A Simple Polynomial-time Rescaling Algorithm for Solving Linear Programs

John Dunagan
Microsoft Research
One Microsoft Way
Redmond, WA 98052
jdunagan@microsoft.com

Santosh Vempala*
Department of Mathematics
Massachusetts Institute of Technology
Cambridge, MA 02139
vempala@math.mit.edu

ABSTRACT

The perceptron algorithm, developed mainly in the machine learning literature, is a simple greedy method for finding a feasible solution to a linear program (alternatively, for learning a threshold function.). In spite of its exponential worst-case complexity, it is often quite useful, in part due to its noise-tolerance and also its overall simplicity. In this paper, we show that a randomized version of the perceptron algorithm with periodic rescaling runs in polynomial-time. The resulting algorithm for linear programming has an elementary description and analysis.

Categories and Subject Descriptors

F.2 [**Theory of Computation**]: Analysis of Algorithms

Keywords

Perceptron, Linear Programming, Polynomial Time

1. INTRODUCTION

Linear programming problems arise in many areas. The standard form is

$$\begin{array}{rcl} \max & c^T x & \\ Ax & \leq & b \\ x & \geq & 0 \end{array}$$

where b, c are in $\mathbb{R}^n$, and A is an $m \times n$ matrix of reals. It is often convenient to view the constraints as halfspaces in $\mathbb{R}^n$ and the problem is to find a point in their intersection of maximum objective value. The *dual* view is to think of the rows of A as points in R^n, and then the goal is to find a threshold function (i.e., a halfspace) that satisfies the given thresholds and maximizes the objective threshold.

*Supported by NSF CCR-6895000 and a Sloan foundation fellowship.

STOC'04, June 13–15, 2004, Chicago, Illinois, USA.

Polynomial-time algorithms for solving linear programs include the Ellipsoid method [13, 15, 22], interior point methods [14, 21] and the random walk method [2].

Another classical algorithm which is quite practical for problems arising in machine learning is the perceptron algorithm [1, 16, 18]. It was developed to solve the problem of learning a half-space. The algorithm is a simple greedy method that is guaranteed to converge to a feasible solution, if one exists. It could take an exponential number of iterations in the worst case. Nevertheless, it has many useful properties, including certain types of noise tolerance [4, 5]. It has also been related to boosting in the PAC model of learning [19, 20]. It was recently shown that the algorithm is polynomial with high probability for randomly perturbed linear programs [3]. It has been an open question as to whether there is a variant of the perceptron algorithm that is polynomial in the worst case. Besides suggesting a useful modification to the basic algorithm in practice, it would give a noise-tolerant (in a specific sense) polynomial-time algorithm.

To simplify matters, we focus on the problem of finding a feasible solution to a given set of linear constraints. Polynomial time reductions of the optimization problem to the feasibility version of the problem are well-known [13]. A typical approach is described in Section 4.

In this paper, we show that a randomized perceptron-like algorithm along with periodic rescaling applied to a feasible linear program will return a feasible point in polynomial time. As in all variants of the perceptron algorithm, it does not use any matrix inversions or barrier functions. Our main Theorem (Theorem 3.5) is that a strictly feasible linear program with m constraints in $\mathbb{R}^n$ is solved in time $O(mn^4 \log n \log(1/\rho))$ where ρ is a parameter that roughly corresponds to the radius of a ball that fits in the feasible region, and $\log(1/\rho)$ is guaranteed to be polynomial in the input description. It should be noted that the complexity of our algorithm is not as good as the current best algorithms [10, 21]. On the other hand, it is rather simple and inherits the noise-tolerance properties of the perceptron algorithm. As a direct consequence, we get a simpler (and faster) solution to the problem of learning noisy linear threshold functions [4, 6].

Here is the main idea. The perceptron algorithm maintains a halfspace that it updates using a constraint that is currently violated. The convergence of the algorithm depends on a separation parameter quantifying the amount of "wiggle room" available for a feasible halfspace. This is in the dual view described above. In the corresponding primal view, when we think of constraints as halfspaces, the wiggle room is the radius of the largest ball that fits in the feasible region. Roughly speaking, the analysis of the perceptron algorithm says that if this ball is large, then the algorithm will converge quickly. The work of [4] shows that even when this ball is small, a simple variant of the perceptron algorithm finds a *nearly* feasible solution, i.e., constraints might be violated but each is not violated by much. Here we show that when the wiggle room is small, such a nearly feasible solution can be used to apply a linear transformation (in fact, a rank 1 update) that expands the wiggle room (enlarges the ball contained in the feasible region) by essentially a $1+1/n$ factor. Thus, in $O(n)$ iterations, we either find a feasible point or double the wiggle room.

2. THE ALGORITHM

In this section we present an algorithm for the linear feasibility problem

$$Ax \geq 0, x \neq 0$$

consisting of m constraints in n dimensions, i.e. $x \in \mathbb{R}^n$ and A is $m \times n$. In Section 4 we show how well-known methods can be used to reduce the standard form to this *homogenized* form.

The algorithm is iterative and each iteration consists of three phases, a *perceptron* phase, a *perceptron improvement* phase, and a *rescaling* phase. The perceptron phase uses the classical perceptron algorithm. The perceptron improvement phase uses a modified version of the basic perceptron algorithm. This modified version was described earlier in [4]. Figure 1 depicts the rescaling phase.

The classical perceptron algorithm for the linear feasibility problem is to find a violated constraint, move the trial solution x one unit in the direction normal to the violated constraint, and repeat if necessary (this is step 2 in the algorithm).

In the rest of the paper, we let $\bar{x}$ denote the unit vector in the direction of x.

Algorithm.

Input: An $m \times n$ matrix A.
Output: A point x such that $Ax \geq 0$ and $x \neq 0$.

1. Let $B = I, \sigma = 1/(32n)$.
2. (Perceptron)
   (a) Let x be the origin in $\mathbb{R}^n$.
   (b) Repeat at most $1/\sigma^2$ times:
   If there exists a row a such that $a \cdot x \leq 0$, set $x = x + \bar{a}$.
3. If $Ax \geq 0$, then output Bx as a feasible solution and stop.
4. (Perceptron Improvement)
   (a) Let x be a random unit vector in $\mathbb{R}^n$.
   (b) Repeat at most $\ln n/\sigma^2$ times:
   If there exists a row a such that $\bar{a} \cdot \bar{x} < -\sigma$, set $x = x - (\bar{a} \cdot x)\bar{a}$.
   If $x = 0$, go back to step (a).
   (c) If there is still a row a such that $\bar{a} \cdot \bar{x} < -\sigma$, restart at step (a).
5. If $Ax \geq 0$, then output Bx as a feasible solution and stop.
6. (Rescaling)
   Set $A = A\left(I + \bar{x}\bar{x}^T\right)$ and $B = B\left(I + \bar{x}\bar{x}^T\right)$.
7. Go back to step 2.

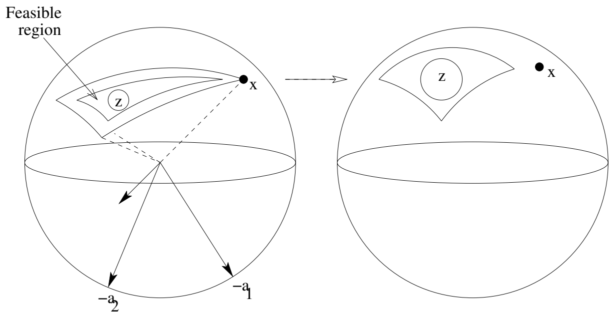

Figure 1: A constraint system before and after rescaling.

3. ANALYSIS

3.1 Ideas

Let A_0 denote the initial A input to the algorithm, and let A_* denote the A that the algorithm has when it terminates. When the algorithm terminates, it produces a non-zero vector $x' = Bx$ such that $A_* x = (A_0 B)x \geq 0$, i.e., $A_0(Bx) \geq 0$, as desired.

During each outer iteration (steps 2-7), the perceptron improvement phase (step 4) may be re-run a number of times, but it terminates quickly with high probability (Lemma 3.2). Thus the main question is, how many iterations does the algorithm make? To answer this, we define the following quantity which measures the "roundness" of the feasible region:

$$\rho(A) = \max_{x:|x|=1, Ax\geq 0} \min_i (\bar{a}_i \cdot x)$$

were a_i denotes the i'th row of A and $\bar{a}_i$ is the unit vector along a_i.

We call ρ the *radius* of A, and z, the unit vector that achieves the maximum, its *center*. Note that ρ is just the radius of the largest ball that fits in the feasible cone, such that the center of the ball is on the unit sphere. To avoid confusion, let $\rho_0 = \rho(A_0)$ denote the roundness of the input. See Section 4 for a further discussion of ρ.

The classical analysis of the perceptron algorithm [16] (repeated in lemma 3.1) shows that the classical perceptron algorithm (our perceptron phase) applied to a linear feasibility problem with radius ρ will terminate in at most $1/\rho^2$ iterations. Further, it is not hard to see that if ρ is initially more than 0, there exists a linear transformation that takes ρ arbitrarily close to 1. This transformation is defined by $A'_\nu = A(I + \nu zz^T)$. As $\nu \to \infty$, a simple calculation shows that $\rho(A)$ goes to 1. So if we knew the transformation, we could just apply it once and then run the standard perceptron algorithm!

But this is equivalent to the original problem, since z is a feasible point. Instead, in our algorithm, we incrementally transform A using near-feasible solutions, so that ρ increases steadily. Our main lemma shows that in any iteration of our algorithm where ρ is small, it will increase in expectation by a multiplicative factor. Combining our main lemma with the classical analysis will yield that if ρ is small, it gets bigger (guaranteed by the perceptron improvement and scaling phases), and if ρ is big, the algorithm finds a feasible point (guaranteed by the perceptron phase).

Each iteration of the algorithm consists a perceptron phase and a perceptron improvement phase. Alternatively, one could do just the perceptron improvement phase for a predetermined number of steps, and then check for completion by running the perceptron phase once.

3.2 Proofs

We begin with the well-known analysis of the standard perceptron algorithm.

LEMMA 3.1 (BLOCK-NOVIKOFF [16]). *The classical perceptron algorithm (our perceptron phase) returns a feasible point in at most $1/\rho^2$ iterations.*

PROOF. Consider the potential function $x \cdot z / \|x\|$. The numerator increases by at least ρ on each step:

$$(x + \bar{a}_i) \cdot z = x \cdot z + \bar{a}_i \cdot z \geq x \cdot z + \rho$$

While the square of the denominator increases by at most 1:

$$(x + \bar{a}_i) \cdot (x + \bar{a}_i) = x \cdot x + 2x \cdot \bar{a}_i + \bar{a}_i \cdot \bar{a}_i \leq x \cdot x + 1$$

since $x \cdot \bar{a}_i \leq 0$. After t iterations, the potential function is at least $\frac{t\rho}{\sqrt{t}}$ and thus the classical perceptron algorithm must terminate before $1/\rho^2$ iterations. If the algorithm terminates, it must have found a feasible point. □

Next, we recall the analysis of the modified perceptron algorithm (our perceptron improvement phase).

LEMMA 3.2 (BFKV [4]). *Let A be the constraint matrix at the beginning of a perceptron improvement phase and let z be any unit vector such that $Az \geq 0$. With probability at least $\frac{1}{8}$, in at most $\ln n/\sigma^2$ steps, the perceptron improvement phase returns a vector x such that*

(a) $\bar{a} \cdot x \geq -\sigma$ *for every row a of A and*

(b) $z \cdot \bar{x} \geq \dfrac{1}{\sqrt{n}}$.

PROOF. The proof of both parts is similar. A standard computation shows that for a random unit vector x, $z \cdot x \geq 1/\sqrt{n}$ with probability at least $1/8$. We now show that if this is the case, we terminate in the desired number of iterations.

Note that in each update step, $z \cdot x$ does not decrease

$$(x - (x \cdot \bar{a}_i)\bar{a}_i) \cdot z = x \cdot z - (x \cdot \bar{a}_i)(\bar{a}_i \cdot z) \geq x \cdot z$$

because $x \cdot \bar{a}_i$ had to be negative in order for $\bar{a}_i$ to be used in an update step, and $\bar{a}_i \cdot z \geq 0$ by assumption. (This also implies that if $z \cdot x \geq 1/\sqrt{n}$ initially, x will never be set to zero.) On the other hand $x \cdot x$ does decrease significantly because

$$\begin{aligned}(x - (x \cdot \bar{a}_i)\bar{a}_i) \cdot (x - (x \cdot \bar{a}_i)\bar{a}_i) &= x \cdot x - 2(\bar{a}_i \cdot x)^2 + (\bar{a}_i \cdot x)^2 \\ &= x \cdot x - (\bar{a}_i \cdot x)^2 \\ &\leq x \cdot x(1 - \sigma^2).\end{aligned}$$

Thus after t iterations $\|x\| \leq (1-\sigma^2)^{t/2}$. If $t > (\ln n)/\sigma^2$, we would have $\frac{x \cdot z}{\|x\|} > 1$, which cannot happen. Therefore, every time we start through the algorithm, we finish with probability at least $1/8$ and we return x for which

$$\frac{x \cdot z}{\|x\|} \geq \frac{1}{\sqrt{n}}.$$

□

We are now ready to prove our main lemma about progress in any iteration that does not find a feasible solution.

LEMMA 3.3. *Suppose that $\rho, \sigma \leq 1/32n$. Let A' be obtained from A by one iteration of the algorithm (one on which the problem was not solved). Let ρ' and ρ be the radii of A' and A respectively. Then,*

(a) $\rho' \geq (1 - \frac{1}{32n} - \frac{1}{512n^2})\rho$.

(b) With probability at least $\frac{1}{8}$, $\rho' \geq (1 + \frac{1}{3n})\rho$.

PROOF. The analysis will use Lemma 3.2. Part (a) of the lemma says that with probability $\frac{1}{8}$, the vector x at the end of step 4(b) satisfies $\bar{a} \cdot \bar{x} \geq -\sigma$. Part (b) of the lemma asserts that the vector x at the end of step 4(b) will simultaneously satisfy $z \cdot \bar{x} \geq 1/\sqrt{n}$ with probability at least 1/8.

Let a_i, $i = 1, \ldots m$ be the rows of A at the beginning of some iteration. Let z be a unit vector satisfying $\rho = \min_i \bar{a}_i \cdot z$, and let $\sigma_i = \bar{a}_i \cdot \bar{x}$. After a perceptron improvement phase, we get a vector x such that

$$\bar{a}_i \cdot \bar{x} = \sigma_i \geq -\sigma.$$

As in the theorem statement, let A' be the matrix obtained after the rescaling step, i.e.

$$a'_i = a_i + (a_i \cdot \bar{x})\bar{x}.$$

Finally, define

$$z' = z + \alpha(z \cdot \bar{x})\bar{x}.$$

where α will be specified shortly. Although z' is not necessarily the center of A', ρ' is a maximum over a set, and so considering one element $(\bar{z}')$ of the set suffices to lower bound ρ'. We have

$$\rho' \geq \min_i \bar{a}'_i \cdot \bar{z}' = \min_i \frac{\bar{a}'_i \cdot z'}{|z'|}.$$

We will first prove that $\bar{a}'_i \cdot z'$ cannot be too small.

$$\begin{aligned}\bar{a}'_i \cdot z' &= \left(\frac{\bar{a}_i + (\bar{a}_i \cdot \bar{x})\bar{x}}{\|\bar{a}_i + (\bar{a}_i \cdot \bar{x})\bar{x}\|}\right) \cdot z' \\ &= \frac{[\bar{a}_i + (\bar{a}_i \cdot \bar{x})\bar{x}][z + \alpha(z \cdot \bar{x})\bar{x}]}{\sqrt{1 + 3(\bar{a}_i \cdot \bar{x})^2}} \\ &\geq \frac{\rho + \sigma_i(z \cdot \bar{x})(1 + 2\alpha)}{\sqrt{1 + 3\sigma_i^2}}\end{aligned}$$

We choose :

$$\alpha = \frac{1}{2}(\frac{\rho}{\bar{x} \cdot z} - 1)$$

so that $2\alpha + 1 = \rho/(\bar{x} \cdot z)$. We have not ensured that $\bar{x} \cdot z \neq 0$, but substituting in for $\alpha(\bar{x} \cdot z)$ in the definition of z' using the value we have chosen for α would remove this boundary case. We have chosen not to do this to aid the exposition. We proceed to calculate

$$\bar{a}'_i \cdot z' \geq \rho \frac{1 + \sigma_i}{\sqrt{1 + 3\sigma_i^2}} \geq \rho \frac{1 - \sigma}{\sqrt{1 + 3\sigma^2}}. \tag{1}$$

where the second inequality follows from $\sigma_i \in [-\sigma, 1]$. Next, observe that

$$|z'|^2 = 1 + (\alpha^2 + 2\alpha)(\bar{x} \cdot z)^2 = 1 + \frac{\rho^2}{4} + (z \cdot \bar{x})\left(\frac{\rho}{2} - \frac{3}{4}(z \cdot \bar{x})\right).$$

We consider two cases:

1. $|z \cdot \bar{x}| < \frac{1}{\sqrt{n}}$. This happens with probability at most $7/8$.

   Viewing $|z'|^2$ as a quadratic polynomial in $(z' \cdot x)$, we see that it is maximized when $(z' \cdot x) = \frac{\rho}{3}$. In this case, we have

   $$|z'|^2 \leq 1 + \frac{\rho^2}{4} + \frac{\rho^2}{12} \leq 1 + \frac{\rho^2}{3}.$$

   Using the elementary inequality $\frac{1}{\sqrt{1+\beta}} \geq 1 - \frac{\beta}{2}$ for $\beta \in (-1, 1)$, we find

   $$\begin{aligned}\rho' &\geq \rho \frac{1 - \sigma}{\sqrt{1 + 3\sigma^2}|z'|} \\ &\geq \rho(1 - \sigma)(1 - \frac{3\sigma^2}{2})(1 - \frac{\rho^2}{6}) \\ &\geq \rho(1 - \sigma - \frac{3\sigma^2}{2} - \frac{\rho^2}{6}) \\ &\geq \rho(1 - \frac{1}{32n} - \frac{1}{512n^2})\end{aligned}$$

   since $\sigma, \rho \leq 1/32n$.

2. $|z \cdot \bar{x}| \geq \frac{1}{\sqrt{n}}$. This happens with probability at least $1/8$.

   In this case,

   $$|z'|^2 = 1 + \frac{\rho^2}{4} + (z \cdot \bar{x})\frac{\rho}{2} - \frac{3}{4}(z \cdot \bar{x})^2 \leq 1 - \frac{3}{4n} + \frac{\rho}{2\sqrt{n}} + \frac{\rho^2}{4}.$$

   Using the elementary inequality again, we find

   $$\begin{aligned}\rho' &\geq \rho(1 - \sigma)(1 - \frac{3\sigma^2}{2})(1 + \frac{3}{8n} - \frac{\rho}{4\sqrt{n}} - \frac{\rho^2}{8}) \\ &\geq \rho\left(1 - \sigma - \frac{3\sigma^2}{2} + \frac{3}{8n} - \frac{\rho}{4\sqrt{n}} - \frac{\rho^2}{8}\right) \\ &\geq \rho\left(1 + \frac{1}{3n}\right).\end{aligned}$$

This proves both parts of the Lemma. ☐

We can now bound the number of iterations. Recall that $\rho_0 = \rho(A_0)$ is the ρ value for the initial input matrix.

THEOREM 3.4. *With high probability, the algorithm finds a feasible solution in $O(n \log(1/\rho_0))$ iterations.*

PROOF. It suffices to show that ρ will be at least $1/32n$ in $O(n \log(1/\rho_0))$ iterations. Let X_i be a random variable for the i'th iteration, with value 1 if ρ grows by a factor of $(1 + 1/3n)$ or more and value 0 otherwise. Let X be the sum of the X_i's over T iterations. Then by Lemma 3.3(b)

$$\mathbf{E}[X] \geq \frac{T}{8}.$$

The Chernoff bound gives,

$$\Pr(X < (1 - \epsilon)\mathbf{E}[X]) \leq e^{-\epsilon^2 E[X]/2}.$$

Let ρ_T be the ρ value after T iterations. Let $T = 2048n \log(1/\rho_0)$ and $\epsilon = 1/16$. Then

$$e^{-\epsilon^2 T/8} \leq e^{-n}$$

which satisfies the definition of high probability. Analyzing ρ_T in the case that X is within ϵ of its expectation, we have

$$\begin{aligned}\rho_T &\geq \rho\left(1 + \frac{1}{3n}\right)^X \left(1 - \frac{1}{32n} - \frac{1}{512n^2}\right)^{T-X} \\ &\geq \rho\left(1 + \frac{1}{3n}\right)^{\frac{T}{8} - \epsilon\frac{T}{8}} \left(1 - \frac{1}{32n} - \frac{1}{512n^2}\right)^{\frac{7T}{8} + \epsilon\frac{T}{8}} \\ &\geq \rho\left(1 + \frac{1}{3n}\right)^{\frac{15T}{128}} \left(1 - \frac{1}{32n} - \frac{1}{512n^2}\right)^{\frac{113T}{128}} \\ &\geq \rho e^{T/1000n}.\end{aligned}$$

We summarize: with probability at least $1 - e^{-n}$, in at most T iterations, ρ grows to at least $1/(32n)$, at which point the perceptron phase succeeds in finding a feasible point. ☐

Finally, the time complexity is a matter of accounting.

THEOREM 3.5. *With high probability, the algorithm finds a feasible solution in time $O(mn^4 \log n \log(1/\rho_0))$.*

PROOF. The inner loop of either the perceptron phase or the perceptron improvement phase requires at most one matrix-vector multiplication (time $O(mn)$), and a constant number of vector manipulations (time $O(n)$). The number of times we repeat the inner loop is at most $16n^2$ in the perceptron phase, and at most $\log n/\sigma^2 = O(n^2 \log n)$ in the perceptron improvement phase. The bound on the total number of times through the perceptron improvement phase is probabilistic, but it can be shown to be close to its expectation using a Chernoff bound just as in the previous theorem. The scaling phase takes time $O(n^2)$. Calculating Bx similarly takes time $O(n^2)$.

By the previous theorem, the number of iterations is at most $O(n \log(1/\rho_0))$. This yields the overall time bound $O(mn^4 \log n \log(1/\rho_0))$. ☐

4. THE STANDARD FORM

In this section we discuss how to reduce the standard linear programming problem to the one solved in this paper. A typical approach to reduce optimization to feasiblity is to replace max $c^T x$ by the constraint $c^T x \geq c_0$ for some value of c_0. Binary search can be used to determine a nearly optimal value of c_0. A solution that is feasible for the problem with c_0 sufficiently close to optimal can be *rounded* to an optimal vertex solution through basis reduction [13].

Next, we show how to reduce the standard form linear feasibility problem

$$Ax \leq b, x \geq 0$$

into the linear feasibility problem we studied earlier. This technique is typically referred to as *homogenization*. We also relate ρ, the condition number, and the "bit-length" of the problem before and after homogenization. We conclude from this that the algorithm is polynomial in the traditional sense [13].

Introduce the variable x_0 and consider the problem

$$Ax \leq bx_0, x \geq 0, x_0 > 0$$

To convert a solution for the standard form to one for the homogenized form, set $x_0 = 1$. To convert a solution from the homogenized form to a solution for the standard form, divide x by x_0. To rewrite the homogenized form as just

$$A'x' \geq 0, x' \neq 0$$

Let $x' = [x\ x_0]$ and

$$A' = \begin{bmatrix} -A & b \\ I & \end{bmatrix}$$

A valid solution to $A'x' \geq 0$ might have $x_0 = 0$. However, because the classic perceptron algorithm (our perceptron phase) always produces solutions in the strict interior of the feasible region, our algorithm will always return a solution with $x_0 > 0$.

The traditional measure of the difficulty of a linear programming problem in the Turing model of computation is called the "bit-length" and denoted by L. The quantity L is essentially the input length of the linear program. The condition number of a linear program was defined by Renegar as the normalized distance to ill-posedness [17]. By applying a small perturbation one can ensure that the log of the condition number is upper-bounded by a polynomial in L, and so linear programming algorithms that are polynomial in the log of the condition number are also polynomial in L. The quantity ρ was related to the condition number in [11, 7] (it was called the minimum width by Freund) and more recently in [9]. In particular, ρ of the homogenized program is no more than n times the condition number of a slight perturbation of the original program. We conclude from this that the algorithm is polynomial in L.

5. ACKNOWLEDGEMENTS

We thank Dan Stefankovic and Adam Smith for helping these ideas through their formative stages, and Adam Klivans and Rocco Servedio for bringing [12] to their attention. We are grateful to Rob Freund and Mario Szegedy for many useful comments.

6. REFERENCES

[1] S. Agmon, The relaxation method for linear inequalities, *Canadian J. of Math.*, 6(3), 382–392, 1954.

[2] D. Bertsimas and S. Vempala, Solving convex programs by random walks, *Proceedings of the Annual ACM Symposium on the Theory of Computing*, 109–115, 2002.

[3] A. Blum and J. Dunagan, Smoothed Analysis of the Perceptron Algorithm for Linear Programming, *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms*, 905–914, 2002.

[4] A. Blum, A. Frieze, R. Kannan, and S. Vempala, A polynomial-time algorithm for learning noisy linear threshold functions, *Algorithmica*, 22(1), 35–52, 1998.

[5] T. Bylander, Learning linear threshold functions in the presence of classification noise, *Proceedings of the Workshop on Computational Learning Theory*, 1994.

[6] E. Cohen, Learning noisy perceptrons by a perceptron in polynomial time, *Proceedings of the Annual IEEE Symposium on the Foundations of Computer Science*, 514–523, 1997.

[7] F. Cucker and D. Cheung, A new condition number for linear programming, *Mathematical Programming, Series A*, 91(1), 163–174, 2001.

[8] V. Chvatal, *Linear Programming*, W.H. Freeman, 1983.

[9] J. Dunagan, S. Teng, and D. A. Spielman, Smoothed Analysis of Renegar's Condition Number for Linear Programming, *SIAM Conference on Optimization*, 2002.

[10] R. M. Freund and S. Mizuno: "Interior Point Methods: Current Status and Future Directions," in High Performance Optimization, H. Frenk et al. (eds.), Kluwer Academic Publishers, pp. 441-466, 2000.

[11] R. M. Freund and J. R. Vera, Some characterizations and properties of the "distance to ill-posedness" and the condition measure of a conic linear system, *Math. Programming*, 86, 225–260, 1999.

[12] J. Forster, A Linear Lower Bound on the Unbounded Error Probabilistic Communication Complexity, *Sixteenth Annual IEEE Conference on Computational Complexity 2001*, http://citeseer.nj.nec.com/forster01linear.html

[13] L. Grotchel, L. Lovasz, and A. Schrijver, *Geometric algorithms and Combinatorial Optimization*, Springer-Verlag, Berlin, 1988.

[14] N. Karmarkar, A new polynomial-time algorithm for linear programming, *Combinatorica*, 4, 373–396, 1984.

[15] L. G. Khachiyan, A polynomial algorithm in linear programming, (in Russian), *Doklady Akedamii Nauk SSSR*, 244, 1093–1096, 1979 (English translation: *Soviet Mathematics Doklady*, 20, 191–194, 1979).

[16] M. L. Minsky and S. A. Papert, *Perceptrons: An introduction to computational geometry*, 1969.

[17] J. Renegar, Incorporating condition measures into the complexity theory of linear programming, *SIAM Journal on Optimization*, 5(3), 506–524, 1995.

[18] F. Rosenblatt, *Principles of Neurodynamics*, Spartan Books, 1962.

[19] R. Servedio. On PAC Learning using Perceptron, Winnow and a Perceptron-Like Algorithm. *SIAM Journal on Computing*, 31(5), 1358-1369, 2002.

[20] R. Servedio, Smooth Boosting and Learning with Malicious Noise. Fourteenth Annual Conference on Computational Learning Theory, 473-489, 2001.

[21] P. M. Vaidya, A new algorithm for minimizing convex functions over convex sets, *Mathematical Programming*, 291–341, 1996.

[22] D. B. Yudin and A. S. Nemirovski, Evaluation of the information complexity of mathematical programming problems, (in Russian), Ekonomika i Matematicheskie Metody 12, 128-142, 1976 (English translation: Matekon 13, 2, 3-45, 1976).