END-TO-END DEEP KERNEL MAP DESIGN FOR IMAGE ANNOTATION

Mingyuan Jiu^{*,†}

Hichem Sahbi[‡]

^{*} School of Information Engineering, Zhengzhou University, Zhengzhou, China

[†] Zhengzhou University Research Institute of Industrial Technology Co., Ltd., Zhengzhou, China

[‡] CNRS LIP6 Lab, UPMC, Sorbonne University, Paris, France

ABSTRACT

Deep kernel map networks have shown excellent performances in various classification problems including image annotation. Their general recipe consists in aggregating several layers of singular value decompositions (SVDs) – that map data from input spaces into high dimensional spaces – while preserving the similarity of the underlying kernels. However, the potential of these deep map networks has not been fully explored as the original setting of these networks focuses mainly on the approximation quality of their kernels and ignores their discrimination power.

In this paper, we introduce a novel “end-to-end” design for deep kernel map learning that balances the approximation quality of kernels and their discrimination power. Our method proceeds in two steps; first, layerwise SVD is applied in order to build initial deep kernel map approximations and then an “end-to-end” supervised learning is employed to further enhance their discrimination power while maintaining their efficiency. Extensive experiments, conducted on the challenging ImageCLEF annotation benchmark, show the high efficiency and the out-performance of this two-step process with respect to different related methods.

Index Terms— Deep kernel networks, deep map networks, supervised end-to-end learning, image annotation.

1. INTRODUCTION

Kernel learning has been an active research field in the last two decades with many applications ranging from support vector classification [1] to regression, through dimensionality reduction [2]. More recently, an extension of kernels known as deep kernel networks (DKNs) has attracted a particular attention [3–8] following the resurgence of neural networks [9, 10]. These deep kernels – defined as nonlinear and recursive combinations of standard positive semi-definite

(p.s.d) kernels [8] – are proven to be successful in describing and comparing highly nonlinear data. However, the downside of DKNs resides in their limited efficiency; indeed, the computational complexity of these networks scales linearly w.r.t. their depth and quadratically w.r.t. the size of training data, and this makes their evaluation clearly intractable for large (and even mid) scale problems.

According to the kernel theory (see for instance [11]), any p.s.d kernel admits an implicit or explicit map in a high (possibly infinite) dimensional Hilbert space. Considering this property, an interesting alternative to kernels, is to design their associated maps explicitly. In the related work, kernel map approximation techniques include (i) the Nyström expansion [12] which obtains low-rank kernel maps using uniformly sampled data without replacement, (ii) the random Fourier sampling (proposed by Rahimi and Recht [13] for gaussian kernels and extended to group-invariant kernels in [14]) and (iii) the explicit kernel map design for additive homogeneous kernels [15]. Other solutions rely rather on the strength of deep learning (such as [16]) and aim at designing explicit kernel representations using deep map networks (DMNs); the latter make it possible to build maps whose inner products approach the original DKNs while being deep and highly efficient. A more recent extension, in [17], further enhances the approximation quality of DMNs using unsupervised learning. Nevertheless, the discrimination power of DMNs has not been fully investigated in these works; indeed, these networks are biased towards the underlying DKNs and their design *ignores* the targeted classification tasks.

Following the tremendous success of deep representation learning [18] particularly in image classification (see for instance [19–23]), we introduce in this paper an “end-to-end” framework that further enhances the discrimination power of the learned DMN representations. Our DMN inputs (associated to standard kernels such as polynomial kernel) are forwarded to output layers through intermediate projection and activation operations. In contrast to the aforementioned kernel approximation techniques which are mainly unsupervised, our formulation is both “deep and supervised” and proceeds by a greedy (layerwise) SVD decomposition followed by an “end-to-end” supervised learning that improves the discrimination power of the DMNs while maintaining their approx-

This work was supported by the National Natural Science Foundation of China (No. 61806180), Key Research Projects of Henan Higher Education Institutions in China (No. 19A520037) and Zhengzhou Municipal Science Technology Innovation Project (No. 2019CXZX0037), and also in part by a grant from the research agency ANR (Agence Nationale de la Recherche) of France under the MLVIS project (ANR-11-BS02-0017).

Fig. 1. This figure shows a three-layer deep map network (DMN); each rectangular block corresponds to the explicit map of an – input, intermediate or output – layer in the underlying deep kernel network (DKN). The last fully connected layers are used for classification (**better to zoom the PDF version**).

imation quality w.r.t. the underlying DKNs. Note that our work is related to the method in [7, 24] which approximates the maps of gaussian kernels using convolutional neural networks and the recent work of deep neural mapping proposed by Li and Ting [25] which combines SVMs and kernel maps. Our method is also related to [26] which proposes deep radial kernel networks to approximate gaussian kernel SVMs and also the method in [27] that considers a deep hybrid neural-network based on random Fourier features combining neural networks and kernel machines. However, these related methods as well as those cited earlier, are either shallow or restricted to a sub-class of kernels (including gaussians) while *our solution in this paper is deep and targeted to a more general class of deep nonlinear kernels*.

2. DEEP KERNEL MAP NETWORKS

A deep kernel is defined as a multilayered network whose units – denoted as $\{\kappa_p^{(l)}\}_{l,p}$ – correspond to (input or intermediate) kernels with $\kappa_p^{(l)}(\cdot, \cdot) = g(\sum_{q=1}^{n_{l-1}} \mathbf{w}_{p,q}^{(l-1)} \kappa_q^{(l-1)}(\cdot, \cdot))$; here g is a nonlinear activation, p refers to the p -th unit of the l -th layer, with $l \in \{1, \dots, L\}$, $p \in \{1, \dots, n_l\}$ and $q \in \{1, \dots, n_{l-1}\}$. Considering p.s.d input kernels $\{\kappa_p^{(1)}\}_p$ (such as linear, polynomial and gaussian), provided that the weights $\{\mathbf{w}_{p,q}^{(l-1)}\}_{l,p,q}$ are positive and resulting from the closure of the p.s.d w.r.t. sum and product, any intermediate kernel is at least conditionally p.s.d for a particular class of activation functions (including hyperbolic tangent and exponential). Extra details about the setting of these weights together with the activation functions that guarantee the conditional positive semi-definiteness can be found in [8].

Considering the above definition, a conditionally p.s.d kernel $\kappa_p^{(l)}$ admits an explicit (either exact or approximate) map $\hat{\phi}_p^{(l)}(\cdot)$ s.t. $\kappa_p^{(l)}(\mathbf{x}, \mathbf{x}') \simeq \langle \hat{\phi}_p^{(l)}(\mathbf{x}), \hat{\phi}_p^{(l)}(\mathbf{x}') \rangle$; here $\hat{\phi}_p^{(l)}(\mathbf{x})$ is a mapping that takes $\mathbf{x}$ from an input space into a high dimensional Hilbert space. Let $\mathcal{S} = \{\mathbf{x}_i\}_{i=1}^N$ denote N samples taken from our training set; assuming the kernel maps $\{\hat{\phi}_p^{(1)}(\cdot)\}_p$ of the first layer known (either exactly or tightly

approximated [28]), we recursively define the explicit map $\hat{\phi}_p^{(l)}(\cdot)$ of the p -th unit and the l -th layer as

$$\hat{\phi}_p^{(l)}(\mathbf{x})^\top = \left(g(\langle \hat{\phi}_p^{l,c}(\mathbf{x}), \hat{\phi}_p^{l,c}(\mathbf{x}_1) \rangle) \dots g(\langle \hat{\phi}_p^{l,c}(\mathbf{x}), \hat{\phi}_p^{l,c}(\mathbf{x}_N) \rangle) \right) \mathbf{U}_p^{(l)}, \quad (1)$$

where $^\top$ denotes the matrix transpose operator and $g(\cdot)$ stands for an activation function taken, in practice, as hyperbolic for intermediate layers and exponential for the final layer, and

$$\hat{\phi}_p^{l,c}(\mathbf{x}) = \left(\sqrt{\mathbf{w}_{p,1}^{(l-1)}} \hat{\phi}_1^{(l-1)}(\mathbf{x})^\top \dots \sqrt{\mathbf{w}_{p,n_{l-1}}^{(l-1)}} \hat{\phi}_{n_{l-1}}^{(l-1)}(\mathbf{x})^\top \right)^\top. \quad (2)$$

In Eq. (1), $\mathbf{U}_p^{(l)} = \mathbf{V}\mathbf{\Lambda}^{-1/2}$ is a transformation matrix obtained by solving the following eigenproblem

$$\mathbf{K}_p^l \mathbf{V} = \mathbf{V} \mathbf{\Lambda}, \quad (3)$$

here $\mathbf{K}_p^l$ is the kernel matrix of $\kappa_p^{(l)}$ on $\mathcal{S}$. Fig. 1 shows an example of a DMN architecture obtained using Eqs. (1) – (3). Considering the maps of these equations one may introduce the following proposition.

Proposition 1 *Let $\mathcal{S} = \{\mathbf{x}_i\}_{i=1}^N$ be a subset of N samples and let $\mathbf{K}_p^l$ be a gram-matrix whose entries are defined on $\mathcal{S}$. Let $\mathbf{U}_p^{(l)} = \mathbf{V}\mathbf{\Lambda}^{-1/2}$ with $\mathbf{V}, \mathbf{\Lambda}$ being respectively the matrices of eigenvectors and eigenvalues obtained by solving Eq. (3). Considering $\|\cdot\|_2$ as the ℓ_2 (matrix) norm and $\hat{\mathbf{K}}_p^l$ as the gram-matrix associated to $\{\langle \hat{\phi}_p^{(l)}(\mathbf{x}), \hat{\phi}_p^{(l)}(\mathbf{x}') \rangle\}_{\mathbf{x}, \mathbf{x}' \in \mathcal{S}}$ with Eq. (1) and Eq. (2), then the following property is satisfied*

$$\|\hat{\mathbf{K}}_p^l - \mathbf{K}_p^l\|_2 = 0. \quad (4)$$

Details of the proof are omitted and can be found online¹. More importantly, this proposition shows that the inner products obtained using kernel maps in Eqs. (1) – (3) are equal to the original kernel values if data belong to $\mathcal{S}$; otherwise one may at least show that $\|\hat{\mathbf{K}}_p^l - \mathbf{K}_p^l\|_2 \rightsquigarrow 0$ when N is sufficiently large.

3. END-TO-END LEARNING

As designed above, DMNs significantly reduce the computational complexity of DKNs especially on large scale datasets. However, with this particular unsupervised setting, DMNs fit their DKN counterparts but ignore the underlying classification tasks. Besides, kernel map design as shown in Eqs. (1) and (2) considers only $\{\hat{\phi}_p^{l,c}(\mathbf{x}_i)\}_{i,l,p}$ and $\{\mathbf{U}_p^{(l)}\}_{l,p}$ as training parameters of DMNs while $\{\mathbf{w}_{p,q}^{(l)}\}_{l,p,q}$ are fixed and taken from DKNs. As a result, the potential of DMNs is not fully explored for supervised classification. In what follows, we propose an end-to-end supervised framework that further enhances the discrimination power of DMNs.

¹<https://www.dropbox.com/s/pdpixj73xwxjevz/suppIcIp2020.pdf?dl=0>

Algorithm 1: End-to-end supervised DMN learning

Input: Parameter setting; set the learning rate $\eta > 0$.
Initialization: $\{\mathbf{w}_{p,q}^{(l)}\}, \{\hat{\phi}_p^{l,c}(\mathbf{x}_i)\}_i, \{\mathbf{U}_{p,q}^{(l)}\}, \{\omega_k\}_k$,
 $l \in \{1, \dots, L\}, i \in \{1, \dots, N\}, k \in \{1, \dots, K\}$.

- 1 **repeat**
- 2 Optimize $\{\omega_k\}_k$ by LIBLINEAR toolbox;
- 3 Compute the gradients $\frac{\partial E}{\partial \hat{\phi}_1^L(\mathbf{x}_i)}$ by Eq. (7);
- 4 Compute the gradients $\Delta \mathbf{U}_p^{(l)}, \Delta \hat{\phi}_p^{l,c}(\mathbf{x}_i)$ and $\Delta \mathbf{w}_{p,q}^{(l)}, \forall l \in \{L-1, \dots, 1\}$;
- 5 Update these parameters by gradient descent:
 $\mathbf{U}_p^{(l)} \leftarrow \mathbf{U}_p^{(l)} - \eta \Delta \mathbf{U}_p^{(l)}$;
- 6 $\hat{\phi}_p^{l,c}(\mathbf{x}_i) \leftarrow \hat{\phi}_p^{l,c}(\mathbf{x}_i) - \eta \Delta \hat{\phi}_p^{l,c}(\mathbf{x}_i)$;
- 7 $\mathbf{w}_{p,q}^{(l)} \leftarrow \mathbf{w}_{p,q}^{(l)} - \eta \Delta \mathbf{w}_{p,q}^{(l)}$;
- 8 **until** Convergence;

Let $\mathcal{T} = \{(\mathbf{x}_i, \mathbf{y}_i^k)\}_i$ be a training set whose samples belong to K classes; here $\mathbf{x}_i$ is a training data and $\mathbf{y}_i^k$ its class membership, with $\mathbf{y}_i^k = +1$ iff $\mathbf{x}_i$ belongs to class k , otherwise $\mathbf{y}_i^k = -1$. For classification, a fully connected layer with K units is stacked on top of the DMN whose parameters $\{\omega_k\}_{k=1}^K$ correspond to the normals of K hyperplane classifiers. The decision function of each classifier f_k is given by

$$f_k(\mathbf{x}_i) = \omega_k^\top \hat{\phi}_1^L(\mathbf{x}_i). \quad (5)$$

With this decision function, a concept k is declared as present in $\mathbf{x}_i$ iff the score $f_k(\mathbf{x}_i)$ is positive. In order to learn the DMN, we minimize a squared hinge loss criterion rather than the widely used logistic loss, due to the fact that the former has Lipschitz continuous gradients and shows good discrimination ability for classification [29]. Therefore, the loss E to minimize is written as

$$\min_{\hat{\phi}_p^{l,c}, \mathbf{U}_p^{(l)}, \mathbf{w}_{p,q}^{(l)}, \omega_k} \sum_{k=1}^K \frac{1}{2} \|\omega_k\|_2^2 + C_k \sum_i \max(0, 1 - \mathbf{y}_i^k f_k(\mathbf{x}_i))^2, \quad (6)$$

where the first term is an ℓ_2 penalization, the second one is an empirical loss on training data and C_k controls the influence of these two terms.

In order to minimize Eq. (6), we adopt an alternating optimization strategy that makes training tractable: first, we optimize the classifier weights by LIBLINEAR [30] while fixing the parameters of DMN and then we update the latter while fixing the classifier weights. In the second step, when classifier weights are fixed, the gradient of E w.r.t. the output of DMN (i.e. $\hat{\phi}_1^L(\cdot)$) is given as

$$\frac{\partial E}{\partial \hat{\phi}_1^L(\mathbf{x}_i)} = -2 \sum_{k=1}^K C_k \mathbf{y}_i^k \omega_k \max(0, 1 - \mathbf{y}_i^k f_k(\mathbf{x}_i)). \quad (7)$$

Then, we employ the chain rule [31] and we back-propagate the above gradient to the preceding layers in DMN to obtain the gradients of E w.r.t. $\{\mathbf{w}_{p,q}^{(l)}\}_{l,p,q}, \{\hat{\phi}_p^{l,c}(\mathbf{x}_i)\}_{i,l,p}$ and

$\{\mathbf{U}_p^{(l)}\}_{l,p}$. Finally, we update the DMN parameters using gradient descent. The learning procedure is repeated till convergence or when the maximum number of iterations is reached (see more details in Algorithm 1).

4. EXPERIMENTS

In this section, we evaluate the performance of the proposed algorithm on the challenging and widely used ImageCLEF annotation benchmark. The goal of image annotation (also known as multi-label classification) is to predict the presence of semantic concepts into images; note that concepts in this task are not exclusive, so each image may be annotated by one or multiple concepts and this makes image annotation a highly challenging task.

ImageCLEF dataset [10] contains more than 250k (training, dev and test) images belonging to 95 different concepts. We only use the dev set of 1000 images in our experiments, as the ground truth is released only for this subset. The dev set is split into two subsets: the first one is used for training, and the other for testing. The discrimination power for annotation task are evaluated using F-measures (harmonic means of recalls and precisions) both at the image and the concept levels (denoted as MF-S and MF-C respectively) as well as mean average precision (mAP). Higher values of these measures imply better performances.

Firstly, we study the performance of ten handcrafted visual features (provided by the ImageCLEF challenge organizers). Four input kernels (linear, polynomial, RBF and histogram intersection) are considered for each feature, so we have 40 input kernels in total. Then we learn a three-layer DKN with a hidden layer of 80 units (i.e. twice the size of the input as set in [8]), next we build an initial DMN by using the method in [16]. Finally, we update this DMN using the proposed end-to-end algorithm. The classification weights in the final layer are randomly initialized, the trade-off parameter C_k is initially set using 3-fold cross-validation on the training subset and the learning rate is empirically set to 10^{-6} to guarantee convergence. In these experiments, we observe that 500 iterations are sufficient in order to converge to a stable solution. Tab. 1 shows the comparison between the DKN, the initial DMN and the enhanced (end-to-end) DMNs on the handcrafted features.

Secondly, we conduct another set of experiments by taking into account the deep features from the pre-trained VGG model on the ImageNet database (“imagenet-vgg-m-1024”) [32], containing five convolutional layers and three fully-connected layers. We use the outputs of the second fully-connected layer in order to describe images. Similarly, we consider four input kernels on top of the deep features, and we repeat the experiments as described above. The comparative results are again shown in Tab. 1; the latter shows performances for different settings including the original DKN, its two initial DMNs (trained using SVD, with and

	GT	DMN-I	DMN-S	DDMN-I	DDMN-S
Union building	*		*	*	*
castle	*			*	*
church		*	*	*	*
cloud	*		*	*	*
cloudless		*	*	*	*
daytime	*	*	*	*	*
grass	*		*	*	*
outdoor	*	*	*	*	*
person		*	*	*	*
plant	*		*	*	*
sculpture		*	*	*	*
shadow		*	*	*	*
sign	*	*	*	*	*
sky	*	*	*	*	*
tree	*	*	*	*	*

Fig. 2. This figure shows examples of annotation results using the original and the end-to-end DMNs (resp. denoted as “DMN-I”, “DMN-S” for handcrafted features and “DDMN-I”, “DDMN-S” for deep features, “I” for initialization and “S” for supervised end-to-end learning.). “GT” stands for ground truth annotation while the stars mean the presence of a concept in a test image.

Features	Method	MF-S	MF-C	MAP
Handcrafted feat.	GMKL([33])	41.3	24.3	49.1
	2LMKL([6])	45.0	25.8	54.0
	LDMKL ([8])	47.8	30.0	58.6
	DKN ([8])	46.2	30.0	55.7
	Ini. DMN ([16])	47.7	29.4	53.2
	Sup. DMN (Proposed)	49.6	31.9	58.5
	Uns. DMN ([17])	48.0	29.8	53.3
+Deep feat.	Ft. DMN (Proposed)	50.2	32.2	59.2
	DKN ([8])	56.3	38.9	66.6
	Ini. DMN ([16])	56.7	39.7	66.4
	Sup. DMN (Proposed)	56.8	40.4	67.2
	Uns. DMN ([17])	56.4	39.3	66.5
	Ft. DMN (Proposed)	56.2	40.5	67.4

Table 1. Comparison of annotation performances (in %) of different methods, using handcrafted and deep features. In this table, “Ini.” stands for initialization, “Sup.” for supervised, “Uns.” for unsupervised and “Ft.” for fine-tuned.

without unsupervised learning), and the two underlying “end-to-end” DMN variants whose weights are initially taken from these two initial DMNs respectively.

From Tab. 1, we observe that our proposed “end-to-end” learning framework is able to further boost the discrimination power of DMNs compared to initial and unsupervised DMNs as well as the original DKNs on both handcrafted and deep features. Recall that “Ini/Uns” DMNs shown in Tab. 1 are designed only to fit the underlying DKNs without taking into account any label information. In contrast, “Sup/Ft” DMNs make it possible to retrain their parameters while also maximizing classification performances. Extra comparisons against other kernel-based methods are also shown in Tab. 1 and they validate the effectiveness of the proposed algorithm. Fig. 2 shows several annotation results in the test set for the initial DMNs and the fine-tuned ones using supervised end-to-end learning on handcrafted and deep features.

Tab. 2 shows a comparison of the average runtime in order to process (classify) any given sample using DKNs vs. DMNs. When using DKNs, this process requires evaluating and propagating M kernel values through L layers between a given sample and all training data in $\mathcal{T}$ with a complexity

Framework	Sample size $ \mathcal{T} $	Average runtime (in seconds)
DKN	500	0.305
	1000	0.826
	2000	2.822
	5000	16.188
DMN	500	0.569
	1000	0.566
	2000	0.595
	5000	0.594

Table 2. This table shows the average runtime in order to classify any given sample using DKNs vs DMNs. When classifying a sample with DKNs, all the kernel values between that sample and $\mathcal{T}$ should be evaluated prior to classification using the dual SVM form; in contrast, DMNs rely on *efficient* explicit kernel map (+ primal SVM) evaluations. These performances were obtained on a workstation with four Xeon CPUs of 3.2GHz. In all these experiments $|\mathcal{S}| = N = 1000$.

$O(M^2L|\mathcal{T}|)$; here $M = \max_l n_l$ (i.e., upper bound on the width of DKN). When using instead DMNs, the complexity of evaluating the kernel maps in Eqs. (1) and (2) is independent from $\mathcal{T}$ and equal to $O(MLN^2)$; see also Tab. 2. Hence, when $|\mathcal{T}| \leq N$ and provided that $M < N$, DKNs are more efficient while larger values of $|\mathcal{T}|$ make DMNs more and more efficient with respect to DKNs.

5. CONCLUSION

We introduced in this paper an “end-to-end” design of deep map networks that effectively approximate the underlying deep kernel networks while being highly efficient. The strength of our method resides in its ability to fit not only the original DKNs but also the targeted classification task. Our method proceeds in two steps: first, an SVD step is achieved in order to build the initial DMN architecture, followed by an “end-to-end” supervised training step that further enhances the discrimination power of our DMNs and their parameters. Experiments conducted on the challenging ImageCLEF benchmark show a clear and a consistent gain of our “end-to-end” DMN design compared to other different settings.

6. REFERENCES

- [1] B. Caputo, C. Wallraven, and M. Nilsback, "Object categorization via local kernels," in *ICPR*, 2004, vol. 2, pp. 132–135.
- [2] J. Shawe-Taylor and N. Cristianini, "Kernel methods for pattern analysis," 2004, Cambridge University Press.
- [3] Y. Cho and L. Saul, "Kernel methods for deep learning," in *NIPS*, 2009, vol. 28 (1), pp. 342–350.
- [4] K. Yu, W. Xu, and Y. Gong, "Deep learning with kernel regularization for visual recognition," in *Advances in Neural Information Processing Systems (NIPS)*, 2008, pp. 1889–1896.
- [5] M. Jiu and H. Sahbi, "Semi supervised deep kernel design for image annotation," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 1156–1160.
- [6] J. Zhuang, I. Tsang, and S. Hoi, "Two-layer multiple kernel learning," in *ICML*, 2011, pp. 909–917.
- [7] J. Mairal, P. Koniusz, Z. Harchaoui, and C. Schmid, "Convolutional kernel networks," in *NIPS*, 2014, pp. 2627–2635.
- [8] M. Jiu and H. Sahbi, "Nonlinear deep kernel learning for image annotation," *IEEE Transactions on Image Processing*, vol. 26, no. 4, pp. 1820–1832, 2017.
- [9] G. Hinton, S. Osindero, and Y.-W. Teh, "A fast learning algorithm for deep belief nets," *Neural Computation*, 2006.
- [10] M. Villegas, R. Paredes, and B. Thomee, "Overview of the imageclef 2013 scalable concept image annotation subtask," in *CLEF 2013 Evaluation Labs and Workshop*, 2013.
- [11] V. Vapnik, "Statistical learning theory," 1998, Wiley.
- [12] C.K.I. Williams and M. Seeger, "Using the nyström method to speed up kernel machines," in *NIPS*, 2001, pp. 682–688.
- [13] A. Rahimi and B. Recht, "Random features for large-scale kernel machines," in *NIPS*, 2007, vol. 20, pp. 1177–1184.
- [14] F. Li, C. Ionescu, and C. Sminchisescu, "Random fourier approximations for skewed multiplicative histogram kernels," in *DAGM conference Pattern Recognition*, 2010, pp. 262–271.
- [15] A. Vedaldi and A. Zisserman, "Efficient additive kernels via explicit feature maps," *PAMI*, pp. 480–492, 2012.
- [16] M. Jiu and H. Sahbi, "Deep kernel map networks for image annotation," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 1571–1575.
- [17] M. Jiu and H. Sahbi, "Deep representation design from deep kernel networks," *Pattern Recognition*, vol. 88, pp. 447–457, 2019.
- [18] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *TPAMI*, vol. 35, no. 8, pp. 1798–1828, 2013.
- [19] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *NIPS*, 2012, vol. 60, pp. 1097–1105.
- [20] C. Farabet, C. Couprie, L. Najman, and Y. LeCun, "Learning hierarchical features for scene labeling," *TPAMI*, vol. 35, no. 8, pp. 1915–1929, 2013.
- [21] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *CVPR*, 2015, pp. 1–9.
- [22] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016, pp. 770–778.
- [23] G. Huang, Z. Liu, L. v. d. Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *CVPR*, 2017, pp. 2261–2269.
- [24] Julien Mairal, "End-to-end kernel learning with supervised convolutional kernel networks," in *Advances in Neural Information Processing Systems* 29, pp. 1399–1407. 2016.
- [25] Y. Li and T. Zhang, "Deep neural mapping support vector machines," *Neural Networks*, vol. 93, pp. 185–194, 2017.
- [26] B. McCane and L. Szymanski, "Deep radial kernel networks: Approximating radially symmetric functions with deep networks," *CoRR:1703.03470*, 2017.
- [27] S. Mehrkanoon and J. A. K. Suykens, "Deep hybrid neural-kernel networks using random fourier features," *Neurocomputing*, vol. 298, pp. 46–54, 2018.
- [28] H. Sahbi, "Imageclef annotation with explicit context-aware kernel maps," *International Journal of Multimedia Information Retrieval*, vol. 4, pp. 113–128, 2015.
- [29] Y. Tang, "Deep learning using linear support vector machines," in *Challenges in Representation Learning W.*, 2013, pp. 1–6.
- [30] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin, "Liblinear: A library for large linear classification," *JMLR*, vol. 9, pp. 1871–1874, 2008.
- [31] Y. LeCun, L. Botto, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [32] K. Chatfield, K. Simonyan, A. Vedaldi, and A. Zisserman, "Return of the devil in the details: Delving deep into convolutional nets," in *BMVC*, 2014.
- [33] M. Varma and B. Babu, "More generality in efficient multiple kernel learning," in *ICML*, 2009.