On the size of sets of permutations with bounded VC-dimension

Josef Cibulka and Jan Kynčl

Department of Applied Mathematics
Charles University, Prague

11 August 2010 / PP 2010

VC-dimension

Set systems

- *VC-dimension* of a family $\mathcal{C}$ of sets over $[n] = \{1, \ldots, n\}$: size of the largest subset of $[n]$ *shattered* by $\mathcal{C}$
- Sauer's lemma: $\mathrm{VCdim}(\mathcal{C}) = \mathrm{k} \Rightarrow |\mathcal{C}| \leq \mathrm{O}(\mathrm{n}^{\mathrm{k}})$

Sets of permutations

- $\mathcal{P}$ . . . set of n-permutations
- $\mathcal{P}$ has *VC-dimension* k if k is the largest number for which there is a k-tuple of elements such that restriction of permutations of $\mathcal{P}$ on these elements gives all k-permutations
- In other words: For every $(k+1)$-tuple of elements, some $(k+1)$-permutation is missing (is *avoided*).

Forbidden Permutation Questions

- All $(k+1)$-tuples of elements avoid the same permutation.

Theorem (Marcus, Tardos(2004), using result of Klazar (2000))

The number of n-permutations avoiding a fixed permutation is $2^{\Theta(n)}$.

- Was a long–standing conjecture of Stanley and Wilf.

Permutation Sets Arising in Discrete Geometry

Arrangements of pseudolines

- In how many different ways can we place a new one?
- Placement $\leftrightarrow$ permutation of the pseudolines

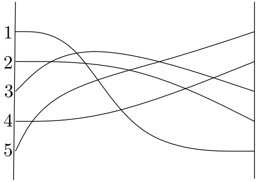

Permutation Sets Arising in Discrete Geometry

Arrangements of pseudolines

- In how many different ways can we place a new one?
- Placement $\leftrightarrow$ permutation of the pseudolines
- Fix the leftmost point $\longrightarrow$ VC-dimension is 2

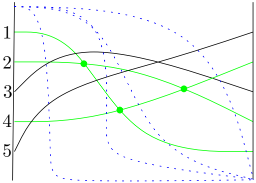

Graph drawing

- Upper bound on the number of weakly nonisomorphic complete topological graphs (Kynčl, 2010+)

Bounds on the Size of Sets of VC-dimension k

Theorem (Raz 2000)

Any set $\mathcal{P}$ of n-permutations with VC-dimension 2 has size $2^{O(n)}$.

Theorem (Our Main Result)

For a fixed $k \geq 3$, a set $\mathcal{P}$ of n-permutations with VC-dimension k has size $2^{O(n \log^\star(n))}$.
There is a set $\mathcal{P}$ of n-permutations with VC-dimension 3 and size $\alpha(n)^{\Omega(n)}$.

Matrix point of view

- permutations $\rightarrow$ permutation matrices
- each $(k+1)$-tuple of columns avoids some $(k+1)$-permutation matrix

Proof of the Upper Bound — Flattening

- Same as one step in Alon, Friedgut (1999)
- Contract layers of $n/h(n)$ consecutive rows of $P \in \mathcal{P} \rightarrow h(n) \times n$ function matrix F.
- Each of the layers $\rightarrow$ permutation matrices $P_1 \ldots P_{h(n)}$.
- F and P_i's uniquely determine P.
- Set of F's has VC-dimension at most k.
- For each F, sets of P_1's, P_2's ... have VC-dimension at most k.

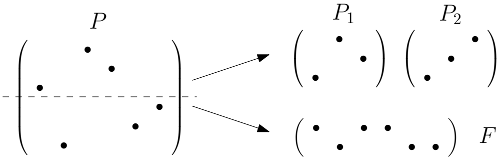

Flat Function Matrices

- $h(n) := cn/\log^6(n)$

Lemma

A set $\mathcal{F}$ of $h(n) \times n$ function matrices with VC-dimension k has size $2^{O(n)}$.

- Thus, by induction, a set of n-permutation matrices with VC-dimension k has size $2^{O(n \log^\star(n))}$.

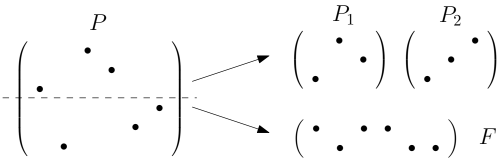

Proof of Lemma - Basic Idea

- Similar to Raz (2000)
- $M \ldots h(n) \times n$ $(0,1)$-matrix with 1's on positions where some matrix of $\mathcal{F}$ has 1
- $|M| \ldots$size of M $\ldots$number of 1-entries
- $v(M) := |M|/n$
- Decreasing $v(M)$ while not decreasing $|\mathcal{F}|$ too much $\ldots$find 1-entries not contained in many function matrices.
- End when $v(M) = O(1)$ and so $|\mathcal{F}'| \leq v(M)^n = 2^{O(n)}$
- Simple case: column with at least $v(M)\log^2(n)$ 1-entries $\ldots$ remove half of its 1-entries

Finding 1-entries to Remove

- No column has more than $v(M)\log^2(n)$ 1-entries.
- Thus $\Omega(n/log^2(n))$ columns have at least $v(M)/2$ 1-entries
- Find a large set of $(k+1)$-*splittable* columns ... $k+1$ layers; each of the columns has a 1-entry in each layer.

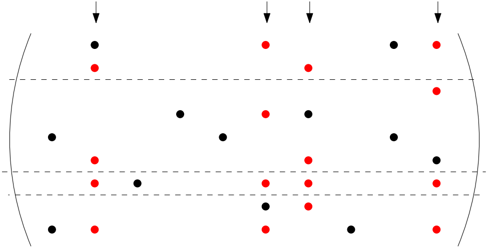

Finding 1-entries to Remove — Splittable Columns

Lemma (Nivasch 2009)

Let M be an $m \times n$ matrix with at least $v \geq v_{d,k}$ 1-entries in each column. If $n \geq c_{d,k} s m \alpha_d(m)^{k-2}$, then M contains an $(k+1)$-splittable s-tuple of columns.

- $s \geq \log^2(n)$
- $S \ldots m \times s$ matrix consisting of the splittable s-tuple of columns of M

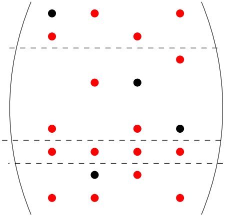

Finding 1-entries to Remove — Criss-crossing

- Take i-tuple of columns of S.
- Assign one layer to each of the columns, pairwise distinct.
- Consider function matrices that visit the assigned layer in each of the columns.
- The i-tuple of columns is *criss-crossed* if, for each assignment of one layer to each column, the number of function matrices is at least $|\mathcal{F}|/n^{2i}$.
- No criss-crossed $(k+1)$-tuple of columns - all $(k+1)$-permutation matrices would appear.

Finding 1-entries to Remove

- Criss-crossed i-tuple of columns, but no $(i+1)$-tuple.
- For each of the remaining columns, take the assignment due to which it cannot be added to the i-tuple.
- Constant number $\left(\binom{k+1}{i+1}(i+1)!\right)$ of different assignments.
- Take the most frequent assignment, fix its 1-entries (i. e., remove all the other 1's) in the first i columns and remove the ones in the assigned layer of all the possible last columns.
- $\Omega(\log^2(n))$ removed 1's; $\quad |\mathcal{F}| \to |\mathcal{F}|/(2n^{2i})$

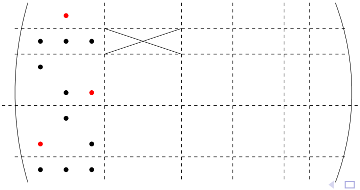

Extremal Problems on Forbidden Matrices

- $(0,1)$-matrices
- Matrix A *contains* $l \times k$ matrix B if the 1-entries of B appear in the intersection of some k-tuple of columns and some l-tuple of rows of A.

$$B = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix} \qquad A = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}$$

- Otherwise A *avoids* B
- $f(n; B)$ ... maximum number of 1-entries in an $n \times n$ matrix A avoiding B
- B is *forbidden*

Spectrum of Growth Rates

$f(n;B) =$

- $\Theta(n^{3/2})$ for $B = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$ (Turán-type result)
- $\Theta(n\log(n))$ for $B = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}$ (Füredi, 1990)
- $\Theta(n\log(n)\log\log(n))$ for a 4×5 acyclic pattern (Pettie, 2010)
- $\Theta(n\alpha(n))$ for $B = \begin{pmatrix} 0 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 \end{pmatrix}$ (Füredi and Hajnal, 1992, from DS-sequences)
- $O(n2^{\alpha^{O(1)}(n)})$ if B is a *function matrix* ...exactly 1 1-entry in each column (from generalized DS-sequences)
- $\Theta(n)$ if B is a *permutation matrix* (Marcus and Tardos, 2004)

Different Forbidden Matrices

- Different $(k+1)$-tuples of columns can have a different forbidden matrix
- Forbidden matrices are permutation
- Reformulation: No $(k+1)$-tuple of columns contains all $(k+1)$-permutation matrices
- Is $f_k(n)$ linear?
- YES if $k \leq 2$ (Raz 2000)
- NO if $k \geq 3$

Theorem

For every $k \geq 3$:

$$\Omega(n\alpha(n)) \leq f_k(n) \leq O(n2^{\alpha^{O(1)}(n)}).$$

Different Forbidden Matrices

- Different $(k+1)$-tuples of columns can have a different forbidden matrix
- Forbidden matrices are permutation
- Reformulation: No $(k+1)$-tuple of columns contains all $(k+1)$-permutation matrices
- Is $f_k(n)$ linear?
- YES if $k \leq 2$ (Raz 2000)
- NO if $k \geq 3$

Theorem

For every $k \geq 3$:

$$\Omega(n\alpha(n)) \leq f_k(n) \leq O(n2^{\alpha^{O(1)}(n)}).$$

Different Forbidden Matrices

- Different $(k+1)$-tuples of columns can have a different forbidden matrix
- Forbidden matrices are permutation
- Reformulation: No $(k+1)$-tuple of columns contains all $(k+1)$-permutation matrices
- Is $f_k(n)$ linear?
- YES if $k \leq 2$ (Raz 2000)
- NO if $k \geq 3$

Theorem

For every $k \geq 3$:

$$\Omega(n\alpha(n)) \leq f_k(n) \leq O(n2^{\alpha^{O(1)}(n)}).$$

Superlinear Lower Bound

Lemma

If A contains $P_1 = \begin{pmatrix} & & & \bullet \\ & & \bullet & \\ & \bullet & & \\ \bullet & & & \end{pmatrix}$ *and* $P_2 = \begin{pmatrix} & \bullet & & \\ \bullet & & & \\ & & & \bullet \\ & & \bullet & \end{pmatrix}$ *on the same quadruple of columns, then it contains* $DS_3 = \begin{pmatrix} & \bullet \\ \bullet & \\ & \bullet \\ \bullet & \end{pmatrix}$

Proof.

If the second row of P_1 is higher than the second row of the other matrix, then we find $\begin{pmatrix} & & & \bullet \\ & & \bullet & \\ & & & \bullet \\ & & \bullet & \end{pmatrix}$, otherwise $\begin{pmatrix} & \bullet & & \\ \bullet & & & \\ & \bullet & & \\ \bullet & & & \end{pmatrix}$ □

Corollary

A with $\Theta(n\alpha(n))$ *1-entries avoiding* DS_3 *has no quadruple of columns with all permutation matrices.*

Superlinear Lower Bound

Lemma

If A contains $P_1 = \begin{pmatrix} & & & \bullet \\ & & \bullet & \\ & \bullet & & \\ \bullet & & & \end{pmatrix}$ *and one of* $\begin{pmatrix} & \bullet & & \\ \bullet & & & \\ & & & \bullet \\ & & \bullet & \end{pmatrix}$, $\begin{pmatrix} & \bullet & & \\ \bullet & & & \bullet \\ & & \bullet & \end{pmatrix}$ *on the same quadruple of columns, then it contains* $DS_3 = \begin{pmatrix} & \bullet \\ \bullet & \\ & \bullet \\ \bullet & \end{pmatrix}$

Proof.

If the second row of P_1 is higher than the second row of the other matrix, then we find $\begin{pmatrix} & & \bullet \\ & \bullet & \\ & & \bullet \\ & \bullet & \end{pmatrix}$, otherwise $\begin{pmatrix} & \bullet & \\ \bullet & & \\ & \bullet & \\ \bullet & & \end{pmatrix}$ □

Corollary

A with $\Theta(n\alpha(n))$ 1-entries avoiding DS_3 has no quadruple of columns with all permutation matrices.

Quasilinear Upper Bound

- If A contains $k^2 \times k$ matrix $KI_k = \begin{pmatrix} \bullet & & & \\ & \bullet & & \\ & & \ddots & \\ & & & \bullet \\ \bullet & & & \\ & \bullet & & \\ & & \ddots & \\ & & & \bullet \\ & & \vdots & \\ \bullet & & & \\ & \bullet & & \\ & & \ddots & \\ & & & \bullet \end{pmatrix}$, then the k-tuple of columns contains all permutation matrices.
- $O(n2^{\alpha^{O(1)}(n)})$ 1-entries
- On the other hand, if some k^2-tuple of columns of A contains all permutation matrices, then it contains KI_k.

Back to VC-dimension — Lower Bound Construction

- Take A with $\Theta(n\alpha(n))$ 1-entries that avoids $DS_3 = \begin{pmatrix} & & \bullet \\ & \bullet & \\ & & \bullet \\ & \bullet & \end{pmatrix}$.
- All but $cn/\alpha(n)$ rows empty.
- Half of its columns have at least $c'\alpha(n)$ 1's each.
- Contains at least $(c'\alpha(n))^{n/2}$ function matrices of size $cn/\alpha(n) \times n/2$.
- Expand each row of the function matrices with the diagonal matrix $\rightarrow$ permutation matrices.

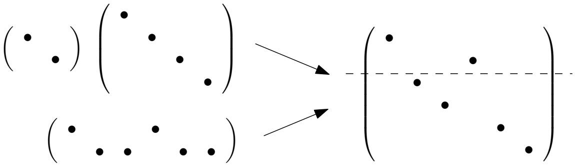

Correctness of the Construction

- Resulting permutation matrix could be obtained from at most $\binom{n/2+cn/\alpha(n)}{cn/\alpha(n)} = 2^{O(n)}$ different function matrices.
- Thus we have $\alpha(n)^{\Omega(n)}$ permutation matrices.
- VC-dimension is 3
  - Assume that the resulting permutation matrices contain on some quadruple of columns both $P_1 = \begin{pmatrix} & & & \bullet \\ & & \bullet & \\ & \bullet & & \\ \bullet & & & \end{pmatrix}$ and $P_2 = \begin{pmatrix} & \bullet & & \\ \bullet & & & \\ & & & \bullet \\ & & \bullet & \end{pmatrix}$
  - Then we find in A on this quadruple of columns: P_1 and one of $\begin{pmatrix} & \bullet & & \\ \bullet & & & \\ & & & \bullet \\ & & \bullet & \end{pmatrix}$, $\begin{pmatrix} & \bullet & & \\ \bullet & & & \bullet \\ & & \bullet & \end{pmatrix}$
  - This is impossible since A avoids $DS_3 = \begin{pmatrix} & \bullet \\ \bullet & \\ & \bullet \\ \bullet & \end{pmatrix}$.