



# Representation Disentanglement for Multi-task Learning with Application to Fetal Ultrasound

Qingjie Meng<sup>(✉)</sup>, Nick Pawlowski, Daniel Rueckert, and Bernhard Kainz


Department of Computing, BioMedIA, Imperial College London, London, UK  
q.meng16@imperial.ac.uk

**Abstract.** One of the biggest challenges for deep learning algorithms in medical image analysis is the indiscriminate mixing of image properties, *e.g.* artifacts and anatomy. These entangled image properties lead to a semantically redundant feature encoding for the relevant task and thus lead to poor generalization of deep learning algorithms. In this paper we propose a novel representation disentanglement method to extract semantically meaningful and generalizable features for different tasks within a multi-task learning framework. Deep neural networks are utilized to ensure that the encoded features are maximally informative with respect to relevant tasks, while an adversarial regularization encourages these features to be disentangled and minimally informative about irrelevant tasks. We aim to use the disentangled representations to generalize the applicability of deep neural networks. We demonstrate the advantages of the proposed method on synthetic data as well as fetal ultrasound images. Our experiments illustrate that our method is capable of learning disentangled internal representations. It outperforms baseline methods in multiple tasks, especially on images with new properties, *e.g.* previously unseen artifacts in fetal ultrasound.

## 1 Introduction

Image interpretation using convolutional neural networks (CNNs) has been widely and successfully applied to medical image analysis during recent years. However, in contrast to human observers, CNNs exhibit weaknesses of being generalized to tackle previously unseen entangled image properties (*e.g.* shape and texture) [6]. In Ultrasound (US), the image property entanglement can be observed when acquisition-related artifacts (*e.g.* shadows) obfuscate the underlying anatomy (see Fig. 1). A CNN simultaneously learns anatomical features and artifacts features for either anatomy classification or artifacts detection [15]. As a result, the model trained by images with certain entangled properties (*e.g.* images without acoustic shadows) can hardly handle images with new entangled properties which are unseen during training (*e.g.* images with shadows).

Approaches for representation disentanglement have been proposed in order to learn semantically disjoint internal representations for improving image interpretation [12]. These methods pave a way for improving the generalization of



**Fig. 1.** Examples of fetal US data. Green framed images are shadow-free and red framed images contain acoustic shadows. (Color figure online)

CNNs in a wide range of medical image analysis problems. Specifically for a practical application in this work, we want to disentangle anatomical features from shadow features so that to generalize anatomical standard plane analysis for a better detection of abnormality in early pregnancy.

**Contribution:** In this paper, we propose a novel, end-to-end trainable representation disentanglement model that can learn distinct and generalizable features through a multi-task architecture with adversarial training. The obtained disjoint features are able to improve the performance of multi-task networks, especially on data with previously unseen properties. We evaluate the proposed model on specific multi-task problems, including shape/background-color classification tasks on synthetic data and standard-plane/shadow-artifacts classification tasks on fetal US data. Our experiments show that our model is able to disentangle latent representations and, in a practical application, improves the performance for anatomy analysis in US imaging.

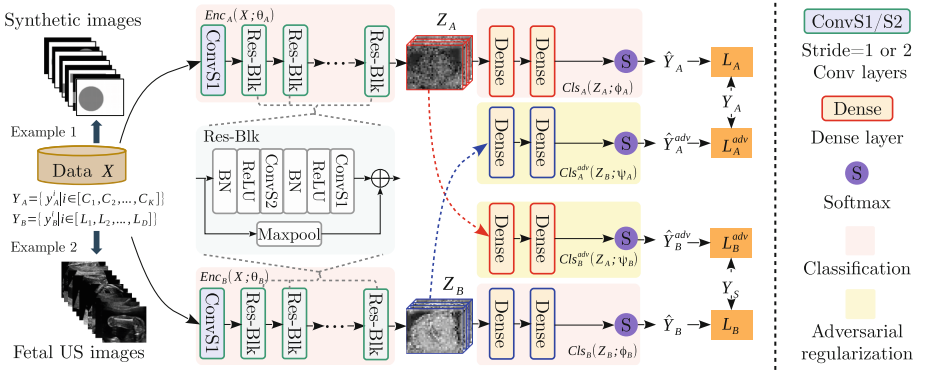
**Related work:** Representation disentanglement has been widely studied in the machine learning literature, ranging from traditional models such as Independent Component Analysis (ICA) [10] and bilinear models [18] to recent deep learning-based models such as InfoGAN [4] and  $\beta$ -VAE [3, 9]. Disentangled representations can be utilized to interpret complex interactions of underlying factors within data [2, 5] and enable deep learning models to manipulate relevant information for specific tasks [7, 8, 13]. Particularly related to our work is the work by Mathieu et al. [14], which proposed a conditional generative model with adversarial networks to disentangle specific and unspecific factors of variation in deep representations without strong supervision. Compared to [14], Hadad et al. [8] proposed a simpler two-step method with the same aim. Their network directly utilizes the encoded latent space without assuming the underlying distribution, which can be more efficient for learning various unspecified features. Different from their aim – disentangling one *specific* representation from *unspecific* factors – our work focuses on disentangling several *specific* factors. Further related to our research question is to learn only unspecific invariant features, for example, for domain adaptation [11]. However, unlike learning invariant features, which ignores task-irrelevant information [2], our method aims to preserve information for multiple tasks while enhancing feature generalizability.

In the medical image analysis community, few approaches have focused on disentangling internal factors of representations in discriminative tasks. Ben-Cohen et al. [1] proposed a method to disentangle lesion type from image appearance and use disentangled features to generate more training samples for data

augmentation. Their work improves liver lesions classification. In contrast, our work aims to utilize disentangled features for generalization of deep neural networks in medical image analysis.

## 2 Method

Our goal is to disentangle latent representations  $Z$  of the data  $X$  into distinct feature sets ( $Z_A, Z_B$ ) that separately contain relevant information for corresponding different tasks ( $T_A, T_B$ ). The main motivation of the proposed method is to learn feature sets that are maximally informative about their corresponding task (e.g.  $Z_A \rightarrow T_A$ ) but minimally representative for irrelevant tasks (e.g.  $Z_A \rightarrow T_B$ ). While our approach scales to any number of classification tasks, in this work we focus on two tasks as a proof of concept. The proposed method consists of two classification tasks ( $T_A, T_B$ ) with an adversarial regularization. The classification aims to map the encoded features to their relevant class identities, and is trained to maximize  $I(Z_A, Y_A)$  and  $I(Z_B, Y_B)$ . The adversarial regularization penalizes the mutual information between the encoded features and their irrelevant class identities, in other words, minimizes  $I(Z_A, Y_B)$  and  $I(Z_B, Y_A)$ . The training architecture of our method is shown in Fig. 2.



**Fig. 2.** Training framework for the proposed method. Res-Blk refers to residual-blocks. Example 1/2 are two data set examples used in Sect. 3. The classifications enables the encoded features  $Z_A, Z_B$  to be maximally informative about related tasks while the adversarial regularization encourages these features to be less informative about irrelevant tasks.

**Classification** is used to learn the encoded features that enable high prediction performance for the class identity of the relevant task. Each of the two classification networks is composed of an encoder and a classifier for a defined task. Given data  $X = \{x_i \mid i \in [1, N]\}$ , the matching labels are  $Y_A = \{y_A^i \mid y_A^i \in \{C_1, C_2, \dots, C_K\}, i \in [1, N]\}$  for  $T_A$  and  $Y_B = \{y_B^i \mid y_B^i \in$

$\{L_1, L_2, \dots, L_D\}, i \in [1, N]\}$  for  $T_B$ .  $N$  is the number of images and  $K, D$  are the number of class identities in each task. Two independent encoders map  $X$  to  $Z_A$  and  $Z_B$  with parameters  $\theta_A$  and  $\theta_B$  respectively, yielding  $Z_A = Enc_A(X; \theta_A)$  and  $Z_B = Enc_B(X; \theta_B)$ . Two classifiers are used to predict class identity for the corresponding task, where  $\hat{Y}_A = Cls_A(Z_A; \phi_A)$  and  $\hat{Y}_B = Cls_B(Z_B; \phi_B)$ .  $\phi_A$  and  $\phi_B$  are the parameters of the corresponding classifiers. We define the cost functions  $\mathcal{L}_A$  and  $\mathcal{L}_B$  as the softmax cross-entropy between  $Y_A$  and  $\hat{Y}_A$  and between  $Y_B$  and  $\hat{Y}_B$  respectively. The classification loss  $\mathcal{L}_{cls} = \mathcal{L}_A + \mathcal{L}_B$  is minimized to train the two encoders and the two classifiers ( $\min_{\{\theta_A, \theta_B, \phi_A, \phi_B\}} \mathcal{L}_{cls}$ ) for obtaining  $Z_A$  and  $Z_B$  that are maximally related to their relevant task.

**Adversarial regularization** is used to force the encoded features to be minimally informative about irrelevant tasks, which results in disentanglement of internal representations. The adversarial regularization is implemented by using an adversarial network for each task as shown in Fig. 2. These adversarial networks are utilized to map the encoded features to class identity of the irrelevant task, yielding  $\hat{Y}_A^{adv} = Cls_A^{adv}(Z_B; \psi_A)$  and  $\hat{Y}_B^{adv} = Cls_B^{adv}(Z_A; \psi_B)$ . Here,  $\psi_A$  and  $\psi_B$  are the parameters of the corresponding adversarial networks. By referring to  $\mathcal{L}_A^{adv}$  and  $\mathcal{L}_B^{adv}$  as the softmax cross-entropy between  $Y_A$  and  $\hat{Y}_A^{adv}$  and between  $Y_B$  and  $\hat{Y}_B^{adv}$ , the adversarial loss is defined as  $\mathcal{L}_{adv} = \mathcal{L}_A^{adv} + \mathcal{L}_B^{adv}$ . During training, the adversarial networks are trained to minimize  $\mathcal{L}_{adv}$  while two encoders and two classifiers are trained to maximize  $\mathcal{L}_{adv}$  ( $\min_{\{\psi_A, \psi_B\}} \max_{\{\theta_A, \theta_B, \phi_A, \phi_B\}} \mathcal{L}_{adv}$ ). This competition between the encoders/classifiers and the adversarial networks encourages the encoded features to be invalid for irrelevant tasks.

By combining the two classifications with the adversarial regularization, the whole model is optimized iteratively during training. The training objective for optimizing the two encoders and the two classifiers can be written as

$$\min_{\{\theta_A, \theta_B, \phi_A, \phi_B\}} \{\mathcal{L}_A + \mathcal{L}_B - \lambda * (\mathcal{L}_A^{adv} + \mathcal{L}_B^{adv})\}, \lambda > 0. \quad (1)$$

Here,  $\lambda$  is the trade-off parameter of the adversarial regularization. The training objective for the optimization of the adversarial networks thus follows as

$$\min_{\{\psi_A, \psi_B\}} \{\mathcal{L}_A^{adv} + \mathcal{L}_B^{adv}\}. \quad (2)$$

**Network Architectures:**  $Enc_A(X; \theta_A)$  and  $Enc_B(X; \theta_B)$  both consist of six residual-blocks implemented as proposed in [17] to reduce the training error and to support easier network optimization.  $Cls_A(Z_A; \phi_A)$  and  $Cls_B(Z_B; \phi_B)$  both contain two dense layers with 256 hidden units. The adversarial networks  $Cls_A^{adv}(Z_B; \psi_A)$  and  $Cls_B^{adv}(Z_A; \psi_B)$  have the same architecture as  $Cls_A(Z_A; \phi_A)$  and  $Cls_B(Z_B; \phi_B)$  respectively.

**Training:** Our model is optimized for 400 epochs and  $\lambda$  is chosen heuristically and independently for each data set using validation data. For more stable optimization [8], in each iteration, we train the encoders and classifiers once, followed by five training steps of the adversarial networks. Similar to [8], we use the Adam

optimizer (beta = 0.9, learning rate =  $10^{-5}$ ) to train the encoders and classifiers based on Eq. 1, and use Stochastic Gradient Descent (SGD) with momentum optimizer (momentum = 0.9, learning rate =  $10^{-5}$ ) to update the parameters of the adversarial networks in Eq. 2. We apply L2 regularization (scale =  $10^{-5}$ ) to all weights during training to prevent over-fitting. The batch size is 50 and the images in each batch have been randomly flipped as data augmentation. Our model is trained on a Nvidia Titan X GPU with 12 GB of memory.

### 3 Evaluation and Results

**Evaluation on Synthetic Data:** We use synthetic data as a proof of concept example to verify our model. This data set contains a randomly located gray circle or rectangle on a black or white background. We split the data into 1200/300/300 images for train/validation/test and these images consist of circles on white background, rectangles on black background and rectangles on white background. To keep the balance between image properties in the training split, we use circle:rectangle = 1:1 and black:white = 7:5. In this case,  $T_A$  is a background color classification task and  $T_B$  is the a shape classification task. We implement our model as outlined in Sect. 2 and choose  $\lambda = 0.01$ . We evaluate our model on the test data. The experimentation illustrates that the encoded features successfully identify the class identities of the relevant task (*e.g.*  $Z_A \rightarrow T_A : \text{OA}_{acc} = 100\%$ ,  $Z_B \rightarrow T_B : \text{OA}_{acc} = 99.67\%$ ) but fail to handle irrelevant task (*e.g.*  $Z_A \rightarrow T_B : \text{OA}_{acc} = 62\%$ ,  $Z_B \rightarrow T_A : \text{OA}_{acc} = 59.67\%$ ). Here,  $\text{OA}_{acc}$  is the overall accuracy. To show the utility of the proposed method on images with previously unseen entangled properties, we additionally compare the shape classification performance of our model and a baseline (our model without the adversarial regularization) on images with a previously unseen entangled properties (circles on black background). The proposed model achieves  $\text{OA}_{acc} = 99\%$  and outperforms the baseline which achieves  $\text{OA}_{acc} = 10\%$ . We use PCA to examine the learned embedding space at the penultimate dense layer of the classifiers. The top row of Fig. 3 illustrates that the extracted features is able to identify class identities for relevant tasks (see (a, c)) but unable to predict correct class identities for irrelevant tasks (see (b, d)).

**Evaluation on Fetal US Data:** We verify the applicability of our method on fetal US data. Here, we refer to an anatomical standard plane classification task as  $T_A$  and an acoustic shadow artifacts classification task as  $T_B$ . We want to learn the corresponding disentangled features  $Z_A$  for all anatomical information, separated from  $Z_B$  containing only information about shadow artifacts.  $Y_A$  is the label for different anatomical standard planes while  $Y_B^i = 0$  and  $Y_B^i = 1$  are the labels of the shadow-free class and the shadow-containing class respectively.

**Data Set:** The fetal US data set contains 8.4k images sampled from 4120 2D US fetal anomaly screening examinations with gestational ages between 18–22 weeks. These sequences consist of eight standard planes defined in the UK FASP handbook [16], including three vessel view (3VV), left ventricular outflow tract

**Table 1.** Data split. “Others” contains standard planes 4CH, femur, kidneys, lips and RVOT. Test\_seen, LVOT(W\_S) and Artifacts (OTHS) are used for testing.

		Train	Validation	Test_seen	LVOT(W_S)	Artifacts(OTHS)
3VV	W/O_S (W_S)	180 (320)	50 (50)	334 (41)	– (–)	– (–)
LVOT	W/O_S (W_S)	500 (–)	50 (–)	79 (–)	– (418)	– (–)
Abd	W/O_S (W_S)	125 (375)	50 (50)	190 (220)	– (–)	– (–)
Others	W/O_S (W_S)	– (–)	– (–)	– (–)	– (–)	3159 (2211)

(LVOT), abdominal (Abd.), four chamber view (4CH), femur, kidneys, lips and right ventricular outflow tract (RVOT), and are classified by expert observers as shadow-containing (W\_S) or shadow-free (W/O\_S) (Fig. 1). We split the data as shown in Table 1. Train, Validation and Test\_seen are separate data sets. Test\_seen contains the same entangled properties (but different images) as used for the training data set, while LVOT (W\_S) and Artifacts (OTHS) contain new combinations of entangled properties.

**Evaluation Approach:** We refer to *Std plane only* as the networks for standard plane classification only (consists of  $Enc_A$  and  $Cls_A$ ), and *Artifacts only* as the networks for shadow artifacts classification only (consists of  $Enc_B$  and  $Cls_B$ ).  $Proposed_{w/o\_adv}$  refers to the proposed method without the adversarial regularization and  $Proposed$  is our method in Fig. 2.

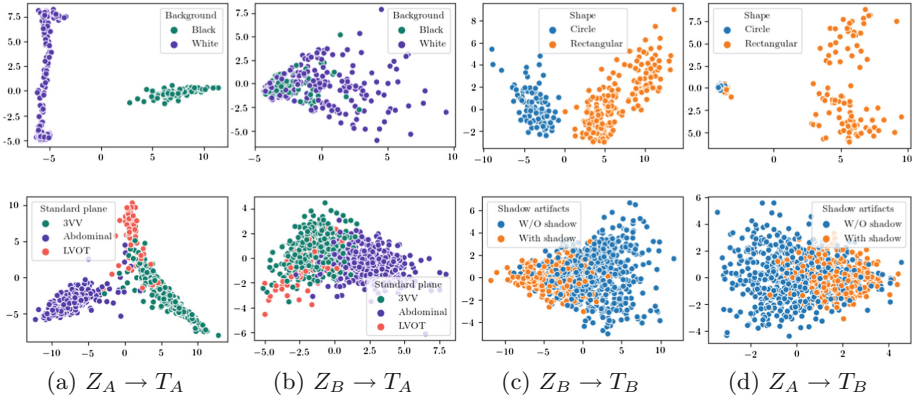
The proposed method is implemented as outlined in Sect. 2 choosing  $\lambda = 0.1$ .  $Cls_A(Z_A; \phi_A)$  contains three dense layers with 256/256/3 hidden units while  $Cls_B(Z_B; \phi_B)$  contains two dense layers with 256/2 hidden units. We choose a bigger network capacity for  $Cls_A(Z_A; \phi_A)$  by assuming that anatomies have more complex structures than shadows to be learned.

Table 2 shows that our method improves the performance of standard plane classification by 16.08% and 13.19% on Test\_seen when compared with the *Std plane only* and the  $Proposed_{w/o\_adv}$  method (see  $OA_{acc}$  in Col. 5). It achieves minimal improvement (*Artifacts only*: +0.35% and  $Proposed_{w/o\_adv}$ : +0.31% classification accuracy) for shadow artifacts classification (see  $OA_{acc}$  in Col. 8). We also demonstrate the utility of the proposed method on images with previously unseen entangled properties. Table 2 shows that the proposed method achieves 73.68% accuracy of standard plane classification on LVOT (W\_S) ( $\sim 36\%$  higher than other comparison methods) while it performs similar to other methods on Artifacts (OTHS) for shadow artifacts classification.

We evaluate the performance of disentanglement by using the encoded features for the irrelevant task on Test\_seen, e.g.  $Z_A \rightarrow T_B$  and  $Z_B \rightarrow T_A$ . Here,  $Z_A$  and  $Z_B$  are encoded features of the proposed method.  $Proposed_{irr\_task}$  in Table 2 indicates that  $Z_B$  contains much less anatomical information for standard plane classification ( $OA_{acc} = 94.44\%$  in proposed vs.  $OA_{acc} = 64.35\%$  in  $Proposed_{irr\_task}$ ), while  $Z_A$  contains less shadow features information ( $OA_{acc} = 79.05\%$  in proposed vs.  $OA_{acc} = 72.57\%$  in  $Proposed_{irr\_task}$ ). We additionally use PCA to show the embedded test data on the penultimate dense layer. The

**Table 2.** The classification accuracy (%) of different methods for the standard classification ( $T_A$ ) and shadow artifacts classification ( $T_B$ ) on Test<sub>seen</sub> data set and data sets with unseen entangled properties (LVOT(W\_S) and Artifacts(OTHS)). “Proposed” uses encoded features for relevant tasks, namely,  $Z_A \rightarrow T_A$  and  $Z_B \rightarrow T_B$ . “Proposed<sub>irr.task</sub>” uses encoded features for irrelevant tasks, namely,  $Z_A \rightarrow T_B$  and  $Z_B \rightarrow T_A$ . OA<sub>acc</sub> is the overall accuracy.

Col. 1	Col. 2	Col. 3	Col. 4	Col. 5	Col. 6	Col. 7	Col. 8	Col. 9	Col. 10
Methods	Test <sub>seen</sub>							LVOT (W_S)	Artifacts (OTHS)
	3VV	LVOT	Abd.	OA <sub>acc</sub>	W/O_S	W_S	OA <sub>acc</sub>		
Std plane only	60.80	96.59	67.09	78.36	—	—	—	34.93	—
Artifacts only	—	—	—	—	77.94	80.46	78.70	—	69.26
Proposed <sub>w/o_adv</sub>	63.73	97.80	78.48	81.25	78.77	77.78	78.74	37.56	69.50
<b>Proposed</b>	<b>93.87</b>	<b>97.56</b>	<b>81.01</b>	<b>94.44</b>	<b>87.89</b>	<b>58.62</b>	<b>79.05</b>	<b>73.68</b>	<b>68.49</b>
Proposed <sub>irr.task</sub>	39.20	83.90	82.28	64.35	68.49	81.99	72.57	—	—



**Fig. 3.** Visualization of the embedded data on the penultimate dense layer. The top row shows embedded synthetic test data while the bottom row shows embedded fetal US Test<sub>seen</sub> data. (a, c) are the results of using encoded features for relevant tasks, e.g.  $Z_A$  for  $T_A$  and  $Z_B$  for  $T_B$ ; separated clusters are desirable here. (b, d) are the results of using encoded features for irrelevant tasks, namely,  $Z_A$  for  $T_B$  and  $Z_B$  for  $T_A$ ; mixed clusters are desirable in this case.

bottom row in Fig. 3 shows that encoded features are more capable of classifying class identities in the relevant task than the irrelevant task (e.g. (a) vs. (d)).

**Discussion:** Acoustic shadows are caused by anatomies which block the propagation of sound waves or by destructive interference. With this dependency between anatomy and artifacts, separating shadow features from anatomical features may lead to decreased performance of artifacts classification (Table 2, Col. 7, Proposed). However, this separation enables feature generalization so that the model is less limited to certain image formation and able to tackle new combinations of entangled properties (Table 2, Col. 9, Proposed).



Generalization of supervised neural networks can also be achieved by extensive data collection across domains and in a limited way by artificial data augmentation. Here, we propose an alternative through feature disentanglement, which requires less data collection and training effort. Figure 3 shows PCA plots for the penultimate dense layer. Observing entanglement in earlier layers reveals that disentanglement occurs in this very last layer. This is due to the definition of our loss functions and is partly influenced by the dense layers interpreting the latent representation for classification. Finally, perfect representation disentanglement is likely infeasible because image features are rarely totally isolated in reality. In this paper we have shown that even imperfect disentanglement is able to provide great benefits for artifact-prone image classification in medical image analysis.

## 4 Conclusion

In this paper, we propose a novel disentanglement method to extract generalizable features within a multi-task framework. In the proposed method, classification tasks lead to encoded features that are maximally informative with respect to these tasks while the adversarial regularization forces these features to be minimally informative about irrelevant tasks, which disentangles internal representations. Experimental results on synthetic and fetal US data show that our method outperforms baseline methods for multiple tasks, especially on images with entangled properties that are unseen during training. Future work will explore the extension of this framework to multiple tasks beyond classification.

**Acknowledgments.** We thank the Wellcome Trust IEH Award [102431], Nvidia (GPU donations) and Intel.

## References

1. Ben-Cohen, A., Mechrez, R., Yedidia, N., Greenspan, H.: Improving CNN training using disentanglement for liver lesion classification in CT. [arXiv:1811.00501](#) (2018)
2. Bengio, Y., Courville, A., Vincent, P.: Representation learning: a review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.* **35**(8), 1798–1828 (2013)
3. Burgess, C.P., et al.: Understanding disentangling in  $\beta$ -VAE. [arXiv:1804.03599](#) (2018)
4. Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., Abbeel, P.: InfoGAN: interpretable representation learning by information maximizing generative adversarial nets. In: *NeurIPS 2016*, pp. 2180–2188. Curran Associates Inc., USA (2016)
5. Chen, X., et al.: Variational lossy autoencoder. In: *ICLR 2017* (2017)
6. Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W.: ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. [arXiv:1811.12231](#) (2018)
7. Gonzalez-Garcia, A., van de Weijer, J., Bengio, Y.: Image-to-image translation for cross-domain disentanglement. In: *NeurIPS 2018*, pp. 1287–1298. Curran Associates, Inc. (2018)



8. Hadad, N., Wolf, L., Shahar, M.: A two-step disentanglement method. In: CVPR 2018 (2018)
9. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., Lerchner, A.: beta-VAE: learning basic visual concepts with a constrained variational framework. In: ICLR 2017 (2017)
10. Hyvärinen, A., Oja, E.: Independent component analysis: algorithms and applications. *Neural Netw.* **13**(4–5), 411–430 (2000)
11. Kamnitsas, K., et al.: Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. In: Niethammer, M., et al. (eds.) IPMI 2017. LNCS, vol. 10265, pp. 597–609. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-59050-9\\_47](https://doi.org/10.1007/978-3-319-59050-9_47)
12. Kim, H., Mnih, A.: Disentangling by factorising. CoRR arXiv/1802.05983 (2018)
13. Liu, A.H., Liu, Y.C., Yeh, Y.Y., Wang, Y.C.F.: A unified feature disentangler for multi-domain image translation and manipulation. In: NeurIPS, pp. 2590–2599. Curran Associates, Inc. (2018)
14. Mathieu, M.F., Zhao, J.J., Zhao, J., Ramesh, A., Sprechmann, P., LeCun, Y.: Disentangling factors of variation in deep representations using adversarial training. In: NeurIPS 2016, pp. 5040–5048 (2016)
15. Meng, Q., et al.: Weakly supervised estimation of shadow confidence maps in fetal ultrasound imaging. *IEEE Trans. Med. Imaging* (2019). <https://ieeexplore.ieee.org/document/8698843>
16. NHS: Fetal anomaly screening programme: programme handbook June 2015. Public Health England (2015)
17. Pawlowski, N., et al.: DLTK: state of the art reference implementations for deep learning on medical images. [arXiv:1711.06853](https://arxiv.org/abs/1711.06853) (2017)
18. Tenenbaum, J.B., Freeman, W.T.: Separating style and content with bilinear models. *Neural Comput.* **12**(6), 1247–1283 (2000)