

A MILP Formulation for Setwise Max-Margin Learning

Paolo Viappiani¹

Abstract. In max-margin learning, the system aims at establishing a solution as robust as possible. In this paper we extend the idea of max-margin learning to cases where the goal is to produce a set of solutions, instead of a single one. In particular, we focus on the problem of preference elicitation, but we believe our idea could be adapted to other situations. We present a MILP formulation for our “setwise” max margin learner and discuss current ongoing works.

1 Max-margin Learning for Preference Optimization

The notion of margin is important in several machine learning algorithms. By maximizing the margin, the learner ensures robustness in the solution. In particular, a max-margin approach has been proposed for optimization with a partially specified preference profile and elicitation of further preferences [2]. Automatic assessment of preferences is an important component of many artificial intelligence systems [3].

In this paper, we show how to extend the idea of max-margin for optimizing a set and we provide a MILP formulation. The result of the computation can be used to generate a set of recommendations and to generate a “diverse” set of feasible utility functions. This can be useful for further elicitation as we may ask the user a choice query based on this set (“Among these items, which one do you prefer?”).

First of all, we review the standard (singleton) max-margin learning problem. We assume a multi attribute feature space \mathbf{X} (we assume possible options are specified by configuration constraints) and that the user preferences are encoded by a vector of weights $\mathbf{w} = (w_1, \dots, w_m)$; the utility of an option \mathbf{y} is then $\mathbf{w} \cdot \mathbf{y}$.

The weight vector \mathbf{w} is unknown but we are given constraints (that encode a “feasible region”). A pairwise comparison between a more preferred product \mathbf{y}_+ and a less preferred product \mathbf{y}_- is encoded by a constraint $\mathbf{w} \cdot (\mathbf{y}_+ - \mathbf{y}_-) \geq 0$. Given a number of such pairwise comparisons, represented by a set D (the “learning set”), and, possibly, other additional constraints (representing some kind of “prior” knowledge, as for example $w_2 > 0.3$) we want to find a vector \mathbf{w} of parameters such that all constraints are satisfied and does that by the largest margin.

As shown by [2] the problem can be easily formulated as a linear program. The shared margin is represented by the decision variable M , that we want to maximize. We require,

for each pairwise statement in D , to have $\mathbf{w} \cdot (\mathbf{y}_+^h - \mathbf{y}_-^h) \geq M$.

$$\begin{aligned} \max_{M, \mathbf{w}} \quad & M \\ \text{s.t.} \quad & \mathbf{w} \cdot (\mathbf{y}_+^h - \mathbf{y}_-^h) \geq M \quad \forall (\mathbf{y}_+^h, \mathbf{y}_-^h) \in D \end{aligned}$$

The optimization can be made tolerant to contradictory or hard-to-satisfy constraints; this can be done by adding explicit slack variables ε_h .

$$\max_{M, \mathbf{w}} \quad M - \alpha \sum_h \varepsilon_h \quad (1)$$

$$\text{s.t.} \quad \mathbf{w} \cdot (\mathbf{y}_+^h - \mathbf{y}_-^h) \geq M - \varepsilon_h \quad \forall (\mathbf{y}_+^h, \mathbf{y}_-^h) \in D \quad (2)$$

Where α is a parameter that controls the tolerance to a small number of violated constraints or constraints satisfied by a margin smaller than M . If $\alpha=0$, then all original constraints must be strictly satisfied.

Once the optimization has been solved, the learned \mathbf{w} can be used to choose a recommendation \mathbf{x}^* by maximizing $\mathbf{w} \cdot \mathbf{x}$. If the space of feasible options (or products) is a configuration space encoded by linear constraints, the recommendation can be computed by solving a linear program. We refer to $\mathbf{x} \in \text{Feasible}(\mathbf{X})$ to specify that an option must satisfy the configuration constraints.

In interactive settings, one should decide the query to ask next. [2] suggests asking queries whose associated hyperplane divide the feasible region in roughly equal part, but to do so they require considering all possible pairs of options; this approach can only be feasible in small datasets and do not apply in configuration domains. Instead, a viable method for large configuration spaces should be able to retrieve a set of options in a single optimization.

2 Setwise Max-Margin

We extend *maxmargin* optimization of preference weights to sets. Given, as before, a number of pairwise statements D and possibly additional constraints (the “learning set”) our goal is to find a number of utility weight vectors $\mathbf{w}^1, \dots, \mathbf{w}^k$ (with k given), consistent with our knowledge about the user, but that are representative of different parts of the weights’ feasible region.

In order to do that, we require, as before, a shared margin M to hold for all constraints associated with prior knowledge and responses. We also use the same margin to enforce each vector \mathbf{w}^i to be as “farther away” from each other as possible. This insight is realized by introducing a set of options $\{\mathbf{x}^1, \dots, \mathbf{x}^k\}$ as decision variables (of the same cardinality k). For each i , we impose \mathbf{x}^i to be the best option among them

¹ Aalborg University, Denmark, email: paolo@cs.aau.dk




Figure 1. Pairwise max-margin ($k = 2$). The feasible region in the vector space (black constraints) is divided in two parts by the red hyperplane. The two weight vectors w_1 and w_2 (red circles), together with the hyperplane, are chosen to maximize the minimum of the margins (gray arrows).

(offering higher utility) when evaluated according to \mathbf{w}^i - and we require these constraints to hold by at least a margin M . The optimization is the following:

$$\max_{M, \mathbf{w}^i, \mathbf{x}^i} M - \alpha_1 \sum_h \varepsilon_h - \alpha_2 \sum_{i,j} \varepsilon'_{i,j} \quad (3)$$

$$\text{s.t. } \mathbf{w}^i \cdot (\mathbf{y}_+^h - \mathbf{y}_-^h) \geq M - \varepsilon_h \quad \forall (\mathbf{y}_+^h, \mathbf{y}_-^h) \in D, \forall i \in [1, k] \quad (4)$$

$$\mathbf{w}^i \cdot (\mathbf{x}^i - \mathbf{x}^j) \geq M - \varepsilon'_{i,j} \quad \forall j \neq i; i, j \in [1, k] \quad (5)$$

$$\mathbf{x}^i \in \text{Feasible}(\mathbf{X}) \quad \forall i \in [1, k]$$

Note that, we are choosing the options $\mathbf{x}^1, \dots, \mathbf{x}^k$ and the weights $\mathbf{w}^1, \dots, \mathbf{w}^k$ simultaneously: since we want to maximize M , the optimizer will be better off by choosing a set of outcomes \mathbf{x}^i that divide the weight space roughly equally, and the utility functions such each \mathbf{w}^i should lie (intuitively) near the centre of each subregion.

This initial formulation is problematic, as we have quadratic terms. However, there is a solution: in fact, by using integer programming tricks, the problem can be formulated as a mixed integer linear program (MILP).

$$\max M - \alpha_1 \sum_h \varepsilon_h - \alpha_2 \sum_{i,j} \varepsilon'_{i,j} \quad (6)$$

$$\text{s.t. } \mathbf{w}^i \cdot (\mathbf{y}_+^h - \mathbf{y}_-^h) \geq M - \varepsilon_h \quad \forall (\mathbf{y}_+^h, \mathbf{y}_-^h) \in D, \forall i \in [1, k] \quad (7)$$

$$\sum_z A_z^i - B_z^{i,j} \geq M - \varepsilon'_{i,j} \quad \forall j \neq i; i, j \in [1, k] \quad (8)$$

$$A_z^i \leq w \uparrow x_z^i \quad \forall i \in [1, k], z \in [1, m] \quad (9)$$

$$A_z^i \leq w_z^i \quad \forall i \in [1, k], z \in [1, m] \quad (10)$$

$$B_z^{i,j} \geq w_z^i - C \cdot (1 - x_z^j) \quad \forall j \neq i \in [1, k], z \in [1, m] \quad (11)$$

$$B_z^{i,j} \geq 0 \quad \forall j \neq i \in [1, k], z \in [1, m] \quad (12)$$

$$\mathbf{x} \in \text{Feasible}(\mathbf{X}) \quad \forall i \in [1, k]$$

In the optimization, the decision variables are the following:

- M is the size of the shared margin
- $\mathbf{w}^1, \dots, \mathbf{w}^k$ is a set of utility vectors; each vector defined over m attributes (features): $\mathbf{w}^i = (w_1^i, \dots, w_m^i)$
- $\mathbf{x}^1, \dots, \mathbf{x}^k$ is a set of configurations (options) with each configuration $\mathbf{x}^i = (x_1^i, \dots, x_m^i)$; each element x_z^i is binary
- ε slack variables to represent cost of unsatisfied constraints in D ; ε' slack variables to represent cost of not respecting the margin M when choosing the hyperplanes
- \mathbf{A}^i encodes the vector $(w_1^i x_1^i, \dots, w_m^i x_m^i)$, the element-by-element product of \mathbf{w}^i and \mathbf{x}^i : $A_z^i = w_z^i x_z^i$

- $\mathbf{B}^{i,j}$ encodes $(w_1^i x_1^j, \dots, w_m^i x_m^j)$, the element-by-element product of \mathbf{w}^i and \mathbf{x}^j

In addition to the decision variables, we make use of the following parameters:

- k is the size of the set
- α_1, α_2 control the tolerance with respect to violated constraints (we can differentiate the tolerance for the two different types of constraints).
- D is the *learning set*: a set of pairwise comparisons known to the system
- $w \uparrow$ is any upper bound of the value of the utility weights
- C is an arbitrary large number
- $\mathbf{y}_+^h, \mathbf{y}_-^h$ are pairs of configurations (options) in the learning set D , with \mathbf{y}_+^h preferred to \mathbf{y}_-^h

The solver aims at setting M , that is a decision variable, as large as possible. Constraint 5 of the original program is replaced by constraint 8, enforcing M to be smaller than $\sum_z A_z^i - B_z^{i,j}$ for any i, j . In order to maximize M , the solver will try to keep the A_z^i as large as possible and the $B_z^{i,j}$ as small as possible. The fact that $\sum_z A_z^i - B_z^{i,j}$ evaluates to $\mathbf{w}^i \cdot (\mathbf{x}^i - \mathbf{x}^j)$ is achieved by setting additional constraints. Constraints 9 and 10 together force each \mathbf{A}^i to be the element-by-element product of \mathbf{w}^i and \mathbf{x}^i ($A_z^i = w_z^i x_z^i$): if $x_z^i = 0$ then (constraint 9) also A_z^i must be 0; otherwise (if $x_z^i = 1$) $A_z^i \leq w_z^i$ (constraint 10) but because of the objective function maximizing M , and each A constrain M , the solver will set A_z^i to w_z^i . Similarly constraints 11 and 12 together make it so that $\mathbf{B}^{i,j}$ encodes $(w_1^i x_1^j, \dots, w_m^i x_m^j)$, the element-by-element product of \mathbf{w}^i and \mathbf{x}^j (as C is an arbitrary large constant, constraint 11 is binding only when x_z^j is 1, otherwise it is always satisfied).

3 Discussion

This paper describes an extension of maximum margin learner to sets. We are currently planning to perform experiments to evaluate the efficiency of this approach. We will compare to baseline methods, maximizing product diversity for a single learned weight vector. We are interested in computation time, but also in the quality of the elicitation using our method in an interactive setting. In particular, we are interested in comparing this approach for generating recommendation sets with others based on Bayesian inference [5] and minimax regret [4]. We are also interested in learning approaches ensuring sparsity in the parameter space [1].

REFERENCES

- [1] Paolo Campigotto, Andrea Passerini, and Roberto Battiti. Active learning of combinatorial features for interactive optimization. In Carlos A. Coello Coello, editor, *LION*, volume 6683 of *Lecture Notes in Computer Science*, pages 336–350. Springer, 2011.
- [2] Krzysztof Gajos and Daniel S. Weld. Preference elicitation for interface optimization. In Patrick Baudisch, Mary Czerwinski, and Dan R. Olsen, editors, *UIST*, pages 173–182. ACM, 2005.
- [3] Bart Peintner, Paolo Viappiani, and Neil Yorke-Smith. Preferences in interactive systems: technical challenges and case studies. *AI Magazine*, 29(4):13–24, 2008.
- [4] Paolo Viappiani and Craig Boutilier. Regret-based optimal recommendation sets in conversational recommender systems. In *Proceedings of the 3rd ACM Conference on Recommender Systems (RecSys09)*, pages 101–108, New York, 2009.
- [5] Paolo Viappiani and Craig Boutilier. Optimal bayesian recommendation sets and myopically optimal choice query sets. In *Advances in Neural Information Processing Systems 23 (NIPS)*, Vancouver, 2010.