



Published in final edited form as:

IEEE Trans Med Imaging. 2021 October ; 40(10): 2857–2868. doi:10.1109/TMI.2021.3060634.

Transferable Visual Words: Exploiting the Semantics of Anatomical Patterns for Self-supervised Learning

Fatemeh Haghighi,

School of Computing, Informatics, and Decision Systems Engineering, Arizona State University, Tempe, AZ 85281 USA

Mohammad Reza Hosseinzadeh Taher,

School of Computing, Informatics, and Decision Systems Engineering, Arizona State University, Tempe, AZ 85281 USA

Zongwei Zhou [Member, IEEE],

Department of Biomedical Informatics, Arizona State University, Scottsdale, AZ 85259 USA

Michael B. Gotway,

Radiology Department, Mayo Clinic, Scottsdale, AZ 85259 USA

Jianming Liang [Senior Member, IEEE]

Department of Biomedical Informatics, Arizona State University, Scottsdale, AZ 85259 USA

Abstract

This paper introduces a new concept called “transferable visual words” (TransVW), aiming to achieve annotation efficiency for deep learning in medical image analysis. Medical imaging—focusing on particular parts of the body for defined clinical purposes—generates images of great similarity in anatomy across patients and yields sophisticated anatomical patterns across images, which are associated with rich *semantics* about human anatomy and which are natural *visual words*. We show that these visual words can be automatically harvested according to anatomical consistency via self-discovery, and that the self-discovered visual words can serve as strong yet free supervision signals for deep models to learn semantics-enriched generic image representation via self-supervision (self-classification and self-restoration). Our extensive experiments demonstrate the annotation efficiency of TransVW by offering higher performance and faster convergence with reduced annotation cost in several applications. Our TransVW has several important advantages, including (1) TransVW is a fully autodidactic scheme, which exploits the semantics of visual words for self-supervised learning, requiring no expert annotation; (2) visual word learning is an add-on strategy, which complements existing self-supervised methods, boosting their performance; and (3) the learned image representation is semantics-enriched models, which have proven to be more robust and generalizable, saving annotation efforts for a variety of applications through transfer learning. Our code, pre-trained models, and curated visual words are available at <https://github.com/JLiangLab/TransVW>.

Keywords

Self-supervised learning; transfer learning; visual words; anatomical patterns; computational anatomy; 3D medical imaging; 3D pre-trained models

I. Introduction

A grand promise of computer vision is to learn general-purpose image representation, seeking to automatically discover generalizable knowledge from data in either a supervised or unsupervised manner and transferring the discovered knowledge to a variety of applications for performance improvement and annotation efficiency.

In the literature, convolutional neural networks (CNNs) and bags of visual words (BoVW) [1] are often presumed to be competing methods, but they actually offer complementary strengths. Training CNNs requires a large number of annotated images, but the learned features are transferable to many applications. Extracting visual words in BoVW, on the other hand, is unsupervised in nature, demanding no expert annotation, but the extracted visual words lack transferability. Therefore, the first question that we seek to answer is *how to beneficially integrate the transfer learning capability of CNNs with the unsupervised nature of BoVW in extracting visual words for image representation learning?*

In the meantime, medical imaging protocols, typically designed for specific clinical purposes by focusing on particular parts of the body, generate images of great similarity in anatomy across patients and yield an abundance of sophisticated anatomical patterns across images. These anatomical patterns are naturally associated with rich semantic knowledge about human anatomy. Therefore, the second question that we seek to answer is *how to exploit the deep semantics associated with anatomical patterns embedded in medical images to enrich image representation learning?*

In addressing the two questions simultaneously, we have conceived a new concept: **transferable visual words (TransVW)**, where the sophisticated anatomical patterns across medical images are natural “visual words” associated with deep semantics in human anatomy (see Fig. 1). These anatomical visual words can be automatically harvested from unlabeled medical images and serve as strong yet free supervision signals for CNNs to learn semantics-enriched representation via self-supervision. The learned representation is generalizable and transferable because it is not biased to the idiosyncrasies of pre-training (pretext) tasks and datasets, thereby it can produce more powerful models to solve application-specific (target) tasks via transfer learning. As illustrated in Fig. 2, our TransVW consists of three components: (1) a novel self-discovery scheme that automatically harvests anatomical visual words, exhibiting consistency (in pattern appearances and semantics) and diversity (in organ shapes, boundaries, and textures), directly from unlabeled medical images and assigns each of them with a unique label that bears the semantics associated with a particular part of the body; (2) a unique self-classification scheme that compels the model to learn semantics from anatomical consistency within visual words; and (3) a scalable self-restoration scheme that encourages the model to encode anatomical diversity within visual words.

Our extensive experiments demonstrate the annotation efficiency of TransVW in higher performance, faster convergence, and less annotation cost on the applications where there is a dearth of annotated images. Compared with existing publicly available models, pre-trained by either self-supervision or full-supervision, our TransVW offers several advantages, including (1) TransVW is a fully autodidactic scheme, which exploits the semantics of visual words for self-supervised learning, requiring no expert annotation; (2) visual word learning is an add-on strategy, which complements existing self-supervised methods, boosting their performance; and (3) the learned image representation is semantics-enriched models, which have proven to be more robust and generalizable, saving annotation efforts for a variety of applications through transfer learning. In summary, we make the following contributions:

- An unsupervised clustering strategy, curating a dataset of anatomical visual words from unlabeled medical images.
- An add-on learning scheme, enriching representations learned from existing self-supervised methods [2]-[5].
- An advanced self-supervised framework, elevating transfer learning performance, accelerating training speed, and reducing annotation efforts.

II. Transferable Visual Words

TransVW aims to learn transferable and generalizable image representation by leveraging the semantics associated with the anatomical patterns embedded in medical images (see Fig. 1). For clarity, as illustrated in Fig. 2a, we define a *visual word* as a segment of consistent anatomy recurring across images and the *instances of a visual word* as the patches extracted across different but resembling images for this visual word. Naturally, all instances of a visual word exhibit great similarity in both appearance and semantics. Furthermore, to reflect the semantics of its corresponding anatomical parts, a unique (pseudo) label is automatically assigned to each visual word during the self-discovery process (see Section II-A); consequently, all instances of a visual word share the same label bearing the same semantics in anatomy.

A. TransVW learns semantics-enriched representation

TransVW enriches representation learning with the semantics of visual words, as shown in Fig. 2, through the following three components.

1) Self-discovery—harvesting the semantics of anatomical patterns to form visual words.—The self-discovery component aims to automatically extract a set of C visual words from unlabeled medical images as shown in Fig. 2a. To ensure a high degree of consistency in anatomy among the instances of each visual word, we first identify a set of K patients that display a great similarity in their overall image appearance. To do so, we use the whole patient scans in the training dataset to train a feature extractor $\widehat{\Phi}(\cdot)$, an auto-encoder network, to learn an identical mapping from a whole-patient scan to itself. Once trained, its latent features can be used as an indicator of the similarity among patient scans. As a result, we can form such a set of K patient scans by randomly anchoring a

patient scan as the reference and appending its top $K-1$ nearest neighbors found throughout the entire training dataset based on the $L2$ distance in the latent feature space. Given the great resemblance among the selected K patient scans, the patches extracted at the same coordinate across these K scans are expected to exhibit a high degree of similarity in anatomical patterns. Therefore, for each visual word, we extract its K instances by cropping around a random but fixed coordinate across a set of selected K patients, and assign a unique pseudo label to them. We repeat this process C times, yielding a set of C visual words, each with K instances, extracted from C random coordinates. The extracted visual words are naturally associated with the semantics of the corresponding human body parts. For example, in Fig. 2a, four visual words are defined randomly in a reference patient (top-left), bearing local anatomical information of (1) anterior ribs 2–4, (2) spinous processes, (3) right pulmonary artery, and (4) LV. In summary, our self-discovery automatically generates a dataset of visual words associated with semantic labels, as a free and rich source for training deep models to learn semantics-enriched representations from unlabeled medical images

2) Self-classification—learning the semantics of anatomical consistency from visual words.—Once a set of visual words are self-discovered, representation learning can be formulated as self-classification, a C -way multi-class classification that discriminates visual words based on their semantic (pseudo) labels. As illustrated in Fig. 2b, the self-classification branch is composed of (1) an encoder that projects visual words into a latent space, and (2) a classification head, a sequence of fully-connected layers, that predicts the pseudo label of visual words. It is trained by minimizing the standard categorical cross-entropy loss function:

$$\mathcal{L}_{cls} = -\frac{1}{B} \sum_{b=1}^B \sum_{c=1}^C \mathcal{Y}_{bc} \log \mathcal{P}_{bc} \quad (1)$$

where B denotes the batch size; C denotes the number of visual words; \mathcal{Y} and \mathcal{P} represent the ground truth (one-hot pseudo label vector) and the network prediction, respectively. Through training, the model is compelled to learn features that distinguish the anatomical dissimilarity among instances belonging to different visual words and that recognize the anatomical resemblance among instances belonging to the same visual words, resulting in image representations associated with the semantics of anatomical patterns underneath medical images. Therefore, self-classification is for learning image representation enriched with semantics that can pull together all instances of each visual word, while pushing apart instances of different visual words.

3) Self-restoration—encoding the semantics of anatomical diversity of visual words.—In self-discovery, we intentionally selected patients, rather than patches, according to their resemblance at the whole patient level. Given that no scans of different patients are the same in appearance, behind their great similarity, the instances of a visual word are also expected to display subtle anatomical diversity in terms of organ shapes, boundaries, and texture. Such a balance between the consistency and diversity of anatomical patterns for each visual word is critical for deep models to learn robust image representation. To encode this anatomical diversity within visual words for image representation learning, we

augment our framework with self-restoration, training the model to restore the original visual words from the perturbed ones. The self-restoration branch, as shown in Fig. 2c, is an encoder-decoder with skip connections in between. We apply a perturbation operator $g(\cdot)$ on a visual word x to get the perturbed visual word $\tilde{x} = g(x)$. The encoder takes the perturbed visual word \tilde{x} as an input and generates a latent representation. The decoder takes the latent representation from the encoder and decodes it to produce the original visual word. Our perturbation operator $g(\cdot)$ consists of non-linear, local-shuffling, out-painting, and inpainting transformations, which are proposed by Zhou *et al.* [5], [6], as well as identity mapping (*i.e.*, $x = g(x)$). Restoring visual word instances from their perturbations enables the model to learn image representation from multiple perspectives [5], [6]. The restoration branch is trained by minimizing the $L2$ distance between the original and reconstructed visual words:

$$\mathcal{L}_{rec} = \frac{1}{B} \sum_{i=1}^B \|x_i - x'_i\|_2 \quad (2)$$

where B denotes the batch size, x and x' represent the original visual word and the reconstructed prediction, respectively.

To enable end-to-end representation learning from multiple sources of information and yield more powerful models for a variety of medical tasks, in TransVW, self-classification and self-restoration are integrated together by sharing the encoder and jointly trained with one single objective function:

$$\mathcal{L} = \lambda_{cls} \mathcal{L}_{cls} + \lambda_{rec} \mathcal{L}_{rec} \quad (3)$$

where λ_{cls} and λ_{rec} adjust the weights of classification and restoration losses, respectively. Our *unique* definition of \mathcal{L}_{cls} empowers the model to learn the common anatomical semantics across medical images from a strong discriminative signal—the semantic label of visual words. The definition of \mathcal{L}_{rec} equips the model to learn the anatomical finer details of visual words from multiple perspectives by restoring original visual words from varying image perturbations. We should emphasize that our goal is *not* simply discovering, classifying, and restoring visual words *per se*, rather advocating it as a holistic pre-training scheme for learning semantics-enriched image representation, whose usefulness must be assessed objectively based on its generalizability and transferability to various target tasks.

B. TransVW has several unique properties

1) Autodidactic—exploiting semantics in unlabeled data for self supervision.—Due to the lack of sufficiently large, curated, and labeled medical datasets, self-supervised learning holds a great promise for representation learning in medical imaging because it does not require manual annotation for pre-training. Unlike existing self-supervised methods for medical imaging [4], [5], [14], our TransVW explicitly employs the (pseudo) labels that bear the semantics associated with the sophisticated anatomical patterns embedded in the unlabeled images to learn more pronounced representation for medical applications. Particularly, TransVW benefits from a large, diverse set of anatomical visual words discovered by our self-discovery process, coupled with a training scheme integrating

both self-classification and self-restoration, to learn semantics-enriched representation from unlabeled medical images. With *zero* annotation cost, our TransVW not only outperforms other self-supervised methods but also surpasses publicly-available, fully-supervised pre-trained models, such as I3D [15], NiftyNet [16], and MedicalNet [17].

2) Comprehensive—blending consistency with diversity for semantics richness.—Our self-discovery component secures both consistency and diversity within each visual word, thereby offering a lucrative source for pre-training deep models. More specifically, our self-discovery process computes similarity at the patient level and selects the top nearest neighbors of the reference patient. Extracting visual word instances from these similar patients, based on random but fixed coordinates, strikes a balance between consistency and diversity of the anatomical pattern within each visual word. Consequently, our self-classification exploits the semantic consistency by classifying visual words according to their pseudo labels, resulting in class-level feature separation among visual word classes. Furthermore, our self-restoration leverages the fine-grained anatomical information—the subtle diversity of intensity, shape, boundary, and texture enables instance-level feature separation among instances within each visual word. As a result, our TransVW projects visual words into more comprehensive feature space in both class-level and instance-level by blending consistency with diversity.

3) Robust—preventing superficial solutions for deep representation.—Self-supervised learning is notorious for learning shortcut solutions in tackling pretext tasks, leading to less generalizable image representations [18], [19]. However, our method, especially our self-classification component, is discouraged from learning superficial solutions since our self-discovery process imposed substantial diversity among instances of each visual word. Furthermore, we follow two well-known techniques to further improve the diversity of visual words. Firstly, following Doersch *et al.* [19] and Mundhenk *et al.* [20], we extract multi-scale instances for each visual word within a patient, in which each instance is randomly jittered by a few pixels. Consequently, having various scale instances in each class enforces self-classification to perform more semantic reasoning by preventing easy matching of simple features among the instances of the same class [20]. Secondly, during pre-training, we augment visual words with various image perturbations to increase the diversity of data. Altogether, the substantial diversity of visual words coupled with various image perturbations enforce our pretext task to capture semantic-bearing features, resulting in a compelling and robust representation obtained from anatomical visual words.

4) Versatile—complementing existing self-supervised methods for performance enhancement.—TransVW boasts an innovative add-on capability, a versatile feature that is unavailable in other self-supervised learning methods. Unlike existing self-supervised methods that build supervision merely from the information within individual images of training data, our self-discovery and self-classification leverage the anatomical similarities present across different images (reflected in visual words) to learn common anatomical semantics. Consequently, incorporating visual word learning into existing self-supervised methods enforces them to encode semantic structures of visual words into their learned embedding space, resulting in more versatile representations.

Therefore, our self-discovery and self-classification components together can serve as an add-on to boost existing self-supervised methods, as evidenced in Fig. 3.

III. Experiments

A. Pre-training TransVW

TransVW models are pre-trained solely on unlabeled images. Nevertheless, to ensure no test-case leaks from pretext tasks to target tasks, any images that will be used for validation or test in target tasks (see Table I) are excluded from pre-training. We have released two pre-trained models: (1) TransVW Chest CT in 3D, which was pre-trained from scratch using 623 chest CT scans in LUNA16 [7] (the same as the publicly released Models Genesis¹), and (2) TransVW Chest X-rays in 2D, which was pre-trained on 76K chest X-ray images in ChestX-ray14 [12] datasets. We set $C=45$ for TransVW Chest CT and $C=100$ for TransVW Chest X-rays (see Appendix I-B for the ablation studies on the impact of the number of visual words on performance). We empirically set $K=200$ and $K=1000$ for TransVW Chest CT and TransVW Chest X-rays, respectively, to strike a balance between diversity and consistency of the visual words. For each instance of a visual word, multi-scale cubes/patches for 3D/2D images are cropped, and then all are resized to $64\times 64\times 32$ and 224×224 for TransVW Chest CT and TransVW Chest X-rays, respectively (see Appendix I-C for samples of the discovered visual words).

B. Fine-tuning TransVW

The pre-trained TransVW can be used for a variety of target tasks through transfer learning. To do so, we utilize (1) the encoder for target classification tasks by appending a target task-specific classification head, and (2) the encoder and decoder for target segmentation tasks by replacing the last layer with a $1\times 1\times 1$ convolutional layer. In this study, we have evaluated the generalization and transferability of TransVW by fine-tuning all the parameters of target models on seven diverse target tasks, including image classification and segmentation tasks in both 2D and 3D. As summarized in Table I and detailed in Appendix I-A, these target tasks offer the following two advantages:

- **Covering a wide range of diseases, organs, and modalities.** This allows us to verify the add-on capability of TransVW on various 3D target tasks, covering a diverse range of diseases (*e.g.*, nodule, embolism, tumor), organs (*e.g.*, lung, liver, brain), and modalities (*e.g.*, CT and MRI). It also enables us to verify the generalizability of TransVW in not only the target tasks on the same dataset as pre-training (NCC and NCS), but also target tasks with a variety of domain shifts (ECC, LCS, and BMS) in terms of modality, scan regions, or dataset. To our best knowledge, we are among the first to investigate cross-domain self-supervised learning in medical imaging.
- **Enjoying a sufficient amount of annotation.** This paves the way for conducting annotation reduction experiments to verify the annotation efficiency of TransVW.

¹Models Genesis: github.com/MrGiovanni/ModelsGenesis

C. Benchmarking TransVW

For a thorough evaluation, in addition to the training from scratch (the lower-bound baseline), TransVW is compared with a whole range of transfer learning baselines, including both self-supervised and (fully-)supervised methods.

Self-supervised baselines: We compare TransVW with Models Genesis [5], the state-of-the-art self-supervised learning method for 3D medical imaging, as well as Rubik’s cube [14], the most recent multi-task self-supervised learning method for 3D medical imaging. Since most self-supervised learning methods are initially proposed in the context of 2D images, we also have extended three most representative ones [4], [3], [2] into their 3D version for a fair comparison.

Supervised baselines: We also examine publicly available *fully-supervised* pre-trained models for 3D transfer learning in medical imaging, including NiftyNet [16] and Medical-Net [17]. Moreover, we fine-tune Inflated 3D (I3D) [15] in our 3D target tasks since it has been successfully utilized by Ardila *et al.* [22] to initialize 3D models for lung nodule detection.

We utilize 3D U-Net² for 3D applications, and U-Net³ with ResNet-18 as the backbone for 2D applications. For our pretext task, we modify those architectures by appending fully-connected layers to the end of the encoders for the classification head. In pretext tasks, we set the weights of losses as $\lambda_{rec} = 1$ and $\lambda_{cls} = 0.01$. All the pretext tasks are trained using Adam optimizer, with a learning rate of 0.001, where $\beta_1 = 0.9$ and $\beta_2 = 0.999$. Regular data augmentation techniques including random flipping, transposing, rotating, elastic transformation, and adding Gaussian noise are utilized in target tasks. We use the early-stop technique on the validation set to prevent over-fitting. We run each method ten times on each target task and report the average, standard deviation, and further provide statistical analyses based on independent two-sample *t*-test.

IV. Results

This section presents the cornerstones of our results, demonstrating the significance of our self-supervised learning framework. We first integrate our two novel components, self-discovery and self-classification of visual words, into four popular self-supervised methods [2]-[5], suggesting that these two components can be adopted as an add-on to enhance the existing self-supervised methods. We then compare our TransVW with the current state-of-the-art approaches in a triplet of aspects: transfer learning performance, convergence speedup, and annotation cost reduction, concluding that TransVW is an annotation-efficient method for medical image analysis.

A. TransVW is an add-on scheme

Our self-discovery and self-classification components can readily serve as an add-on to enrich existing self-supervised learning approaches. In fact, by introducing these two

²3D U-Net: github.com/ellisdg/3DUnetCNN

³Segmentation Models: github.com/qubvel/segmentation_models

components, we open the door for the existing image-based self-supervision approaches to capture a more high-level and diverse visual representation that reduces the gap between pre-training and semantic transfer learning tasks.

Experimental setup: To study the add-on capability of our self-discovery and self-classification components, we incorporate them into four representative self-supervised methods, including (1) Models Genesis [5], which restores the original image patches from the transformed ones; (2) Inpainting [3], which predicts the missing parts of the input images; (3) Context restoration [4], which restores the original images from the distorted ones that are obtained by shuffling small patches within the images; and (4) Rotation [2], which predicts the rotation angles that are applied to the input images. Since all the reconstruction-based self-supervised methods under study [3]-[5] utilize encoder-decoder architecture with skip connections in between, we append additional fully-connected layers to the end of the encoder, enabling models to learn image representation simultaneously from classification and restoration tasks. For Rotation, the network only includes an encoder, followed by two classification heads to learn representations from rotation angle prediction as well as the visual words classification tasks. Note that the original self-supervised methods in [2], [3], and [4] were implemented in 2D, but we have extended them into 3D.

Observations: Our results in Fig. 3 demonstrate that incorporating visual words with existing self-supervised methods consistently improves their performance across five 3D target tasks. Specifically, visual words significantly improve Rotation by 3%, 1.5%, 2%, 1.5%, and 5%; Context restoration by 1.75%, 1%, 5%, 2%, and 8%; Inpainting by 2%, 1.5%, 3%, 3%, and 6% in NCC, NCS, ECC, LCS, and BMS applications, respectively. Moreover, visual words significantly advance Models Genesis by 1%, 1.5%, and 1% in NCC, LCS, and BMS, respectively.

Discussion: How can TransVW improve representation learning?—Most existing self-supervised learning methods, such as predicting contexts [3]-[5], [27] or discriminating image transformations [2], [14], [18], [28], concentrate on learning the visual representation of each image individually, thereby, overlooking the notion of semantic similarities across different images [29], [30]. In contrast, appreciating the recurrent anatomical structure in medical images, our self-discovery can extract meaningful visual words across different images and assign unique semantic pseudo labels to them. By classifying the resultant visual words according to their pseudo labels, our self-classification is enforced to explicitly recognize these visual words across different images grouping similar visual words together while separating dissimilar ones apart. Consequently, integrating our two novel components into existing self-supervised methods empowers the model to not only learn the local context within a single image but also learn the semantic similarities of the consistent and recurring visual words across images. It is worth noting that self-discovery and self-classification should be considered a significant add-on in terms of methodology. As shown in Fig. 3, simply adding these two components on top of four popular self-supervised methods can noticeably improve their fine-tuning performance.

B. TransVW is an annotation-efficient method

To ensure that our method addresses annotation scarcity challenge in medical imaging, we first precisely define the annotation-efficiency term. We consider a method as **annotation-efficient** if (1) it achieves superior performance using the same amount of annotated training data, (2) it reduces the training time using the same amount of annotated data, or (3) it offers the same performance but requires less annotated training data. Based on this definition, we adopt a rigorous **three-pronged** approach in evaluating our work, by considering not only the transfer learning performance but also the acceleration of the training process and label efficiency on a variety of target tasks, demonstrating that TransVW is an annotation-efficient method.

1) TransVW provides superior transfer learning performance: A generic pre-trained model transfers well to many different target tasks, indicated by considerable performance improvements. Thus, we first evaluate the generalizability of TransVW in terms of improving the performance of various medical tasks across diseases, organs, and modalities.

Experimental setup: We fine-tune TransVW on five 3D target tasks, as described in Table I, covering classification and segmentation. We investigate the generalizability of TransVW not only in the target tasks with the pre-training dataset (NCC and NCS), but also in the target tasks with a variety of domain shifts (ECC, LCS, and BMS).

Observations: Our evaluations in Table II suggest three major results. Firstly, TransVW significantly outperforms training from scratch in all five target tasks under study by a large margin and also stabilizes the overall performance.

Secondly, TransVW surpasses all self-supervised counterparts in the five target tasks. Specifically, TransVW significantly outperforms Models Genesis, state-of-the-art self-supervised 3D models pre-trained using image restoration, in three applications, *i.e.*, NCC, LCS, and BMS, and offers equivalent performance in NCS and ECC. Moreover, TransVW yields remarkable improvements over Rubik's cube, the most recent 3D multi-task self-supervised method, in all five applications. Particularly, Rubik's cube formulates a multi-task learning objective solely based on contextual cues within single images while TransVW benefits from semantic supervision of anatomical visual words, resulting in more enhanced representations.

Thirdly, TransVW achieves superior performance in comparison with publicly available fully-supervised pre-trained 3D models, *i.e.*, NiftyNet, MedicalNet, and I3D, in all five target tasks. It is noteworthy that our TransVW does not solely depend on the architecture capacity to achieve the best performance since it has much fewer model parameters than its counterparts. Specifically, TransVW is trained on basic 3D U-Net with 23M parameters, while MedicalNet with ResNet-101 as the backbone (reported as the best performing model in [17]) carries 85.75M parameters, and I3D contains 25.35M parameters in the encoder. Although NiftyNet model is offered with 2.6M parameters, its performance is not as good as its supervised counterparts in any of the target tasks.

Interestingly, although TransVW is pre-trained on chest CT scans, it is still beneficial for different organs, diseases, datasets, and even modalities. In particular, the pulmonary embolism false positive reduction (ECC) is on Contrast-Enhanced CT scans, which may appear differently from the normal CT scans that are used for pre-training; yet, according to Table II, TransVW obtains a 7% improvement over training 3D models from scratch in this task. Additionally, fine-tuning from TransVW provides a substantial gain in liver segmentation (LCS) accuracy despite the noticeable differences between pretext and target domains in terms of organs (lung vs. liver) and datasets (LUNA 2016 vs. LiTS 2017). We further examine the transferability of TransVW in brain tumor segmentation on MRI Flair images (BMS). Referring to Table II, despite the marked differences in organs, datasets, and even modalities between the pretext and BMS target task, we still observe a significant performance boost from fine-tuning TransVW in comparison with learning from scratch. Moreover, TransVW significantly outperforms Models Genesis in the most distant target domains from the pretext task, *i.e.*, LCS, and BMS.

Discussion: How can TransVW improve the performance of cross-domain target tasks?: Learning universal representations that can transfer effectively to a wide range of target tasks is one of the supreme goals of computer vision in medical imaging. The best known existing examples of such representations are pre-trained models on ImageNet dataset. Although there are marked differences between natural and medical images, the image representations learned from ImageNet can be beneficial not only for natural imaging [31] but also for medical imaging [32], [33]. Therefore, rather than developing pre-trained models specifically for a particular dataset/organ, TransVW aims to develop generic pre-trained 3D models for medical image analysis that are not biased to idiosyncrasies of the pre-training task and dataset and generalize effectively across organs, datasets, and modalities.

As is well-known, CNNs trained on large scale visual data form feature hierarchies; lower layers of deep networks are in charge of general features while higher layers contain more specialized features for target domains [33]-[35]. Due to generalizability of low and mid-level features, they can lead to significant benefits of transfer learning, even when there is a substantial domain gap between the pretext and target task [33], [35]. Therefore, while TransVW is pre-trained solely on chest CT scans, it still elevates the target task performance in different organs, datasets, and modalities since its low and mid-level features can be reused in the target tasks.

2) TransVW accelerates the training process: Although accelerating the training of deep neural networks is arguably an influential line of research [36], its importance is often underappreciated in the medical imaging literature. Transfer learning provides a warm-up initialization that enables target models to converge faster and mitigates the vanishing and exploding gradient problems. In that respect, we argue that a good pre-trained model should yield better target task performance with less training time. Hence, we further evaluate TransVW in terms of accelerating the training process of various medical tasks.

Experimental setup: We compare the convergence speedups obtained by TransVW with training from scratch (the lower bound baseline) and Models Genesis (the state-of-the-art

baseline) in our five 3D target tasks. For conducting fair comparisons in all experiments, all methods benefit from the same data augmentation and use the same network architecture while we endeavor to optimize each model with the bestperforming hyper-parameters.

Observations: Fig. 4 presents the learning curves for training from scratch, fine-tuning from Models Genesis, and fine-tuning from TransVW. Our results demonstrate that initializing 3D models from TransVW remarkably accelerates the training process of target models in comparison with not only learning from scratch but also Models Genesis in all five target tasks. These results imply that TransVW captures representations that are more aligned with the subsequent target tasks, leading to faster convergence of the target models.

Putting the transfer learning performance on five target tasks in Table II and the training times in Fig. 4 together, TransVW demonstrates significantly better or equivalent performance with remarkably less training time in comparison to its 3D counterparts. Specifically, TransVW significantly outperforms Models Genesis in terms of both performance and saving training time in three out of five applications, *i.e.*, NCC, LCS, and BMS, and achieves equivalent performance in NCS and ECC but in remarkably less time. Altogether, we believe that TransVW can serve as a primary source of transfer learning for 3D medical imaging applications to boost the performance and accelerate the training of target tasks.

3) TransVW reduces the annotation cost: Transfer learning yields more accurate models by reusing the previously learned knowledge in target tasks with limited annotations. This is because a good representation should not need many samples to learn about a concept [21]. Thereby, we conduct experiments on partially labeled data to investigate transferability of TransVW in small data regimes.

Experimental setup: We compare the transfer learning performance of TransVW using partial labeled data with training from scratch and fine-tuning from Models Genesis in five 3D target tasks. For clarity, we determine the minimum required data for training from scratch and fine-tuning Models Genesis to meet the comparable performance (based on independent two-sample *t*-test) when training using the entire training data. Moreover, we investigate the minimum required data for TransVW to meet the equivalent performance that training from scratch and fine-tuning Models Genesis can achieve.

Observations: Fig. 5 illustrates the results of using the partial amount of labeled data during the training of five 3D target tasks. As an illustrative example, in lung nodule false positive reduction (NCC), our results demonstrate that using only 35% of training data, TransVW achieves equivalent performance to training from scratch using 70% of data. Therefore, around 50% of the annotation cost in NCC can be reduced by fine-tuning models from TransVW compared with training from scratch. In comparison with Models Genesis in the same application (NCC), TransVW with 75% of data achieves equal performance with Models Genesis using 90% of data. Therefore, about 17% of the annotation cost associated with fine-tuning from Models Genesis in NCC is recovered by fine-tuning from TransVW. In general, transfer learning from TransVW reduces the annotation cost by 50%, 50%, 57%, 60%, and 80% in comparison with training from scratch in NCC, NCS, ECC,

LCS, and BMS applications, respectively. In comparison with Models Genesis, TransVW reduces the annotation efforts by 17%, 24%, and 50% in NCC, LCS, and BMS applications, respectively, and both models performs equally in the NCS and ECC applications. These results suggest that TransVW achieves state-of-the-art or comparable performance over other self-supervised approaches while being more efficient, *i.e.*, less annotated data is required for training high-performance target models.

Summary: We demonstrate that TransVW provides more generic and transferable representations compared with self-supervised and supervised 3D competitors, confirmed by our evaluations on a triplet of transfer learning performance, optimization speedup, and annotation cost. To further illustrate the effectiveness of our framework, we adopt TransVW to the nnU-Net [26], a state-of-the-art segmentation framework in medical imaging, and evaluate it on liver tumor segmentation task from the Medical Segmentation Decathlon [37]. Our results demonstrate that TransVW obtains improvements in segmentation accuracy over training from scratch and Models Genesis by 2.5% and 1%, respectively (details in Appendix I-E). These results, in line with our previous results, reinforce our main insight that TransVW provides an annotation-efficient solution for 3D medical imaging.

V. Ablation experiments

In this section, we first conduct ablation experiments to illustrate the contribution of different components to the performance of TransVW. We then train a 2D model using chest X-ray images, called TransVW 2D, and compare it with state-of-the-art 2D models. In the 2D experiments, we consider three target tasks: thorax diseases classification (DXC), pneumothorax segmentation (PXS), and lung nodule false positive reduction (NCC). We evaluate NCC in a 2D slice-based solution, where the 2D representation is obtained by extracting axial slices from the volumetric dataset.

A. Comparing individual self-supervised tasks

Our TransVW takes the advantages of two sources in representation learning: self-classification and self-restoration of visual words. Therefore, we first directly compare the two isolated tasks and then investigate whether joint-task learning in TransVW produces more transferable features compared with isolated training schemes. In Table II, the last three rows show the transfer learning results of TransVW and each of the individual tasks in five 3D applications. According to the statistical analysis results, self-restoration and self-classification reveal no significant difference (p -value > 0.05) in three target tasks, NCS, ECC, and LCS, and self-restoration achieves significantly better performance in NCC and BMS. Despite the success of self-restoration in encoding fine-grained anatomical information from individual visual words, it neglects the semantic relationships across different visual words. In contrast, our novel self-classification component explicitly encodes the semantic similarities that presents across visual words into the learned embedding space. Therefore, as evidenced by Table II, integrating self-classification and self-restoration into a single framework yields a more comprehensive representation that can guarantee the highest target task performance. In particular, TransVW outperforms each

isolated task in four applications, *i.e.*, NCC, ECC, LCC, and BMS, and provides comparable performance with self-restoration in NCS.

B. Evaluating 2D applications

We evaluate our TransVW 2D with Models Genesis 2D (self-supervised) and ImageNet (fully-supervised) pre-trained models in two experimental settings: (1) linear evaluation on top of the fixed features from the pre-trained network, and (2) full fine-tuning of the pre-trained network for target tasks.

Linear evaluation: To evaluate the quality of the learned representations, we follow the common practice in [21], [29], which trains linear classifiers on top of the fixed features obtained from various layers of the pre-trained networks. Specifically, for ResNet-18 backbone, we extract image features from the last layer of every residual stage (denoted as res1, res2, etc) [21], and then evaluate them in two classification target tasks (DXC and NCC). Based on our results in Fig. 6a, TransVW 2D representations are transferred better across all the layers on both target tasks in comparison with Models Genesis 2D and ImageNet, demonstrating the generalizability of TransVW 2D representations. Specifically, in thorax diseases classification (DXC), which is in the same dataset as the pretext task, the best performing features are extracted from res4 in the last layer of the TransVW 2D network. This indicates that TransVW 2D encourages the models to squeeze out high-level representations, which are aligned with the target task, in the deeper layers of the network. Moreover, in lung nodule false positive reduction (NCC), which presents a domain shift compared with the pretext task, TransVW 2D remarkably reduces the performance gap between res3 and res4 features compared with Models Genesis 2D and ImageNet. This suggests that TransVW reduces the overfitting of res4 features to the pretext task and dataset, resulting in more generic features.

Full fine-tuning: We evaluate the initialization provided by our TransVW 2D via fine-tuning it for three 2D target tasks, covering classification (DXC and NCC) and segmentation (PXS) in X-ray and CT. As evidenced by our statistical analysis in Fig. 6b, TransVW 2D: (1) significantly surpasses training from scratch and Models Genesis 2D in all three applications, and (2) achieves equivalent performance with ImageNet in NCC and PXS, which is a significant achievement because to date, all self-supervised approaches lag behind fully supervised training [38]-[40]. Taken together, these results indicate that our TransVW 2D, which comes at zero annotation cost, generalizes well across tasks, datasets, and modalities.

Discussion: Is there any correspondence between linear evaluation and fine-tuning performance?—As seen in Fig. 6, ImageNet models underperform in linear evaluations with fixed features; however, full fine-tuning of ImageNet features yield higher (in DXC) or equal (in NCC and PXS) performance compared with TransVW 2D. We surmise that although ImageNet models leverage large-scale annotated data during pre-training, due to the marked domain gap between medical and natural images, the fixed features of ImageNet models may not be aligned with medical applications. However, transfer learning from ImageNet models can still provide a good initialization point for

the CNNs since its low and mid-level features can be reused in the target tasks. Thus, fine-tuning their features on a large-scale dataset such as ChestX-ray14 could mitigate the discrepancy between natural and medical domains, yielding good target task performance.

Our observations in line with [41] suggest that although linear evaluation is informative for utilizing fixed features, it may not have a strong correlation to the fine-tuning performance.

VI. Related works

Bag-of-Visual-Words (BoVW):

The BoVW model [1] represents images by local invariant features [42] that are condensed into a single vector representation. BoVW and its extensions have been widely used in various tasks [43]. A major drawback of BoVW is that the extracted visual words cannot be transferred and fine-tuned for new tasks and datasets like CNN models. To address this challenge, recently, a few works have integrated BoVW in the training pipeline of CNNs [44], [45]. Among them, Gidaris *et al.* [44] proposed a pretext task based on the BoVW pipeline, which first discretizes images to a set of spatially dense visual words using a pre-trained network, and then trains a second network to predict the BoVW histogram of the images. While this approach displays impressive results for natural images, the extracted visual words may not be intuitive and explainable from a medical perspective since they are automatically determined in feature space. Moreover, using K-means clustering in creating the visual vocabulary may lead to imbalanced clusters of visual words (known as cluster degeneracy) [46]. Our method differs from previous works in (1) automatically discovering visual words that carry explainable semantic information from medical perspective, (2) bypassing the clustering, reducing the training time and leading to the balanced classes of visual words, and (3) proposing a novel pretext rather than predicting the BoVW histogram.

Self-supervised learning:

Self-supervised learning methods aim to learn general representations from unlabeled data. In this paradigm, a neural network is trained on a manually designed (pretext) task for which ground-truth is available for free. The learned representations can be later fine-tuned on numerous target tasks with limited annotated data. A broad variety of self-supervised methods have been proposed for pre-training CNNs in natural images domain, solving jigsaw puzzles [28], predicting image rotations [2], inpainting of missing parts [3], clustering images and then predicting the cluster assignments [29], [47], and removing noise from noisy images [27]. However, self-supervised learning is a relatively new trend in medical imaging. Recent methods, including colorization of colonoscopy images [48], anatomical positions prediction within cardiac MR images [49], context restoration [4], and Rubik's cube recovery [14], which were developed individually for specific target tasks, have a limited generalization ability over multiple tasks. TransVW distinguishes itself from all other existing works by explicitly employing the strong yet free semantic supervision signals of visual words, leading to a generic pre-trained model effective for various target tasks. Recently, Zhou *et al.* [5] proposed four effective image transformations for learning generic autodidactic models through a restoration-based task for 3D medical imaging. While our method derives the transformations from Models Genesis [5], it shows **three significant**

advancements. First, Models Genesis has only one self-restoration component, while we introduce two more novel components: self-discovery and self-classification, which are sole factors in the performance gain. Second, our method learns semantic representation from the consistent and recurring visual words discovered during our self-discovery phase, but Models Genesis learns representation from random sub-volumes with no semantics, since no semantics can be discovered from random sub-volumes. Finally, our method serves as an add-on for boosting other self-supervised methods while Models Genesis do not offer such advantage.

Our previous work:

Haghighi *et al.* [50] first proposed to utilize consistent anatomical patterns for training semantics-enriched pre-trained models. The current work presents several extensions to the preliminary version: (1) We introduce a new concept: **transferable visual words**, where the recurrent anatomical structures in medical images are anatomical visual words, which can be automatically discovered from unlabeled medical images, serving as strong yet free supervision signals for training deep models; (2) We extensively investigate the add-on capability of our self-discovery and self-classification, demonstrating that they can boost existing self-supervised learning methods (Sec. IV-A); (3) We expand our 3D target tasks by adding pulmonary embolism false positive reduction, indicating that our TransVW generalizes effectively across organs, datasets, and modalities (Sec. III-B); (4) We extend Rotation [2] into its 3D version as an additional baseline (Sec. IV-B1); (5) We adopt a rigorous three-pronged approach in evaluating the transferability of our work, including transfer learning performance, convergence speedups, and annotation efficiency, highlighting that TransVW is an annotation-efficient solution for 3D medical imaging. As part of this endeavor we illustrate that transfer learning from TransVW provides better optimization and accelerates the training process (Sec. IV-B2) and dramatically reduces annotation efforts (Sec. IV-B3). (6) We conduct linear evaluations on top of the fixed features, showing that TransVW 2D provides more generic representations by reducing the overfitting to the pretext task (Sec. V-B); and (7) We conduct ablation studies on five 3D target tasks to search for an effective number of visual word classes (Appendix I-B).

VII. Conclusion

A key contribution of ours is designing a self-supervised learning framework that not only allows deep models to learn common visual representation from image data directly, but also leverages the semantics associated with the recurrent anatomical patterns across medical images, resulting in generic semantics-enriched image representations. Our extensive experiments demonstrate the annotation-efficiency of TransVW by offering higher performance and faster convergence with reduced annotation cost in comparison with publicly available 3D models pre-trained by not only self-supervision but also full supervision. More importantly, TransVW can be used as an add-on scheme to substantially improve other self-supervised methods. We attribute these outstanding results to the compelling deep semantics derived from recurrent visual words resulted from consistent anatomies naturally embedded in medical images.

Acknowledgment

This research has been supported partially by ASU and Mayo Clinic through a Seed Grant and an Innovation Grant, and partially by the NIH under Award Number R01HL128785. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH. This work has utilized the GPUs provided partially by the ASU Research Computing and partially by the Extreme Science and Engineering Discovery Environment (XSEDE) funded by the National Science Foundation (NSF) under grant number ACI-1548562. We thank Zuwei Guo for implementing Rubik's cube [14], Md Mahfuzur Rahman Siddiquee for examining NiftyNet [16], Jiaxuan Pang for evaluating I3D [15], Shivam Bajpai for helping in adopting TransVW to nnU-Net [26], and Shrikar Tatapudi for helping improve the writing of this paper. The content of this paper is covered by patents pending.

References

- [1]. Sivic J and Zisserman A, "Video google: a text retrieval approach to object matching in videos," in Proceedings Ninth IEEE International Conference on Computer Vision, 10 2003, pp. 1470–1477 vol.2.
- [2]. Gidaris S, Singh P, and Komodakis N, "Unsupervised representation learning by predicting image rotations," arXiv preprint arXiv:1803.07728, 2018.
- [3]. Pathak D, Krahenbuhl P, Donahue J, Darrell T, and Efros AA, "Context encoders: Feature learning by inpainting," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 2536–2544.
- [4]. Chen L, Bentley P, Mori K, Misawa K, Fujiwara M, and Rueckert D, "Self-supervised learning for medical image analysis using image context restoration," *Medical image analysis*, vol. 58, p. 101539, 2019. [PubMed: 31374449]
- [5]. Zhou Z, Sodha V, Rahman Siddiquee MM, Feng R, Tajbakhsh N, Gotway MB, and Liang J, "Models genesis: Generic autodidactic models for 3d medical image analysis," in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*. Cham: Springer International Publishing, 2019, pp. 384–393.
- [6]. Zhou Z, Sodha V, Pang J, Gotway MB, and Liang J, "Models genesis," *Medical Image Analysis*, vol. 67, p. 101840, 2021. [PubMed: 33188996]
- [7]. Setio AAA, Traverso A, De Bel T, Berens MS, van den Bogaard C, Cerello P, Chen H, Dou Q, Fantacci ME, Geurts B et al. , "Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images: the luna16 challenge," *Medical image analysis*, vol. 42, pp. 1–13, 2017. [PubMed: 28732268]
- [8]. Armato III SG, McLennan G, Bidaut L, McNitt-Gray MF, Meyer CR, Reeves AP, Zhao B, Aberle DR, Henschke CI, Hoffman EA et al. , "The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans," *Medical physics*, vol. 38, pp. 915–931, 2011. [PubMed: 21452728]
- [9]. Tajbakhsh N, Gotway MB, and Liang J, "Computer-aided pulmonary embolism detection using a novel vessel-aligned multi-planar image representation and convolutional neural networks," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2015, pp. 62–69.
- [10]. Bilic P, Christ PF, Vorontsov E, Chlebus G, Chen H, Dou Q, Fu C-W, Han X, Heng P-A, Hesser J et al. , "The liver tumor segmentation benchmark (lits)," arXiv preprint arXiv:1901.04056, 2019.
- [11]. Bakas S, Reyes M, Jakab A, Bauer S, Rempfler M, Crimi A, Shinohara RT, Berger C, Ha SM, Rozycki M et al. , "Identifying the best machine learning algorithms for brain tumor segmentation, progression assessment, and overall survival prediction in the brats challenge," arXiv preprint arXiv:1811.02629, 2018.
- [12]. Wang X, Peng Y, Lu L, Lu Z, Bagheri M, and Summers RM, "Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 2097–2106.
- [13]. (2019) Siim-acr pneumothorax segmentation. [Online]. Available: <https://www.kaggle.com/c/siim-acr-pneumothorax-segmentation/>

- [14]. Zhuang X, Li Y, Hu Y, Ma K, Yang Y, and Zheng Y, “Self-supervised feature learning for 3d medical images by playing a rubik’s cube,” in Medical Image Computing and Computer Assisted Intervention – MICCAI 2019, Shen D, Liu T, Peters TM, Staib LH, Essert C, Zhou S, Yap P-T, and Khan A, Eds. Cham: Springer International Publishing, 2019, pp. 420–428.
- [15]. Carreira J and Zisserman A, “Quo vadis, action recognition? a new model and the kinetics dataset,” in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 6299–6308.
- [16]. Gibson E, Li W, Sudre C, Fidon L, Shkir DI, Wang G, Eaton-Rosen Z, Gray R, Doel T, Hu Y et al. , “Niftynet: a deep-learning platform for medical imaging,” Computer methods and programs in biomedicine, vol. 158, pp. 113–122, 2018. [PubMed: 29544777]
- [17]. Chen S, Ma K, and Zheng Y, “Med3d: Transfer learning for 3d medical image analysis,” arXiv preprint arXiv:1904.00625, 2019.
- [18]. Jenni S, Jin H, and Favaro P, “Steering self-supervised feature learning beyond local pixel statistics,” in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 6 2020.
- [19]. Doersch C, Gupta A, and Efros AA, “Unsupervised visual representation learning by context prediction,” in Proceedings of the IEEE International Conference on Computer Vision, 2015, pp. 1422–1430.
- [20]. Mundhenk TN, Ho D, and Chen BY, “Improvements to context based self-supervised learning.” in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 9339–9348.
- [21]. Goyal P, Mahajan D, Gupta A, and Misra I, “Scaling and benchmarking self-supervised visual representation learning,” in Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2019.
- [22]. Ardila D, Kiraly AP, Bharadwaj S, Choi B, Reicher JJ, Peng L, Tse D, Etemadi M, Ye W, Corrado G et al. , “End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography,” Nature medicine, vol. 25, pp. 954–961, 2019.
- [23]. Gibson E, Giganti F, Hu Y, Bonmati E, Bandula S, Gurusamy K, Davidson B, Pereira SP, Clarkson MJ, and Barratt DC, “Automatic multi-organ segmentation on abdominal ct with dense v-networks,” IEEE transactions on medical imaging, vol. 37, no. 8, pp. 1822–1834, 2018. [PubMed: 29994628]
- [24]. Wu B, Zhou Z, Wang J, and Wang Y, “Joint learning for pulmonary nodule segmentation, attributes and malignancy prediction,” in 2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018). IEEE, 2018, pp. 1109–1113.
- [25]. Zhou Z, Shin J, Zhang L, Gurudu S, Gotway M, and Liang J, “Fine-tuning convolutional neural networks for biomedical image analysis: actively and incrementally,” in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 7340–7349.
- [26]. Isensee F, Jger PF, Kohl SAA, Petersen J, and Maier-Hein KH, “Automated design of deep learning methods for biomedical image segmentation,” 2020.
- [27]. Vincent P, Larochelle H, Bengio Y, and Manzagol P-A, “Extracting and composing robust features with denoising autoencoders,” in Proceedings of the 25th International Conference on Machine Learning, ser. ICML 08. Association for Computing Machinery, 2008, p. 10961103.
- [28]. Noroozi M and Favaro P, “Unsupervised learning of visual representations by solving jigsaw puzzles,” in European Conference on Computer Vision. Springer, 2016, pp. 69–84.
- [29]. Zhan X, Xie J, Liu Z, Ong Y-S, and Loy CC, “Online deep clustering for unsupervised representation learning,” in Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020.
- [30]. Chaitanya K, Erdil E, Karani N, and Konukoglu E, “Contrastive learning of global and local features for medical image segmentation with limited annotations,” in Neural Information Processing Systems (NeurIPS), 2020.
- [31]. Kornblith S, Shlens J, and Le QV, “Do better imagenet models transfer better?” in 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2019, pp. 2656–2666.

- [32]. Tajbakhsh N, Shin JY, Gurudu SR, Hurst RT, Kendall CB, Gotway MB, and Liang J, "Convolutional neural networks for medical image analysis: Full training or fine tuning?" *IEEE transactions on medical imaging*, vol. 35, no. 5, pp. 1299–1312, 2016. [PubMed: 26978662]
- [33]. Neyshabur B, Sedghi H, and Zhang C, "What is being transferred in transfer learning?" in *Neural Information Processing Systems (NeurIPS)*, 2020.
- [34]. Yosinski J, Clune J, Bengio Y, and Lipson H, "How transferable are features in deep neural networks?" in *Advances in neural information processing systems*, 2014, pp. 3320–3328.
- [35]. Raghu M, Zhang C, Kleinberg J, and Bengio S, "Transfusion: Understanding transfer learning with applications to medical imaging," *arXiv preprint arXiv:1902.07208*, 2019.
- [36]. Wang Y, Jiang Z, Chen X, Xu P, Zhao Y, Lin Y, and Wang Z, "E2-train: Training state-of-the-art cnns with over 80
- [37]. Simpson AL, Antonelli M, Bakas S, Bilello M, Farahani K, van Ginneken B, Kopp-Schneider A, Landman BA, Litjens G, Menze BH, Ronneberger O, Summers RM, Bilic P, Christ PF, Do RKG, Gollub M, Golia-Pernicka J, Heckers S, Jarnagin WR, McHugo M, Napel S, Vorontsov E, Maier-Hein L, and Cardoso MJ, "A large annotated medical image dataset for the development and evaluation of segmentation algorithms," *CoRR*, vol. abs/1902.09063, 2019. [Online]. Available: <http://arxiv.org/abs/1902.09063>
- [38]. Hendrycks D, Mazeika M, Kadavath S, and Song D, "Using self-supervised learning can improve model robustness and uncertainty," in *Advances in Neural Information Processing Systems*, 2019, pp. 15 637–15 648.
- [39]. Caron M, Bojanowski P, Mairal J, and Joulin A, "Unsupervised pretraining of image features on non-curated data," in *Proceedings of IEEE International Conference on Computer Vision*, 2019, pp. 2959–2968.
- [40]. Zhang L, Qi G-J, Wang L, and Luo J, "Aet vs. aed: Unsupervised representation learning by auto-encoding transformations rather than data," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 2547–2555.
- [41]. Newell A and Deng J, "How useful is self-supervised pretraining for visual tasks?" in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6 2020.
- [42]. Lowe DG, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, pp. 91–110, 2004.
- [43]. Yue-Hei Ng J, Yang F, and Davis LS, "Exploiting local features from deep networks for image retrieval," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2015.
- [44]. Gidaris S, Bursuc A, Komodakis N, Perez P, and Cord M, "Learning representations by predicting bags of visual words," in *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [45]. Arandjelovic R, Gronat P, Torii A, Pajdla T, and Sivic J, "Netvlad: Cnn architecture for weakly supervised place recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 6 2016.
- [46]. Gansbeke WV, Vandenhende S, Georgoulis S, Proesmans M, and Gool LV, "Scan: Learning to classify images without labels," 2020.
- [47]. Caron M, Bojanowski P, Joulin A, and Douze M, "Deep clustering for unsupervised learning of visual features," in *Proceedings of the European Conference on Computer Vision*, 2018, pp. 132–149.
- [48]. Ross T, Zimmerer D, Vemuri A, Isensee F, Wiesenfarth M, Bodenstedt S, Both F, Kessler P, Wagner M, Müller B et al. , "Exploiting the potential of unlabeled endoscopic video data with self-supervised learning," *International journal of computer assisted radiology and surgery*, vol. 13, no. 6, pp. 925–933, 2018. [PubMed: 29704196]
- [49]. Bai W, Chen C, Tarroni G, Duan J, Guitton F, Petersen SE, Guo Y, Matthews PM, and Rueckert D, "Self-supervised learning for cardiac mr image segmentation by anatomical position prediction," in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, Shen D, Liu T, Peters TM, Staib LH, Essert C, Zhou S, Yap P-T, and Khan A, Eds. Cham: Springer International Publishing, 2019, pp. 541–549.

- [50]. Haghighi F, Hosseinzadeh Taher MR, Zhou Z, Gotway MB, and Liang J, “Learning semantics-enriched representation via self-discovery, self-classification, and self-restoration,” in *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*. Cham: Springer International Publishing, 2020, pp. 137–147.

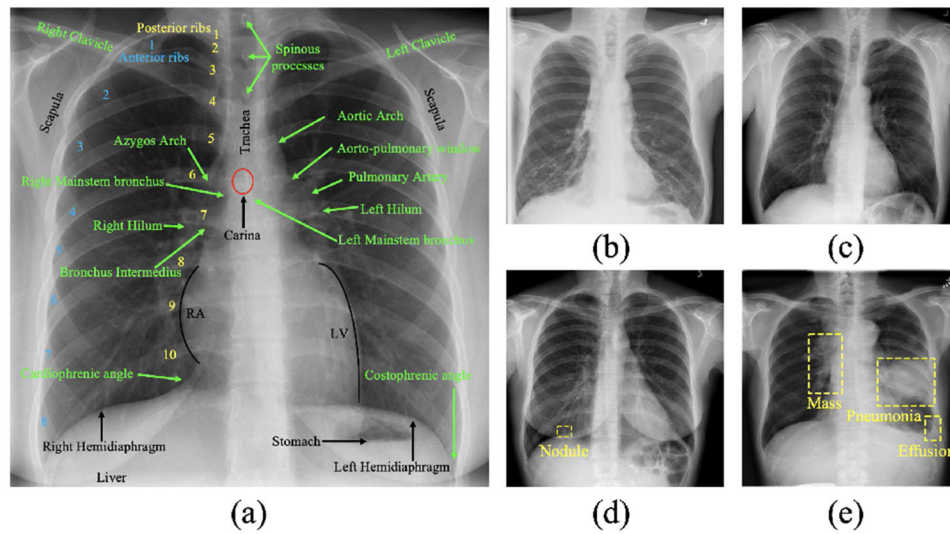


Fig. 1. Without loss of generality, we illustrate our idea in 2D with chest X-rays. The great similarity of the lungs in anatomy, partially annotated in (a), across patients yields complex yet consistent and recurring anatomical patterns across X-rays in healthy (a, b, and c) or diseased (d and e), which we refer to as anatomical *visual words*. Our proposed TransVW (transferable visual words) aims to learn *generalizable* image representation from the anatomical visual words *without* expert annotations, and *transfer* the learned deep models to create powerful application-specific target models. TransVW is general and applicable across organs, diseases, and modalities in both 2D and 3D.

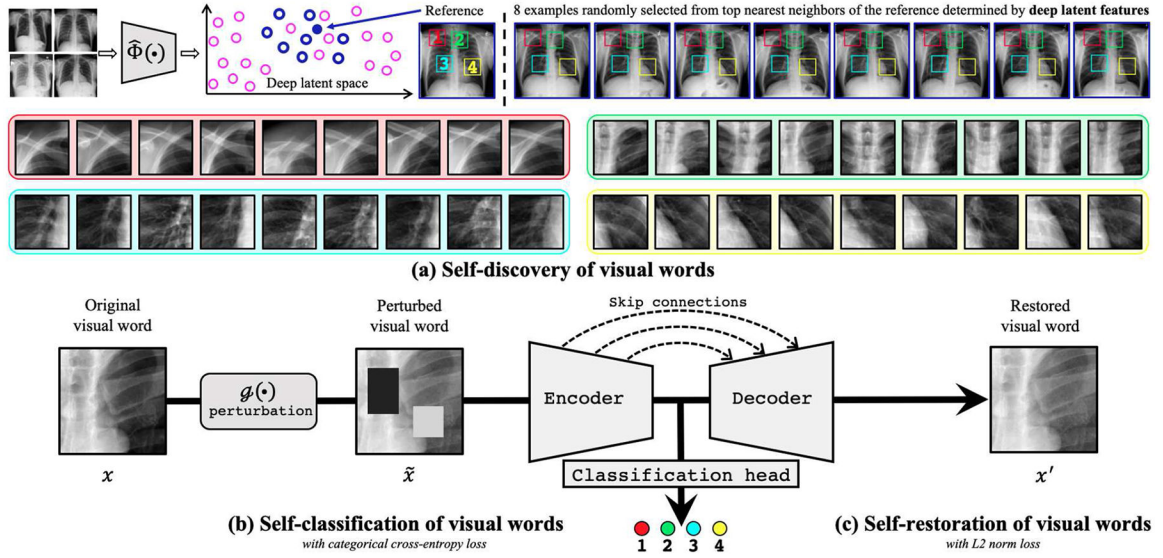
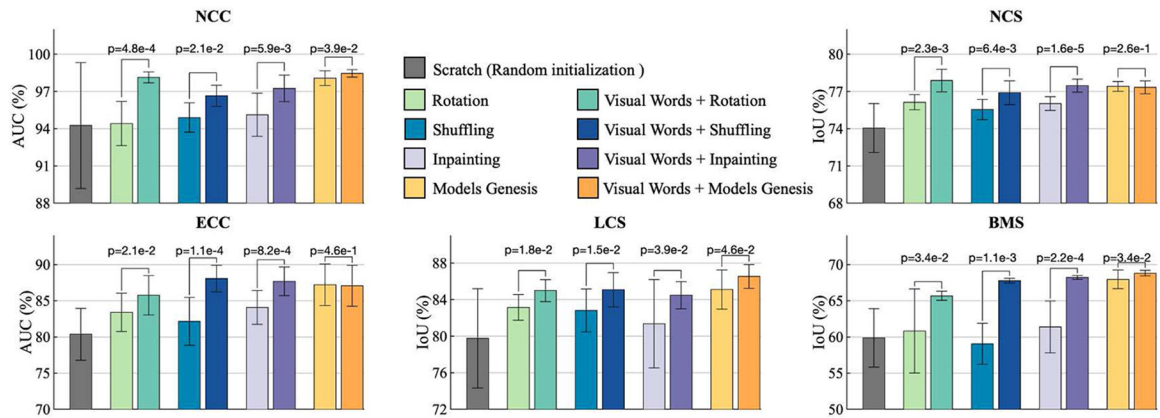


Fig. 2.

Our proposed self-supervised learning framework TransVW is for learning general-purpose image representation enriched with the semantics of anatomical visual words by (a) self-discovery, (b) self-classification, and (c) self-restoration. First, to discover anatomically consistent instances for each visual word across patients, we train a feature extractor $\hat{\Phi}(\cdot)$ (e.g., auto-encoder) with unlabeled images, so that images of great resemblance can be automatically identified based on its deep latent features. Second, after selecting a random reference patient and using the feature extractor to find patients similar in appearance, to extract instances of a visual word, we crop image patches at a random yet fixed coordinate across all selected patients and assign a unique (pseudo) label to the extracted patches (instances). For simplicity and clarity, we have shown instances of four visual words extracted at four different random coordinates to illustrate the similarity and consistency among the discovered instances of each visual word. Our self-discovery automatically curates a set of visual words associated with semantically meaningful labels, providing a free and rich source for training deep models to learn semantic representations. Finally, we perturb instances of the visual words with $g(\cdot)$ and give them as input to an encoder-decoder network with skip connections in between and a classification head at the end of the encoder. Our self-classification and self-restoration of visual words empower the deep model to learn anatomical semantics from the visual words, resulting in image representation, which has proven to be more generalizable and transferable to a variety of target tasks.

**Fig. 3.**

Our self-supervised learning scheme serves as an add-on, which can be added to enrich existing self-supervised learning methods. By introducing self-discovery and self-classification of visual words, we empower four representative self-supervised learning advances (*i.e.*, Inpainting [3], Context restoration (Shuffling) [4], Rotation [2], and Models Genesis [5]) to capture more high-level and diverse representations, resulting in substantial ($p < 0.05$) performance improvements on five 3D target tasks.




Fig. 4.

Fine-tuning from TransVW provides better optimization and accelerates the training process in comparison with training from scratch as well as state-of-the-art Models Genesis [5], as demonstrated by the learning curves for the five 3D target tasks. All models are evaluated on the validation set, and the average accuracy and dice-coefficient over ten runs are plotted for the classification and segmentation tasks, respectively.




Fig. 5.

Fine-tuning TransVW reduces the annotation cost by 50%, 50%, 57%, 60%, and 80% in NCC, NCS, ECC, LCS, and BMS applications, respectively, when comparing with training from scratch. Moreover, TransVW reduces the annotation efforts by 17%, 24%, and 50% in NCC, LCS, and BMS applications, respectively, compared with state-of-the-art Models Genesis [5]. The horizontal gray and orange lines show the performance achieved by training from scratch and Models Genesis, respectively, when using the entire training data. The gray and orange bars indicate the minimum portion of training data that is required for training models from scratch and Models Genesis to achieve the comparable performance (based on the statistical analyses) with the corresponding models when training with the entire training data.




Fig. 6.

We compare the learned representation of TransVW 2D with Models Genesis 2D (self-supervised) and ImageNet (fully-supervised) by (a) training linear classifiers on top of fixed features, and (b) full fine-tuning of the models on 2D applications. In the linear evaluations (a), TransVW representations are transferred better across all the layers on both DXC and NCC in comparison with Models Genesis 2D and ImageNet, demonstrating more generalizable features. Based on the fine-tuning results (b), TransVW 2D significantly surpasses training from scratch and Models Genesis 2D, and achieves equivalent performance with ImageNet in NCC and PXS.

Table I

Target tasks for Transfer Learning. We transfer the learned representations by fine-tuning it for seven medical imaging applications including 3D and 2D image classification and segmentation tasks. To evaluate the generalization ability of our TransVW, we select a diverse set of applications ranging from the tasks on the same dataset as pre-training to the tasks on the unseen organs, datasets, or modalities during pre-training. For each task, ✓ denotes the properties that are in common between the pretext and target tasks.

Code*	Application	Object	Modality	Dataset	Common with pretext task	
					Organ	Modality
NCC	Lung nodule false positive reduction	Lung Nodule	CT	LUNA16 [7]	✓	✓
NCS	Lung nodule segmentation	Lung Nodule	CT	LIDC-IDRI [8]	✓	✓
ECC	Pulmonary embolism false positive reduction	Pulmonary Emboli	CT	PE-CAD [9]	✓	✓
LCS	Liver segmentation	Liver	CT	LITS-2017 [10]		✓
BMS	Brain Tumor Segmentation	Brain Tumor	MRI	BraTS2018 [11]		
DXC	Fourteen thorax diseases classification	Thorax Diseases	X-ray	ChestX-Ray14 [12]	✓	✓
PXS	Pneumothorax Segmentation	Pneumothorax	X-ray	SIIM-ACR-2019 [13]	✓	✓

*The first letter denotes the object of interest ("N" for lung nodule, "B" for brain tumor, "L" for liver, etc); the second letter denotes the modality ("C" for CT, "X" for X-ray, "M" for MRD); the last letter denotes the task ("C" for classification, "S" for segmentation).

TransVW significantly outperforms training from scratch, and achieves the best or comparable performance in five 3D target applications over five self-supervised and three publicly available supervised pre-trained 3D models. We evaluate the classification (*i.e.*, NCC and ECC) and segmentation (*i.e.*, NCS, LCS, and BMS) target tasks under auc and iou metrics, respectively. For each target task, we show the average performance and standard deviation across ten runs, and further perform independent two sample t-test between the best (bolded) vs. others and highlighted boxes in green when they are not statistically significantly different at $p = 0.05$ level.

Table II

Method	Pre-training task		Target tasks					
	Supervised	Dataset	NCC (%)	NCS ¹ (%)	ECC ² (%)	LCS ³ (%)	BMS ⁴ (%)	
Random	×	N/A	94.25±5.07	74.05±1.97	80.36±3.58	79.76±5.42	59.87±4.04	
NiftyNet [16]	✓	Pancreas-CT, BTCV [23]	94.14±4.57	52.98±2.05	77.33±8.05	83.23±1.05	60.78±1.60	
MedicalNet [17]	✓	3Dseg-8 [17]	95.80±0.51	75.68±0.32	86.43±1.44	85.52±0.58	66.09±1.35	
I3D [15]	✓	Kinetics [15]	98.26±0.27	71.58±0.55	80.55±1.11	70.65±4.26	67.83±0.75	
Inpainting [3]	×	LUNA16 [7]	95.12±1.74	76.02±0.55	84.08±2.34	81.36±4.83	61.38±3.84	
Context restoration [4]	×	LUNA16 [7]	94.90±1.18	75.55±0.82	82.15±3.3	82.82±2.35	59.05±2.83	
Rotation [2]	×	LUNA16 [7]	94.42±1.78	76.13±0.61	83.40±2.71	83.15±1.41	60.53±5.22	
Rubik's Cube [14]	×	LUNA16 [7]	96.24±1.27	72.87±0.16	80.49±4.64	75.59±0.20	62.75±1.93	
Models Genesis [5]	×	LUNA16 [7]	98.07±0.59	77.41±0.40	87.2±2.87	85.1±2.15	67.96±1.29	
VW classification	×	LUNA16 [7]	97.49±0.45	76.93±0.87	84.25±3.91	84.14±1.78	64.02±0.98	
VW restoration	×	LUNA16 [7]	98.10±0.19	77.70±0.59	86.20±3.21	84.57±2.20	67.78±0.57	
TransVW	×	LUNA16 [7]	98.46±0.30	77.33±0.52	87.07±2.83	86.53±1.30	68.82±0.38	

¹ [24] holds a Dice of 74.05% vs. 75.85 ± 0.83% (ours)

² [25] holds an AUC of 87.06% vs. 87.07%±2.83% (ours)

³ [26] holds a Dice of 95.76% vs. 95.84% ± 0.07% (ours using nnU-Net framework)

⁴ MR Flair images are only utilized for segmenting brain tumors, so the results are not submitted to BraTS-2018.