

UC Irvine

UC Irvine Previously Published Works

Title

Detection of sclerotic spine metastases via random aggregation of deep convolutional neural network classifications

Permalink

<https://escholarship.org/uc/item/7bs0z8jh>

ISBN

9783319141473

Authors

Roth, HR
Yao, J
Lu, L
[et al.](#)

Publication Date

2015

DOI

10.1007/978-3-319-14148-0_1

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Detection of Sclerotic Spine Metastases via Random Aggregation of Deep Convolutional Neural Network Classifications

Holger R. Roth^{*1}, Jianhua Yao¹, Le Lu¹, James Stieger¹, Joseph E. Burns², and Ronald M. Summers¹

¹Imaging Biomarkers and Computer-Aided Diagnosis Laboratory,, Radiology and Imaging Sciences, National Institutes of Health Clinical Center, Bethesda, MD 20892-1182, USA.

²Department of Radiological Sciences, University of California-Irvine, Orange, CA 92868, USA.

Abstract

Automated detection of sclerotic metastases (bone lesions) in Computed Tomography (CT) images has potential to be an important tool in clinical practice and research. State-of-the-art methods show performance of 79% sensitivity or true-positive (TP) rate, at 10 false-positives (FP) per volume. We design a two-tiered coarse-to-fine cascade framework to first operate a highly sensitive candidate generation system at a maximum sensitivity of $\sim 92\%$ but with high FP level (~ 50 per patient). Regions of interest (ROI) for lesion candidates are generated in this step and function as input for the second tier. In the second tier we generate N 2D views, via scale, random translations, and rotations with respect to each ROI centroid coordinates. These random views are used to train a deep Convolutional Neural Network (CNN) classifier. In testing, the CNN is employed to assign individual probabilities for a new set of N random views that are averaged at each ROI to compute a final per-candidate classification probability. This second tier behaves as a highly selective process to reject difficult false positives while preserving high sensitivities. We validate the approach on CT images of 59 patients (49 with sclerotic metastases and 10 normal controls). The proposed method reduces the number of FP/vol. from 4 to 1.2, 7 to 3, and 12 to 9.5 when comparing a sensitivity rates of 60%, 70%, and 80% respectively in testing. The Area-Under-the-Curve (AUC) is 0.834. The results show marked improvement upon previous work.

^{*}holger.roth@nih.gov, h.roth@ucl.ac.uk

1 Introduction

Early detection of sclerotic bone metastases plays an important role in clinical practice. Their detection can assess the staging of the patient’s disease, and therefore has the potential to alter the treatment regimen the patient will undergo (Msaouel et al., 2008). Approximately 490,000 patients per year are affected by metastatic diseases of the skeletal structures in the United States alone (Hitron and Adams, 2009). More than 80% of these bone metastases are thought to originate from breast and prostate cancer (Coleman, 2001). As a ubiquitous screening and staging modality employed for disease detection in cancer patients, Computed Tomography (CT) is commonly involved in the detection of bone metastases. Both lytic and sclerotic metastatic diseases change or deteriorate the bone structure and bio-mechanically weaken the skeleton. Sclerotic metastases grow into irregularly mineralized and disorganized “woven” bone (Saylor and Smith, 2010, Keller and Brown, 2004, Lee et al., 2011, Guise and Mundy, 1998). Typical examples of sclerotic metastases are shown in Fig. 1. The detection of sclerotic metastases often occurs during manual prospective visual inspection of every image (of which there may be thousands) and every section of every image in each patient’s CT study. This is a complex process that is performed under time restriction and which is prone to error. Furthermore, thorough manual assessment and processing is time-consuming and has potential to delay the clinical workflow. Computer-Aided Detection (CADe) of sclerotic metastases has the potential to greatly reduce the radiologists’ clinical workload and could be employed as a second reader for improved assessment of disease (Wiese et al., 2012, Burns et al., 2013, Hammon et al., 2013).

The CADe method presented here aims to build upon an existing system for sclerotic metastases detection and focuses on reducing the false-positive (FP) number of its outputs. We make use of recent advances in computer vision, in particular deep Convolutional Neural Networks (CNNs), to attain this goal. Recently, the availability of large annotated training sets and the accessibility of affordable parallel computing resources via GPUs has made it feasible to train “deep” CNNs (also popularized under the keyword: “deep learning”) for computer vision classification tasks. Great advances in classification of natural images have been achieved (Krizhevsky et al., 2012, Zeiler and Fergus, 2013). Studies that have tried to apply deep learning and CNNs to medical imaging applications also showed promise, e.g. (Prasoon et al., 2013, Roth et al., 2014). In particular, CNNs have been applied successfully in biomedical applications such as digital pathology (Cireřan et al., 2013). In this work, we apply CNNs for the reduction of FPs using random sets of 2D CNN observations. Our motivation is partially inspired by the spirit of hybrid sys-

tems using both parametric and non-parametric models for hierarchical coarse-to-fine classification (Lu et al., 2011).

2 Methods

2.1 Sclerotic Metastases Candidate Detection

We use a state-of-the-art CADe method for detecting sclerotic metastases candidates from CT volumes (Burns et al., 2013, Wiese et al., 2011). The spine is initially segmented by thresholding at certain attenuation levels and performing region growing. Furthermore, morphological operations are used to refine the segmentation and allow the extraction of the spinal canal. For further information on the segmentation refer to (Yao et al., 2006b). Axial 2D cross sections of the spinal vertebrae are then divided into sub-segments using a watershed algorithm based on local density differences (Yao et al., 2006a). The CADe algorithm then finds initial detections that have a higher mean attenuation than neighboring 2D sub-segments. Because the watershed algorithm can cause over-segmentation of the image, similar 2D sub-segments detections are merged by performing an energy minimization based on graph cuts and attenuation thresholds. Finally, 2D detections on neighboring cross sections are combined to form 3D detections using a graph-cut-based merger. Each 3D detection acts as a seed point for a level-set segmentation method that segments the lesions in 3D. This step allows the computation of 25 characteristic features, including shape, size, location, attenuation, volume, and sphericity. A committee of SVMs (Yao et al., 2005) is then trained on these features. The trained SVMs further classify each 3D detection as ‘true’ or ‘false’ bone lesion. Example of bone lesions candidates using this detection scheme are shown in Fig. 1. Next, true bone lesions from this step are used as candidate lesions for a second classification based on CNNs as proposed in this paper. This is a coarse-to-fine classification approach somewhat similar to other CADe schemes such as Lu et al. (2011).




Figure 1: Examples of sclerotic metastases as detected by the CADe candidate generation step (red mark).

2.2 CNN training on 2D Image Patches

A Region-of-Interest (ROI) in a CT image is extracted at each bone lesion candidate location (see Fig. 2). In order to increase the variation of the training data and to avoid overfitting analogous to the data augmentation approach in Krizhevsky et al. (2012), each ROI is translated along a random vector v in axial space. Furthermore, each translated ROI is rotated around its center N_r times by a random angle $\alpha = [0^\circ, \dots, 360^\circ]$. These translations and rotations for each ROI are computed N_s times at different physical scales s (the edge length of each ROI), but with fixed numbers of pixels. This procedure results in $N = N_s \times N_t \times N_r$ random observation of each ROI – an approach similar to Göktürk et al. (2001). Note that 2.5-5 mm thick-sliced CT volumes are used for this study. Due to this relative large slice thickness, our spatial transformations are all drawn from within the axial plane. This is in contrast to other approaches that use CNNs which sample also sagittal and/or coronal planes (Roth et al., 2014, Prasoon et al., 2013). Following this procedure, both the training and test data can be easily expanded to better scale to this type of neural net application. A CNN’s predictions on these N random observations $\{P_1(x), \dots, P_N\}$ can then be simply averaged at each ROI to compute a per-candidate probability:

$$p(x|\{P_1(x), \dots, P_N(x)\}) = \frac{1}{N} \sum_{i=1}^N P_i(x). \quad (1)$$

Here, $P_i(x)$ is the CNN’s classification probability computed one individual image patch. In theory, more sophisticated fusion rules can be explored but we find that simple averaging works well. This proposed random resampling is an approach to effectively and efficiently increase the amount of available training data. In computer vision, translational shifting and mirroring of 2D image patches is often used for this purpose (Krizhevsky et al., 2012). By averaging the N predictions on random 2D views as in Eq. 1, the robustness and stability of CNN can be further increased as shown in Sec. 3.

2.3 CNN Architecture

A CNN derives its name from the convolutional filters that it applies to the input images. Typically, several layers of convolutional filters are cascaded to compute image features. Other layers of a CNN often perform max-pooling operations or consist of fully-connected neural networks. Our CNN ends with a final 2-way softmax layer for ‘true’ and ‘false’ classification (see Fig. 3). The fully connected layers are typically constrained in order to avoid overfitting. We use “DropConnect” for this




Figure 2: Image patches are generated from CADe candidates using different scales, 2D translations (along a random vector v) and rotations (by a random angle α) in the axial plane.

purpose. “DropConnect” is a method that behaves as a regularizer when training the CNN (Wan et al., 2013). It can be seen as a variation of the earlier developed “DropOut” method (Hinton et al., 2012). GPU acceleration allows efficient training of the CNN. We use an open-source implementation by Krizhevsky et al. (2012) with the DropConnect extension by Wan et al. (2013). Further execution speed-up for both training and evaluation is achieved by using rectified linear units as the neuron model instead of the traditional neuron model $f(x) = \tanh(x)$ or $f(x) = (1 + e^{-x})^{-1}$ (Krizhevsky et al., 2012). At the moment, it is still difficult to design a theoretically optimal CNN architecture for a particular image classification task (Zeiler and Fergus, 2013). We evaluate several CNNs with slightly different layer architectures (independently to the later evaluations) in order to find a suitable CNN architecture for our classification task, using a small number of CT cases within the training data subset. We find relatively stable behavior over model variations and hence fix the CNN architecture for subsequent experiments performed in this study. Recently, approaches have been proposed that aim to visualize the feature activations of CNNs in order to allow better CNN design (Zeiler and Fergus, 2013). Potentially, these approaches allow better understanding of how CNNs behave at a given task. This could lead to improved CNN architecture design compared to the heuristic approach applied in this work.

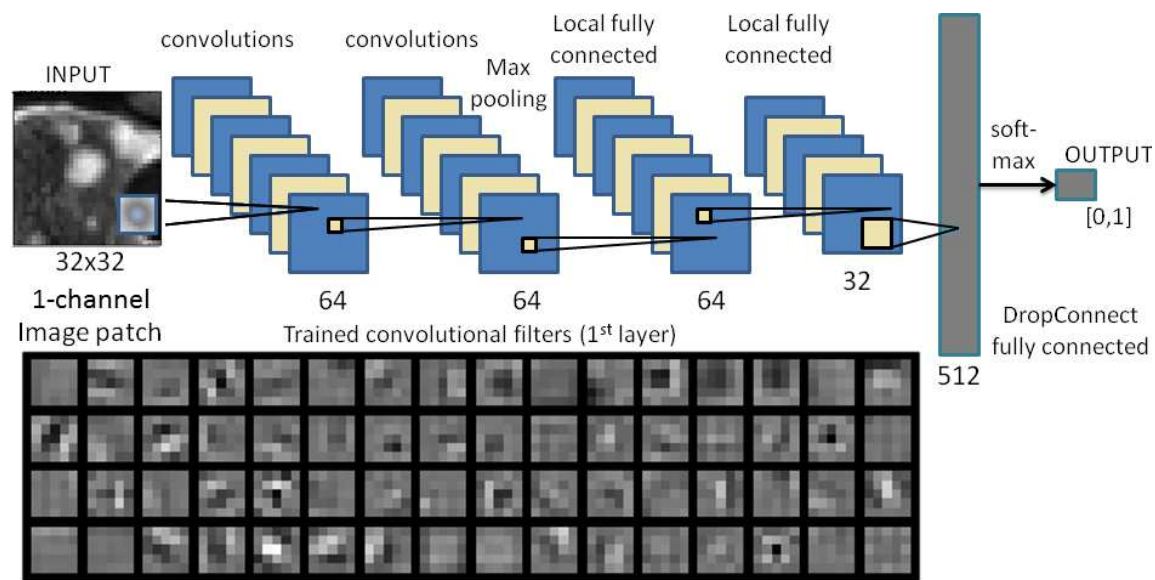


Figure 3: The proposed convolution neural network consists of two convolutional layers, max-pooling layers, locally fully-connected layers, a DropConnect layer, and a final 2-way softmax layer for classification. The number of filters, connections for each layer, and the first layer of learned convolutional kernels are shown.

3 Evaluation and Results on Sclerotic Metastases

In our evaluation, radiologists label a total of 532 sclerotic metastases (‘positives’) in CT images of 49 patients (14 female, 35 male patients; mean age, 57.0 years; range, 12-77 years). A lesion is only labeled if its volume is greater than 300 mm^3 . The CT scans have reconstruction slice thicknesses ranging between 2.5 mm to 5 mm. Furthermore, we include 10 control cases (4 female, 6 male patients; mean age, 55.2 years; range, 1970 years) without any spinal lesions.

Any false-positive detections from the candidate generation step on these patients are used as ‘negative’ candidate examples for training the CNN. All patients were randomly split into five sets at the patient level in order to allow a 5-fold cross-validation. We adjust the sample rates for positive and negative image patches in order to generate a balanced data set for training (50% positives and 50% negatives). This proves to be beneficial for training the CNNs – no balancing was done during testing. Each three-channel image patch was centered at the CAde coordinate with 32×32 pixels in resolution. All patches were sampled at 4 scales of $s = [30, 35, 40, 45]$ mm ROI edge length in physical image space, after iso-metric resampling of the CT

image (see Fig. 2). We used a bone window level of $[-250, 1250 \text{ HU}]$. Furthermore, all ROIs were randomly translated (up to 3 mm) and rotated at each scale ($N_s = 4$, $N_t = 5$ and $N_r = 5$), resulting in $N = 100$ image patches per ROI. Training each CNN model took 12-15 hours on a NVIDIA GeForce GTX TITAN, while running 100 2D image patches at each ROI for classification of one CT volume only took circa 30 seconds. Image patch extraction from one CT volume took around 2 minutes on each scale.

We now apply the trained CNN to classify image patches from the test data sets. Figure 4 and Fig. 5 show typical classification probabilities on random subsets of positive and negative ROIs in the test case. Averaging the N predictions at




Figure 4: Test probabilities of the CNN for being sclerotic metastases on ‘true’ sclerotic metastases candidate examples (close to 1.0 is good).

each sclerotic metastases candidate allows us to compute a per-candidate probability $p(x)$, as in Eq. 1. Varying thresholds on probability $p(x)$ are used to compute Free-Response Receiver Operating Characteristic (FROC) curves. The FROC curves are compared in Fig. 6 for varying amounts of N . It can be seen that the classification performance saturates quickly with increasing N . This means the run-time efficiency of our second layer detection could be further improved without losing noticeable performance by decreasing N . The proposed method reduces the number of FP/vol.

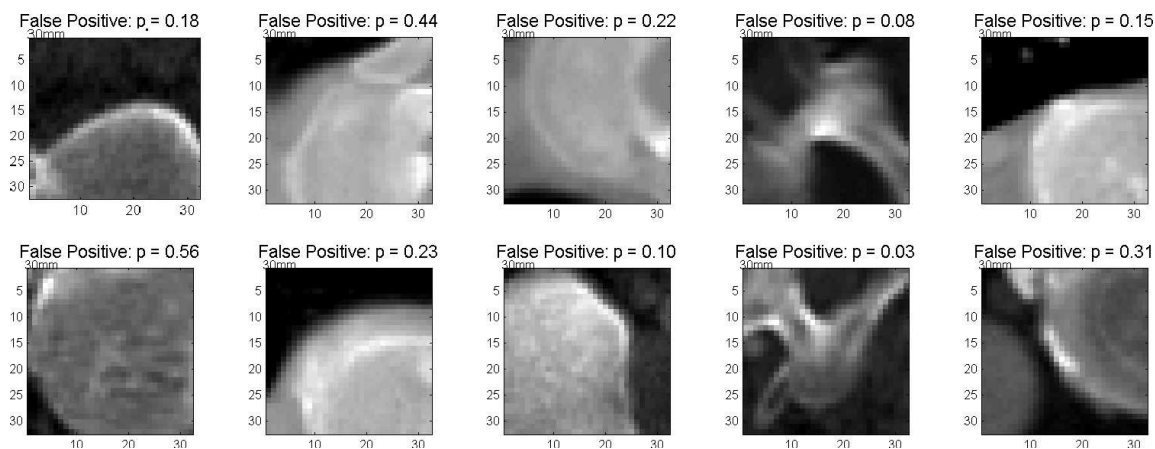


Figure 5: Test probabilities of the CNN for being sclerotic metastases on ‘false’ sclerotic metastases candidate examples (close to 0.0 is good).

of the existing sclerotic metastases CADe systems (Burns et al., 2013) from 4 to 1.2, 7 to 3, and 12 to 9.5 when comparing a sensitivity rates of 60%, 70%, and 80% respectively in cross-validation testing (at $N = 100$). This has the potential to greatly reduce radiologists’ clinical workload when employing the proposed CADe system as a second reader. The Area-Under-the-Curve (AUC) shows a value of 0.834 at this number of N .

Figure 7 compares the FROCs from the initial (first layer) CADe system (Burns et al., 2013) and illustrates the progression towards the proposed coarse-to-fine two tiered method in both training and testing datasets. This clearly demonstrates a marked improvement in performance.

4 Discussion and Conclusions

This work demonstrates that deep CNNs can be generalized to tasks in medical image analysis, such as effective FP reduction in Computer-aided Detection (CADe) systems. This is especially true, since the main drawback of current CADe developments often generates too many false positive detections at clinically relevant sensitivity levels. We show that a random set of CNN classifications can be used to reduce FPs when operating an existing method for CADe of sclerotic metastases (bone lesions) at a particular point its FROC curve. Different scales, random translations, and rotations around each of the CADe detections can be utilized to increase the CNN’s classification performance. The FROC curves show a marked reduction

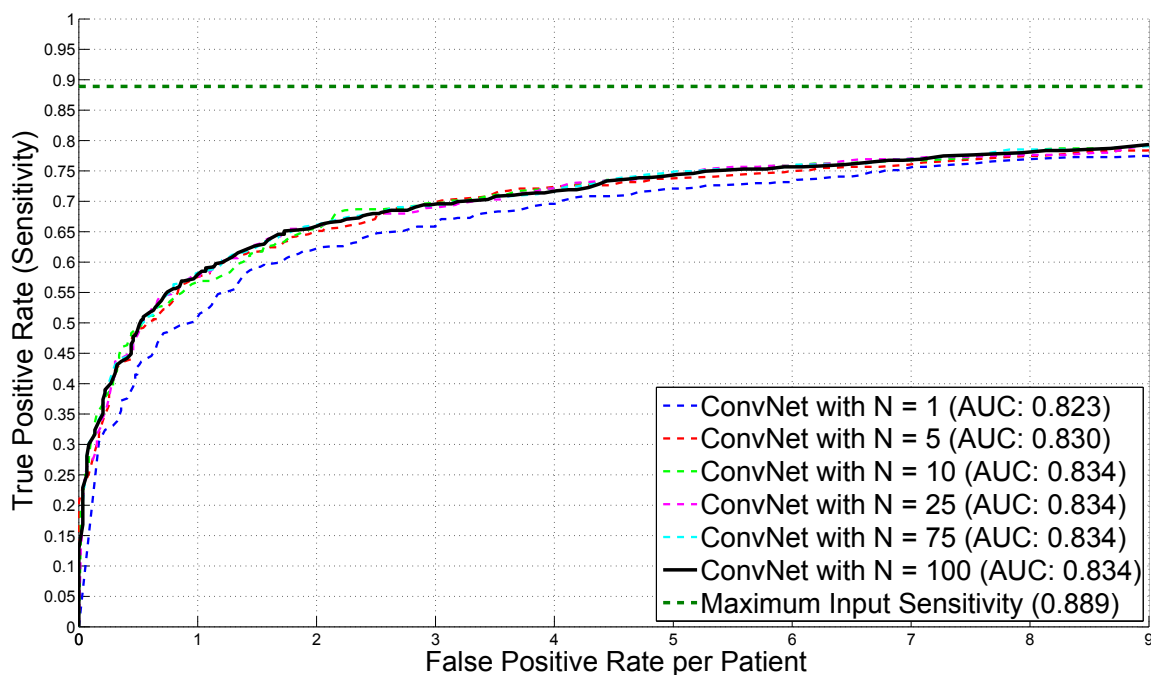


Figure 6: Free-Response Receiver Operating Characteristic (FROC) curves for a 5-fold cross-validation using a varying number of N random CNN observers in 59 patients (49 with sclerotic metastases and 10 normal controls). AUC values are computed for corresponding ROC curves.

of the FP/vol. rate at clinically useful levels of sensitivity. These results improve upon the state-of-the-art.

The average of CNN classification probabilities was chosen in this work for simplicity (see Eq. 1), but this approach shows to be very efficient and effective. Future work will investigate more sophisticated methods of label fusion from the CNNs. A similar 2.5D generalization of CNNs also shows promise in the detection of lymph nodes in CT images (see Roth et al. (2014)). In this work, we decide against a 2.5D or full 3D approach due to the relative large slice thicknesses of ~ 5 mm in the used CT data. This prevents reformatting the data in sufficient detail in any other than the axial plane. However, the improvements achieved in this study and other methods utilizing CNNs in medical image computing show promise for a variety of applications in computer-aided detection of 2D and 3D medical images. Our mainly 2D approach may adapt and generalize particularly well to the current trend of low-dose, low-resolution (slice thickness) CT imaging protocols, compared to direct 3D based methods that require volumetric medical images of higher resolution.

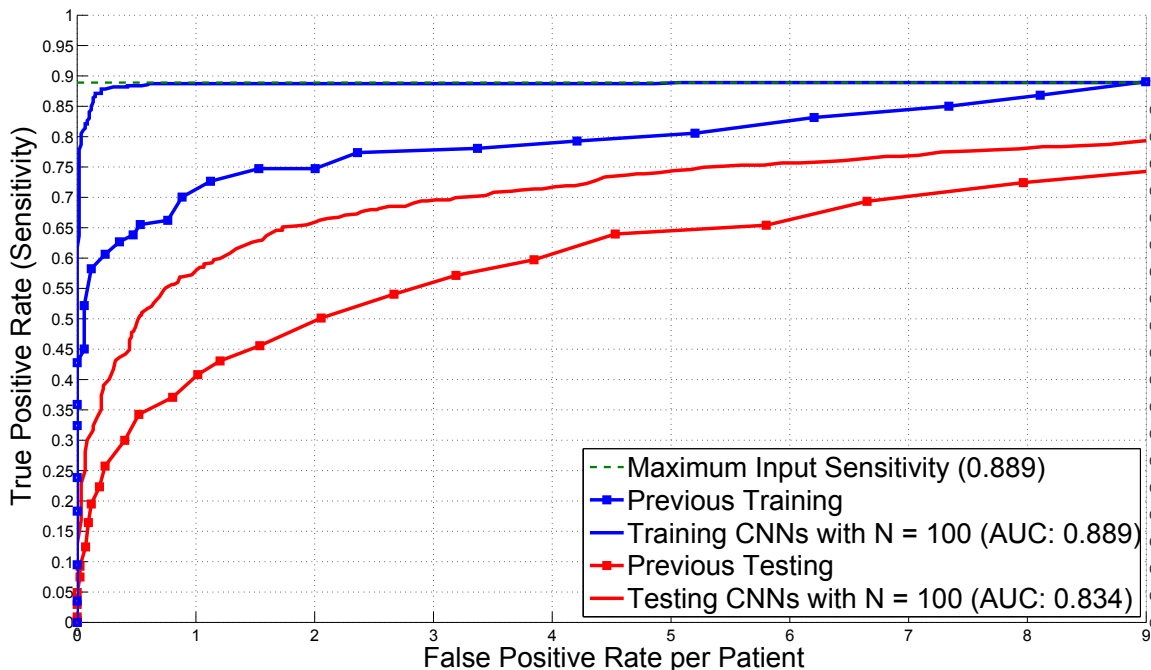


Figure 7: Comparison of Free-Response Receiver Operating Characteristic (FROC) curves of the first layer bone lesion candidate generation (squares) and the second layer classification using $N = 100$ CNN observers (lines) for both training and testing cases. Results are computed using a 5-fold cross-validation in 59 patients (49 with sclerotic metastases and 10 normal controls).

Acknowledgments This work was supported by the Intramural Research Program of the NIH Clinical Center.

References

- Burns, J. E., J. Yao, T. S. Wiese, H. E. Muñoz, E. C. Jones, and R. M. Summers (2013). Automated detection of sclerotic metastases in the thoracolumbar spine at ct. *Radiology* 268(1), 69–78. 2, 3, 8
- Cireşan, D. C., A. Giusti, L. M. Gambardella, and J. Schmidhuber (2013). Mitosis detection in breast cancer histology images with deep neural networks. *MICCAI*. 2

- Coleman, R. (2001). Metastatic bone disease: clinical features, pathophysiology and treatment strategies. *Cancer treatment reviews* 27(3), 165–176. 2
- Göktürk, S. B., C. Tomasi, B. Acar, C. F. Beaulieu, D. S. Paik, R. B. Jeffrey, J. Yee, and Y. Napel (2001). A statistical 3-d pattern processing method for computer-aided detection of polyps in ct colonography. *IEEE Trans. on Med. Imag.* 20, 1251–1260. 4
- Guise, T. A. and G. R. Mundy (1998). Cancer and bone 1. *Endocrine Reviews* 19(1), 18–54. 2
- Hammon, M., P. Dankerl, A. Tsymbal, M. Wels, M. Kelm, M. May, M. Suehling, M. Uder, and A. Cavallaro (2013). Automatic detection of lytic and blastic thoracolumbar spine metastases on computed tomography. *European radiology* 23(7), 1862–1870. 2
- Hinton, G. E., N. Srivastava, A. Krizhevsky, I. Sutskever, and R. R. Salakhutdinov (2012). Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*. 5
- Hitron, A. and V. Adams (2009). The pharmacological management of skeletal-related events from metastatic tumors. *Orthopedics* 32(3), 188. 2
- Keller, E. T. and J. Brown (2004). Prostate cancer bone metastases promote both osteolytic and osteoblastic activity. *Journal of cellular biochemistry* 91(4), 718–729. 2
- Krizhevsky, A., I. Sutskever, and G. Hinton (2012). Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems* 25. 2, 4, 5
- Lee, R. J., P. J. Saylor, and M. R. Smith (2011). Treatment and prevention of bone complications from prostate cancer. *Bone* 48(1), 88–95. 2
- Lu, L., M. Liu, X. Ye, S. Yu, and H. Huang (2011). Coarse-to-fine classification via parametric and nonparametric models for computer-aided diagnosis. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, pp. 2509–2512. ACM. 3
- Msaouel, P., N. Pissimissis, A. Halapas, and M. Koutsilieris (2008). Mechanisms of bone metastasis in prostate cancer: clinical implications. *Best Practice & Research Clinical Endocrinology & Metabolism* 22(2), 341–355. 2

- Prasoon, A., K. Petersen, C. Igel, F. Lauze, E. Dam, and M. Nielsen (2013). Deep feature learning for knee cartilage segmentation using a triplanar convolutional neural network. *MICCAI*. 2, 4
- Roth, H. R., L. Lu, A. Seff, K. M. Cherry, J. Hoffman, S. Wang, J. Liu, E. Turkbey, and R. M. Summers (2014). A new 2.5D representation for lymph node detection using random sets of deep convolutional neural network observations. *MICCAI 2014 (in-print)*, *arXiv:1406.2639 [cs.CV]*. 2, 4, 9
- Saylor, P. and M. Smith (2010). Bone health and prostate cancer. *Prostate cancer and prostatic diseases* 13(1), 20–27. 2
- Wan, L., M. Zeiler, S. Zhang, Y. L. Cun, and R. Fergus (2013). Regularization of neural networks using dropconnect. *Proc. Int. Conf. Machine Learning (ICML-13)*. 5
- Wiese, T., J. Burns, J. Yao, and R. M. Summers (2011). Computer-aided detection of sclerotic bone metastases in the spine using watershed algorithm and support vector machines. In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, pp. 152–155. IEEE. 3
- Wiese, T., J. Yao, J. E. Burns, and R. M. Summers (2012). Detection of sclerotic bone metastases in the spine using watershed algorithm and graph cut. In *SPIE Medical Imaging*, pp. 831512–831512. International Society for Optics and Photonics. 2
- Yao, J., S. D. O’Connor, and R. Summers (2006a). Computer aided lytic bone metastasis detection using regular ct images. In *Medical Imaging*, pp. 614459–614459. International Society for Optics and Photonics. 3
- Yao, J., S. D. O’Connor, and R. M. Summers (2006b). Automated spinal column extraction and partitioning. In *Biomedical Imaging: Nano to Macro, 2006. 3rd IEEE International Symposium on*, pp. 390–393. IEEE. 3
- Yao, J., R. M. Summers, and A. K. Hara (2005). Optimizing the support vector machines (svm) committee configuration in a colonic polyp cad system. In *Medical Imaging*, pp. 384–392. International Society for Optics and Photonics. 3
- Zeiler, M. D. and R. Fergus (2013). Visualizing and understanding convolutional neural networks. *arXiv:1311.2901*. 2, 5