

Brain Image Synthesis with Unsupervised Multivariate Canonical $CSC\ell_4$ Net

Yawen Huang¹, Feng Zheng^{2*}, Danyang Wang¹, Weilin Huang¹, Matthew R. Scott¹, Ling Shao³

¹Malong LLC, ²Southern University of Science and Technology, ³Inception Institute of Artificial Intelligence
{yawhuang, danwang, whuang, mscott}@malongtech.com, zhengf@sustech.edu.cn, ling.shao@ieee.org

Abstract

Recent advances in neuroscience have highlighted the effectiveness of multi-modal medical data for investigating certain pathologies and understanding human cognition. However, obtaining full sets of different modalities is limited by various factors, such as long acquisition times, high examination costs and artifact suppression. In addition, the complexity, high dimensionality and heterogeneity of neuroimaging data remains another key challenge in leveraging existing randomized scans effectively, as data of the same modality is often measured differently by different machines. There is a clear need to go beyond the traditional imaging-dependent process and synthesize anatomically specific target-modality data from a source input. In this paper, we propose to learn dedicated features that cross both inter- and intra-modal variations using a novel $CSC\ell_4$ Net. Through an initial unification of intra-modal data in the feature maps and multivariate canonical adaptation, $CSC\ell_4$ Net facilitates feature-level mutual transformation. The positive definite Riemannian manifold-penalized data fidelity term further enables $CSC\ell_4$ Net to reconstruct missing measurements according to transformed features. Finally, the maximization ℓ_4 -norm boils down to a computationally efficient optimization problem. Extensive experiments validate the ability and robustness of our $CSC\ell_4$ Net compared to the state-of-the-art methods on multiple datasets.

1. Introduction

Cranio-cerebral examination can be carried out using a multitude of imaging techniques with varying degrees of specificity and invasiveness, each directly or indirectly quantifying the structure, function and pathology of the brain. These multi-modal neuroimaging techniques, such as magnetic resonance imaging (MRI) and positron emission tomography (PET), offer diverse and complementary information to investigate human cognitive activities, population

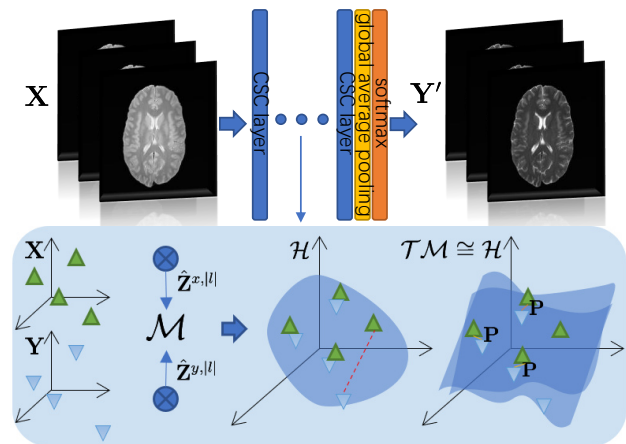


Figure 1. Architecture of our $CSC\ell_4$ Net. $CSC\ell_4$ Net is constructed by repeatedly stacking multivariate canonical CSC layers. The blue bar denotes the CSC layer, the yellow bar is the global average pooling, and the orange bar shows the softmax. $\hat{Z}^{x,|l|}$ and $\hat{Z}^{y,|l|}$ are feature maps of the l -th layer. \mathcal{M} represents the manifold on each tangent space \mathcal{T} having \mathcal{TM} . \mathbf{P} is the associator and \mathcal{H} denotes a Hilbert space.

imaging cohorts, neurodegeneration, and certain pathology. However, acquiring a full library of multi-modal images is impractical since the collection faces several constraints, including long acquisition times (e.g. a normal MRI scan can take as long as an hour), high examination cost, or even worse, image corruption in the event of artifacts from patient motion. Missing data is a critical problem in neurological studies and clinical diagnosis [11], and thus there is a clear need to obtain the absent data through beyond simple scanning.

Recently, there has been a surge of interest in synthesizing target-modality medical images by transferring information across different appearances [12, 28, 30, 37]. One early and noteworthy model for this was joint sparse representation [11, 24], which allows multi-modal data to be mapped in a common space rather than in separate ways to obtain a linear approximation. Based on this, the convolutional sparse coding (CSC) model replaces the local optimization with a global shift-invariant one, achieving significant improvements [8, 10, 26]. Further, deep learning has

*Corresponding author

obtained promising results in multi-modal image synthesis, mostly with convolutional neural networks (CNNs) [4] and generative adversarial networks (GANs) [34, 37].

While synthesis methods have had significant impact on research, there is now a debate regarding whether such synthetic images can substitute real acquisitions in clinical analyses. In general, clinical diagnosis requires multiple biomarkers to identify the disease and its status. When the target acquisitions are missing, accurate synthesis is essential. Although this challenge has been tackled by generating objective modality byproducts, current results remain unacceptable for clinical diagnosis.

In addition to differences in modalities, the complexity, high-dimensionality and heterogeneity of medical data remains another key challenge in leveraging existing randomized scans effectively. Specifically, imaging using machines developed by different manufacturers (*e.g.* Philips, Siemens, GE, *etc.*), the abundance of various physical parameters, and the presence of temporal dependency, all introduce conflicted and inconsistent features, thus preventing the complete use of real acquired data. It is nontrivial to harmonize all the different information and construct their correlations, but efforts to address the above challenges and develop a reliable algorithm to effectively utilize data are extremely necessary for both research and clinical decision support.

In this paper, we propose an unsupervised multivariate canonical CSC ℓ_4 Net, a novel approach to crossing both intra-modal (*i.e.* more than one measurement from the same data modality) and inter-modal (*i.e.* more than one data modality) heterogeneities. Our model synthesizes anatomically specific target modality data from a source modality, and makes efficient use of real acquisitions. CSC ℓ_4 Net works well under multiple datasets, despite the high dimensionality, temporal dependency and irregularity of neuroimaging, making it possible to combine acquisitions from different scanner manufacturers by initially normalizing differences between features, and then mapping them into the Hilbert space for multivariate canonical adaptation. Both cross-modal geometry transformations and a neuroimaging-specific positive definite conditions are incorporated within a Riemannian manifold. Finally, solving an ℓ_4 -maximization instead of an ℓ_1 -minimization problem enables us to employ the lowest sample complexity for high computational medical data. An overview of our CSC ℓ_4 Net is shown in Fig. 1.

To summarize, this paper provides the following contributions:

- To the best of our knowledge, this is the first work to generate anatomically meaningful images, by modeling an unsupervised multivariate canonical CSC ℓ_4 Net.
- We propose a novel intra-modal unit normalization for the initial unification of variate data of the same modal-

ity to guarantee a unique convolutional sparse solution.

- The multivariate canonical feature mapping is formulated over the multi-layer CSC to optimize the inter-modal structure.
- We introduce a Riemannian manifold-penalized transformation data fidelity term under the positive definite condition, for which we show how a reformulation based on CSC is crucial to empirical success.
- We prove that maximizing the ℓ_4 -norm instead of minimizing the ℓ_1 -norm leads to the lowest complexity and the highest robustness.

2. Related Work

2.1. Image Synthesis

Image synthesis (*a.k.a.* image-to-image translation) is commonly performed via appearance transformations, such as linear regression or distribution transformation. Previous works with stand-alone image pairs tend to focus on constructing a linear relationship in different contrasts [24]. Limited by few off-the-shelf paired data, Vemulapalli *et al.* [28] relaxed the invariable supervision by matching similarities across different modalities of image patches, and then jointly maximizing both global mutual information and local spatial consistency. Huang *et al.* [11] proposed a data-efficient synthesis method by mapping both a few pairs and large amounts of unpaired patches into a high-dimensional space, and then adopting Laplacian eigenmaps for geometric co-regularization. Neural style transfer [4] is another popular strategy for content-fixed image style translation, which computes the distance via the Gram matrix statistics of pre-trained deep features. GAN [6] was introduced to generate images using a random noise vector with discriminator judgment. Subsequently, many improvements and task-oriented generative models have been proposed. CycleGAN [39] uses a simple yet efficient strategy with cycle-consistent adversarial networks for unsupervised image-to-image translation. To synthesize high-resolution images from semantic labels, Wang *et al.* [30] proposed an adversarial learning objective which leverages GANs in a conditional setting and discards hand-crafted losses or pre-trained networks. FUNIT [16] contains a content and class-level encoder and an overall decoder for synthesizing the analogous image in a desirable category from the given class input. Leveraging the U-Net architecture, a data augmentation method [35] was established by learning both spatial deformation fields and intensity transforms to generate samples. TrGAN [29] focus on improving unsupervised image synthesis and representation learning.

2.2. Medical Image Analysis

The existing studies on medical imaging [12, 37, 13], *e.g.* synthesis, segmentation and registration, have shown

great promise for either research purposes or clinical analysis, mostly toward a macro objective, *i.e.* computer-aided diagnosis (CAD). In [14], a probabilistic model was proposed for joint registration and synthesis with cross-modality alignment. Uzunova *et al.* [27] used a multi-scale GAN to generate large amounts of high-quality medical images by learning a growing resolution conditioned on front scales. Shao *et al.* [25] presented a diagnosis-guided multi-modal feature selection method for prognostic prediction of a specific disease. Ravi *et al.* [23] employed adversarial learning in their proposed DaniNet, a degenerative adversarial neuroimage network that allows neurodegeneration to be modeled.

One of the fundamental purposes of medical image processing is to implement CAD, which in turn allows clinicians to make accurate decisions or provide treatment. Toward generalization and practicability, our work is complementary to the aforementioned approaches, where we improve the usage of large-scale heterogeneous medical data in an unsupervised manner.

2.3. Convolutional Sparse Coding

Convolutional sparse coding (CSC) has been successfully used in a wide range of computer vision and medical image processing problems [3, 8, 10]. CSC deals with the suboptimality of conventional local-independent representations by introducing a global shift-invariant filter. Fast CSC model was proposed by Bristow *et al.* [3], where the quad-decomposition solves the CSC objective in the Fourier domain, resulting in fast training. CSC-SR, introduced by Gu *et al.* [8], uses CSC with improved consistency to super-resolve natural images. Heide [10] tackled feature learning with fast and flexible CSC. Huang *et al.* [12] constructed two invertible mappings based on CSC for cross-modality synthesis and super-resolution of brain images. To reconstruct clean images, while avoiding adversarial attacks, Sun *et al.* [26] presented a stratified CSC algorithm with the benefit of an input transformation-based defense.

3. Mathematical Description of CSC

Given samples $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_S\}$ in \mathbb{R}^N , the problem of learning a set of convolutions of sparse feature maps $\mathbf{z}_k \in \mathbb{R}^N$ by filters $\mathbf{f}_k \in \mathbb{R}^M$, $\forall k = \{1, \dots, K\}$ can be expressed as minimizing the optimization function that combines the least-squares error and the ℓ_1 -norm penalty on the representations:

$$\arg \min_{\mathbf{f}, \mathbf{z}} \frac{1}{2} \left\| \mathbf{x} - \sum_{k=1}^K \mathbf{f}_k * \mathbf{z}_k \right\|_2^2 + \lambda \sum_{k=1}^K \|\mathbf{z}_k\|_1 \quad (1)$$

$$s.t. \|\mathbf{f}_k\|_2^2 \leq 1 \quad \forall k = \{1, \dots, K\},$$

where $*$ represents the 2D convolution operator and λ denotes a regularization parameter for the sparsity of the ℓ_1 -

norm. The objective of Eq. (1) is not jointly convex with respect to \mathbf{f} and \mathbf{z} , but is convex in optimizing one variable while fixing the other [31]. On the basis of such a convex optimization theory, the alternating direction method of multipliers (ADMM) was presented to solve the augmented Lagrangian formulation by introducing more proxies solving in the Fourier domain. ADMM transforms the convolution operation to an element-wise multiplication in order to speed up the spatial dominated convolution, *i.e.*, from $\mathcal{O}(KNM)$ time complexity in the spatial space to $\mathcal{O}(KN \log N)$ in the frequency domain. Despite the accelerated computation attained by FCSC [10], the auxiliary variables are still deemed heavy training parts, especially for tuning. More recent works [19, 40] alleviate such a limitation by forming the coordinate descent solution in a local greedy fashion.

Once the biconvex problem of Eq. (1) has been solved by alternating between learning filters and learning feature maps, the reconstruction can be obtained by a summation of the convolution outputs (*i.e.* convolving every row of \mathbf{f} by \mathbf{z}), leading to $\mathbf{x}' = \mathbf{f}_k * \mathbf{z}_k$.

4. Multivariate Canonical CSC ℓ_4 Net

4.1. Intra-Modal Unit Normalization

Let $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_S\}$ be a source domain training set containing S source modality images, and $\mathbf{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_T\}$ be a target domain training set containing T target modality examples. We define \mathbf{F} as the filter of the convolutional feature maps \mathbf{Z} . When standard training following the independence assumption illustrated in Eq. (1) is applied to two modalities, this leads to two separate filters of shift-invariant atoms $\mathbf{F}^x, \mathbf{F}^y$, with their corresponding feature maps \mathbf{Z}^x and \mathbf{Z}^y . However, when the same modality data are sampled from random measurements (*i.e.* intra-modal variation), the new features of these variables are no longer a CSC solution. To overcome this problem, we attempt to regularize the intra-modal data to guarantee a unique convolutionally sparse solution. Considering that a unit normalization of data is taken from the random measurements of one modality, one possibility is to normalize the column elements of \mathbf{Z}^x and \mathbf{Z}^y to unity, *i.e.*, $\|\mathbf{Z}_i^x\|_2^2 = 1, \|\mathbf{Z}_j^y\|_2^2 = 1, \forall i = 1, \dots, S, \forall j = 1, \dots, T$. However, using the unit normalization leads to "erased" modality information, which is especially the case in the context of neuroimaging data. To circumvent this issue, we eliminate the intra-modal scaling ambiguity by first computing the maximum ℓ_2 norm of \mathbf{Z}^x and \mathbf{Z}^y , respectively, and then performing our intra-modal unit normalization (IUN). The IUN (termed as Υ) of the convolutional feature maps becomes

$$\hat{\mathbf{z}}_i^x = \frac{\mathbf{Z}_i^x}{\max(\|\mathbf{Z}_i^x\|_2) \sqrt{1 - \|\mathbf{Z}_i^x\|_2^2}}, \forall i = 1, \dots, S,$$

$$\hat{\mathbf{Z}}_j^y = \frac{\mathbf{Z}_j^y}{\max(\|\mathbf{Z}_j^y\|_2, \sqrt{1 - \|\mathbf{Z}_j^y\|_2^2})}, \forall j = 1, \dots, T. \quad (2)$$

The restricted elements of \mathbf{Z}^x and \mathbf{Z}^y need to be unified by IUN and satisfy the general unit normalization $\|\hat{\mathbf{Z}}_i^x\|_2^2 = 1, \forall i$, and $\|\hat{\mathbf{Z}}_j^y\|_2^2 = 1, \forall j$.

4.2. Multivariate Canonical Feature Mapping

The CSC model of Eq. (1) has the advantage of computational efficiency when compared to the network-relevant methods. However, the pure penalty under ℓ_1 -norm ignores the diversity and complexity of data, leading to unsatisfactory results. Recent studies [6, 36] have shown that, by stacking multiple layers on top of each other, the extracted features are deeper and thus the performance of applications (*e.g.* classification or reconstruction) can be further boosted. Inspired by the success of extensive efforts on networks, we construct a CSC network architecture with a novel multivariate canonical adapted feature mapping layer, which has cross-modal learning power with synthetic potential.

Without loss of generality, we consider the CSC over IUN, which has multiple layers (termed as CSC-Net). We denote a function $f(\cdot, \cdot)$ for the feature mapping of each layer, such that $\mathbf{Z}^{x'} = f(\mathbf{X}, \Psi^x)$, $\mathbf{Z}^{y'} = f(\mathbf{Y}, \Psi^y)$, where $\Psi^x = \{\mathbf{F}^x, \Upsilon, \lambda\}$, $\Psi^y = \{\mathbf{F}^y, \Upsilon, \lambda\}$. Let l be the network layers, where $l \in [1, L]$. In CSC-Net, the feature maps $\hat{\mathbf{Z}}$ can be defined as the representation of the layer l with tensor properties of height h and width w : $\hat{\mathbf{Z}}_i^{x,|l|} \in \mathbb{R}^{N^{|l|} \times h^{|l|} \times w^{|l|}}, \forall i, l$, and $\hat{\mathbf{Z}}_j^{y,|l|} \in \mathbb{R}^{N^{|l|} \times h^{|l|} \times w^{|l|}}, \forall j, l$. To hierarchically approximate the convolutional feature maps, we construct a sparse intermediate representation which imposes the same structure on the upper layer of the representation. Intuitively, the l -th layer of \mathbf{Z} for \mathbf{X} and \mathbf{Y} can be estimated as $\hat{\mathbf{Z}}_i^{x,|l|} = f(\hat{\mathbf{Z}}_i^{x,|l-1|}, \Psi^{x,|l-1|})$ and $\hat{\mathbf{Z}}_j^{y,|l|} = f(\hat{\mathbf{Z}}_j^{y,|l-1|}, \Psi^{y,|l-1|})$, respectively.

In addition to the multi-layer sparse and deeper representation, another core module of the CSC-Net is the multivariate canonical adaptation (MCA). The multivariate formulation is initially normalized by IUN for intra-modal unity, and then we optimize the inter-modal (*i.e.* cross-modal data) structure using the constructed space, allowing us to utilize the feature maps learned in the previous step. Many domain-related studies [11, 38, 17] are based on the concept of feature shift, which can be summarized as learning a domain-invariant feature representation between the source and target domains for objective transformation. Motivated by the state-of-the-art domain adaptation works [21, 5], we cross-transfer both intra-modal and inter-modal features to a high-level projective space to handle the multivariate heterogeneous medical data. Specifically, we begin by importing a reproducing kernel Hilbert space (RKHS), where the

learned features between the CSC-Net layers and RKHS can be formed to minimize the maximum mean discrepancy (MMD) [7] between the source and target domains optimally, using a kernel $k(x_i, y_j) = \langle \hat{\mathbf{Z}}_i^x, \hat{\mathbf{Z}}_j^y \rangle_{\mathcal{H}}$ under a Hilbert space \mathcal{H} . Note that, $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ represents the inner product, and k is defined on the vector. Following the virtue of the MMD function in Eq. (3),

$$\begin{aligned} \mathcal{L}(X, Y) &= \frac{1}{S^2} \sum_{i=1}^S \sum_{j=1}^S k(\hat{\mathbf{Z}}_i^x, \hat{\mathbf{Z}}_j^x) + \frac{1}{T^2} \sum_{i=1}^T \sum_{j=1}^T k(\hat{\mathbf{Z}}_i^y, \hat{\mathbf{Z}}_j^y) \\ &\quad - \frac{2}{ST} \sum_{i=1}^S \sum_{j=1}^T k(\hat{\mathbf{Z}}_i^x, \hat{\mathbf{Z}}_j^y), \end{aligned} \quad (3)$$

we adapt the multiple feature layers using the multilayer MMD penalty (termed as $\mathcal{L}_{\mathcal{H}}(X, Y)$) for adapting the cross-modal CSC-Net, which is defined as

$$\begin{aligned} \mathcal{L}_{\mathcal{H}}(X, Y) &= \frac{1}{S^2} \sum_{i=1}^S \sum_{j=1}^S \prod_{l \in L} k^{||l|}(\hat{\mathbf{Z}}_i^{x,|l|}, \hat{\mathbf{Z}}_j^{x,|l|}) \\ &\quad + \frac{1}{T^2} \sum_{i=1}^T \sum_{j=1}^T \prod_{l \in L} k^{||l|}(\hat{\mathbf{Z}}_i^{y,|l|}, \hat{\mathbf{Z}}_j^{y,|l|}) \\ &\quad - \frac{2}{ST} \sum_{i=1}^S \sum_{j=1}^T \prod_{l \in L} k^{||l|}(\hat{\mathbf{Z}}_i^{x,|l|}, \hat{\mathbf{Z}}_j^{y,|l|}). \end{aligned} \quad (4)$$

Here, the previously described kernel k is computed by the Gaussian kernel function defined on the vectorization of tensors $\hat{\mathbf{Z}}_i^{x,|l|} \in \mathbb{R}^{N^{|l|} \times h^{|l|} \times w^{|l|}}$ and $\hat{\mathbf{Z}}_j^{y,|l|} \in \mathbb{R}^{N^{|l|} \times h^{|l|} \times w^{|l|}}$ of the layer $l \in L$. The calculation in a RKHS space can be expressed as $e^{-\|\hat{\mathbf{Z}}_i^{x,|l|} - \hat{\mathbf{Z}}_j^{y,|l|}\|_p^2}$, where p is a bandwidth parameter reflected in the Gaussian kernel function. $\prod_{l \in L} k^{||l|}(\hat{\mathbf{Z}}_i^{x,|l|}, \hat{\mathbf{Z}}_j^{y,|l|})$ gives the multiplicative results of the inner products.

4.2.1 Geometric Matching

While $\mathcal{L}_{\mathcal{H}}(X, Y)$ is endowed with attractive properties for domain invariant representations, its inability to preserve the natural vector structure of a domain leads to the geometric mismatching problem. In addition, deriving an approach that both has an adequate transformation capability and satisfying the neuroimaging constraints is a key challenge. For medical imaging, the positive definite (PD) condition is necessary for diffeomorphisms [1]. This regulation (*i.e.* PD) is omitted in most image synthesis works. As a result, the image being approximated is superficially consistent, but the underlying tissue information or structures are incorrect. To ensure the correctness of medical image synthesis, both information geometry and the PD condition are considered within a Riemannian manifold (RM) [22].

Building on the multilayer MMD introduced earlier, suppose the RM is also a smooth manifold constructed in the Hilbert space \mathcal{H} . Theoretically, the RM is equipped with an inner product, denoted as $\langle \cdot, \cdot \rangle_{\mathcal{T}}$, on each tangent space \mathcal{T} having manifold \mathcal{M} , which can be defined as $\mathcal{T}\mathcal{M}$. As a consequence, the norm in the tangent space is equivalent to that of the Hilbert space $\mathcal{T}\mathcal{M} \cong \mathcal{H}$.

To make the sparse feature maps respect their intrinsic geometries, we assume that all learned maps comply with the original data properties, *i.e.* the distances between data structures, as well as their corresponding maps, are closed in the RM. Particularly for measurements under high-level feature maps, the main point of RM is to make sense of the image transformation in the manifold setting. Inspired by the benefits of joint learning [11], we follow such a strategy and model an associated function for converting \mathbf{Z}^x to \mathbf{Z}^y . This is done by a linear projection \mathbf{P} of \mathbf{Z}^x and \mathbf{Z}^y using a least square problem: $\left\| \hat{\mathbf{Z}}^y - \mathbf{P}\hat{\mathbf{Z}}^x \right\|_2$. We then recall how the RM can be defined analogously for preserving data fidelity over the image transformation term. We begin by giving the Riemannian metric $\langle \hat{\mathbf{Z}}_i^x, \hat{\mathbf{Z}}_j^y \rangle_{\mathcal{T}}$ for $\hat{\mathbf{Z}}_i^x$ and $\hat{\mathbf{Z}}_j^y$ in $\mathcal{T}\mathcal{M}$, and then we rewrite the associated loss in the above least square problem by computing distances on the manifold as $d_{\mathcal{M}}(\hat{\mathbf{Z}}^y, \mathbf{P}\hat{\mathbf{Z}}^x) = \left\| \log(\hat{\mathbf{Z}}^y)^{-\frac{1}{2}} \mathbf{P}\hat{\mathbf{Z}}^x \log(\hat{\mathbf{Z}}^y)^{-\frac{1}{2}} \right\|_2$, where \log denotes the matrix logarithm and $d_{\mathcal{M}}$ is affine invariant. Note that $d_{\mathcal{M}}$ is computed on RA by projecting the symmetric data from \mathcal{T} onto the manifold with the positive definite property [2]. For the l -layer CSC-Net, the coordinate representation of $d_{\mathcal{M}}$ is

$$\mathcal{L}_{\mathcal{M}}(X, Y) = \sum_{i=1}^S \sum_{j=1}^T \prod_{l \in L} d_{\mathcal{M}}^{(l)}(\hat{\mathbf{Z}}_i^{y,|l|}, \mathbf{P}_{i,j}^{(l)} \hat{\mathbf{Z}}_j^{x,|l|}). \quad (5)$$

Here, $\mathcal{L}_{\mathcal{M}}(X, Y)$ denotes the loss function for the multi-layer cross-modal data fidelity of the RM penalty under the positive definite condition.

4.3. ℓ_4 Maximization

Generally, CSC consists of minimizing a convolutional model-fitted least-square system and a sparse regularizer by adopting an ℓ_0 or ℓ_1 penalty to promote the expected sparsity for recovering objectives. The ℓ_0 -norm is NP-hard when finding a local minimum of a nonconvex function, while the ℓ_1 -norm provides a unique solution through a convex program in the polynomial time. Although existing CSC algorithms can be separately optimized by alternating subproblems to update sparse feature maps and filters under the non-smooth ℓ_1 penalty and the ℓ_2 constraint, when applied to neuroimaging (which has a high-dimensional character), this leads to a high computational complexity due to the quasi-polynomial problem.

Pioneering research [32] shows that maximizing the element-wise ℓ_4 -norm enables an interpretative, stable and robust algorithm for reducing the per-iteration cost. Moreover, the global geometry of the ℓ_4 -norm over a unit ℓ_2 -sphere guarantees a randomly initialized first-order gradient descent algorithm [33], since each saddle point has negative curvature. Based on the essence of the ℓ_4 -penalized heuristic formulation, *i.e.* replacing $\min \|\cdot\|_1$ as $\max \|\cdot\|_4^4$, we attempt to address the shortcomings of CSC by modeling ℓ_4 maximization directly as the sparsity penalty within the objective of CSC ℓ_4 Net, which can be formulated as

$$\begin{aligned} \min_{\mathcal{L}_{\mathcal{M}}, \mathcal{L}_{\mathcal{H}}, \mathcal{L}_{\mathcal{R}}} \max_S \lambda (\mathcal{S}(X) + \mathcal{S}(Y)) + \mathcal{L}_{\mathcal{R}}(X) + \mathcal{L}_{\mathcal{R}}(Y) \\ + \mathcal{L}_{\mathcal{M}}(X, Y) + \mathcal{L}_{\mathcal{H}}(X, Y). \end{aligned} \quad (6)$$

Here, $\mathcal{S}(X) = \sum_{i=1}^S \left\| \hat{\mathbf{Z}}_i^{x,|l|} \right\|_4^4$ and $\mathcal{S}(Y) = \sum_{j=1}^T \left\| \hat{\mathbf{Z}}_j^{y,|l|} \right\|_4^4$ represent the sparsity cost function for penalizing \mathbf{Z}^x and \mathbf{Z}^y to be sparse. Mathematically, $\|\cdot\|_4^4$ is the element-wise ℓ_4 -norm with an expression $\|\mathbf{Z}\|_4^4 = \sum_{i,j} \mathbf{Z}_{i,j}^4, \forall \mathbf{Z} \in \mathbb{R}$. $\mathcal{L}_{\mathcal{R}}(X)$ denotes the reconstruction loss of \mathbf{X} with $\frac{1}{2} \left\| \mathbf{X} - \sum_{i=1}^S \mathbf{F}_i^{x,|l|} * \hat{\mathbf{Z}}_i^{x,|l|} \right\|_2^2$ and $\mathcal{L}_{\mathcal{R}}(Y)$ is the reconstruction loss of \mathbf{Y} having $\frac{1}{2} \left\| \mathbf{Y} - \sum_{j=1}^T \mathbf{F}_j^{y,|l|} * \hat{\mathbf{Z}}_j^{y,|l|} \right\|_2^2$.

Accompanied with the reconstruction procedure and associated formulation to obtain the target-modality data from the available source domain, we should solve the two respective problems: \mathbf{Z} and \mathbf{F} . In this work, to motivate ℓ_4 -based formulations, we employ the matching, stretching and projection (MSP) [32] optimization method to solve our objective. Simply, the single-layer feature matrix \mathbf{Z} is optimized by introducing MSP on an orthogonal matrix $\mathbf{F}_o \in \mathbb{O}$ with $\mathbf{Z}^{x'} = \arg \max f(\mathbf{X}, \Psi^x)$, $\mathbf{Z}^{y'} = \arg \max f(\mathbf{Y}, \Psi^y)$. Tackling the problem of a multi-layer CSC ℓ_4 Net, the l -th layer of \mathbf{Z} can be expressed as

$$\begin{aligned} \hat{\mathbf{Z}}_i^{x,|l'|} &\leftarrow \arg \max \left\| \hat{\mathbf{Z}}_i^{x,|l|} \right\|_4^4 \text{ s.t. } \mathbf{X}_i^{(l)} = \mathbf{F}_i^{x,|l-1|} * \hat{\mathbf{Z}}_i^{x,|l|}, \\ \hat{\mathbf{Z}}_j^{y,|l'|} &\leftarrow \arg \max \left\| \hat{\mathbf{Z}}_j^{y,|l|} \right\|_4^4 \text{ s.t. } \mathbf{Y}_j^{(l)} = \mathbf{F}_j^{y,|l-1|} * \hat{\mathbf{Z}}_j^{y,|l|}, \end{aligned} \quad (7)$$

where $\mathbf{X}_i^{(l)} = \phi(\hat{\mathbf{Z}}_i^{x,|l|})$ and $\mathbf{Y}_j^{(l)} = \phi(\hat{\mathbf{Z}}_j^{y,|l|})$, and $\phi(\cdot)$ denotes the concatenation of all i -th features from the previous layer $l-1$. We summarize our CSC ℓ_4 Net in Algorithm 1.

4.4. Synthesis

Once the training stage is complete, we can get two sets of filters, \mathbf{F}^x and \mathbf{F}^y , and the associator \mathbf{P} . Given a test image \mathbf{T}^x with the source modality, the desirable target modality version of \mathbf{T}^x can be treated as $\mathbf{T}^y = \sum (\mathbf{Z}^{tx} \mathbf{P}) * \mathbf{F}_y$, where $\mathbf{Z}^{tx} = f(\mathbf{T}^x, \Psi^x)$.

Algorithm 1 CSC ℓ_4 Net

Input: Training data \mathbf{X}, \mathbf{Y} , parameter λ

- 1: Initialize: $\mathbf{Z}_0^x, \mathbf{Z}_0^y, \mathbf{F}_0^x \in \mathbb{O}, \mathbf{F}_0^y \in \mathbb{O}, \mathbf{P}_0$
- 2: Perform IUN by Eq. (2), $\mathbf{Z}_0^x \rightarrow \hat{\mathbf{Z}}_0^x, \mathbf{Z}_0^y \rightarrow \hat{\mathbf{Z}}_0^y$
- 3: **for** $l = 1, \dots, L$ **do**
- 4: $\hat{\mathbf{Z}}_i^{x,|l|'} \leftarrow \arg \max \left\| \hat{\mathbf{Z}}_i^{x,|l|} \right\|_4^4$ s.t. $\mathbf{X}_i^{x,|l|} = \mathbf{F}_i^{x,|l-1|} * \hat{\mathbf{Z}}_i^{x,|l|}$
 $\hat{\mathbf{Z}}_j^{y,|l|'} \leftarrow \arg \max \left\| \hat{\mathbf{Z}}_j^{y,|l|} \right\|_4^4$ s.t. $\mathbf{Y}_j^{y,|l|} = \mathbf{F}_j^{y,|l-1|} * \hat{\mathbf{Z}}_j^{y,|l|}$
- 5: Perform MSP, update \mathbf{P}
- 6: **end for**

Output: $\mathbf{F}^x, \mathbf{F}^y, \mathbf{P}$

5. Experiments

5.1. Experimental Setup

The proposed CSC ℓ_4 Net is evaluated on three datasets: IXI,¹ NAMIC Multimodality,² and BraTS.³ IXI dataset contains 578 healthy subjects, NAMIC dataset includes 10 normal controls and 10 schizophrenic cases, and BraTS dataset has 220 brain tumor subjects. We conduct extensive experiments on three scenarios: (1) Proton Density (PD) \rightleftharpoons T2 on IXI dataset, (2) T1 \rightleftharpoons T2 on NAMIC dataset, (3) FLAIR \rightleftharpoons T1 on BraTS. In each dataset, the existing well-aligned paired images are removed, with half of them per a side for a strict unsupervised setting. Specifically, 150 unpaired PD-w and T2-w MRI are selected from IXI dataset, 6 unpaired T1-w, T2-w acquisitions are picked from NAMIC dataset, and 70 unpaired T1-w and FLAIR data are chosen from BraTS dataset for training. The split 50 samples from IXI, 2 samples from NAMIC and 30 samples from BraTS are used for validation. The remaining data: 100 (IXI), 4 (NAMIC), and 40 (BraTS) are used for testing. It is worth noting that, in our experimental datasets, IXI only contains healthy images while both NAMIC and BraTS include pathological data. In other words, the first setting in our experiments considers healthy cases, the second setting involves a mix of healthy and pathological examples, and the third setting tests our method on common diseases. We tune the hyper-parameters of our model on the validation set. To verify whether the synthesized results can replace the ground truths with diagnostic acceptability, we feed both real scans and all generations into a classic and commonly used segmentation algorithm, FM-RIB software library (FSL) [15]. We classify cerebrospinal fluid (CSF), gray matter (GM), white matter (WM) and optional tumor lesions to obtain the average quantification of


¹<http://brain-development.org/ixi-dataset>²<http://insight-journal.org/midas/collection/view/190>³<https://www.med.upenn.edu/sbia/brats2018/data.html>

Figure 2. Visual comparisons of MIMeCS, V-US, cycleGAN, 3D-cGAN, MSGAN and ours for PD \rightarrow T2 on IXI dataset.

the whole brain volume. FSL⁴ is a comprehensive library of analysis, specific for functional and structural brain imaging data. We obtain the segmented results under the tissue prior probability templates in a default image space, and therefore there is no guarantee that our segmentation will exactly follow other methods. For evaluation criteria, we use PSNR, SSIM and Dice score (which measures segmentation overlap, with a higher value meaning a better result) to assess the performance of different methods. We demonstrate that the proposed CSC ℓ_4 Net exhibits competitive performance with fewer layers (and the corresponding parameters) compared to other deep learning methods.

5.2. Network Architecture

Inspired by ResNet [9], CSC ℓ_4 Net includes seven CSC layers with layered MMD and RM constraints, where the spatial subsampling operation is performed with a stride of 2 in the last two bottleneck layers. Batch normalization is incorporated layer-wise, following each CSC layer, for faster convergence. The last CSC layer lies on top of a global spatial average pooling layer. The bottleneck layer in CSC ℓ_4 Net outputs 64, 128, 256 with 3, 2, 2 times repeated stacks, respectively. We use a learning rate of 0.0002, a batch size of 16, a total of 200 epochs, and a sparse regularization parameter of 0.02. We work in 2D slices and employ the Adam optimizer. For the layered MMD parameters, we use the Gaussian kernel with bandwidths equipped as median pairwise squared distances. We use the same network configurations for all datasets.

5.3. Baseline Methods

We compare our CSC ℓ_4 Net with the most relevant and state-of-the-art synthesis methods: 1) MIMeCS [24]; 2) V-US [28]; 3) cycleGAN [39]; 4) 3D-cGAN [20]; 5) MSGAN [18]. Note that, MIMeCS and MSGAN are supervised methods, so we input originally paired data for their

⁴<https://fsl.fmrib.ox.ac.uk/fsl/fslwiki>

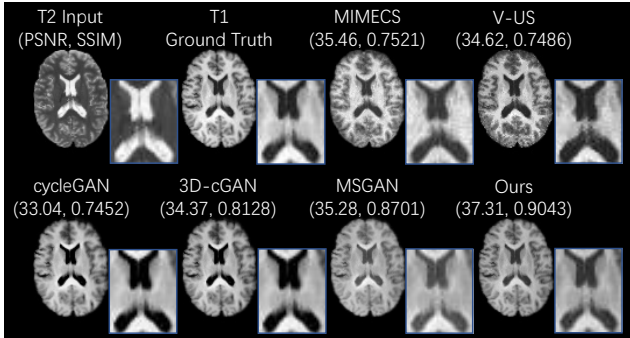


Figure 3. Visual comparisons of MIMeCS, V-US, cycleGAN, 3D-cGAN, MSGAN and ours for T2→T1 on NAMIC dataset.

Metric(avg.)	MIMeCS	V-US	cycleGAN	3D-cGAN	MSGAN	CSCl ₄ Net
IXI: T2-w → PD-w						
PSNR (dB)	30.49	32.53	32.66	32.75	33.11	36.45
SSIM	0.7801	0.8322	0.8345	0.8517	0.8557	0.9065
Dice (in %)	72.59	67.73	68.89	76.87	72.92	80.78
IXI: PD-w → T2-w						
PSNR (dB)	31.64	33.94	34.34	34.97	35.29	38.18
SSIM	0.8170	0.9022	0.9085	0.9035	0.8971	0.9596
Dice (in %)	76.52	69.87	70.26	80.63	79.78	87.90
NAMIC: T1-w → T2-w						
PSNR (dB)	35.97	35.50	36.41	35.44	35.78	37.82
SSIM	0.7690	0.7638	0.8291	0.8846	0.9033	0.9613
Dice (in %)	70.47	69.98	69.23	71.67	71.33	85.21
NAMIC: T2-w → T1-w						
PSNR (dB)	34.50	34.38	34.43	34.34	35.26	37.00
SSIM	0.7389	0.7517	0.8135	0.8673	0.8883	0.9222
Dice (in %)	72.19	70.09	70.07	74.62	72.69	83.64
BraTS: T1-w → FLAIR						
PSNR (dB)	30.50	31.87	31.37	33.74	31.81	37.41
SSIM	0.7944	0.8341	0.8129	0.8763	0.8653	0.9344
Dice (in %)	70.55	69.23	71.96	73.87	73.92	83.98
BraTS: FLAIR → T1-w						
PSNR (dB)	30.04	31.81	31.25	32.87	33.69	36.46
SSIM	0.8125	0.8421	0.8495	0.8798	0.8607	0.9112
Dice (in %)	71.59	69.83	73.21	78.87	76.99	82.63

Table 1. The performance of synthesis and synthesis-based segmentation on IXI, NAMIC and BraTS datasets.

training. V-US, cycleGAN and 3D-cGAN are unsupervised approaches, so we input the selected unpaired images for training. For fair comparison, we empirically set all parameters of the compared methods following their recommended data to reach their best performance. Notably, we exactly followed the data processing steps implemented by MIMeCS, V-US and 3D-cGAN, where a standard intensity correction and skull-stripping were preformed. Therefore, no extra data processes are implemented.

5.4. Results

We evaluate the effectiveness of our method by synthesizing images on different datasets. The results of CSCl₄Net and the five baseline methods are compared in Table 1. As further validation, we also apply the synthesized results to segment major tissues in each dataset and summarize the average performances in Table 1. First, we compare our CSCl₄Net against several state-of-the-art syn-




Figure 4. Visual comparisons of MIMeCS, V-US, cycleGAN, 3D-cGAN, MSGAN and ours for T1→FLAIR on BraTS dataset.

thesis methods on IXI for transforming PD-w MRI to T2-w MRI and *vice versa*. For better interpretation, we provide the visualization results in Fig. 2. From both Table 1 (top) and Fig. 2, we observe that CSCl₄Net achieves much better performance (visually and quantitatively) than the other five baseline methods in all cases. The average synthesis performances of CSCl₄Net on IXI dataset are {PSNR: 36.45dB, SSIM: 0.9065} and {PSNR: 38.18dB, SSIM: 0.9596} for T2-w → PD-w and PD-w → T2-w, respectively. CSCl₄Net gains a significant {PSNR, SSIM} performance improvement of {3.34dB, 0.0508} and {2.89dB, 0.0625}, respectively, compared to the best baseline MSGAN. In addition, we report the segmentation accuracy (*Dice scores*), which are based on the synthesized results, in Table 1. As we intuitively expected, if the model can provide results with data fidelity, the synthesis task is able to deliver practical usable results for further diagnosis. For medical image synthesis-driven segmentation, our method again outperforms the other methods, *i.e.*, MIMeCS, V-US, cycleGAN, 3D-cGAN and MSGAN, by a large margin. Our results are more representative than those of other state-of-the-art image synthesis methods on IXI dataset, showing dice improvements of 3.91% and 7.27% for T2-w → PD-w and PD-w → T2-w respectively, *w.r.t* the best baseline.

In addition to healthy cases, we also evaluate our algorithm on the pathological datasets (*i.e.* NAMIC, BraTS). We adopt the same evaluation criteria as on IXI dataset. We provide the qualitative results of our method in Figs. 3-4. These figures demonstrate how BrainGAN can handle schizophrenic and brain tumor synthesis on NAMIC and BraTS dataset, respectively. The proposed method consistently obtains the best performance, particularly when the lesions in the T1-w data appear with much lower contrast than in the FLAIR brain MRI (shown in Fig. 4). In Table 1 (middle and bottom), similarly for PD ⇌ T2 synthesis, CSCl₄Net outperforms the compared methods with improvements of {1.41dB, 0.0580} and {1.74dB, 0.0339} for T1 ⇌ T2 on NAMIC dataset, and {3.67dB, 0.0582} and {2.77dB, 0.0314} for T1 ⇌ FLAIR on BraTS dataset, respectively. Using the defined settings (refer to Section 5.1)

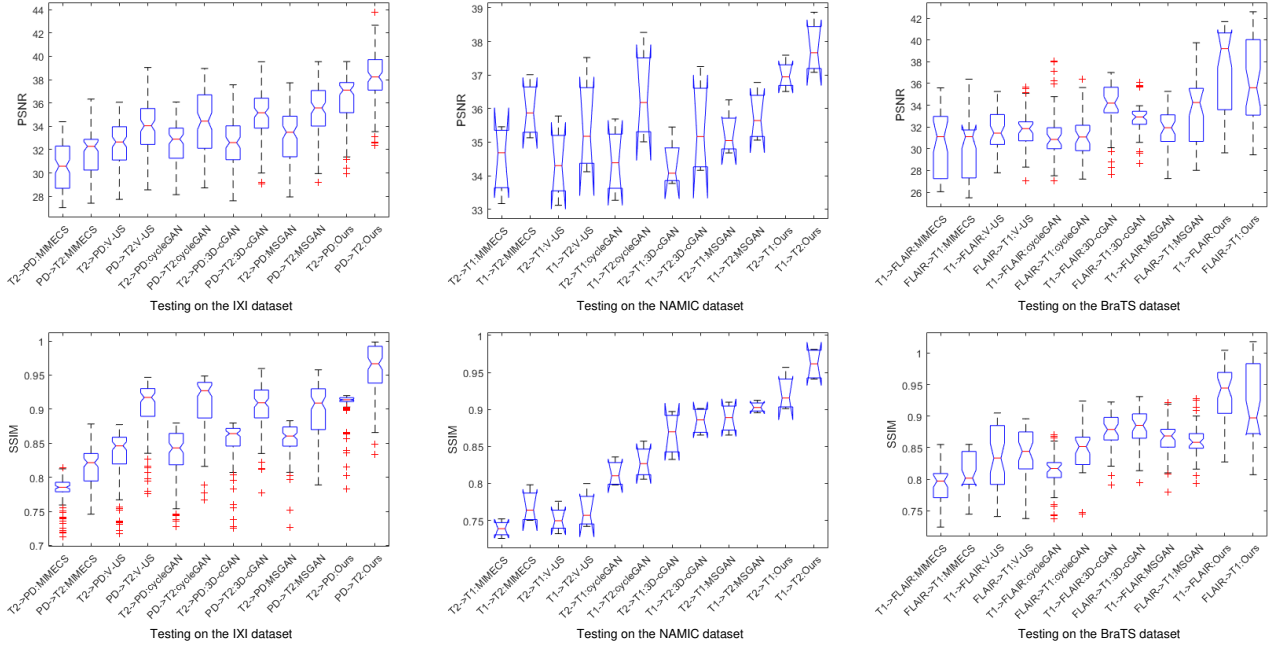


Figure 5. Quantitative evaluations including PSNRs and SSIMs of all methods distributed across IXI, NAMIC and BraTS, respectively.

and all synthesized data from both the forward and backward transformation, Table 1 compares the segmentation results. As with the Dice scores reported in Table 1 (middle and bottom), our method again gains significant performance enhancements, *i.e.*, increasing by 13.54% for $T1 \Rightarrow T2$, 9.02% for $T2 \Rightarrow T1$, 9.72% for $T1 \Rightarrow \text{FLAIR}$, and 3.76% for $\text{FLAIR} \Rightarrow T1$. For clarity, Fig. 5 provides detailed comparison results of all methods distributed across different datasets.

5.5. Ablation Studies

Since $\text{CSC}_{\ell_4}\text{Net}$ comprises a combination of several components, we perform an extensive ablation study to better understand why it is able to obtain state-of-the-art results. Considering the number of experiments in the above studies, we focus on the case of $\text{PD-w} \rightarrow \text{T2-w}$ from IXI dataset and report our ablation results in Table 2. We separately adopt $\{\text{IUN}, \mathcal{L}_{\mathcal{H}}, \mathcal{L}_{\mathcal{M}}\}$ as the additional module under the baseline CSC, and evaluate the effects in terms of image quality and synthesis-based segmentation performance. When each of the components is separately combined with CSC, the PSNR, SSIM, and Dice scores are improved by $\{2.42\text{dB}, 3.73\text{dB}, 3.54\text{dB}\}$, $\{0.05, 0.07, 0.07\}$, $\{0.35\%, 2.19\%, 7.79\%\}$, respectively. With the assistance of $\mathcal{L}_{\mathcal{H}}$ and $\mathcal{L}_{\mathcal{M}}$, the image quality and synthesis-based segmentation have the greatest impact. CSC with different combinations of two modules achieves $\{7.42\text{dB}, 7.05\text{dB}, 9.33\text{dB}\}$, $\{0.16, 0.16, 0.19\}$, and $\{13.44\%, 17.7\%, 26.43\%\}$ improvements over the baseline.

CSC	IUN	$\mathcal{L}_{\mathcal{H}}$	$\mathcal{L}_{\mathcal{M}}$	PSNR(dB)	SSIM	Dice(%)
✓				27.54	0.7357	58.73
✓	✓			29.96	0.7864	59.08
✓		✓		31.27	0.8022	60.92
✓			✓	31.08	0.8070	66.52
✓	✓	✓		34.96	0.8958	72.17
✓	✓		✓	34.59	0.8996	76.43
✓		✓	✓	36.87	0.9214	85.16

Table 2. Ablation study showing the effect of our modules/loss functions in improving the performance for $\text{PD-w} \rightarrow \text{T2-w}$ on IXI.

6. Conclusion

We have proposed a novel multivariate canonical $\text{CSC}_{\ell_4}\text{Net}$ approach for the cross-modal synthesis of medical images. $\text{CSC}_{\ell_4}\text{Net}$ aims at unifying multivariate datasets across both intra-modal and inter-modal through layer-wise feature adaptation and manifold transformation. $\text{CSC}_{\ell_4}\text{Net}$ is robust against both appearance variation and irregular machines. In addition, the proposed method extends the general CSC to a multi-layer CSC and imposes multivariate canonical feature mapping under the ℓ_4 -maximization to account for the high-dimensional and heterogeneous nature of neuroimaging. Comprehensive experiments show that $\text{CSC}_{\ell_4}\text{Net}$ is effective for a variety of cross-modality medical image synthesis problems, with segmentation quality, and can significantly outperform state-of-the-art methods even for very difference datasets. In the future, we plan to extend our $\text{CSC}_{\ell_4}\text{Net}$ to various image formats to investigate its generality.

Acknowledgment: This work is supported by the National Natural Science Foundation of China under Grant 61972188.

References

- [1] Peter J Basser, James Mattiello, and Denis LeBihan. Mr diffusion tensor spectroscopy and imaging. *Biophysical journal*, 66(1):259–267, 1994. 4
- [2] Rajendra Bhatia. *Positive definite matrices*, volume 24. Princeton university press, 2009. 5
- [3] Hilton Bristow, Anders Eriksson, and Simon Lucey. Fast convolutional sparse coding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 391–398, 2013. 3
- [4] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016. 2
- [5] Behnam Gholami, Pritish Sahu, Ognjen Rudovic, Konstantinos Bousmalis, and Vladimir Pavlovic. Unsupervised multi-target domain adaptation: An information theoretic approach. *IEEE Transactions on Image Processing*, 2020. 4
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 2, 4
- [7] Arthur Gretton, Dino Sejdinovic, Heiko Strathmann, Sivaraman Balakrishnan, Massimiliano Pontil, Kenji Fukumizu, and Bharath K Sriperumbudur. Optimal kernel choice for large-scale two-sample tests. In *Advances in neural information processing systems*, pages 1205–1213, 2012. 4
- [8] Shuhang Gu, Wangmeng Zuo, Qi Xie, Deyu Meng, Xiangchu Feng, and Lei Zhang. Convolutional sparse coding for image super-resolution. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1823–1831, 2015. 1, 3
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [10] Felix Heide, Wolfgang Heidrich, and Gordon Wetzstein. Fast and flexible convolutional sparse coding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5135–5143, 2015. 1, 3
- [11] Yawen Huang, Ling Shao, and Alejandro F Frangi. Cross-modality image synthesis via weakly coupled and geometry co-regularized joint dictionary learning. *IEEE transactions on medical imaging*, 37(3):815–827, 2017. 1, 2, 4, 5
- [12] Yawen Huang, Ling Shao, and Alejandro F Frangi. Dote: Dual convolutional filter learning for super-resolution and cross-modality synthesis in mri. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 89–98. Springer, 2017. 1, 2, 3
- [13] Yawen Huang, Feng Zheng, Runmin Cong, Weilin Huang, Matthew R Scott, and Ling Shao. Mcmt-gan: Multi-task coherent modality transferable gan for 3d brain image synthesis. *IEEE Transactions on Image Processing*, 29:8187–8198, 2020. 2
- [14] Juan Eugenio Iglesias, Marc Modat, Loïc Peter, Allison Stevens, Roberto Annunziata, Tom Vercauteren, Ed Lein, Bruce Fischl, Sebastien Ourselin, Alzheimers Disease Neuroimaging Initiative, et al. Joint registration and synthesis using a probabilistic model for alignment of mri and histological sections. *Medical image analysis*, 50:127–144, 2018. 3
- [15] Mark Jenkinson, Christian F Beckmann, Timothy EJ Behrens, Mark W Woolrich, and Stephen M Smith. Fsl. *Neuroimage*, 62(2):782–790, 2012. 6
- [16] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation. *arXiv preprint arXiv:1905.01723*, 2019. 2
- [17] Mingsheng Long, Jianmin Wang, Yue Cao, Jiaguang Sun, and S Yu Philip. Deep learning of transferable representation for scalable domain adaptation. *IEEE Transactions on Knowledge and Data Engineering*, 28(8):2027–2040, 2016. 4
- [18] Qi Mao, Hsin-Ying Lee, Hung-Yu Tseng, Siwei Ma, and Ming-Hsuan Yang. Mode seeking generative adversarial networks for diverse image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1429–1437, 2019. 6
- [19] Thomas Moreau, Laurent Oudre, and Nicolas Vayatis. Dico: Distributed convolutional coordinate descent for convolutional sparse coding. In *International Conference on Machine Learning*, pages 3623–3631, 2018. 3
- [20] Yongsheng Pan, Mingxia Liu, Chunfeng Lian, Tao Zhou, Yong Xia, and Dinggang Shen. Synthesizing missing pet from mri with cycle-consistent generative adversarial networks for alzheimers disease diagnosis. In *MICCAI*, pages 455–463. Springer, 2018. 6
- [21] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1406–1415, 2019. 4
- [22] Peter Petersen, S Axler, and KA Ribet. *Riemannian geometry*, volume 171. Springer, 2006. 4
- [23] Daniele Ravi, Daniel C Alexander, Neil P Oxtoby, Alzheimers Disease Neuroimaging Initiative, et al. Degenerative adversarial neuroimage nets: Generating images that mimic disease progression. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 164–172. Springer, 2019. 3
- [24] Snehashis Roy, Aaron Carass, and Jerry L Prince. Magnetic resonance image example-based contrast synthesis. *IEEE transactions on medical imaging*, 32(12):2348–2363, 2013. 1, 2, 6
- [25] Wei Shao, Tongxin Wang, Zhi Huang, Jun Cheng, Zhi Han, Daoqiang Zhang, and Kun Huang. Diagnosis-guided multimodal feature selection for prognosis prediction of lung squamous cell carcinoma. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 113–121. Springer, 2019. 3
- [26] Bo Sun, Nian-hsuan Tsai, Fangchen Liu, Ronald Yu, and Hao Su. Adversarial defense by stratified convolutional sparse coding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11447–11456, 2019. 1, 3

- [27] Hristina Uzunova, Jan Ehrhardt, Fabian Jacob, Alex Frydrychowicz, and Heinz Handels. Multi-scale gans for memory-efficient generation of high resolution medical images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 112–120. Springer, 2019. 3
- [28] Raviteja Vemulapalli, Hien Van Nguyen, and Shaohua Kevin Zhou. Unsupervised cross-modal synthesis of subject-specific scans. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 630–638, 2015. 1, 2, 6
- [29] Jiayu Wang, Wengang Zhou, Guo-Jun Qi, Zhongqian Fu, Qi Tian, and Houqiang Li. Transformation gan for unsupervised image synthesis and representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 472–481, 2020. 2
- [30] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018. 1, 2
- [31] Matthew D Zeiler, Dilip Krishnan, Graham W Taylor, and Rob Fergus. Deconvolutional networks. In *2010 IEEE Computer Society Conference on computer vision and pattern recognition*, pages 2528–2535. IEEE, 2010. 3
- [32] Yuexiang Zhai, Zitong Yang, Zhenyu Liao, John Wright, and Yi Ma. Complete dictionary learning via ℓ_4 -norm maximization over the orthogonal group. *arXiv preprint arXiv:1906.02435*, 2019. 5
- [33] Yuqian Zhang, Han-Wen Kuo, and John Wright. Structured local optima in sparse blind deconvolution. *IEEE Transactions on Information Theory*, 66(1):419–452, 2019. 5
- [34] Zizhao Zhang, Lin Yang, and Yefeng Zheng. Translating and segmenting multimodal medical volumes with cycle-and shape-consistency generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9242–9251, 2018. 2
- [35] Amy Zhao, Guha Balakrishnan, Fredo Durand, John V Guttag, and Adrian V Dalca. Data augmentation using learned transformations for one-shot medical image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8543–8553, 2019. 2
- [36] Feng Zheng, Cheng Deng, Xing Sun, Xinyang Jiang, Xiaowei Guo, Zongqiao Yu, Feiyue Huang, and Rongrong Ji. Pyramidal person re-identification via multi-loss dynamic training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8514–8522, 2019. 4
- [37] Yi Zhou, Xiaodong He, Shanshan Cui, Fan Zhu, Li Liu, and Ling Shao. High-resolution diabetic retinopathy image synthesis manipulated by grading and lesions. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 505–513. Springer, 2019. 1, 2
- [38] Fan Zhu and Ling Shao. Weakly-supervised cross-domain dictionary learning for visual recognition. *International Journal of Computer Vision*, 109(1-2):42–59, 2014. 4
- [39] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 2, 6
- [40] Ev Zisselman, Jeremias Sulam, and Michael Elad. A local block coordinate descent algorithm for the csc model. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8208–8217, 2019. 3