Fairness in Assignment Markets with Dual Decomposition

Bert Huang BERT@CS.UMD.EDU
University of Maryland, College Park, MD 20742 USA

Abstract

In this work, we present a market design for assignment problems that computes a globally optimal solution by adjusting incentives. Such markets can help in settings such as the assignment of peer-reviewers to submitted academic articles, assignment of tutors to students, or online matchmaking services. In these settings, each assignment has some reward value, and existing strategies for achieving high global reward involve either adjustments to greedy choices by agents or global optimization of estimated reward values. The benefit of maintaining a market is that we combine benefits from these methods in a principled way. The agents make incentivized greedy decisions, which is ideal because they understand their reward functions best, while the incentives push their decisions toward the global optimum. We update the incentives by relating the assignment market to the standard dual of the (b-) matching linear programming relaxation. We evaluate our proposed system on simulations, demonstrating that the market quickly improves the global reward.

1. Introduction

Many real-world situations involve the general problem of *assignment*, where entities are assigned to other entities. The assignments are associated with different reward values and are constrained. For example, the academic peer-review process involves the assignment of volunteer reviewers to submitted articles, and the associated reward value could be the overall quality of the review process. Since the volunteer reviewers have limited time, and since the articles need a fair panel of reviewers, there are constraints on the number of assignments to each entity. Existing strategies for these assignments combine greedy, free-market ideas, such as allowing agents to bid on assignments, with more global objectives, such as solving a maximum reward matching with estimated rewards, estimated using the combination of collected and actively queried data (Charlin et al., 2011; 2012). Greedy strategies allow agents to decide which assignments have highest reward but tend to find poor global solutions, in part because of the tendency in natural systems for rewards to be skewed. On the other hand, global matching-based strategies depend on the quality of estimated rewards and suffer when these estimates are inaccurate. In this preliminary work, we present a method for setting incentives based on a global optimization algorithm that allows agents to make iterative greedy decisions, but such that the incentives force these greedy decisions toward the global optimum. The method is based on a dual decomposition for the generalized matching linear programming relaxation, and we use a simple subgradient update rather than running the full auction algorithm update (Bertsekas, 2009). Our simulations show that the global reward significantly improves with only a small number of iterations.

One important aspect of the method we propose is that it shares computation between human agents and a market maker. Loosely speaking, the scheme separates the computation into parts that each component is able to do well. Humans are good at making greedy decisions over complicated reward functions, and computers have trouble modeling these decisions or their reward functions. The market we propose allows humans to make incentivized greedy decisions, while the (approximate) global optimization is done by adjusting the incentives. In the subgradient method we describe later, we further do not require the human agents to report quantitative scores of their reward functions. This aspect is also useful, since humans are better at making decisions than quantifying a scale for their preferences.

1.1. Related Work

The technical portion of this work stems from the original motivational explanation of the famous auction algorithm (Bertsekas, 2009) and its variants, which are too many to fully reference here. Various advances in recent years found an efficient algorithm for computing maximum weight b-matchings by encoding the problem as a probability distribution and running loopy belief propagation to find the most likely state (Bayati et al., 2005; 2006; 2008; 2011; Huang & Jebara, 2007; 2011; Sanghavi et al., 2007). Among the various results produced in this line of research, Sanghavi et al. (2007) and Bayati et al. (2006; 2011) identify deep connections between the belief propagation algorithms and linear programming relaxations of the problems including the auction algorithm, relating the convergence guarantees for belief propagation to the tightness of the linear programming relaxations. Manshadi et al. (2013) recently developed a dual linear program method for solving maximum generalized matchings in the MapReduce framework, allowing solution of massive scale problems. In this work, our aim is not to reach the global solution, but to move incentives such that we take a few steps toward it, improving on the base greedy solution. Since we expect to run our market updates for a limited number of rounds, we propose our approach as an improvement on a loose approximation, so the tightness of the linear programming relaxation does not necessarily play an important role in the usefulness of our approach.

Finally, though the auction algorithm and the belief propagation approaches have guaranteed convergence in a bounded number of iterations, their update rules require the reward values of the decisions made at each round. In this work, we focus on the simpler subgradient update, which does not require this information. In future work, we will explore the design of a market that can support acquisition of the raw reward values, but this design will need mechanisms to avoid asking too much work from the agents. The subgradient update we propose here only requires the agents to make greedy decisions, something humans tend to do easily, while quantification of the reward for their decisions is significantly more complex.

Various studies analyze using incentives to induce greedy agents to produce a global computation. For example, Judd et al. (2011) conduct behavioral experiments in social network formation games, identifying parameters that affect the global properties of the resulting solutions found by players.

Finally, our proposed ideas are complementary to another important thread in the study of methods for improving peer-review assignments. Charlin et al. (2011) developed a now-prominent system for learning and inferring review quality using analysis of reviewers' publications. They similarly propose solving the global assignment optimization with matching-like constraints. They also provide techniques for incorporating human feedback into their matching (Charlin et al., 2012). Their work has noticeably improved reviewer assignments, and we aim to build upon their improvements. Our proposed idea can similarly be interpreted as inserting a human into the loop of this matching process.

2. An Incentivized Assignment Market

In our assignment market setting, each entity to be matched is or has an agent that makes decisions to achieve maximum reward. The rewards combine the hidden, latent reward for the assignment and additional incentives we add as the market maker. Each agent has a cardinality requirement for its assignments. For example, in the peer review setting, reviewers estimate the quality of the review they can provide for each paper by considering the overlap in the paper's topics and their areas of expertise, as well as various factors such as personal interests. We expect that the reward function for review quality is complex and difficult to quantify, even for the agents themselves. Reviewers also have time constraints, so they have some ideal number of papers that they can confidently devote time to review, and deviating from this time constraint can be detrimental to the overall quality of reviews.

We formalize this setting as follows. Let the entities be indexed from 1 to n.[1] The desired cardinality, or degree, of each entity is stored in vector $\mathbf{b}$, such that b_i is the desired degree for entity i. A cost matrix $\mathbf{W}$ contains the cost of each assignment, such that the cost of assigning i to j is W_{ij}. Given full knowledge of the cost matrix, the globally optimal assignment is

$$\begin{aligned} \min_{\mathbf{A}\in\{0,1\}^{n\times n}} \quad & \sum_{ij} W_{ij} A_{ij} \qquad (1)\\ \text{s.t.} \quad & b_j = \sum_i A_{ij}, \forall i, \text{ and } b_i = \sum_j A_{ij}, \forall j. \end{aligned}$$

In words, the optimization finds the minimum cost binary assignment matrix $\mathbf{A}$ such that the sums along the rows and the sums along the columns are equal to the desired degree vector $\mathbf{b}$. The problem is translation invariant and symmetric, so we can negate $\mathbf{W}$ or add an arbitrary constant to all entries to obtain a maximum-reward optimization. For mathematical convention, we pose the objective as a cost minimization.

We consider the linear programming (LP) relaxation

[1] If the entities have a bipartite structure, as in the reviewer-paper setting, the indices may be ordered such that reviewers are numbered 1 through m and papers are $m + 1$ through n.

of this integer linear program. Note that for bipartite matchings, if the optimum is unique, then the LP relaxation is tight (Bayati et al., 2008; 2011; Bertsekas, 2009; Sanghavi et al., 2007). The objective simply becomes $\min_{\mathbf{A}\in[0,1]^{n\times n}} \sum_{ij} W_{ij}A_{ij}$ subject to the same constraints as above.

2.1. Dual Decomposition for Matching

The objective in Equation 1 is decomposable in an elegant manner that corresponds to a market setting. By using Lagrangian relaxation for only one direction of the degree constraint (w.l.o.g. we relax the row sum constraint), we obtain the dual objective,

$$\begin{aligned} \max_{\Lambda} \min_{\mathbf{A}\in[0,1]^{n\times n}} \quad & \sum_{ij} W_{ij}A_{ij} + \sum_i \lambda_i \left(b_i - \sum_j A_{ij} \right) \\ \text{s.t.} \quad & b_j = \sum_i A_{ij}, \forall j, \end{aligned}$$

which can be simplified to

$$\begin{aligned} \max_{\Lambda} \min_{\mathbf{A}\in[0,1]^{n\times n}} \quad & \sum_{ij} \left(W_{ij} - \lambda_i \right) A_{ij} + \Lambda^\top \mathbf{b} \\ \text{s.t.} \quad & b_j = \sum_i A_{ij}, \forall j. \end{aligned} \tag{2}$$

For a fixed dual variable Λ, the inner minimization is a simple, greedy optimization: each i'th agent chooses its b_i least incentivized-cost assignments, where the incentivized cost is $W_{ij} - \lambda_i$. Once all agents make their greedy decisions, we have an assignment matrix $\mathbf{A}$ that satisfies the column-sum constraint, but not necessarily the row-sum constraint.

To solve the outer optimization, a straightforward approach is to perform subgradient ascent. A subgradient for Λ is

$$\nabla_{\lambda_i} \quad = \quad b_i - \sum_j A_{ij}. \tag{3}$$

The standard subgradient ascent update at iteration t is

$$\lambda_i^t \leftarrow \lambda_i^{t-1} + \alpha_t \left(b_i - \sum_j A_{ij} \right),$$

where α_t represents a decaying learning rate schedule such as $1/t$.

In the bipartite setting, one useful observation is that the dual variable update depends only on the "incoming" assignment selections. That is, if we consider the peer-review setting, the greedy decisions by the reviewers affect the dual variables for the papers, and the "greedy decisions by the papers" (which do not make much sense in this context) have no effect on the relevant updates.[2] We will take advantage of this simplification in our experiments, where we will only simulate the reviewer selections.

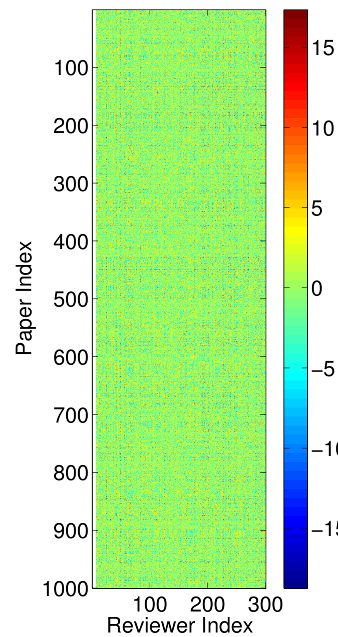

Figure 1. Simulated quality scores for reviews based on random feature vectors representing papers and reviewers.

Another important aspect of this dual objective is that it is translation invariant to constant shifts in the dual variables. This invariance allows the market maker to adjust the incentives depending on its goal. One can shift incentives so that the market maker pays for fairness, ensuring that incentives are always in the agents' favor. Alternatively, one can shift the incentives such that they create a zero-sum game for the agents, providing fairness at no cost to the market maker, and requiring some agents to pay for assignments, while others receive payment. Finally, if the market maker is performing a for-profit assignment service, one could shift the dual variables such that the incentives contribute to a profit margin.

2.2. Real-world Implementation Discussion

In this section, we discuss implementation ideas. We focus on how an assignment market can benefit peer

[2]This is a side effect of the general form we have used to describe assignment problems. We could have alternatively used a strictly bipartite formulation, which would make this more obvious, but would limit the applicable settings of this method.

reviewing systems, but we will briefly discuss other potential settings as well.

The idea we consider most attractive about our method is that the agents themselves compute the reward-decisions. Since human reward functions can be arbitrarily complex, and humans themselves tend to be inconsistent in quantifying their reward functions, the incentivized market setting changes the cognitive task for agents from reporting reward values to simply making decisions.

To implement such a system in a volunteer scientific review process, one needs a commodity to offer as incentive. In many settings, money is an obvious incentive but is likely infeasible for scientific conferences and journals. Instead, simulated currency may be a plausible option, creating a "gamification" of the review process. For example, a market can assign a relatively inexpensive prize to the agent who has the most simulated currency at the end of the review process. Prizes could be awarded via a raffle or some other payout system. In the scientific community, we should expect that, since all volunteer reviewers are working toward the common goal of improving the quality of the publication, only a small amount of incentive is necessary to skew reviewers' paper bids.

Matchmaking is another interesting potential application for this method. Tutor-student matching services, dating services, and online freelancing services matching employees to job opportunities include processes that can be improved with an iterative incentive market. In these settings, customers may be willing to pay money to participate in a fair assignment, or some customers may take payment for choosing the less desirable assignments. Economically, these matchmaking services may want to pay highly desirable customers, because their inclusion in the assignment pool makes the service attractive to new customers.

3. Simulation Experiments

In our simulations, we test the subgradient incentive update, showing that a few rounds of incentive updates and re-selection by the agents strongly improves the overall (simulated) review quality. We generate ten-dimensional feature vectors for 1,000 papers and 300 reviewers. We require that all papers have three assigned reviewers and all reviewers read ten papers. We assume true review quality is measured by taking an inner product of the feature vectors. Figure 1 is a visualization of the resulting review quality values.

We assume each reviewer knows the quality of review he or she can provide for each paper but has no knowledge of any other reviewer's quality scores. Thus, each reviewer simply selects the ten papers for which he or she will produce the highest-quality reviews. We plot the resulting assignment and the degree distribution for papers in Figure 2. The degree distribution is rather skewed, with a few papers that many reviewers want to read, and many papers that few or no reviewers want to read. We believe this simulation is fairly realistic in this manner.

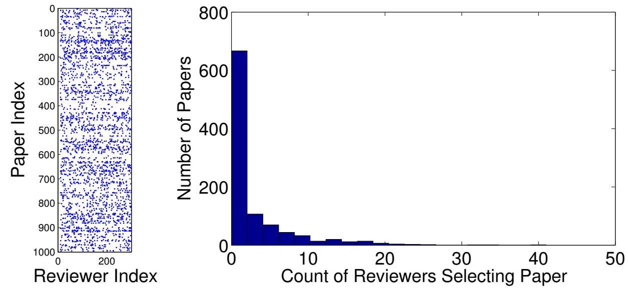

Figure 2. Results from greedy bidding by reviewers. Left: assignment matrix selected by simulated reviewers using true review quality. Right: degree distribution for papers from greedy assignment. Many papers have too few reviewers and a small number of papers interest too many reviewers.

To correct this skew, we consider a few strategies. Recall that we should have no knowledge of the true review quality. In each strategy, we compute a surrogate quality score and solve the global optimization to find a feasible assignment. The simplest surrogate is to randomly set weights as an arbitrary tie-breaker. We generate random unit variance Gaussian scores for all pairings and add 10.0 to the score of any agent-selected pairing. The second approach is to use an approximate estimate of review quality and compute a global solution based on that, i.e., a simulation of the Toronto review matching system (Charlin et al., 2011). We implement this idea by adding random Gaussian noise of varying amounts (variance in $\{1.0, 2.0, 5.0\}$) to the feature vectors of reviewers and papers, then computing the estimated review quality and solving the assignment problem using these noisy review qualities. We consider using only the noisy review quality estimates, as well as combining them with bids by adding 10.0 to the estimated quality score of each agent-selected assignment.

Finally, we simulate the market approach by asking reviewers to select papers greedily, updating incentives, and then re-polling the reviewers. Since we do not expect to reach a global optimum, we consider all the correction strategies listed above, noting that the market

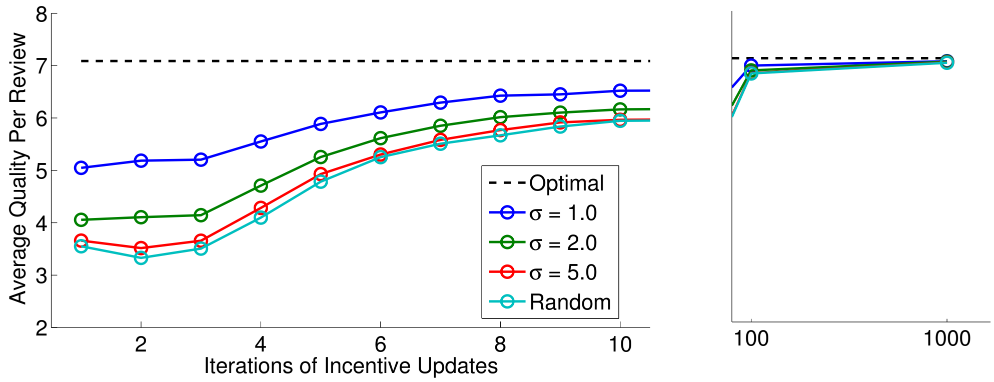

Figure 3. Review quality after iterations of incentive updates. Iteration 1 is equivalent to asking reviewers to bid once. The different solid lines represent different surrogate quality matrices used to complete the assignment, including noisy estimates simulated here by adding different levels of independent Gaussian noise, and a random matrix that simply serves as an arbitrary tiebreaker. We include the result at high iteration counts 100 and 1,000 for reference, though such high counts are unlikely to be feasible in a real system.

approach should require less correction. We report results using various numbers of iterations of re-polling.

3.1. Results

We measure the average quality over all 3,000 reviews using the different strategies. The optimal average obtained by solving the matching with full quality information is 7.09. Using the noisy quality estimates alone without reviewer bidding scores an average quality of 3.53, 1.43, and 0.34, respectively, for noise levels 1.0, 2.0, and 5.0. Adding reviewer bids improves the overall quality, producing average quality scores of 3.5 with random tie breaking, and 5.05, 4.06, and 3.66 for the different noisy surrogates. Running five iterations of incentive updates improves each by 0.84, 1.2, and 1.27, and running ten iterations improves by 1.47, 2.1, and 2.3, respectively. For reference, we include the scores when running 100 and 1,000 iterations of incentive updates, which are very close to the optimal score; all methods are within 0.08 of the optimal score at 1,000 iterations. These results are displayed in Figure 3.

Overall, updating incentives in this simple scheme seems to significantly lift overall quality. For context, consider that the scale for average quality score of reviews is 0.0 for randomly assigned reviews. Comparing with real-world intuition of the quality of reviews we would expect from random assignment, even within a scientific subfield, and noting that the best possible single review in our simulation scores 17.3, the improvement seems quite significant.

4. Discussion

In this preliminary work, we present a market making strategy for enforcing fairness and improving overall reward in an assignment problem. Like the classical auction algorithm, the method is based on a dual decomposition of the maximum weight b-matching LP. However, since our goal is to run a few iterations of simple updates, we use the naive subgradient update for the dual variables. We present simulation results demonstrating that this strategy can produce significant improvement of overall review quality in peer-review assignments.

The idea for this type of market-based computation partially stems from the notion that humans tend to be very good at greedy selection and not very good at global optimization. Human reward functions are extremely complicated and difficult to model, and humans themselves are not able to quantify them, despite their ability to make decisions using them. Our market-based idea isolates the simple piece of the algorithm that is needed to enforce global optimality, and absolves itself from the more complex reward function, allowing the human agents themselves to do that part of the computation.

While the primary example we present here is in the context of volunteer reviewing, this market-based approach has many more applications in assignment settings, as discussed in Section 2.2.

We are exploring various directions of future work. We are exploring ways to design interfaces and a market system that can obtain the necessary information for the true auction algorithm update, which is guaranteed to improve the dual and has fast convergence. The guarantees of the auction algorithm update suggest it will produce more dramatic improvements with fewer iterations of the market, but since it uses the actual reward values, requires careful design of a market mechanism that does not produce too much cognitive strain on the agents. We are also exploring different variants of the objective function, such as ones that incorporate inequality constraints or models that include the effect of deviation from the ideal number of assignments on review quality.

References

Bayati, M., Shah, D., and Sharma, M. Maximum weight matching via max-product belief propagation. In *IEEE International Symposium on Information Theory*, 2005.

Bayati, M., Shah, D., and Sharma, M. A simpler max-product maximum weight matching algorithm and the auction algorithm. In *IEEE International Symposium on Information Theory*, 2006.

Bayati, M., Shah, D., and Sharma, M. Max-product for maximum weight matching: Convergence, correctness, and LP duality. *IEEE Trans. Inf. Theor.*, 54(3):1241–1251, 2008.

Bayati, M., Borgs, C., Chayes, J., and Zecchina, R. Belief-propagation for weighted b-matchings on arbitrary graphs and its relation to linear programs with integer solutions. *SIAM Journal on Discrete Mathematics*, 25:989–1011, 2011.

Bertsekas, D. Auction algorithms. In Floudas, C. and Pardalos, P. (eds.), *Encyclopedia of Optimization*, pp. 128–132. Springer, 2009. ISBN 978-0-387-74758-3.

Charlin, L., Zemel, R., and Boutilier, C. A framework for optimizing paper matching. In *Uncertainty in Artificial Intelligence*, 2011.

Charlin, L., Zemel, R., and Boutilier, C. Active learning for matching problems. In Langford, J. and Pineau, J. (eds.), *Proceedings of the International Conference on Machine Learning (ICML)*, ICML '12, 2012.

Huang, B. and Jebara, T. Loopy belief propagation for bipartite maximum weight b-matching. In *Artificial Intelligence and Statistics*, 2007.

Huang, B. and Jebara, T. Fast b-matching via sufficient selection belief propagation. In *Artificial Intelligence and Statistics*, 2011.

Judd, S., Kearns, M., and Vorobeychik, Y. Behavioral conflict and fairness in social networks. In *Internet and Network Economics*, 2011.

Manshadi, F., Awerbuch, B., Gemulla, R., Khandekar, R., Mestre, J., and Sozio, M. A distributed algorithm for large-scale generalized matching. In *Proceedings of the VLDB Endowment*, 2013.

Sanghavi, S., Malioutov, D., and Willsky, A. Linear programming analysis of loopy belief propagation for weighted matching. In *Neural Information Processing Systems 20*, 2007.