# Near-Optimal Active Learning of Halfspaces via Query Synthesis in the Noisy Setting 

Lin Chen, ${ }^{1,2}$ Hamed Hassani, ${ }^{3}$ Amin Karbasi ${ }^{1,2}$<br>${ }^{1}$ Department of Electrical Engineering, ${ }^{2}$ Yale Institute for Network Science, Yale University<br>${ }^{3}$ Computer Science Department, ETH Zürich<br>\{lin.chen, amin.karbasi\}@yale.edu, hamed@inf.ethz.ch


#### Abstract

In this paper, we consider the problem of actively learning a linear classifier through query synthesis where the learner can construct artificial queries in order to estimate the true decision boundaries. This problem has recently gained a lot of interest in automated science and adversarial reverse engineering for which only heuristic algorithms are known. In such applications, queries can be constructed de novo to elicit information (e.g., automated science) or to evade detection with minimal cost (e.g., adversarial reverse engineering). We develop a general framework, called dimension coupling (DC), that 1) reduces a $d$-dimensional learning problem to $d-1$ lowdimensional sub-problems, 2) solves each sub-problem efficiently, 3) appropriately aggregates the results and outputs a linear classifier, and 4) provides a theoretical guarantee for all possible schemes of aggregation. The proposed method is proved resilient to noise. We show that the DC framework avoids the curse of dimensionality: its computational complexity scales linearly with the dimension. Moreover, we show that the query complexity of DC is near optimal (within a constant factor of the optimum algorithm). To further support our theoretical analysis, we compare the performance of DC with the existing work. We observe that DC consistently outperforms the prior arts in terms of query complexity while often running orders of magnitude faster.


## 1 Introduction

In contrast to the passive model of supervised learning, where all the labels are provided without any interactions with the learning mechanism, the key insight in active learning is that the learning algorithm can perform significantly better if it is allowed to choose which data points to label. This approach has found far-reaching applications, including the classical problems in AI (e.g., classification (Tong and Koller 2002), information retrieval (Tong and Chang 2001), speech recognition (Hakkani-Tur, Riccardi, and Gorin 2002)) as well as the modern ones (e.g., interactive recommender systems (Karbasi, Ioannidis, and Massoulie 2012), optimal decision making (Javdani et al. 2014)), and optimal information gathering (Chen et al. 2015). In all the above applications, the unlabeled data are usually abundant and easy to obtain, but training labels are either time-

[^0]consuming or expensive to acquire (as they require asking an expert).

Throughout this paper, our objective is to actively learn an unknown halfspace $H^{*}=\left\{x \in \mathbb{R}^{d}:\left\langle h^{*}, x\right\rangle>0\right\}$ via query synthesis (a.k.a. membership queries), where $\langle\cdot, \cdot\rangle$ denotes the standard inner product of the Euclidean space and $h^{*}$ is the unit normal vector of the halfspace we want to learn. We would like to note that learning the halfspace $H^{*}$ is mathematically equivalent to learning its unit normal vector $h^{*}$; therefore we focus on learning $h^{*}$ hereinafter. In addition, it should be noted that using the kernel trick we can easily extend the halfspace learning to more complex (e.g., nonlinear) decision boundaries (Shawe-Taylor and Cristianini 2004).

The hypothesis space $\mathcal{H}$, which consists of all possibilities of unit normal vectors, is the unit sphere $S^{d-1}=\left\{x \in \mathbb{R}^{d}\right.$ : $\|x\|=1\}$, where $\|\cdot\|$ denotes the standard Euclidean norm.

In active learning of halfspaces via query synthesis, the algorithm is allowed to query whether any point $x$ in $\mathbb{R}^{d}$ resides in the true halfspace. When the algorithm queries $x$, the true outcome is $\boldsymbol{\operatorname { s i g n }}\left(\left\langle h^{*}, x\right\rangle\right) \in\{1,-1\}$. When $\boldsymbol{\operatorname { s i g n }}\left(\left\langle h^{*}, x\right\rangle\right)=1$, it means that $x \in H^{*}$; otherwise, $x \notin H^{*}$. We should note here that the only information we obtain from a query is the sign of the inner product rather than the value. For example, the queries of the form $\operatorname{sign}\left(\left\langle h^{*}, e_{i}\right\rangle\right)$, where $e_{i}$ is the $i$ th standard basis vector, will only reveal the sign of the $i$ th component of $h^{*}$ (and nothing further about its value).

In the noiseless setting, we observe the true outcome of the query, i.e. $\operatorname{sign}\left\langle h^{*}, x\right\rangle \in\{1,-1\}$. In the noisy setting, the outcome is a flipped version of the true sign with independent flip probability $\rho$. That is, denoting the outcome by $y$ we have $y \in\{-1,1\}$ and $\operatorname{Pr}\left[y \neq \boldsymbol{\operatorname { s i g n }}\left\langle h^{*}, x\right\rangle\right] \triangleq \rho<1 / 2$.

Since the length of the selected vector $x$ will not affect the outcome of the query, we only query the points on the unit sphere $S^{d-1}=\left\{x \in \mathbb{R}^{d}:\|x\|=1\right\}$. Hence, we term $\mathcal{X}=S^{d-1}$ as the query space.

Given $\epsilon, \delta>0$, we would like to seek an active learning algorithm that (i) adaptively selects vectors $x_{1}, x_{2}, \ldots \in \mathcal{X}$, (ii) observes the (noisy) responses to each query $\operatorname{sign}\left\langle h^{*}, x_{i}\right\rangle$, (iii) and outputs, using as few queries as possible, an estimate $\hat{h}$ of $h^{*}$ such that $\left\|\hat{h}-h^{*}\right\|<\epsilon$ with probability at least $1-\delta$.

Our main contribution in this paper is to develop a noise
resilient active learning algorithm that has access to noisy membership queries. To the best of our knowledge, we are the first to show a near-optimal algorithm that outperforms in theory and practice the naive repetition mechanism and the recent spectral heuristic methods (Alabdulmohsin, Gao, and Zhang 2015). Specifically, we develop a framework, called Dimension Coupling (DC), with the following guarantees. Its query complexity is $\tilde{O}\left(d\left(\log \frac{1}{\epsilon}+\log \frac{1}{\delta}\right)\right)^{1}$ and its computational complexity is $\tilde{O}\left(d\left(\log \frac{1}{\epsilon}+\log \frac{1}{\delta}\right)^{2}\right)$. In particular, in the noiseless setting $(\rho=0)$, both its computational complexity and query complexity are $\tilde{O}\left(d \log \frac{1}{\epsilon}\right)$. Note that in both settings the computational complexity scales linearly with the dimension. Moreover, the query complexity in both settings is near-optimal. Our empirical experiments demonstrate that DC runs orders of magnitude faster than the existing methods.

The rest of the paper is structured as follows. In Section 2, we start with investigating this problem in the 2dimensional case and present an algorithm called $\mathrm{DC}^{2}$. Then in Section 3 we generalize it to the $d$-dimensional case and present a general framework called DC. Empirical results are shown in Section 4. We extensively review related literature in Section 5. All proofs are provided in the extended version (Chen, Hassani, and Karbasi 2016).

## 2 DC ${ }^{2}$ : Solving the 2-Dimensional Problem

To gain more intuition before studying the general $d$ dimensional problem, it might be beneficial to study a special case where the dimension is two. In other words, we study in this section how to learn the normalized projection of the true unit normal vector $h^{*} \in \mathbb{R}^{d}$ onto $\operatorname{span}\left\{e_{1}, e_{2}\right\}$, where $e_{1}, e_{2} \in \mathbb{R}^{d}$ are two orthonormal vectors and $\operatorname{span}\left\{e_{1}, e_{2}\right\}$ is the linear subspace spanned by $e_{1}$ and $e_{2}$. We should note here that the underlying space is still $d$-dimensional (i.e., $\mathbb{R}^{d}$ ) but our goal is not to learn $h^{*}$ per se but its normalized projection onto a 2 -dimensional subspace.

Formally, given two orthonormal vectors $e_{1}, e_{2}$ we denote the (normalized) projection of $h^{*}$ onto $\operatorname{span}\left\{e_{1}, e_{2}\right\}$ by $h^{\perp}$, i.e.,

$$
\begin{equation*}
h^{\perp}=\frac{\left\langle h^{*}, e_{1}\right\rangle e_{1}+\left\langle h^{*}, e_{2}\right\rangle e_{2}}{\left\|\left\langle h^{*}, e_{1}\right\rangle e_{1}+\left\langle h^{*}, e_{2}\right\rangle e_{2}\right\|_{2}} \tag{1}
\end{equation*}
$$

Our objective is to find a unit vector $\hat{e} \in \operatorname{span}\left\{e_{1}, e_{2}\right\}$ such that $\left\|\hat{e}-h^{\perp}\right\|<\epsilon$. In fact, we require the latter to hold with probability at least $1-\delta$.

We should emphasize that noise, characterized by independent flip probability $\rho$, is generally present. In the $2-$ dimensional problem, one may propose to use the simple binary search (a detailed discussion with examples is presented in the extended version) to find a unit vector $\hat{e}$ that resides $\epsilon$-close to $h^{\perp}$. To make it noise-tolerant, when the binary search algorithm queries a point, say $x_{i}$, we query it $R$ times to obtain $R$ noisy versions of $\operatorname{sign}\left\langle h^{*}, x_{i}\right\rangle$ and view the majority vote of the noisy versions as the true outcome (Kääriäinen 2006; Karp and Kleinberg 2007; Nowak 2011). We call this method repetitive querying. However, its

[^1]```
Algorithm 1 DC \({ }^{2}\)
Input: orthonormal vectors \(e_{1}, e_{2}\), estimation error at most
    \(\epsilon\), success probability at least \(1-\delta\).
Output: a unit vector \(\hat{e}\) which is an estimate for the normal-
    ized orthogonal projection of \(h^{*}\) onto \(\operatorname{span}\left\{e_{1}, e_{2}\right\}\).
    Set \(p_{0}(h)\) to be uniform, i.e., \(\forall h \in S^{1}: p_{0}(h)=1 / 2 \pi\).
    for \(m=1\) to \(T_{\epsilon, \delta}\) do
        Find a vector \(x_{m} \in S^{1}\) which is a solution to the fol-
        lowing equation: \(\int_{S^{1}} \operatorname{sign}\langle x, h\rangle p_{m-1}(h) d h=0\). If
        there are multiple solutions, choose one arbitrarily.
        Ask from the oracle the value of \(\operatorname{sign}\left\langle x_{m}, h^{*}\right\rangle\).
        Based on the (noisy) response obtained from the ora-
        cle, update the distribution \(p_{m-1}(h)\) to \(p_{m}(h)\).
    end for
    return \(\hat{e}=\arg \max _{h \in S^{1}} p_{T_{\epsilon, \delta}}(h)\).
```

query complexity is $O(\log (1 / \epsilon)(\log \log (1 / \epsilon)+\log (1 / \delta))$, which is suboptimal both theoretically (proof in the extended version) and empirically (referred to as RepetitiveDC in Section 4).

As a result, instead, we will present a Bayesian algorithm termed $\mathrm{DC}^{2}$ that solves this 2 -dimensional problem with query complexity $O(\log (1 / \epsilon)+\log (1 / \delta))$. Recall that any unit vector inside $\operatorname{span}\left\{e_{1}, e_{2}\right\}$, e.g., $h^{\perp}$, can equivalently be represented as a pair $\left(c_{1}, c_{2}\right)$ on the two-dimensional unit circle $S^{1}$ (e.g., $h^{\perp}=c_{1} e_{1}+c_{2} e_{2}$ and $c_{1}^{2}+c_{2}^{2}=1$ ). To simplify notation, we use a point $\left(c_{1}, c_{2}\right) \in S^{1}$ and its corresponding unit vector $c_{1} e_{1}+c_{2} e_{2}$ interchangeably. In this setting, it is easy to see that for any $x \in \operatorname{span}\left\{e_{1}, e_{2}\right\}$

$$
\begin{equation*}
\boldsymbol{\operatorname { s i g n }}\left\langle x, h^{*}\right\rangle=\boldsymbol{\operatorname { s i g n }}\left\langle x, h^{\perp}\right\rangle . \tag{2}
\end{equation*}
$$

We take a Bayesian approach. In the beginning, when no queries have been performed, $\mathrm{DC}^{2}$ assumes no prior information about the vector $h^{\perp}$. Therefore, it takes the uniform distribution on $S^{1}$ (with pdf $p_{0}(h)=\frac{1}{2 \pi}$ ) as its prior belief about $h^{\perp}$. After performing each query, the posterior (belief) about $h^{\perp}$ will be updated according to the observation. We let $p_{m}(h)$ denote the (pdf of the) posterior after performing the first $m$ queries. In this manner, $\mathrm{DC}^{2}$ runs in total of $T_{\epsilon, \delta}$ rounds, where in each round a specific query is selected and posed to the oracle. The number $T_{\epsilon, \delta}$ will be specified later (see Theorem 1). Upon the completion of round $T_{\epsilon, \delta}$, the algorithm returns as its final output a vector $\hat{e} \in S^{1}$ that maximises the posterior pdf $p_{T_{\epsilon, \delta}}(h)$. If there are multiple such maximisers, it picks one arbitrarily. We now proceed with a detailed description of $\mathrm{DC}^{2}$ (a formal description is provided in Algorithm 1).

As shown in Algorithm 1, at each round, say round $m+1$, the algorithm maintains and updates the distribution $p_{m}$ that encodes its current belief in the true location of $h^{\perp}$. We should note here that these distributions can be stored efficiently and as a result the vector $x_{m+1}$ can be computed efficiently. Indeed, (the pdf of) $p_{m}$ is piecewise constant on the unit circle (see Figure 1). More precisely, at any round $m$, there are at most $2 m$ points $u_{1}, u_{2}, \cdots, u_{2 m}$ that are ordered clock-wise on the unit-circle and $p_{m}$ is constant when restricted to each of the sectors $\left[u_{i}, u_{i+1}\right)$. The piecewise
constant property of the pdf of $p_{m}$ can be established by induction on $m$. Recall that the initial distribution $p_{0}$ is uniform and thus piecewise constant. The Bayesian update step (line 5 of Algorithm 1) preserves this property when the algorithm updates the distribution $p_{m}(h)$ to $p_{m+1}(h)$. We will show why this is true when we discuss the Bayesian update step in detail.

At round $m+1$, in order to find $x_{m+1}$ (see line 3 of Algorithm 1), $\mathrm{DC}^{2}$ first finds a line that passes through the centre of $S^{1}$ and cuts $S^{1}$ into two "halves" which have the same measure with respect to $p_{m}$. Note that finding such a line can be done in $O(m)$ steps because $p_{m}$ has the piecewise constant property. Once such a line is found, it is then easy to see that $x_{m+1}$ can be any of the two points orthogonal to the line. As a result, $\mathrm{DC}^{2}$ at round $m+1$ can find $x_{m+1}$ in $O(m)$ operations. We denote the half-circle containing $x_{m+1}$ by $R^{+}$and the other half by $R^{-}$. We refer to Figure 1 for a schematic illustration.

The key step in Algorithm 1 is the Bayesian update (line 5). Once a noisy response to the query $\operatorname{sign}\left\langle x_{m+1}, h^{*}\right\rangle$ is obtained (line 4)), the probability distribution $p_{m}$ will be updated to $p_{m+1}$ in the following way. First, consider the event that the outcome of $\boldsymbol{\operatorname { s i g n }}\left\langle x_{m+1}, h^{*}\right\rangle$ is +1 . We have $p_{m}\left(\boldsymbol{\operatorname { s i g n }}\left\langle x_{m+1}, h^{*}\right\rangle=+1\right)=(1-\rho) p_{m}\left(R^{+}\right)+\rho p_{m}\left(R^{-}\right)=$ $1 / 2$, and similarly $p_{m}\left(\boldsymbol{\operatorname { s i g n }}\left\langle x_{m+1}, h^{*}\right\rangle=-1\right)=1 / 2$. Therefore, by Bayes theorem we obtain the following update rules for $p_{m+1}$. If we observe that the outcome of $\operatorname{sign}\left\langle x_{m+1}, h^{*}\right\rangle$ is +1 , then for $h \in R^{+}$we have $p_{m+1}(h)=2(1-\rho) p_{m}(h)$ and for $h \in R^{-}$we have $p_{m+1}(h)=(2 \rho) p_{m}(h)$. Also, if we observe that the outcome of $\operatorname{sign}\left\langle x_{m+1}, h^{*}\right\rangle$ is -1 , then for $h \in R^{+}$we have $p_{m+1}(h)=(2 \rho) p_{m}(h)$ and for $h \in R^{-}$ we have $p_{m+1}(h)=2(1-\rho) p_{m}(h)$. Note that the factor of 2 here is due to the normalization. It is easy to verify that $p_{m+1}$ is also a piecewise constant distribution (now on $2(m+1)$ sectors; see Figure 1).

Theorem 1 shows that after $T_{\epsilon, \delta}=O\left(\log \frac{1}{\epsilon}+\log \frac{1}{\delta}\right)$ rounds, with probability at least $1-\delta, \mathrm{DC}^{2}$ outputs a unit vector $\hat{e} \in \operatorname{span}\left\{e_{1}, e_{2}\right\}$ such that $\left\|\hat{e}-h^{\perp}\right\|<\epsilon$. Also, as discussed above, the computational complexity of $\mathrm{DC}^{2}$ is $O\left(T_{\epsilon, \delta}^{2}\right)$, i.e., $O\left(\left(\log \frac{1}{\epsilon}+\log \frac{1}{\delta}\right)^{2}\right)$.
Theorem 1. When the independent flip probability is $\rho$, having

$$
\begin{equation*}
T_{\epsilon, \delta} \geq M+\max \left\{T_{0}, T_{1}, T_{2}, T_{3}\right\}=O\left(\log \frac{1}{\epsilon}+\log \frac{1}{\delta}\right) \tag{3}
\end{equation*}
$$

is sufficient to guarantee that $\mathrm{DC}^{2}$ outputs with probability at least $1-\delta$ a vector that is within a distance $\epsilon$ of $h^{\perp}$. Here, we have $M=\left\lceil\frac{2 \log \frac{2}{\delta}}{-\log (4 \rho(1-\rho))}\right\rceil, T_{0}=\frac{8 \log \frac{2}{\delta}}{\log (2(1-\rho))}, T_{1}=$ $\frac{8 \log \frac{1}{8 \pi \epsilon}}{\log (2(1-\rho))}, T_{2}=\frac{8}{\log (2(1-\rho))}\left(\log (2 M)+\log \left(\frac{4}{\log (2(1-\rho))}\right)\right)$ and $T_{3}=\frac{24 \rho \log ^{2} \frac{1-\rho}{\rho}}{\log ^{2}(2(1-\rho))}\left(\log (M)+\log \left(\frac{4}{\delta}\right)\right)$.

We would like to remark that when the independent flip probability $\rho$ is 0 (i.e., in the noiseless case), the algorithm $\mathrm{DC}^{2}$ reduces to the binary search. If we let $T_{\epsilon, \delta}=\left\lceil\log _{2} \frac{\pi}{\epsilon}\right\rceil$, then $\mathrm{DC}^{2}$ outputs a vector that is within a distance $\epsilon$ of $h^{\perp}$. We present a detailed discussion with examples in the extended version.

round $m$

$$
p_{m}\left(R^{+}\right)=p_{m}\left(R^{-}\right)=\frac{1}{2}
$$


round $m+1$

Figure 1: Upon the completion of round $m$ (left figure), the distribution (pdf of) $p_{m}$ is constant over each of the sectors [ $u_{i}, u_{i+1}$ ). In the next round (right figure), in order to find $x_{m+1}, \mathrm{DC}^{2}$ first finds a diagonal line (red line) which separates two half-circles ( $R^{+}$and $R^{-}$) that each has measure $1 / 2$ w.r.t $p_{m}$. The vector $x_{m+1}$ will then be one of the two points on the unit circle that are orthogonal to this line. For updating $p_{m}$ to $p_{m+1}$, we note that all the points inside $R^{+}$get the same factor (either $2 \rho$ or $2(1-\rho)$ depending on the outcome of the query). The same is true for $R^{-}$. Thus, $p_{m+1}$ is again a piecewise constant pdf but now on $2(m+1)$ sectors.

A few comments are in order: The above guarantee for $\mathrm{DC}^{2}$ holds with probability one and thus the parameter $\delta$ is irrelevant in the noiseless setting. Furthermore, during each round of $\mathrm{DC}^{2}$, the distribution $p_{m}$ can be represented by only two numbers (the starting and ending points of the version space after $m$ rounds $R_{m}$, which is a sector), and the vector $x_{m}$ can be computed efficiently (it is the orthogonal vector to the midpoint of $R_{m}$ ). Therefore, assuming one unit of complexity for performing the queries, $\mathrm{DC}^{2}$ can be implemented with complexity $O\left(T_{\epsilon, \delta}\right)$, i.e., $O(\log (1 / \epsilon))$.

## 3 Dimension Coupling Based Framework

In Section 2, we devise an algorithm, called $\mathrm{DC}^{2}$ ( $e_{1}, e_{2}, \epsilon, \delta$ ), that takes as input two orthonormal vectors $e_{1}, e_{2}$, uses noisy responses to queries of the form $\boldsymbol{\operatorname { s i g n }}\left\langle x, h^{*}\right\rangle$, and outputs with probability at least $1-\delta$ a vector $\hat{e}$ with the following three properties:

$$
\hat{e} \in \operatorname{span}\left\{e_{1}, e_{2}\right\},\|\hat{e}\|=1,\left\|\hat{e}-\frac{\left\langle h^{*}, e_{1}\right\rangle e_{1}+\left\langle h^{*}, e_{2}\right\rangle e_{2}}{\left\|\left\langle h^{*}, e_{1}\right\rangle e_{1}+\left\langle h^{*}, e_{2}\right\rangle e_{2}\right\|}\right\|<\epsilon
$$

In other words, the unit vector $\hat{e}$ is within a distance $\epsilon$ to the (normalized) projection of $h^{*}$ onto the subspace $\operatorname{span}\left\{e_{1}, e_{2}\right\}$. In the current section, we explain a framework DC that estimates $h^{*}$ using at most $d-1$ calls to $\mathrm{DC}^{2}$ (a formal description will be given in Algorithm 2 later).

Let us begin our discussion with a motivating example. Let $\left\{e_{1}, e_{2}, \ldots, e_{d}\right\}$ be an orthonormal basis of $\mathbb{R}^{d}$. Suppose that $h^{*}$ has the form $h^{*}=\sum_{i=1}^{d} c_{i} e_{i}$, where $\left\{e_{i}\right\}_{i=1}^{d}$ is an arbitrarily chosen orthonormal basis for $\mathbb{R}^{d}$. We assume w.l.o.g. that $h^{*}$ is normalized (i.e., $\sum_{i=1}^{d} c_{i}^{2}=1$ ). Our objective is then to learn the coefficients $\left\{c_{i}\right\}_{i=1}^{d}$ within a given precision by using the noisy responses to the selected sign queries. The key insight here is that this task can be partitioned in

(b) Scheme 2: an unbalanced full binary tree

Figure 2: Two possible divide-and-conquer schemes for a 4dimensional problem. Each scheme can be represented by a full binary tree of 4 leaf nodes.
a divide-and-conquer fashion into many smaller tasks, each involving a few dimensions. The final answer (the values of $\left\{c_{i}\right\}_{i=1}^{d}$ ) will then be obtained by aggregating the answers of these subproblems.

In Figure 2, we present two possibilities of divide-andconquer schemes for a 4-dimensional problem. In fact, each scheme corresponds to a full binary tree of 4 leaf nodes.

Assume $h^{*}=c_{1} e_{1}+c_{2} e_{2}+c_{3} e_{3}+c_{4} e_{4}$, where $e_{i}$ 's are the standard basis vectors for $\mathbb{R}^{4}$. Define $e_{12}=\frac{c_{1} e_{1}+c_{2} e_{2}}{\sqrt{c_{1}^{2}+c_{2}^{2}}}, e_{34}=$ $\frac{c_{3} e_{3}+c_{4} e_{4}}{\sqrt{c_{3}^{2}+c_{4}^{2}}}$. Note here that $e_{12}$ is the (normalized) orthogonal projection of $h^{*}$ onto $\operatorname{span}\left\{e_{1}, e_{2}\right\}$ and $e_{34}$ is the (normalized) orthogonal projection of $h^{*}$ onto $\operatorname{span}\left\{e_{3}, e_{4}\right\}$. Consider the following procedure to learn $h^{*}$ : first find out what $e_{12}$ and $e_{34}$ are, and then use the relation $h^{*}=$ $\sqrt{c_{1}^{2}+c_{2}^{2}} e_{12}+\sqrt{c_{3}^{2}+c_{4}^{2}} e_{34}$ to find $h^{*}$ based on the orthonormal vectors $e_{12}, e_{34}$. By this procedure, the original "fourdimensional" problem has been broken into three "twodimensional" problems.

This procedure is illustrated in Figure 2a. We first call $\mathrm{DC}^{2}\left(e_{1}, e_{2}\right)$ to obtain an estimate $\hat{e}_{12}$ for $e_{12}$; then we call $\mathrm{DC}^{2}\left(e_{3}, e_{4}\right)$ to obtain an estimate $\hat{e}_{34}$ for $e_{34}$; finally we call DC ${ }^{2}\left(\hat{e}_{12}, \hat{e}_{34}\right)$ to obtain an estimate $\hat{h}$ for $h^{*}$. Another unbalanced scheme is illustrated in Figure 2b.

For general $d$, the idea is similar: We break the problem into at most $d-1$ "two-dimensional" problems that each can be solved efficiently. Again, each divide-and-conquer
scheme corresponds to a full binary tree of $d$ leaf nodes.
Consider the decomposition $h^{*}=\sum_{i=1}^{d} c_{i} e_{i}$. Without loss of generality, suppose that the first two leaf nodes to be combined are $e_{1}$ and $e_{2}$. We can write

$$
\begin{equation*}
h^{*}=\sum_{i=1}^{d} c_{i} e_{i}=\hat{c}_{12} \frac{c_{1} e_{1}+c_{2} e_{2}}{\sqrt{c_{1}^{2}+c_{2}^{2}}}+\sum_{i=3}^{d} c_{i} e_{i} \tag{4}
\end{equation*}
$$

where in the last step we have taken $\hat{c}_{12} \triangleq \sqrt{c_{1}^{2}+c_{2}^{2}}$. Now, note that $\hat{e}_{12} \triangleq \frac{c_{1} e_{1}+c_{2} e_{2}}{\sqrt{c_{1}^{2}+c_{2}^{2}}}$ is the normalized orthogonal projection of $h^{*}$ onto span $\left\{e_{1}, e_{2}\right\}$. Hence, by using DC ${ }^{2}$ $\left(e_{1}, e_{2}, \epsilon, \delta\right)$ we can obtain, with probability at least $1-\delta$, a good approximation $\hat{e}_{12}$ (within a distance $\epsilon$ ) of this projection. Therefore, for small enough $\epsilon$ we have $h^{*} \approx \hat{c}_{12} \hat{e}_{12}+$ $\sum_{i=3}^{d} c_{i} e_{i}$. Since $h^{*}$ is now expressed (approximately) in terms of $d-1$ orthonormal vectors $\left\{\hat{e}_{12}, e_{3}, e_{4}, \ldots, e_{d}\right\}$, we have effectively reduced the dimensionality of problem from $d$ to $d-1$. The idea is then to repeat the same procedure as in (4) to the newly obtained representation of $h^{*}$. Hence, by repeating this procedure $d-1$ times in total we will reach a vector which is the final approximation of $h^{*}$.

We present this general method in Algorithm 2.

```
Algorithm 2 Dimension Coupling (DC)
Input: an orthonormal basis \(E=\left\{e_{1}, e_{2}, \ldots, e_{d}\right\}\) of \(\mathbb{R}^{d}\).
Output: a unit vector \(\hat{h}\) which is an estimate for \(h^{*}\).
    for \(j \leftarrow 1\) to \(d-1\) do
        Replace any two vectors \(e^{\prime}\) and \(e^{\prime \prime}\) in \(E\) with the vec-
        tor \(\mathrm{DC}^{2}\left(e^{\prime}, e^{\prime \prime}, \epsilon, \delta\right)\).
    end for
    Let \(\hat{h}\) be the only remaining vector in \(E\).
    return \(\hat{h}\)
```

Theorem 2. For DC (outlined in Algorithm 2) and any of its divide-and-conquer scheme represented by a full binary tree, we have:

1. DC will call the two-dimensional subroutine $\mathrm{DC}^{2} d-1$ times.
2. Provided that the output of $\mathrm{DC}^{2}$ is with probability $1-\delta$ within distance $\epsilon$ of the true value and $\epsilon \leq 5 / 18$, DC ensures an estimation error of at most $5 \epsilon(d-1)$ with probability at least $1-\delta(d-1)$.
As a result of Theorem 2, if we desire the framework DC to estimate $h^{*}$ within distance $\tilde{\epsilon}$ and with probability at least $1-\tilde{\delta}$, then it is enough to fix the corresponding parameters of $\mathrm{DC}^{2}$ to $\epsilon=\frac{\tilde{\epsilon}}{5(d-1)}$ and $\delta=\frac{\tilde{\delta}}{d-1}$.

Theorem 2 indicates that DC requires $\tilde{O}\left(d\left(\log \frac{1}{\epsilon}+\log \frac{1}{\delta}\right)\right)$ queries, since each call to $\mathrm{DC}^{2}$ needs $O\left(\log \frac{1}{\epsilon}+\log \frac{1}{\delta}\right)$ queries. Recall that the computational complexity of $\mathrm{DC}^{2}$ is $O\left(\left(\log \frac{1}{\epsilon}+\log \frac{1}{\delta}\right)^{2}\right)$. Hence, DC has computational complexity $\tilde{O}\left(d\left(\log \frac{1}{\epsilon}+\log \frac{1}{\delta}\right)^{2}\right)$. As a special case, if in absence of noise, both the query complexity and time complexity of DC are $\tilde{O}\left(d \log \frac{1}{\epsilon}\right)$.


Figure 3: Figures 3a and 3b show the estimation error in the noiseless setting as we increase the number of queries, for $d=25$ and 100 , respectively. Figure 3 c shows the corresponding execution times. Figure 3 d and 3 e show the scatter plots of the execution time and the estimation error of different methods for $d=25,50$ and the noise level $\rho=0.1$. We allow each algorithm to use a budget of 800 and 1800 queries in Figure 3d and 3e, respectively. Figure 3 f presents the estimation error of DC and Repetitive-DC as we increase the number of queries for $d=1000$ and noise levels $\rho=0.01,0.1,0.2$.

## 4 Empirical Results

In this section, we extensively evaluate the performance of DC against the following baselines:

RANDOM-SAMPLING: Queries are generated by sampling uniformly at random from the unit sphere $S^{d-1}$.

UnCERTAINTY-SAMPLING: Queries are sampled uniformly at random from the orthogonal complement of $w$, where $w$ is the vector learned by linear SVM.

Query-by-Bagging: The bag size is set to 20 and 1000 queries are generated at each iteration. The query with the largest disagreement is picked (Abe and Mamitsuka 1998).

Spectral: The version space is approximated by the largest ellipsoid consistent with all previous query-label pairs. Then, at each iteration a query is selected to approximately halve the ellipsoid (Alabdulmohsin, Gao, and Zhang 2015).

Repetitive-DC: In the noisy setting, one easy way to apply DC is to query each point $R$ times and use the majority rule to determine its label; i.e., the combination of repetitive querying (Section 2) and the DC framework (Section 3).

Our metrics to compare different algorithms are: a) estimation error, b) query complexity, and c) execution time. In particular, as we increase the number of queries we measure the average estimation errors and execution times for all the baselines (with $90 \%$ confidence intervals). By nature, in active learning via query synthesis, all data points and queries
are generated synthetically. For all the baselines, we used the fastest available implementations in MATLAB.

Noiseless setting: Figures 3a and 3b (with dimension $d=25$ and 50, respectively) show that in terms of estimation error, DC outperforms all other baselines, and significantly outperforms Random-SAmpling, UncertaintySampling and Query-by-Bagging. Note that the estimation errors are plotted in log-scales. In terms of execution times, we see in Fig. 3c that DC runs three orders of magnitude faster than other baselines. Training an SVM at each iteration for Random-Sampling, UncertaintySAmpling and Query-by-Bagging comes with a huge computational cost. Similarly, Spectral requires solving a convex optimization problem at each iteration; thus its performance drastically deteriorates as the dimension increases, which makes it infeasible for many practical problems.

Noisy setting: We set the noise level to $\rho=0.1$ and compare the performance of DC against RANDOM-SAMPLING, Uncertainty-Sampling, Query-by-Bagging, and Repetitive-DC. As mentioned in (Alabdulmohsin, Gao, and Zhang 2015), and we have also observed in our experiments, Spectral does not work even for small amounts of noise as it incorrectly shrinks the version space and misses the true linear separator; therefore it is excluded here. We see again in Figures 3d and 3e (for $d=25$ and 50) that DC significantly outperforms all other methods in terms of estimation error. More precisely, using the same number of
queries, the estimation error of DC is around two orders of magnitude smaller than other baselines. We can also observe from these two figures that DC still runs around 100 times faster than Random-Sampling, UncertaintySAmpling, and Query-by-Bagging. Clearly, DC has a higher computational cost than Repetitive-DC, as DC performs a Bayesian update after each query. Finally, as we increase the dimension to $d=1000$, RANDOM-SAMPLING, Uncertainty-Sampling, and Query-by-BagGing become significantly slower. Hence, in Figure 3f we only show how the estimation error (for noise levels $\rho=0.01,0.1,0.2$ ) decreases for DC and REPETITIVE-DC with more queries. It can be observed from Figure $3 f$ that consuming the same number of queries, DC can achieve an estimation error from one order (when the noise intensity is very small) to three orders of magnitude (when the noise intensity is 0.2 ) smaller than that of Repetitive-DC.

## 5 Related Work

The sample complexity of learning a hypothesis was traditionally studied in the context of probably approximately correct (PAC) learning (Valiant 1984). In PAC learning theory, one assumes that a set of hypotheses $\mathcal{H}$ along with a set of unlabeled data points $\mathcal{X}$ are given, where each data point $x \in \mathcal{X}$ is drawn i.i.d. from some distribution $D$. Classical PAC bounds then yield the sample complexity (i.e., the number of required i.i.d. examples) from $D$ to output a hypothesis $h \in \mathcal{H}$ that will have estimation error at most $\epsilon$ with probability at least $1-\delta$, for some fixed $\epsilon, \delta>0$. Here, the estimation error is defined as $\epsilon=\operatorname{Pr}_{x \sim D}\left[h(x) \neq h^{*}(x)\right]$, where $h^{*}$ is the unknown true hypothesis. In the realizable case of learning a halfspace, i.e., when $h^{*} \in \mathbb{R}^{d}$ perfectly separates the data points into positive and negative labels, it is known that with $\tilde{O}(d / \epsilon)$ i.i.d. samples one can find a linear separator with an estimation error $\epsilon$. The main advantage of using active learning methods, i.e., sequentially querying data points, is to reduce the sample complexity exponential fast, ideally to $\tilde{O}(d \log (1 / \epsilon))$. In fact, a simple counting argument based on sphere packing shows that any algorithm needs $\Omega(d \log (1 / \epsilon))$ examples to achieve an estimation error of $\epsilon$ (Dasgupta, Kalai, and Monteleoni 2009).

For $d=2$ and when the distribution is uniform over the unit sphere $S^{1}$ it is very easy to see that the halving or bisection leads to $\tilde{O}(\log (1 / \epsilon))$. By using the same halving method, one can in principle extend the result to any dimension $d$. To do so, we need to carefully construct the version space (i.e., the set of hypotheses consistent with the queries and outcomes) at each iteration and then find a query that halves the volume (in the uniform case) or the density (in the general case if the distribution is known) (Dasgupta 2004). Finding such a query in high dimension is very challenging.

One very successful approach that does not suffer from the aforementioned computational challenge is pool-based active learning (Settles 2010), where instead of ideally halving the space, effective approaximations are performed. Notable algorithms are uncertainty sampling (Lewis and Gale 1994) and query-by-committee (QBC) (Freund et al. 1997). In fact, our problem is closely related to learning homo-
geneous linear separators under the uniform distribution in the pool-based setting. This problem is very well understood and there exist efficient pool-based algorithms (Balcan, Broder, and Zhang 2007; Dasgupta, Kalai, and Monteleoni 2005; Dasgupta and Hsu 2008). In particular, Dasgupta et al. (Dasgupta, Kalai, and Monteleoni 2009) presented an efficient perceptron-based algorithm that achieve a near-optimal query complexity. Similar results can be obtained under log-concave distributions (Balcan and Long 2013). Most of the pool-based methods require to have access to $\tilde{O}(1 / \epsilon)$ number of unlabeled samples in each iteration or otherwise they perform very poorly (Balcan, Broder, and Zhang 2007; Dasgupta, Kalai, and Monteleoni 2009). This means that in order to have exponential guarantee in terms of sample complexity, we need to grow the pool size exponentially fast (note that there is no need to store all of these points). Moreover, with a few exceptions (Awasthi, Balcan, and Long 2014; Balcan, Beygelzimer, and Langford 2006) pool-based learning of linear separators in the noisy setting has been much less studied and the dependency of sample complexity on noise is not very well understood.

An attractive alternative to the pool-based framework is query synthesis where we have access to membership queries (Angluin 1988)): a learner can request for any unlabeled data instance from the input space, including queries that the learner synthesizes from scratch. This way the pool size limitation is entirely eliminated. In many recent applications, ranging from automated science (King et al. 2009), to robotics (Cohn, Ghahramani, and Jordan 1996), and to adversarial reverse engineering (Lowd and Meek 2005), query synthesis is the appropriate model. For instance, in securitysensitive applications (e.g., spam filters and intrusion detection systems) that routinely use machine learning tools, a growing concern is the ability of adversarial attacks to identify the blind spots of the learning algorithms. Concretely, classifiers are commonly deployed to detect miscreant activities. However, they are attacked by adversaries who generate exploratory queries to elicit information that in return allows them to evade detection (Nelson et al. 2012). In this work, we show how an adversary can use active learning methods by making synthetically de novo queries and thus identify the linear separator used for classification. We should emphasize that in active learning via synthesized queries the learning algorithm can query the label of any points in order to explore the hypothesis space. In the noiseless setting (if we ignore the dependency of the pool size on $\tilde{O}(\log (1 / \epsilon)))$, one can potentially use the pool-based algorithms (under the uniform distribution). Our main contribution in this paper is to develop a noise resilient active learning algorithm that has access to noisy membership queries. To the best of our knowledge, we are the first to show a near optimal algorithm that outperforms in theory and practice the naive repetition mechanism and the recent spectral heuristic methods (Alabdulmohsin, Gao, and Zhang 2015).

## Acknowledgements

LC and AK were supported by DARPA Young Faculty Award (D16AP00046). HH was supported partially by ERC

StG 307036. LC thanks Zheng Wei for his support.

## References

Abe, N., and Mamitsuka, H. 1998. Query learning strategies using boosting and bagging. In ICML, volume 1.
Alabdulmohsin, I.; Gao, X.; and Zhang, X. 2015. Efficient active learning of halfspaces via query synthesis. In $A A A I$, 2483-2489.
Angluin, D. 1988. Queries and concept learning. Machine learning 2(4):319-342.
Awasthi, P.; Balcan, M. F.; and Long, P. M. 2014. The power of localization for efficiently learning linear separators with noise. In STOC, 449-458. ACM.

Balcan, M.-F., and Long, P. M. 2013. Active and passive learning of linear separators under log-concave distributions. In COLT, 288-316.

Balcan, M.-F.; Beygelzimer, A.; and Langford, J. 2006. Agnostic active learning. In ICML, 65-72. ACM.

Balcan, M.-F.; Broder, A.; and Zhang, T. 2007. Margin based active learning. In Learning Theory. Springer. 35-50.

Chen, Y.; Hassani, S. H.; Karbasi, A.; and Krause, A. 2015. Sequential information maximization: When is greedy nearoptimal? In COLT, volume 40.
Chen, L.; Hassani, H.; and Karbasi, A. 2016. Nearoptimal active learning of halfspaces via query synthesis in the noisy setting. arXiv preprint arXiv:1603.03515. https://arxiv.org/abs/1603.03515.
Cohn, D. A.; Ghahramani, Z.; and Jordan, M. I. 1996. Active learning with statistical models. JAIR.
Dasgupta, S., and Hsu, D. 2008. Hierarchical sampling for active learning. In ICML, 208-215. ACM.
Dasgupta, S.; Kalai, A. T.; and Monteleoni, C. 2005. Analysis of perceptron-based active learning. In COLT, 249-263.
Dasgupta, S.; Kalai, A. T.; and Monteleoni, C. 2009. Analysis of perceptron-based active learning. JMLR 10(Feb):281299.

Dasgupta, S. 2004. Analysis of a greedy active learning strategy. In NIPS, 337-344.
Freund, Y.; Seung, H. S.; Shamir, E.; and Tishby, N. 1997. Selective sampling using the query by committee algorithm. Machine learning 28(2-3):133-168.
Hakkani-Tur, D.; Riccardi, G.; and Gorin, A. 2002. Active learning for automatic speech recognition. In ICASSP, volume 4, IV-3904. IEEE.
Javdani, S.; Chen, Y.; Karbasi, A.; Krause, A.; Bagnell, J. A.; and Srinivasa, S. 2014. Near optimal bayesian active learning for decision making. In AISTAT, 430-438.
Kääriäinen, M. 2006. Active learning in the non-realizable case. In ALT, 63-77. Springer.
Karbasi, A.; Ioannidis, S.; and Massoulie, L.
2012. Comparison-based learning with rank nets. ICML.
Karp, R. M., and Kleinberg, R. 2007. Noisy binary search and its applications. In SODA, 881-890.

King, R. D.; Rowland, J.; Oliver, S. G.; Young, M.; Aubrey, W.; Byrne, E.; Liakata, M.; Markham, M.; Pir, P.; Soldatova, L. N.; et al. 2009. The automation of science. Science 324(5923):85-89.
Lewis, D. D., and Gale, W. A. 1994. A sequential algorithm for training text classifiers. In SIGIR, 3-12.
Lowd, D., and Meek, C. 2005. Adversarial learning. In KDD, 641-647.
Nelson, B.; Rubinstein, B. I.; Huang, L.; Joseph, A. D.; Lee, S. J.; Rao, S.; and Tygar, J. 2012. Query strategies for evading convex-inducing classifiers. JMLR 13(May):1293-1332.
Nowak, R. D. 2011. The geometry of generalized binary search. IEEE Transactions on Information Theory 57(12):7893-7906.
Settles, B. 2010. Active learning literature survey. University of Wisconsin, Madison 52(55-66):11.
Shawe-Taylor, J., and Cristianini, N. 2004. Kernel methods for pattern analysis. Cambridge university press.
Tong, S., and Chang, E. 2001. Support vector machine active learning for image retrieval. In Proc. 9th ACM international conference on Multimedia, 107-118.
Tong, S., and Koller, D. 2002. Support vector machine active learning with applications to text classification. JMLR 2:4566.

Valiant, L. G. 1984. A theory of the learnable. Communications of the ACM 27(11):1134-1142.


[^0]:    Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[^1]:    ${ }^{1}$ We use the $\tilde{O}$ notation to ignore logarithmic factors. In Section 5, the $\tilde{O}$ notation also ignores terms dependent on $\delta$.

