Beyond Covariance: Higher-order Tensor Descriptors and Applications in Computer Vision

P. Koniusz

National ICT Australia (NICTA)

In collaboration with: A. Cherian, F. Porikli, P. H. Gosselin, F. Yan, K, Mikolajczyk.

piotr.koniusz@nicta.com.au

July 26, 2016

The Content

- Recap of Bag-of-Words (universal baseline to study).
- Second- and higher-order occurrence pooling.
- Intuitive explanation: uncertainty in max-pooling.
- Evaluations and comparisons to fisher vector encoding.
- Region covariance descriptors.
- Embedding into RKHS+Third-order Super-symmetric Tensor descriptors.
- Details of linearization process.
- Sparse coding for TOSST descriptor.
- Evaluations of TOSST on texture recognition.
- Action recognition from 3D skeletons: Sequence and Dynamics Compatibility Kernels.
- Details of eigenvalue Power Normalisation.
- Results on 3D skeleton action classification.
- Natural Inner Product on Gaussians and Deep Architectures.

Bag-of-Words (First-order Approaches)



First-order Occurrence Pooling

- The local descriptors **x** are extracted from an image and coded by f that operates on columns.
- Pooling g aggregates visual words from the mid-level features ϕ along rows:



• Let me remind the following three steps (without Pyramid Matching):

$$\begin{split} \phi_n &= f(\mathbf{x}_n, \mathbf{D}), \ \forall n \in \mathcal{N} \quad (\text{encode}) \quad (1) \\ \hat{h}_k &= g\left(\{\phi_{kn}\}_{n \in \mathcal{N}}\right) \quad (\text{pool}) \quad (2) \\ \mathbf{h} &= \hat{\mathbf{h}} / \|\hat{\mathbf{h}}\|_2 \quad (\text{normalise}) \quad (3) \end{split}$$

Higher-order Occurrence Pooling

- Note that Fisher Vector Encoding and Vector of Locally Aggregated Tensors use the second-order statistics and Power Normalisation.
- BoW can employ the second-order statistics with $\uparrow \otimes_r$, e.g. $\uparrow \otimes_2 \phi = \phi \phi^T$:



• Formally, this can be expressed in four steps:

P. Koniu

$oldsymbol{\phi}_n = f(oldsymbol{x}_n, oldsymbol{D}), \ orall n \in \mathcal{N}$	(encode)	(4)
$oldsymbol{\psi}_n = \otimes_r oldsymbol{\phi}_n$	(co-occurrences)	(5)
$oldsymbol{\psi}_n := u_{:}\left(oldsymbol{\psi}_n ight)$	(vectorise)	(6)
$\hat{h}_{k} = g\left(\left\{\psi_{kn}\right\}_{n\in\mathcal{N}}\right)$	(pool)	(7)
$oldsymbol{h} = oldsymbol{\hat{h}} / \ oldsymbol{\hat{h}}\ _2$	(normalise)	(8)
sz (NICTA) Beyond Covariance: Higher-order	Tensor Desc July 26, 2016	5/2

Uncertainty in max-pooling



• Draw randomly descriptors from this interval and apply max-pooling.



0.8

0.6

0.4

0.2

• Position of the descriptor x = 1.5 in (a) can be uniquely retrieved from ϕ because $f^{-1}([0.5, 0.5]^T) = 1.5$.

• Position of $x_3 = 1.5$ in (b) is lost as $f^{-1}([1,1]^T) \in [1;2]$.

Uncertainty in max-pooling



- Two linear slopes (LLC) display coding values ϕ_1 and ϕ_2 for any $1 \le x \le 2$.
- Co-occurrences improve on the making effect, e.g. take co-occurrence $\phi_1\phi_2$.
- It results in a new maximum for x = 1.5 masking region u is now split in two smaller regions u_1 and u_2 .

Second-order Occurrence Pooling: results

PascalVOC07, Flower102, Opponent SIFT, Sparse Coding



(a) VOC07

(b) Flower102

- K^* length of the image signature.
- r=2 Second-order Occurrence Pooling FV/VLAT - Fisher Kernels/Vector of Locally Aggregated Tensors.
- Fusions: early descriptor level, late kernel level, bi-modal tensor.
- DoPM a first-order method.

TOSST Texture Descriptors

- Region covariance descriptors (co-occurrences).
- They typically use outer-product of low-level feature vectors $\uparrow \otimes_2 \mathbf{u} = \mathbf{u} \mathbf{u}^T$.



Low-level features u (linear). Non-linear features are better.
 We embed u into RKHS/linearise the RBF kernel by feature maps.

TOSST Texture Descriptors

• Non-linear co-occurrence descriptors.



Can we form more informative co-occurrences?
 Yes, we extend ↑⊗₂ to third-order outer product ↑⊗₃.

TOSST Texture Descriptors

 Non-linear third-order descriptors +eigenvalue Power Normalization (ePN).



- Eigenvalue Power Normalisation prevents correlated signal bursts. Imagine the largest eigenvalue repesents the count of pattern of the prevents the tree bark. Surely, the amount of brick and bark patterns should not affect the prediction. But it does!
- Higher-order models can be derived analytically.

Higher-order Occurrence Pooling: derivation

- Assume a kernel, e.g. RBF and its linearisation given by: $ker(\mathbf{u}, \bar{\mathbf{u}}) \approx \langle \phi, \bar{\phi} \rangle.$
- Assume the dot product $\langle \phi, \bar{\phi} \rangle$ on a pair of features and polynomial kernel: $\langle \phi, \bar{\phi} \rangle^r$, $r \ge 2$.
- Define a sum kernel between two sets of features $\boldsymbol{U} = \{\mathbf{u}_n\}_{n \in \mathcal{N}}$ and $\bar{\boldsymbol{U}} = \{\bar{\mathbf{u}}_n\}_{\bar{n} \in \bar{\mathcal{N}}}$ for two images/regions/sequences (anything you like):

$$\operatorname{Ker}\left(\boldsymbol{U}, \bar{\boldsymbol{U}}\right) = \frac{1}{|\mathcal{N}|} \sum_{n \in \mathcal{N}} \frac{1}{|\bar{\mathcal{N}}|} \sum_{\bar{n} \in \bar{\mathcal{N}}} \ker\left(\mathbf{u}_{n}, \bar{\mathbf{u}}_{\bar{n}}\right)^{r}$$
$$\approx \frac{1}{|\mathcal{N}|} \sum_{n \in \mathcal{N}} \frac{1}{|\bar{\mathcal{N}}|} \sum_{\bar{n} \in \bar{\mathcal{N}}} \left\langle \phi_{n}, \bar{\phi}_{\bar{n}} \right\rangle^{r}$$
$$= \frac{1}{|\mathcal{N}|} \sum_{n \in \mathcal{N}} \frac{1}{|\bar{\mathcal{N}}|} \sum_{\bar{n} \in \bar{\mathcal{N}}} \left(\sum_{k=1}^{K} \phi_{kn} \bar{\phi}_{k\bar{n}}\right)^{r}$$
(9)

Higher-order Occurrence Pooling: derivation

 The rightmost summation can be re-expressed as a dot product of two outer-products of order r on φ:

$$\left(\sum_{k=1}^{K}\phi_{kn}\bar{\phi}_{k\bar{n}}\right)^{r} = \sum_{k^{(1)}=1}^{K}\dots\sum_{k^{(r)}=1}^{K}\phi_{k^{(1)}}\bar{\phi}_{k^{(1)}}\cdot\dots\cdot\phi_{k^{(r)}}\bar{\phi}_{k^{(r)}} = \left\langle \otimes_{r}\phi_{n}, \otimes_{r}\bar{\phi}_{\bar{n}}\right\rangle_{F}$$

• Now, the problem is further simplified:

$$\begin{split} & \operatorname{Ker}\left(\boldsymbol{U}, \bar{\boldsymbol{U}}\right) \approx \operatorname{Ker}'\left(\boldsymbol{\Phi}, \bar{\boldsymbol{\Phi}}\right) = \frac{1}{|\mathcal{N}|} \sum_{n \in \mathcal{N}} \frac{1}{|\bar{\mathcal{N}}|} \sum_{\bar{n} \in \bar{\mathcal{N}}} \left\langle \otimes_{r} \phi_{n}, \otimes_{r} \bar{\phi}_{\bar{n}} \right\rangle_{F} \\ & = \left\langle \frac{1}{|\mathcal{N}|} \sum_{n \in \mathcal{N}} \otimes_{r} \phi_{n}, \frac{1}{|\bar{\mathcal{N}}|} \sum_{\bar{n} \in \bar{\mathcal{N}}} \otimes_{r} \bar{\phi}_{\bar{n}} \right\rangle_{F} = \left\langle \operatorname{Avg}_{n \in \mathcal{N}} \left(\otimes_{r} \phi_{n} \right), \operatorname{Avg}_{\bar{n} \in \bar{\mathcal{N}}} \left(\otimes_{r} \bar{\phi}_{\bar{n}} \right) \right\rangle_{F} \end{split}$$

• We introduce operator \mathcal{G} (similarity for matrices/tensors, e.g. ePN): $\mathcal{K}er^{*}\left(\Phi, \bar{\Phi}\right) = \left\langle \mathcal{G}\left(\frac{1}{|\mathcal{N}|} \sum_{n \in \mathcal{N}} \otimes_{r} \phi_{n}\right), \mathcal{G}\left(\frac{1}{|\bar{\mathcal{N}}|} \sum_{\bar{n} \in \bar{\mathcal{N}}} \otimes_{r} \bar{\phi}_{\bar{n}}\right) \right\rangle_{F}$ (11)

(10)

Sparse Coding for Third-order Tensor Descriptors (TSC)

- However, X is cubic w.r.t. size of features.
 We propose Sparse Coding for Third-order Tensor Descriptors.
- We can learn a dictionary to encode TOSST:

$$\arg\min_{\substack{\boldsymbol{\mathcal{B}}_{1},\ldots,\boldsymbol{\mathcal{B}}_{K}\\\alpha^{1},\ldots,\alpha^{N}}} \sum_{n=1}^{N} \left\| \boldsymbol{\mathcal{X}}_{n} - \sum_{k=1}^{K} \boldsymbol{\mathcal{B}}_{k} \alpha_{k}^{n} \right\|_{F}^{2} + \lambda \left\| \boldsymbol{\alpha}^{n} \right\|_{1}.$$
(12)

• However, *B* have three modes (overparametrised model), so we learn instead low-rank dictionary:

$$\arg\min_{\substack{\mathbf{B}_{1},\ldots,\mathbf{B}_{K}\\\mathbf{b}_{1},\ldots,\mathbf{b}_{K}\\\boldsymbol{\alpha}^{1},\ldots,\boldsymbol{\alpha}^{N}}} \sum_{n=1}^{N} \left\| \boldsymbol{\mathcal{X}}_{n} - \sum_{k=1}^{K} \left(\mathbf{B}_{k} \uparrow \otimes \mathbf{b}_{k} \right) \boldsymbol{\alpha}_{k}^{n} \right\|_{F}^{2} + \lambda \left\| \boldsymbol{\alpha}^{n} \right\|_{1}.$$
(13)

• Resulting sparse codes lpha are pooled and used for SVM training.

P. Koniusz (NICTA)

Sparse Coding for Third-order Tensor Descriptors (TSC)

- We use training set of TOSST descriptors $\mathcal{X}_1, ..., \mathcal{X}_N$.
- We learn low-rank dictionary atoms B₁↑⊗ b₁, ..., B_K↑⊗ b_K (outer product of matrices with vectors).
- They approximate full-rank tensor atoms $\mathcal{B}_1, ..., \mathcal{B}_K$.





- Brodatz textures; 99.9% accuracy (the state of the art); others score \sim 98.72%.
- UIUC materials recognition; 58.0% accuracy.
- PASCAL VOC07 descriptor compression: 61.2% mAP (25K signature) vs. 61.3% mAP (176K signature).

Action Recognition from 3D Skeletons

Sequence Compatibility Kernel



- Components $\phi(x)$, $\phi(y)$, $\phi(z)$, $\phi(\frac{t}{T})$ denoted as (\circ , \Box , \forall , +).
- \mathcal{V} captures all triplets: ($\circ \Box \lor$), ($\circ \Box +$), ($\circ \triangledown +$), ($\Box \triangledown +$).
- ePN evens out counts of these co-occurrences.
- Tensors ${m {\cal X}}$ are the samples for training SVM.

Action Recognition from 3D Skeletons

Dynamics Compatibility Kernel



- Enumerate all unique joint displacement vectors x_{it}−x_{jt',i≤j,t≤t'}.
- Embed displacements into RKHS and linearise to obtain $\phi(\mathbf{x}_{it} \mathbf{x}_{jt'})$.
- Embed start-/end-times into RKHS, linearise to obtain $\phi(\frac{t}{T})$, $\phi(\frac{t'}{T})$.
- Take outer products $\phi(\mathbf{x}_{it} \mathbf{x}_{jt'})\phi(\frac{t}{T}) \uparrow \otimes \phi(\frac{t'}{T})$, aggregate+ePN.

P. Koniusz (NICTA)

Beyond Covariance: Higher-order Tensor Desc

• Four simple steps in MATLAB:

$$(\boldsymbol{\mathcal{E}}; \boldsymbol{A}_1, ..., \boldsymbol{A}_r) = \mathsf{HOSVD}(\boldsymbol{\mathcal{V}})$$
(14)

$$\hat{\boldsymbol{\mathcal{E}}} = \mathsf{Sgn}(\boldsymbol{\mathcal{E}}) |\boldsymbol{\mathcal{E}}|^{\gamma}$$
 (15)

$$\hat{\boldsymbol{\mathcal{V}}} = ((\hat{\boldsymbol{\mathcal{E}}} \otimes_1 \boldsymbol{A}_1) \dots) \otimes_r \boldsymbol{A}_r$$
(16)

$$\mathcal{X} = \mathcal{G}(\mathcal{V}) = \operatorname{Sgn}(\hat{\mathcal{V}}) |\hat{\mathcal{V}}|^{\gamma^*}$$
 (17)

- Perform Higher Order SVD (equivalent of SVD for more than 2 modes).
- Obtain the core tensor ${\cal E}$ (equivalent of singular values).
- Power-normalise this spectrum (values \mathcal{E} can be negative).
- Assemble back tensor, if needed, perform additionally standard PN.

Action Recognition from 3D Skeletons: results

• Eigenvalue Power Normalisation w.r.t. γ :



Action Recognition from 3D Skeletons: results

• Florence3D-Action (state-of-the-art):

	SCK	DCK		SCK+DCK	
accuracy	92.98%	93.03%	92.77%	95.47%	
size	26565	9450	16920	43485	
Bag-of-Poses 82.00% SE(3) 90.88%					

• Using fewer less noisy key-joints may be better (SCK):



• UTKinect-Action (state-of-the-art):

	SCK	DCK	SCK+DCK	
accuracy	96.08%	97.69%	98.39%	
size	40480	16920	57400	
3D joints hist. 90.92% SE(3) 97.08%				

• MSR-Action3D:

	SCK+DCK	SE(3)
accuracy, standard protocol	92 .7%	89.48%
accuracy, specific classes/subjectss	96%	92.46%
size	57400	-

Natural Inner Product on Gaussians

• Gaussian kernel between $\mathbf{u} \in \mathbb{R}^{d'}$ and $\mathbf{\bar{u}} \in \mathbb{R}^{d'}$ can be simply rewritten as:

$$G_{\sigma}(\mathbf{u}-\bar{\mathbf{u}}) = e^{-\|\mathbf{u}-\bar{\mathbf{u}}\|_{2}^{2}/2\sigma^{2}} = \left(\frac{2}{\pi\sigma^{2}}\right)_{\boldsymbol{\zeta}\in\mathbb{R}^{d'}}^{\frac{d'}{2}} \int G_{\sigma/\sqrt{2}}(\mathbf{u}-\boldsymbol{\zeta}) G_{\sigma/\sqrt{2}}(\bar{\mathbf{u}}-\boldsymbol{\zeta}) \,\mathrm{d}\boldsymbol{\zeta}.$$
(18)

• Finite approximation by $\zeta_1,...,\zeta_Z$ pivots is given by:

$$\phi(\mathbf{u}) = \left[G_{\sigma/\sqrt{2}}(\mathbf{u} - \zeta_1), ..., G_{\sigma/\sqrt{2}}(\mathbf{u} - \zeta_Z) \right]^T, \quad (19)$$

and $G_{\sigma}(\mathbf{u} - \bar{\mathbf{u}}) \approx \left\langle \frac{\phi(\mathbf{u})}{\|\phi(\mathbf{u})\|_2}, \frac{\phi(\bar{\mathbf{u}})}{\|\phi(\bar{\mathbf{u}})\|_2} \right\rangle. \quad (20)$

• As few as 6 pivots yield $\leq 0.8\%$ approximation error.

Action Recognition from 3D Skeletons: results

• Florence3D-Action w.r.t. kernel radii and pivot numbers:



- Not too sensitive to parameter variations.
- More pivots provide better RBF approximation but SVM has to learn more parameters (overfitting).

Extensions to Deep Architectures

• Convolutional Kernel Networks: future talk.



• The simplest extension - instead of Fisher Vectors apply HOSVD+ePN to aggregate over multiple CNN-based sub-patches.

P. Koniusz (NICTA)

Beyond Covariance: Higher-order Tensor Desc

July 26, 2016 25 / 27

Conclusions

- It may look difficult, but it requires few easy steps only in practice:
 - A. Extract your favourite features.
 - B. Embed them into RKHS/linearise.
 - C. Form outer products of desired order.
 - D. Aggregate and apply ePN.
 - E. Train SVM (or any favourite classifier).
- Interested in any related ideas? Talk to me to see if we can collaborate :-) Ideas take know-how, time and people to develop them.
- References:



P. Koniusz, F. Yan, P. H. Gosselin, K. Mikolajczyk.

Higher-order Occurrence Pooling for Bags-of-Words: Visual Concept Detection, TPAMI, 2016.



P. Koniusz, A. Cherian.

Sparse Coding for Third-order Super-symmetric Tensor Descriptors with Application to Texture Recognition, CVPR, 2016.



P. Koniusz, A. Cherian, F. Porikli.

Tensor Representations via Kernel Linearlization for Action Recognition from 3D Skeletons, ECCV, 2016.

J. Mairal, P. Koniusz, Z. Harchaoui, C. Schmid.

Convolutional Kernel Networks, NIPS, 2014.

Thank You

