## Expressive Efficiency and Inductive Bias of Convolutional Networks:

## Analysis \& Design via Hierarchical Tensor Decompositions

Amnon Shashua

The Hebrew University of Jerusalem

## Sources

## Deep SimNets

N. Cohen, O. Sharir and A. Shashua

Computer Vision and Pattern Recognition (CVPR) 2016
On the Expressive Power of Deep Learning: A Tensor Analysis
N. Cohen, O. Sharir and A. Shashua

Conference on Learning Theory (COLT) 2016
Convolutional Rectifier Networks as Generalized Tensor Decompositions
N. Cohen and A. Shashua

International Conference on Machine Learning (ICML) 2016
Inductive Bias of Deep Convolutional Networks through Pooling Geometry
N. Cohen and A. Shashua

International Conference on Learning Representations (ICLR) 2017
Tractable Generative Convolutional Arithmetic Circuits
O. Sharir. R. Tamari, N. Cohen and A. Shashua
arXiv preprint 2017
On the Expressive Power of Overlapping Operations of Deep Networks
O. Sharir and A. Shashua
arXiv preprint 2017
Boosting Dilated Convolutional Networks with Mixed Tensor Decompositions
N. Cohen, R. Tamari and A. Shashua
arXiv preprint 2017
Deep Learning and Quantum Entanglement: Fundamental Connections with Implications to Network
Design Y. Levine, D. Yakira, N. Cohen and A. Shashua
arXiv preprint 2017

## Students



Nadav Cohen


Or Sharir



David Yakira


Yoav Levine

# Classic vs. State of the Art Deep Learning 

## Classic


hidden layer 1 hidden layer 2

## Multilayer Perceptron (MLP)

Architectural choices:

- depth
- layer widths
- activation types


## Classic vs. State of the Art Deep Learning


hidden layer 1 hidden layer 2


## Multilayer Perceptron (MLP)

Architectural choices:

- depth
- layer widths
- activation types


## Convolutional Networks (ConvNets)

Architectural choices:

- depth
- layer widths
- activation types
- pooling types
- convolution/pooling windows
- convolution/pooling strides
- dilation factors
- connectivity
- and more...


## Classic vs. State of the Art Deep Learning


hidden layer 1 hidden layer 2

## Multilayer Perceptron (MLP)

Architectural choices:

- depth
- layer widths
- activation types


## Convolutional Networks (ConvNets)

Architectural choices:

- depth
- layer widths
- activation types
- pooling types
- convolution/pooling windows
- convolution/pooling strides

> Can the architectural choices of state of the art ConvNets be theoretically analyzed?

## Outline

(1) Expressiveness
(2) Expressiveness of Convolutional Networks - Questions
(3) Convolutional Arithmetic Circuits
(4) Efficiency of Depth (Cohen+SharirtShashua@cOLT'16, Cohen+Shashua@ICML'16)
(5) Inductive Bias of Pooling Geometry (Cohen+Shashua@ICLR'17)
(3) Efficiency of Overlapping Operations (Sharir+Shashua@arXiv'17)
(7) Efficiency of Interconnectivity (Cohen+Tamari+Shashua@arXiv'17)
(8) Inductive Bias of Layer Widths (Levine+Yakira+Cohen+Shashua@arXiv'17)

## Expressiveness

Fundamental theoretical questions:

- What kind of functions can different network architectures represent?
- Why are these functions suitable for real-world tasks?
- What is the representational benefit of depth?
- Can other architecture features deliver representational benefits?
- What does it mean to have a "representational benefit"?


## Efficiency

Expressive efficiency compares network architectures in terms of their ability to compactly represent functions

## Efficiency

Expressive efficiency compares network architectures in terms of their ability to compactly represent functions

Let:

- $\mathcal{H}_{A}$ - space of func compactly representable by network arch $A$
- $\mathcal{H}_{B}-$ network arch $B$


## Efficiency

Expressive efficiency compares network architectures in terms of their ability to compactly represent functions

Let:

- $\mathcal{H}_{A}$ - space of func compactly representable by network arch $A$
- $\mathcal{H}_{B}-$ network arch $B$
$A$ is efficient w.r.t. $B$ if $\mathcal{H}_{B}$ is a strict subset of $\mathcal{H}_{A}$

$$
\mathcal{H}_{\mathrm{A}} \bigcirc \mathcal{H}_{\mathrm{B}}
$$

## Efficiency

Expressive efficiency compares network architectures in terms of their ability to compactly represent functions

Let:

- $\mathcal{H}_{A}$ - space of func compactly representable by network arch $A$
- $\mathcal{H}_{B}-$ network arch $B$
$A$ is efficient w.r.t. $B$ if $\mathcal{H}_{B}$ is a strict subset of $\mathcal{H}_{A}$

$$
\mathcal{H}_{\mathrm{A}} \mathcal{H}_{\mathrm{B}}
$$

$A$ is completely efficient w.r.t. $B$ if $\mathcal{H}_{B}$ has zero "volume" inside $\mathcal{H}_{A}$

$$
\mathcal{H}_{\mathrm{A}} \quad \mathcal{H}_{\mathrm{B}}
$$

## Inductive Bias

Networks of reasonable size can only realize a fraction of all possible func Efficiency does not explain why this fraction is effective


## Inductive Bias

Networks of reasonable size can only realize a fraction of all possible func Efficiency does not explain why this fraction is effective


To explain the effectiveness, one must consider the inductive bias:

- Not all functions are equally useful for a given task
- Network only needs to represent useful functions


## Outline

(1) Expressiveness
(2) Expressiveness of Convolutional Networks - Questions
(3) Convolutional Arithmetic Circuits
(4) Efficiency of Depth (Cohen+Sharir+Shashua@COLT'16, Cohen+Shashua@ICML'16)
(5) Inductive Bias of Pooling Geometry (Cohen+Shashua@CLR'17)
(6) Efficiency of Overlapping Operations (Sharir+Shashua@arXiv'17)
(5) Efficiency of Interconnectivity (Cohen+Tamari+Shashua@arXiv'17)
(8) Inductive Bias of Layer Widths (Levine+Yakira+Cohen+Shashua@arXiv'17)

## Questions about Efficiency and Inductive Bias

- Depth Efficiency: deep ConvNets are (exponentially) Efficient compared to shallow networks
- Pooling scheme affects inductive bias in an Efficient manner
- ConvNets with Overlapping convolution are Efficient compared to non-overlapping ones.
- Modern connectivity schemes (split/merge/skip) are Efficient compared to standard feed-forward (LeNet, AlexNet,..).
- Layer width distribution affects inductive bias in an Efficient manner.


## Efficiency of Depth

Longstanding conjecture, proven for MLP:
deep networks are efficient w.r.t. shallow ones


Q: Can this be proven for ConvNets?
Q: Is their efficiency of depth complete?

## Inductive Bias of Convolution/Pooling Geometry

ConvNets typically employ square conv/pool windows


## Inductive Bias of Convolution/Pooling Geometry

ConvNets typically employ square conv/pool windows


Recently, dilated windows have also become popular


Q: Conv/Pooling Scheme $\leftrightarrow$ Set of functions modeled per network size $\leftrightarrow$ Suitability per task

## Efficiency of Overlapping Operations

Modern ConvNets employ both overlapping and non-overlapping conv/pool operations


Q: ConvNets with Overlapping conv are expressively Efficient w.r.t. those without (stride $=$ kernel size)?

## Efficiency of Connectivity Schemes

Nearly all state of the art ConvNets employ elaborate connectivity schemes: layers in parallel, split/merge/skip connections..


Q: Connectivity schemes are Efficient compared to standard feed-forward (LeNet, Alexnet,..)?

## Inductive Bias of Layer Widths

No clear principle for setting widths (\# of channels) of ConvNet layers


Q: What is the inductive bias of one layer's width vs. another's?
Q: Can the widths be tailored for a given task?

## Outline

(1) Expressiveness
(2) Expressiveness of Convolutional Networks - Questions
(3) Convolutional Arithmetic Circuits
4) Efficiency of Depth (Cohen+Sharir+Shashua@COLT'16, Cohen +Shashua@ICML'16)
(5) Inductive Bias of Pooling Geometry (Cohen+Shashua@ICLR'17)
6) Efficiency of Overlapping Operations (Sharir+Shashua@arXiv'17)
(7) Efficiency of Interconnectivity (Cohen+Tamari+Shashua@arXiv'17)
8) Inductive Bias of Layer Widths (Levine+Yakira+Cohen+Shashua@arXiv'17)

## Convolutional Arithmetic Circuits: Baseline Architecture



Baseline ConvAC architecture:

- Linear activation $(\sigma(z)=z)$, product pooling $\left(P\left\{c_{j}\right\}=\prod_{j} c_{j}\right)$
- $1 \times 1$ convolution windows (non-overlapping convolution: stride $=$ kernel size).

Intimate relationship to math machinery: tensor analysis, measure theory, functional analysis and graph theory.

## Coefficient Tensor



Function realized by output $y$ :

$$
h_{y}\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{N}\right)=\sum_{d_{1} \ldots d_{N}=1}^{M} \mathcal{A}_{d_{1}, \ldots, d_{N}}^{y} \prod_{i=1}^{N} f_{\theta_{d_{i}}}\left(\mathbf{x}_{i}\right)
$$

- $\mathbf{x}_{1} \ldots \mathbf{x}_{N}$ - input patches
- $f_{\theta_{1}} \ldots f_{\theta_{M}}$ - representation layer functions
- $\mathcal{A}^{y}$ - coefficient tensor ( $M^{N}$ entries, polynomials in weights $\mathbf{a}^{I, j, \gamma}$ )


## Shallow Convolutional Arithmetic Circuit $\longleftrightarrow$ CP (CANDECOMP/PARAFAC) Decomposition

Shallow network (single hidden layer, global pooling):


Coefficient tensor $\mathcal{A}^{y}$ given by classic $\mathbf{C P}$ decomposition:

$$
\begin{gathered}
\mathcal{A}^{y}=\sum_{\gamma=1}^{r_{0}} a_{\gamma}^{1,1, y} \cdot \underbrace{\mathbf{a}^{0,1, \gamma} \otimes \mathbf{a}^{0,2, \gamma} \otimes \cdots \otimes \mathbf{a}^{0, N, \gamma}}_{\text {rank-1 tensor }} \\
\left(\operatorname{rank}\left(\mathcal{A}^{y}\right) \leq r_{0}\right)
\end{gathered}
$$

## Deep Convolutional Arithmetic Circuit $\longleftrightarrow$ Hierarchical Tucker Decomposition

Deep network ( $L=\log _{2} N$ hidden layers, size-2 pooling windows):


Coefficient tensor $\mathcal{A}^{y}$ given by Hierarchical Tucker decomposition:

$$
\begin{aligned}
\phi^{1, j, \gamma} & =\sum_{\alpha=1}^{r_{0}} a_{\alpha}^{1, j, \gamma} \cdot \mathbf{a}^{0,2 j-1, \alpha} \otimes \mathbf{a}^{0,2 j, \alpha} \\
& \cdots \\
\phi^{\prime, j, \gamma} & =\sum_{\alpha=1}^{r_{l-1}} a_{\alpha}^{I, j, \gamma} \cdot \phi^{I-1,2 j-1, \alpha} \otimes \phi^{I-1,2 j, \alpha} \\
& \cdots \\
\mathcal{A}^{y} & =\sum_{\alpha=1}^{r_{L-1}} a_{\alpha}^{L, 1, y} \cdot \phi^{L-1,1, \alpha} \otimes \phi^{L-1,2, \alpha}
\end{aligned}
$$

## Universality

## Fact: <br> CP decomposition can realize any tensor $\mathcal{A}^{y}$ given $M^{N}$ terms Implies: <br> Shallow network can realize any function given $M^{N}$ hidden channels <br> Fact: <br> Hierarchical Tucker decomposition is a superset of CP decomposition if each level has matching number of terms <br> Implies: <br> Deep network can realize any function given $M^{N}$ channels in each of its hidden layers

Convolutional arithmetic circuits are universal

## Outline

(1) Expressiveness
(2) Expressiveness of Convolutional Networks - Questions
(3) Convolutional Arithmetic Circuits
4) Efficiency of Depth (Cohen+Sharir+Shashua@COLT'16, Cohen+Shashua@ICML'16)
(5) Inductive Bias of Pooling Geometry (Cohen+Shashua@ICLR'17)
6) Efficiency of Overlapping Operations (Sharir+Shashua@arXiv'17)
(7) Efficiency of Interconnectivity (Cohen+Tamari+Shashua@arXiv'17)
8) Inductive Bias of Layer Widths (Levine+Yakira+Cohen+Shashua@arXiv'17)

## Tensor Matricization

Let $\mathcal{A}$ be a tensor of order $(\operatorname{dim}) N$
Let $(I, J)$ be a partition of $[N]$, i.e. $/ \cup J=[N]:=\{1, \ldots, N\}$

## Tensor Matricization

Let $\mathcal{A}$ be a tensor of order $(\operatorname{dim}) N$
Let $(I, J)$ be a partition of $[N]$, i.e. $/ \cup J=[N]:=\{1, \ldots, N\}$
$\llbracket \mathcal{A} \rrbracket_{I, J}$ - matricization of $\mathcal{A}$ w.r.t. $(I, J)$ :

- Arrangement of $\mathcal{A}$ as matrix
- Rows correspond to modes (axes) indexed by I
- Cols



## Exponential \& Complete Efficiency of Depth

## Claim

Tensors generated by CP decomposition $w / r_{0}$ terms, when matricized under any partition $(I, J)$, have rank $r_{0}$ or less

## Exponential \& Complete Efficiency of Depth

## Claim

Tensors generated by CP decomposition $w / r_{0}$ terms, when matricized under any partition $(I, J)$, have rank $r_{0}$ or less

## Theorem

Consider the partition $I_{\text {odd }}=\{1,3, \ldots, N-1\}, J_{\text {even }}=\{2,4, \ldots, N\}$. Besides a set of measure zero, all param settings of HT decomposition give tensors that when matricized w.r.t. ( $\left.l_{\text {odd }}, J_{\text {even }}\right)$, have exponential ranks.

## Exponential \& Complete Efficiency of Depth

## Claim

Tensors generated by CP decomposition $w / r_{0}$ terms, when matricized under any partition $(I, J)$, have rank $r_{0}$ or less

## Theorem

Consider the partition $I_{\text {odd }}=\{1,3, \ldots, N-1\}, J_{\text {even }}=\{2,4, \ldots, N\}$. Besides a set of measure zero, all param settings of HT decomposition give tensors that when matricized w.r.t. ( $\left.l_{\text {odd }}, J_{\text {even }}\right)$, have exponential ranks.

Since \# of terms in CP decomposition corresponds to \# of hidden channels in shallow ConvAC:

## Corollary

Almost all func realizable by deep ConvAC cannot be replicated by shallow ConvAC with less than exponentially many hidden channels

## Exponential \& Complete Efficiency of Depth

## Claim

Tensors generated by CP decomposition $w / r_{0}$ terms, when matricized under any partition $(I, J)$, have rank $r_{0}$ or less

## Theorem

Consider the partition $I_{\text {odd }}=\{1,3, \ldots, N-1\}, J_{\text {even }}=\{2,4, \ldots, N\}$. Besides a set of measure zero, all param settings of HT decomposition give tensors that when matricized w.r.t. ( $\left.l_{\text {odd }}, J_{\text {even }}\right)$, have exponential ranks.

Since \# of terms in CP decomposition corresponds to \# of hidden channels in shallow ConvAC:

## Corollary

Almost all func realizable by deep ConvAC cannot be replicated by shallow ConvAC with less than exponentially many hidden channels

W/ConvACs efficiency of depth is exponential and complete!

## Depth Efficiency Theorem - Proof Sketch

- $\llbracket \mathcal{A} \rrbracket$ - arrangement of tensor $\mathcal{A}$ as matrix (matricization)
- Relation between tensor and Kronecker products: $\llbracket \mathcal{A} \otimes \mathcal{B} \rrbracket=\llbracket \mathcal{A} \rrbracket \odot \llbracket \mathcal{B} \rrbracket$
- $\odot$ - Kronecker product for matrices. Holds: $\operatorname{rank}(A \odot B)=\operatorname{rank}(A) \cdot \operatorname{rank}(B)$
- Implies: $\mathcal{A}=\sum_{z=1}^{Z} \lambda_{z} \mathbf{v}_{1}^{(z)} \otimes \cdots \otimes \mathbf{v}_{2 L}^{(z)} \Longrightarrow \operatorname{rank} \llbracket \mathcal{A} \rrbracket \leq Z$
- By induction over $I=1$...L, almost everywhere w.r.t. $\left\{\mathbf{a}^{1, j, \gamma}\right\}_{\ell, j, \gamma}$ :

$$
\forall j \in\left[N / 2^{\prime}\right], \gamma \in\left[r_{1}\right]: \operatorname{rank} \llbracket \phi^{\prime, j, \gamma} \rrbracket \geq\left(\min \left\{r_{0}, M\right\}\right)^{2^{\prime} / 2}
$$

- Base: "SVD has maximal rank almost everywhere"
- Step: $\operatorname{rank} \llbracket \mathcal{A} \otimes \mathcal{B} \rrbracket=\operatorname{rank}(\llbracket \mathcal{A} \rrbracket \odot \llbracket \mathcal{B} \rrbracket)=\operatorname{rank} \llbracket \mathcal{A} \rrbracket \cdot \operatorname{rank} \llbracket \mathcal{B} \rrbracket$, and "linear combination preserves rank almost everywhere"


## A Note about Measure Zero

- Depth Efficiency occurs with probability 1, i..e, besides a set of measure zero, all functions that can be implemented by a deep network of polynomial size, require exponential size in order to be realized (or even approximated) by a shallow network.


## A Note about Measure Zero

- Depth Efficiency occurs with probability 1, i..e, besides a set of measure zero, all functions that can be implemented by a deep network of polynomial size, require exponential size in order to be realized (or even approximated) by a shallow network.
- The set is a zero set of a certain polynomial (based on determinants).
- The zero set of a polynomial is closed, i.e., cannot approximate anything that is not included in the set.
- In other words, the closure of the set is also of measure zero.


## A Note about Measure Zero

- Depth Efficiency occurs with probability 1, i..e, besides a set of measure zero, all functions that can be implemented by a deep network of polynomial size, require exponential size in order to be realized (or even approximated) by a shallow network.
- The set is a zero set of a certain polynomial (based on determinants).
- The zero set of a polynomial is closed, i.e., cannot approximate anything that is not included in the set.
- In other words, the closure of the set is also of measure zero.
- For example, the set of Rational numbers is of measure zero, but the closure of the set is not of measure zero. It actually fills the entire space.


## A Note about Measure Zero

- Depth Efficiency occurs with probability 1, i..e, besides a set of measure zero, all functions that can be implemented by a deep network of polynomial size, require exponential size in order to be realized (or even approximated) by a shallow network.
- The set is a zero set of a certain polynomial (based on determinants).
- The zero set of a polynomial is closed, i.e., cannot approximate anything that is not included in the set.
- In other words, the closure of the set is also of measure zero.
- For example, the set of Rational numbers is of measure zero, but the closure of the set is not of measure zero. It actually fills the entire space.
- Therefore, the set of functions that do not satisfy depth efficiency should be viewed as a low-dimensional manifold rather than a scattered set in space.


## From Convolutional Arithmetic Circuits to Convolutional Rectifier Networks



Transform ConvACs into convolutional rectifier networks (R-ConvNets):

$$
\begin{array}{ll}
\text { linear activation } \longrightarrow \quad \operatorname{ReLU} \text { activation: } \sigma(z)=\max \{z, 0\} \\
\text { product pooling } \longrightarrow \quad \max / \text { average pooling: } P\left\{c_{j}\right\}=\max \left\{c_{j}\right\} / \operatorname{mean}\left\{c_{j}\right\}
\end{array}
$$

## Generalized Tensor Decompositions

ConvACs correspond to tensor decompositions based on tensor product $\otimes$ :

$$
(\mathcal{A} \otimes \mathcal{B})_{d_{1}, \ldots, d_{P+Q}}=\mathcal{A}_{d_{1}, \ldots, d_{P}} \cdot \mathcal{B}_{d_{P+1}, \ldots, d_{P+Q}}
$$

## Generalized Tensor Decompositions

ConvACs correspond to tensor decompositions based on tensor product $\otimes$ :

$$
(\mathcal{A} \otimes \mathcal{B})_{d_{1}, \ldots, d_{P+Q}}=\mathcal{A}_{d_{1}, \ldots, d_{P}} \cdot \mathcal{B}_{d_{P+1}, \ldots, d_{P+Q}}
$$

For an operator $g: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, the generalized tensor product $\otimes_{g}$ :

$$
\begin{gathered}
\left(\mathcal{A} \otimes_{g} \mathcal{B}\right)_{d_{1}, \ldots, d_{P+Q}}:=g\left(\mathcal{A}_{d_{1}, \ldots, d_{P}}, \mathcal{B}_{d_{P+1}, \ldots, d_{P+Q}}\right) \\
(\text { same as } \otimes \text { but with } g(\cdot) \text { instead of multiplication) }
\end{gathered}
$$

Generalized tensor decompositions are obtained by replacing $\otimes$ with $\otimes_{g}$

## Convolutional Rectifier Networks <br> $\longleftrightarrow$ Generalized Tensor Decompositions

Define the activation-pooling operator:

$$
\rho_{\sigma / P}(a, b):=P\{\sigma(a), \sigma(b)\}
$$

${ }^{1}$ Sum and average pooling are equivalent in terms of expressiveness

## Convolutional Rectifier Networks $\longleftrightarrow$ Generalized Tensor Decompositions

Define the activation-pooling operator:

$$
\rho_{\sigma / P}(a, b):=P\{\sigma(a), \sigma(b)\}
$$

- ReLU activation: $\sigma(z)=[z]_{+}:=\max \{z, 0\}$
- max/average pooling: $P\left\{c_{j}\right\}=\max \left\{c_{j}\right\} / \operatorname{mean}\left\{c_{j}\right\}$
${ }^{1}$ Sum and average pooling are equivalent in terms of expressiveness


## Convolutional Rectifier Networks $\longleftrightarrow$ Generalized Tensor Decompositions

Define the activation-pooling operator:

$$
\rho_{\sigma / P}(a, b):=P\{\sigma(a), \sigma(b)\}
$$

- ReLU activation: $\sigma(z)=[z]_{+}:=\max \{z, 0\}$
- max/average pooling: $P\left\{c_{j}\right\}=\max \left\{c_{j}\right\} / \operatorname{mean}\left\{c_{j}\right\}$

Corresponding activation-pooling operators associative and commutative:

- $\rho_{\text {ReLU/max }}(a, b):=\max \left\{[a]_{+},[b]_{+}\right\}=\max \{a, b, 0\}$
- $\rho_{\text {ReLU/sum }}(a, b):=[a]_{+}+[b]_{+}{ }^{1}$
${ }^{1}$ Sum and average pooling are equivalent in terms of expressiveness


## Exponential But Incomplete Efficiency of Depth

By analyzing matricization ranks of tensors realized by generalized CP and HT decompositions w/g(•) $\equiv \rho_{\sigma / P}(\cdot)$, we show:

## Claim

There exist func realizable by deep $R$-ConvNet requiring shallow $R$-ConvNet to be exponentially large

## Exponential But Incomplete Efficiency of Depth

By analyzing matricization ranks of tensors realized by generalized CP and HT decompositions w/g(•) $\equiv \rho_{\sigma / P}(\cdot)$, we show:

## Claim

There exist func realizable by deep $R$-ConvNet requiring shallow $R$-ConvNet to be exponentially large

On the other hand:

## Claim

A non-negligible (positive measure) set of the func realizable by deep $R$ ConvNet can be replicated by shallow R-ConvNet w/few hidden channels

## Exponential But Incomplete Efficiency of Depth

By analyzing matricization ranks of tensors realized by generalized CP and HT decompositions w/g(•) $\equiv \rho_{\sigma / P}(\cdot)$, we show:

## Claim

There exist func realizable by deep $R$-ConvNet requiring shallow $R$-ConvNet to be exponentially large

On the other hand:

## Claim

A non-negligible (positive measure) set of the func realizable by deep $R$ ConvNet can be replicated by shallow R-ConvNet w/few hidden channels

W/R-ConvNets efficiency of depth is exponential but incomplete!

## Outline

(1) Expressiveness
(2) Expressiveness of Convolutional Networks - Questions
(3) Convolutional Arithmetic Circuits
(4) Efficiency of Depth (Cohen+Sharir+Shashua@COLT'16, Cohen+Shashua@ICML'16)
(5) Inductive Bias of Pooling Geometry (Cohen+Shashua@ICLR'17)
(6) Efficiency of Overlapping Operations (Sharir+Shashua@arXiv'17)
(7) Efficiency of Interconnectivity (Cohen+Tamari+Shashua@arXiv'17)
8) Inductive Bias of Layer Widths (Levine+Yakira+Cohen+Shashua@arXiv'17)

## Separation Rank - A Measure of Input Correlations

ConvNets realize func over many local structures:

$$
f\left(\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots, \mathbf{x}_{N}\right)
$$

$\mathbf{x}_{i}$ - image patches (2D network) / sequence samples (1D network)

## Separation Rank - A Measure of Input Correlations

ConvNets realize func over many local structures:

$$
f\left(\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots, \mathbf{x}_{N}\right)
$$

$\mathbf{x}_{i}$ - image patches (2D network) / sequence samples (1D network) Important feature of $f(\cdot)$ - correlations it models between the $\mathbf{x}_{i}$ 's

## Separation Rank - A Measure of Input Correlations

ConvNets realize func over many local structures:

$$
f\left(\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots, \mathbf{x}_{N}\right)
$$

$\mathbf{x}_{i}$ - image patches (2D network) / sequence samples (1D network) Important feature of $f(\cdot)$ - correlations it models between the $\mathbf{x}_{i}$ 's

## Separation rank:

Formal measure of these correlations


Partition A


Partition B


Sep rank of $f(\cdot)$ w.r.t. input partition $(I, J)$ measures dist from separability (sep rank $\nearrow \Longrightarrow$ more correlation between $\left(\mathbf{x}_{i}\right)_{i \in I}$ and $\left.\left(\mathbf{x}_{j}\right)_{j \in J}\right)$

## Deep Networks Favor Some Correlations Over Others

## Claim

W/ConvAC sep rank w.r.t $(I, J)$ is equal to rank of $\llbracket \mathcal{A}^{y} \rrbracket_{I, J}$ - matricized w.r.t. (I, J)

## Deep Networks Favor Some Correlations Over Others

## Claim

W/ConvAC sep rank w.r.t $(I, J)$ is equal to rank of $\llbracket \mathcal{A}^{y} \rrbracket I_{, J}$ - matricized w.r.t. (I, J)

## Theorem

Maximal rank of tensor generated by HT decomposition, when matricized w.r.t. (I, J), is:

- Exponential for "interleaved" partitions
- Polynomial for "coarse" partitions


## Deep Networks Favor Some Correlations Over Others

## Claim

W/ConvAC sep rank w.r.t $(I, J)$ is equal to rank of $\llbracket \mathcal{A}^{y} \rrbracket_{I, J}$ - matricized w.r.t. (I, J)

## Theorem

Maximal rank of tensor generated by HT decomposition, when matricized w.r.t. (I, J), is:

- Exponential for "interleaved" partitions
- Polynomial for "coarse" partitions


## Corollary

Deep ConvAC can realize exponential sep ranks (correlations) for favored partitions, polynomial for others

## Pooling Geometry Controls the Preference

contiguous pooling

alternative pooling

local correlations

alternative correlations


> Pooling geometry of deep ConvAC determines which partitions are favored - controls the correlation profile (inductive bias)!

## Outline

(1) Expressiveness
(2) Expressiveness of Convolutional Networks - Questions
(3) Convolutional Arithmetic Circuits
(4) Efficiency of Depth (Cohen+Sharir+Shashua@COLT'16, Cohen+Shashua@ICML'16)
(5) Inductive Bias of Pooling Geometry (Cohen+Shashua@ICLR'17)
6) Efficiency of Overlapping Operations (Sharir+Shashua@arXiv'17)
(7) Efficiency of Interconnectivity (Cohen+Tamari+Shashua@arXiv'17)
(8) Inductive Bias of Layer Widths (Levine+Yakira+Cohen+Shashua@arXiv'17)

## Overlapping Operations

Baseline ConvAC arch has non-overlapping conv and pool windows:


## Overlapping Operations

Baseline ConvAC arch has non-overlapping conv and pool windows:



Replace those by (possibly) overlapping generalized convolution:


## Exponential Efficiency

## Theorem

Various ConvACs w/overlapping GC layers realize func requiring ConvAC $w /$ no overlaps to be exponentially large

## Examples

- Network starts with large receptive field:

- Typical scheme of alternating $B \times B$ "conv" and $2 \times 2$ "pool":



## Exponential Efficiency

## Theorem

Various ConvACs w/overlapping GC layers realize func requiring ConvAC $w /$ no overlaps to be exponentially large

## Examples

- Network starts with large receptive field:

- Typical scheme of alternating $B \times B$ "conv" and $2 \times 2$ "pool":


W/ConvACs overlaps lead to exponential efficiency!

## Outline

(1) Expressiveness
(2) Expressiveness of Convolutional Networks - Questions
(3) Convolutional Arithmetic Circuits
(4) Efficiency of Depth (Cohen+Sharir+Shashua@COLT'16, Cohen+Shashua@CML'16)
(5) Inductive Bias of Pooling Geometry (Cohen+Shashua@ICLR'17)
6) Efficiency of Overlapping Operations (Sharir+Shashua@arXiv'17)
(7) Efficiency of Interconnectivity (Cohen+Tamari+Shashua@arXiv'17)
(8) Inductive Bias of Layer Widths (Levine+Yakira+Cohen+Shashua@arXiv'17)

## Dilated Convolutional Networks

Study efficiency of interconnectivity w/dilated convolutional networks:


- 1D ConvNets (sequence data)
- Dilated (gapped) conv windows
- No pooling

Underlie Google's WaveNet \& ByteNet - state of the art for audio \& text!

## Mixing Tensor Decompositions $\longrightarrow$ Interconnectivity

With dilated ConvNets, mode (axes) tree underlying corresponding tensor decomposition determines dilation scheme


Mixed tensor decomposition blending different mode (axes) trees corresponds to interconnected networks with different dilations

## Efficiency of Interconnectivity

## Theorem

Mixed tensor decomposition generates tensors that can only be realized by individual decompositions if these grow quadratically

## Corollary

Interconnected dilated ConvNets realize func that cannot be realized by individual networks unless these are quadratically larger

## Efficiency of Interconnectivity

## Theorem

Mixed tensor decomposition generates tensors that can only be realized by individual decompositions if these grow quadratically

## Corollary

Interconnected dilated ConvNets realize func that cannot be realized by individual networks unless these are quadratically larger

## W/dilated ConvNets interconnectivity brings efficiency!

## Outline

(1) Expressiveness
(2) Expressiveness of Convolutional Networks - Questions
(3) Convolutional Arithmetic Circuits
4) Efficiency of Depth (Cohen+Sharir+Shashua@COLT'16, Cohen+Shashua@ICML'16)
(5) Inductive Bias of Pooling Geometry (Cohen+Shashua@ICLR'17)
6) Efficiency of Overlapping Operations (Sharir+Shashua@arXiv'17)
(7) Efficiency of Interconnectivity (Cohen+Tamari+Shashua@arXiv'17)
8) Inductive Bias of Layer Widths (Levine+Yakira+Cohen+Shashua@arXiv'17)

## Convolutional Arithmetic Circuits $\longleftrightarrow$ Contraction Graphs

Computation of ConvAC can be cast as a contraction graph $G$, where:

- Edge weights hold layer widths (\# of channels)
- Degree-1 nodes correspond to input patches



## Correlations $\longleftrightarrow$ Min-Cut over Layer Widths

## Theorem

For input partition $(I, J)$, the rank of $\mathcal{A}^{y}$ matricized w.r.t. $(I, J)$ is upper-bounded by the min-cut in $G$ separating the degree-1 nodes of I from those of J.

## Corollary

To model interactions between input regions represented by a specific bi-partition, it is required to set layer widths such that the min-cut is of high value. A low value represents "bottlenecks" in expressivity.


## The Quantum Many-Body Wave Function

A state of a system (interchangeably its wave function) is denoted by:

$$
|\psi\rangle \in \mathcal{H}
$$

- $\mathcal{H}$ - the relevant Hilbert Space
- $|\psi\rangle$ - vector in the Hilbert Space ('ket' notation)


## The Quantum Many-Body Wave Function

A state of a system (interchangeably its wave function) is denoted by:

$$
|\psi\rangle \in \mathcal{H}
$$

- $\mathcal{H}$ - the relevant Hilbert Space
- $|\psi\rangle$ - vector in the Hilbert Space ('ket' notation)

For a single particle with a WF in an $M$ dimensional Hilbert space $\mathcal{H}_{1}$ :

$$
|\psi\rangle=\sum_{d=1}^{M} \underbrace{v_{d}}_{\begin{array}{c}
\text { coefficients } \\
\text { vector }
\end{array}}\left|\psi_{d}\right\rangle
$$

## The Quantum Many-Body Wave Function

A state of a system (interchangeably its wave function) is denoted by:

$$
|\psi\rangle \in \mathcal{H}
$$

- $\mathcal{H}$ - the relevant Hilbert Space
- $|\psi\rangle$ - vector in the Hilbert Space ('ket' notation)

For a single particle with a WF in an $M$ dimensional Hilbert space $\mathcal{H}_{1}$ :

$$
|\psi\rangle=\sum_{d=1}^{M} \underbrace{v_{d}}_{\begin{array}{c}
\text { coefficients } \\
\text { vector }
\end{array}}\left|\psi_{d}\right\rangle
$$

The quantum many-body WF: $\quad\left(|\psi\rangle \in \mathcal{H}=\otimes_{j=1}^{N} \mathcal{H}_{j}\right)$

$$
|\psi\rangle=\sum_{d_{1} \ldots d_{N}=1}^{M} \underbrace{\mathcal{A}_{d_{1} \ldots d_{N}}}_{\begin{array}{c}
\text { coefficients } \\
\text { tensor }
\end{array}}\left|\psi_{d_{1}}\right\rangle \otimes \cdots \otimes\left|\psi_{d_{N}}\right\rangle
$$

## A Tailored Product State

Consider a single tensor product of local states $\left|\phi_{j}\right\rangle \in \mathcal{H}_{j}$ :

$$
\left|\psi^{\mathrm{ps}}\right\rangle=\left|\phi_{1}\right\rangle \otimes \cdots \otimes\left|\phi_{N}\right\rangle
$$

## A Tailored Product State

Consider a single tensor product of local states $\left|\phi_{j}\right\rangle \in \mathcal{H}_{j}$ :

$$
\left|\psi^{\mathrm{ps}}\right\rangle=\left|\phi_{1}\right\rangle \otimes \cdots \otimes\left|\phi_{N}\right\rangle
$$

By expanding each local state in the respective basis,
$\left|\phi_{j}\right\rangle=\sum_{d_{j}=1}^{M} v_{d_{j}}^{(j)}\left|\psi_{d_{j}}\right\rangle$, the product state assumes the form:

$$
\left|\psi^{\mathrm{ps}}\right\rangle=\sum_{d_{1} \ldots d_{N}=1}^{M} \mathcal{A}_{d_{1} \ldots d_{N}}^{\mathrm{ps}}\left|\psi_{d_{1}}\right\rangle \otimes \cdots \otimes\left|\psi_{d_{N}}\right\rangle
$$

$\mathcal{A}_{d_{1} \ldots d_{N}}^{\mathrm{ps}}=\prod_{j=1}^{N} v_{d_{j}}^{(j)}$ is a rank-1 tensor

## A Tailored Product State

Consider a single tensor product of local states $\left|\phi_{j}\right\rangle \in \mathcal{H}_{j}$ :

$$
\left|\psi^{\mathrm{ps}}\right\rangle=\left|\phi_{1}\right\rangle \otimes \cdots \otimes\left|\phi_{N}\right\rangle
$$

By expanding each local state in the respective basis,
$\left|\phi_{j}\right\rangle=\sum_{d_{j}=1}^{M} v_{d_{j}}^{(j)}\left|\psi_{d_{j}}\right\rangle$, the product state assumes the form:

$$
\left|\psi^{\mathrm{ps}}\right\rangle=\sum_{d_{1} \ldots d_{N}=1}^{M} \mathcal{A}_{d_{1} \ldots d_{N}}^{\mathrm{ps}}\left|\psi_{d_{1}}\right\rangle \otimes \cdots \otimes\left|\psi_{d_{N}}\right\rangle
$$

$\mathcal{A}_{d_{1} \ldots d_{N}}^{\mathrm{ps}}=\prod_{j=1}^{N} v_{d_{j}}^{(j)}$ is a rank-1 tensor
We compose each local state $\left|\phi_{j}\right\rangle$ s.t. its projection on the local basis vector equals $v_{d}^{(j)}=\left\langle\psi_{d} \mid \phi_{j}\right\rangle=f_{\theta_{d}}\left(\mathbf{x}_{j}\right)$
$\longrightarrow \mathcal{A}_{d_{1} \ldots d_{N}}^{\mathrm{ps}}=\prod_{j=1}^{N} f_{\theta_{d_{j}}}\left(\mathbf{x}_{j}\right)$

## Equivalence to a ConvAC

- Many-body WF:

$$
|\psi\rangle=\sum_{d_{1} \ldots d_{N}=1}^{M} \mathcal{A}_{d_{1} \ldots d_{N}}\left|\psi_{d_{1}}\right\rangle \otimes \cdots \otimes\left|\psi_{d_{N}}\right\rangle
$$

- Constructed product state:

$$
\left|\psi^{\mathrm{ps}}\right\rangle=\sum_{d_{1} \ldots d_{N}=1}^{M} \prod_{j=1}^{N} f_{\theta_{d_{j}}}\left(\mathbf{x}_{j}\right)\left|\psi_{d_{1}}\right\rangle \otimes \cdots \otimes\left|\psi_{d_{N}}\right\rangle
$$

## Equivalence to a ConvAC

- Many-body WF:

$$
|\psi\rangle=\sum_{d_{1} \ldots d_{N}=1}^{M} \mathcal{A}_{d_{1} \ldots d_{N}}\left|\psi_{d_{1}}\right\rangle \otimes \cdots \otimes\left|\psi_{d_{N}}\right\rangle
$$

- Constructed product state:

$$
\left|\psi^{\mathrm{ps}}\right\rangle=\sum_{d_{1} \ldots d_{N}=1}^{M} \prod_{j=1}^{N} f_{\theta_{d_{j}}}\left(\mathbf{x}_{j}\right)\left|\psi_{d_{1}}\right\rangle \otimes \cdots \otimes\left|\psi_{d_{N}}\right\rangle
$$

$$
\longrightarrow\langle\psi \mathrm{ps} \mid \psi\rangle=\sum_{d_{1} \ldots d_{N}=1}^{M} \mathcal{A}_{d_{1} \ldots d_{N}} \prod_{j=1}^{N} f_{\theta_{d_{j}}}\left(\mathbf{x}_{j}\right)=\mathbf{h}_{y}\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{N}\right)
$$

## Exactly reproducing the form of the function realized by a ConvAC!

conv weights tensor $\longleftrightarrow$ coefficients tensor
rep. functions on the inputs $\longleftrightarrow$ constructed product state

## Quantum Entanglement

Lend means of quantifying physical correlations:
"Quantum Entanglement"


## Quantum Entanglement

Lend means of quantifying physical correlations:

## "Quantum Entanglement"



Many-body WF:
$|\psi\rangle=\sum_{\alpha=1}^{\operatorname{dim}\left(\mathcal{H}^{\prime}\right)} \sum_{\beta=1}^{\operatorname{dim}\left(\mathcal{H}^{J}\right)}\left(\llbracket \mathcal{A} \rrbracket_{I, J}\right)_{\alpha, \beta}\left|\psi_{\alpha}^{\prime}\right\rangle \otimes\left|\psi_{\beta}^{J}\right\rangle$
$\llbracket \mathcal{A} \rrbracket_{I, J}-$ matricization of $\mathcal{A}$ according to $(I, J)$

## Quantum Entanglement

Lend means of quantifying physical correlations:

## "Quantum Entanglement"



Many-body WF:
$|\psi\rangle=\sum_{\alpha=1}^{\operatorname{dim}\left(\mathcal{H}^{\prime}\right)} \sum_{\beta=1}^{\operatorname{dim}\left(\mathcal{H}^{J}\right)}\left(\llbracket \mathcal{A} \rrbracket \rrbracket_{I, J}\right)_{\alpha, \beta}\left|\psi_{\alpha}^{\prime}\right\rangle \otimes\left|\psi_{\beta}^{J}\right\rangle$
$\llbracket \mathcal{A} \rrbracket_{I, J}-$ matricization of $\mathcal{A}$ according to $(I, J)$
Change of basis (SVD) $\longrightarrow|\psi\rangle=\sum_{\alpha=1}^{r} \lambda_{\alpha}\left|\phi_{\alpha}^{\prime}\right\rangle \otimes\left|\phi_{\alpha}^{J}\right\rangle$
$\lambda_{\alpha}$ - singular values of $\llbracket \mathcal{A} \rrbracket /, J$

## Measures of Entanglement

Using the singular values of $\llbracket \mathcal{A} \rrbracket \rrbracket_{I, J}$, we can define several Measures of Entanglement between I and J.

## Measures of Entanglement

Using the singular values of $\llbracket \mathcal{A} \rrbracket \nmid, J$, we can define several
Measures of Entanglement between I and J.
Examples:

- Entanglement Entropy: the entropy of singular values

$$
-\sum_{\alpha}\left|\lambda_{\alpha}\right|^{2} \ln \left|\lambda_{\alpha}\right|^{2}
$$

- Geometric Measure: the $L^{2}$ distance of $|\psi\rangle$ from the set of separable states $\min _{\left|\psi^{\mathrm{sp}(I, J)}\right\rangle}\left|\left\langle\psi^{\mathrm{sp}(I, J)} \mid \psi\right\rangle\right|^{2}$ (shown to be related to separation rank)
- Schmidt Number: the number of non-zero singular values $\operatorname{rank}(\llbracket \mathcal{A} \rrbracket,, J)$


## Measures of Entanglement

Using the singular values of $\llbracket \mathcal{A} \rrbracket I, J$, we can define several
Measures of Entanglement between / and J.
Examples:

- Entanglement Entropy: the entropy of singular values

$$
-\sum_{\alpha}\left|\lambda_{\alpha}\right|^{2} \ln \left|\lambda_{\alpha}\right|^{2}
$$

- Geometric Measure: the $L^{2}$ distance of $|\psi\rangle$ from the set of separable states $\min _{\left|\psi^{\mathrm{sp}(I, J)}\right\rangle}\left|\left\langle\psi^{\mathrm{sp}(I, J)} \mid \psi\right\rangle\right|^{2}$ (shown to be related to separation rank)
- Schmidt Number: the number of non-zero singular values $\operatorname{rank}(\llbracket \mathcal{A} \rrbracket,, J)$

All measures of entanglement:

- minimal for a separable state
- increase as the dependance between $/$ and $J$ becomes more complicated


## Measures of Entanglement - Convolutional Network

Can now use entanglement measures to describe the correlations supported by a ConvAC:


Interleaved


Left-Right


The network should support high entanglement measures for the partitions which correspond to input correlations.

## Tensor Networks

Physicists' approach for efficient representation of many-body WFs:

## Tensor Networks (TNs)

## Tensor Networks

Physicists' approach for efficient representation of many-body WFs:

## Tensor Networks (TNs)

The basic building blocks of a TN are tensors - nodes in the network:
vector:
$\left\lvert\, \begin{aligned} & d \in[r] \\ & \mathbf{v} \in \mathbb{R}^{r}\end{aligned}\right.$
matrix:


## Tensor Networks

Physicists' approach for efficient representation of many-body WFs:

## Tensor Networks (TNs)

The basic building blocks of a TN are tensors - nodes in the network:
vector:

matrix: order $N$ tensor:


## Tensor Networks

Physicists' approach for efficient representation of many-body WFs:

## Tensor Networks (TNs)

The basic building blocks of a TN are tensors - nodes in the network:


- Internal indices are summed upon
- External indices belong to the resultant tensor



## Tensor Networks

Physicists' approach for efficient representation of many-body WFs:

## Tensor Networks (TNs)

The basic building blocks of a TN are tensors - nodes in the network:


- Internal indices are summed upon
- External indices belong to the resultant tensor



## Convolutional Arithmetic Circuits $\longleftrightarrow$ Tensor Networks

Computation of ConvAC can be cast as a Tensor Network:

- Edge weights hold layer widths (\# of channels)
- Degree-1 nodes correspond to input patches



## Correlations $\longleftrightarrow$ Min-Cut over Layer Widths

## Theorem

For input partition $(I, J)$, the rank of $\mathcal{A}^{y}$ matricized w.r.t. $(I, J)$ is upper-bounded by the min-cut in $G$ separating the degree-1 nodes of I from those of J.

## Corollary

To model interactions between input regions represented by a specific bi-partition, it is required to set layer widths such that the min-cut is of high value. A low value represents "bottlenecks" in expressivity.


## Implications of the Quantum-min-cut on Layer Width



Interleaved


Left-Right


$$
\begin{equation*}
W_{C}^{\text {left-right }}=\min \left(r_{L-1}, r_{L-2}, \ldots, r_{I}^{2^{(L-2-I)}}, \ldots, r_{0}^{N / 4}, M^{N / 2}\right), \tag{1}
\end{equation*}
$$

whereas the minimal weight of a cut w.r.t. the interleaved partition is guaranteed to be exponential in $N$ and obeys:

$$
\begin{equation*}
W_{C}^{\text {interleaved }}=\min \left(r_{0}^{N / 4}, M^{N / 2}\right) . \tag{2}
\end{equation*}
$$

## Outline

(1) Expressiveness
(2) Expressiveness of Convolutional Networks - Questions
(3) Convolutional Arithmetic Circuits
4) Efficiency of Depth (Cohen+Sharir+Shashua@COLT'16, Cohen+Shashua@CML'16)
(5) Inductive Bias of Pooling Geometry (Cohen+Shashua@ICLR'17)
6) Efficiency of Overlapping Operations (Sharir+Shashua@arXiv'17)
(7) Efficiency of Interconnectivity (Cohen+Tamari+Shashua@arXiv'17)
8) Inductive Bias of Layer Widths (Levine+Yakira+Cohen+Shashua@arXiv'17)

## Conclusion

- Expressiveness - the driving force behind deep networks
- Formal concepts for treating expressiveness:
- Efficiency - network arch realizes func requiring alternative arch to be much larger
- Inductive bias - prioritization of some func over others given prior knowledge on task at hand
- We analyzed efficiency and inductive bias of ConvNet arch features:
- depth
- pooling geometry
- overlapping operations
- interconnectivity
- layer widths
- Fundamental tool underlying all of our analyses:

ConvNets $\longleftrightarrow$ hierarchical tensor decompositions

## Thank You

