My primary research interests are in Machine Learning, Algorithmic Game Theory, and Algorithms, as well as on the interactions between these areas and other areas of Computer Science and Economics. The common thread of my work is a rigorous study of real world problems, involving both the development of mathematical models and the design of algorithms and tools for their analysis.

1 Brief Overview

New Frameworks and Algorithms for Machine Learning I am particularly excited by problems that explore new frontiers of learning, and as a consequence a significant part of my work has focused on providing theoretical models for important new learning paradigms which are not captured by existing theoretical frameworks. Over the years, Machine Learning has grown into a broad discipline that has produced fundamental theories of learning processes, as well as learning algorithms that are routinely used in commercial systems. The primary theoretical advances have been for *supervised* learning problems [24], where a target function (e.g., a classifier) is estimated using only *labeled examples*. For example, in spam detection an automatic classifier to label emails as spam or not would be trained using a sample of previous emails labeled by a human user. However, for most contemporary practical problems there is often useful additional information available in form of cheap and plentiful *unlabeled* data: e.g., unlabeled emails for the spam detection problem. As a consequence, there has recently been substantial interest in using unlabeled data to improve learning. Several different algorithmic approaches have been developed and numerous successful experimental results have been reported. However, the underlying assumptions of these methods are quite different and their effectiveness cannot be explained by standard learning models. In my work I provide *foundational* theoretical understanding for these new learning techniques. For example, I have developed a comprehensive theoretical framework [4, 5] that provides a unified way for thinking about Semi-Supervised Learning (a learning paradigm in which the learning algorithm can use unlabeled examples in addition to labeled examples). My model appears in a recent book about Semi-Supervised Learning [5] and it can explain when and why unlabeled data can help in many of the specific methods in all the other chapters of the book. In the context of Active Learning (a paradigm in which the algorithm can interactively ask for the labels of unlabeled examples of its own choosing), I have presented theoretical justification for margin-based algorithms [14] which have proven quite successful in practical applications, e.g., in text classification [26].

Another important component of my work is the development of more intuitive and more operational explanations for well-established learning paradigms, for which a solid theory did exist, but it was too abstract and disconnected from practice. In particular, in the context of Kernel methods (a state of the art technique for supervised learning), I have developed a theory of learning with similarity functions [3] that provides theoretical justification for the common intuition that a good kernel function is one that acts as a good measure of similarity. This theory is more general and involves more tangible quantities than those used by the traditional analysis.

Finally, I am also enthusiastic about exploring new perspectives on classic learning problems. For instance, in the context of Clustering I have developed new algorithms and a powerful model [9] that directly addresses what kind of information a clustering algorithm needs in order to produce a highly accurate clustering of the data.

Algorithmic Game Theory A second major focus of my research has been on Algorithmic Game Theory. My work here includes an exciting application of machine learning techniques to *automate* aspects of Mechanism Design and formally address the problem of market analysis, as well as development of pricing algorithms with improved guarantees over previous methods.

Developing algorithms for a highly distributed medium such as the Internet requires considering the objectives of the various parties in the system. As a consequence, Mechanism Design has become an increasingly important part of algorithmic research and Computer Science more generally in recent years. Mechanism Design can be thought of as a distinct form of algorithm design, where a central entity must perform some computation (e.g., resource allocation) under the constraint that the agents supplying the inputs have their own interest in the outcome of the computation. As a result, it is desirable that the employed procedure be *incentive compatible*, meaning that it should be in each agent's best interest to report truthfully, or to otherwise act in a well-behaved manner. Typical examples of such mechanisms are auctions of products (e.g., software packages) where the central entity would use inputs (bids) from the agents in order to allocate goods in a way that maximizes its revenue. Most of the previous work on incentive compatible mechanism design for revenue maximization has been focused on very restricted settings [19, 25] (e.g., one item for sale and/or single parameter agents), and many of the previous incentive compatible mechanisms have been "hand-crafted" for the specific problem at hand. In my work, I use techniques from Machine Learning to provide a *generic reduction* from the incentive-compatible mechanism design question to more standard algorithmic questions, for a wide variety of revenue-maximization problems, in an unlimited supply setting [6, 11]. These results are featured in a recent book on Algorithmic Game Theory [20].

I also work on applying online learning techniques for designing good online pricing mechanisms [2], as well as on pure algorithmic pricing problems [2, 7, 10]. More broadly, I am interested in systems of interacting self-interested agents in which learning, game theory, and mechanism design all play a role.

I will describe below some of my current work in detail, as well as some future plans.

2 New Frameworks and Algorithms for Machine Learning

My research has been at the forefront of several exciting foundational developments in areas of significant activity and practical importance in Machine Learning. The variety of results in my work relies on a very diverse set of insights and techniques from Algorithms and Complexity, Empirical Processes and Statistics, as well as Geometry and Embeddings.

Semi-Supervised Learning Semi-Supervised Learning has become an important area in Machine Learning and numerous heuristic approaches have been developed [16, 13, 23, 12]. These approaches are each based on different assumptions about how the target function relates with the underlying distribution: for example, that it partitions data by a wide margin as in Transductive SVM [23] or that data contains redundant sufficient information as in Co-training [16, 13]. The main theoretical difficulty, however, is that the standard discriminative learning models (PAC and Statistical Learning Theory frameworks [21, 27]) do not capture the underlying assumptions of these existing methods, neither do they explain how and why unlabeled data can be of help. In generative-model settings, on the other hand, it is relatively easy to accommodate the use of unlabeled data theoretically; however, these results typically require extremely strong assumptions essentially implying that there is only one natural distinction to be made for a given (unlabeled) data distribution. In joint work with Avrim Blum [4, 5], I have provided a PAC-style framework that bridges positions: I have extended the PAC model in a way that allows one to express not only the form of the target function, but also the relationship that one hopes the target function and underlying distribution will possess, but without going as far as is done in generative models. This model captures many of the ways in which unlabeled data is typically used and it provides a unified framework for analyzing why and under what conditions unlabeled data can help.

Active Learning In joint work with Alina Beygelzimer and John Langford [1], I have designed the first Active Learning procedure that works in the presence of arbitrary forms of noise. This procedure relies only upon the assumption that samples are drawn i.i.d. from some underlying distribution and it makes no assumptions about the mechanism producing noise (e.g., class/target misfit, fundamental randomization, etc.). In joint work with Tong Zhang and Andrei Broder [14], I have provided a generic analysis of a natural margin-based active learning strategy (that queries points near the hypothesized decision boundary), with a more detailed analysis (both in the realizable case and in a specific noisy setting) for a few important scenarios. Most recently, I have been investigating an asymptotic model [15], and have shown that contrary to the common belief Active Learning does help asymptotically in most natural learning problems.

Learning with Kernels and More General Similarity Functions Kernel methods have become especially popular in recent years, both because they are very useful in practice for dealing with many different kinds of data, and because they have a solid theoretical foundation. This theory is based on the fact that the *convergence rates* of all kernelizable learning algorithms depend only on the *margin* of the best separator in the implicit ϕ-space, and not on the dimension of the ϕ-space itself. As a consequence, kernel functions have been perceived as implicit mappings to a high-dimensional space that allow one to "magically" get the power of that space without having to pay for it.

In joint work with Avrim Blum and Santosh Vempala [8], I have used Random Projection techniques to help "demystify" kernel functions. I have additionally shown that in the presence of a large margin, a kernel can instead be *efficiently* converted into a mapping to a low dimensional space; in particular, I have given an efficient procedure that, given black-box access to the kernel and unlabeled data, generates a small number of features that approximately preserve both separability and margin. This result has two major implications. Conceptually, it suggests that designing a good kernel function is much like designing a good feature space. From a practical perspective it provides an alternative to "kernelizing" a learning algorithm: rather than modifying the algorithm to use kernels, one can instead construct a mapping into a low-dimensional space using the kernel and the data distribution, and then run an un-kernelized algorithm over examples drawn from the mapped distribution.

I have further developed a theory of learning with *more general similarity functions* [3] that are not necessarily legal kernels. This theory provides conditions on the suitability of a similarity function for a given learning problem in terms of more tangible and more operational quantities than those used by the standard theory of kernel functions. These conditions are both sufficient for learning and satisfied by the usual large-margin notion of a good kernel function; however, they do not require reference to implicit high dimensional spaces, nor that the similarity function be positive semi-definite. This framework provides the first rigorous explanation for why a kernel function that is good in the large-margin sense can also formally be viewed as a good measure of similarity, thereby giving formal justification to a common intuition about kernels.

Clustering via Similarity Functions Problems of clustering data from pairwise similarity information are ubiquitous in science. A typical example task is to cluster a set of emails or documents according to some criterion (say, by topic) by making use of a pairwise similarity measure among data objects (perhaps for measuring the fraction of important words that two documents have in common). In joint work with Avrim Blum and Santosh Vempala [9], I present a theoretical study of the clustering problem that directly addresses the fundamental question of how good the similarity measure must be in terms of its relationship to the desired ground-truth clustering (e.g., clustering by topic) in order to allow an algorithm to cluster well. Very strong properties are needed if the goal is to produce a single approximately-correct clustering; however, I show that if we relax the objective and allow the algorithm to produce a hierarchical clustering such that desired clustering is close to some *pruning* of this tree (which a user could navigate), then we can develop a general theory of natural properties that are sufficient for clustering via various kinds of algorithms. This framework is as an analogue of the PAC learning model for clustering, where the natural object of study, rather than being a concept class, is instead a property of the similarity information with respect to the desired ground-truth clustering.

Future Work on Machine Learning

A short term goal is to fully characterize the settings where active learning does help (asymptotically) in the existing models. A long term direction is to *broaden* the range of real world learning problems solvable with *interactive* learning by designing models and algorithms that allow more powerful types of queries than those considered in traditional active learning models.

One natural ultimate goal of my work on learning with similarity functions is to understand what are the weakest conditions for a similarity function that are sufficient for learning, and more generally to understand the space of conditions that suffice for learning from limited labeled data. Another interesting related direction is analyzing algorithms that augment the native features (say the words in a document) with a small number of additional features representing the similarity of the current example with each of a pre-selected set of initial documents. Such algorithms have proven extremely useful in practice [18], and they can be theoretically motivated by our current work, but the existing theory does not fully capture the approaches that seem to work best in practice.

I also plan to further develop my Clustering framework and to analyze new properties of similarity functions motivated by real world clustering problems arising in Vision or Biology, as well as to extend the model to accommodate alternate forms of limited but interactive feedback. The ultimate long term goal of this line of work is to provide a unified framework for data clustering that will allow researchers from many different areas to convert their intuition about their own problem into a choice of the most appropriate clustering algorithm for their needs.

I am also interested in developing new algorithms and theory for other important new paradigms of machine learning, (e.g., Transfer Learning), as well as new connections between machine learning and other areas (as described in Sections 3 and 4).

3 Algorithmic Game Theory

Mechanism Design and Machine Learning In joint work with Avrim Blum, Jason Hartline, and Yishay Mansour [11], I have shown how Model Selection and Sample Complexity techniques in Machine Learning can be used to convert difficult mechanism design problems to more standard algorithmic questions for a wide range of pricing problems. In doing so, I obtained a unified approach for considering a variety of profit maximizing mechanism design problems, such as the problem of auctioning a digital good, the attribute auction problem (which includes many discriminatory pricing problems), and the problem of item pricing in unlimited supply combinatorial auctions. My results substantially generalize the previous work on Random Sampling mechanisms [19] by both broadening the applicability of such mechanisms (e.g., to multi-parameter settings), and by simplifying and refining the analysis. From a learning perspective, these settings present several unique challenges and particularities: the loss function is discontinuous and asymmetric, and the range of bidders' valuations may be large.

Algorithms for Pricing Problems My work on pricing problems has been focused on developing better algorithms for pricing items for sale in order to maximize revenue in various settings [2, 7, 10]. For example, I have provided an $O(k)$ approximation algorithm for item pricing in unlimited supply combinatorial auctions with single minded bidders who each want at most k items [2, 7]; this algorithm improves the $O(k^2)$ bound of Briest and Krysta [17], and it can naturally be adapted to the online setting.

In recent work joint with Avrim Blum and Yishay Mansour [10], I completely characterize the power of a natural single posted price algorithm. I show that it achieves revenue within a logarithmic factor of the total social welfare in unlimited supply settings for buyers with general valuation functions (not just single-minded or unit-demand as known before). As no good approximation is possible for general valuations in the limited supply setting, I study the natural case of subadditive valuation functions instead. I show that here a random single price achieves revenue within a $2^{O\left(\sqrt{\log n \log\log n}\right)}$ factor of the total social welfare, thus providing the best revenue guarantees known for subadditive buyers for any item pricing scheme. I also complement this result with a surprising nearly matching lower bound.

Future Work on Algorithmic Game Theory

I plan to continue working on designing pricing mechanisms and algorithms with good revenue or efficiency guarantees. I am particularly excited by the connections between Machine Learning and Economic theory. Extending my work in reducing incentive compatible mechanism design to standard algorithmic questions, I plan to investigate under what conditions active learning techniques can be used to provide even better guarantees. I also want to explore the problem in an online setting, where bidders arrive one at a time. Standard online learning methods for "combining expert advice" [22] and generalizations can be applied for simple problems such as the auction of a single digital good. However, a generic reduction that works for such broad classes of problems as our reduction in the batch setting [11] is an intriguing open question that I plan to attack. Such methods would have a variety of applications, e.g., they could be used to provide better pricing schemes for various transactions on the Internet.

I am also interested in exploring connections between Online Learning and Equilibrium concepts. From a mechanism design perspective an interesting generic question is the following: given a repeated game that a mechanism designer has some power to adjust (e.g., say in an auction there is flexibility to smooth out the rules in some way), can one construct a game such that regret-minimizing learning algorithms or other adaptive natural strategies will produce desirable behavior? Moreover, in a complex auction setting, can a mechanism designer make adaptive strategies available to the bidders in a way that is beneficial both for the auctioneer and for the bidders? In a different direction, it is known that in general regret-minimizing behavior need not perform any better than the worst Nash or correlated equilibrium. However, it is conceivable that in natural settings by providing new feature information to the algorithms about the game being played one can produce behavior that avoids bad equilibria and yields more socially optimal results.

4 Final Words

I am broadly interested in systems of interacting self-interested agents in which learning, game theory, and mechanism design all play a role. A challenging long term goal (which might involve both learning and game theoretic techniques) is to investigate the development of learning algorithms that act as modules of a larger system. Most machine learning models view learning as a standalone process, focusing on prediction accuracy as the measure of performance. However, when a learning algorithm is placed in a larger system (perhaps in settings dominated by other adaptive algorithms acting in their own users' interests), other important aspects might have to be considered.

I intend to contribute to the advancement of Machine Learning and Algorithmic Game Theory by creatively developing new algorithms and theory for important problems. Moreover, I am determined to explore connections between my main areas of expertise and other research areas. I look forward to expanding my circle of collaborators to researchers working in completely new fields. I do believe that science can only progress if we teach each other the lessons each of us has learned.

References

[1] M.-F. Balcan, A. Beygelzimer, and J. Langford. Agnostic Active Learning. Proceedings of the 23rd International Conference on Machine Learning (ICML) 2006. Invited to appear in a special issue (on Learning Theory) of the *Journal of Computer and System Sciences*, 2007.

[2] M.-F. Balcan and A. Blum. Approximation Algorithms and Online Mechanisms for Item Pricing. *Theory of Computing*, 2007. Preliminary version appeared in the Proceedings of the 7th ACM Conference on Electronic Commerce (EC) 2006.

[3] M.-F. Balcan and A. Blum. On a Theory of Learning with Similarity Functions. Proceedings of the 23rd International Conference on Machine Learning (ICML) 2006.

[4] M.-F. Balcan and A. Blum. A PAC-style Model for Learning from Labeled and Unlabeled Data, Proceedings of 18th Annual Conference on Learning Theory (COLT) 2005.

[5] M.-F. Balcan and A. Blum. An Augmented PAC Model for Semi-supervised Learning, Book Chapter in *Semi-Supervised Learning*, Chapelle, O., Zien, A., and Scholkopf, B. (Eds.), MIT Press, 2006.

[6] M.-F. Balcan and A. Blum. Mechanism Design, Machine Learning, and Pricing Problems. SIGecom Exchanges 2007.

[7] M.-F. Balcan, A. Blum, H. Chan, and M. T. Hajiaghayi. A Theory of Loss Leaders: Making Money by Pricing Bellow Cost. Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE) 2007.

[8] M.-F. Balcan, A. Blum, and S. Vempala. Kernels as Features: On Kernels, Margins, and Low-dimensional Mappings, *Machine Learning Journal*.Preliminary version appeared in the Proceedings of the 15th International Conference on Algorithmic Learning Theory (ALT) 2004.

[9] M.-F. Balcan, A. Blum, and S. Vempala. A Learning Theoretic Framework for Clustering via Similarity Functions, Technical Report, CMU-CS-07-142, 2007.

[10] M.-F. Balcan, A. Blum, Y. Mansour. Item Pricing for Revenue Maximimization. Technical Report, CMU-CS-07-111, 2007.

[11] M.-F. Balcan, A. Blum, J. Hartline, and Y. Mansour. Reducing Mechanism Design to Algorithm Design via Machine Learning. *Journal of Computer and System Sciences*, invited to the Special Issue on Learning Theory, 2007. Preliminary version appeared in the Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS) 2005, under the title “Mechanism Design via Machine Learning”.

[12] M.-F. Balcan, A. Blum, P. Choi, J. Lafferty, B. Pantano, M. R. Rwebangira, and X. Zhu. Person Identification in Webcam Images: An Application of Semi-Supervised Learning. ICML 2005 Workshop on Learning with Partially Classified Training Data.

[13] M.-F. Balcan, A. Blum and K. Yang. Co-Training and Expansion: Towards Bridging Theory and Practice. Proceedings of the Eighteenth Annual Conference on Neural Information Processing Systems (NIPS) 2004.

[14] M.-F. Balcan, A. Broder, and T. Zhang. Margin Based Active Learning. Proceedings of 20th Annual Conference on Learning Theory (COLT) 2007.

[15] M.-F. Balcan, E. Even-Dar, S. Hanneke, M. Kearns, Y. Mansour, and J. Wortman Asymptotic Active Learning, NIPS 2007 Workshop on Principles of Learning Problem Design.

[16] A. Blum and T. Mitchell. Combining Labeled and Unlabeled Data with Co-Training, Proceedings of the 11th Annual Conference on Computational Learning Theory (COLT) 1998.

[17] P. Briest and P. Krysta. Single-Minded Unlimited Supply Pricing on Sparse Instances, Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) 2006.

[18] E. Gabrilovich and S. Markovitch. Feature Generation for Text Categorization Using World Knowledge. Proceedings of The 19th International Joint Conference on Artificial Intelligence (IJCAI) 2005.

[19] A. Goldberg, J. Hartline, A. Karlin, M. Saks, A. Wright. Competitive Auctions. *Games and Economic Behavior*, 2006.

[20] A. Karlin and J. Harltine. Profit Maximization in Mechanism Design. Book Chapter in Algorithmic Game Theory, Nisan N., Roughgarden T., Tardos E., and Vizarani V. (Eds.), MIT Press, 2007.

[21] M. Kearns and U. Vazirani. An Introduction to Computational Learning Theory. MIT Press, 1994.

[22] N. Littlestone and M. K. Warmuth. “The Weighted Majority Algorithm”. *Information and Computation*, 1994.

[23] T. Joachims. Transductive Inference for Text Classification using Support Vector Machines. Proceedings of the International Conference on Machine Learning (ICML) 1999.

[24] T. Mitchell. The Discipline of Machine Learning. Essay, 2006.

[25] R. Myerson. Mathematics of Optimal Auction Design. *Operations Research*, 1981.

[26] S. Tong and D. Koller. Support Vector Machine Active Learning with Applications to Text Classification. *Journal of Machine Learning Research*, 2001.

[27] V. N. Vapnik. Statistical Learning Theory. John Wiley and Sons Inc., 1998.