Aspects of Submodularity Workshop

Titles and Abstracts

1. **Jeff Bilmes, University of Washington**

Title: Learning Submodular Mixtures; and Active/Semi-Supervised Learning

Abstract: We discuss several recent applications of submodularity to machine learning. First, we present a class of submodular functions useful for document summarization. We show the best ever results on for both generic and query-focused document summarization on widely used and standardized evaluations. We then further improve on these results using a new method to learn submodular mixtures in a large-margin setting. The method uses a submodular loss function, and has risk-bound guarantees in terms of the loss function optimization's approximation ratio. Second, we present recent results for joint active semi-supervised learning with submodular functions. Here, a submodular function parameterizes an objective that consists of an active-learning and a semi-supervised learning part, and that when optimized minimizes a deterministic upper bound on the error. We present an algorithm that approximately minimizes this upper bound.

Joint work with Andrew Guillory and Hui Lin.

2. **Chandra Chekuri, UIUC**

Title: Multicommodity Flows and Cuts in Polymatroidal Networks

Abstract: The well-known maxflow-mincut theorem for single-commodity flows in directed networks is of fundamental importance in combinatorial optimization. Flow-cut equivalence does not hold in the multicommodity case. A substantial amount of research in the algorithms community has focused on understanding the gap between multicommodity flows and associated cuts. Poly-logarithmic gap results have been established in various settings via tools from network decomposition and metric embeddings.

In this talk we describe work that obtains flow-cut gap results in *polymatroidal* networks. In these networks there are submodular capacity constraints on the edges incident to a node. Such networks were introduced by Lawler & Martel and Hassin in the single-commodity setting and are closely related to the submodular flow model of Edmonds and Giles. The well-known maxflow-mincut theorem for a single-commodity holds in polymatroidal networks, however, the multicommodity case has not been previously explored. Our work is primarily motivated by applications to information flow in wireless networks although the technical results we obtain are of independent interest as they generalize known results in standard networks. The results highlight the utility of line embeddings suggested by Matousek and Rabinovich to round the dual of the flow-relaxation rewritten with a convex objective function (obtained from the Lovasz-extension of a submodular function).

Joint work with Sreeram Kannan, Adnan Raja and Pramod Viswanath from UIUC ECE Department.

3. **Satoru Fujishige, Kyoto University**

Title: Dual Consistency and Cardinality Constrained Polytopes

Abstract: We introduce a concept of dual consistency of systems of linear inequalities with full generality.

We show that a cardinality constrained polytope is represented by a certain system of linear inequalities if and only if the systems of linear inequalities associated with the cardinalities are dual consistent.

Typical dual consistent systems of inequalities are those which describe polymatroids, generalized polymatroids, and dual greedy polyhedra with certain choice functions.

We show that the systems of inequalities for cardinality-constrained ordinary bipartite matching polytopes are not dual consistent in general, and give additional inequalities to make them dual consistent.

Moreover, we show that ordinary systems of inequalities for the cardinality-constrained (poly)matroid intersection are not dual consistent, which disproves a conjecture of Maurras, Spiegelberg, and Stephan about a linear representation of the cardinality-constrained polymatroid intersection.

(Joint work with Jens Ma{\ss}berg (Research Institute for Discrete Mathematics, University of Bonn))

4. **Gagan Goel, Google Research**

Title: Polyhedral Clinching Auctions and the AdWords Polytope

Abstract: Motivated by ad auctions, a mechanism design problem that has come to the fore is that of designing auctions in the presence of budget-constrained bidders. Recent results in this direction extend Ausubel's clinching auction to give Pareto-optimal auctions for specific allocation scenarios such as multi-unit supply (Dobzinski et al) and certain matching markets (Fiat et al). A natural question one must ask is: For what all allocation scenarios can we design a Pareto-optimal auction in the presence of budget-constrained bidders? In this work, we give a Pareto-optimal auction for any allocation scenario that can be described using a polymatroid. Our auction extends Ausubel's clinching auction. We also show that a very general model of Adwords that includes multiple slots and multiple keywords can be described using a polymatroid. Finally we give some impossibility results for more general allocation scenarios. As a byproduct, we also get an impossibility result for multi-unit auctions with decreasing marginal utilities, thus resolving a conjecture by Ausubel stating that a variation of the clinching auction should work for this setting. This conjecture was later reinforced by Dobzinski et al.

This is a joint work with Vahab Mirrokni and Renato Paes Leme.

5. **Michel X. Goemans, MIT**

Title: (Poly)matroids and Integrality Gaps for Hypergraphic Steiner Tree Relaxations

Abstract: The undirected Steiner tree problem — the problem of connecting a given set of vertices (possibly using additional, Steiner vertices) at minimum total cost — is a classical, hard, combinatorial optimization problem. In a major development last year, Byrka, Grandoni, Rothvoss and Sanita provided the first approximation algorithm based on repeatedly solving a strong linear programming relaxation of the problem. Their randomized algorithm has a performance guarantee of ln(4)~1.39 in general and 73/60 for quasi-bipartite instances. Their analysis also shows that the integrality gap of their "component" LP relaxation is at most 1.55, the first bound less than 2 on any integrality gap for the problem.

In this talk, we derive a deterministic approximation algorithm with the same performance guarantees and an improved analysis providing a proof that the integrality gap is at most ln(4) in general and at most 73/60 for quasi-bipartite graphs. Our algorithm requires solving the relaxation only once, and both the algorithm and analysis rely heavily on (poly)matroidal properties of the relaxation. Various other features of the relaxation, including the sparsity of extreme solutions and an efficient separation algorithm, are also given.

This is joint work with Neil Olver, Thomas Rothvoss and Rico Zenklusen.

6. Satoru Iwata, Kyoto University

Title: Minimum Linear Ordering with Submodular/Supermodular Functions

Abstract: This talk addresses the Minimum Linear Ordering Problem (MLOP). Given a nonnegative set function on a finite set, find a linear ordering on the ground set such that the sum of the function values for all the suffixes is minimized. This problem generalizes well-known problems such as the Minimum Linear Arrangement, Min Sum Set Cover, Minimum Latency Set Cover, and Multiple Intents Ranking.

Extending a result of Feige, Lovasz, and Tetali (2004) on Min Sum Set Cover, we show that the greedy algorithm provides a factor 4 approximate optimal solution when the cost function is supermodular.

We also present a factor 2 rounding algorithm for MLOP with a monotone submodular cost function, using the convexity of the Lovasz extension.

These are among very few constant factor approximation algorithms for NP-hard problems formulated in terms of submodular/supermodular functions.

This is a joint work with Prasad Tetali and Pushkar Tripathi.

7. Kamal Jain, e-Bay

Title: Verification of Li & Li Conjecture on a Single Graph.

Abstract: Li and Li conjectured in 2004 that network coding does not give higher throughput in multiple unicast communication in undirected graphs.

This is a beautiful conjecture and still vastly open. In 2006, in a joint work with Vijay Vazirani, Raymond Yeung, and Gideon Yuval, we confirmed the conjecture on the smallest hard graph, namely K-3,2; i.e. a complete bipartite graph on 3 vertices on one side and 2 on the other side. The verification of the conjecture uses submodularity properties of entropy function. It is a hope that people who are interested in submodular functions can help make a significant progress on the conjecture.

8. Andreas Krause, ETH Switzerland

Title: Submodular Function Optimization in Sensor and Social Networks

Abstract: Many applications in sensor and social networks involve discrete optimization problems. In recent years, it was discovered that many such problems have submodular structure. These problems include optimal sensor placement, informative path planning, active learning, influence maximization, online advertising and structure learning. In contrast to most previous approaches, submodularity allows to efficiently find provably (near-)optimal solutions. In this tutorial, I will give examples of submodular optimization problems arising in sensor and social networks, discuss algorithms for solving these problems and present results on real applications. I will also discuss recent work in online and adaptive optimization of submodular functions in these domains.

9. Andrei Krokhin, Durham University (UK)

Title: Submodularity and the Complexity of Constraint Satisfaction

Abstract: The constraint satisfaction problem (CSP) is a general framework in which one can express, in a natural way, a wide variety of combinatorial problems.

The basic aim in a CSP is to decide whether there is an assignment of values to a given set of variables, subject to constraints on the values which can be assigned simultaneously to certain specified subsets of variables. There are several natural optimization versions of CSP, which all amount to optimizing certain functions over a direct power of fixed finite set.

Since the general problem is NP-hard, there is much research into identifying tractable cases of various versions of CSP.

In this talk, I will give a brief introduction into CSP, and then explain how submodularity and related more general conditions play a crucial role in classifying the complexity of optimization CSPs, and how this opens a new direction in submodularity research.

10. Jon Lee, U. Michigan

Title: Maximum Entropy Sampling

Abstract: I will survey exact mathematical-programming based algorithms for the NP-Hard problem of choosing a subset, of prespecified cardinality, from a finite set of random variables, so as to maximize the differential entropy (which is a submodular function).

11. Tom McCormick, UBC

Title: A Primal-Dual Algorithm for Weighted Abstract Cut Packing

Abstract: Hoffman and Schwartz developed the Lattice Polyhedron model and proved that it is totally dual integral (TDI), and so has integral optimal solutions. The model generalizes many important combinatorial optimization problems such as polymatroid intersection, cut covering polyhedra, min cost aborescences, etc., but has lacked a combinatorial algorithm. The problem can be seen as the blocking dual of Hoffman's Weighted Abstract Flow (WAF) model, or as an abstraction of ordinary Shortest Path and its cut packing dual, so we call it Weighted Abstract Cut Packing (WACP). We develop the first combinatorial algorithm for WACP, based on the Primal-Dual Algorithm framework. The framework is similar to that used by Martens and McCormick for WAF, in that both algorithms depend on a relaxation by a scalar parameter, and then need to solve an unweighted "restricted" subproblem. The subroutine to solve WACP's restricted subproblem is quite different from the corresponding WAF subroutine. The WACP subroutine uses an oracle to solve a restricted abstract cut packing/shortest path subproblem using greedy cut packing, breadth-first search, and an update that achieves complementary slackness. This plus a standard scaling technique yields a polynomial combinatorial algorithm.

(Joint with Britta Peis, TU Berlin)

12. Aranyak Mehta, Google Research

Title: Online Matching and Allocations

Abstract: An important problem in the area of search and display advertising is that of Ad allocation: matching ad slots to advertisers, online, under demand and supply constraints. The core algorithmic problem is that of online matching, and its generalization to various special cases of online submodular allocation. I will present an overview of some of the key algorithmic results in this area, and touch upon how these ideas can be applied in practice.

13. Vahab Mirrokni, Google Research

Title: Page-based ad allocation and online submodular welfare maximization

Abstract: As an important component of any ad serving system, online capacity (or budget) planning is a central problem in online ad allocation. Many variants of this problem are special cases of the online submodular welfare maximization problem. In this talk, I first discuss a page-based ad allocation problem and describe its relation with the submodular welfare maximization problem with online bidders. Then I will survey primal-based and dual-based techniques borrowed from online stochastic matching literature and report theoretical approximation guarantees and practical evaluations of these algorithms on real data sets. I will conclude the talk by several open problems about online stochastic variants of the welfare maximization problem.

14. Kazuo Murota, University of Tokyo

Title: (1) Introduction to Discrete Convex Analysis

Title: (2) Minimization and Maximization Algorithms in Discrete Convex Analysis

Abstract: Discrete convex analysis is a theory that aims at a discrete analogue of convex analysis for nonlinear discrete optimization. Technically it is a nonlinear generalization of matroid/submodular function theory; matroids are generalized to M-convex functions and submodular set functions to L-convex function.

This talk consists of two parts:

In the first part fundamental concepts and theorems in discrete convex analysis are explained with reference to familiar combinatorial optimization problems like minimum spanning tree, shortest path, and matching.

In view of the recent development in submodular function maximization, minimization and maximization algorithms in discrete convex analysis are explained in the second part of the talk. In particular, M-natural concave functions form a subclass of submodular functions that can be maximized in polynomial time.

15. Kiyohito Nagano, University of Tokyo

Title: Exact Approaches to Constrained Submodular Minimization Problems

Abstract: A number of combinatorial optimization problems in machine learning, including clustering problems, can be described as constrained submodular minimization problems. It is known that an unconstrained submodular function minimization problem can be solved in strongly polynomial time. However, additional constraints make the problem intractable in many settings.

In this talk, we discuss a submodular function minimization problem under a size constraint and a partitioning problem of a submodular system. Both of these problems are NP-hard. In our approaches, we do not give approximation algorithms. Instead, we propose new methods, each of which computes a portion of exact optimal solutions. Additionally, we evaluate the performance of the proposed algorithms through computational experiments.

16. Maurice Queyranne, UBC

Title: On Optimum k-way Partitions with Submodular Costs and Minimum Part-size Constraints.

Abstract: We investigate the computational complexity of finding an exact optimum for the following class of set partitioning problems: for a submodular set function f given by a value oracle, and small integers (i.e., constants) k and s1,…,sk, find a partition of the ground set into exactly k parts P1,..,Pk, such that each part Pi is of size (cardinality) at least si and with minimum total cost f(P1)+...+ f(Pk).

Such problems generalize k-way cut problems in graphs and hypergraphs, and arise in VLSI design (netlist partitioning), clustering and network analysis.

This talk will report on current joint work with Flavio Guiñez (Universidad de Chile).

17. Prasad Raghavendra, Georgia Tech

Title: Towards a Generalization of Submodular Minimization.

Abstract: In this talk, we present a possible avenue to generalize submodular minimization. Specifically, we present a notion of generalized submodularity (that is closely along the lines of the notion of multimorphisms). We will present an algorithm to minimize functions satisfying this notion and an additional condition of separability. A function is separable if it can be represented as a sum of several functions depending on constant number of inputs. For example, the cut function is submodular and separable, since it is equal to the sum of the edges cut, and each edge being cut depends only on two bits. In particular, a separable function is like an instance of CSP (constraint satisfaction problem). We show that any separable function and satisfying our generalized submodularity notion, can be minimized in polynomial time.

We believe that the condition on separability of the function is non-essential, and in particular, any function that satisfies the generalized notion could possibly be minimized efficiently.

18. Akiyoshi Shioura, Tohoku University

Title: Computing the Convex Closure of Discrete Convex Functions

Abstract: We consider computational aspect of the convex closure of discrete convex functions.

More precisely, given a discrete convex function defined on the integer lattice, we consider an algorithm for computing the function value and a subgradient of the convex closure at a given point.

Such an algorithm is required when the continuous relaxation approach is applied to nonlinear integer programs. In the theory of discrete convex analysis, two classes of discrete convex functions called M-convex and L-convex functions play primary roles; an M-convex function is a function defined on an integral polymatroid, while an L-convex function can be seen as an extension of a submodular set function.

It is known that both of M-convex and L-convex functions can be extended to ordinary convex functions in real variables.

While the convex closure of an L-convex function can be expressed by a simple formula, it is not clear how to compute the convex closure of an M-convex function.

In this talk, we show that the function values and subgradients of the convex closure of an M-convex function can be computed efficiently.

This result is shown by making full use of conjugacy results of M-convex functions.

19. Maxim Sviridenko, IBM

Title: New and Improved Bounds for the Minimum Set Cover Problem

Abstract: We study the relationship between the approximation factor for the Set-Cover problem and the parameters D: the maximum cardinality of any subset, and k: the maximum number of subsets containing any element of the ground set. We show an LP rounding based approximation of (k-1)(1-e^{-ln D/(k-1)}) +1, which is substantially better than the classical algorithms when k is approximately ln D, and also improves on related previous works. For the interesting case when k = \theta(log D) we also exhibit an integrality gap which essentially matches our approximation algorithm.

We also prove a hardness of approximation factor of \Omega (log D/ (log log D)^2) when k = theta(log D)$. This is the first study of the hardness factor specifically for this range of k and D, and improves on the only other such result implicitly proved before.

In the end of the talk we will discuss various submodular generalizations of the minimum set cover problem and results known for these settings.

20. Zoya Svitkina, Google

Title: Approximating Combinatorial Minimization Problems with Submodular Objective Functions.

I will discuss several generalizations of classical computer science problems obtained by replacing simpler objective functions with general submodular functions. These problems include submodular load balancing, which generalizes load balancing, or minimum-makespan scheduling, submodular sparsest cut and submodular balanced cut, which generalize their respective graph cut problems, as well as submodular function minimization with a cardinality lower bound. I will present upper and lower bounds for the approximability of these problems with a polynomial number of queries to a function-value oracle. The approximation guarantees that most of our algorithms achieve are of the order of sqrt(n/ln n), but, as the lower bounds show, this is the inherent difficulty of these problems. The talk is based on joint work with Lisa Fleischer that is published in FOCS 2008 and SICOMP 2011.

21. Vijay V. Vazirani, Georgia Tech

Title: Dispelling an Old Myth about an Ancient Algorithm

Abstract: Myth — and grapevine — has it that the Micali-Vazirani maximum matching algorithm is "too complicated".

The purpose of this talk is to show the stunningly beautiful structure of minimum length alternating paths and to convince the audience that our algorithm exploits this structure in such a simple and natural manner that the entire algorithm can fit into one's mind as a single thought.

For all practical purposes, the Micali-Vazirani algorithm, discovered in 1980, is still the most efficient known algorithm (for very dense graphs, more efficient algorithms have been given; however, they use fast matrix multiplication algorithms as subroutines). Yet, no more than a handful of people have understood this algorithm. By the end of this talk, we hope to quadruple the number of people in the latter class.

This talk is based on the following two papers:

1. http://www.cc.gatech.edu/~vazirani/MV.pdf
2. http://www.cc.gatech.edu/~vazirani/Theory-of-matching.pdf

22. Jan Vondrak, IBM

Title: (1) Optimization of Submodular Functions: Relaxations and Algorithms

Abstract: In this lecture, we cover the basic variants of optimization problems involving submodular functions, and known algorithms for solving them. Some particular problems that will be mentioned: the unconstrained minimization/ maximization of a submodular function, submodular optimization subject to packing/covering constraints, and welfare maximization in combinatorial auctions. Starting from greedy and local search algorithms, we will move on to continuous relaxations of these problems and their role in deriving efficient algorithms. Two important concepts that we will cover here are the Lovasz extension (for minimization), and the multilinear extension (for maximization) of submodular functions.

Title: (2) Optimization of Submodular Functions: Hardness and Optimality

Abstract: In this lecture, we describe techniques that allow us to prove the impossibility of achieving certain approximations for submodular optimization problems, and in some cases the optimality of known approximations. A key concept here is the notion of “symmetry gap” for submodular functions, which is also related to the integrality gap of a certain linear program. We will explain how this technique unifies some earlier results and leads to new proofs of optimality for several approximation algorithms. Finally, we will discuss some future challenges and questions in this area.

23. Laszlo Vegh, Georgia Tech

Title: Concave Generalized Flows with Applications to Market Equilibria

Abstract: We consider a nonlinear extension of the generalized network flow model, with the flow leaving an arc being an increasing concave function of the flow entering it, as proposed by Truemper and by Shigeno. We give the first polynomial time combinatorial algorithm for solving corresponding flow maximization problems, finding an epsilon-approximate solution in O(m(m+log n) log(MUm/epsilon)) arithmetic operations and value oracle queries, where M and U are upper bounds on simple parameters.

We show that this general convex programming model serves as a common framework for several market equilibrium problems, including the linear Fisher market model and its various extensions. Our result immediately extends these market models to more general settings. We also obtain a combinatorial algorithm for nonsymmetric Arrow-Debreu Nash bargaining, settling an open question by Vazirani.