On Kinetic Range Spaces and their Applications

Jean-Lou De Carufel* Matthew J. Katz† Matias Korman‡ André van Renssen§ ¶
Marcel Roeloffzen§ ¶ Shakhar Smorodinsky‖

Abstract

We study geometric hypergraphs in a kinetic setting and show that for many of the static cases where the VC-dimension of the hypergraph is bounded the kinetic counterpart also has bounded VC-dimension. Among other results we show that for any set of n moving points in $\mathbb{R}^d$ and any parameter $1 < k < n$, one can select a non-empty subset of the points of size $O(k \log k)$ such that each cell of the Voronoi diagram of this subset is "balanced" at any given time (i.e., it contains $O(n/k)$ of the other points). We also show that the bound is near optimal even for the one-dimensional case in which points move linearly.

1 Introduction

Geometric hypergraphs (also called range-spaces) are central objects in computational geometry, statistical learning theory, combinatorial optimization, linear programming, discrepancy theory, data bases and several other areas in mathematics and computer science.

In most of these cases, we have a finite set P of points in $\mathbb{R}^d$ and a family of simple geometric regions, such as the family of all halfspaces in $\mathbb{R}^d$. We then consider the combinatorial structure of the set system $(P, \{h \cap P\})$ where h is any halfspace. Many optimization problems can be formulated on such structures. A key property of such hypergraphs is a bounded VC-dimension (see Section 2 for exact definitions). In this paper we study a more complex structure by allowing the underlying set of points to move along some "reasonable" trajectories. Even though the static case is well-known, little research has been done for the case in which the points move. We show that those more complex hypergraphs defined as the union of all hypergraphs obtained at all possible times still have a bounded VC-dimension. As a result, many deep results that hold for arbitrary hypergraphs with bounded VC-dimension readily apply to such kinetic hypergraphs. By adding several other ingredients, we are able to prove our main result about points moving with bounded description complexity (see below for the exact definitions):

Theorem 1 *Let $P = \{p_1, \ldots, p_n\}$ be any set of n moving points in $\mathbb{R}^d$ with bounded description complexity. For any integer $2 \le k \le n$, there exists a subset $N \subset P$ of cardinality $O(k \log k)$, such that for any time $t \ge 0$, each cell of the Voronoi diagram $Vor(N(t))$ contains at most $O(n/k)$ points of $P(t)$.*

The paper is organized as follows: in Section 2 we introduce several key concepts as well as review known results that hold for static range spaces. In Section 3 we extend these results to the kinetic case. In Section 4 we show several applications. Due to lack of space, proofs in this paper are omitted or sketched. Details can be found in the extended version of this paper [2].

2 Preliminaries and Previous Work

We consider the following families of geometric hypergraphs: Let P be a set of points in $\mathbb{R}^d$ and let $\mathcal{R}$ be a family of regions in the same space. We refer to the hypergraph $H = (P, \{P \cap r : r \in \mathcal{R}\})$ as the hypergraph induced by P with respect to $\mathcal{R}$. In the literature, such kind of hypergraphs are also referred to as *range spaces*.

Our aim is to show that many properties that hold for static range spaces extend to their kinetic counterparts. We start by introducing some concepts that are frequently used in (static) range spaces.

Recall the definition of an ε-net for a hypergraph: let $H = (V, \mathcal{E})$ be a hypergraph with V finite. Let $\varepsilon \in [0, 1]$ be a real number. A set $N \subset V$ (not necessarily in $\mathcal{E}$) is called an *ε-net* for H if for every hyperedge $S \in \mathcal{E}$ with $|S| \ge \varepsilon |V|$ we have $S \cap N \neq \emptyset$.

A closely related concept of ε-net is the so called Vapnik-Chervonenkis dimension [8]: let $H = (V, \mathcal{E})$ be a hypergraph. A subset $X \subset V$ (not necessarily in $\mathcal{E}$) is said to be *shattered* by H if $\{X \cap S : S \in \mathcal{E}\} = 2^X$. The *Vapnik-Chervonenkis dimension*, also

*School of Computer Science, Carleton University, Ottawa, Canada.
†Department of Computer Science, Ben-Gurion University, Beer-Sheva, Israel.
‡Tohoku University, Sendai, Japan. Work supported by the ELC project (MEXT KAKENHI No. 24106008).
§National Institute of Informatics (NII), Tokyo, Japan.
¶JST, ERATO, Kawarabayashi Large Graph Project.
‖Department of Mathematics, Ben-Gurion University of the NEGEV, Beer-Sheva, Israel and EPFL, Switzerland. Research partially supported by Swiss National Science Foundation Grants 200020144531 and 200021-137574, and by Grant 1136/12 from the Israel Science Foundation.

denoted the VC-*dimension* of H, is the maximum size of a subset of V shattered by H.

Theorem 2 (ε-net theorem [4]) *Let $H = (V, \mathcal{E})$ be a hypergraph with VC-dimension d. For every $\varepsilon \in (0, 1)$, there exists an ε-net $N \subset V$ with cardinality at most $O(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon})$.*

It is known that whenever range spaces are defined through semi-algebraic sets of constant description complexity (i.e., sets defined as a Boolean combination of a constant number of polynomial equations and inequalities of constant maximum degree), the resulting hypergraph has finite VC-dimension. Halfspaces, balls, boxes, etc. are examples of ranges of this kind; see, e.g., [6, 7] for more details.

Thus, by Theorem 2, these hypergraphs admit "small" size ε-nets. Kómlos *et al.* [5] proved that the bound $O(\frac{d}{\varepsilon} \log \frac{1}{\varepsilon})$ on the size of an ε-net for hypergraphs with VC-dimension d is best possible.

3 Kinetic hypergraphs

In this section we extend the above results to the kinetic model. Let $P = \{p_1, \ldots, p_n\}$ denote a set of n moving points in $\mathbb{R}^d$, where each point is moving along some "simple" trajectory. That is, each p_i is a function $p_i : [0, \infty) \to \mathbb{R}^d$ of the form $p_i(t) = (x^i_1(t), \ldots, x^i_d(t))$. For a given real number $t \geq 0$ and a subset $P' \subset P$, we denote by $P'(t)$ the fixed set of points $\{p(t) \colon p \in P'\}$. For a family of ranges $\mathcal{R}$, we define its induced *kinetic hypergraph* as follows:

Definition 1 (kinetic hypergraph) *Let P be a set of moving points in $\mathbb{R}^d$ and let $\mathcal{R}$ be a family of ranges. Let $(P, \mathcal{E})$ denote the hypergraph where $\mathcal{E}$ consists of all subsets $P' \subseteq P$ for which there exists a time t and a range $r \in \mathcal{R}$ such that $P'(t) = P(t) \cap r$. We call $(P, \mathcal{E})$ the* kinetic hypergraph *induced by $\mathcal{R}$.*

As in the static case we abuse the notation and denote the hypergraph by $(P, \mathcal{R})$. In order to apply our techniques, we need the following "bounded description complexity" assumption concerning the movement of the points of P. We say that a point $p_i = p_i(t) = (x^i_1(t), \ldots, x^i_d(t)) \in P$ moves with *description complexity* $s > 0$ if for each $1 \leq j \leq d$ it holds that $x^i_j(t)$ is a univariate polynomial of degree at most s. In the remainder of this paper, we assume that all points of P move with bounded description complexity, that is, the description complexity s is a constant.

3.1 VC-Dimension of kinetic hypergraphs

In this section we prove that for many of the static range spaces that have small VC-dimension, their kinetic counterparts also have small VC-dimension. We start with the family $\mathcal{H}_d$ of all halfspaces in $\mathbb{R}^d$.

Theorem 3 *Let $P \subset \mathbb{R}^d$ be a set of moving points with bounded description complexity s. Then, the kinetic-range space $(P, \mathcal{H}_d)$ has VC-dimension bounded by $O(d \log d)$.*

To prove Theorem 3, we need the following known definition and lemma (see, e.g., [6]). The *primal shatter function* of a hypergraph $H = (V, \mathcal{E})$ denoted by π_H is a function:

$$\pi_H : \{1, \ldots, |V|\} \to \mathbb{N}$$

defined by $\pi_H(i) = max_{V' \subseteq V, |V'|=i} |H[V']|$, where $|H[V']|$ denotes the number of hyperedges in the sub-hypergraph $H[V']$.

Lemma 4 *Let $H = (V, \mathcal{E})$ be a hypergraph whose primal shatter function π_H satisfies $\pi_H(m) = O(m^c)$ for some constant $c \geq 2$. Then the VC-dimension of H is $O(c \log c)$.*

Proof. [Proof of Theorem 3] By Lemma 4 it suffices to bound the primal shatter function $\pi_{\mathcal{H}_d}(m)$ by a polynomial of constant degree. It is a well known fact that the number of combinatorially distinct half-spaces determined by n (static) points in $\mathbb{R}^d$ is $O(n^d)$. This can be easily seen by charging hyperplanes to d-tuples of points (using rotations and translations) and observing that each tuple can be charged at most a constant (depending on the dimension d) number of times. Thus, at any given time, the number of hyperedges is bounded by $O(n^d)$. Next, note that as t varies, a combinatorial change in the hypergraph $(P(t), \mathcal{R})$ can occur only when $d+1$ points $p_1(t), \ldots, p_{d+1}(t)$ become affinely dependent. Indeed, a hyperedge is defined by a hyperplane that contains d points of $P(t)$, and that hyperedge changes when an additional point of $P(t)$ crosses the hyperplane (and thus $d+1$ points become affinely dependent). This happens if and only if the following determinant condition holds:

$$\begin{vmatrix} x^1_1(t) & x^1_2(t) & \cdots & x^1_d(t) & 1 \\ x^2_1(t) & x^2_2(t) & \cdots & x^2_d(t) & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ x^{d+1}_1(t) & x^{d+1}_2(t) & \cdots & x^{d+1}_d(t) & 1 \end{vmatrix} = 0 \qquad (1)$$

where $x^j_i(t)$ denotes the i'th coordinate of $p_j(t)$. The left side of the equation is a univariate polynomial of degree at most ds. By our general position assumption this polynomial is not zero for some instant of time. In particular, it cannot be identically zero and thus it can have at most ds solutions.

That is, a tuple of $d+1$ points of $P(t)$ generates at most ds events. Hence, the total number of such events is bounded by $O(\binom{n}{d+1}) = O(n^{d+1})$. Between any two events we have $O(n^d)$ distinct halfspaces, thus we can have $O(n^{2d+1})$ distinct ranges distributed along all instants of time.

Since each hyperedge is defined by the points on its boundary, this property is hereditary. That is, for any subset $P' \subseteq P$ the hypergraph $H[P']$ has at most $O(|P'|^{2d+1})$ hyperedges. Thus, the shatter function satisfies $\pi_H(m) = O(m^{2d+1})$. Then by Lemma 4, $(P, \mathcal{H}_d)$ has bounded constant VC-dimension, where the constant depends only on d and s. ☐

Theorem 3 can be further generalized to arbitrary ranges with so-called bounded description complexity as defined below:

Theorem 5 *Let $\mathcal{R}$ be a collection of semi-algebraic subsets of $\mathbb{R}^d$, each of which can be expressed as a Boolean combination of a constant number of polynomial equations and inequalities of maximum degree c (for some constant c). Let P be a set of moving points in $\mathbb{R}^d$ with bounded description complexity. Then the kinetic range-space $(P, \mathcal{R})$ has bounded* VC-*dimension.*

Proof. The proof combines Lemma 4 with Theorem 3 and the so-called Veronese lifting map from Algebraic Geometry. Proof is similar to the proof for the static case. See, e.g., [6]. ☐

4 Applications

Balanced Voronoi cells for moving points

Given a set P of moving points or *clients*, we are interested in locating k *facilities* so that at each instant of time no facility serves too many clients (assuming that each client goes to its nearest facility).

Proof. [Sketch of the proof of Theorem 1] First we show that the kinetic hypergraph $H = (P, \mathcal{W})$ induced by all bounded cones (i.e. the intersection of an infinite cone with a ball centered at the apex of the cone) has constant VC-dimension, and construct an ε-net N for H with $\varepsilon = \frac{1}{k}$ (this net is certified to exist thanks to Theorems 2 and 3). We claim that N satisfies the desired property.

Let C_d be the minimum number of sixty degrees caps that are needed to cover the unit sphere $\mathcal{S}^{d-1}$ (such a constant always exists and depends only on d). Assume to the contrary that the Voronoi cell of $q(t)$ contains a subset $P'(t) \subset P(t)$ of more than $C_d n/k$ points. By the pigeonhole principle, at least $n/k + 1$ of the points of $P'(t)$ lie in a 60-degrees infinite cone W and has $q(t)$ as its apex. Consider the (inclusionwise) smallest bounded cone $W' \in \mathcal{W}$ such that $W' \subset W$ and both cones contain the same points of $P(t)$. Translate W' infinitesimally so that $q(t)$ is not in the cone, but no other point enters or leaves the cone. This cone has more than n/k points of $P'(t)$, so it must contain another point $q'(t)$ of $N(t)$. By the triangle inequality, there must be a point in that cone whose closest point is $q'(t)$, contradicting the fact that it belongs to Voronoi cell of $q(t)$. ☐

Corollary 6 *Let $N \subset P$, $|N| = O(k \log k)$, as in Theorem 1. Then, for any finite point set $S \subset \mathbb{R}^d$, and for any $t \geq 0$, the cell of any $q \in S$ in the Voronoi diagram Vor$(N(t) \cup S)$ contains at most $O(n/k)$ points of $P(t)$.*

Remark We note that the bound of $O(k \log k)$ in Theorem 1 is near optimal already for $d = 1$ and points moving linearly. This follows from a recent lower-bound construction of Alon [1] for ε-nets for static hypergraphs consisting of points with respect to strips in the plane. See more details in [2].

Low interference for moving transmitters

In the following we define the concept of (receiver-based) *interference* of a set of ad-hoc sensors [9]. Let $P = \{p_1, \ldots, p_n\}$ be a set of n points in $\mathbb{R}^d$ and let $r_1, \ldots, r_n$ be n non-negative reals representing the transmission radii assigned to the points $p_1, \ldots, p_n$, respectively. Let $G = (P, E)$ be the graph associated with this power assignment, where $E = \{\{p, q\} : d(p, q) \leq min\{r_p, r_q\}\}$. Let $D = \{d_1, \ldots, d_n\}$ denote the family of balls where d_i is the ball centered at p_i and having radius r_i.

Let $I(D)$ denote the maximum depth of the arrangement of the balls in D. We call $I(D)$ *the interference* of D. Note that both G and $I(D)$ are determined by P and $r_1, \ldots, r_n$. Given a set P of points in $\mathbb{R}^d$, the *interference minimization problem* asks for the power assignment with smallest possible interference among the assignments whose underlying graph is connected.

Empirically, it has been observed that networks with high interference have large probability of messages colliding [9]. Thus, a significant amount of research has focused in the creation of connected networks with low interference. Among other results, we highlight the one of Halldórsson and Tokuyama [3] that any pointset P of n points in $\mathbb{R}^d$ has a radius assignment whose associated interference is $O(\sqrt{n})$.

Here, we turn our attention to the kinetic version of the interference problem in arbitrary but fixed dimension. We wish to maintain a connected graph of a set of moving points that always has low interference. Unless the distances between sensors remain constant, no static radii assignment can work for a long period of time since points will eventually be far from each other. Instead, we describe the network in a combinatorial way. That is, we look for a function $f : P \times [0, \infty) \to P$ that determines, for each sensor of P and instant of time, its furthest away sensor that must be reached. Then, at time t the communication radius of a sensor $p \in P$ is simply set equal to the distance between p and $f(p,t)$. Ideally, we would like to construct a network that not only has small interference for all instants of time, but also limits the number of combinatorial changes in the graph along time.

Theorem 7 *Let P be a set of n moving points in $\mathbb{R}^d$ with bounded description complexity s. Then there is a power assignment such that at any given time t we have $I(P(t)) = O(\sqrt{n \log n})$. Moreover, the total number of combinatorial changes in the network is at most $O^*(n^{1.5}\sqrt{\log n})$ where the O^* notation hides a term involving the inverse Ackermann function that depends on d and s.*

Proof. [Sketch] As in the static case [3], we use Theorem 1 for $k = \sqrt{n/\log n}$ to obtain a set N of size $O(k \log k)$ with the properties guaranteed by Theorem 1. The elements of N are called *hubs*, and we map to each point of N its furthest point in P. We map non-hub points $p \in P \setminus N$ to their nearest hub. Equivalently, if we consider the Voronoi diagram with sites $N(t)$, function $f(p,t)$ will match point $p(t)$ with the site associated to the Voronoi cell that contains $p(t)$ at time t (the mapping for points of N is similar, but we would use the farthest point Voronoi diagram instead).

Each point of $P \setminus N$ has radius large enough to reach one point of N, and all points of N form a clique, thus the network is connected. Regarding interference, by Corollary 6, we know that no point $q \in \mathbb{R}^d$ can be reached by more than $O(n/k)$ points of $P \setminus N$ at any instant of time. Thus, the total interference of any point $q \in \mathbb{R}^d$ is at most $O(k \log k)$ from hubs (since $|H| = O(k \log k)$), and at most $O(n/k)$ from non-hubs. Since $k = \sqrt{n/\log n}$ we obtain the claimed bound.

The bound on the number of combinatorial changes follows from the fact that we are tracking changes in upper and lower envelopes of a family of one-dimensional functions (i.e., pairwise distances along time). For any two points $p, p' \in P$, the function $d(p(t), p'(t))$ is algebraic of bounded degree. Thus, any two such functions cross $O(s)$ times. This allows us to use the Davenport-Schinzel Theorem to bound the overall number of combinatorial changes of the network. $\square$

5 Conclusions

We showed that for many range-spaces with bounded VC-dimension, the kinetic version of such range-spaces, a more complex and rich structure, still has a bounded VC-dimension. We believe that the boundedness of the VC-dimension of the kinetic hypergraphs is of independent interest and hope that further research will reveal more applications (some of which can be found in [2]).

References

[1] N. Alon. A non-linear lower bound for planar epsilon-nets. *Discrete & Computational Geometry*, 47(2):235–244, 2012.

[2] J. D. Carufel, M. Katz, M. Korman, A. van Renssen, M. Roeloffzen, and S. Smorodinsky. On kinetic range spaces and their applications. *CoRR*, abs/1507.02130, 2015.

[3] M. Halldórsson and T. Tokuyama. Minimizing interference of a wireless ad-hoc network in a plane. *Theoretical Computer Science*, 402(1):29–42, 2008.

[4] D. Haussler and E. Welzl. Epsilon-nets and simplex range queries. *Discrete & Computational Geometry*, 2:127–151, 1987.

[5] J. Komlós, J. Pach, and G. Woeginger. Almost tight bounds for epsilon-nets. *Discrete & Computational Geometry*, 7:163–173, 1992.

[6] J. Matoušek. *Lectures on Discrete Geometry.* Springer-Verlag New York, Inc., 2002.

[7] J. Pach and P. K. Agarwal. *Combinatorial Geometry.* Wiley Interscience, 1995.

[8] V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications*, 16(2):264–280, 1971.

[9] P. von Rickenbach, R. Wattenhofer, and A. Zollinger. Algorithmic models of interference in wireless ad hoc and sensor networks. *IEEE/ACM transactions on sensor networks*, 17(1):172–185, 2009.