspects of the game (*differential games*). Modern research is focused on methods that introduce probabilistic models into games of multiple realizations and/or multiple stages [33].

Summary:

- In nonzero-sum games, there may be non-competitive (cooperative) options that are mutually beneficial to all players.
- Under some general rationality principles, *Nash's bargaining theorem* ensures that these cooperative outcomes may indeed become the game solution, provided that strict competitiveness yields lower gains for all.
- The procedure of structuring the "common ground" of cooperation between the players, normally conducted over several iterations, is the bargaining process.

2.2 Coalitions, stable sets, the Core

Nash's work on the Nash equilibrium and bargaining theorem provides the necessary means to study n-person non-cooperative and cooperative games under a unifying point of view. Specifically, a nonzero-sum game can be realized as a strictly competitive or a possibly cooperative form, according to the game's rules and restrictions. Therefore, the cooperative option can be viewed as a generalization to the strictly competitive mode of gaming.

When players are allowed to cooperate in order to agree on a mutually beneficial solution of game, they essentially choose one strategy over the others and bargain this option with all the others in order to come to an agreement. For symmetrical games, this is like each player chooses to join a group of other players with similar preference over their initial choice. Each of these groups is called a *coalition* and it constitutes the basic module in this new type of gaming: the members of each coalition act as cooperative players joined together and at the same time each coalition competes over the others in order to impose its own position and become the *winning coalition*. This setup is very common when modeling voting schemes, where the group that captures the relative majority of the votes becomes the winner.

Coalition Theory is closely related to the classical Game Theory, especially the cooperating gaming [33, 28]. In essence, each player still tries to maximize its own expectations, not individually any more but instead as part of a greater opposing term. Therefore, the individual gains and capabilities of each player is now considered in close relation to the coalition this player belongs, as well as how its individual decision to join or leave a coalition affects this coalition's winning position. As in classic nonzero-sum games, the notion of equilibrium points and solutions is considered under the scope of domination or not in the game at hand. Furthermore, the theoretical implications of having competing coalitions of cooperative players is purely combinatorial in nature, thus making its analysis very complex and cumbersome. There are also special cases of collective decision schemes where a single player is allowed to *abstain* completely from the voting procedure, or prohibit a contrary outcome of the group via a *veto* option.

In order to study the properties of a single player participating in a game of coalitions, it is necessary to analyze the wining conditions of each coalition. Usually each player is assigned a fixed value of "importance" or "weight" when participating in this type of games and each coalition's power is measured as a sum over the individual weights of all players participating in this coalition. The coalition that ends up with the highest cumulative value of power is the winning coalition. Therefore, it is clear that, while each player's power is related to its individual weight, this relation is *not* directly mapped on how the participation in any arbitrary coalition may affect this coalition's winning or losing position. As this process stands true for all possible coalitions that can be formed, this competitive type of "claiming" over the available pool of players/voters by each coalition suggests that there are indeed configurations that marginally favor the one or the other coalition, i.e., a set of "solutions".

The notion of solution in coalition games is somewhat different from the one suggested for typical nonzero-sum games, as it identifies minimal settings for coalitions that dominate all the others. In other words, they do not identify points of maximal gain for a player or even a coalition, but equilibrium "points" that determine which of the forming coalitions is the winning one. This type of "solutions" in coalition games is defined in close relation to *domination* and *stability* of such points and they are often referred to as *the Core*. Von Neumann and Morgenstern have defined a somewhat more relaxed definition of such conditions and the corresponding solutions are called *stable sets* [33, 28]. It should be noted that, in contrast to Nash's theorems and the Minimax assertion of solutions, there is generally no guarantee that solutions in the context of the Core and stable sets need to exist in an arbitrary coalition game.

Summary:

- Players of similar preferences and mutual benefits may join in groups or *coalitions*; these coalitions may be competing with each other, similarly to competitive games between single players.
- The study of games between coalitions is inherently more complex than with single players, as in this case every player contributes to the collective "power" and enjoys a share of the wins.
- In general, coalitions are formed and structured under the scope of *voting ensembles*, where the voting weight of each individual player contributes to the combined weight of the coalition.

2.3 Indices of power in committees

The notion of the Core and stable sets in coalition gaming is of vital importance when trying to identify the winning conditions and the relative power of each individual player in affecting the outcome of the game. The observation that a player's weight in a weighted system may not intuitively correspond to its voting "power" goes back at least to Shapley and Shubik (1954). For example, a specific weight distribution to the players may make them relatively equivalent in terms of voting power, while only a slight variation of the weights may render some of them completely irrelevant on determining the winning coalition [45].

Shapley and Shubik (1954) and later Banzhaf and Coleman (1965, 1971) suggested a set of well-defined equations for calculating the relative power of each player, as well as each forming coalitions as a whole [33, 28]. The *Shapley-Shubik index of power* is based on the calculation of the actual contribution of each player entering a coalition, in terms of improving a coalition's gain and winning position. Similarly, the *Banzhaf-Coleman index of power* calculates how an individual player's decision to join or leave a coalition ("swing vote") results in a winning or loosing position for this coalition, accordingly. Both indexes are basically means of translating each player's individual importance or weight within the coalition game into a quantitative measure of power in terms of determining the winner. While both indices include combinatorial realizations, the Banzhaf index is usually easier to calculate, as it is based on the sum of "shifts" on the winning condition a player can incur [5]. Furthermore, its importance in coalition games is made clearer when the Banzhaf index is viewed as the direct result of calculating the derivatives of a *weighted majority game* (WMG).

Seminal work by L. S. Penrose [35], as well as more recent studies with computer simulations [8], have shown that this discrepancy between voting weights and actual voting power is clearly evident when there is large variance in the weighting profile and/or when the voting group has less than 12-15 members. Even in large voting pools, the task of designing optimal voting mechanisms and protocols with regard to some *collective efficiency* criterion is one of the most challenging topics in Decision Theory.

Summary:

- *Weighted majority games* (WMG) are the typical theoretical structures of the process of formulating the collective decision within a coalition.
- In voting ensembles, each player's voting weight is *not* directly proportional to his/her true *voting power* within the group, i.e., the level of steering the collective decision towards its own choices.

2.4 Voting ensembles and majority winners

In most cases, majority functions that are employed in practice very simplistic when it comes to weighting distribution profile or they imply a completely uniform weight distribution. However, a specific weighting profile usually produces better results, provided that is simple enough to be applied in practice and attain a consensus in accepting it as "fair" by the voters. Taylor and Zwicker [45] have defined a voting system as *trade robust* if an arbitrary series of trades among several winning coalitions can never simultaneously render them losing. Furthermore, they proved that a voting system is trade robust if and only if it is weighted. This means that, if appropriate weights are applied, at least one winning coalition can benefit from this procedure.

As an example, institutional policies usually apply a non-uniform voting scheme when it comes to collective board decisions. This is often referred to as the “inner cabinet rule”. In a hospital, senior staff members may attain increased voting power or the chairman may hold the right of a tie-breaking vote. It has been proven both in theory and in practice that such schemes are more efficient than simple majority rules or any restricted versions of them like trimmed means. Nitzan and Paroush [32] have studied the problem of optimal *weighted majority rules* (WMR) extensively and they have proved that they are indeed the optimal decision rules for a group of decision makers in dichotomous choice situations. This proof was later extended by Ben-Yashar and Paroush, from dichotomous to polychotomous choice situations [3]; hence, the optimality of the WMR formulation has been proven theoretically for any n-label voting task.

The weight optimization procedure has been applied experimentally in trained or other types of combination rules, but analytical solutions for the weights is not commonly used. However, Shapley and Grofman [42] have established that an analytical solution for the weighting profile exists and it is indeed related to the individual player skill levels or *competencies* [23]. Specifically, if decision independence is assumed for the participating players, the optimal weights in a WMR scheme can be calculated as the log-odds of their respective skill probabilities, i.e.:

$$w_k = \log\left(O_k\right) = \log\left(\frac{p_k}{1-p_k}\right) \tag{2.1}$$

where p_k is the competency of player k and w_k is its corresponding voting weight. Interestingly enough, this is exactly the solution found by analytical Bayesian-based approaches in the context of decision fusion of independent experts in Machine Learning [24]. The optimality assertion regarding the WMR, together with an analytical solution for the optimal weighting profile, provides an extremely powerful tool for designing theoretically optimal collective decision rules. Even when the independence assumption is only partially satisfied in practice, studies have proved that WMR-based models employing log-odds weighting profiles for combining pattern classifiers confirm these theoretical results [19, 18].

Summary:

- *Weighted majority rules* (WMR) have been proven theoretically as the optimal decision-making structures in weighted majority games.
- The *log-odds* model has been proven both as the theoretically optimal way to weight the individual player’s votes, provided that they decide independently.
- The optimality of the log-odds weighting method has also been proven experimentally, even when the independence assumption is only partially satisfied.

2.5 Collective efficiency

Condorcet (1785) [9] was the first to address the problem of how to design and evaluate an efficient voting system, in terms of fairness among the people that participating in the voting process, as well as the optimal outcome for the winner(s). This first attempt to create a probabilistic model of a voting body is known today as the *Condorcet Jury Theorem* [51]. In essence, this theorem says that if each of the voting individuals is somewhat more likely than not to make the "better" choice from a set of alternative options; and if each individual makes its own choice *independently* from all the others, then the probability that the group majority is "correct" is greater than the individual probabilities of the voters. Moreover, this probability of correct choice by the group increases as the number of independent voters increases. In practice, this means that if each voter decides independently and performs marginally higher than 50%, then a group of such voters is *guaranteed* to perform better than each of the participating individuals. This assertion has been used in Social sciences for decades as a proof that decentralized decision making, like in a group of juries in a court, performs better than centralized expertise, i.e., a sole judge. The Condorcet Jury Theorem and its implications have been used as one guideline for estimating the efficiency of any voting system and decision making in general [51]. Under this context, the coalition games are studied by applying quantitative measures on *collective competence* and *optimal distribution of power* in the ensemble, e.g. tools like the Banzhaf or Shapley indices of power. The degree of consistency of such a voting scheme on establishing the pair-wise winner(s), as the Condorcet Jury Theorem indicates, is often referred to as the *Condorcet criterion*.

Shapley-Shubik and Banzhaf-Coleman are only two of several formulations for the indices of power in voting ensembles, each defining different payoff distributions or realizations among the members of winning coalitions. In general, these formulations are collectively referred to as *semivalue functions* or *semivalues* and they are considered more or less equivalent in principle, although may be different in exact values. Almost all of them are based on combinatorial functions (inclusion-exclusion operations in subsets) and, as a result, there is no easy way to formulate proper inverse functions that can be calculated in polynomial time. Therefore, the design of exact voting profiles with weights based on semivalues, instead of competencies as described above (log-odds), is generally impractical even for ensembles of small sizes.

For further insight on weighted majority games, weighted majority voting, collective decision efficiency and Condorcet efficiency, as well as applications to Machine Learning for designing pattern classifiers, see [17, 19, 18].

Tab. 1: Generic 2x2 zero-sum game in analytical form.

Game example		Player-2 y	$1-y$
Player-1	x	a	b
	$1-x$	c	d

Summary:

- Under the assumption of independent voters and that each decides "correctly" marginally higher than 50% of the time, then their collective decision as a group is theoretically proven to be asymptotically better any single member of the ensemble.
- Furthermore, as the size of the ensemble increases, its collective competency is guaranteed to increase too.
- In the other hand, the problem of formulating an analytical solution for the optimal distribution of voting power within such a group, i.e., the design of theoretically optimal *voting mechanisms*, is still an open research topic.

3 Game analysis & solution concepts

One of the most important factors in understanding and analyzing games correctly is the way they are represented. Games can be represented and analyzed in two generic formulations: (a) the *analytical* or *normal* form, where each player is manifested as one dimension and its available choices (strategies) as offsets on it, and (b) the *extensive* or *tree-graph* form, where each player's "move" correspond to a node split in a tree representation. Each one of them has its own advantages and disadvantages, but theoretically they are equivalent.

3.1 Games in analytical (matrix) form

In Table 1, an example of a zero-sum game in analytical form is presented. Player-1 is usually referred to as the "max" player and Player-2 is referred to as the "min" player, while rows and columns correspond to each player's available strategies, respectively. Since this is a zero-sum game and one player's gains is the other player's losses, the "max" player tries to maximize the game value (outcome) while the "min" player tries to minimize it. In the context of the Minimax theorem, Player-2 chooses the *maximum-of-minimums*, while Player-2 chooses the *minimum-of-maximums*. The x and y correspond to the weight or probability of choosing the first strategy and, since this is a 2x2 game, the other strategies are attributed with the complementary probabilities, 1-x and 1-y.

The exact Minimax solution for x and y depends solely on the values of the individual payoffs for each of the four outcomes. Here, it is assumed that there is no *domination* in strategies, i.e., there is no row/column that is strictly "better"

Tab. 2: Example 2x2 zero-sum game in analytical form.

Game example		Player-2 (0)	(1)
Player-1	(0)	0	-3
	(1)	4	**1**

Tab. 3: Example of a 2x2 nonzero-sum game in analytical form.

Game example		Player-2 y	$1-y$
Player-1	x	(a_1,a_2)	(b_1,b_2)
	$1-x$	(c_1,c_2)	(d_1,d_2)

than another row/column (column-wise/row-wise, respectively, all payoffs). For example, Player-1 would have a dominating strategy in the first row if and only if $a \geq c$ and $b \geq d$. Based on this generic setup, this is a typical 2x2 system of linear equations and, if no domination is present, its solution can be determined analytically as [44, 14, 26]:

$$[x, 1-x] = \left[\frac{d-c}{a-b-c+d}, \frac{a-b}{a-b-c+d}\right] \tag{3.1}$$

$$[y, 1-y] = \left[\frac{d-b}{a-b-c+d}, \frac{a-c}{a-b-c+d}\right] \tag{3.2}$$

$$u = \frac{ad-bc}{a-b-c+d} \tag{3.3}$$

The Minimax solution $[x, y]$ determines the saddle-point, i.e., the equilibrium that is reached when both opponents play optimally in the Minimax sense, when the game has no pure (non-mixed) solution. In this case, the expected payoff or *value* of the game for both players is calculated by u (remember, this is a zero-sum game). If the game has a pure solution, then it is determined as either 0 or 1 for each probability x and y. Table 2 illustrates a zero-sum game and the corresponding pure Minimax solution, by selecting the appropriate strategies for each player. In this case, "max" Player-1 chooses the the maximum {1} between the two minimum values {-3,1} from its own two possible worst-case outcomes, while "min" Player-2 chooses the the minimum {1} between the two maximum values {4,1} from its own two possible worst-case outcomes. Hence, the pure solution [1,1] is the Minimax outcome.

In nonzero-sum games, the analytical form is still a matrix, but now the payoffs for each player are separate, as illustrated in Table 3. Here, since the payoffs are separated, both players are treated as "max" and the Minimax solution for each one is calculated by selecting the maximum-of-minimums as described before for zero-sum games, focused solely on its own payoffs from each value pair.

Although a (pure) Minimax solution can always be calculated for nonzero-sum games, the exact Nash equilibrium solution is a non-trivial task that cannot be solved analytically in the general case. However, pure Nash equilibrium outcomes can be identified by locating any payoff pairs (z, w) such that z is

Tab. 4: Example of a 2x2 nonzero-sum game with one Nash equilibrium at [A,B]:(2,4).

Game example		Player-2 A	Player-2 B
Player-1	A	(3,3)	(2*,4*)
Player-1	B	(4*,1)	(1,2*)

the maximum of its column and w is the maximum of its row. In other words, every row for Player-1 is scanned and every entry in it is compared to the values in the same column, marking it if it is the maximum among them; the same process is conducted for every column for Player-2, scanning each value row-wise for its maximum; any payoff pair that has both values marked as maximums is a Nash equilibrium in the game. Table 4 illustrates such an example, where asterisk (*) marks the identified max-values and the single Nash equilibrium for [A,B] at (2,4). Here, although the strategies are the same for both players, their (separated) payoffs are not, hence the game is referred to as *asymmetric*. According to the *oddness theorem* by Wilson (1971), the Nash equilibria almost always appear in odd numbers [44, 33], at least for *non-degenerate* games, where mixed strategies are calculated upon k linearly independent pure strategies.

Summary:

- Game representation in *analytical* form introduces a game matrix, with row and column positions associated to the strategies available to the players and contents associated to the corresponding payoffs.
- Analytical-form representation introduces very convenient ways to identify Minimax solutions and Nash equilibria in games.
- However, they are appropriate mostly for 2-player simultaneous games, since any other configuration cannot be fully illustrated.

3.2 Games in extensive (tree-graph) form

In the extensive form the game is represented as a tree-graph, where each node is a state labeled by a player's number and each (directed) edge is a player's choice or "move". Strictly speaking, this is a form of state-transition diagram that illustrates how the game evolves as the players choose their strategies. Figure 3.1 shows such a 2x2 nonzero-sum game of perfect information, while Figure 3.2 shows a similar 2x2 game of imperfect information [46, 28, 49, 41, 16, 14]. Nodes with numbers indicate players, edges with letters indicate chosen strategies (here, symmetric) and separated payoffs (in parentheses) indicate the game outcome after one full round. The dashed line between the two nodes for Player-2 indicate that its current true state is not clearly defined due to imperfect information regarding Player-1's move. In practice, these two states form an *information set* for Player-2, which has no additional information to differentiate between them. This is also valid in the case of simultaneous moves,

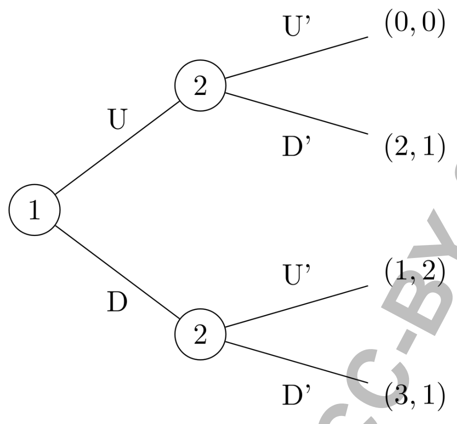

Fig. 3.1: Example of a 2x2 nonzero-sum game of *perfect* information.

where Player-2 cannot observe Player-1's move in advance of its own, and vice versa. In extensive form, an information set is indicated by a dotted line or by a loop, connecting all nodes in that set.

The extensive form of game is usually the preferred way to represent the tree-graph of simple 2-player board games, where each node is clearly a state and each edge is a player's move. Even in single-player games, where a puzzle has to be solved through a series of moves (e.g. Rubik's cube)[2], the tree-graph is a very effective way to organize the game under an algorithmic perspective, in order to program a "solver" in a computer. In practice, the problem is structured as sequences of states and transitions in a tree-graph manner and the "game" is explored as it is evolving, move after move, expanding the tree-graph from every terminal node. The tree-graph can be expanded either by full a level ("breadth-first"), or from a branch all the way down to non-expandable terminal nodes ("depth-first"), or some hybrid scheme between these two alternatives.

As described above, small games like Checkers can be structured and expanded fully, with their tree-graph having only internal (already expanded) and terminal nodes; however, in larger games like Chess or Go this is practically unfeasible even with super-computers. In such cases, the algorithm should assess the "optimality" of each expandable terminal node with regard to relevance towards the predefined goal ("win" or "solution"), sort all these nodes according to their ranking and choose the "best" ones for expansion in the next iteration. This way, the search is sub-optimal but totally feasible with almost any memory constraints - this is exactly how most computer players are programmed for playing board games or solving complex puzzle games. In Artificial Intelligence, algorithms like $A*$ and AB solve this type of problems as a path-finding optimization procedure towards a specified goal [40, 31].

[2] The combinatorial analysis of the classic 3x3x6 Rubik's cube should take into account tile permutations that can only be reached by the available shifts and turns of the slices of the device. Therefore, a totally "free" permutation scheme would produce: $8! \cdot 3^8 \cdot 12! \cdot 2^{12} = 519,024,039,293,878,272,000$ cube instances, while in practice the possible permutations are only: $8! \cdot 3^7 \cdot (12!/2) \cdot 2^{11} = 43,252,003,274,489,856,000$ cube instances (about 12 times fewer) [50].

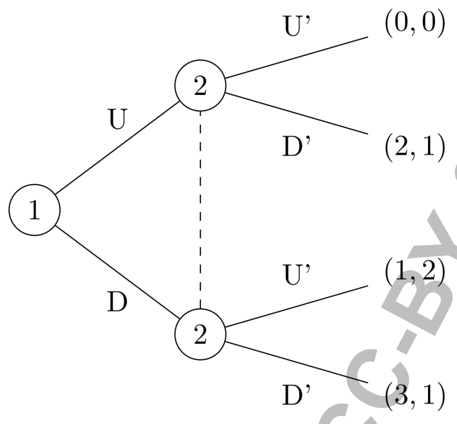

Fig. 3.2: Example of a 2x2 nonzero-sum game of *imperfect* information.

Figure 3.3 illustrates the way a path-finding algorithm like A* would work in expanding a tree-graph as described above. The "root" node is the starting state in a puzzle game (single-player) and each node represents a new state after a valid move. The numbers indicate the sequence in which the nodes are expanded, according to some optimality-ranking function (not relevant here). For example, node "4" in the 3rd level is expanded before node "5" in the 2nd level, node "21" in the 5th level is expanded before node "22" in the 3rd level, etc. Here, node "30" in the 5th level is the last and most relevant terminal node (still expandable) towards the goal, hence the optimal path from the "root" state is currently the: "5"→"7"→"11"→"30" and the next "best" single-step move is the one towards "5". The tree-graph can be expanded in an arbitrary number of levels according to the current memory constraints for the program, but the same path-finding procedure has to be reset and re-applied after the realization of each step when two or more players are involved, since every response from the opponent effectively nullifies every other branch of the tree-graph.

It should be mentioned that, although the extensive form of game representation is often inefficient for large games like Chess, it can be used as a tool in the proof of the existence of an optimal solution [15, 46]. Specifically, in every such game of complete and perfect information (all board games), each player knows its exact position in the graph-tree prior to choosing the next move. In other words, each player is not only aware of the complete structure of the game but also knows all the past moves of the game, including the ones of random choice. Hence, since there is no uncertainty in the moves, each player can remove the dominated strategies and subsequently identify the optimal choice, which is always a pure strategy, i.e., the one that corresponds to the saddle-point of the game. This proof actually ensures the existence of a (pure) optimal strategy in every typical board game, no matter how large or complex it is. Examples include Tic-Tac-Toe, Chess, Backgammon, etc.

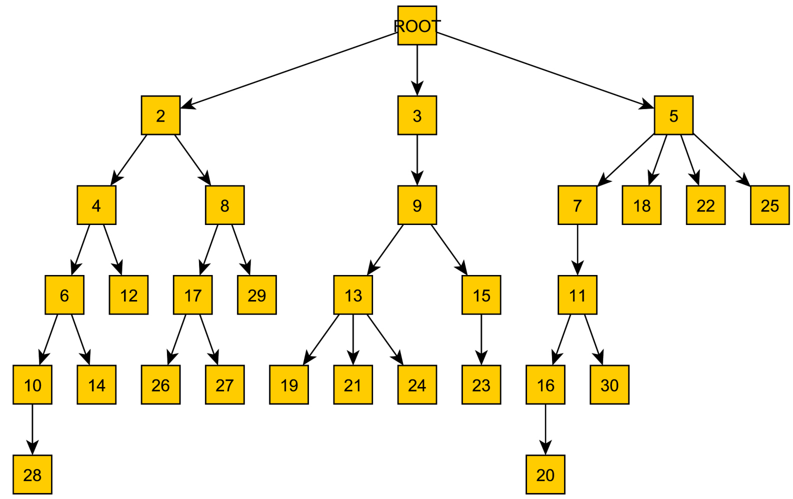

Fig. 3.3: Example of the way a path-finding algorithm like A* would work in expanding the tree-graph of a single-player "puzzle" game like Rubik's cube.

Summary:

- Game representation in *extended* form introduces a tree-graph, with nodes associated to individual players and (directed) edges associated to selected strategies ("moves").
- Extended-form representation introduces very convenient ways to identify chains of moves and solution paths.
- However, the calculation of Minimax solutions and Nash equilibria is not straight-forward.

4 The four interesting cases

In the real world, games may be either zero-sum or nonzero-sum by nature. As described previously, the case of zero-sum games can be considered simpler and much easier to solve analytically, since it can be formulated as a typical algebraic set of linear equations that define the Minimax solution, regardless if it contains pure or mixed strategies [44, 14]. However, nonzero-sum games are inherently much more complex and require non-trivial solution approaches, usually via some Linear Programming (constraint) optimization procedure, e.g. see: [20, 43]. In fact, it has been proven that the general task of finding the Nash equilibria is algorithmically intractable[3] [12, 10, 11, 34] - something that

[3] In their seminal works, Daskalakis, Goldberg and Papadimitriou have shown that the task of finding a Nash equilibrium is PPAD-complete; informally, PPAD is the class of all search

Tab. 5: The general analytical (matrix) form of a 2x2 nonzero-sum symmetric game.

Game template		Player-2 C	Player-2 D
Player-1	C	(R,R)	(S,T)
	D	(T,S)	(P,P)

puts into a "philosophical" question the very nature and practical usefulness of having proof of game solutions (i.e., stable outcomes) that we may not be able to calculate.

Some cases of nonzero-sum games are particularly interesting, especially when they involve symmetric configurations. The players can switch places, the actual payoff values are usually of much less importance than their relative ordering as a simple preference list, the Minimax and Nash equilibria can be easily identified, yet these simple games seem to capture the very essence of bargaining and strategic play in a vast set of real-world conflict situations with no trivial outcomes.

Table 5 shows a generic template for such very simple symmetric nonzero-sum games, employing only two strategies and four payoff values to completely define such games in analytical (matrix) form. Here, the game is symmetric because the players can switch roles without any effect in their corresponding payoff pairs. Furthermore, they share two common strategies C and D, named typically after the choices of "cooperate" or "defect", while constants P, R, S and T are the real-valued payoffs in each case [7].

In practice, a player's preference of strategies (and hence, the equilibria) depends only on the relative ordering of the corresponding payoffs and not their exact values, which become of real importance only when the actual payoff value of the game solution is to be calculated for each player. There is a finite number of rank combinations, i.e., permutations, of these four constants, which produce all the possible unique game matrices of this type. Specifically, there are $4! = 24$ different ways to order these four numbers, 12 of which can be discarded as qualitatively equivalent to other game configurations. Out of the 12 remaining games, eight of them possess optimal pure strategies for both players, therefore they can be considered trivial in terms of calculating their solution. The four remaining configurations are the most interesting ones, as they do not possess any optimal pure strategy. These are the following:

- *Leader*: $T > S > R > P$.
- *Battle of the Sexes*: $S > T > R > P$.
- *Chicken*: $T > R > S > P$.
- *Prisoner's Dilemma*: $T > R > P > S$.

These four qualitatively unique games seem to capture the essence of most of the majority real-world conflict situations historically. Although they have

problems which always have a solution and whose proof is based on the parity argument for directed graphs. Due to the proof of intractability, the existence of Nash equilibrium in all nonzero-sum games somewhat loses its credibility as a predictor of behavior.

Tab. 6: The typical setup of the *Leader* game with two players. Nash equilibria are marked with paired asterisks and the Minimax solution with bold numbers.

Leader game		Player-2 *C*	*D*
Player-1	*C*	(**2**,**2**)	(3*,4*)
	D	(4*,3*)	(1,1)

been studied extensively in the past, there are still many open research topics regarding the feasibility, tractability and stability of the theoretical solutions.

4.1 Leader

The *Leader* or *Coordination* game [28, 33, 46, 44, 7, 16] is named after the typical problem of two drivers attempting to enter a stream of increased traffic from opposite sides of an intersection. When the road is clear, each driver has to decide whether to move in immediately or concede and wait for the other driver to move first. If both drivers move in (i.e., choose D), they risk crashing onto each other, while if they both wait (i.e., choose C), they will waste time and possibly the opportunity to enter the traffic. The former case is the worst, hence the payoff of (1,1), while the later case is slightly more preferable with a payoff of (2,2). The best outcome is for one driver to become the “leader” and move first, while the other becomes the “follower” and move second. There is still some difference in their absolute gains, but now the deadlock is resolved in the best possible way, no matter who is actually the leader and who is the follower.

Table 6 illustrates the analytical form of this game setup, where numbers indicate relative preferences rather than absolute gain values. There are two pure Nash equilibria, (3,4) and (4,3), which correspond to the proper assignment of roles to the players, explicitly or implicitly, such that coordination is achieved. Since the game is symmetric the two players can switch roles, with only marginal increase/decrease to their payoffs. In terms of Minimax strategies, each player is free to choose the strategy that guarantees the maximum-of-the-minimums without any concern about the opponent's payoffs, since this is a nonzero-sum game. Hence, the Minimax solution is $[C,C]$ at (2,2) marked in bold.

In the real world, the assignment of roles as leader/follower is more effective when applied explicitly, typically by some external mechanism or a predefined set of rules. Street signs, traffic policemen and highway code for driving properly are all such mechanisms for explicit resolution of deadlocks via priority assignment in traffic.

4.2 Battle of the Sexes

In the *Battle of the Sexes* game [28, 33, 46, 7, 16], a married couple has to decide between entertainment options for the evening. The husband prefers one choice, while the wife prefers another. The problem is that they would both prefer to concede to the same choice together even if it is not their own, rather than follow their own choices alone. For example, of he wants to watch a sports match on

Tab. 7: The typical setup of the *Battle of the Sexes* game with two players. Nash equilibria are marked with paired asterisks and the Minimax solution with bold numbers.

Battle of the Sexes		Player-2 *C*	Player-2 *D*
Player-1	*C*	(1,1)	(3*,4*)
	D	(4*,3*)	(**2**,**2**)

TV and she wants to go out for dinner, they both prefer either watching TV or going out for dinner as long as they are together.

Table 7 illustrates the analytical form of the game, where strategy C is for conceding to the other's preference and D is for defecting to his/her own choice. If they both concede the payoff (1,1) is the worst outcome, since they both end up miserable and bored. If they both defect the payoff (2,2) is marginally better for both, but they end up being alone. The two other cases of someone following the other yields the best payoffs for both, since the game is symmetric and they can switch places. The outcomes (3,4) and (4,3) are actually the two Nash equilibria, similarly to the *Leader* game; however, the Minimax solution (2,2) here corresponds to both players choosing D (not C as in *Leader*) as their best Minimax strategy.

4.3 Chicken

One of the most well-known strategic games is *Chicken* [15, 26, 28, 33, 46, 7], dating back at least as far as the Homeric era. Two or more adversaries engage in a very dangerous or even lethal confrontation, each having a set of choices at his/her disposal and each of these choices producing more or less damage to all players if their choice is the same. Typically, this translates to the Hollywood's favorite version of two cars speeding towards each other, the drivers can choose to turn and avoid collision or keep the course and risk death if the other driver do not turn either. The game seems simple enough, but there are several theoretical implications that make it one of the most challenging situations, appearing in many real-world conflicts throughout History.

Table 8 illustrates the typical *Chicken* game setup with two players and two strategic choices. Option C corresponds to turning away ("swerve") and losing the game, while option D corresponds to keeping the course and risk death. The worst possible outcome is at (1,1) when players persist in keeping course and eventually crashing against each other. The mutually beneficial outcome or "draw" is at (3,3) when both players decide to play safe and turn away; this is actually the Minimax solution of the game, i.e., the most conservative and "rational" outcome if the game is a one-off round. On the other hand, there are two Nash equilibria for the two outcomes when only one player turns away and one persists.

One particularly interesting feature of the *Chicken* game is that it is impossible to avoid playing it with some insistent adversary, since refusing to play is effectively equivalent to choosing C (swerve). Furthermore, the player who succeeds in making his/her commitment to D adequately convincing is always the one that can win at the expense of the other player, assuming that the other

Tab. 8: The typical setup of the *Chicken* game with two players. Nash equilibria are marked with paired asterisks and the Minimax solution with bold numbers.

Chicken game		Player-2 *C*	Player-2 *D*
Player-1	*C*	(**3**,**3**)	(2*,4*)
Player-1	*D*	(4*,2*)	(1,1)

player is rational and would inevitably decide to avoid disaster. In other words, the player that is somehow bounded to avoid losing *at any cost* and makes this commitment very clear to the opponent, is the one that will always win against any rational player.

This aspect of credible commitment is closely related to the notion of *reputation*, as well as the strange conclusion that in this game the most effectively "rational" strategy is the manifestation of "irrational" commitment to lethal risk. This becomes especially relevant in cases where the game is played a number of times repeatedly and previous behaviors directly affect the players' strategic choices in the future: once the risky player starts winning he/she may maintain or even improve this advantage, as confidence and prior "risky" behavior makes it more and more difficult for future opponents to decide and deviate from their cautious Minimax choice of swerving. The *Chicken* game is perhaps the most descriptive and simple case where players' previous behavior (i.e., *reputation*) is of such importance for predicting the actual outcome.

4.4 Prisoner's Dilemma

This forth basic type of non-trivial, nonzero-sum game is by far the most interesting one. The *Prisoner's Dilemma* game [15, 26, 28, 33, 46, 44, 7, 16] typically involves two prisoners who are accused of a crime. Each of them has the option of remaining silent and withholding any information or confessing to the police and accusing the other by disclosing details about the crime. The first choice C is effectively the cooperative option, while the second choice D corresponds to purely competitive behavior in order to reduce he/her own damages.

Table 9 illustrates the typical *Prisoner's Dilemma* game setup with two players and two strategic choices. The payoffs here correspond simply to preferences and not real gain/cost values, but the essence and the strategic properties of the game remain intact. In practice, what the game matrix says is that if the two prisoner's remain silent, i.e., mutually cooperate, they will not be freed but they will share an equal, relatively mild conviction. If they both talk and accuse each other, i.e., mutually defect, they will share and equal but more severe conviction. If only one of them talks to the police and the other remains silent, the one that talked is freed and the other serves a full-time conviction for both. It is of course imperative that the two prisoners are immediately separated upon capture and no communication between them is allowed; this does not nullifies any prior arrangements they may have, but isolation after being captured means that neither of them can confirm they loyalty of the other. This is one of the main reasons why police always isolates suspects prior and during any similar investigation.

Tab. 9: The typical setup of the *Prisoner's Dilemma* game with two players. Nash equilibria are marked with paired asterisks and the Minimax solution with bold numbers.

Prisoner's Dilemma		Player-2 *C*	Player-2 *D*
Player-1	*C*	(3,3)	(1,4*)
Player-1	*D*	(4*,1)	(**2***,**2***)

The real beauty and singularity of the *Prisoner's Dilemma* is that it implies a paradox. A quick analysis of the payoffs in Table 9 yields two extremes at (1,4) and (4,1), corresponding to the two interchangeable cases one player cooperating (C) and one not (D), but in contrast to the three previous games these are *not* Nash equilibria. There is only one Nash equilibrium at (2,2), which is in fact the Minimax solution too. This means that under the solution concepts of both Minimax strategy and Nash equilibrium, theory suggests that the two prisoner's will probably choose to betray one another, despite any previous arrangements. It is clearly evident that the outcome (3,3) is *mutually beneficial* and at the same time unattainable due to lack of communication. However, in therms of strict personal gain, defecting (D) is the dominant strategy for both and neither of them has any incentive to deviate from it. In other words, it appears that defecting is always the optimal choice *regardless* of what the other prisoner does - but if both adopt the same rationale, they will end up at (2,2) which is clearly worse than the (3,3) that they could have gotten if they had chosen mutual cooperation.

The essence of the paradox of *Prisoner's Dilemma* lies in the inherent conflict between *individual* and *collective rationality*. While individual rationality is well-understood, collective rationality deals with the scope of optimizing the *mutual* gain of the players. This is not a default behavior in strictly competitive situations, as in zero-sum games, or nonzero-sum games that do not imply cooperation. However, nonzero-sum games permit the idea of *mutually* optimal gains as a combination of *simultaneously* optimal separate payoffs. Under this broader scope, even (4,1) and (1,4) are worse than (3,3) since they yield a sum of 5 in gain value rather than 6, respectively.

It should also be noted that the single Nash equilibrium in *Prisoner's Dilemma* is stable, while the corresponding pairs of Nash equilibria in the three previous games are inherently unstable, since the players are not in agreement as to which of the two equilibria is preferable. Furthermore, in the three previous games the worst possible outcome comes when both players choose their non-Minimax strategy; in *Prisoner's Dilemma* this is not so. In fact, *Prisoner's Dilemma* has produced lengthy academic debates and hundreds of studies in a wide range of disciplines, from Game Theory and Mathematics to Sociology and Evolutionary Biology. The paradox of this game (as described above) has been illustrated as a notorious example where theory often fails to predict the true "gaming" outcomes in the real world: cooperation can emerge spontaneously, even though theory says it should not [1, 2, 27, 7].

Summary:

- There are four basic nonzero-sum game types of particular interest namely: *Leader* (or *Coordination*), *Battle of the Sexes*, *Chicken* and *Prisoner's Dilemma*.
- Three of these games (except *Prisoner's Dilemma*) have two "mirrored" pure Nash equilibria and players receive the worst possible payoff when they choose to deviate from their optimal Minimax strategy.
- *Prisoner's Dilemma* is a very unique type of game, since neither Minimax solution or Nash equilibrium (single one in this case) point to the best *mutually* beneficial outcome; this is informally labeled as the *paradox* of this game.

5 Signals, mechanisms & rationality

Game formulation and representation in analytical or extensive form are imperative for proper analysis and identification of equilibria. However, they fail to capture many elements of *gaming* as a multi-aspect process, especially in relation to *strategic moves*; these are actions performed by the players at different places and times, even before the realization of the current game, with the goal of enhancing strategic advantages and increasing the effectiveness of chosen strategies. Sometimes the "moves" are no more than message exchanges between the players, explicit or implicit, or simply tracking the history of previous choices in iterated games. Formulating these factors into a proper mathematical model can be very difficult, but nevertheless they are matters of great importance in real-world conflict situations.

5.1 Signals, carriers & bluffs

The exchange of messages between the players is a very useful option when a player is trying to model or even predict the behavior of its opponent(s). A message or *signal* from one player to another may be voluntary or involuntary, direct or indirect, explicit or implicit [46, 44]. In any case, it carries some sort of strategic information, which is always valuable to the other player if it can be asserted as credible with a high degree of confidence. On the other hand, if this credibility can be manipulated and falsely asserted as such, the source player may gain some strategic advantage by means of deceiving its opponent.

Strategic *signaling* is the process of information exchange between two or more players in a game, using any means or intermediate third-parties as carriers. If the source player does this deliberately, the purpose is to project some strategic preference or *stance* ("posturing") in the game without making any actual "move", in order to intimidate or coordinate with the opponent(s). This is particularly useful in situations where mutually beneficial equilibria are achievable but lack of preference ranking can lead to disastrous lack of coordination. The *Leader* and *Battle of the Sexes* games are such examples (see Tables 6 and

7). On the other hand, if the source player signals its opponent unintentionally, this strategic information could be a "leak" of such importance that may determine the actual outcome of the game.

Explicit signaling means that the source player sends out a clear message with undeniable association and content. An explicit signal may be voluntary or involuntary; in the later case, the message is simply a "leak" with very clear origin and content. *Implicit* signaling happens when the origin or (most commonly) the content of the message is somehow inconclusive or "plausibly deniable" as to the intentions of the source player. A signal exchange may occur directly between the players or via a third-party that performs the role of a carrier. A number of combinations of these attributes are possible in practice, employing direct/indirect messaging, voluntary/involuntary information exchange, with explicit/implicit messages. For example, a third-party carrier may share an implicit signal or "leaked" (involuntary) information about a player's stance with another player, participating in the game only as a mediator, coordinator or "referee", rather than an actively involved player.

A very special type of signaling is when the message exchange involves false information, i.e., a *bluff*. This kind of signals is a very common practice in games of imperfect and/or incomplete information, where the players do not have a complete view of the game structure itself and/or the opponents' choices, respectively. In this case, *false signaling* or *bluffing* is usually a strategic option by itself, exploiting this uncertainty regarding the true status of the game to enhance advantages or mitigate disadvantages. A very common example of such games is Poker, where a player with weaker deck of cards can project a *false stance* to its opponents, in order to avoid defeat or even secure a victory against players with better decks of cards [46, 44]. Bluffing can be realized directly between players or indirectly via a third-party carrier. In the later case, especially when the signaling is implicit and assumed involuntary, the credibility of the assertion is strongly associated with the credibility of the carrier itself. In other words, even if the source player could not project a successful bluff on its own, a credible third-party carrier might be the necessary intermediate to achieve such a move. The role of third-party mediators in signaling is a special topic in the study of strategic moves and how they affect the final outcome in games.

Summary:

- A *signal* between players is a voluntary or involuntary, direct or indirect, explicit or implicit exchange of a message; it is usually a declaration of *stance* (“posture”) in the game, i.e., intent to include or exclude a strategy from a set of open options.
- *Strategic moves*, e.g. signaling, project some strategic preference without making any actual “move”, in order to intimidate or coordinate with the other player(s).
- A *bluff* is a projection of false information, i.e., exploiting the incomplete/imperfect information structure of a game to gain some strategic advantage that could not be achievable if the game was of complete/perfect information.

5.2 Credibility, reputation, promises & threats

The effectiveness of projecting a strategic stance via signaling, regardless if it is true or bluff, depends heavily on the *credibility* of that signal, as well as the credibility of the player itself [46, 44]. When it comes to a single signal or stance, the credibility is closely linked to the level of compatibility of that signal or stance with the rationality of the player. Although rationality per se may be only an assumption with regard to one’s opponent, in general terms it is fairly easy to examine the matrix or the tree-graph representation of a game and establish whether a declared stance is beneficial or not to the associated player. In other words, if that player is assumed to behave rationally, Minimax strategies and Nash equilibria can be used to filter out choices that are clearly excluded, at least with a high probability.

The set of previous stances and/or moves, as well as their associated credibility values, can be used as the history or *reputation* of that player, which is in fact the a priori probability for any future stance and/or move of being consistent with its previous behavior [27]. Since games of complete and perfect information, e.g. Chess, are not compatible with false signaling and bluffs, the true theoretical aspect of credibility and reputation is relevant only in games of incomplete and/or imperfect information. Hence, Poker players are indeed characterized as being cautious or risk-takers according to their reputation on using bluffs in lower or higher frequency, respectively.

A player with a specific reputation can signal a specific stance to the others, projecting either a *promise* or a *threat*. A promise is a signal that usually declares the intent to cooperate, i.e., choose the less aggressive approach. This is particularly useful when the players need to coordinate in order to avoid much worse outcomes, as in the games *Leader* and *Battle of the Sexes* (see Tables 6 and 7). On the other hand, a threat is a signal that usually declares the intent to compete, i.e., choose the more aggressive approach. This is still useful as the means to enforce some kind of coordination, now in the form of extortion rather than willful cooperation. The *Chicken* game is such any example (see Table 8), where one player must force the other to swerve, in order to naturally end up in one of the two Nash equilibria and avoid the worst outcome of crash.

As it was mentioned earlier, *Prisoner's Dilemma* is a very special type of game, since neither Minimax solution or Nash equilibrium points to the mutually beneficial option of cooperation; however, if signaling between the prisoners is possible, i.e., if they are allowed to communicate with each other, cooperation becomes much more plausible: all they have to do is to promise each other to remain silent and threat to accuse the other as a retaliation if they see the other doing such thing. One of the most interesting topics in modern Game Theory is the study and analytical formulation of the conditions, the constraints and the exact processes of the *evolution of cooperation* in games like *Prisoner's Dilemma*, where typical theory fails to predict optimal strategies, although such strategies seem to exist, usually in accordance to some *Tit-for-Tat* variation [1, 2, 27, 7].

In any case, whether it is a promise or a threat, the signal or stance is labeled as credible or not. Hence, a *credible promise* is one that comes from a player with a reputation of being consistently reliable in fulfilling that promise, i.e., actually choosing less aggressive strategies when signaling intent to cooperate. Similarly, a *credible threat* is one that comes from a player with a reputation of being consistently reliable in fulfilling that threat, i.e., actually choosing more aggressive strategies when signaling intent to compete [28, 33].

Summary:

- *Promise* is a signal that usually declares the intent to cooperate, i.e., choose the less aggressive approach; it is useful when players need to coordinate in order to avoid much worse outcomes.
- *Threat* is a signal that usually declares the intent to compete, i.e., choose the more aggressive approach; it is useful a player wants to enforce some kind of coordination, in the form of extortion.
- *Credibility* is closely linked to the level of compatibility of a signal or stance with the rationality of the player; in practice, it is a measure (probability) of whether the player will fulfill a promise or a threat, if necessary.
- *Reputation* of a player is the a priori probability for any future stance and/or move of being consistent with its previous behavior.
- *Credible promises* and *credible threats* are associated with the reputation and credibility of each player, as well as the actual payoffs in the corresponding game matrix.

5.3 Utility, incentives & "rational irrationality"

As it was mentioned earlier, if that player is assumed to behave rationally, i.e., trying to minimize losses and maximize gains in terms of actual payoffs in each outcome, the credibility of a promise or a threat can be easily established with a high probability. Nevertheless, the fact that this is just a probability and not a perfect forecast comes from the fact that, in turn, the level of rationality of

Tab. 10: The typical setup of the *Hostage Situation* game with two players. Player-1 is the assaulter and Player-2 is the rescuer-protector.

Hostage Situation		Player-2 *C*	Player-2 *D*
Player-1	*C*	(2,3)	(1,4*)
	D	(4*,**2***)	(**3***,1)

that player cannot be evaluated perfectly and in exact terms.

Rationality and *incentives* of a player emerge naturally from the exact formulation of its own *utility function*, which is nothing more than a generalization of the loss/gain function that is described by the matrix or the tree-graph of the game [33, 16, 28]. If the formulation of the game's payoff matrix is perfect, then it is clear when a strategy is optimal for a player and when it is not. However, the truth is that these payoff values may not reflect the exact *utility*, i.e., overall loss/gain value for that player, usually due to some "hidden" outcomes or side-effects. For example, a game matrix may describe the payoffs for each outcome and each player correctly, but with the assumption that these players are rational in the same way: winning over their opponent; this may not be true, e.g. when one player cares more about securing that their opponent does not win, rather than securing their own win. In other words, when the players' rationality is not symmetrically the same, then they do not share the same utility function and the true payoffs in the game matrix may actually be quite different.

A very classic example of such games, assumed to be symmetric when they are actually asymmetric by nature, is the *Hostage Situation*, described in analytical form by Table 10. If the two opponents are treated as similarly rational, i.e., symmetric in terms of incentives and behavior, then the game is not much different than the classic *Chicken*, where one must convince the other to swerve first, in order to avoid the crash. This translates to either the authorities give in to the assaulter's demands or the assaulter eventually surrenders to the authorities, both outcomes assumed to be equally rational, correspondingly, to each player. However, if for some reason the assaulter is more determined than initially presumed, preferring to fight to the death rather than surrendering and ending up in jail, then the game is inherently asymmetric and the payoff matrix is quite different, as illustrated in Table 10. What the matrix shows is that now Player-1, i.e., the assaulter, has a dominant strategy of always choosing the most aggressive stance, no matter what the authorities choose to do. There is no pure Minimax solution here, since there is no pure saddle-point (see payoffs "3" and "2" in bold); however, there is now a single Nash equilibrium at (4,2), i.e., aggressive assaulter and passive authorities - this is in fact the standard approach internationally in all hostage situations: the authorities start with trying to establish a communication link and negotiate with the assaulter, rather than choosing a rescue operation by direct action that could put the hostages in danger.

As it is evident from the *Hostage Situation* game of Table 10, the authorities are normally guided to a more passive and cooperative approach of negotiating rather than using force, because the incentive is to protect the hostages at all costs. This effectively translates to employing a utility function that includes

Tab. 11: The typical setup of the *Kamikaze* game with two players. Player-1 is the "kamikaze" and Player-2 is the defender.

Kamikaze		Player-2	
		C	D
Player-1	C	(2,3)	(1,4*)
	D	(4*,*1*)	(**3***,***2****)

a high priority on the hostages' lives, higher than the immediate capture or incapacitation of the assaulter. Hence, the rationality of Player-2 dictates a more passive, cooperative stance. This changes drastically if, during this evolution, the lives of hostages are put in severe danger, e.g. when the assaulter poses a very credible threat or actually harms a hostage (assuming there are more). In this case, the authorities should change stance and employ the more aggressive option, because this is now the optimal response.

Table 11 illustrates the *Kamikaze* game, which is actually a slightly modified *Hostage Situation* game in terms of payoff matrix. The game is still asymmetric and the only variation is the swapping of payoff values {2} and {1} for Player-2 (marked in italics), which illustrates the new fact that at this point it is more harmful for the hostages to remain idle rather than using direct force to rescue them, even if this too poses some danger to them - again, this is exactly the standard approach internationally in all hostage situations: the authorities follow strict rules-of-engagement which state that, once it is established that the lives of hostages is in clear and severe danger, direct action is to be employed immediately. The same setup emerges when the *Kamikaze* game is studied according to its name: when one player (assaulter) is more concerned about damaging the opponent (defender) rather than protecting itself, then there is indeed a dominant strategy of always choosing the most aggressive stance, no matter what the defender chooses to do. Likewise, the defender is now forced to choose between its two worst outcomes and naturally chooses the less damaging one, i.e., direct counter-action rather than swerve. Here, the passive stance is established as more damaging than all-out-conflict, exactly as in *Hostage Situation* with a very aggressive assaulter. In terms of game analysis, now there is indeed a pure Minimax solution at (3,2), which is also the single Nash equilibrium of the game. This explains why there is practically no other rational (strategically optimal) way to defend against a murderous hostage-taker or a desperate kamikaze than employing equally aggressive response.

The concepts described along the strategic analysis and "rationalization" of the players in games like *Hostage Situation* and *Kamikaze* illustrate how a seemingly irrational course of actions can be easily explained and even classified as rational behavior, if the proper utility functions are employed. In other words, if the utility of each and every player is defined correctly, then all players in any game can be labeled as "rational" ones. This proposition is often referred to as "rational irrationality" (valid/explainable behavior), rather than "irrational rationality" (incomprehensible behavior) [27].

Summary:

- *Utility* is the generalized cost/gain function of a player in a specific game, depending on the outcomes but including any "hidden" regards and side-effects.
- Given a specific *utility function*, a player's *incentives* emerge naturally as the rational behavior of the underlying payoff-optimization process.
- A player's behavior may seem "irrational" if its utility function is incomplete; given a properly defined utility function, a player's behavior can always be labeled as rational per se.
- *Hostage Situation* and *Kamikaze* are two examples of (asymmetric) stand-off games where the notion of "rational irrationality" is fully explained via proper definition of the corresponding utility functions for the assaulter.

6 The frontier

This paper included only some of the most basic concepts of Game Theory, including solution methods and representations of typical games of special interest, like *Chicken* and *Prisoner's Dilemma*. However, these are only a scratch on the surface of what lies beneath, the rigorous mathematical theory and the complex, some still unsolved, problems in this extremely interesting and useful scientific area.

All the games and setups presented thus far was somewhat "too perfect", too simple compared to real-world situations of conflict. There are few cases where only two players are involved, their moves are full observable and their incentives clear and consistent. In most conflicts, groups of players are spiraling in alternating rounds competing and cooperating, each knowing its own utility function and very little about the others', while signaling, third-party credibility assertions and continuous bargaining are common things. Is there really a way Game Theory can address all these aspects in the same clarity, mathematical robustness and universality as is does with simple cases of zero-sum and nonzero-sum games like the ones presented previously?

The short answer is "No". Game Theory is the mathematical way to approach some of the most complex problems the human mind has ever encountered. For example, what are the prerequisites, the dynamics and the survivability of the evolution of cooperation as a strategy, in human or animal societies? What is the asymptotic behavior of such "cooperative" groups? Can they survive in an environment of pure competition? These issues are addressed in other aspects of the theory, namely the *Evolutionary Stable Strategies* (ESS), not analyzed in this study. In short, ESS are patterns of behavior in games of pure competition and/or possible cooperation, such as the *Prisoner's Dilemma*, that not only may emerge spontaneously but also survive as optimal strategies in iterative games. *Tit-for-Tat* [1, 2] is such an example of ESS in iterated *Prisoner's Dilemma*: cooperation can emerge spontaneously given a set of conditions, primarily (a) players "start nicely", (b) continue with reciprocity, (c) don't know when the game finishes. Although it seems simple enough, spontaneous cooperation in conflict situations is one of the most intriguing and theoretically complex problems in Game Theory today.

In a slightly simpler scenario, a player may be involved in a game with another player, while *at the same time* its strategic choices are relevant to a second game, with some other player. For example, a politician may be in a "bargain" with voters, trying to gain their support by promising specific actions if elected, while at the same time a second "bargain" may be taking place in parallel with the party's main policies and governmental plan if it comes to power. If some of that politician's promises are on conflict with the party's main lines, then as a player is involved in what is called a *two-level game* [39, 38]. This form of gaming was first proposed by Putnam in the late '70s and it models two-level or multi-level conflict situations in general, where the strategic choices of a player affect two or more simultaneous games. The solution concepts and equilibria are not much different than those of simple games, but now a strategy is optimal and produces a stable outcome only if it is such *simultaneously* in all these games.

Another very interesting aspect of gaming in general is the evolution of strategies and each player's behavior as each observes the others' moves. In single-step games, the Minimax solution (pure or mixed) is the one that dictates the optimal strategy for each player. The concept of iterative gaming is much more general, since it includes cases where the same players may face one another in the same single-step games many times in the future. In this case, Nash equilibria predict the most probable outcomes with much better accuracy. But the knowledge that there will be a "next round", especially when players alternate moves and one can observe the other before making its own (e.g. in Chess), then the game analysis can expand to two or more steps ahead. In practice, the player does not only take into account the strategic choices available to the opponent(s) but also the "what if" combinations of moves-countermoves. Hence, the corresponding game matrix includes these combinations of composite states on the opponent(s) side and the payoffs are estimated accordingly. This type of composite multi-step setup is often referred to as a *metagame* [46]. The extended-form representation of metagames is more natural than the analytical (matrix) form, but the identification of equilibria and solutions is somewhat less straight-forward.

Some games involve elements of chance regarding the game's state or partial information regarding the observability of each player's moves. In such games of imperfect information, modeling via a game matrix or a tree-graph can be problematic, since many of the paths may be mutually exclusive and not just alternative choices. In the '60s, very early on in the history of Game Theory, Harsanyi introduced the so-called *Harsanyi transformation* [21, 22, 28] for transforming a game of incomplete information to an equivalent game of complete but imperfect information. This may not seem much, but in reality there is a very distinct and important difference between them. If a random event dictates the exact structure and payoffs of the games, perhaps even the strategic behavior of the players, then the analysis of such a game is inherently a very difficult task. On the other hand, the Harsanyi transformation models this random event as a deterministic one, removing the element of chance and introducing the notion of "hidden" information about it. In practice, this results in creating multiple variations of the game, one for each possible configuration, and treating them separately. After they are individually analyzed, solutions and equilibria are combined together within a probabilistic framework, introducing the more generalized concept of Bayesian Nash equilibria [28].

In real-world conflict situations it is not uncommon that one or some of the players have a different knowledge or "view" of the game structure, its payoffs and the other players' preferences. This means that each player acts upon its own payoff matrix, possibly very different in structure and values than the one used by the other players. Of course, all players are involved in the same, single game and the payoffs on each outcome is effectively a single one, despite each player's unique view of the game. This is extremely important if some of the players have a more complete view of the game, i.e., when they address the game as one of (almost) complete information, while some opponents address it as one of incomplete information. These special types of conflict are often referred to as *hypergames* [47, 4]. Introduced by Bennett and Dando in late '70s and later revised in the '00s by Vane and others, hypergames is a very efficient way to describe games of asymmetric information between players by employing different variations of the game matrix or tree-graph, according to each player's view. In practice, hypergames are treated the same way as simple games, with each player deciding its strategic choices according to its own view and, subsequently, combining the (partial) outcomes together.

Game Theory is a vast scientific and research area, based almost entirely on Mathematics and some experimental methods, with applications that vary from simple board games and auctions to Evolutionary Psychology and Sociology-Biology in group behavior of humans and animals. Although real-world situations reveal that sometimes its predictive value is limited, the robust theoretical framework and solution concepts provide an extremely valuable set of tools that clarifies the inner workings and dynamics of conflict situations.

Summary:

- In accordance to Nash's bargaining theorem, cooperation can *emerge spontaneously*, even in competitive games, when a specific set of prerequisites are satisfied.
- *Evolutionary stable strategies* (ESS) are patterns of behavior in games of pure competition and/or possible cooperation that survive as optimal strategies in iterative games.
- In *two-level games*, a player may be involved in a game with another player, while *at the same time* its strategic choices are relevant to a second game, with some other player.
- *Metagames* are multi-step game setups where the corresponding game matrix includes combinations of "what if" composite states, regarding the future strategic choices of the opponent(s).
- The *Harsanyi transformation* is used in games of incomplete information, e.g. when the game structure and payoffs depend on some random event, to transform it to an equivalent game of complete but imperfect information.
- *Hypergames* is a very efficient way to describe games of asymmetric information between players by employing different variations of the game matrix or tree-graph, according to each player's view.
- In general, Game Theory is a vast scientific and research area with robust theoretical foundation that can be used as a predictive tool, as well as (mostly) an extremely valuable approach to analyze conflict situations.

Acknowledgement: *This work is dedicated to John F. Nash, pioneer and mathematical genius, who was killed earlier this month on May 23th 2015 in a car accident along with his wife Alicia. His inspirational work and breakthrough ideas has changed Game Theory and Economics forever.*

References

[1] R. Axelrod. *The Evolution of Cooperation.* Basic Books, NY: 1984, 1984.

[2] R. Axelrod and D. Dion. The further evolution of cooperation. *Science, 242 (1988), 1385-1390*, 1988.

[3] R. Ben-Yashar and S. Nitzan. Optimal decision rules for fixed-size committees in polychotomous choice situations. *Social Choice and Welfare, 18 (2001) 737-746*, 2001.

[4] P. G. Bennett and M. R. Dando. Complex strategic analysis: A study of the fall of france. *Journal of Operational Research Society, 33, pp.41-50*, 1979.

[5] S. Berg. Indirect voting systems: Banzhaf numbers, majority functions and collective competence. *Eur. J. Pol. Econ., 13 (1997), pp.557-573*, 1997.

[6] C. Camerer, T. Ho, and K. Chong. Behavioral game theory: Thinking, learning, and teaching. *Learning and Teaching (November 13, 2001) - Caltech Working Paper*, 2002.

[7] J. L. Casti. *Reality Rules II: Picturing the world in Mathematics - The Frontier.* John Wiley & Sons, 1997.

[8] P.-L. Chang, V. Chua, and M. Machover. L s penrose's limit theorem: Tests by simulation. *Mathematical Social Sciences, 51 (2006) 90-106*, 2006.

[9] Marquis de Condorcet. An essay on the application of probability theory to plurality decision making: An election between three candidates. *Sommerland and Mclean (1989) 66-80*, 1989.

[10] C. Daskalakis, P. W. Goldberg, and Ch. Papadimitriou. The complexity of computing a nash equilibrium. *Communications of the ACM, 52(2):89-97, 2009*, 2009.

[11] C. Daskalakis, P. W. Goldberg, and Ch. Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing, 39(1):195-259, 2009*, 2009.

[12] C. Daskalakis and Ch. Papadimitriou. The complexity of computing a nash equilibrium. In *38th ACM Symposium on Theory of Computing (STOC 2006)*, 2006.

[13] A. Dixit and B. Nalebuff. Game theory (2nd/ed.). *The Concise Encyclopedia of Economics*, 2008.

[14] M. Dresher. *Games of Strategy – Theory and Applications.* Prentice-Hall / RAND Corp, USA: 1961, 1961.

[15] T. S. Ferguson. *Game Theory (2nd/Ed).* Mathematics Department, UCLA: 2014, 2014.

[16] D. Fudenberg and J. Tirole. *Game Theory.* MIT Press, 1991.

[17] H. Georgiou. Collective decision efficiency and optimal voting mechanisms: A comprehensive overview for multi-classifier models. *arXiv.org preprint (en)(arXiv:1502.02191v1 [cs.GT])*, 2015.

[18] H. Georgiou and M. Mavroforakis. A game-theoretic framework for classifier ensembles using weighted majority voting with local accuracy estimates. *arXiv.org preprint (en)(arXiv:1302.0540v1 [cs.LG])*, 2013.

[19] H. Georgiou, M. Mavroforakis, and S. Theodoridis. A game-theoretic approach to weighted majority voting for combining svm classifiers. In *Int. Conf. on ANN (ICANN), 10-13 Sept. 2006, Athens, Greece, in: S.Kollias, et al. (Eds.): ICANN 2006, Part I, LNCS 4131, 2006, pp. 284-292*, 2006.

[20] S. Gu. Game theory and linear programming. *Duke University, CS lecture presentation (2008)*, 2008.

[21] J. C. Harsanyi. Bargaining in ignorance of the opponent's utility function. *Journal of Conflict Resolution (1962)*, 1962.

[22] J. C. Harsanyi. Games with incomplete information played by bayesian players, i-iii. *Management Science, 14, pp.159-182,320-332,468-502 (1967)*, 1967.

[23] D. Karotkin. The network of weighted majority rules and weighted majority games. *Games and Econ.Beh., 22 (1998), pp.299-315*, 1998.

[24] L. Kuncheva. *Combing Pattern Classifiers - Methods and Algorithms*. Wiley, 2004.

[25] R. D. Luce and H. Raiffa. *Games and Decisions*. Wiley, New York, 1957.

[26] M. Maschler, E. Solan, and S. Zamir. *Game Theory*. Cambridge University Press: 2013, 2013.

[27] L. Mero. *Moral Calculations - Game Theory, Logic and Human Frailty*. Springer-Verlag, New York, 1998.

[28] C. Montet and D. Serra. *Game Theory and Economics*. Palgrave Macmillan, New York, 2003.

[29] J.F. Nash. Non-cooperative games. *Annals of Math., 54 (1950) pp.286-295*, 1950.

[30] J.F. Nash. Non-cooperative games (phd dissertation). *PhD dissertation, Princeton University*, 1950.

[31] N. J. Nilsson. *Artificial Intelligence: A New Synthesis*. Elsevier / Morgan Kaufmann, 1998.

[32] S. Nitzan and J. Paroush. Optimal decision rules in uncertain dichotomous choice situations. *International Economic Review, 23 (1982) 289-297*, 1982.

[33] G. Owen. *Game Theory (3rd/Ed)*. Academic Press, 1995.

[34] Ch. Papadimitriou. Complexity of finding a nash equilibrium. *Berkeley University (EE-CS), lecture notes (2011)*, 2011.

[35] L. S. Penrose. *On the objective study of crowd behavior.* HK Lewis & Co, London, 1952.

[36] W. H. Press and F. J. Dyson. Iterated prisoner's dilemma contains strategies that dominate any evolutionary opponent. *PNAS, 109 (26)(2012) 10409-10413*, 2012.

[37] J. Prywes. The mathematics of magic: The gathering - a study in probability, statistics, strategy, and game theory. *(draft paper)*, 1999.

[38] R. D. Putnam. Diplomacy and domestic politics: The logic of two-level games. *International Organization, 42 (3) 1988*, 1988.

[39] R. D. Putnam and C. R. Henning. *The Bonn summit of 1978: How does international economic policy coordination actually work?* Brookings Institution, USA: 1986, 1986.

[40] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach (3rd/Ed).* Prentice Hall, 2009.

[41] A. Schalk. *The Theory of Games and Game Models (CS3191 - lecture notes).* Department of Computer Science, University of Manchester, 2003.

[42] L. Shapley and B. Grofman. Optimizing group judgemental accuracy in the presence of interdependencies. *Public Choice, 43 (1984) 329-343*, 1984.

[43] G. Sierksma. *Linear and Integer Programming: Theory and Practice (2nd/Ed).* CRC Press: 2001, 2001.

[44] S. Stahl. *A gentle Introduction to Game Theory.* Mathematical World, vol.13, American Mathematical Society, 1999.

[45] A. Taylor and W. Zwicker. Weighted voting, multicameral representation and power. *Games and Econ. Beh., 5 (1993), pp.170-181*, 1993.

[46] L. C. Thomas. *Games, Theory and Applications.* Ellis Horwood Ltd, UK, 1984.

[47] R. Vane. Advances in hypergame theory. *Workshop on Game Theoretic and Decision Theoretic Agents – Conference on Autonomous Agents and Multi-Agent Systems, 2006*, 2006.

[48] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior.* Princeton Univ. Press, Princeton, 1947.

[49] Wikipedia.org. Extended-form game (article).

[50] Wikipedia.org. Rubik's cube (article).

[51] P. Young. Condorcet's theory of voting. *Amer. Pol. Sci. Review, 82 (4)(1988) 1231-1244*, 1988.

Harris Georgiou received his B.Sc. degree in Informatics from University of Ioannina, Greece, in 1997, and his M.Sc. degree in Digital Signal Processing & Computer Systems and Ph.D. degree in Machine Learning & Medical Imaging, from National & Kapodistrian University of Athens, Greece, in 2000 and 2009, respectively. Since 1998, he has been working with the Signal & Image Processing Lab (SIPL) in the Department of Informatics & Telecommunications (DIT) at National & Kapodistrian University of Athens (NKUA/UoA), Greece, in various academic and research projects. In 2013-2015 he worked as a post-doctorate associate researcher with SIPL in sparse models for distributed analysis of functional MRI (fMRI) signals. He has been actively involved in several national and EU-funded research & development projects, focusing on new and emerging technologies in Biomedicine and applications. He has also worked in the private sector as a consultant in Software Engineering and Quality Assurance (SQA, EDP/IT), as well as a faculty professor in private institutions in various ICT-related subjects, for more than 17 years. His main research interests include Machine Learning, Pattern Recognition, Signal Processing, Medical Imaging, Soft Computing, Artificial Intelligence and Game Theory. He has published more than 65 papers and articles (43 peer-reviewed) in various academic journals & conferences, open-access publications and scientific magazines, as well as two books in Biomedical Engineering & Computer-Aided Diagnosis and several contributions in seminal academic textbooks in Machine Learning & Pattern Recognition. He is a member of the IEEE and the ACM organizations and he has given several technical presentations in various countries.

Section 2: Collective decision efficiency and optimal voting mechanisms

Summary:

A new game-theoretic approach for combining multiple classifiers is proposed. A short introduction in basic Game Theory and coalitions illustrate the way any collective decision scheme can be viewed as a competitive game of coalitions that are formed naturally when players state their preferences. The winning conditions and the voting power of each player are studied under the scope of Banzhaf and Shapley numbers, as well and the collective competence of the group in terms of correct collective decision. Coalitions and power indices are presented in relation to the Condorcet criterion of optimality in voting systems, and weighted Borda count models are asserted as a way to implement them in practice. A special case of coalition games, the weighted majority games (WMG) are presented as a restricted realization in dichotomy choice situations. As a result, the weighted majority rules (WMR), an extended version of the simple majority rules, are asserted as the theoretically optimal and complete solution to this type of coalition gaming. Subsequently, a generalized version of WMRs is suggested as the means to design a voting system that is optimal in the sense of both the correct classification criterion and the Condorcet efficiency criterion. In the scope of Pattern Recognition, a generalized risk-based approach is proposed as the framework upon which any classifier combination scheme can be applied. A new fully adaptive version of WMRs is proposed as a statistically invariant way of adjusting the design process of the optimal WMR to the arbitrary non-symmetrical properties of the underlying feature space. SVM theory is associated with properties and conclusions that emerge from the game-theoretic approach of the classification in general, while the theoretical and practical implications of employing SVM experts in WMR combination schemes are briefly discussed. Finally, a summary of the most important issues for further research is presented. This report is a compact introduction to the theoretical material upon which a new expert fusion model can be designed.

Citation:

"Collective decision efficiency and optimal voting mechanisms: A comprehensive overview for multi-classifier models", H. Georgiou, ArXiv.org: 7-Feb-2015 (arXiv: 1502.02191v1 [cs.GT]).

Collective decision efficiency and optimal voting mechanisms: A comprehensive overview for multi-classifier models

Harris V. Georgiou (MSc, PhD)

Dept. of Informatics & Telecommunications,

National Kapodistrian Univ. of Athens (NKUA/UoA)

Email: harris@xgeorgio.info -- URL: http://xgeorgio.info

Last Updated: 8 February 2015

Abstract

A new game-theoretic approach for combining multiple classifiers is proposed. A short introduction in basic Game Theory and coalitions illustrate the way any collective decision scheme can be viewed as a competitive game of coalitions that are formed naturally when players state their preferences. The winning conditions and the voting power of each player are studied under the scope of Banzhaf and Shapley numbers, as well and the collective competence of the group in terms of correct collective decision. Coalitions and power indices are presented in relation to the Condorcet criterion of optimality in voting systems, and weighted Borda count models are asserted as a way to implement them in practice. A special case of coalition games, the weighted majority games (WMG) are presented as a restricted realization in dichotomy choice situations. As a result, the weighted majority rules (WMR), an extended version of the simple majority rules, are asserted as the theoretically optimal and complete solution to this type of coalition gaming. Subsequently, a generalized version of WMRs is suggested as the means to design a voting system that is optimal in the sense of both the correct classification criterion and the Condorcet efficiency criterion. In the scope of Pattern Recognition, a generalized risk-based approach is proposed as the framework upon which any classifier combination scheme can be applied. A new fully adaptive version of WMRs is proposed as a statistically invariant way of adjusting the design process of the optimal WMR to the arbitrary non-symmetrical properties of the underlying feature space. SVM theory is associated with properties and conclusions that emerge from the game-theoretic approach of the classification in general, while the theoretical and practical implications of employing SVM experts in WMR combination schemes are briefly discussed. Finally, a summary of the most important issues for further research is presented. This report is a compact introduction to the theoretical material upon which a new expert fusion model can be designed.

Keywords: game theory, coalition games, multi-classifiers, weighted majority, voting mechanisms

Part I – Theoretical Background

1.1 Preface

The ultimate goal of any pattern recognition system is to achieve the best possible classification performance for the specific problem at hand. This objective has led to the development of sophisticated algorithms and classification models in a way that best captures and enhances the underlying structure of the input space. As the complexity and intercorrelation between the classes increases, more robust and efficient classifiers have to be employed in order to provide adequate adaptability and generalization. However, the sensitivity and specificity of each classifier model can prove efficient in one case and inefficient in another, thus there is no clear indication of one single classifier design that can be considered as a universal pattern recognition solver.

Recent studies have focused in the possibility of taking advantage of this complementary performance of various classifier designs, in order to produce combination schemes for optimal fusion of multiple classifiers. Specifically, each classifier is considered as a trained expert that participates along with others in a "committee", which produces a collective decision according to some well-specified rule. The most common combination rules include the min rule, the max rule, the median rule, the majority voting rule, the averaging rule, etc.

It has been proven that all these combination rules are special realizations of two basic combinations schemes, namely the product rule and the sum rule [01]. The product rule essentially combines the classifiers' estimations on a-posteriori probabilities in a way that is consistent with the classic probability theory on independent events. When the classifiers are considered independent, i.e., when the decision of each classifier is not affected in any way by the corresponding decision made by all the other classifiers, then the join probability of the combined result can be calculated as a product of all the individual probabilities. Although the product rule is based on solid theoretical background, specifically the Bayes theory, it has provided only moderate results in practice in pattern recognition problems. One reason is that, in practice, well-trained classifiers tend to produce similar predictions, therefore the hypothesis on independent decisions is not established, although the classifiers produce their estimations separately. The other reason for poor performance is the fact that, when the combined output of all the experts is based on a product of their individual estimations, one single severe error or poor training on one expert could drive the whole group into similarly poor performance.

In contrast, the sum rule utilizes the individual experts' estimation in some additive form. All the popular combination rules, including the *majority voting*, the *median* and the *weighted averaging* rules, can be derived from the generic sum rule. It has been proven both theoretically and experimentally that these sum-based rules outperform the product rule in terms of minimum classification error (MCE) and error sensitivity. The reason for this enhanced performance in comparison to the product rule is the fact that the sum form of these rules effectively nullifies any single outliner estimations that may be produced by a poor classifier. This means that any single severe error has only minimal effect on the final combined estimation. Therefore, any combination scheme that utilizes a sum-based rule when calculating the experts' collective decision is relatively resilient to individual expert errors and this property of resiliency is increased as the number of combined experts increases.

The exact theoretical analysis of any of these combination schemes has proven noticeably cumbersome and complex over the years. Recent studies have established some general properties on total error magnitude and sensitivity, but only under strict preconditions on number of classes and their distributions [02]. For K classes, current rules cannot guarantee an increase of overall performance for the combined decision when each expert exhibits accuracy only at $p=1/K$. On the contrary, averaging rules usually establish increase in performance when each expert exhibits accuracy at least $p>0.5$. Much work has been conducted experimentally for identifying whether the most important factor for the final performance of such a combination scheme depends more on the exact form of the combination rule or the diversity of the experts themselves. Ensemble design methods, such as ADAboost, Random Forests and Arcing [12], focus exactly on enhancing the experts' diversity, either by means of amplifying the classification training process in areas close to the decision boundary, or by employing dataset splitting algorithms for improving the independency criterion. However, conceptualizing and quantifying diversity between classifier outputs is very challenging and it is usually based on experimental results, rather than a solid theoretical basis. As a result, most combination schemes employ optimization heuristics when designing the exact form and parameters of such a combination rule, e.g. calculating the optimal weights in a weighted averaging model [04].

Recently, more generic approaches have been proposed for designing the combination stage of such multi-classifier methods. Specifically, instead of employing a fixed rule for combining the individual outputs from the experts, a new fully functional expert node is introduced in the form of a meta-classifier: using the outputs of all the previous K classifiers as input, it produces an arbitrary mapping between the K-dim individual decision space into the final output space. This meta-classifier can essentially be any linear or non-linear model that is usually trained in the same way any of the other K classifiers is trained on the base data. Experimental results for models using neural networks and meta-classification nodes have been proven very efficient in many practical problems, justifying the practical gain in introducing classifier combination schemes in cases where the complexity of a pattern recognition problem requires the use of multiple, highly specialized experts [03].

Despite the fact that many practical solutions have been proposed and tested for combining multiple classifiers, the core issues of the combination problem still remains:

(a) Is there a generic and simple way to describe current combination rules?

(b) Is there an optimal realization of this rule for combining multiple classifiers?

(c) If such an optimal rule exists, does it cover the non-linear combination schemes too?

Before these questions can be answered, a brief introduction on special aspects on collective decision theory has to be made first.

1.2 Rank and response combination from classifier confidence transformation

Any classifier combination scheme is restricted by the type and form of the individual participating classification models, as their outputs must be compatible and suited for using them in the selected combination rule. Classifiers can be generally categorized into three types, according to their output [04]:

- Type-I : the classifier produces a simple statement on class selection
- Type-II : the classifier produces a ranked list of decreasing preferability of each classes
- Type-III : the classifier produces assignments of estimated probabilities for each class

Most combination schemes use Type-I or Type-III classifiers and most of them do not allow mixed types. Nevertheless, it is possible to convert between these types in some cases, e.g. assigning simple class selections (Type-I) for the maximum estimated probability (Type-III) or for the first preference in a ranked list (Type-II). In practice, Type-II classifiers are more generic than Type-I and Type-III classifiers are more generic than Type-II, however it is possible to calculate posterior information about the estimated probabilities for each class in Type-II classifiers by examining the produced confusion matrix. In fact, many combination schemes that exploit the confusion matrix data have the advantage of using already trained classifiers with known performance, instead of applying complete re-training of all the K experts when optimizing the combination node [05].

The confusion matrix can be considered as an experimental estimation of each classifier's performance on the specific task at hand and it can be used as a measure of confidence related to the predictions produced by this particular classifier. Specifically, the elements of the confusion matrix can be used directly in a transformation model, where individual correct or incorrect counts can be rescaled and mapped in a way that is consistent with a predetermined probabilistic model. Usually, such a formulation includes a scaling or normalization function and an activation function that is consistent with a specific type of confidence measure. The log-likelihood (linear), the likelihood (exponential) and sigmoid formulations have been proposed among others as candidate functions for confidence transformation [04]. Essentially, this process ensures that the outputs of the individual classifiers are consistent and compatible with each other and with the combination rule.

When using classifiers of Type-II, i.e., classification outputs that include a ranked list of preferences to the available classes, simple counts in the form of a confusion matrix are not possible. Instead, the ranked lists have to be converted into a simple metric that defines the overall "preferability" of each class, according to multiple classification outputs. In other words, instead of counting the number of times each class is selected as the best candidate in a simple class prediction (Type-I), a measure of "desirability" is calculated by assigning desirability points or "ranks" in every sorted list of preferences and then summing them for each class separately. This scheme is known as the *Borda count method* of combining ranked lists of predictions, i.e., the outputs of Type-II classifiers. Usually, the Borda count is calculated using ranks equivalent to the position of each class index in a sorted list, which means that in a N-class problem the first rank position receives N-1 points, the second receives N-2 points and the last position receives 0. It is possible to allow a weighted scheme when assigning these ranking points, so that the distribution is not uniform but arbitrary. In this case, the scheme is called *weighted Borda count* or *wBorda*.

The Borda and the more generic wBorda count methods are based on the presumption that the class selection at the rank position (i+1) is the second most probable candidate when classifying at class selection of rank (i). However, due to the absence of explicit probability estimations, it is not possible to directly extract "closure" measurements between these two choices. The weights in such a combination scheme can be designed in a way that produces this closure measurement in an optimized way.

One obvious question is whether ranked list combination produces better results that simple majority selections. Indeed, in many cases the simple class selection rule and the corresponding majority selection for the final output produces different results than the one produced by Borda count [05]. However, in the case of weighted Borda count, if weights are assigned in a specific way that is consistent with a required criterion, i.e., the majority rule, then the same combination result can be achieved. In the case of the majority rule, the weights that need to be assigned are w=1 for the first rank position and w=0 for the rest of the ranks. In general, a weighting scheme can be applied in the rank positions in a way that satisfies all the requirements needed by the typical combination rules, like the sum rule and its specialized versions (min, max, median average, majority, trimmed means, spread combiner). The product rule can also be applied in the same way [05].

It has been proposed that the weights of the rank items can be variable and proportional to the measured properties of each classifier performance. When probability estimations are available (Type-III), like when using a neural network classifier, they can be used instead. In the more restricted case of simple class selections (Type-I), it has been suggested that the confusion matrix itself can be used as the basis for calculating ranked classifications, translating the a posteriori probabilities for each class into ranked list of preferences [05]. In this case, the first assumption is that the behavior of each classifier is known it is characterized by its confusion matrix, and second that this prior behavior is representative of its future behavior. The second assumption is also applied for Type-II or Type-III classifiers, i.e., even when the confusion matrix is of no importance and its validity is increased proportionally with the size of the datasets on which the classifiers are tested.

But why restrict the classifier combination scheme into a Type-II, i.e., ranked classifications? It is true that the lack of explicit probability estimation from the classifier itself produces an inherent lack of information about the classification itself. However, many widely used classification models, including Support Vector Machines (SVM), are inherently designed to produce simple class selections as output. Furthermore, the use of simple class selections or ranked lists is required when a specific weighting profile has to be applied uniformly throughout an entire "committee" of experts, like in voting schemes. In other words, experts are weighted according to their competence in an adaptive way, but the calculation of these weights is not a subject each classifier's own performance. Instead, these weights are the realization of a collective decision rule, like in a wBorda scheme. This issue and its implications will be discussed later on under the scope on weighted majority rules (WMR).

1.3 Cascaded versus joint parameter optimization in combination schemes

There are three general groups of combination rules that can be applied when creating a mixture of experts: (i) the fixed rules, (ii) the trained rules, and (iii) weighted combinations of confidence transformation. The fixed rules group contains all the typical rules discussed thus far, including the product rule, the sum rule and its specialized version, etc. Trained rules refer to the case of meta-classifier nodes, where an arbitrary expert is trained experimentally upon the best way to combine the outputs of K experts against a given training dataset. Finally, the third group refers to models that are based on weighted order statistics, where a weight is assigned for each rank of confidence measure, rather than each classifier as in the case of weighted combination of classifier outputs. An example of this third group of combination rules is the Borda and, more specifically, the wBorda count models.

As mentioned earlier, the parameters of the combination rule itself, i.e., the weights in a weighted average or in a wBorda count, are a subject of optimization against a specific criterion, normally the minimum classification error (MCE). Similarly, for trained rules, the meta-classifier is trained according to the same optimality criterion. Since this combined classification process can be realized as either two separate stages in a cascaded model or a unified modular architecture, the optimization process can address each stage separately or jointly together. The latter case is often used for trained rules, as the meta-classifier rule can be trained jointly together with the K classifiers of the first stage.

For linear trained rules, the optimal weights correspond to the relative confidence attributed to each classifier, as in the fixed rules. However, an optimization process determines the best weighting profile based on a specific training dataset, instead of using a pre-defined weighting profile as the fixed rules suggest. For rules based on typical linear discriminant functions, the optimization process can be realized as parameter estimation via regression or by applying any other formulation of typical linear classifier model. In the case of rules based on *weighted order statistics*, like the wBorda count, the most commonly used method is also regression or some other linear optimization approach.

It is worth noticing that, as mentioned in the previous section, there is an inherent relationship between applying a weighting profile to the classifiers themselves and applying a weighting profile to the rank items that they produce as output. Classifiers with high accuracy rates, i.e., a high confidence value and a proportionally high weighting factor, will produce rank classifications where the first (top) item, i.e., the one with the largest rank weight, is usually correct. This evident correlation will be explained more clearly later on, within the context of "winning coalitions" and their realization in weighted majority games (WMG).

While joint optimization of parameters should be able to provide more generic solutions, it is not yet clear if the joint approach produces better results than the cascaded scheme. In fact, recent studies have shown that the joint optimization does not improve the combination accuracy of the validation data as compared to the two-stage strategy [04]. In most cases, a simple weighted averaging rule upon trained classifiers produces the best results, even in problems of high complexity. This is attributed to the fact that, regardless of the complexity of the initial input space, the classifiers transform it to a highly restricted subspace with dimensionality equal to the number of classes available. Therefore, it is evident why optimized linear solutions, like the WMR model, may be the answer to this problem, especially when robust classifiers like SVM are considered.

1.4 Elements of Game Theory

In principle, the mathematical theory of games and gaming was first developed as a model for situations of conflict. Since the early 1940's, the work of John Von Neumann and Oskar Morgenstern has provided a solid foundation for the most simple types of games, as well as analytical forms for their solutions, with many applications to Economics, Operations Research and Logistics. The "zero-sum" games are able to model situations of conflict between two or more "players", where one's gain is the other's loss and vice versa. Furthermore, if all players are full informed about their opponents' decisions the game is called of "perfect information". Such games are all board games like chess and it has been proven that there is at least one optimal plan of decisions or "strategy" for each player, as well as a "solution" to the game that comes naturally as a result of all players following their optimal strategies. At the game's solution, each player can guarantee that the maximum gain an opponent can gain is kept under a specific minimal limit, defined only by this player's own strategy. Von Neumann and Morgenstern proved this assertion as a theorem called "Minimax" and in the simple case of two opposing players the solution of the game can be calculated analytically as a solution of a 2x2 set of linear equations. The consequences of the Minimax theorem have been thoroughly studied for many years after its proof. As an example, it mathematically proves the assertion that all board games, including the most complex ones like chess, have at least one solution, optimal for both players that can be analytically calculated, at least in theory.

Although the Minimax theorem provided a solid base for solving many types of games, it is only applicable in practice for the zero-sum type of games. In reality, it is common that in a conflict not all players receive their opponents' looses as their own gain and vice versa. In other words, it is very common a specific combination of decisions between the players to result in a certain amount of "loss" to one and a corresponding "gain", not of equal magnitude, to another. In this case, the game is called "nonzero-sum" and it requires a new set of rules for estimating optimal strategies and solutions. As each player's gains and losses are not directly related to the opponents', the optimal solution is only based on the assertion that it should be the one that ensures that the player has "no regrets" when choosing between possible decision options. This essentially means that, since each player is now interested in his/her own gains and losses, the optimal solution should focus on maximizing each player's own "expectations" [11]. The Minimax property can still be applied in principle when the single most "secure" option must be identified, but the solution of the game has now a new meaning.

During the early 1950's, John Nash has focused primarily on the problem of finding a set of "equilibrium points" in nonzero-sum games, where the players eventually settle after a series of competitive rounds of the game. In strict mathematical terms, these equilibrium points would not the same in essence with the Minimax solutions, as they would come as a result of the players' competitive behavior and not as an algebraic solution of the games' mathematical formulation. In 1957 Nash has successfully proved that indeed such equilibrium points exist in all nonzero-sum games*, in a way that is analogous to the Minimax theorem assertion. However, although the Nash theorem ensures that at least one such "Nash equilibrium" exists in all nonzero-sum games, there is no clear indication on how the game's solution can be analytically calculated at this point. In other words, although a solution is known to exist, there is no closed form for nonzero-sum games until today.

It should be noted that players participating in a nonzero-sum game may or may not have the same options available as alternative course of action, or the same set of options may lead to different payoffs between the players. When players are fully exchangeable and their ordering in the game makes not difference to the game setup and its solutions, the game is called "symmetrical". Otherwise, if exchanging players' position does not yield a proportional exchange in their payoffs, then the game is called "asymmetrical". Naturally, symmetrical games lead to Nash equilibrium points that appear in pairs, as an exchange between players creates its symmetrical counterpart.

But the Nash equilibrium points are not always the globally optimal option for the players. In fact, the Nash equilibrium is optimal only when players are strictly competitive, i.e., when there is no chance for a mutually agreed solution that benefits them more. These strictly competitive forms of games are called "non-cooperative games". The alternative option, the one that allows communication and prior arrangements between the players, is called a "cooperative game" and it is generally a much more complicated form of nonzero-sum gaming. Naturally, there is no option of having cooperative zero-sum games, since the game structure itself prohibits any other settlement between the players other than the Minimax solution.

The problem of cooperative or possibly cooperative gaming is the most common form of conflict in real life situations. Since nonzero-sum games have at least one equilibrium point when studied under the strictly competitive form, Nash has extensively studied the cooperative option as an extension to it. However, the possibility of finding and mutually adopting a solution that is better for both players than the one suggested by the Nash equilibrium, essentially involves a set of behavioral rules regarding the players' stance and "mental" state, rather than strict optimality procedures [11]. Nash named this process as "bargain" between the players, trying to mutually agree on one solution between multiple candidates within a "bargaining set". In practice, each player should enter a bargaining procedure if there is a chance that a cooperative solution exists and it provides at least the

* Seminal works by C. Daskalakis & Ch. Papadimitriou in 2006-2007 and on have proved that, while Nash equilibria exist, they may be unattainable and/or practically impossible to calculate due to the inherent algorithmic complexity of this problem; see e.g. "The Complexity of Computing a Nash Equilibrium", 38th ACM Symposium on Theory of Computing, STOC 2006.

same gain as the best strictly competitive solution, i.e., the best Nash equilibrium. In this case, if such a solution is agreed between the players, it is called "bargaining solution" of the game.

As mentioned earlier, each player acts upon the property of no regrets, i.e., follow the decisions that maximize their own expectations. Nevertheless, the game setup itself provides means of improving the final gain in an agreed solution. In some cases, the bargaining process may involve the option of "threats", that is a player may express the intention to follow a strategy that is particularly costly for the opponent. Of course, the opponent can do the same, focusing on a similar "threat". This procedure is still a cooperative bargaining process, with the threshold of expectations raised for both players. The result of such a process may be a mutually "deterring" solution, which in this case is called a "threating solution". Nash has formulated all these bargaining situations into a set of relatively logical axioms, under which a solution (equilibrium) exists. As in the general case of non-cooperative games, Nash's "bargaining theorem" does not provide analytical means of finding such solutions.

The notion of "bargaining sets" and "threat equilibrium" is often extended in special forms of games that include iterative or recursive steps in gaming, either in the form of multi-step analysis (metagames) or focusing on the transitional aspects of the game (differential games). Modern research is focused on methods that introduce probabilistic models into games of multiple realizations and/or multiple stages [11].

1.5 Coalitions, Stable Sets and Indices of Power

Nash's work on the "Nash equilibrium" and "bargaining theorem" provides the necessary means to study n-person non-cooperative and cooperative games under a unifying point of view. Specifically, a nonzero-sum game can be realized as a strictly competitive or a possibly cooperative form, according to the game's rules and restrictions. Therefore, the cooperative option can be viewed as a generalization to the strictly competitive mode of gaming.

When players are allowed to cooperate in order to agree on a mutually beneficial solution of game, they essentially choose one strategy over the others and bargain this option with all the others in order to come to an agreement. For symmetrical games, this is like each player chooses to join a group of other players with similar preference over their initial choice. Each of these groups is called a "coalition" and it constitutes the basic module in this new type of gaming: the members of each coalition act as cooperative players joined together and at the same time each coalition competes over the others in order to impose its own position and become the "winning coalition". This setup is very common when modeling voting schemes, where the group that captures the relative majority of the votes becomes the winner.

Coalition Theory is closely related to the classical Game Theory, especially the cooperating gaming [11]. In essence, each player still tries to maximize its own expectations, not individually any more but instead as part of a greater opposing term. Therefore, the individual gains and capabilities of each player is now considered in close relation to the coalition this player belongs, as well as how its individual decision to join or leave a coalition affects this coalition's winning position. As in classic nonzero-sum games, the notion of equilibrium points and solutions is considered under the scope of dominating or not in the game at hand. Furthermore, the theoretical implications of having competing coalitions of cooperative players is purely combinatorial in nature, thus making its analysis very complex and cumbersome. There are also special cases of collective decision schemes where a single player is allowed to "abstain" completely from the voting procedure, or prohibit a contrary outcome of the group via a "veto" option.

In order to study the properties of a single player participating in a game of coalitions, it is necessary to analyze the wining conditions of each coalition. Usually each player is assigned a fixed value of "importance" or "weight" when participating in this type of games and each coalition's power is measured as a sum over the individual weights of all players participating in this coalition. The coalition that ends up with the highest value of power is the winning coalition. Therefore, it is clear that, while each player's power is related to its individual weight, this relation is not directly mapped on how the participation in any arbitrary coalition may affect this coalition winning position. As this process stands true for all possible coalitions that can be formed, this competitive type of "claiming" over the available players by each coalition suggests that there are indeed configurations that marginally favor the one or the other coalition, i.e., a set of "solutions". The notion of solution in coalition games is somewhat different from the one suggested for typical nonzero-sum games, as it identifies minimal settings for coalitions that dominate all the others. In other words, they do not identify points of maximal gain for a player or even a coalition, but equilibrium points that determine which of the forming coalitions is the winning one. This type of "solutions" in coalition games is defined in close relation to "domination" and "stability" of such points and they are often referred to as "the Core". Von Neumann and Morgenstern have defined a somewhat more relaxed definition of such conditions and the corresponding solutions are called "stable sets" [11]. It should be noted that, in contrast to Nash's theorems and the Minimax assertion of solutions, there is generally no guarantee that solutions in the context of the Core and stable sets need to exist in an arbitrary coalition game.

The notion of the Core and stable sets in coalition gaming is of vital importance when trying to identify the winning conditions and the relative power of each individual player in affecting the outcome of the game. The observation that a player's weight in a weighted system may not intuitively correspond to its voting "power" goes back at least to Shapley and Shubik (1954). For example, a specific weight distribution to the players may make them relatively equivalent in terms of voting power or, while only a slight variation of the weights may render some of them completely irrelevant on determining the winning coalition [06]. Shapley and Shubik (1954) and later Banzhaf and Coleman (1965, 1971) suggested a set of well-defined equations for calculating the relative power of each player, as well as each forming coalitions as a whole [11]. The "Shapley index of power" is based on the calculation of the actual contribution of each player entering a coalition, in terms of improving the coalition's gain and winning position. Similarly, the "Banzhaf index of power" calculates how an individual player's decision to join or leave a coalition results in a winning or loosing position for this coalition, accordingly. Both indexes are basically means of translating each player's individual importance or weight within the coalition game into a quantitative measure of power in terms of determining the winner. While both indices include combinatorial realizations, the Banzhaf index is usually easier to calculate, as it is based on the sum of "shifts" on the winning condition a player can incur [07]. Furthermore, its importance in coalition games will be made clearer later on, where the Banzhaf index will come as a direct result when calculating the derivatives of a weighted majority game.

1.6 Collective competence and the Condorcet criterion

The transformation of cooperative *n*-person games into coalition games essentially brings the general problem closer to a voting scheme. Each player casts a vote related to its own choice or strategy, thus constituting him/her as a member of a coalition of players with similar choices. The coalition that gains more votes becomes the winner.

Condorcet (1785) was the first to address the problem of how to design and evaluate an efficient voting system, in terms of fairness among the people that participating in the voting process, as well as the optimal outcome for the winner(s). This first attempt to create a probabilistic model of a voting body is known today as the "Condorcet Jury Theorem". In essence, this theorem says that if each of the voting individuals is somewhat more likely than not to make the "better" choice between some pair of alternative options, and if each individual makes its own choice independently from all the others, then the probability that the group majority is "correct" is greater than the individual probabilities of the voters. Moreover, this probability of correct choice by the group increases as the number of independent voters increases [07]. In practice, this means that if each voter decides independently and performs marginally higher than 50%, then a group of such voters is guaranteed to perform better than each of the participating individuals. This assertion has been used in Social sciences for decades as a proof that decentralized decision making, like in a group of juries in a court, performs better than centralized expertise, i.e., a sole judge.

The Condorcet Jury Theorem and its implications have been used as one guideline for estimating the efficiency of any voting system and decision making in general. The study of effects like diversely informed voters or situations of conflicting interests have provided several aspects of possible applications in social and economical models. In the context of collective decision-making via voting schemes, the theorem provides a mean to test the "fairness" and effectiveness of such a system, as it usually constitutes the outcome that yields the best possible degree of consensus among the voting participants. Specifically, the interest is focused on how aggregate competence of the whole voting group, measured by the probability of making a correct collective decision, depends on the defining properties of the decision-making process itself, such as different coalition sizes, team setups and possible overlapping memberships. Under this context, the coalition games are studied by applying quantitative measures on "collective competence" and optimal distribution of power, e.g. tools like the Banzhaf or Shapley indices of power. The degree of consistency of such a voting scheme on establishing the pair-wise winner(s), as the Condorcet Jury Theorem indicates, is often referred to as the "Condorcet criterion". This criterion is not the only possible measure of collective competence in a voting scheme, but as it will be explained in the next section, it is very generic and it is directly linked to optimal wBorda models.

1.7 Optimal scoring rules and Condorcet efficiency

Let us consider a typical voting situation where an n-person voting group is required to cast their votes regarding a set of M classes, not as simple class selections but rather in the more general sense of ranking all the available options in a list of strict preference by each of the voters. Clearly there is a fixed set of possible ranking permutations and each voter essentially chooses one of them as his/her vote. The problem is to decide upon the exact combination procedure for these votes, in order to produce a result that exhibits the highest possible degree of consensus between the voters, not only in the first place as in simple majority rules, but throughout the final sorted list. One of the more widely accepted criteria for choosing the exact permutation that best reflects the cumulative will of the voting group is the Condorcet criterion [07]. As it is based exclusively on pair-wise comparisons between the voting options, i.e., every possible pair of subsequent classes in a sorted list, a system that exhibits a high degree of consistency with the Condorcet criterion should provide an aggregate ranking result that represents the best consensus solution. The degree in which such a voting scheme maximizes the consistency with the Condorcet criterion is often called "Condorcet efficiency" of the system.

It is clear that voting systems as the one described above fit the specifications of a wBorda scheme, as an optimized wBorda should also be able to produce an aggregate ranking consistent with the cumulative will of the complete voting group. Therefore, the next obvious question is whether there is a way to design such a wBorda voting scheme that maximizes the Condorcet criterion. The direct answer is that this problem is classified as NP-complete by nature, which means that due to its combinatorial nature it is not possible to be solved with algorithms of polynomial complexity. In fact, this is the reason why simple realizations of the weighting profile or even a simple majority rule is often applied in practice, in order to keep such a system simple and widely accepted by the voting group [06].

Using the notion of Condorcet efficiency of a voting system, the real problem can be focused on the exact scoring rules, often called "weighted scoring rules", than must be applied on each rank of the vote, in order to produce a result that maximizes this criterion. Two scoring systems are particularly worth noting in this context: the plurality voting, where only the top-position rank is awarded with one point (w=1) and all the other positions with nothing (w=0), and the classic Borda count, where the top-position rank receives maximum points (w=1), the bottom-position rank the minimum (w=0) and all intermediate positions a value proportional to the exact rank [08].

Under the scope of weighted scoring rules and the more general theory of weighted order statistics, it has been proven theoretically that for three classes and n-person voting group, the scheme that maximizes the probability that any pair-wise contest in the final ranked list will be consistent with the pairwise majority rule, is in fact the Borda count model. This means that if the system should sort the list of winners in a way that is consistent with the pairwise majority rule, then the Borda count scheme can accomplish this. Similarly, if the majority criterion is instead replaced with the more generic Condorcet criterion, a specific wBorda model with non-uniform weighting profile is the optimal solution in this case [08]. This diversity between the two optimized scoring rules comes from the fact that the Condorcet criterion suggests a stricter rule of optimality than the simple majority and this is why the existence of a Condorcet winner is not always guaranteed [05]. In [08], a geometrical realization of the wBorda design process has been suggested and results have shown that for three candidate classes the middle-position rank has to be assigned with a weight value of less than zero, i.e., different than the classic Borda rule, in order to obtain maximum Condorcet efficiency. Current theoretical results are not sufficient to support any generic statement regarding the design properties of such optimal schemes. Furthermore, the high degree of complexity prohibits the analytical theoretical study of such systems, in order to produce generic constructive methods for wBorda of maximum Condorcet efficiency.

1.8 Majority functions and Banzhaf numbers

Let us now focus in the case of dichotomous choice situations where there are only two candidate classes to vote for. This is clearly a simpler problem in terms of pattern recognition, since an input has to be classified in either one of the two available choices, "true" or "false", "positive" or "negative", "benign" or "malignant". As there are only two available class choices, the Borda count, which is used when class rankings are considered, reduces to the simple class selection scheme and the resulting majority rule that is used in practice.

Dichotomy choice situations have been the center of many analytical probabilistic studies within the scope of voting systems, primarily because of the simplicity of the probabilistic formulations of such models. A dichotomy choice can be easily modeled as a binomial distribution and the combined result in a *n*-player voting game becomes a product of the corresponding "skill" probabilities of the individual players. Then, a combined decision rule can be formulated according to the aggregate choice that is supported by the largest combined probability, i.e., the choice that corresponds to the maximum degree of consensus among the players. These decision rules are called "majority functions" and in the special case where all players and classes are accompanied with the same weight, the simple majority rule emerges as a natural result.

In terms of coalition games, the simple majority functions are modeled in a way which is much more trivial than the generic wBorda scheme that was presented in previously. Again, if special voting situations like abstains and veto are not allowed, the choice of either one of the two available classes automatically assigns every participating player into one of the two possible coalitions. Classes may or may not be weighted with the same "value" or "importance", while the players themselves may be accompanied with a weight or "reliability" value too. In any case, if a linear rule is applied to accumulate and combine all individual choices in order to make a final collective estimation, a weighted majority rule emerges. The threshold of the majority decision may also be altered in a way that requires not only relative majority, but a majority value higher than a specific decision threshold. In practice, this means that a bias may be used in the weighted majority function in order to ensure that the final majority outcome is valid only if it attains a specific confidence level.

The analysis of the majority functions is often restricted to the non-weighted case, as they are much easier to analyze within the scope of classic probabilistic theory. In fact, this special case of majority functions can be easily related to the Banzhaf power index [07]. Specifically, if the collective efficiency is to be calculated as a function of the individual "skill" probabilities of the players, the partial derivatives of the majority functions against these probabilities are calculated. These derivatives essentially estimate the number of "shifts" that a player with a specific skill probability can cause in the winning position of any winning coalition, i.e., it is exactly what the Banzhaf power index stands for. This assertion can be extended for the weighted majority functions as well, in a slightly more complex probabilistic form. Based in this very important conclusion, it is possible to translate many of the properties of coalition games into properties that are directly linked with each player's skill.

The first and extremely important conclusion from studying the Banzhaf numbers as the derivatives of a majority function is the fact that the maximum of these derivatives should point to the configuration where maximum Banzhaf power occurs for all the voting players. Indeed, it can be easily proven that maximum Banzhaf values correspond to individual skill probabilities close to 0.5 for all voting players. That is, the vote of each voting member reaches its maximum "value" when all players have the same average skill for making correct estimations. Interestingly enough, that is exactly what the Condorcet Jury Theorem suggests in a more generic way. An electorate system with high Banzhaf number corresponds to a high level of collective competence, which in turn is obtained for a high level of "democracy" in the sense of an equitable distribution of decisional power among the voters [07]. As high Banzhaf numbers indicate a high degree of democracy among the voting members, the decentralized option for making collective decisions is, again, asserted as the optimal way – this time in a more strict mathematical statement. This is perhaps a sufficient justification for using ensembles of independent experts with only moderate efficiency, rather than one single expert of the very high efficiency.

But what about the distribution of power within the voting group itself? Using the same formulation of Banzhaf numbers as derivatives of the corresponding majority functions, Berg [07] and Taylor and Zwicker [06] have stated some very interesting results regarding the optimal structure and distribution of voting power of such a system. Specifically, it has been proven that in any non-weighted majority function, the sub-division of the players into teams that subsequently take part in a second-stage indirect voting scheme results in a loss of individual decisional power, in terms of how a single player can affect the final result with his/her vote [07]. This essentially means that it is better to have *K*x*N* voters in one voting group than to split them into *K* teams of *N*/*K* members each for voting via *K* representatives. Boland (1989) has proved that although the Condorcet Jury Theorem also stands true for indirect voting systems, splitting a voting group into teams essentially lowers the majority threshold necessary for a coalition to become a winning one, thus reducing the probability for a collectively correct decision [07]. This implies that for the same overall voting group size and individual skill probabilities p>0.5, an indirect voting scheme always has less reliability than a corresponding direct system, which in turn favors combination methods with the least possible integration stages. This effect is more evident for systems that employ a relatively large number of voters, rather than small-sized systems where this difference is expected to be minimal [07].

In terms of team sizes versus number of teams, Boland (1989) has also proved that in an indirect system, a large number of small teams are collectively more effective than a small number of large teams. In the extreme case where each team includes only one voter, the indirect system becomes direct, i.e., with no representatives, which is the strictly more efficient voting structure as noted earlier. Special studies have been carried out for situations of teams with unequal number of members or for overlapping memberships. Again, as in the case of single voting players, it has been proven that the collectively more efficient choice is splitting the voters into teams of equal size and distinct memberships, i.e., in way that favors equal distribution of voting power in every case [07].

The overlapping membership case can be viewed as a situation where some of the players are allowed to participate in more than one representative team, in other words to affect the final outcome with more than one votes. This is essentially equivalent to having an increased reliability or weight assigned to these players. As mentioned earlier, the weighted majority functions are a generalized version of the ones that have been studied within this scope. As it turns out, the collective decision efficiency may benefit from such an overlapping membership, i.e., a weighted voting scheme, only when the players with multiple votes exhibit a skill probability higher than a specific threshold. This higher than the average skill level *P* of the rest of the voting players threshold and it depends on this average skill level *P* and the total number of *N* voters, but not on the number of *K* teams [07]. This conclusion favors the application of weighted versus non-weighted majority functions in theory, but it does not specify an optimal way to find out which players to favor, in other words how to calculate these weights. This issue will be addressed later on within the context of weighted majority games (WMG). Generally speaking, analyses in terms of voting games and distribution of power are not common in the literature on the Condorcet Jury Theorem. Austen-Smith and Banks (1996), as well as Berg [07], stress the importance of a game-theoretic approach to collective decision making.

1.9 Weighted Majority Games and Weighted Majority Rules

In most cases, majority functions that are employed in practice very simplistic when it comes to weighting distribution profile or they imply a completely uniform weight distribution. However, a specific weighting profile usually produces better results, provided that is simple enough to be applied in practice and attain a consensus in accepting it as "fair" by the voters. Taylor and Zwicker (1991) have defined a voting system as "trade robust" if an arbitrary series of trades among several winning coalitions can never simultaneously render them losing [06]. Furthermore, they prove that a voting system is trade robust if and only if it is weighted. This means that, if appropriate weights are applied, at least one winning coalition can benefit from this procedure.

As an example, institutional policies usually apply a non-uniform voting scheme when it comes to collective board decisions. This is often referred to as the "inner cabinet rule". In a hospital, senior staff members may attain increased voting power or the chairman may hold the right of a tie-breaking vote. It has been proven both in theory and in practice that such schemes are more efficient than simple majority rules or any restricted versions of them like trimmed means.

Nitzan and Paroush (1982) have studied the problem of optimal weighted majority rules (WMR) extensively and they have proved that they are indeed the optimal decision rules for a group of decision makers in dichotomous choice situations [09]. This proof was later (2001) extended by Ben-Yashar and Paroush, from dichotomous to polychotomous choice situations [13]; hence, the optimality of the WMR formulation has been proven theoretically for any n-label voting task.

A WMR is in fact a realization of a weighted majority game (WMG), where a group of players with arbitrary skill levels form coalitions of similar interests but different estimations. The WMGs are a well-known subgroup of coalition games where there are only two possible coalitions, each related to one of the two class options available. In this form of gaming, there is no need of class ranking schemes as in wBorda, thus the classification problem reduces to the optimal design of a combination rule between two extreme options. These optimal combination rules are similarly called WMR and the proof that they are linear in nature limits the problem to the estimation of an optimal (non-negative) weighting profile for the voters.

The weight optimization procedure has been applied experimentally in trained or other types of combination rules, but analytical solutions for the weights is not commonly used. However, Shapley and Grofman (1984) have established that an analytical solution for the weighting profile exists and it is indeed related to the individual player competencies or skill levels [09]. Specifically, if decision independency is assumed for the participating players, the optimal weights in a WMR scheme can be calculated as the *log-odds* of their respective skill probabilities, i.e., $W\text{i}=\log(O\text{i})=\log(P\text{i}/(1-P\text{i}))$. Interestingly enough, this is exactly the solution found by analytical Bayesian-based approaches in the context of decision fusion of independent experts [12]. The optimality assertion regarding the WMR, together with an analytical solution for the optimal weighting profile, provides an extremely powerful tool for designing theoretically optimal collective decision rules. Usually, the winning coalition in a WMR is the one that accumulates the relative weighted majority, which is more than half the sum of weights. If a bias is also applied as a confidence threshold, then the simple weighted majority rule becomes the "cogent" weighted majority rule. In this case, a region of "stalemate" or "no-decision" is created and the existence of a winning coalition in a WMR is not guaranteed.

There is an equivalence relation on the WMRs, whereby two WMRs are equivalent if they produce the same decision function, i.e., the same outcome for each decision profile produced by the participating players. Therefore, from all the possible realizations of a WMR of a given dimension, it is possible to identify a closed set of unique WMRs that are able to produce all the possible combination outcomes with a normalized version of the weights [09]. As an example of such closed set of solutions, the unique WMRs for n=4 voting players are: S1={1,0,0,0}, S2={2,1,1,1} and S3={1,1,1,0}. The S1 and S3 solutions clearly implement the restricted majority rules with odd number of voters, while S2 implements a simple majority rule with a tie-breaking option for one of the players. It can be proven that all other realizations of 4-player WMRs can be mapped into one of these three unique WMRs.

The calculation of such a set of WMRs is cumbersome and it is generally an NP-complete problem. However, there have been analytical studies for WMRs of dimension up to seven players. For example, in the set of all 84 WMGs between five players, only 7 of them are not transformations of others. Von Neuman and Morgenstern (1944) identified the 21 unique WMGs for six players, while

Isbel (1959) and Fishburn and Gehrlien (1977) identified the 135 unique WMGs for seven players. Karotkin (1994) coded a special algorithm to identify the WMRs or the WMGs for any group size and proposed a graph-based method for illustrating the decision-based "closeness" of such WMRs in a three-dimensional space, called the "network of WMR" [09]. Using directed edges, the network of WMR can identify the node that is optimal for a given set of players' skill probabilities profile, i.e., the WMG solution that is in fact the one that the log-odds calculates for the corresponding weights. The practical use of such a graph-based representation of WMRs is that it can suggest optimal substitutions of the theoretically optimal WMR with best sub-optimal simplified realization [10]. This is extremely important in real life problems where simple collective decision rules with straightforward application are needed.

1.10 Weighted Majority Rules and Condorcet efficiency

The efficiency of a WMR is defined as the likelihood that it will resolve in the correct choice, given the skill probabilities of the participating voters. These likelihood functions are quite difficult to calculate in practice due to the fact that the number of possible decision profiles is a combinatorial enumeration problem. As a result, it is also difficult to compare relative efficiencies between different WMRs. However, since a WMR is proved to be the optimal structure in the sense of collective decision competence, the corresponding weighting profile that is optimal for a given set of skill probabilities of the participating voters should be the actual realization of the theoretically most efficient voting scheme.

As mentioned earlier, in the case of multiple class options where class rankings are necessary, it is possible to find a wBorda scheme that maximizes the Condorcet efficiency of such a voting system, although this problem is generally NP-complete. When this setup is reduced into the dichotomous choice situation where there are only two classes available, this model becomes the theoretically optimal formulation of the WMRs. However, in this case there is an analytical solution for the optimal weighting profile that is not NP-complete, although the complete enumeration of all the unique WMRs that can be implemented in practice is of that complexity. As a result, the next question is at what degree a WMR can be viewed as an optimal solution to a WMG in the Condorcet sense.

To answer this question, the notion of "bias" or "confidence threshold" in a WMR has to be reviewed under a new perspective. In wBorda schemes, each class ranking position is scored with a specific weight and the corresponding scoring rule is considered optimal in the Condorcet sense if it maximizes the Condorcet criterion. Similarly, in the case of two class problem, the simple choice between the one or the other choice essentially implies a similar preference ranking regarding the classes. Therefore, both the first (proposed) class choice and the second (rejected) class choice can be assigned with a scoring value, i.e., a weight, that can be incorporated into the standard WMR formulation. These scoring values are not a subject of the players' skill probabilities, since the efficiency of each player affects only the corresponding weight it receives within the WMR function, not the scoring result of selecting or rejecting a class. In a sense, the scoring of class selection or rejection adds a weighting scheme in the second dimension of the WMRs, that of the classes.

Using this new more generalized formulation of WMRs, it can be easily proven through linear transformations that this class scoring essentially produces a "positive" or "negative" bias to the accumulated result of the standard WMR. Therefore, a decision threshold can be shifted towards the one or the other class accordingly, based not only on which exactly of the players selected it but also the mere (weighted) count of the times it was proposed or rejected by all the players. If all players' votes are weighted exactly the same, then this new scheme is a two-class realization of a wBorda count model. But the wBorda model has been already proven as adequate for providing an optimal voting realization in the Condorcet sense. In fact, adding a bias to both classes according to their selection count is mathematically equivalent to setting the WMR decision threshold at a value other than half the sum of the weights, i.e., "biased" towards one of the classes. Not surprisingly, this new generalized version of the WMR can also be considered adequate for implementing voting schemes that maximize the Condorcet criterion, i.e., the exhaustive pair-wise ranking contest between the coalitions.

The assertion that WMRs are optimal realization of combinations schemes in dichotomy choice situations has some extremely significant implications in the way the WMRs can be used as a unified template model for creating optimal collective decision systems. These linear formulations of WMGs are optimal in the MCE sense but additionally they can be designed to be optimal in the Condorcet sense. Dichotomy choice situations are simple enough so that a Condorcet winner, that is the overall top-ranked class, is also the majority winner, which is simple the class that received the most votes. If weights are applied to the players, then the simple majority rule becomes a weighted one. If scores are also applied in the "support" or "reject" options (ranks) of the classifications, then a two-class wBorda count model can be realized in a way that maximizes the Condorcet criterion. In practice, this second case is equivalent to imposing a collective decision threshold other than half the sum of the weights. While the players' optimal weighting profile in the WMR solves the problem of how to combine their individual decisions in an MCE-optimal way, the class scoring provides the means the design the voting system in a way that is also optimal in the Condorcet sense. The conditions under which these two properties can be satisfied simultaneously remains an open issue.

Part II – Applications to Pattern Recognition†

2.1 Basic framework

In part-I of this report it was suggested that classifiers providing "hard" decisions, or other types of classifiers with translated output to a set of distinct choices, can be used as the basis of a general combination procedure for providing a collective decision system. Furthermore, it was described how this model can be effectively fused into a game-theoretic approach of the combination problem that finally leads to coalition games and collective decision theory. This section describes how these models can be realized and implemented as practical systems within the scope of Pattern Recognition.

In order to combine classifier outputs in an optimized way, first it is necessary to convert their posterior accuracy probabilities into quantitative measures of evidence regarding their past performance. For classifiers of "hard" decisions, the confusion matrix is a very descriptive and perhaps the most practical way to do this. Specifically, the confusion matrix itself can be translated into class rankings and conditional probabilities estimations, as it was suggested earlier. Furthermore, the use of the confusion matrix is completely compatible with any extension that introduces the notion of "risk" into the classification process.

Decision-critical applications, like in medical diagnostics, require strict distinction between the various cases of correct and incorrect predictions. This means that a specific weight is assigned for every such classification case, in the form of a positive "gain" for correct predictions or a negative "loss" for misclassifications. When combined with the corresponding posteriori probabilities of the classifier, it is possible to calculate the expected statistical "risk", i.e., the average gain or loss that this particular classifier can produce. If the classifier is trained from the start by applying optimization criteria based on risk factors, rather than simply the classification accuracy, then the process is a "risk-based" rather than "error-based" training of the classifier. Not all classifier architectures are fit to be implemented as risk-based models, primarily due to the fact that the introduction of risk factors within the feedback process of the training may result in severe instability and failure. However, the notion of risk embodies a much more generalized viewpoint of the classification problem and it is very important in real-world applications.

Using risk-based models for the classifiers, the game-theoretic approach of collective decision systems becomes much more comprehensible. The efficiency of each participating "player" is now measured not simply in a sense of absolute accuracy but in the scope of average "gain" in each run of the game. Therefore, every combination scheme also embodies the same notion of maximizing the collective "gain" or, equivalently, minimizing the collective "loss", by employing an optimal combination rule. This risk-based approach is also valid for a coalition's winning stance against the others, as well as the expected payoff from the whole game, since a winning coalition's gain coincides with the overall gain

† Comments in this section are subject of own study and experimental verification, conducted during the author's PhD work, 2001-2008 and on. Almost all of the proposed items have been addressed, experimentally tested and subsequently published in various conference, journal and open-access papers. For detailed description of the theoretical and practical aspects of applying these ideas in the context of novel classifier combination architectures, see e.g. [14-16]:

- "A Game-Theoretic Approach to Weighted Majority Voting for Combining SVM Classifiers", Harris Georgiou, Michael Mavroforakis, Sergios Theodoridis. Int. Conf. on ANN (ICANN), 10-13 September 2006 @ Athens, Greece. Ref: S.Kollias et al. (Eds): ICANN 2006, Part I, LNCS 4131, pp. 284-292, 2006.
- "A game-theoretic framework for classifier ensembles using weighted majority voting with local accuracy estimates", H. Georgiou, M. Mavroforakis, arXiv.org preprint (en)(arXiv:1302.0540v1 [cs.LG]).
- "Algorithms for Image Analysis and Combination of Pattern Classifiers with Application to Medical Diagnosis", H. Georgiou, PhD thesis summary (en)(arXiv:0910.3348v1 [cs.CV]).

of the collective decision rule. Therefore, the formulation of this type of gaming under the scope of WMGs and the corresponding WMRs comes very naturally.

Since the WMR have been proven as the optimal combination rules in dichotomy choice situations, it is very interesting to examine the conditions under which these optimality assertions stand true for trained classifiers in the place of players. It is expected that, as these classifiers are more or less dependent with each other due to similar architectures or training datasets, the design of optimal WMRs for combining them can not be realized in completely closed form. Instead, the calculation of the exact weighting profile requires the exploitation of various statistical and structural properties of the feature space, as well as the correlation of input patterns and classes. Most ensemble techniques exploit these properties by enhancing classification regions of special interest, like points close to the decision boundary. This issue has been noted earlier within the scope of voting systems, specifically in relation to the diversity between the experts. In practice, many ensemble methods that employ maximization of diversity essentially increase the degree of independency between the participants. Since an increased level of independency provides the means for a collectively efficient decision, it is not surprising to see that the results from the Pattern Recognition viewpoint coincide with the ones inferenced by the game-theoretic approach, where Banzhaf and Shapley indices of power can be considered as measures of diversity among the participants.

2.2 Adaptive realizations of WMRs

As it was mentioned previously, assumptions of complete independency between the classes and their corresponding coalitions in WMGs, as well as between the experts are never completely true. Thus, it is necessary to enhance the combination process in a sense that takes into account these types of correlations. Tresp and Taniguchi (1996) have suggested that a combination scheme which uses a fixed or a weighted majority rule should exploit the properties of the statistical distributions of the classes at hand. Using a Gaussian approximation, they have shown how the efficiency of such a combination rule can be improved if the mean and variance of each class are used when calculating the parameters (weights) of the combination rule. However, the standard WMR approach for optimal combination of experts does not include such an adaptive scheme. Furthermore, the variance-based weighting method of Tresp and Taniguchi impose further assumptions and restrictions to the distributions of the classes, which can be non-Gaussian in general.

Instead of employing a fixed statistical approximation for the complete class distribution, a new fully adaptive approach can be designed on a lower level. Specifically, since the distribution of each class, i.e., the topological couplings between the individual training samples of the class, affects the exact weighting profile of the classifiers that is optimal in some error-based or risk-based criterion, then a topological measure of "closure" between any arbitrary pair of samples should be used instead of the statistical approximation of their distribution. In other words, instead of checking how well an unclassified sample fits the statistical distribution of the one or the other class, it should be checked under a criterion that measures how close it is with the identified members of each class, preferably with the most representative ones. As in the case of the variance-based method, this measure should be used as a quantitative guideline regarding the degree of "responsibility" that each class manifests over this particular point in the feature space. In terms of classification, it is a statistical method of measuring how much a class is accountable for this new sample, but in a more invariant way than that of using Gaussian approximations for the class' distributions.

When this fully adaptive version is adopted within the scope of the WMRs, each classifier's posterior probability or "skill", already known from the training process, can be adjusted according to how "far" or "close" an unclassified sample lies with respect to the already known members of this class. From a clearly topological point of view, this is a way to minimize the "structural risk" of the classification problem by introducing a bias or "preference" towards the class that seems to be more responsible for this region of the feature space. Although this is what each of the participating classifiers does on its own, this adaptation process essentially adjusts the process of evaluating the optimal WMR to the non-symmetrical structural properties of the feature space. Therefore, it is also expected that the weighting profile calculated for the WMR design would be optimal, not in a global but rather in a more local sense.

2.3 A link with the theory of SVM

SVM architectures provide the necessary foundations for a theoretical sound framework of optimal classifiers. The use of special form of kernel functions essentially makes them equivalent almost to any type of linear and non-linear pattern classification model. But the solid theoretical background of this type of classifiers makes them ideal in situations where their performance and consistency is required for studying collective decision rules.

The linear form of the WMRs makes the combination process very simple, not only in terms of calculating the final outcome of the group decision, but also in the scope of statistical properties of this decision rule in relation to each classifier's own properties. SVM theory states that the structure of the SVM classifier permits the linear transformation of a number of kernel functions into one combined linear form. Furthermore, if each of the kernel functions is well-defined under the typical constraints for SVM kernels, then their linear combination is also a well-defined SVM kernel function. This assertion is of extreme importance when viewed under the perspective of WMGs and WMRs. In essence, if all participating classifiers are assumed to be SVM realizations, then the optimal WMR itself defines a new compound SVM kernel, i.e., an SVM meta-classifier.

This conclusion is an adequate justification on why such a combination rule does not need to be more complex than a linear transformation of each expert's assessment: if every expert is of adequate skill and acts independently from the others. This means that such an expert can moderately efficient in the complete feature space or, alternatively, well-adapted to only a part of the complete feature space. In the first case, the WMR is optimal in the global sense in a way that combines the group of experts in the most promising manner, while the second case corresponds to the fully adaptive WMR realization that was proposed in the previous section.

It should be noted that, although SVM classifiers are generally design for "hard" decision classifications in dichotomy choice problems, it is not difficult to design a set of SVM classifiers that are "specialized" in one of $N>2$ classes if the problem requires it – this is essentially the one-versus-all classification mode when applying binary pattern classifiers in multi-class tasks. Furthermore, their internal structure that is based on support vectors, i.e., class members that primarily define the classification outcome, is well-suited for the design of robust topological measures of "closure" between a class and a new unclassified sample, based on distance transformations from the support vectors of each candidate class. In this sense, even a two-class problem that is solved by a single SVM classifier can be viewed as a coalition game in the form of WMG, solved by an optimal WMR with weights and bias proportional to the class' distribution characteristics, and an SVM kernel function that effectively transforms the original non-linear feature space to a linear space of higher dimension that can be solved by this WMR.

2.4 Conclusion and further work

The material presented briefly in part-I of this report clearly define a solid background for a game-theoretic approach to the problem of classifier combination in its general form. A set of theoretic formalizations lead to some very intuitive and simple solutions to this problem in the general sense, especially in dichotomy choice situations.

An extension of these theories to the area of Pattern Recognition can be easily inferenced. Specifically, there are three main issues of special interest:

1. The introduction of a theoretically solid model for using transformations of posterior probabilities of the classifiers, e.g. by the confusion matrices, in combination with the general framework of risk minimization, either in a post-training sense or within the training process itself (risk-based training).
2. The formulation of a complete and fully adaptive realization of the WMR model that incorporates the non-symmetrical properties of the underlying feature space, when calculating the optimal weighting profile for the combination rule.
3. The study of theoretical and practical implications of introducing SVM classifier architectures as voting players in a WMG, primarily in the scope of completeness and optimality of such a solution in the general sense.

A study that addresses all these three issues should first focus on the theoretical aspects and formal definitions of any new models and algorithms, and subsequently conduct experimental tests on well-known classification problems where comparative results are available for other typical classifier combination schemes. Based in the theoretical assessment presented in this study, it is expected that such a game-theoretic approach of collective decision, along with the application of SVM classifiers, will produce results of at least the same degree of success as the best ensemble methods available today.

References (short)

[01] J. Kittler, M. Hatef, R. Duin, "Combining Classifiers", IEEE Proc. ICPR '96, pp.897-901.

[02] L. Kuncheva, "A theoretical study on six classifier fusion strategies", IEEE Trans. PAMI, Vol.24, No.2, Feb.2002, pp.281-286

[03] N. Ueda, "Optimal linear combination of neural networks for improving classification performance", IEEE Trans. PAMI, Vil.22, No.2, Feb.2000, pp.207-215.

[04] C.-L. Liu, "Classifier combination based on confidence transformation", Pat.Rec., 38 (2005), pp.11-28.

[05] J.R. Parker, "Rank and response combination from confusion matrix data", Inf.Fusion, 2 (2001), pp.113-120.

[06] A. Taylor, W. Zwicker, "Weighted voting, multicameral representation and power", Games and Econ.Beh., 5 (1993), pp.170-181.

[07] S. Berg, "Indirect voting systems: Banzhaf numbers, majority functions and collective competence", Eur.J.Pol.Econ., 13 (1997), pp.557-573.

[08] D. Cervone, W. Gehrlein, W. Zwicker, "Which scoring rule maximizes Condorcet efficiency?"

[09] D. Karotkin, "The network of weighted majority rules and weighted majority games", Games and Econ.Beh., 22 (1998), pp.299-315.

[10] D. Karotkin, M. Schaps, "The network of weighted majority rules and its geometric realizations", Games and Econ.Beh., 42 (2003), pp.75-90.

[11] G. Owen, "Game Theory", 3rd/Ed., Academic Press, 1995.

[12] L. Kuncheva, "Combing Pattern Classifiers – Methods and Algorithms", Wiley, 2004.

[13] R. Ben-Yashar, J. Paroush, "Optimal decision rules for fixed-size committees in polychotomous choice situations", Social Choice and Welfare, 18 (2001) 737-746.

[14] H. Georgiou, M. Mavroforakis, S. Theodoridis, "A Game-Theoretic Approach to Weighted Majority Voting for Combining SVM Classifiers", Int. Conf. on ANN (ICANN), 10-13 September 2006 @ Athens, Greece. Ref: S.Kollias et al. (Eds): ICANN 2006, Part I, LNCS 4131, pp. 284-292, 2006.

[15] H. Georgiou, M. Mavroforakis, "A game-theoretic framework for classifier ensembles using weighted majority voting with local accuracy estimates", arXiv.org preprint (en)(arXiv:1302.0540v1 [cs.LG]).

[16] H. Georgiou, "Algorithms for Image Analysis and Combination of Pattern Classifiers with Application to Medical Diagnosis", PhD thesis summary (en)(arXiv:0910.3348v1 [cs.CV]).

Section 3: A Game-Theoretic Approach to Weighted Majority Voting for Combining SVM Classifiers

Summary:

A new approach from the game-theoretic point of view is proposed for the problem of optimally combining classifiers in dichotomous choice situations. The analysis of weighted majority voting under the viewpoint of coalition gaming, leads to the existence of analytical solutions to optimal weights for the classifiers based on their prior competencies. The general framework of weighted majority rules (WMR) is tested against common rank-based and simple majority models, as well as two soft-output averaging rules. Experimental results with combined support vector machine (SVM) classifiers on benchmark classification tasks have proven that WMR, employing the theoretically optimal solution for combination weights proposed in this work, outperformed all the other rank-based, simple majority and soft-output averaging methods. It also provides a very generic and theoretically well-defined framework for all hard-output (voting) combination schemes between any type of classifier architecture.

Citation:

"A Game-Theoretic Approach to Weighted Majority Voting for Combining SVM Classifiers", Harris Georgiou, Michael Mavroforakis, Sergios Theodoridis. Int. Conf. on ANN (ICANN), 10-13 September 2006 @ Athens, Greece. Ref: S.Kollias et al. (Eds): ICANN 2006, Part I, LNCS 4131, pp. 284-292, 2006.

A Game-Theoretic Approach to Weighted Majority Voting for Combining SVM Classifiers

Harris Georgiou[1], Michael Mavroforakis[1], Sergios Theodoridis[1]

[1] Dept. of Informatics and Telecommunications, Division of Communications and Signal Processing, University of Athens, Greece
http://www.di.uoa.gr/~dsp
157 84, Panepistimioupolis, Ilissia, Athens, Greece
Harris Georgiou, xgeorgio@di.uoa.gr

Abstract. A new approach from the game-theoretic point of view is proposed for the problem of optimally combining classifiers in dichotomous choice situations. The analysis of weighted majority voting under the viewpoint of coalition gaming, leads to the existence of *analytical solutions* to optimal weights for the classifiers based on their prior competencies. The general framework of weighted majority rules (WMR) is tested against common rank-based and simple majority models, as well as two soft-output averaging rules. Experimental results with combined support vector machine (SVM) classifiers on benchmark classification tasks have proven that WMR, employing the theoretically optimal solution for combination weights proposed in this work, outperformed all the other rank-based, simple majority and soft-output averaging methods. It also provides a very generic and theoretically well-defined framework for all hard-output (voting) combination schemes between any type of classifier architecture.

1 Introduction

1.1 Classifier Combination and Game Theory

In the discipline of collective decision-making, a group of N experts with moderate performance levels are combined in an optimal way, in order to produce a collective decision that is better than the best estimate of each individual expert in the group. According to the famous Condorcet Jury Theorem [1], if the experts' individual decisions are independent and their corresponding estimations are more likely to be correct than incorrect ($p_{correct}>0.5$), then an increase in the collective performance, as a group, is guaranteed when the individual estimations are combined. Moreover, this increase in performance continues to increase asymptotically as the size N of the group increases.

In the case where each expert selects only one out of M available options, the collective group decision can be estimated by the majority voting scheme, i.e., the choice selected is the one gathering the majority of votes. When the simple majority rule is employed, each of the *N* experts acts with the same common interest of reaching the optimal collective decision. However, their individual choices place them in possibly contradicting estimations, with each expert trying *to impose* its decision to the others and to the group. This is a typical competitive situation, which can be modeled by the well-studied theory of non-zero sum competitive gaming in classic Game Theory [2]. In reality, each subgroup of consentient experts essentially represents an opposing assembly to all the other similar subgroups with different consensus of choice. It is clear that this second type of cooperative, instead of purely competitive (per expert), gaming reflects the problem of collective decision-making in the most generic way. Special sections of Game Theory, namely the Coalitions and Stable Sets in cooperative gaming [2], have studied the effects of introducing "weights" to the choice of each expert according to their competencies, in order to optimize the final decision of the group.

1.2 Weighted Majority Games and Weighted Majority Rules

The case of a dichotomous situation, where there are only two symmetrical choices for each expert (i.e., M=2) to vote for, then this restricted form is known as the weighted majority game (WMG) [2]. It has been proven by Nitzan and Paroush (1982) [3] and Shapley and Grofman (1984) [4], that the optimal decision rules, in terms of collective performance, are the weighted majority rules (WMR); this is in fact a different name for the well-known weighted majority voting schemes [5], which are often used in pattern recognition for combining hard-output classifiers. The same assertion has also been verified by Ben-Yashar and Nitzan [6] as the optimal aggregation rule for committees under the scope of informative voting in Decision Theory. Although there is in fact an exponential number of such WMR for each WMG, only a few of them can be proven to be well-defined or *qualified* combination rules and even fewer can be proven to be unique, i.e., not producing exactly the same decision profile with others [7]. For example, in the 2^{32} possible[1] voting games of five experts, there are exactly 85 qualified WMR if only positive integer weights are permitted, of which only seven are unique in terms of their decision profile [7].

In this paper, the notion of modeling dichotomous choice situations for a group of experts via the theory of WMG and WMR is for the first time applied for combining hard-output classifiers. Under the conditional independence assumption, a *closed form solution* for the voting weights in the WMR formula exists and it is *directly linked to each expert's competency*. This optimal weight profile for the voting experts is the log of the odds of their individual competencies [3], [4], [7], [8].

In this paper, this particular type of game-theoretic analytical solution for optimal expert combinations in dichotomous choice situations is tested for the first time against other popular combination schemes. The possibility of having a weighted voting scheme that is based only on the prior capabilities of the experts in the group, as well as on the theoretical assertion that this analytical solution is optimal, in terms

[1] For five experts with two choices each there are 2^5=32 decision profiles, each of which can be generally mapped in any of the two possible outputs of the combination rule. See [7].

of collective competency (at least for all non-trained, i.e., iteratively optimized, weights), is extremely attractive as an option of designing very simple yet effective combination models for an arbitrary pool of classifiers.

2 Datasets and Methods

2.1 SVM Classifier Model

The SVM classifier was used as the base model for creating a pool of classifiers for each combination scheme. Specifically, a geometric nearest point algorithm (NPA) [9], based on the notion of reduced convex hulls (RCH) [10], was used for training a standard SVM architecture with radial-basis function (RBF) as the kernel of the non-linear mapping. In previous studies [11] have shown experimental evidence that optimal combinations of SVM classifiers can be achieved through linear combination rules, i.e., the same category of combination rules examined in this study. In the two averaging combination rules that use the soft-output of the individual classifiers, the distances from the decision boundary were used instead of the (thresholded) hard-output of the SVM classifier, as they are indicative of the corresponding classification confidence [12], [13].

2.2 Datasets and Feature Grouping

In order to assess the performance of each classifier combination method, a number of publicly available test datasets [14], with known single-classifier accuracy rates for this specific SVM training model, were used. These datasets are: 1) Diabetis, 2) Flare-Solar, 3) German, 4) Heart and 5) Waveform.

Each base dataset was randomly separated into a base training set and a validation set of samples. In order to make individually trained classifiers as "independent" as possible, the method of training them in different subspaces was employed. As it has been reported previously, e.g., [13], [15], this is an effective approach towards independence among classifiers. To this end, the training set was partitioned into K distinct segments of feature groupings, i.e., containing only some of the features (dimensions) of the initial dataset. Each group of features was created in a way that satisfied two constraints: (a) each group to be distinct, i.e., no feature is included in two or more groups, and (b) each group to contain a subset of features that can describe the classification task equally well as the other feature groups, i.e., employ a "fair" distribution of the available features into K groups. Satisfaction of the second constraint required a method for ranking all the features in terms of discrimination power against the two classes, as well as their statistical independency to all the other features in the initial training set. Thus, the MANOVA method [16] was used to assign a multivariate statistical significance value to each one of the features and then produce a sorted list based on (the log of) this value.

In order to create a "fair" partitioning of this list into equally efficient segments, features were selected in pairs from the top and bottom positions, putting the currently "best" and "worst" features in the same group. Furthermore, the efficiency of each group was measured in terms of summing the log of the statistical significance value, assigned by MANOVA, of all the features contained in this group. The log was employed in order to avoid excessive differences between the values assigned by MANOVA, thus creating more even subset sums of these values. Essentially, every such pair of features was assigned in groups sequentially, in a way that all groups contained features with approximately equal sum of the log of the values assigned by MANOVA. In other words, the MANOVA-sorted list of features was "folded" once in the middle and then "cut" into *K* subsequent parts of equal sums of log-values, i.e., with every part exhibiting roughly the same sum of the log of the statistical significance values, accompanying each feature included in this part.

Each one of these *K* distinct feature groups was used for training an individual SVM classifier. Thus, each of these *K* classifiers used a different, dimensionally reduced, version of the original (full) training set and therefore learns a totally different classification task.

2.3 Classifier Combination Methods

Nine linear combination rules were examined in this study. Specifically, five hard-output combination methods were employed, namely three standard rank-based methods and two voting-based schemes. These rank-based rules are [8], [13]:

- minimum ("min")
- maximum ("max")
- median ("median")

The two majority rules, including the WMR model, are [8], [13]:

- simple majority voting ("majority")
- weighted majority voting, i.e.:

$$O_{wmr}(\vec{x}) = \sum_{i=1}^{K} w_i D_i(\vec{x}) \quad . \tag{1}$$

where D_i is the hard-output of each of the *K* individual classifiers in the pool, w_i is its assigned weight and O_{wmr} the weighted majority sum. The final hard-output decision D_{wmr} of the WMR is taken against a fixed threshold (T) that defines the decision boundary for the combination rule [7], [8]:

$$D_{wmr}(\vec{x}) = sign\left(O_{wmr}(\vec{x}) - T\right) \quad . \tag{2}$$

Specifically for the weighted majority voting scheme, three different methods for calculating the weight profile were tested for comparative results:

- "direct" weighting profile for WMR ("wmr/direct") [5], [8]:

$$w_i = p_i \quad , \quad p_i = P_i(\theta = \omega_{correct} \mid \vec{x}) \quad . \tag{3}$$

- “odds” weighting profile for WMR (“wmr/odds”) [7], [8]:

$$w_i = \frac{p_i}{1-p_i} \quad , \quad p_i = P_i(\theta = \omega_{correct} \mid \vec{x}) \ . \tag{4}$$

- “logodds” weighting profile for WMR (“wmr/logodds”) [7], [8]:

$$w_i = \log\left(\frac{p_i}{1-p_i}\right) \quad , \quad p_i = P_i(\theta = \omega_{correct} \mid \vec{x}) \ . \tag{5}$$

where w_i is the combination weight assigned for the i-th classifier, p_i is its prior probability for correct classification, measured in the validation set, and θ, ω are the predicted class labels.

Additionally, two soft-output averaging models were included, a non-weighted and a weighted [8]:

- simple average (“average”)
- weighted average (“lsewavg”)

The weights in the weighted average rule were calculated as the optimal weighting profile of the individual classifier outputs against the correct classification tag, in terms of a least-squares error (LSE) minimization criterion [15]. Thus, this method can be considered as an example of “trained” weighting rules of soft-output classifiers. In contrast, the WMR approach employs fixed analytical weighting profile and hard-output classifications (votes) as input, that is, no further training is required.

3 Experiments and Results

The evaluation of the combination models consisted of two phases, namely: (a) the design and training of SVM classifiers, trained in distinctly different subspaces, and (b) the application of the various combination schemes to the outputs of the individual classifiers.

Each of the K classifiers was separately trained and optimized, using a different group of features from the full dataset, and subsequently evaluated using the corresponding validation set. This training/validation cycle was applied three times, for each of the five datasets, each time using a new random partitioning of the full dataset into training and validation sets. The mean values and standard deviations of the success rates of all the individual ($3K$) classifiers for each dataset, as well as the details about the size and dimensionality of each (full) training and validation sets, are presented in Table 1.

The K value, i.e., the number of feature groups for each dataset, was determined experimentally in a way that each of the corresponding K training segments would be adequate to produce a well-trained SVM classifier. Thus, the German training set was split in K=5 segments, while the Flare-Solar training set in K=4 segments.

Table 1. Single versus multiple classifier accuracy percentages per dataset and *K* values (number of dataset partitions)

Dataset	Train set	Vali-dat. set	Data Dim.	Single classifier accuracy	*K* value	Individual classifier mean acc%
diabetis	468	300	8	76.5 ± 1.7	5	68.3 ± 3.9
flare-solar	666	400	9	67.6 ± 1.8	4	55.7 ± 3.6
german	700	300	20	76.4 ± 2.1	5	68.9 ± 1.8
heart	170	100	13	84.0 ± 3.3	5	74.3 ± 2.3
waveform	400	4600	21	90.1 ± 0.4	5	81.1 ± 1.2

The classification outputs of the pool of *K* classifiers from each training/validation cycle were fed as input to all nine combination schemes, producing the corresponding combined classification outputs. Since the output of each of the *K* classifiers in the pool was calculated based on the same (dimensionally reduced) validation set, the corresponding outputs and accuracy of the combination rules also refer to this validation set.

Table 2 illustrates the mean accuracy of each combination rule (each cell corresponds to three training/validation cycles), as well as the mean value and standard deviation of the success rates of all nine combination rules, for each dataset and *K* value employed.

Table 2. Mean accuracy percentages of all the nine combination rules, with optimized decision threshold, per dataset and *K* values (number of feature groups and classifiers)

Combination Rule	Diabetis	Flare-Solar	German	Heart	Waveform
	K=5	*K*=4	*K*=5	*K*=5	*K*=5
average	71.67	66.08	70.67	85.33	88.12
lsewavg	76.11	65.58	71.56	85.00	86.79
min	68.56	55.92	70.67	69.00	72.98
max	69.11	60.42	67.33	76.67	85.95
median	69.00	58.33	69.78	80.00	81.17
majority	73.00	63.75	70.67	82.33	86.59
wmr/direct	74.00	66.58	70.67	82.33	86.59
wmr/odds	75.33	66.58	71.33	84.00	86.70
wmr/logodds	75.33	66.42	71.33	84.00	86.64
Mean	72.46	63.30	70.44	80.96	84.62
Stdev	2.99	4.06	1.28	5.25	4.77

In the sequel, the overall relative performance of each combination rule was determined in terms of ranking position for each case, i.e., according to its corresponding accuracy for each dataset and *K* value employed. Specifically, a weighted Borda scheme (wBorda) [17] was employed to attribute 10 points to the top-ranked combination rule, 9 points to the second, and so on. In case of a "tie" where two combination rules exhibited exactly the same classification accuracy, both got the same wBorda points for the specific ranking position. Using the results from Table 3, regarding the accuracies, Table 4 illustrates the corresponding wBorda ranking points of all nine combination rules, for each dataset and *K* value employed in this study.

Table 3. wBorda value of all combination rules, with optimized decision threshold, per dataset and *K* values. Each cell value represents the ranking weight according to classification accuracies, with 10 points for top position, 9 points for the second and so on. In cases of equal accuracies, the same ranking weight was assigned to the corresponding combination rules

Combination Rule	Diabetis	Flare-Solar	German	Heart	Waveform
	K=5	*K*=4	*K*=5	*K*=5	*K*=5
average	6	8	8	10	10
lsewavg	10	7	10	9	9
min	3	3	8	4	3
max	5	5	6	5	5
median	4	4	7	6	4
majority	7	6	8	7	6
wmr/direct	8	10	8	7	6
wmr/odds	9	10	9	8	8
wmr/logodds	9	9	9	8	7

Table 4. Overall evaluation of all the combination rules, with optimized decision threshold, using the wBorda results for all datasets and *K* values available. The list is sorted according to the wBorda sum and mean ranking position of each combination rule, from the best to the worst combination rule

Combination Rule	wBorda Sum	wBorda Mean	wBorda Stdev
lsewavg	45	9.0	1.22
wmr/odds	44	8.8	0.84
average	42	8.4	1.67
wmr/logodds	42	8.4	0.89
wmr/direct	39	7.8	1.48
majority	34	6.8	0.84
max	26	5.2	0.45
median	25	5.0	1.41
min	21	4.2	2.17

Table 4 presents a summary of the results shown in Table 3, as well as the list of all the combination rules sorted according to their sum of wBorda points, i.e., their overall efficiency throughout the five original datasets. Tables 2 through 4 present the performance and wBorda results for all the combination rules with optimized decision threshold (T). The decision threshold employed by each combination rule was in every case optimized against final accuracy, using a typical Newton-Raphson optimization algorithm [18].

4 Discussion

The results from Tables 3 and 4 clearly demonstrate the superior performance of the WMR model. Specifically, the all versions of the WMR model exhibited the best performance amongst all the other hard-output combination rules. As expected, it has been proven better than the simple majority voting, as well as all the other rank-based methods (max, min, median). The "odds" weighting profile has also been proven marginally better than the "direct"- and the "logodds"-based profiles for designing the optimal WMR formula.

Interestingly, the "odds"-based version of WMR exhibited better performance than the simple averaging rule, e.g., a soft-output combination model, losing only from the weighted averaging rule with LSE-trained weights. Thus, the WMR model, especially with the "odds" and "logodds" weighting profiles, performs equally well or better than simple soft-output averaging combination rules. All four weighted combination rules, i.e., the three WMR and the LSE-trained weighted average, have been clearly proven better than all the non-weighted hard-output combination rules.

Table 4 also demonstrates the robustness and stability of the each combination rule. For small values of standard deviation (less than unity) in the corresponding wBorda mean ranks, the relative ranking position of a combination rule against the others remains more or less the same. Thus, the maximum rule exhibits a consistently lower ranking position than the simple majority rule, while the "odds"- and the "logodds"-based versions of the WMR models perform consistently better than the simple majority and the three rank-based rules. Furthermore, the "odds"- and the "logodds"-based versions of WMR exhibit the same consistency and robustness as the simple majority rule but with higher success rates.

With respect to the overall performance of the combination rules, results from Tables 1 and 2 demonstrate that in all cases the best combination rules increased the overall success rates of the classifier pool, from +2% (German dataset) to +11% (Flare-Solar dataset), in many cases very close to or equal to the corresponding reference performance level of the single SVM classifier results.

The ensemble of these classifiers clearly demonstrates that the combination of multiple simpler models, each using a $1/K$ portion of the feature space of the dataset, instead of a single classifier for the complete feature space, can be used to reduce the overall training effort. Specifically for the SVM model, kernel evaluation employs inner product between vectors, i.e., its complexity is directly proportional to the dimensionality (number of features) in the input vectors. If this feature space reduction, from F to F/K features, results in a proportional increase in the complexity of the new (reduced) input space in terms of new class distributions, then it is expected that the training of each of the K SVM classifiers may be completed up to K times faster on average. A similar approach has also been examined in other studies [11], using an ensemble of SVM classifiers trained in small training sets, instead of one large training set for a single SVM classifier. Furthermore, there is evidence that such ensembles of kernel machines are more stable than the equivalent kernel machines [11]. This reduction in training time, of course, has to be compared to the additional overhead of calculating a combination rule for every output vector from the classifier pool. Consequently, if the optimal design of this combination rule is simple (linear) and efficient, and its weighting profile can be determined analytically with no need for iterative weight optimization, the WMR approach could prove very prominent for this role in classification tasks of high dimensionality and/or dataset sizes.

5 Conclusions

The game-theoretic modeling of combining classifiers in dichotomous choice problems leads to cooperative gaming approaches, specifically coalition gaming in the form of WMG. Theoretically optimal solutions for this type of games are the WMR schemes, often referred to as weighted majority voting. Under the conditional independence assumption for the experts, there exists a closed solution for the optimal weighting profiles for the WMR formula.

In this paper, experimental comparative results have shown that such simple combination models for ensembles of classifiers can be more efficient than all typical rank-based and simple majority schemes, as well as simple soft-output averaging schemes in some cases. Although the conditional independence assumption was moderately satisfied by using distinct partitions of the feature space, results have shown that the theoretical solution is still valid to a considerable extent. Therefore, the WMR can be asserted as a simple yet effective option for combining almost any type of classifier with others in an optimal and theoretically well-defined framework.

References

1. Condorcet, Marquis de: An Essay on the Application of Probability Theory to Plurality Decision Making: An Election Between Three Candidates. In: Sommerlad and Mclean (1989) 66-80
2. Owen, G.: Game Theory. 3rd edn. Academic Press, San Diego USA (1995)
3. Nitzan, S., Paroush, J.: Optimal Decision Rules in Uncertain Dichotomous Choice Situations. Int. Econ. Rev., 23 (1982) 289–297
4. Shapley, L., Grofman, B.: Optimizing Group Judgemental Accuracy in the Presence of Independence. Public Choice, 43 (1984) 329–343
5. Littlestone, N., Warmuth, M. K.: The Weighted Majority Algorithm. Information and Computation, 108 (1994) 212–261
6. Ben-Yashar, R., Nitzan, S.: The Optimal Decision Rule for Fixed-Size Committees in Dichotomous Choice Situations: The General Result. International Economic Review, 38 (1997) 175–186
7. Karotkin, D.: The Network of Weighted Majority Rules and Weighted Majority Games. Games and Econ. Beh., 22 (1998) 299–315
8. Kuncheva, L. I.: Combining Pattern Classifiers. John Wiley and Sons, New Jersey USA (2004)

9. Mavroforakis, M., Theodoridis, S.: Support Vector Machine Classification through Geometry. Proc. XII Eur. Sig. Proc. Conf. (EUSIPCO2005), Antalya, Turkey, Sep. 2005
10. Mavroforakis, M., Theodoridis, S.: A Geometric Approach to Support Vector Machine (SVM) Classification. IEEE Trans. NN, 2006 (in press)
11. Evgeniou, T., Pontil, M., Elisseeff, A.: Leave One Out Error, Stability, and Generalization of Voting Combinations of Classifiers. Machine Learning, 55 (2004) 71–97
12. Cristianini, N., Shawe-Taylor, J.: An Introduction to Support Vector Machines and Other Kernel-Based Learning Methods. Cambridge University Press UK (2000)
13. Jain, A., Duin, R., Mao J.: Statistical pattern recognition: a review. IEEE Trans. PAMI, 22 (2000) 4–37
14. Mavroforakis, M., Sdralis, M., Theodoridis, S.: A novel SVM Geometric Algorithm based on Reduced Convex Hulls. 18th International Conference on Pattern Recognition (ICPR 2006), Hong Kong, Aug. 2006
15. Theodoridis, S., Koutroumbas, K.: Pattern Recognition. 4th edn. Academic Press, San Diego USA (2006)
16. Cooley, W., Lohnes, P.: Multivariate data analysis. John Willey & Sons, New York USA (1971)
17. Parker, J. R.: Rank and Response Combination from Confusion Matrix Data. Information Fusion, 2 (2001) 113–120
18. Forsythe, G., Malcom, M., Moler, C.: Computer Methods for Mathematical Computations. Prentice-Hall, New Jersey USA (1977)

Section 4: A game-theoretic framework for classifier ensembles using weighted majority voting with local accuracy estimates

Summary:

In this paper, a novel approach for the optimal combination of binary classifiers is proposed. The classifier combination problem is approached from a Game Theory perspective. The proposed framework of adapted weighted majority rules (WMR) is tested against common rank-based, Bayesian and simple majority models, as well as two soft-output averaging rules. Experiments with ensembles of Support Vector Machines (SVM), Ordinary Binary Tree Classifiers (OBTC) and weighted k-nearest-neighbor (w/k-NN) models on benchmark datasets indicate that this new adaptive WMR model, employing local accuracy estimators and the analytically computed optimal weights outperform all the other simple combination rules.

Citation:

"A game-theoretic framework for classifier ensembles using weighted majority voting with local accuracy estimates", H. Georgiou, M. Mavroforakis, ArXiv.org: 5-Feb-2013 (arXiv:1302.0540v1 [cs.LG]).

A game-theoretic framework for classifier ensembles using weighted majority voting with local accuracy estimates

Harris V. Georgiou[1] & Michael E. Mavroforakis[2]

Abstract

In this paper, a novel approach for the optimal combination of binary classifiers is proposed. The classifier combination problem is approached from a Game Theory perspective. The proposed framework of adapted weighted majority rules (WMR) is tested against common rank-based, Bayesian and simple majority models, as well as two soft-output averaging rules. Experiments with ensembles of Support Vector Machines (SVM), Ordinary Binary Tree Classifiers (OBTC) and weighted *k*-nearest-neighbor (w/*k*-NN) models on benchmark datasets indicate that this new *adaptive* WMR model, employing *local accuracy estimators* and the *analytically* computed optimal weights outperform all the other simple combination rules.

Keywords

Classifier combination, weighted majority voting, game theory, linear combiners, ensemble learning, decision fusion.

1. Introduction

Classifier combination is one of the most active areas of research in the discipline of Pattern Recognition. The challenging problem of designing optimal aggregation schemes for multi-classifier systems has been addressed by a wide range of methodologies and approaches during the last decade [1]. However, few of them introduce a framework of analytical solutions. Instead, most of them employ either heuristics or iterative optimization procedures.

In this paper, a novel viewpoint is proposed for the problem of optimally combining classifiers using game-theoretic arguments. Specifically, the problem of designing optimal ensembles of voting classifiers is investigated within the context of Game Theory [2, 3], as an analogy to *n*-person games. A special type of *cooperative games*, namely the *coalition games*, is introduced as the natural setting for formalizing the ensemble design problem, within the scope of Coalition Theory [2, 3] and the Weighted Majority Games (WMG) [2, 3]. This new formulation of the problem leads to the development of a theoretical framework of the weighted voting schemes [1]. Furthermore, this approach leads to *optimal analytical solutions* for the two core problems of: (a) designing the aggregation rule in an optimal way, and (b) assigning optimal voting weights in a voting ensemble of experts. For the problem in (a), the theory of WMG states that the *optimal voting aggregation rules* in a fixed-size ensemble for an arbitrary *n*-label classification task is the weighted majority rule (WMR) [2, 3, 4, 62]; while for the problem in (b), the *optimal voting weights* in such WMR schemes are *calculated analytically* from the experts' competencies, under the conditional independence assumption [4, 5].

This particular type of game-theoretic analytical solution is extremely useful in the process of designing optimal classifier ensembles. The use of simple linear combination models that employ single weights for each

[1] (MSc, PhD) Associate researcher (post-doc) at Dept. of Informatics & Telecommunications, National Kapodistrian Univ. of Athens (NKUA/UoA), Greece – Email: `harris {at} xgeorgio {dot} info`

[2] (MSc, PhD) Associate professor at Computational Biomedicine Lab., Dept. of Computer Science, University of Houston (UH), Texas, USA – Email: `mmavrof {at} uh {dot} edu`

classifier, which, however, *do not require iterative training/optimization*, can provide the necessary means to apply multi-classifier schemes in parallel implementations with on-line updating capabilities. In other words, the classifiers can be trained independently and off-line, using any architecture and algorithm available, while *the aggregation scheme involves only direct (analytical) calculation of the voting weight of each classifier*. Another novelty of this paper refers to the notion of the expert's competency, i.e., the prior estimation of the success rate of each individual classifier, as it is required by the WMR optimal rule. The expert's competency is extended to include the *posterior* probabilities associated with each pattern. In practice, the voting weight of each classifier is calculated analytically, in the sense of the WMR formulation, for every sample using the *local accuracy estimates (LAE)*.

This paper is organized as follows. Section 2 describes the core aspects of the classification task and its realization under the concept of multi-classifier systems. Section 3 summarizes some basic concepts of Game Theory and phrases the classifier combination task in game-theoretic terms. Section 4 describes the details regarding the datasets and methods used. Section 5 presents the experiments and results. Section 6 is a discussion on the results. Section 7 presents the conclusions.

2. Problem statement and current practices

2.1 Combining classifiers

The ultimate goal of any pattern recognition system is to design optimally a classifier while at the same time attaining the best generalization performance, for the specific problem at hand. However, even the "best" classifier model can fail on points that other classifiers may succeed in predicting the correct label [6, 1]. Many studies have focused on the possibility of exploiting this complementary nature of the various classifiers, in order to enhance the overall performance. Specifically, each classifier is considered as a trained expert that participates along with others in a "committee", which produces a collective decision according to some well-specified rule.

In the discipline of collective decision-making, a group of *N* experts, each one with moderate performance levels, are combined in order to produce a collective decision that may be better than the estimate of the best among the experts in the group. According to the famous Condorcet Jury Theorem [7], if the experts' individual decisions are *independent* and their corresponding estimates are more likely to be correct than incorrect ($p_{correct}$>0.5), then an increase in the collective performance, as a group, is guaranteed when the individual estimations are combined. Moreover, this performance continues to grow asymptotically as the size *N* of the group increases and under the independence assumption. This assertion has been the base for very active experimental and theoretical research in the discipline of pattern recognition.

Over the last decade or so, a wide range of different approaches have been studied to design aggregations or *ensembles* of experts. These employ either a *selection* or a *fusion* scheme [1] to combine the individual classifiers' outputs into a final collective decision. The combining rules vary from very simple to more sophisticated ones. Typical examples include simple averaging and fusion [8, 9], mixture of experts [10, 11], consensus or majority voting [12], dynamic classifier selection [13], supra Bayesian methods [14, 15], evidence-based [16, 17] or template-based [18] decision models. Most of the methods employ weights upon each member in the pool, essentially dictating a corresponding level of confidence to its individual decisions. Hence, the design of such ensembles reduces to the problem of finding these optimal aggregation parameters, with the goal of improving the final accuracy rates of the ensemble.

To comply with the spirit of Condorcet's theorem, a major research effort has been inserted in designing the individual classifiers to be as independent as possible. More recent approaches, such as boosting [19], bagging [20] and random subspace models [21], employ different techniques to increase the level of diversity, essentially by training individual classifiers in different subsets or subspaces of the original set of data.

This requirement can be implemented, in practice, by separating or *splitting* the original training datasets into a new set of distinct or partially overlapping realizations, with respect to: (a) the data samples, (b) the dimensionality, or (c) both. Random subspace methods, most commonly used in aggregation models like the Random Forests [22], are typical examples of using dimensionally-reduced versions of the original data space. Rotation Forests [23] is an example of using both different subsets and different subspaces simultaneously.

3. A game-theoretic approach to classifier combination

This section presents a brief overview of the basics of Game Theory, the main concepts of cooperative games and Coalition Theory, as well as the formal definition of weighted majority games (WMG) and weighted majority rules (WMR). Furthermore, the core problem of designing optimal weighted voting schemes for multi-classifier ensembles is introduced within the context of WMG and theoretical analytical solutions are presented under the WMR formulation.

3.1 Elements of Game Theory

The mathematical theory of *games* and *gaming* was first developed as a model for situations of conflict. Since the early 1940's, the work of John Von Neumann and Oskar Morgenstern [24] has provided a solid foundation for the most simple types of games, as well as analytical forms for their solutions, with many applications to Economics, Operations Research and Logistics [2, 3]. Each opposing *player* in a game has a set of possible actions to choose from, in the form of *pure* (single choice) or *mixed* (random combination) *strategies*. The set of optimal strategies for all the players is called the *solution* to this game.

The *zero-sum* games are capable of modeling situations of conflict between two or more players, where one's gain is the other's loss and vice versa. In reality, it is common that in a conflict not all players receive their opponents' looses as their own gain and vice versa. In other words, it is very common a specific combination of decisions among the players to result in a certain amount of loss to one and a corresponding gain, not of equal magnitude, to another. In this case, the game is called *nonzero-sum* and it requires a new set of rules for estimating optimal strategies and solutions.

During the early 1950's, John Nash has focused primarily on the problem of finding a set of *equilibrium points* in nonzero-sum games, where the players eventually settle after a series of competitive rounds of the game. In 1957 [25], Nash successfully proved that indeed such equilibrium points exist in all nonzero-sum games, defining what is now known as the *Nash theorem* or *Nash solution* to the *bargaining problem* [2, 3]. However, although the Nash theorem ensures that at least one such *Nash equilibrium* exists in all nonzero-sum games, there is no clear indication on how the game's solution can be analytically calculated at this point. In other words, although a solution is known to exist, there is no closed form for nonzero-sum games until today.

The Nash equilibrium points are not always the globally optimal option for the players. In fact, the Nash equilibrium is optimal only when players are strictly competitive, i.e., when there is no chance for a mutually agreed solution that benefits them more. These strictly competitive forms of games are called *non-cooperative* games. The alternative option, the one that allows communication and prior arrangements between the players, is called a *cooperative* game and it is generally a much more complicated form of nonzero-sum gaming.

3.2 Cooperative games and coalition gaming

The problem of cooperative or possibly-cooperative gaming is the most common form of conflict in real life situations. Since nonzero-sum games have at least one equilibrium point, when studied under the strictly competitive form, Nash has comprehensively studied the cooperative option as an extension to it. However, the possibility of finding and mutually adopting a solution that is better for both players than the one suggested by the Nash equilibrium, essentially involves a set of behavioral rules regarding the players' stance and mental state, rather than strict optimality procedures [2, 3]. Nash named this process as *bargain* between the players, trying to mutually agree on one solution between multiple choices within a *bargaining set*. In practice, each player should enter a bargaining procedure if there is a chance that a cooperative solution exists and it provides at least the same gain as the best strictly competitive solution, i.e., the best Nash equilibrium. In this case, if such a solution is agreed between the players, it is called *bargaining solution* of the game. This new framework provides the necessary means to study *n*-person non-cooperative and cooperative games under a unifying point of view. Specifically, a nonzero-sum game can be realized as a strictly competitive or a possibly cooperative form, according to the game's rules and restrictions. Therefore, the cooperative option can be viewed as a generalization to the strictly competitive mode of gaming.

When players are allowed to cooperate in order to agree on a mutually beneficial solution of game, they essentially choose one strategy over the others and bargain this option with all the others in order to come to an agreement. For *symmetrical* games, i.e., when all players receive the same gains and losses when switching places, this situation is like each player choosing to join a group of other players with similar preference over their initial choice. Each of these groups is called a *coalition* and it constitutes the basic module in this new type of gaming: the members of each coalition act as cooperative players joined together and at the same time each coalition competes over the others in order to impose its own position and become the *winning* coalition. This setup is very common when modeling voting schemes, where the group that captures the relative majority of the votes becomes the winner.

Coalition Theory [2, 3] is closely related to the classical Game Theory and in particular the cooperative gaming. In essence, each player still tries to maximize its own expectations, not individually any more but instead as part of a greater opposing term. Therefore, the individual gains and capabilities of each player is now considered in close relation to the coalition this player belongs to, as well as how its individual decision to join or leave a coalition affects this coalition’s winning position. The theoretical implications of having competing coalitions of cooperative players, instead of single players, is purely combinatorial in nature, thus making its analysis very complex and cumbersome. There are also special cases of collective decision schemes where a single player is allowed to *abstain* completely from the voting procedure, or prohibit a contrary outcome of the group via a *veto* option. Special sections of Game Theory, namely the *coalition gaming* and *stable sets* in cooperative gaming [2, 3], have studied the effects of introducing “weights” to the choice of each expert according to their competencies, in order to optimize the final decision of the group.

3.3 Classifier combination as a game-theoretic problem

The transformation of cooperative *n*-person games into coalition games essentially brings the general problem into a voting situation. Each player casts a vote related to its own choice or strategy, thus constituting him/her as a member of a coalition of players with similar choices. The coalition that gains more votes becomes the winner. In the case where each player selects one out of *M* available options to cast its vote, the collective group decision can be estimated simply by applying the majority voting scheme, i.e., the choice selected is the one gathering the majority of votes. Each subgroup of consentient players essentially represents an opposing assembly to all the other similar subgroups with different consensus of choice.

In the general case where a weight is assigned to each voter and there are *n* available choices to vote for, this form is known in Game Theory as the weighted majority game (WMG) [2, 3]. It has been proven by Nitzan and Paroush (1982) [4] and Shapley and Grofman (1984) [5], that the optimal decision rules for these WMG, in terms of collective performance, are the weighted majority rules (WMR). The same assertion has also been verified independently by Ben-Yashar and Nitzan [26] as the optimal aggregation rule for committees under the scope of informative voting in Decision Theory. This result was later (2001) [62] extended from dichotomous to polychotomous choice situations; hence *the optimality of the WMR formulation has been proven theoretically* for any *n*-label voting task. Furthermore, under the conditional independence assumption, a closed form solution for the voting weights in the WMR formula exists and it is directly linked to each expert’s competency. This *optimal weight profile* for the voting experts is the *log of the odds (“log-odds”) of their individual competencies* [4, 5].

In this paper, the notion of modeling classification tasks for an ensemble of experts via the precise game-theoretic formulation of WMG and WMR is for the first time applied for combining hard-output (voting) classifiers. Specifically, the design of the combination rule is treated as a standard WMG situation, with each classifier participating in a *simple* coalition game, i.e., choosing the final decision based on the maximum votes (sum of weights) casted. The voting weights in this WMR scheme are calculated in an *analytical* way using the log-odds solution [27, 1]:

$$w_i = \log\left(\frac{p_i}{1-p_i}\right) \quad , \quad p_i = P_i\left\{\theta = \omega_{correct}\right\} \quad , \quad i = 1 \ldots K \qquad (1)$$

where w_i is the combination weight assigned to the *i*-th classifier (player), p_i is the respective estimated probability for correct classification, measured in the validation set, θ is the predicted class label and $\omega_{correct}$ is the correct class label for **x** (either ω_1 or ω_2), respectively.

Using this game-theoretic analytical solution, the WMR formula is used as the *optimal voting aggregation scheme*, i.e.:

$$O_{wmr}(\mathbf{x}) = \sum_{i=1}^{K} w_i D_i(\mathbf{x}) \qquad (2)$$

where D_i is the hard-output of each of the *K* individual classifiers in the ensemble, w_i is its assigned weight, O_{wmr} is the weighted majority sum.

In this study, the classification tasks where chosen to include only dichotomous choice situations, for several reasons (explained later on). Hence, there are only two voting options available (*M*=2) and two class labels to choose from (either ω_1 or ω_2), which essentially simplifies the WMR formulation to a sign-assignment problem:

$$D_{wmr}(\mathbf{x}) = sign\left(O_{wmr}(\mathbf{x}) - T\right) \qquad (3)$$

where D_{wmr} is the final decision of the ensemble against a fixed-valued decision threshold T, which is typically half the range of values for O_{wmr} [27, 1], i.e.:

$$T = \frac{1}{2}\left(\max_i\{D_i\} + \min_i\{D_i\}\right) \quad , i = 1 \ldots K \qquad (4)$$

or simply T=1/2 when normalized weights w_i are employed (sum of weights is unity).

Interestingly, although the optimality of this solution under certain conditions has been studied theoretically in the context of many different disciplines, including decision theory and automata theory [28, 29, 30], it is generally considered very limited in terms of optimality, since it does not take into account any dependencies among the trained classifiers on an ensemble.

In this paper, two versions of this WMR-based weighting scheme with respect to the value of p_i are tested: (a) the "static" WMR, using the prior probabilities of correct classification (i.e. "global" competence), and (b) the "adaptive" WMR, using the (estimated) posterior probabilities of correct classification (i.e. "local" competence). In both cases, the success rates are calculated based on a validation set of samples, independently of any training process and any training set used by the classifiers. In this new "adaptive" version of the WMR, which essentially introduces the notion of *local experts* into this framework, the combination weights are calculated so that they reflect the localized (conditional) competencies of the classifiers at each point, i.e.:

$$w_i = \log\left(\frac{p_i}{1-p_i}\right) \quad , \quad p_i = P_i\{\theta = \omega_{correct} \mid \mathbf{x}\} \quad , \quad i = 1 \ldots K \qquad (5)$$

This procedure is presented in section 4.5.

4. Datasets and Methodology

4.1 Selection of benchmark datasets

In order to assess the performance of the various classifier combination methods, publicly available benchmark datasets were considered. Specifically, the Raetch [31, 32, 33] and the ELENA [34] dataset resources were considered and, for the purposes of this study, only 2-class sets with real (non-artificial) data were initially selected. The main reason for employing only dichotomous classification tasks is that multi-class problems essentially add one more layer of complexity in some classification models, especially SVM-based. In practice, the simplification of the classification task itself does not require any second-stage decision, e.g., one-versus-all or pair-wise comparisons. Furthermore, when M>2 choices are available in WMG setups, the corresponding WMR decision requires one additional parameter of non-trivial optimization [67], the majority threshold or *quota* (q), instead of a simple comparison to the half-sum of the voting weights, as in Eq. (3) and (4). Since the goal of this study is the comparative evaluation of combination rules, while keeping all the other factors as simple as possible, using only 2-class benchmark datasets is a natural choice.

A group of 14 datasets were analyzed in terms of class separability and statistical significance of the corresponding results. In order to make individually trained classifiers as diverse as possible, the method of training them in different subspaces was selected. Consequently, datasets of high dimensionality were preferred. The quantitative criteria used for selecting the final datasets from this group included: (a) The inherent dimensionality of the dataset, in order to be able to use a feature subspace method leading to at least five distinct feature groups. (b) The Chernoff Bound and the corresponding Bhattacharyya Distance [6], as a commonly used class separability measure, when Gaussian distributions are assumed for the classes. (c) Guyon's error counting approach [35] for estimating the minimum size of the test data set, based on the results of a leave-one-out [36] error estimation from a simple OBTC model [37]. (d) The ELENA project's proposal [38] for the minimum number of samples necessary for the estimation of the probability density function (pdf) of a Gaussian probability density, using a probability density kernel estimator, with less than 10% error. (e) The intrinsic dimension was calculated in terms of the fractal dimension estimation method [39, 40] and compared to the real number of distinct features of the dataset, in order to get a quantitative measure of the overall complexity of the sample space and the degree of redundant information among the features. Based on these criteria, the final group of the selected datasets included four candidates from the Raetch packages. These datasets are: 1) Ringnorm, 2) Splice, 3) Twonorm, and 4) Waveform.

4.2 Dataset split and subspace method

Each base dataset was randomly separated into a base training set and a test set of samples. Diversity among the classifiers was introduced in the ensemble by training them in different subspaces. Random subspace methods [21] have been successfully used in the past as the means to increase diversity among classifiers. However, in these methods, the grouping of distinct features into subsets is conducted randomly and involves either distinct or overlapping memberships of features in the various groups. In contrast, in this study a non-random subspace procedure was implemented by using a feature ranking method and a subsequent grouping into distinct subsets, in order to achieve more or less equal discrimination power. A much simpler version of this method has been used successfully in the past [41]. These approaches are generally referred to as “ranked” subspace methods, since subsets of features are evaluated and ranked according to some specific statistical criterion, in order to control the discrimination power and the robustness of each subspace in various classification or clustering applications (see e.g. [62]).

In this paper, the training set was partitioned into *K* distinct feature groups. Each group of features was created in a way that satisfied three basic constraints: (a) each group to be distinct, i.e., no feature is common in any two groups, (b) all the features are used “as-is”, i.e., no projection or other complex transformation is applied (e.g. PCA), and (c) each feature group to represent approximately the same class-discrimination potential. The third constraint requires a method for ranking all the features in terms of discrimination power, against the two classes, as well as their statistical independence with regard to the other features in the initial training set. The MANOVA method [42] was used to assign a multivariate statistical significance value to each one of the features and then produce a sorted list, based on (the log of) this value.

A “fair” partitioning of this sorted list of features into equally “accurate” groups, in terms of classification results, was conducted by selecting features in pairs from the top and bottom positions, assigning the currently “best” and “worst” features in the same group. Furthermore, the efficiency of each group was estimated in terms of summing the log of the statistical significance value, assigned by MANOVA, of all the features contained in this group. The log was employed in order to avoid excessive differences between the values assigned by MANOVA, thus creating more even subset sums of these values. Practically, every such pair of features was assigned in groups sequentially, in a way that all groups contained features with approximately equal sum of the log of the values assigned by MANOVA.

Each one of these *K* distinct feature groups was used for training one of the *K* classifiers in the ensemble. As a result, the issue of the desired diversity between the classifiers of the ensemble is addressed independently from the combination rules themselves, making their subsequent comparison easier and more realistic. It should be noted, that the goal of this study is not classifier independence or diversity, but rather to evaluate the performance of the WMR and other combination rules, using *weakly* independent classifiers, i.e., without guaranteed diversity. In fact, the introduction of a feature subspace method only creates *some* diversity, which makes the evaluation of the ensembles more realistic.

4.3 Classifier models

Three, among the most popular, classifier models were selected to form committees of experts, in order to test the various classifier combination schemes. Specifically, the Support Vector Machine (SVM) [43, 44], the (weighted) k-nearest-neighbor (w/*k*-NN) [6] and the Decision Tree (DT) [45] classifiers were employed in this study.

For the SVM architecture, a geometric nearest point algorithm (NPA) [46], based on the notion of reduced convex hulls (RCH) [47], was used for training a standard SVM architecture with radial-basis function (RBF) as the kernel of the non-linear mapping.

The tree-based classification models were selected as a very typical candidate of unstable classifiers, already used successfully in other combination schemes. In their simplest form, each tree node contains a threshold value that is compared to one of the input features and the result dictates which of two possible paths to follow towards the next tree level. This type of decision trees is often referred to as Ordinary Binary Classification Trees (OBCT) [6]. In this study, the classic Classification and Regression Tree (CART) algorithm [45] was employed for designing soft-output (regression) and hard-output (classification) DT, used in conjunction with soft- and hard-output combination rules, respectively. Three splitting criteria were tested separately: (a) the *Gini index* of diversity, a typical choice in CART models that is similar to entropy, (b) the *twoing* criterion, optimizing the criterion of splitting the contents of each node into two disjoint and mutually exclusive subsets, and (c) the *deviance* criterion, which maximizes the variability (variance) reduction within each of the two splits of the node. These three splitting criteria, which are some of the most commonly used choices in typical OBCT models, were tested separately for completeness purposes.

Finally, a modified version of the w/*k*-NN classifier was employed [48, 49, 50]. The typical *k*-NN classifier architecture was enriched with the options of choosing distance functions other than the classic Euclidean, employing a non-constant weighting function to the test samples around the center of the *k*-closure neighborhood. In this study, the weighting functions were fixed (non-trained), symmetric around the center of the *k*-neighborhood and scaled appropriately. In each case, the smallest weight value was assigned to the furthest of the *k* neighbors and the largest weight value to the center of the *k*-neighborhood. In other words, the weighting profile was either constant, in the case of the typical non-weighted *k*-NN implementation, or a constantly decreasing function around the center of the *k*-neighborhood. The distance metrics implemented for this w/*k*-NN classifier were the Euclidean, the city block, the Minkowski, the cosine, the correlation, the Mahalanobis, the Chebychev, and the Hamming kernels [6]. The weighting metrics implemented for this w/*k*-NN classifier were the constant (classic, no weighting), the linearly decreasing and the Gaussian profiles. The introduction of the different distance functions and especially the option of employing weights to the *k*-neighbors according to their distance from the center of the local test set, had little effect to the overall performance of the w/*k*-NN classifier but produced much more stable soft-output profiles, which were used subsequently for the calculation of local accuracy estimates (see: section 4.5).

4.4 Combination rules

A total of eight combination rules were examined in this study. Specifically, four typical hard-output combination methods were employed (namely one classic rank-based method and three voting-based schemes, including the "static" and "adaptive" versions of WMR), two soft-output averaging methods and two Bayesian-based combination rules.

The standard maximum rule was employed as a typical rank-based method [1, 51]:

- maximum ("STD: maximum"):

$$O_{\max}(\mathbf{x}) = \omega_s : p_s = \max_{j=1\ldots2}\left\{\max_{i=1\ldots K}\left\{\mu_{ij}\right\}\right\},\ \mu_{ij} = O_i(\mathbf{x}) \mid \left(\theta(\mathbf{x}) = \omega_j\right) \qquad (6)$$

where: $\mathbf{x}$ is the current input sample to be classified, $O_i(\mathbf{x})$ is output value by the i-th classifier for class label ω_j given $\mathbf{x}$, $\theta(\mathbf{x})$ is the predicted class label, $\{\omega_1, \omega_2\}$ are the two class labels, K is the size of the ensemble, μ_{ij} is the corresponding support value by the i-th classifier for class label ω_j given $\mathbf{x}$, and p_s is the selected support value.

It should be noted that two other typical rank-based combination rules are equivalent to the maximum and the simple majority rules, respectively, in case of dichotomous choice classification [1]. Specifically, the class labels selected by the minimum rule [51] are the same to the ones selected by the corresponding maximum rule that uses the same support values. Similarly, the class labels selected by the median rule [51] are the same to the ones selected by the corresponding simple majority rule that uses the same support values.

- simple majority voting ("STD: simple majority") [1, 52]:

$$O_{maj}(\mathbf{x}) = \sum_{i=1}^{K} D_i(\mathbf{x}) \qquad (7)$$

$$D_{maj}(\mathbf{x}) = sign\left(O_{maj}(\mathbf{x}) - T\right) \qquad (8)$$

where D_i is the hard-output of each of the K individual classifiers in the ensemble, O_{maj} is the majority sum. The final hard-output decision D_{maj} of the simple majority rule is calculated against a fixed-valued decision threshold T, which is typically half the range of values for O_{maj} [1]:

$$T = \tfrac{1}{2}\left(\max_i\{D_i\} + \min_i\{D_i\}\right) \quad , i = 1 \ldots K \qquad (9)$$

which is the same to the one employed for WMR (in Eq.4) but with w_i=1/K for the simple majority rule.

Additionally, two soft-output averaging models were included, a non-weighted and a weighted one [1]:

- simple average ("STD: simple average"):

$$O_{avg}(\mathbf{x}) = \sum_{i=1}^{K} w_i O_i(\mathbf{x}) \quad , w_i = 1/K \tag{10}$$

- weighted average ("STD: LSE-weighted average"):

$$O_{lsewavg}(\mathbf{x}) = \sum_{i=1}^{K} \hat{w}_i O_i(\mathbf{x}) \tag{11}$$

where $O_i(\mathbf{x})$, $O_{avg}(\mathbf{x})$ and $O_{lsewavg}(\mathbf{x})$ are the soft-output value of the single classifier (ensemble member), the simple averaging rule and the weighted average rule, respectively. For the simple average rule, all weights are equal, i.e., $w_i=1/K$. The vector $\hat{\mathbf{w}}$ is the optimal one for the weighted average rule, estimated by a linear regression formula on the individual classifier outputs, against the correct classification tag, in terms of a least-squares error (LSE) minimization criterion [53, 1]. Thus, this method can be considered as an example of a "trained" linear weighting rule of soft-output classifiers. In contrast, *the WMR approach employs fixed analytical optimal weighting profile* and hard-output classifications (votes) as input, *with no need for further training*.

Finally, two Bayesian-based combination rules were employed as a very efficient and simple implementation of non-weighted schemes, which exploit information about local accuracy estimates. Specifically, the method of Dynamic Classifier Selection based on Local Accuracy (DCS-LA) [13, 54, 55] was employed as a typical example of a local accuracy-based non-weighted combination rule, using the notion of overall local accuracy [13]. Two different variants of this model (DCS-LA variants) were implemented:

- employing the full Bayes rule for the conditional probabilities ("STD: DCS-LA (with priors)"):

$$O_{bayesP}(\mathbf{x}) = D_s(\mathbf{x}) : p_s = \max_{i=1\ldots K}\left\{p_i\right\} \quad , p_i = P_i\left\{\theta = \omega_{correct} \mid \mathbf{x}\right\} \cdot P_i\left\{\theta = \omega_{correct}\right\} \tag{12}$$

- or, using only the local accuracy estimate itself ("STD: DCS-LA (no priors)"):

$$O_{bayes}(\mathbf{x}) = D_s(\mathbf{x}) : p_s = \max_{i=1\ldots K}\left\{p_i\right\} \quad , p_i = P_i\left\{\theta = \omega_{correct} \mid \mathbf{x}\right\} \tag{13}$$

The final decision in these types of DCS-LA models is dictated by the expert with the highest conditional probability of success, i.e., highest "confidence". As a result, the model implemented in this study is essentially a direct implementation of the standard Bayes decision theory that is based on maximizing the likelihood of "correct" classification for the current input data **x**. It should be noted that, although the Bayes rule includes division by the pdf against the input data **x**, this factor is irrelevant here since the model is applied to a specific input sample and therefore the corresponding pdf value is always equal to unity.

For all the soft-output combination rules (simple and LSE-weighted averaging, DCS-LA models), the final decision is calculated against a fixed threshold value similarly to the majority voting rules, which is typically half the range of values for the specific combination rule, i.e.:

$$D_{combR}(\mathbf{x}) = sign\left(O_{combR}(\mathbf{x}) - T\right) \tag{14}$$

$$T = 1/2\left(\max_i\{O_i\} + \min_i\{O_i\}\right) \quad , i = 1 \ldots K \tag{15}$$

where O_i is the soft-output of each of the K individual classifiers in the ensemble, O_{combR} is the soft-output of the combination rule and D_{combR} is the final class decision (thresholded value).

Table 1 summarizes all the eight combination rules used in this study.

Table 1: Overview of the eight combination rules used in this study.

	Non-weighted	**Weighted**
Static or Rank-based	simple average	LSE-weighted average
	maximum	WMR (static) logodds
	majority	
Adaptive (using posteriors)	Bayesian (DCS-LA) no priors	WMR (adaptive) logodds
	Bayesian (DCS-LA) with priors	

4.5 Local accuracy estimates method

To compute the local accuracy estimates, as required by the DCS-LA method, as well as the modification of the WMR as pointed out in section 3.3, the "overall" local accuracy method was adopted [13, 55]. To this end, we chose to estimate the error pdf directly by employing a non-parametric density-based method, by means of histogram approximation [6]. For one-dimensional probability functions, the histogram method has been proven more efficient than direct interpolation through isotonic regression functions [56]. In order to avoid non-uniformities in the distribution of the classifiers' soft-output values, the use of *dynamic bin width allocation* [57] was employed instead of equal bin width. This method ensures that every bin contains roughly the same number of samples, i.e., the width of the histogram bins is adjusted appropriately, in order to produce uniform resolution and smoothness of the histogram curve throughout the entire range of values.

In this study, the local accuracy estimation of each classifier was based on the corresponding error pdf, approximated via the histogram method with dynamic bin width allocation. The complete process includes five distinct steps: (a) Each classifier in the ensemble is designed based on a training set of samples and subsequently evaluated on a different set of test samples. (b) The soft-output values of the classifier are then distributed uniformly into bins of dynamic width. (c) The corresponding localized error probability estimation is calculated for every bin in terms of error frequency ratio (errors versus total samples in bin). (d) After the calculation of the error pdf value in each bin, the bin values are interpolated with a piecewise cubic Hermite spline, chosen for its shape-preserving properties [58, 59], in order to produce the final, continuous, error pdf estimator. (e) The local accuracy estimate, i.e., the "success" pdf, for each specific classifier is calculated by one minus its corresponding error pdf (interpolated) value for a given input sample **x**, i.e.:

$$P_i\left\{\theta = \omega_{error} \mid \mathbf{x}\right\} = {Ne_i^m}\big/{N_i^m} \quad , b_i^m \geq O_i(\mathbf{x}) > b_i^{m+1} \tag{16}$$

$$LAE : P_i\left\{\theta = \omega_{correct} \mid \mathbf{x}\right\} = 1 - H_i^m\left(P_i\left\{\theta = \omega_{error} \mid \mathbf{x}\right\}\right) \quad , b_i^m \geq O_i(\mathbf{x}) > b_i^{m+1} \tag{17}$$

where $O_i(\mathbf{x})$ is the *i*-th classifier's output, N_i^m is the number of output values from the *i*-th classifier in (dynamic) histogram bin *m*, Ne_i^m is the number of incorrect classifications committed by the *i*-th classifier in bin *m*, b_i^m and b_i^{m+1} are the boundaries of bin *m* for the *i*-th classifier, and H_i^m is the piecewise cubic Hermite spline for the error pdf in bin *m* for the *i*-th classifier.

Results from all single-classifier tests have shown that the local accuracy estimation based on the dynamic-width histogram method produced very robust and accurate results for all classifiers, even for the OBTC, which is the most unstable of the three base classifiers used in this study. Figure 1 illustrates a real example of the procedure. The plot in (a) illustrates the fixed-width histogram (counts per fixed-width bin) of the SVM kernel output values against correct (green/high) and incorrect (red/low) classifications, for the single-classifier configuration in the "splice" dataset. The plot in (b) illustrates the resulting dynamic-width normalized histogram (dots) and the corresponding local accuracy estimator (interpolated dynamic-width bins) function, which calculates the estimated posterior probability of successful (blue/high) and incorrect (red/low) classification with regard to the current SVM kernel output.

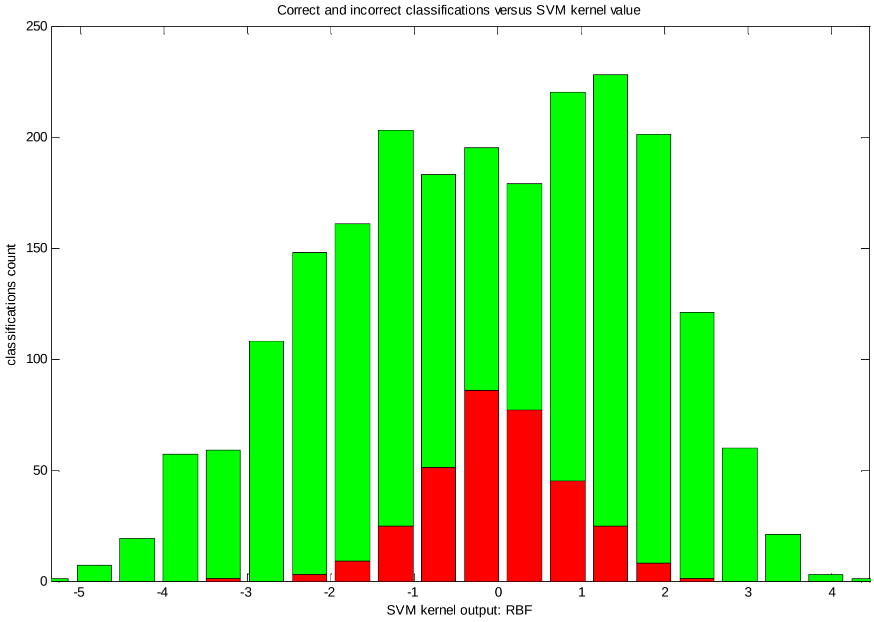

(a)

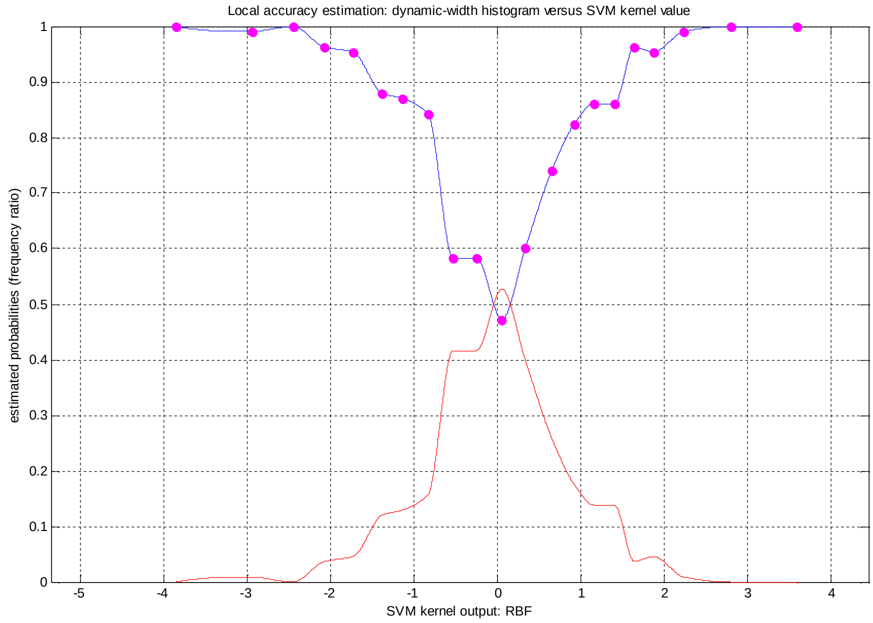

(b)

Figure 1: Example of the local accuracy estimation method. Green (high/lighter) and red (low/darker) portions of the bars in the fixed-width histogram in plot (a) represent number of correct and incorrect classifications, respectively, with regard to the classifier's output value. The marked points in plot (b) correspond to the new bin centers of the dynamic-width histogram for the "success" pdf, while the interpolation curve represents the analytical local accuracy estimator function for correct (blue/high/darker curve) and incorrect (red/low/lighter curve) classification, with regard to the classifier's output value. The results are for the single SVM classifier on the "splice" dataset.

5. Experiments and Results

The experimentation phase of this study involved three consecutive stages: (a) the characterization of all datasets based on single classifier tests, (b) the training of ensembles based on different choices of datasets, number and type of base classifiers, and (c) the comparative evaluation of all the classifier combination rules.

5.1 Datasets and classifiers characterization

In the preliminary analysis stage, every dataset/classifier combination involved five training realizations, each one employing training and parameter optimization according to each classifier model. For the SVM classifier, the optimization included the μ soft margin parameter of the NPA-RCH training algorithm [46, 47], as well as the *epsilon* parameter (convergence accuracy). Furthermore, the Radial Basis Function (RBF) kernel [44] with optimized *sigma* (σ) value was used in the final choice of the SVM classifiers' structure. For the OBTC, each setup included a classification and a regression tree, both trained on the same data and both using the same splitting criterion (one of: *Gini*, *twoing* or *deviance*). In all cases, the OBTC models were optimized against the exact choice of the splitting function and they all employed a minimum limiting threshold of ten samples per splitting node during training/pruning phases [1, 6]. Finally, for the w/*k*-NN classifier, the optimization included the best choice for distance function (Euclidean or other), the *k*-size parameter and the weighting function (constant or other).

Tables 2 and 3 illustrate all the dataset/classifier combinations and the corresponding best-accuracy configurations for the full feature sets (i.e., no feature selection/optimization), which were used as the base for creating and evaluating the corresponding ensembles in the subsequent stage of the experiments.

Table 2: Dataset specifications and single-classifier (reference) accuracies (%). Mean and standard deviation values are based on five training realizations with full feature set.

Dataset	Training set size	Testing set size	Dataset Dimension	SVM accuracy	OBTC accuracy	w/*k*-NN accuracy
ringnorn	400	7000	20	**97.66 ± 0.22**	80.78 ± 2.41	77.00 ± 1.83
splice	1000	2175	60	85.29 ± 1.08	**92.98 ± 0.97**	77.88 ± 2.09
twonorm	400	7000	20	97.70 ± 0.15	76.56 ± 1.63	**97.85 ± 1.52**
waveform	400	4600	21	**90.10 ± 0.40**	80.91 ± 1.65	89.85 ± 1.50

Table 3: Single-classifier best configurations against datasets, based on five training realizations.

Dataset	best SVM configuration	best OBTC configuration	best w/*k*-NN configuration
ringnorn	kernel: RBF (σ=5) μ=0.016 *epsilon*=5.e-4	splt.func=*twoing* splt.limit=10	dist.func=*Euclidean* *k*-size=1 weight.func=*none*
splice	kernel: RBF (σ=42) μ=0.036 *epsilon*=5.e-4	splt.func=*deviance* splt.limit=10	dist.func=*Euclidean* *k*-size=15 weight.func=*Gaussian*
twonorm	kernel: RBF (σ=100) μ=0.008 *epsilon*=5.e-4	splt.func=*deviance* splt.limit=10	dist.func=*Euclidean* *k*-size=17 weight.func=*linear*
waveform	kernel: RBF (σ=20) μ=0.020 *epsilon*=5.e-4	splt.func=*Gini* splt.limit=10	dist.func=*Euclidean* *k*-size=21 weight.func=*none*

5.2 Classifier ensembles training

Using the best-configuration results obtained from the previous experimentation stage, two different base ensemble designs were employed in the second stage. Specifically, the feature subspace method was used to create dataset splits for K=5 and K=7 feature subsets, with the application of MANOVA as the feature ranking method for the complete (original) datasets. Each of these K-splits setups was applied to create ensembles with five or seven classifiers, each employing one of the three base types of classifiers (SVM, OBTC, w/k-NN). Subsequently, for each one of these ensemble setups, ten random realizations of training and testing subsets of the dataset were created, using the training/testing ratio also used for the corresponding single-classifier cases (see: Table 2). This procedure was employed in the same way for all four base datasets, for ensembles of five or seven (same type) classifiers.

For the training of the classifiers in any given ensemble setup, the model parameters used were the same with the ones calculated in the corresponding single-classifier case during the first experimentation phase. The main reason for not employing a full optimization procedure in this second stage was the fact that the subsequent comparison of the different combination rules was based in their *relative* differences in performance between them and not with the corresponding single classifier model (using the full feature set). Moreover, this procedure of optimizing every individual classifier in the ensemble would result in the increase of the total processing time with no actual benefit to the purposes of this particular study, since the goal here is to test the efficiency and robustness of the various combination rules in ensembles of weak or sub-optimally trained classifiers.

Table 4 illustrates the mean and standard deviation values of average and maximum (in parentheses) accuracy rates of ensemble members, for different choices of datasets, K-splits and base classifiers.

Table 4: Average classifier group accuracies (%) for all datasets and splits, against all the ensemble members in ten training realizations. The numbers enclosed in parentheses indicate the mean of the maximum-accuracy members in each corresponding ensemble, i.e., the average over only the top members across the ten training realizations.

Dataset	***K*-splits**	**Classifiers: SVM**	**OBTC**	**w/*k*-NN**
ringnorm	5	77.63 ± 2.13	77.42 ± 0.69	73.64 ± 0.55
		(81.10 ± 1.05)	*(79.29 ± 0.68)*	*(74.85 ± 0.52)*
	7	72.89 ± 1.04	74.30 ± 1.04	69.65 ± 0.43
		(80.53 ± 1.73)	*(78.21 ± 1.25)*	*(73.43 ± 0.91)*
splice	5	66.80 ± 0.98	72.57 ± 0.46	67.33 ± 0.78
		(78.94 ± 0.50)	*(88.39 ± 0.33)*	*(80.84 ± 0.71)*
	7	63.19 ± 1.30	68.70 ± 0.48	66.04 ± 0.34
		(77.17 ± 1.37)	*(82.49 ± 0.49)*	*(80.32 ± 0.41)*
twonorm	5	81.11 ± 0.11	73.82 ± 1.30	79.87 ± 0.16
		(84.11 ± 0.14)	*(75.34 ± 0.93)*	*(82.72 ± 0.34)*
	7	76.83 ± 0.65	71.78 ± 0.44	75.73 ± 0.20
		(82.09 ± 0.65)	*(75.26 ± 0.76)*	*(80.40 ± 0.51)*
waveform	5	75.72 ± 1.28	78.18 ± 0.63	79.92 ± 0.48
		(79.26 ± 0.59)	*(79.57 ± 0.62)*	*(81.43 ± 0.53)*
	7	70.25 ± 1.09	76.65 ± 0.49	77.64 ± 0.38
		(78.54 ± 1.03)	*(79.40 ± 0.46)*	*(80.85 ± 0.45)*

With respect to the dataset split process employed in this study, Table 4 demonstrates that the non-random (MANOVA ranking) feature subspace method produced more or less balanced ensembles, as the standard deviation on the mean and maximum single-member accuracies remained very low in all cases, regardless of the type of base classifier employed.

5.3 Testing of combination rules in the ensembles

In each training/testing cycle, the classification outputs from the pool of K classifiers were fed as input to each of the eight combination schemes (discussed in section 4.4) investigated in this study, producing the corresponding combined classification outputs of the ensemble.

It should be noted that the half-range decision threshold was used in all combination rules, i.e., no analytical optimization was conducted for T. This choice is justified by the results from previous studies [e.g., 41, 1], which support the assertion that the optimized T value rarely lies far from the half-range value. Furthermore, in the case of combination rules that employ local accuracy estimates (i.e., a pdf approximation), information about the shape and properties of the decision boundary of each classifier is already encoded partially in the (estimated) conditional probability of “correct” classification, and is used either directly (in the case of DCS-LA) or indirectly (in the weights of WMR).

In the sequel, the overall relative performance of each combination rule was determined in terms of ranking position for each case, i.e., according to its corresponding improvement over the mean group accuracy, for each dataset and K value employed. Specifically, a *weighted Borda* or *w/Borda* count method [60] was employed to attribute ten points to the top-ranked combination rule (first on the list of eight rules), nine points to the second (second on the list of eight rules), and so on. In case of a “tie” where two combination rules exhibited exactly the same performance, both got the same w/Borda points for the specific ranking position. The Borda and w/Borda count methods are often used in cases when an overall evaluation of classifiers or ensembles is required over a wide range of different configurations, datasets and average “grouped” performances, i.e., when direct aggregation of individual “group” success rates is not valid in terms of statistical context. In this study, as the average accuracies of all the models are more or less compact within-datasets but very different across-datasets (very different classification tasks), their *relative* ranking is much more informative than mean and standard deviation calculations of actual accuracy rates.

Table 5 illustrates the wBorda rankings of each combination rule, as well as the mean increase in accuracy (each cell corresponds to the average over ten training/testing realizations) over the mean accuracy of the individual members in the ensemble of SVM classifiers, for each dataset and K-split value employed.

Table 5: SVM ensemble results for all combination rules, datasets and *K*-splits. Improvements on average group accuracy (%) are presented in decimal numbers, while wBorda ranking values are presented as integers. The underlined wBorda values indicate top-ranking positions (10 points). All accuracy improvements were calculated as the difference of accuracy rates between the ensemble and the corresponding single-classifier configuration. Negative values indicate deterioration in performance. The combination rules are presented sorted against the SUM column, which represents the total sum of wBorda points assigned to each combination rule over all datasets and *K*-splits.

	ringnorm		splice		twonorm		waveform		SUM	MEAN	STDEV
	K=5	*K*=7	*K*=5	*K*=7	*K*=5	*K*=7	*K*=5	*K*=7			
STD: DCS-LA (no priors)	8	8	10	10	9	6	9	10	**70**	8.75	1.39
	16.96	20.25	19.69	22.04	14.55	17.28	8.06	12.87		16.46	4.53
WMR: logodds (adaptive)	7	6	8	8	8	10	10	9	**66**	8.25	1.39
	16.11	18.11	17.68	20.06	14.22	18.80	8.67	12.8		15.81	3.75
STD: DCS-LA (w/priors)	9	9	9	9	9	5	4	3	**57**	7.13	2.64
	17.27	20.43	19.17	21.89	14.55	17.22	-3.31	4.25		13.93	8.83
STD: simple average	10	10	4	4	10	4	7	8	**57**	7.13	2.80
	17.64	21.62	7.76	6.11	15.20	5.82	7.97	11.68		11.73	5.87
STD: LSE-w/average	6	5	7	6	7	7	8	7	**53**	6.63	0.92
	14.81	15.58	16.46	19.04	13.76	17.49	7.98	11.21		14.54	3.56
WMR: logodds (static)	4	4	6	7	7	9	6	5	**48**	6.00	1.69
	14.63	13.62	16.41	19.17	13.76	17.66	7.37	9.77		14.05	3.93
STD: simple majority	5	3	5	5	7	8	6	4	**43**	5.38	1.60
	14.81	12.57	12.28	13.92	13.76	17.64	7.37	7.34		12.46	3.55
STD: maximum	3	7	3	3	6	3	5	6	**36**	4.50	1.69
	14.41	18.68	1.68	-2.18	12.10	1.64	6.07	10.16		7.82	7.21

The same approach was applied to ensembles of OBTC and w/*k*-NN classifiers. Tables 6 and 7 illustrate the wBorda rankings and mean accuracy improvements of OBTC and w/*k*-NN ensembles, accordingly.

Table 6: OBTC ensemble results for all combination rules, datasets and *K*-splits. The adverted notation is the same as in Table 5.

	ringnorm		splice		twonorm		waveform		SUM	MEAN	STDEV
	K=5	*K*=7	*K*=5	*K*=7	*K*=5	*K*=7	*K*=5	*K*=7			
WMR: logodds (adaptive)	10	8	6	7	9	10	10	10	**70**	8.75	1.58
	14.25	15.34	18.93	21.49	15.00	18.44	7.46	9.78		15.09	4.69
WMR: logodds (static)	9	7	8	8	8	8	9	9	**66**	8.25	0.71
	14.17	14.88	19.16	21.58	14.40	17.84	7.10	9.23		14.76	4.86
STD: LSE-w/average	9	9	7	9	8	7	8	8	**65**	8.13	0.83
	14.17	15.40	18.98	21.65	14.40	17.83	7.05	9.17		14.83	4.87
STD: simple majority	9	6	5	5	8	9	9	9	**60**	7.50	1.85
	14.17	14.86	15.14	16.26	14.40	17.87	7.10	9.23		13.63	3.62
STD: simple average	8	10	4	4	10	6	7	7	**56**	7.00	2.33
	13.51	16.07	12.14	15.10	15.32	16.66	5.49	6.60		12.61	4.31
STD: DCS-LA (no priors)	7	4	10	10	7	4	6	6	**54**	6.75	2.31
	9.23	11.00	20.67	21.92	13.34	12.95	3.89	4.02		12.13	6.70
STD: DCS-LA (w/priors)	6	5	9	6	6	5	4	4	**45**	5.63	1.60
	8.34	11.18	20.57	20.90	13.33	12.99	-5.00	-6.71		9.45	10.40
STD: maximum	5	3	3	3	5	3	5	5	**32**	4.00	1.07
	0.52	-1.48	8.13	9.25	2.05	-1.58	-1.37	-3.46		1.51	4.73

Table 7: w/*k*-NN ensemble results for all combination rules, datasets and *K*-splits. The adverted notation is the same as in Table 5.

	ringnorm		splice		twonorm		waveform		SUM	MEAN	STDEV
	K=5	*K*=7	*K*=5	*K*=7	*K*=5	*K*=7	*K*=5	*K*=7			
WMR: logodds (adaptive)	8	8	8	7	10	10	10	10	**71**	8.88	1.25
	17.17	20.76	16.22	19.03	14.16	18.08	6.98	9.12		15.19	4.85
STD: DCS-LA (no priors)	9	9	9	10	7	6	6	6	**62**	7.75	1.67
	19.68	21.41	19.05	21.49	9.48	10.66	4.55	3.65		13.75	7.52
STD: DCS-LA (w/priors)	10	10	10	9	6	5	4	4	**58**	7.25	2.76
	20.00	21.79	19.30	21.24	9.45	10.64	-7.51	-8.96		10.74	12.62
STD: LSE-w/average	7	5	5	5	9	8	9	9	**57**	7.13	1.89
	13.47	16.99	13.91	17.34	14.08	17.85	6.77	8.86		13.66	4.02
WMR: logodds (static)	7	6	4	4	9	9	8	8	**55**	6.88	2.03
	13.47	17.02	13.78	17.09	14.08	17.93	6.76	8.82		13.62	4.01
STD: simple average	6	4	7	8	8	7	7	7	**54**	6.75	1.28
	3.03	-1.72	15.41	19.20	12.82	13.75	6.42	6.16		9.38	7.03
STD: simple majority	7	7	3	3	9	9	8	8	**54**	6.75	2.43
	13.47	17.03	9.39	12.73	14.08	17.93	6.76	8.82		12.53	3.96
STD: maximum	5	3	6	6	6	3	5	5	**39**	4.88	1.25
	-8.01	-9.72	15.29	18.44	5.82	2.07	4.07	2.07		3.75	9.85

Table 8 presents a summary of the wBorda rankings from Tables 5, 6 and 7. The list of all the combination rules is sorted according to their sum of wBorda points, i.e., their overall efficiency throughout all the base datasets (four) and *K*-splits (*K*=5, *K*=7). Additionally, based on the results from Tables 2 and 4 through 7, Table 9 presents a summary of the comparative performance of the best ensemble designs against the corresponding best single-classifier performance, for all datasets.

Table 8: Overall evaluation of all the combination rules, using the wBorda results from all the experiments. The wBorda sum, mean and standard deviation values for each combination rule were calculated across all the datasets, *K*-splits and classifiers. The list is sorted according to the wBorda sum (and mean) ranking position of each combination rule, from the best to the worst combination rule.

Combination rule	**w/Borda SUM**	**w/Borda MEAN**	**w/Borda STDEV**
WMR: logodds (adaptive)	207	8.63	1.38
STD: LSE-w/average	175	7.29	1.40
STD: DCS-LA (no priors)	173	7.21	2.19
WMR: logodds (static)	169	7.04	1.78
STD: simple average	167	6.96	2.14
STD: DCS-LA (w/priors)	160	6.67	2.41
STD: simple majority	157	6.54	2.11
STD: maximum	107	4.46	1.35

Table 9: Overall evaluation of the best ensemble designs against the best single-classifier configurations, for all datasets. The values in the rightmost column refer to the difference between the accuracy (%) of the best ensemble design and the corresponding best single-classifier accuracy for a specific dataset.

Dataset	**Best single-classifier configuration**	**Best ensemble designs**	**Best ensemble accuracy**	**Best accuracy difference**
ringnorm	97.66 (SVM)	SVM: simple average, *K*=5	**95.27**	-2.39
		OBTC: WMR (adaptive) all, *K*=5	91.67	
		w/*k*-NN: DCS-LA (w/priors), *K*=5	93.64	
splice	92.98 (OBTC)	SVM: DCS-LA (no priors), *K*=5	86.49	+0.26
		OBTC: DCS-LA (no priors), *K*=5	**93.24**	
		w/*k*-NN: DCS-LA (no priors), *K*=7	87.53	
twonorm	97.85 (w/*k*-NN)	SVM: simple average, *K*=5	**96.31**	-1.54
		OBTC: WMR (adaptive), *K*=7	90.22	
		w/*k*-NN: WMR (adaptive) all, *K*=5	94.03	
waveform	90.10 (SVM)	SVM: WMR (adaptive) all, *K*=5	84.39	-3.20
		OBTC: WMR (adaptive) all, K=7	86.43	
		w/*k*-NN: WMR (adaptive) all, *K*=5	**86.90**	

6. Discussion

The results from Table 8 clearly demonstrate the overall superior performance of the "adaptive" WMR model. Both the "adaptive" and the "static" versions of the WMR model show improved performance compared to the simple majority voting, as well as the maximum rule, for all three types of classifier ensembles (SVM, OBTC, w/*k*-NN).

The "adaptive" version of WMR also exhibited better performance compared to the simple averaging rule, as well as the weighted averaging rule using LSE-trained weights, i.e., both soft-output combination models. Thus, the "adaptive" WMR model performs equally well or better than simple soft-output averaging combination rules.

With regard to weighted versus non-weighted combination rules, all three weighted combination rules, i.e., the two WMR and the LSE-trained weighted average, have been clearly proven better than the non-weighted hard-output combination rules (i.e., maximum and simple majority). Moreover, in the overall evaluation, the "static" WMR outperformed the Bayesian-based combination rule using priors ("STD: DCS-LA (with priors)"). This essentially means that the WMR model is a very effective way of exploiting information about the classifiers' competencies, even when this information refers to global (i.e., prior) and not localized (i.e., posterior) probabilities. The overall performance of the WMR improved significantly when local accuracy estimates was used in the "adaptive" version of the model, reaching the top-ranking position over all the other combination rules, including the best soft-output ("STD: LSE-weighted average") and the best Bayesian-based ("STD: DCS-LA (no priors)") combination rules.

The w/Borda rankings from Tables 5 through 8 also demonstrate the robustness and stability of the each combination rule. For small values of standard deviation (close to one) in the corresponding w/Borda mean ranks, the relative ranking position of a combination rule against the others remains more or less the same. Thus, the "static" version of the WMR exhibited a consistently lower ranking position compared to the corresponding "adaptive" WMR model in general. Likewise, the "adaptive" WMR model is more stable than almost all the other combination rules (except maximum), including the LSE-weighted average, which exhibits more or less the same consistency and robustness as the "adaptive" WMR but with lower relative ranking.

In terms of the overall performance of the combination rules, results from Tables 5 through 7 demonstrate that in all cases the best combination rules increased the average success rates (Table 4) of the classifier pool significantly, up to +22% (mainly in the "ringnorm" and "splice" datasets), with larger relative improvements as the size of the ensemble increased from five to seven members, for all the three types of base classifiers. Furthermore, Table 9 shows that the performance of the best ensemble designs closely matched the performance of the corresponding best single-classifier configuration and even surpassed it ("splice" dataset). Although the WMR rule was not always selected as the best ensemble design, its overall performance and the top-ranking positions in Table 8 clearly demonstrate that it is inherently robust and consistently efficient.

The general behavior of almost all the ensembles was consistent with the theoretical background and experimentally verified the assertion that combining even *moderately* independent experts results to the improvement of their individual competencies [1]. Previous studies [61] have shown experimental evidence that optimal combination of SVM classifiers can be achieved through linear combination rules. Ensembles of SVM or other type of robust classifiers, as a combination of multiple simpler models, each using a 1/K portion (subspace) of the feature space of the dataset instead of a single classifier of the same type for the complete feature space, can be used to reduce the overall training effort significantly. In particular, for the SVM model case, kernel evaluation employs inner product between vectors, i.e., its complexity is directly proportional to the dimensionality (number of features) of the input vectors. Thus, feature space reduction, from F to F/K features, results in significant decrease in the overall complexity during training. A similar approach has also been examined in [61], where an ensemble of SVM classifiers has been used, trained with small training sets, instead of a single SVM trained with one large training set. Furthermore, there is evidence that such ensembles of kernel machines are more stable than the equivalent kernel machine itself and that their model need not be more complex than a simple linear combination of its member outputs [61], which is consistent with the theoretical assertion of the WMR formulation as the optimal aggregation model for any n-label voting task (see section 2.3). This reduction in training time, of course, has to be compared to the additional overhead of calculating a combination rule for every output vector from the classifier pool, as well as the total training time of the K SVM classifiers. This is one of the main reasons for preferring very simple, linear aggregation schemes with non-trained weights, such as the WMR, for the design of robust classifier ensembles.

It should be noted that using simple linear combination models, including weights that do not require iterative training, can be extremely useful in applications that require parallel and/or on-line updating. In the case of WMR, the combination rule is fully parallelizable, even when using local accuracy estimates in the weighting formula, since they are based on histogram calculations and not on iterative off-line optimization of the weights. Furthermore, the updating of the histogram can also be realized on-line, simply by adding new evaluation results as the model runs on new input data (without any new re-training of classifiers), and only if needed, i.e., when the new data invalidate the statistics of the previous histogram estimations.

7. Conclusions

In this paper, a game-theoretic framework for combining classifiers has been proposed. The adapted WMR has been, for the first time, presented as an alternative approach to design simple and efficient ensembles of voting classifiers, even when the conditional independence assumption is only moderately satisfied via feature subspace methods. Experimental comparative results have shown that such simple combination models for combining classifiers can be more efficient than typical rank-based and simple majority schemes, as well as simple soft-output averaging schemes in some cases. Moreover, when the weighting profiles required in the WMR are associated with the posterior (localized), instead of the prior (global), approximations of the classifiers' accuracies, the resulting ensemble can outperform many commonly used combination methods of similar complexity. The use of simple linear combination models that employ analytically computed weights may provide the necessary means to apply multi-expert classification schemes in parallel implementations with on-line updating capabilities. Therefore, the WMR can be asserted as a simple, yet effective tool in the palette for combining classifiers in an optimal, adaptive and theoretically well-defined framework.

Acknowledgements

The authors wish to thank professor Sergios Theodoridis, Dept. of Informatics & Telecommunications, Univ. of Athens (NKUA/UoA), for his contribution in the early stages of this work.

References

[1] L.I. Kuncheva, Combining Pattern Classifiers, John Wiley and Sons, New Jersey, 2004.

[2] G. Owen, Game Theory, third ed., Academic Press, San Diego, 1995.

[3] C. Montet, D. Serra, Game Theory and Economics, Palgrave Macmillan, New York, 2003.

[4] S. Nitzan, J. Paroush, Optimal decision rules in uncertain dichotomous choice situations, Intl. Econ. Rev. 23 (1982) 289–297.

[5] L. Shapley, B. Grofman, Optimizing group judgemental accuracy in the presence of independence, Public Choice 43 (1984) 329–343.

[6] S. Theodoridis, K. Koutroumbas, Pattern Recognition, third ed., Academic Press, San Diego, 2006.

[7] M. de Condorcet, An essay on the application of probability theory to plurality decision making: An election between three candidates, in: Sommerlad and Mclean, 1989, pp. 66-80.

[8] J. Kittler, M. Hatef, P.W. Duin, J. Matas, On combining classifiers, IEEE Trans. PAMI 20 (3) (1998) 226-239.

[9] I. Bloch, Information combination operators for data fusion: A comparative review with classification, IEEE Trans. SMC-A 26 (1996) 52-67.

[10] R.A. Jacobs, M.I. Jordan, S.J. Nowlan, G.E. Hinton, Adaptive mixtures of local experts, Neural Comp. 3 (1991) 79-87.

[11] R.A. Jacobs, Methods for combining experts' probability assessments, Neural Comp. 7 (1995) 867-888.

[12] J.A. Benediktsson, P.H. Swain, Consensus theoretic classification methods, IEEE Trans. SMC 22 (1992) 688-704.

[13] K. Woods, W.P. Kegelmeyer, K. Bowyer, Combination of multiple classifiers using local accuracy estimates, IEEE Trans. PAMI 19 (4) (1997) 405-410.

[14] R. A. Jacobs, Methods for combining experts' probability assessments, Neural Comp. 7 (1995) 867-888.

[15] V. Tresp, A Bayesian Committee Machine, Neural Comp. 12 (2000) 2719-2741.

[16] G. Rogova, Combining the results of several neural network classifiers, Neural Networks 7 (1994) 777-781.

[17] Y. Lu, Knowledge integration in a multiple classifier system, Appl. Intell. 6 (1996) 75-86.

[18] L. Kuncheva, J. Bezdek, R. Duin, Decision templates for multiple classifier fusion: An experimental comparison, Pat. Rec. 34 (2001) 299-314.

[19] H. Drucker, C. Cortes, L.D. Jackel, Y. LeCun, V. Vapnik, Boosting and other ensemble methods, Neural Comp. 6 (1994) 1289-1301.

[20] D. Tao, X. Tang, X. Li, X. Wu, Asymmetric bagging and random subspace for support vector machines-based relevance feedback in image retrieval, IEEE Trans. PAMI 28 (7) (2006) 1088-1098.

[21] T.K. Ho, The random subspace method for constructing Decision Forests, IEEE Trans. PAMI 20 (8) (1998) 832-844.

[22] L. Breiman, Random Forests, Machine Learning 45 (1) (2001) 5-32.

[23] J.J. Rodriguez, L.I. Kuncheva, C.J. Alonso, Rotation Forest: A new classifier ensemble method, IEEE Trans. PAMI 28 (10) (2006) 1619-1630.

[24] J. von Neumann, O. Morgenstern, Theory of Games and Economic Behavior, Princeton Univ. Press, Princeton, 1947.

[25] J. Nash, Non-cooperative games, Annals of Math. 54 (1950) 286-295.

[26] R. Ben-Yashar, S. Nitzan, The optimal decision rule for fixed-size committees in dichotomous choice situations: The general result, Intl. Econ. Rev. 38 (1997) 175–186.

[27] D. Karotkin, The network of weighted majority rules and weighted majority games, Games and Econ. Beh. 22 (1998) 299–315.

[28] W. Pierce, Improving reliability of digital systems by redundancy and adaptation, Ph.D. dissertation, Dept. Elect. Eng., Stanford Univ., Stanford, CA, 1961.

[29] L. Kuncheva, Fuzzy versus non-fuzzy in combining classifiers designed by boosting, IEEE Trans. Fuzzy Sys. 11 (6) (2003) 729-741.

[30] L. Shapley, B. Grofman, Optimizing group judgemental accuracy in the presence of interdependencies, Public Choice 43 (1984) 329-343.

[31] M. Mavroforakis, M. Sdralis, S. Theodoridis, A novel SVM Geometric Algorithm based on Reduced Convex Hulls, 18th Intl. Conf. on Pat. Rec. (ICPR 2006), Hong Kong, Aug. 2006.

[32] G. Raetch, T. Onoda, K.-R. Muller, Soft margins for AdaBoost, Machine Learning 42 (3) (2001) 287-320.

[33] S. Mika, G. Raetch, J. Weston, B. Scholkopf, K.-R. Muller, Fisher discriminant analysis with kernels, Neural Networks for Signal Processing IX, 1999, in: Proc. 1999 IEEE Signal Processing Society Workshop, 1999, pp. 41-48.

[34] ELENA project: On-line repository of databases for benchmarking machine-learning algorithms, ftp://ftp.dice.ucl.ac.be/pub/neural-nets/ELENA/databases

[35] I. Guyon, J. Makhoul, R. Schwartz, U. Vapnik, What size test set gives good error rate estimates?, IEEE Trans. PAMI 20 (1) (1998) 52-64.

[36] M. Stone, Cross-validatory choice and assessment of statistical predictions, J. R. Stat. Soc. 21 (1) (1974) 111-147.

[37] G.R. Dattatreya, L.N. Kanal, Decision trees in pattern recognition, in: L.N. Kanal, A. Rosenfeld (Eds.), Progress in Pattern Recognition 2, North Holland, 1985.

[38] C. Aviles-Cruz, A. Guerin-Dugue, J.L. Voz, D. Van Cappel, ELENA: Enhanced Learning for Evolutive Neural Architecture, ESPRIT Basic Research Project No.6891, Deliverable R3-B1-P: Databases, 1995, pp. 1-47.

[39] C.Jr. Traina, A.J.M. Traina, L. Wu, C. Faloutsos, Fast feature selection using fractal dimension, In XV Brazilian Database Symposium, Joao Pessoa, PA Brazil, 1 (2000) 158–171.

[40] B. Abrahao, L. Barbosa, Characterizing Datasets Using Fractal Methods, Dept. of Computer Science, Universidade Federal de Minas Gerais, Belo Horizonte, MG Brazil, 2003.

[41] H. Georgiou, M. Mavroforakis, S. Theodoridis , A game-theoretic approach to weighted majority voting for combining SVM classifiers, Int. Conf. on ANN (ICANN), 10-13 Sept. 2006, Athens, Greece, in: S.Kollias, et al. (Eds.): ICANN 2006, Part I, LNCS 4131, 2006, pp. 284-292.

[42] W. Cooley, P. Lohnes, Multivariate Data Analysis, John Willey & Sons, New York, 1971.

[43] V.N. Vapnik, Statistical Learning Theory, John Wiley & Sons, New York, 1998.

[44] N. Cristianini, J. Shawe-Taylor, An Introduction to Support Vector Machines and Other Kernel-Based Learning Methods, Cambridge University Press, UK, 2000.

[45] L. Breiman, J. Friedman, R. Olshen, C. Stone, Classification and Regression Trees, Wadsworth Intl., Belmont, CA, 1984, pp. 226-239.

[46] M. Mavroforakis, S. Theodoridis, Support vector machine classification through geometry, Proc. XII Eur. Sig. Proc. Conf. (EUSIPCO2005), Antalya, Turkey, Sep. 2005.

[47] M. Mavroforakis, S. Theodoridis, A geometric approach to support vector machine (SVM) classification, IEEE Trans. on Neural Networks 17 (3) (2006) 671-682.

[48] R. Paredes, E. Vidal, Weighting Prototypes: A New Editing Approach, Proc. 15th Intl. Conf. Pat. Rec., 2 (2000) 25-28.

[49] R. Short II, K. Fukunaga, The optimal distance measure for nearest neighbor classification, IEEE Trans. Inf. Theory 27 (5) (1981) 622-627.

[50] C. Domeniconi, J. Peng, D. Gunopulos, Locally adaptive metric nearest-neighbor classification, IEEE Trans. PAMI 24 (9) (2002) 1281-1285.

[51] J. Kittler, M. Hatef, P.W. Duin, Combining classifiers, IEEE Proc. 13th Int. Conf. Pat. Rec. (ICPR ’96), 1996, pp. 897-901.

[52] A. Jain, R. Duin, J. Mao, Statistical pattern recognition: a review, IEEE Trans. PAMI 22 (2000) 4–37.

[53] G. Fumera, F. Roli, Performance analysis and comparison of linear combiners for classifier fusion, In: Proc. 16th Intl. Conf. on Pat. Rec., Canada, 2002.

[54] G. Giacinto, F. Roli, Adaptive selection of image classifiers, Proc. of ICIAP ’97, Lecture Notes in Computer Science 1310 (1997) 38-45.

[55] L. Didaci, G. Giacinto, F. Roli, G.L. Marcialis, A study on the performances of dynamic classifier selection based on local accuracy estimation, Pat. Rec. 38 (2005) 2188-2191.

[56] B. Zadrozny, C. Elkan, Transforming classifier scores into accurate multiclass probability estimates, in: Proc. 8th Int. Conf. on Knowledge Discovery and Data Mining (KDD’02), 2002.

[57] M. Li, I.K. Sethi, Confidence-based classifier design, Pat. Rec. 39 (2006) 1230-1240.

[58] C. de Boor, A practical guide to splines, second ed., vol.27, in: Applied Mathematical Sciences, Springer-Verlag, 2001.

[59] M.R. Spiegel, J. Liu, Mathematical handbook of formulas and tables, second ed., McGraw-Hill, USA, 1999.

[60] J.R. Parker, Rank and response combination from confusion matrix data, Information Fusion 2 (2001) 113–120.

[61] T. Evgeniou, M. Pontil, A. Elisseeff, Leave one out error, stability and generalization of voting combinations of classifiers, Machine Learning 55 (2004) 71–97.

[62] R. Ben-Yashar, J. Paroush, Optimal decision rules for fixed-size committees in polychotomous choice situations, Social Choice and Welfare 18 (2001) 737–746.

[63] K. Kailing, H.-P. Kriegel, P. Kroger, S. Wanka, Ranking Interesting Subspaces for Clustering High Dimensional Data, Proc. of Europ. Conf. on Principles and Practice of Knowledge Discovery in Databases (PPKD), Dubrovnic, Croatia, 2003 – Lecture Notes in Artificial Intelligence (LNAI), Vol. 2838, 241-252.

[64] J.A. Swets, R.M. Pickett, Evaluation of Diagnostic Systems: Methods from Signal Detection Theory, Academic Press, New York, 1982.

[65] T. Fawcett, An introduction to ROC analysis, Pat. Rec. Lett., 27 (2006) 861-874.

[66] N.M. Adams, D.J. Hand, Comparing classifiers when the misallocation costs are uncertain, Pat. Rec., 32 (1999) 1139-1147.

[67] M. Zuckerman, P. Faliszewski, Y. Bachrach, E. Elkind, Manipulating the quota in weighted voting games, Artificial Intelligence 180-181 (2012) 1-19.

[68] L. Breiman, Some Properties of Splitting Criteria, Statistics Department, University of California, Berkeley, 1992.