

Imaging without Images: Using Artificial Intelligence for Direct Discovery of Spatial Signatures in the Absence of Image Reconstruction

Sumit Chopra

`sumit.chopra@nyulangone.org`

NYU Grossman School of Medicine

Raghav Singhal

Courant Institute of Mathematical Sciences, NYU

Angela Tong

New York University Grossman School of Medicine <https://orcid.org/0000-0002-4733-3414>

Rajesh Ranganath

NYU

Hersh Chandarana

New York University Grossman School of Medicine

Daniel Sodickson

New York University Grossman School of Medicine

Biological Sciences - Article

Keywords: Machine Learning, Deep Learning, Magnetic Resonance Imaging, Early Disease Detection

Posted Date: June 26th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-8290274/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: There is **NO** Competing Interest.

Imaging without Images: Using Artificial Intelligence for Direct Discovery of Spatial Signatures in the Absence of Image Reconstruction

Raghav Singhal², Angela Tong¹, Rajesh Ranganath²,
Hersh Chandarana¹, Daniel Sodickson¹, Sumit Chopra^{1,2*}

¹Center for Advanced Imaging Innovation and Research (CAI²R),
Department of Radiology, New York University Grossman School of
Medicine, 660 First Avenue, New York, 10016, NY, USA.

²Department of Computer Science, Courant Institute of Mathematical
Sciences, NYU, 60 Fifth Avenue, New York, 10003, NY, USA.

*Corresponding author(s). E-mail(s): sumit.chopra@nyulangone.org;
Contributing authors: rs4070@nyu.edu; angela.tong@nyulangone.org;
rajeshr@cims.nyu.edu; hersh.chandarana@nyulangone.org;
daniel.sodickson@nyulangone.org;

Abstract

The primary goal of scientific discovery or practical diagnostics across multiple disciplines is to discover the presence of entities of interest (EOIs) that are not visible to the human eye and are obscured by potentially distracting signals and noise. In disciplines such as biomedical imaging and astronomy, these objectives are traditionally accomplished by reconstructing images from suitable environment-probing measurements, with the hope that the images will reveal the EOIs. This approach, however, often entails significant cost, complexity, and acquisition time due to the need for elaborate instrumentation and dense sampling. The work reported in this paper challenges a fundamental assumption underlying conventional practice in imaging disciplines: that images are a prerequisite for discovery. We propose an alternative, machine learning-driven paradigm that bypasses the need for image reconstruction altogether, aiming to infer the presence of EOIs directly from a minimal, lower-cost set of measurements that would be insufficient for reconstructing a high-fidelity image. Focusing on disease detection as a representative case study, we apply our framework to simplify the tomographic measurement process in Magnetic Resonance Imaging (MRI).

We accomplish this by learning to map the acquired frequency-space measurements onto a low-dimensional, continuous latent space optimized for downstream inference tasks, rather than forcing the measurements to be represented in an image space. This relaxation of traditional constraints allows us to directly infer the presence or absence of disease from a dramatically reduced set of measurements—measurements that would otherwise be insufficient for reconstructing a diagnostically viable image. Experimental results across multiple diseases and anatomical regions validate our hypothesis: when the goal is accurate inference rather than image formation, substantial reductions in acquisition burden are achievable without sacrificing performance. This paradigm shift has the potential to democratize access to advanced diagnostic tools and accelerate discovery across other scientific domains.

Keywords: Machine Learning, Deep Learning, Magnetic Resonance Imaging, Early Disease Detection

1 Introduction

The goal of medical imaging is to detect diseases within the human body non-invasively. Similarly, one of the goals of astronomy is to identify the existence of cosmic entities, such as exoplanets and black holes, without actually traveling through space. More broadly, many scientific disciplines seek to discover the presence of specific Entities of Interest (EOIs) not visible to the human eye, using environment-probing signals that are admixed with noise. Across disciplines like Magnetic Resonance Imaging (MRI) and radio astronomy, where measured signals represent complex projections of the body or the sky, these signals are traditionally transformed through suitable mathematical operations to extract spatial information in the form of images, in the hope that these images will reveal the EOIs to the human eye. While this image-centric pipeline has been remarkably effective, it comes with a significant burden. Reconstructing high-fidelity images often requires the acquisition of large quantities of high-quality measurements using elaborate and expensive equipment. This makes the discovery process cumbersome, costly, and in many cases, poorly accessible. Our work challenges the fundamental assumption that images are a prerequisite for discovery. We introduce a machine learning-driven framework that detects EOIs directly from intelligently selected raw measurements, bypassing the need for full image reconstruction. Applied to disease detection in MRI, our approach reduces acquisition complexity and cost while preserving diagnostic accuracy, offering a generalizable paradigm for simplified and scalable discovery across imaging-based sciences.

Advanced tomographic imaging, such as Magnetic Resonance Imaging (MRI) and Computed Tomography (CT) have become indispensable tools for accurate disease diagnosis [1–8]. Typically used to confirm clinical hypothesis in differential diagnosis workflows, these modalities serve as the final diagnostic step. At a high level, the tomographic imaging pipeline can be broken into two parts. The **first** part involves the acquisition of measurements (a.k.a., “raw measurements”) that capture some information about the underlying physical phenomena by probing the human body (“the

environment”) non-invasively. In the **second** part the acquired measurements are represented in a suitable form for diagnosing the disease. Since a human (a sub-specialized radiologist) is in charge of making the final diagnosis, these measurements are transformed into a format interpretable to the human eye, typically a cross-sectional image. The diagnostic success of these imaging modalities then lies in their ability to acquire suitable quantities of high-quality raw measurements and represent the information in them as high-fidelity images with high tissue contrast, aiding radiologists in disease identification. For instance, MRI (Figure 1 Part A (top)), first acquires measurements in frequency space (a.k.a., k -space), encoding the electromagnetic response of selected nuclear spins in the human body to applied magnetic fields and radio-frequency pulses. These measurements are then reconstructed into cross-sectional images using an inverse Fourier transform or other suitable mathematical operation. Similarly, CT (Figure 1 Part A (bottom)), collects a series of 2-dimensional X-ray projections from different angles in the form of Sinograms and use them to synthesize cross-sectional images.

Reconstructing high-fidelity images introduces stringent constraints across the imaging pipeline, including shaping hardware requirements, constraints on data acquisition protocols, and reliance on costly, specialized human expertise. These factors collectively contribute to a resource-intensive and cost-prohibitive process, limiting scalability and broader accessibility. In MRI, for example, reconstructing high-fidelity images influences the pipeline in the following three major ways (Figure 1 Part A (top)):

- R1:** Use of sophisticated and expensive scanners installed in special-purpose rooms and operated by expert technicians to acquire high-quality k -space data;
- R2:** Execution of complex and time-consuming protocols to acquire large quantities of k -space data resulting in long scan times;
- R3:** Reliance on sub-specialized radiologists to interpret the reconstructed images.

Similarly, in CT (Figure 1 Part A (bottom)), in addition to **R1**, **R2**, and **R3**, each X-ray projection needed to reconstruct high-fidelity cross-sectional images exposes the patient to radiation, such that considerations of cumulative radiation dose can offset health benefits in any process, like population-level disease surveillance, which requires repeated scans at regular time intervals.

While image-based interpretation is essential for executing many routine radiological tasks, such as multifaceted evaluation of uncertain anomalies or detailed guidance of contemporaneous or subsequent therapy, the high burden of disease discovery associated with advanced imaging modalities restricts their integration into broader areas of the healthcare continuum where their impact could be transformative. A particularly critical example is the early, population-level detection of life-threatening conditions, including various cancers [9–11] and neurological disorders [12], where timely identification could dramatically improve clinical outcomes [13] and reduce the cost of care [14–16]. Though multiple relatively easily accessible tests exist for some diseases, such as prostate cancer [17] and breast cancer [18], they often suffer from low specificity [19, 20], leading to over-diagnosis, unnecessary invasive tissue biopsies, patient trauma, and wasteful spending [19, 21–23]. On the other hand, multiple recent studies have

139 demonstrated the utility of advanced imaging modalities, such as MRI and CT, for
 140 early-stage disease identification by using them as diagnostic tools either in place of
 141 [24–29] or in conjunction with the conventional screening tests [30, 31]. Despite their
 142 proven benefits, the high burden of disease discovery associated with these modalities
 143 constrains their use as tools that can stratify high-risk patients at a population scale.
 144 *There is an unmet need for diagnostic tools that have the accuracy and rich infor-*
 145 *mation content of tomographic imaging and can be deployed near the Point of Care*
 146 *(POC) for early-stage disease screening that is accessible to broad populations.*

147 Unlike traditional downstream radiologic studies, which may involve comprehen-
 148 sive image interpretation, effective population-level disease surveillance often requires
 149 nothing more than a binary decision – whether an individual is likely to have a partic-
 150 ular disease or not – enabling rapid triage either toward further diagnostic evaluation
 151 or away from unnecessary follow-up. In this work, we argue that when such binary out-
 152 comes suffice, representing the raw measurements in pixel space may not be the only
 153 reliable approach to obtaining the solution. Instead, we propose a novel approach that
 154 learns task-specific, non-pixel-space representations of raw measurements optimized
 155 for disease detection by downstream machine learning (ML) models. We hypothesize
 156 that when high-fidelity image reconstruction is not required, diagnostically informative
 157 representations can be extracted directly from low-quality, parsimonious raw measure-
 158 ments – data that would be inadequate for generating diagnostic-quality images but
 159 may still contain sufficient information for binary classification tasks. *Specifically, we*
 160 *demonstrate that a small, carefully selected subset of raw measurements – what we*
 161 *term a “disease signature” in the measurement space – can be identified and leveraged*
 162 *to make accurate disease/no-disease determinations. Furthermore, we can “learn” to*
 163 *identify these minimal sets of raw measurements that are most suitable for disease*
 164 *detection.* To this end, we introduce an ML framework called **Learning Data to Learn**
 165 **Representations (LDLR)**, which: **(1)** learns to identify compact, diagnostically salient
 166 subsets of raw measurements that have sufficient signal to ascertain the presence/ab-
 167 sence of a particular disease (“the disease signature”); **(2)** learns low-dimensional,
 168 task-specific representations of the raw measurements associated with the disease sig-
 169 nature; and **(3)** uses these representations to perform binary disease detection via a
 170 downstream ML model, sidestepping image reconstruction. We contend that LDLR
 171 offers a principled pathway to reducing the complexity and cost of data acquisition
 172 in medical imaging pipelines, potentially enabling their deployment earlier in the
 173 diagnostic workflow than is currently feasible.

174 We validate our hypothesis by applying LDLR to MR scanning to detect multiple
 175 diseases across several anatomies. These include detecting ACL sprains and meniscal
 176 tears in knee MR scans [32, 33], enlarged ventricles and masses in brain scans [32, 33],
 177 and the presence of clinically significant prostate cancer (CSPCA) in abdominal scans.
 178 Our experimental evaluation is designed to address the following five questions:

- 179 Q1 Is it possible to learn an alternative task-specific distributed representation directly
 180 from raw k -space data that preserves the same binary diagnostic information as a
 181 high-fidelity image reconstructed from fully sampled k -space data?
 182 Q2 Can we identify subsets of the k -space with equivalent (binary) diagnostic infor-
 183 mation as in an image reconstructed from full k -space data? Can task-specific
 184

representations be learned from these subsets to support disease detection? What is the minimum subset size required to maintain diagnostic performance?	185
Q3 What are the merits of identifying subsets of raw data and learning their representations, whose information content is optimized for the disease detection task?	186
Q4 For a fixed sampling rate, how do the raw data subsets optimized for disease detection differ from those optimized for conventional image reconstruction tasks?	187
Q5 How does the disease detection performance of a DL classifier trained on the learned task-specific representations compare against an expert radiologist detecting the disease from images reconstructed using the same raw data?	188
Q1 is answered in Section 2.3, where we benchmark the performance of the LDLR trained on fully sampled k -space data against a DL classifier that uses high-fidelity images reconstructed from the same data (i.e., pixel-space representations of the full k -space data). This sets an upper bound on the disease detection accuracy using image-based methods. Results show that the performance of the DL classifier trained on the learned distributed representations is similar to the DL classifier trained on the high-fidelity images. Q2 is answered in Section 2.4, where, for multiple diseases, we show that, using 5% of the data for knee scans and 8% for brain and prostate scans, LDLR matches the accuracy of a DL classifier trained on images reconstructed from fully sampled k -space data. Q3 is answered in Section 2.5, where for sampling rates of 5%, 8%, 10%, and 12.5%, we compare the performance of LDLR with DL classifiers that use images reconstructed using the state-of-the-art image reconstruction models from subsets of k -space data [34, 35]. Results reveal that, for the same disease detection accuracy, LDLR requires a significantly smaller fraction of the k -space data than is needed when detecting the disease from the reconstructed images, underscoring its task-specific efficiency. Section 2.6 answers Q4, where we show that the subsets of k -space identified by LDLR to generate diagnostically rich representations are different from the subsets needed to create representations in pixel space (reconstructing diagnostic quality images). Lastly, Q5 is addressed in Section 2.7, where we compare the disease detection performance of LDLR at a 10% sampling factor to that of board-certified radiologists interpreting images of variable quality. Results show that LDLR outperforms human readers on images reconstructed from undersampled data and matches the human reader performance on fully sampled images, further validating its clinical potential.	189
While our results focus on MR data, the LDLR framework is broadly applicable to other imaging modalities such as CT, ultrasound, and radio astronomy, where direct learning from raw measurements can yield substantial benefits over traditional image reconstruction pipelines.	190
2 Results	191
2.1 The LDLR Pipeline	192
The inference pipeline of the LDLR framework is depicted in Figure 1 Part B. At a high level, LDLR first utilizes a DL model that identifies the minimal subset of	193

k -space containing the maximum diagnostic information specific to a disease (the ‘disease signature’). It then learns a low-dimensional distributed representation of the data associated with the disease signature. Finally, it uses the learnt representations to estimate disease presence/absence (a binary answer) using another DL classifier. Training of the LDLR pipeline is accomplished by executing the following three steps in order (Figure 1 Part C):

1. **Score Function Learning:** We first train a model to estimate the mutual information (the ‘score’) between any subset of the k -space data and the disease label. This is Step 1 in Figure 1 Part C.
2. **Disease Signature Identification:** Next (Step 2 in Figure 1 Part C), we use this score function to score a set of candidate k -space subsets and select the one with the highest score (highest mutual information with the disease label). This subset represents the disease signature. The candidate subsets are chosen using a simple heuristic.
3. **Representation Learning:** Finally (Step 3 in Figure 1 Part C), we train a model that generates the distributed representations of the k -space data associated with the disease signature, which are then used to detect disease presence/absence.

Once the LDLR pipeline is trained, during inference, LDLR requires only the k -space lines corresponding to the learned disease signature to detect the presence of disease, enabling fast and efficient detection from highly undersampled data (Figure 1 Part A).

The core engine of the LDLR pipeline is a novel DL model, which we proposed, called KSPACEENCODER, that takes the frequency-domain data associated with any subset of k -space and generates its low-dimensional distributed representation (Figure 1 Part D). The representations it generates can support multiple downstream tasks, such as estimating the presence or absence of the disease (Step 3 of Figure 1 Part C), or estimating the mutual information between the data of any k -space subset and the disease label (Step 1 in Figure 1 Part C). Implementing KSPACEENCODER required a departure from standard Convolutional Neural Networks (CNNs) [36–39], which rely on inputs that exhibit strong local spatial structure, a property not inherent in the k -space matrix, where each element encodes global image information. A full description of KSPACEENCODER’s architecture, its role in implementing the score function, and the final classifier, along with the training of the LDLR pipeline, is provided in the Methods section.

2.2 Datasets and Anatomies

Table 1 presents statistics related to the datasets used in our experiments.

Knee Scans: We used k -space data from Knee MRI scans in the FASTMRI dataset [32], along with slice-level annotations from FASTMRI+ [33]. The dataset included multi-coil and single-coil coronal proton-density-weighted scans, with and without fat suppression, from NYU Langone Health. This dataset was split by patients into three disjoint subsets for training (816 exams), hyper-parameter tuning (a.k.a., validation) (176 exams), and testing (175 exams) the models. We applied LDLR to detect Meniscal Tears and ACL Sprains.

Table 1: Number of slices in train, validation, and test sets for each dataset. Numbers in parentheses are the percentages of slices in which the disease is visible. csPCA: Clinically Significant Prostate Cancer; EV: Enlarged Ventricle.

	Knee		Abdominopelvic	Brain	
	Meniscal Tear	ACL Sprain	csPCA	EV	Mass
Train slices	29100 (11%)	29100 (3.6%)	6649 (5.0%)	11002 (1.6%)	11002 (2.0%)
Val slices	6298 (11%)	6298 (2.4%)	1431 (4.5%)	2362 (1.5%)	2362 (2.0%)
Test slices	6281 (11%)	6281 (3%)	1462 (6.0%)	2366 (2.6%)	2366 (2.7%)

Brain Scans: We used the annotated k -space data from the brain MR scans in the FASTMRI+ [33] data set, consisting of 1001 volumes, which are a subset of the full FASTMRI [32] datasets (5847 volumes). Each exam included a multi-coil axial series: T2-weighted FLAIR, T1-weighted (pre- and post-contrast) sequences. The training, validation, and test sets were of sizes 700, 150, and 151, respectively. We detected Enlarged Ventricles and Mass (including both ‘Mass’ and ‘Extra-axial Mass’ labels from [33]).

Abdominopelvic Scans: In abdominopelvic MR scans, we detected the presence of clinically significant prostate cancer (csPCA), which we defined as a lesion within the prostate for which a radiologist assigns a Prostate Imaging Reporting And Data System (PIRADS) score [40] of 3 or more. We used bi-parametric prostate MR scans from 313 patients from NYU Langone Health. Scans were acquired on a 3T Siemens scanner using a 30-element body coil. Each exam included axial T2-weighted and diffusion-weighted sequences. For our experiments, we only used T2-weighted data. A board-certified radiologist annotated each slice with PIRADS scores. Slices with PIRADS ≥ 3 were labeled as having csPCA. The data were split into training (218 volumes), validation (48), and test (47) sets, ensuring no patient overlap across the sets.

2.3 Efficacy of Learning Low-Dimensional Distributed Representations of the K-Space Data

We first evaluated whether low-dimensional distributed representations of fully sampled k -space data could support accurate disease detection. To that end, we trained the proposed KSPACEENCODER model, which uses fully sampled composite k -space matrices as input and generates their low-dimensional distributed representations as output. The composite k -space matrices were generated by combining multi-coil k -space matrices using an emulated single coil algorithm [41]. Since different coils have different spatial distributions of sensitivity to MR signal, this coil combination constitutes a modest reduction in information content, which nevertheless preserves full k -space sampling. These representations were then fed into a binary classifier to infer the presence/absence of the disease. To benchmark performance, we compared this approach against a state-of-the-art image-based DL model trained on high-fidelity composite images reconstructed from multi-coil k -space data. Specifically, we used root sum of squares (RSS) combinations of component coil images [42, 43] to generate diagnostic-quality composite images, similar to those typically stored in picture archiving and

communication systems (PACS) and reviewed by radiologists. These reconstructed images were used as input to the pre-activation ResNet-50 [44], whose output is the conditional probability of the presence of the disease. During training, we conducted a rigorous hyperparameter search to identify the best-performing model configuration. We call this model $\text{BASELINE}_{\text{OPTIMG}}$, representing the upper bound achievable by any classifier operating on diagnostic-quality images.

Table 2: The disease detection performance (Area Under the Receiver Operating Statistics Curve - AUROC) using the representations learned by KSPACEENCODER of the fully-sampled raw k -space data in comparison to the performance of the state-of-the-art DL model trained on diagnostic quality images (optimal representations in pixel-space). The results show that it is feasible to learn low-dimensional distributed representations of raw k -space data with equivalent information content to that of diagnostic quality images for binary disease detection.

	Knee		Abdomen	Brain	
	Meniscal Tear	ACL Sprain	CSPCA	Enlarged Ventricle	Mass
$\text{BASELINE}_{\text{IMGOPT}}$	92.1 ± 1.0	90.6 ± 1.0	83.1 ± 1.6	93.8 ± 1.3	88.4 ± 0.5
KSPACEENCODER	93.4 ± 0.7	90.8 ± 1.5	84.1 ± 0.4	92.3 ± 2.0	91.5 ± 1.0
p-values	0.14	0.85	0.39	0.34	0.39

Table 2 compares the disease detection performance of low-dimensional representations learned by KSPACEENCODER and a state-of-the-art DL model that uses diagnostic-quality images (i.e., optimal pixel-space representations) as input. The results demonstrate that the representations learned by KSPACEENCODER are equally effective for binary disease classification. A paired two-sided t-test showed no statistically significant difference between the two approaches. Notably, the emulated single-coil input to KSPACEENCODER has slightly lower SNR than the RSS-based images used by $\text{BASELINE}_{\text{OPTIMG}}$. In future work, we will extend the KSPACEENCODER to ingest multi-coil data directly. While our current experiments utilize data from only a single contrast sequence, incorporating data from multiple contrast sequences is a promising direction for further improved performance. Still, the results presented here are sufficient to demonstrate parity between raw-data-based and image-based inference. Figure 2a shows ROC curves comparing both models across multiple diseases across multiple anatomies.

Lastly, we qualitatively compared the information content present in the distributed representations learned by the KSPACEENCODER with that in the image-based representations. The image-based representations corresponded to the last layer of the pre-activation ResNet-50 model. We projected both sets of representations onto a two-dimensional space using the t-SNE algorithm [45]. The results are shown in Figure 2b, wherein the top row shows the t-SNE projections of the image data, and the bottom row shows the t-SNE projections of representations learned by KSPACEENCODER . Each point in the plot corresponds to the projection associated with one example, colored by class: red for disease-positive and blue for disease-negative. Well-separated clusters indicate stronger discriminative structure. The visualizations reveal that both

models produce comparably separable representations, reinforcing the conclusion that KSPACEENCODER’s raw-data-derived embeddings carry disease-relevant information comparable to image-based approaches.

2.4 Robustness of LDLR Representations to Undersampling

Table 3: AUROC of the LDLR model across different sampling rates for different pathologies. When compared to the performance using fully sampled data, classification performance even with drastic undersampling (sampling factors as low as 5%) shows minimal deterioration.

Sampling factor	5%	8%	10%	12.5%	100%
Meniscal Tear	90.3 $\pm$ 1.0	91.3 $\pm$ 1.01	90.9 $\pm$ 1.6	91.7 $\pm$ 1.3	93.4 $\pm$ 0.7
ACL Sprain	88.7 $\pm$ 0.7	88.5 $\pm$ 1.5	89.3 $\pm$ 0.4	89.4.3 $\pm$ 2.0	90.8 $\pm$ 1.5
csPCA	82.7 $\pm$ 0.7	83.4 $\pm$ 1.5	83.5 $\pm$ 0.4	83.8 $\pm$ 2.0	84.1 $\pm$ 0.4
Enlarged Ventricle	90.1 $\pm$ 0.7	91.8 $\pm$ 1.5	90.3 $\pm$ 0.4	91.3 $\pm$ 2.0	92.3 $\pm$ 2.0
Mass	87.6 $\pm$ 0.7	89.6 $\pm$ 1.5	89.3 $\pm$ 0.4	89.1 $\pm$ 2.0	88.7 $\pm$ 1.0

Next, we assessed how aggressively k -space data can be undersampled without compromising disease detection performance. Using the LDLR framework, we trained separate models at sampling rates $\alpha \in \{5\%, 8\%, 10\%, 12.5\%\}$, each time identifying the most diagnostic subset of k -space and generating task-specific representations (see the Methods section for details on how LDLR identifies the most diagnostic subset of the k -space). Figure 3a shows the ROC curves for multiple diseases using the LDLR-derived representations at different sampling rates, and Table 3 provides the corresponding AUC values. Across all pathologies, the performance remains remarkably stable, even when using drastically undersampled k -space data (5 – 8%). Qualitative inspection of the learned representations further supports this conclusion of robustness. Figure 3b, shows two-dimensional t-SNE projections of the LDLR-derived representations across different diseases for each k -space sampling rate. Discriminative structure, evident from clear clustering of positive (red) and negative (blue) cases, persists across all levels of undersampling, underscoring the efficiency of LDLR in extracting diagnostically rich information from minimal raw data.

Table 4 quantifies Sensitivity, Specificity, negative predictive value (NPV), and positive predictive value (PPV) of the disease detection task using the LDLR-derived representations at a sampling rate of 5% for the knee scans and 8% for the abdominal and brain scans. We benchmarked the LDLR-based performance against that of BASELINE_{IMG_{OPT}}, which utilizes fully sampled diagnostic-quality images. Across all metrics, despite the aggressive undersampling, LDLR achieves disease detection performance comparable to image-based models trained on diagnostic-quality images. This contrasts with findings from the FastMRI challenge [35], where models failed to reliably detect pathologies at sampling rates below 25%. We verified this claim independently by reconstructing images using a state-of-the-art reconstruction method [34] from the

Table 4: Sensitivity/Specificity/NPV/PPV on the test set for the disease detection task using the LDLR learned representations for a sampling rate of only $\alpha = 5\%$ for the knee scans and 8% for brain and prostate scans. These values are obtained using an operating point with 85% Sensitivity on the validation set. Comparison against the $\text{BASELINE}_{\text{IMG OPT}}$ shows that even at such dramatically low sampling rates, the disease detection performance of LDLR representations remain robust.

	Sensitivity	Specificity	NPV	PPV
Knee: Meniscal Tear				
LDLR (5%)	80.8 ± 3.4	84.1 ± 1.1	97.0 ± 0.3	39.8 ± 1.3
$\text{BASELINE}_{\text{IMG OPT}}$	82.8 ± 2.2	86 ± 2.27	97.4 ± 0.3	44.9 ± 4.7
Knee: ACL Sprain				
LDLR (5%)	84.8 ± 3.5	78.1 ± 2.5	99.1 ± 0.1	13.3 ± 0.7
$\text{BASELINE}_{\text{IMG OPT}}$	81.1 ± 4.4	82.2 ± 2.2	99.2 ± 0.2	14.5 ± 0.9
Abdominopelvic: csPCA				
LDLR (8%)	91.3 ± 5.3	60.8 ± 2.1	98.7 ± 0.6	13.8 ± 2.1
$\text{BASELINE}_{\text{IMG OPT}}$	93.3 ± 0.5	60.0 ± 5.4	99.2 ± 0.0	14.5 ± 1.6
Brain: Enlarged Ventricle				
LDLR (8%)	86.2 ± 4.5	75.4 ± 7.4	99.5 ± 0.1	9.1 ± 2.2
$\text{BASELINE}_{\text{IMG OPT}}$	84.9 ± 7.0	85.8 ± 7.9	99.6 ± 0.2	18.3 ± 9.7
Brain: Mass				
LDLR (8%)	87.8 ± 7.7	66.3 ± 7.2	99.5 ± 0.3	7.0 ± 1.2
$\text{BASELINE}_{\text{IMG OPT}}$	86.2 ± 6.9	72.4 ± 3.8	99.5 ± 0.3	8.1 ± 0.9

sparse subsets of k -space identified by LDLR. The results are shown in Figure 4a. Various anatomical features visible in images reconstructed from full k -space data become progressively difficult to discern in images reconstructed from increasingly undersampled data. Specifically, the prostate reconstructions exhibit increasing over-smoothing of the transition zone and blurring of the peripheral/transition zone junction, progressing to outright hallucination of prostate features at high under-sampling levels. In the knees, there is a progressive loss of osseous definition (lines in the cortex). Brain images show progressive blurring of the sulci. These findings reinforce the core insight: when the clinical goal is binary disease classification, direct inference from raw data enables far more efficient use of acquisition resources than traditional image reconstruction pipelines.

2.5 Merits of Learning Task-Specific Data and Representations

The LDLR framework bypasses image reconstruction by directly learning: (1) the subset of k -space measurements most informative for disease detection, (2) a low-dimensional representation of this data, and (3) a classifier to infer the diagnosis from that representation. An alternative image-based approach would learn to identify the subsets of k -space measurements that yield the best possible reconstructed image, and then infer the disease from that image. In our third experiment, we compared the disease detection accuracy of the two approaches. Specifically, we evaluated LDLR against a baseline classifier, denoted by $\text{BASELINE}_{\text{LOUPE:VARNET}}$, which learns a sampling

pattern optimized for image quality using LOUPE [46], reconstructs images with the state-of-the-art End-to-End VarNet [34], and then classifies disease from the reconstructed images. Results show that task-optimized data selection and representation, as done in LDLR, yield better disease detection accuracy than pipelines focused on image reconstruction, highlighting the value of learning directly for the downstream task rather than as a byproduct of imaging.

Figure 4b compares AUROC scores of LDLR and $\text{BASELINE}_{\text{LOUPE:VARNET}}$ across different sampling rates. The trend is consistent across all anatomies and diseases: while both models perform similarly when using fully sampled data, LDLR significantly outperforms the image-based baseline as sampling rates decrease. For low sampling factors (5% for knees and 8% for abdomen and brain), Table 5 compares the sensitivity/specificity/NPV/PPV of LDLR and $\text{BASELINE}_{\text{LOUPE:VARNET}}$, using an operating point calibrated to 85% sensitivity on the validation set. Here again, a significant performance gap is observed for selected tasks, with $\text{BASELINE}_{\text{LOUPE:VARNET}}$ never exceeding LDLR performance for any task.

Table 5: Sensitivity/Specificity/NPV/PPV on the test set for the disease detection task using the LDLR learned representations for a sampling rate of only $\alpha = 5\%$ for the knee scans and 8% for brain and prostate scans. These values were obtained using an operating point with 85% Sensitivity on the validation set. Comparison against the $\text{BASELINE}_{\text{LOUPE:VARNET}}$ clearly shows the superiority, especially at low sampling rates, of learning to identify and represent the data best suited for the disease detection task and using it to infer the disease as opposed to first optimizing the image reconstruction and then detecting the disease.

	Sensitivity	Specificity	NPV	PPV
Knee: Meniscal Tear				
LDLR (5%)	80.8 ± 3.4	84.1 ± 1.1	97.0 ± 0.3	39.8 ± 1.3
$\text{BASELINE}_{\text{LOUPE:VARNET}}$	81.0 ± 3.2	78.9 ± 2.8	96.8 ± 0.5	34.4 ± 2.9
Knee: ACL Sprain				
LDLR (5%)	84.8 ± 3.5	78.1 ± 2.5	99.1 ± 0.1	13.3 ± 0.7
$\text{BASELINE}_{\text{LOUPE:VARNET}}$	73.9 ± 8.2	78.1 ± 6.3	98.8 ± 0.3	11.5 ± 1.6
Abdominopelvic: csPCA				
LDLR (8%)	91.3 ± 5.3	60.8 ± 2.1	98.7 ± 0.6	13.8 ± 2.1
$\text{BASELINE}_{\text{LOUPE:VARNET}}$	89.0 ± 5.3	54.3 ± 2.1	96.0 ± 0.3	11.8 ± 2.1
Brain: Enlarged Ventricle				
LDLR (8%)	86.2 ± 4.5	75.4 ± 7.4	99.5 ± 0.1	9.1 ± 2.2
$\text{BASELINE}_{\text{LOUPE:VARNET}}$	84.8 ± 7.9	63.1 ± 14.8	99.4 ± 0.2	6.5 ± 2.2
Brain: Mass				
LDLR (8%)	87.8 ± 7.7	66.3 ± 7.2	99.5 ± 0.3	7.0 ± 1.2
$\text{BASELINE}_{\text{LOUPE:VARNET}}$	73.4 ± 3.9	55.2 ± 7.5	98.7 ± 0.2	4.4 ± 0.6

2.6 Difference in Learned Data for Different Tasks

Next, we demonstrate that, for a fixed sampling rate, the k -space subset identified by LDLR for maximizing diagnostic signal differs markedly from the one chosen by LOUPE, which optimizes image reconstruction. Figure 5 compares the two: each vertical line represents one k -space line, with central lines encoding low frequency measurements and peripheral ones capturing high frequency measurements. LOUPE (bottom row) prioritizes low-frequency lines, consistent with its pixel-wise loss objective. In contrast, LDLR (top row) selects lines, regardless of frequency, that best support diagnostic inference, highlighting its task-specific focus.

2.7 Qualitative Comparison and Reader Study

To further test our hypothesis that pixel-space representations may be suboptimal for diagnostic inference, especially under conditions of aggressive undersampling, we conducted a structured reader study with board-certified radiologists. We intentionally selected a particularly challenging clinical task: detecting prostate cancer from a single axial slice of a single-contrast image, removing the benefits of volumetric context and multi-modal cues, thereby isolating the diagnostic signal available in a constrained pixel-space representation. Specifically, we utilized only T2-weighted images, which are widely acknowledged to be less informative than diffusion-weighted imaging for assessing prostate cancer. The difficulty of the task is reflected in the relatively low sensitivity achieved by expert radiologists. This challenge is further exacerbated by under-sampling the k -space, which degrades image quality and makes the limitations of human interpretation in pixel-space even more pronounced (see Table 6).

Two fellowship-trained abdominal radiologists with 17 and 10 years of experience in interpreting abdominal MRI, respectively, were recruited to assess a set of 60 prostate MR images independently. The dataset was balanced, comprising an equal number of cases with clinically significant lesions ($\text{PIRADS}_{v2.1} \geq 3$) and without ($\text{PIRADS}_{v2.1} < 3$), with images randomly selected from a held-out test set. Each image was reconstructed from a 10% subset of k -space data identified by LDLR, using a state-of-the-art deep learning-based reconstruction model [47]. Radiologists responded to three questions for each image:

1. *Diagnostic Quality Assessment*: Rate the overall diagnostic quality of the image on a five-point Likert scale (1: non-diagnostic, 2: poor, 3: moderate, 4: good, 5: excellent).
2. *Lesion Detection*: Indicate the presence/absence of a lesion with a binary response.
3. *cSPCA Inference*: Estimate the $\text{PIRADS}_{v2.1}$ score for each identified lesion. To simplify the task, we instructed the readers to identify only a binary category for the $\text{PIRADS}_{v2.1}$ score: whether the score is 1 – 2 or between 3 – 5.

After a one-month washout, the same 60 images were presented to the radiologists, now reconstructed from fully sampled k -space data using the standard algorithm. They repeated the diagnostic tasks, blinded to both the reconstruction method and the image duplication. Before the study, readers received brief training with unlabeled examples from both techniques. Study results are summarized in Table 6.

Diagnostic Quality Assessment: A one-sided paired Wilcoxon signed-rank test was employed to compare the perceived diagnostic quality of images reconstructed from fully sampled k -space data with those reconstructed using the LDLR-selected k -space subset. The null hypothesis posited that the diagnostic quality of LDLR-based reconstructions was non-inferior to that of the standard, fully sampled reconstructions. Statistical significance was defined as $p < 0.05$. As summarized in Table 6 (‘Quality Assessment’ row), both radiologists reported statistically significant differences in image quality assessments, with LDLR-based reconstructions (‘Partial Sampling’ column in the table) consistently rated lower. These results indicate a clear preference for reconstructions derived from fully sampled k -space data.

Lesion Detection: The results of this task are presented in the Lesion Detection rows of Table 6. The proposed LDLR model substantially outperforms both readers in lesion-detection accuracy (‘Acc’) on images reconstructed from partially sampled data. For images reconstructed from fully sampled data, LDLR matches the performance of one reader and surpasses that of the second. On partially sampled data, LDLR also achieves markedly higher sensitivity while maintaining a comparable level of specificity to both readers. For fully sampled data, its sensitivity and specificity fall within the readers’ margins of error. Notably, these performance gains are achieved despite LDLR being trained on only 10% of the data, whereas the readers evaluated clinical-grade images.

A similar pattern emerges for both expert readers and the LDLR model in the csPCA inference task (see the csPCA Inference row in Table 6). The diagnostic accuracy of both readers declines markedly when interpreting images reconstructed from partially sampled k -space data compared with fully sampled acquisitions. In contrast, the performance of LDLR remains comparable to that of human readers evaluating fully sampled images, underscoring its robustness to undersampling. On partially sampled data, LDLR shows higher sensitivity while maintaining specificity comparable to both readers. For fully sampled data, when matched for sensitivity, LDLR exhibits slightly lower specificity than the readers. Although LDLR is not able to match the specificity of an experienced subspecialized radiologist interpreting clinical-grade images, the model’s specificity remains substantially higher than that of the PSA test, the current standard-of-care screening tool for prostate cancer. Importantly, this specificity can be further improved by incorporating patient-specific variables, such as PSA density and additional clinical information.

3 Discussion

In pursuit of democratizing scientific discovery across domains, we introduce a machine learning framework that significantly lowers the barrier to detecting Entities of Interest (EOIs) in complex environments. Rather than relying on traditional image reconstruction pipelines, our approach infers the presence or absence of EOIs directly from intelligently selected subsets of raw environment-probing measurements. Using MRI as a case study, we demonstrated that accurate detection of disease (the EOIs) can be achieved with substantially reduced data requirements when the goal is diagnosis, rather than anatomical reconstruction and visualization. Our method learns to isolate

Table 6: Summary of reader study outcomes. Readers interpreted images reconstructed from 10% undersampled k -space data identified by LDLR (‘Partial Sampling’) as well as images reconstructed using the fully sampled k -space data (‘Full Sampling’). For the quality assessment task, reader scores and P-values between scores for partial and full sampling are reported. For lesion detection and csPCA inference, accuracy, sensitivity, and specificity are compared between readers and LDLR.

Task		Reader 1		Reader 2		LDLR
		Partial Sampling	Full Sampling	Partial Sampling	Full Sampling	
Quality Assessment	Score	1.8 ± 0.6	2.9 ± 0.8	1.9 ± 0.6	3.2 ± 0.9	
	P-value	5.6×10^{-7}		1.2×10^{-7}		
Lesion Detection	Acc (%)	44 ± 6	58 ± 7	47 ± 6	47 ± 7	57
	Sen.	17 ± 7	50 ± 9	17 ± 7	23 ± 8	42
	Spec.	70 ± 8	66 ± 9	77 ± 8	71 ± 8	68
csPCA Inference	Acc (%)	43 ± 6	63 ± 6	48 ± 6	60 ± 6	59
	Sen.	16 ± 7	33 ± 9	13 ± 6	27 ± 6	30
	Spec.	70 ± 8	93 ± 5	83 ± 7	94 ± 7	75

the most diagnostically informative subsets of the measurement space (referred to as the “disease signature”), and encodes them into a compact latent representation. These representations enable downstream classifiers to achieve high detection performance while bypassing the computational burden and potential information loss of full image reconstruction. This paradigm shift — from reconstructing data to directly interpreting it — unlocks new efficiencies in medical imaging pipelines and may be generalized to other sensing-based domains. We validated our approach on a large-scale retrospective dataset and through a reader study with sub-specialized radiologists. Our results highlight the potential of learning-based signal interpretation to transform medical diagnosis by significantly reducing the resource burden within the diagnostic pipeline.

While this work marks an important step towards democratizing disease discovery in medical imaging, several challenges remain before realizing a widely accessible, real-world clinical system that operates outside specialized imaging environments at a population scale. First, most clinical MR scanners acquire data in parallel using multiple coils to accelerate acquisition and increase signal-to-noise (SNR) ratio. In this study, we simplified the data by simulating a single-coil setup using the ESC algorithm [41]. Future efforts will extend LDLR to handle native multi-coil data, which we expect will further improve performance due to the higher effective SNR associated with the signal. Second, clinical MRI typically relies on multiple sequences (e.g., T2-weighted and diffusion-weighted for prostate scans) to capture different tissue contrasts and provide the radiologist with multiple visual facets of the same underlying anatomy. Our current approach uses data from only a single contrast sequence. Expanding LDLR to incorporate multi-sequence k -space data will better reflect clinical practice and likely enhance diagnostic performance. Lastly, the current study utilized high-quality k -space data from an expensive, high-field 3 T scanner (the current standard of care), housed in specialized imaging environments. True democratization of disease detection requires not only reducing data demands from high-field imaging systems

but also robustness to lower-quality data from affordable, portable alternatives like ultra-low-field MRI scanners with in homogeneous fields. These devices, which can possibly operate outside specialized imaging environments, offer broad accessibility but produce noisier, more distorted data. Crucially, by eliminating the requirement for high-fidelity image reconstruction, our framework is well-positioned to generalize to these challenging data regimes. We are actively extending our approach to accommodate such low-quality acquisition settings, and this represents a key direction for future work aimed at expanding the reach and equity of AI-driven medical diagnostics.

While this study demonstrated the effectiveness of our framework within the MR imaging pipeline, the core principles are broadly applicable across a range of medical imaging modalities and beyond. In low-dose CT, for example, it can be adapted to detect disease signatures in the sinogram domain. In this context, the disease signature would correspond to the minimal set of projection-dependent radiation doses required to enable reliable inference, independent of whether the reconstructed images meet diagnostic quality standards. Similarly, in ultrasound (US) imaging, our approach could be extended to analyze raw, temporal multi-channel acoustic data, enabling the detection of subtle pathologies that may elude visual identification in conventional images. Beyond medicine, this framework has compelling relevance in other scientific domains that rely on indirect signal measurements to infer latent phenomena. In radio astronomy, for instance, which similarly involves signal acquisition in a Fourier-like (UV) space, our methodology could facilitate the detection of astrophysical events using observatories with suboptimal spatial configurations or lower-cost, reduced-sensitivity detectors, thereby broadening access and operational flexibility without compromising the ability to identify signals of interest in interstellar observations.

4 Methods

4.1 Background and Notation

MR imaging is an indirect process whereby spatially resolved images of a human subject’s anatomy are reconstructed from the frequency space (a.k.a., k -space) measurements of the electromagnetic activity inside the subject’s body under the influence of a magnetic field and radio-frequency pulses. These measurements are captured by an instrument called a *receiver coil*, which is kept in the vicinity of the part of the body whose image is sought. The k -space measurements from a single-coil are represented as a 2-dimensional complex-valued matrix $\mathbf{x} \in \mathbb{C}^{r \times c}$, where r is the number of rows and c is the number of columns. The spatial image $\mathbf{z}$ is reconstructed from the k -space matrix by a multi-dimensional inverse Fourier transform, $\mathbf{z} = \mathcal{F}^{-1}(\mathbf{x})$. We denote by $\mathbf{y} \in \{1, \dots, K\}$ the clinically relevant response. In our case, $\mathbf{y}$ will be a binary response variable ($\mathbf{y} \in \{1, 0\}$) indicating the presence/absence of the disease.

4.1.1 Multi-Coil Data

In practice, most clinical MR scanners acquire measurements in parallel using multiple receiver coils to speed up data acquisition and enhance signal-to-noise ratio (SNR). In case of multi-coil acquisition, the k -space matrix $\mathbf{x}^{mc}$ is 3-dimensional: $\mathbf{x}^{mc} \in$

691 $\mathbb{C}^{d_c \times r \times c}$ [32], where d_c is the number of coils. The image produced by each coil shows a
692 slightly different view of the anatomy since each coil is sensitive to signals arising from
693 different spatial locations. A variety of methods have been proposed that aggregate
694 data from multiple coils to make them suitable for human interpretation or ingestion
695 by a downstream ML model. A commonly used method that combines images from
696 different coils into a single aggregate image is called the root-sum-of-squares (RSS)
697 [43]. Given the multi-coil k -space matrix $\mathbf{x}^{mc}$, the RSS method first computes the
698 inverse Fourier transform of each coil’s k -space matrix $\tilde{\mathbf{m}}_j = \mathcal{F}^{-1}(\mathbf{x}_j^{mc})$, where $\mathbf{x}_j^{mc}$ is
699 the k -space matrix of the j -th coil. It then generates the image $\tilde{\mathbf{m}}$, by computing the
700 following quantity:

$$\tilde{\mathbf{m}} = \sqrt{\sum_{j=1}^{N_c} |\tilde{\mathbf{m}}_j|^2}.$$

706 In contrast, another method called Emulated Single Coil (ESC) [41] aggregates this
707 multi-coil data in the original k -space and emulates it to come from a single coil. This
708 process reduces the dimension of the matrix $\mathbf{x}^{mc} \in \mathbb{C}^{d_c \times r \times c}$ to a matrix $\tilde{\mathbf{x}}^{mc} \in \mathbb{C}^{r \times c}$.
709 In the subsequent discussion about LDLR, we will assume that we are working with
710 the ESC data matrix $\tilde{\mathbf{x}}^{mc}$ of dimensions $r \times c$.

711

712 4.1.2 Under-Sampled Data

713

714 The notion of “under-sampling” referred to measuring only a subset of entries in the
715 k -space matrix $\mathbf{x}$. We represented the *sampling pattern* using a binary mask matrix
716 $\mathbf{s} \in \{0, 1\}^{r \times c}$ (a.k.a., *sampling mask*), where $\mathbf{s}_{ij} = 1$ if and only if the measurement $\mathbf{x}_{ij}$
717 had been acquired. The under-sampled k -space matrix was represented as $\mathbf{x}_s = \mathbf{x} \circ \mathbf{s}$,
718 where $\circ$ is element-wise multiplication. In this work, we constrained the sampling
719 pattern to “Cartesian”, which consisted of sampling the k -space matrix line-wise. More
720 specifically, for a Cartesian sampling pattern, all the elements of some lines of the
721 matrix $\mathbf{s}$ were set to 0, and all the elements of other lines were set to 1. The k -space
722 matrix had its origin in the center of the matrix. The sampling rate α was defined as
723 the total percentage of measurements.

724

725 4.2 Image-Based DL Models for Disease Identification

726

727 Conventionally, using DL models within the MR pipeline for disease detection involved
728 taking the MR images as input and generating the final answer. Ideally, the images
729 were reconstructed using the full k -space data. The image-based DL method trained
730 on images reconstructed from full k -space data was denoted by $\text{BASELINE}_{\text{IMGOPT}}$.
731 This established an upper bound on the disease detection accuracy. Since generating
732 high-fidelity cross-sectional images required acquiring large quantities of high-quality
733 k -space measurements, researchers proposed to train image-based DL classifiers using
734 images reconstructed from undersampled k -space data. This approach required one to
735 make decisions along two dimensions, namely 1) choosing the sampling pattern over
736 the k -space, the data corresponding to which would be used to reconstruct the image,

and 2) choosing a method to reconstruct the image, given any sampling pattern. Multiple methods have been proposed to learn the sampling pattern [46, 48, 49], and to reconstruct images using the under-sampled k -space data [34, 35, 50–55]. We denoted the class of these models by $\text{BASELINE}_{\langle \text{SAMP} \rangle : \langle \text{RECON} \rangle}$, where $\langle \text{SAMP} \rangle$ refers to the method used to select the sampling pattern and $\langle \text{RECON} \rangle$ refers to the image reconstruction method of choice. Specifically, we compared LDLR’s performance against a pipeline that uses LOUPE [46] to learn the sampling pattern and VarNet [56] to reconstruct the image. This pipeline was denoted by $\text{BASELINE}_{\text{LOUPE:VARNET}}$ (see the Results section).

4.3 The kSpaceEncoder Architecture

Let $\mathbf{x} \in \mathbb{C}^{d \times d}$ denote the complex-valued matrix storing the k -space data acquired using a single contrast sequence processed using the ESC algorithm [41]. Our objective was to map this high-dimensional complex-valued k -space matrix $\mathbf{x}$ to a point on a low-dimensional embedding space, such that the resulting captures maximal diagnostic information contained in $\mathbf{x}$. We achieved this objective by proposing a novel DL model that took the k -space data $\mathbf{x}$ as input and generated its low-dimensional distributed representation. Our novel DL approach, which we call the kSPACEENCODER , was necessary as standard Convolutional Neural Networks (CNNs) [36–39] expect input $\mathbf{x}$ to exhibit strong local spatial structure, a property not inherent in k -space matrix where each element encodes some information of the entire spatial image.

The kSPACEENCODER model, depicted in Figure 1 Part D, used a complex learnable kernel matrix $Z'_{\mathcal{F}}$, which was of the same size as $\mathbf{x}$ and performed element-wise multiplication with $\mathbf{x}$. The size of $Z'_{\mathcal{F}}$ was the same as the matrix X , and it was derived by taking a $k \times k$ kernel matrix Z in image space, zero padding it to create a matrix Z' , of the same size as $\mathbf{x}$, and then passing it through a Fourier transform operation to generate the complex-valued kernel matrix $Z'_{\mathcal{F}}$ in the frequency space. The output of the element-wise multiplication operation between $\mathbf{x}$ and $Z'_{\mathcal{F}}$, denoted by I , underwent an inverse Fourier transform to produce a set of feature maps M . These feature maps served as input to a downstream deep neural network to generate the low-dimensional distributed representations $\mathbf{h}$. Finally, these representations $\mathbf{h}$ were passed into another classifier to generate the disease score $\hat{\mathbf{y}}$. The works of LeCun et al., inspired our network architecture, [57, 58], where they used a similar sequence of operations to speed computations in CNNs by performing the convolution operation in the frequency domain.

4.4 Pipeline for Learning Data to Learn Representations

For a specific disease, the LDLR framework (depicted in Figure 1 Part C) first learned to identify the diminutive subset of the k -space that stores the maximal information pertaining to the disease diagnosis. We call this the “disease signature” in k -space. It then learned compact distributed representations of the k -space data associated with this identified disease signature. Finally, it used these representations to infer the presence/absence of the disease. This is in contrast to the image reconstruction approaches, where the requirement to generate a high-quality cross-sectional image

of the entire anatomy necessitates the sampling of a large portion of the k -space. To the best of our knowledge, LDLR is the first method to propose classifying diseases directly from a carefully chosen (learned) subset of the k -space. This pipeline was realized by executing the following three steps (Figure 1 Part C):

1. Learning a “score function” that, for any given subset of the k -space, estimated the mutual information between the k -space measurements associated with the subset and the presence/absence of the disease.
2. Identifying the “disease signature”, by using the trained “score function” to identify the subset(s) of the k -space with the highest score (highest mutual information between the k -space data associated with it and the presence of disease).
3. Learning the distributed representations of the k -space data associated with the “disease signature” identified in the previous step, and using it to detect the disease.

4.4.1 Learning the “Score Function”

The objective of the “Score Function”, denoted by V , was to accurately and efficiently estimate the mutual information between the clinically relevant variable $\mathbf{y}$ (indicating the presence/absence of the disease) and the k -space data $\mathbf{x}_s$ associated with any subset of the k -space $\mathbf{s}$. This mutual information between $\mathbf{y}$ and $\mathbf{x}_s$ was used as a surrogate for the information content in $\mathbf{x}_s$ for disease inference.

Specifically, let $I(\mathbf{y}; \mathbf{x}_s)$ denote the mutual information between $\mathbf{y}$ and $\mathbf{x}_s$. Then according to [59], $I(\mathbf{y} | \mathbf{x}_s)$ can be defined as:

$$I(\mathbf{y}; \mathbf{x}_s) = \mathbb{E}_{\mathbf{x}_s} \text{KL} (p(\mathbf{y} | \mathbf{x}_s) || p(\mathbf{y})) \quad (1)$$

$$= \mathbb{E}_{\mathbf{x}_s} \mathbb{E}_{\mathbf{y} | \mathbf{x}_s} \log p(\mathbf{y} | \mathbf{x}_s) - \log p(\mathbf{y}) \quad (2)$$

$$= \mathbb{E}_{\mathbf{x}_s} \mathbb{E}_{\mathbf{y} | \mathbf{x}_s} \log p(\mathbf{y} | \mathbf{x}_s) - C, \quad (3)$$

where C is a constant independent of the sampling pattern $\mathbf{s}$, and $p(\mathbf{y} | \mathbf{x}_s)$ and $p(\mathbf{y})$ are the conditional and the marginal distribution of the response variable, respectively. According to equation 3, we could estimate $I(\mathbf{y} | \mathbf{x}_s)$ if we were able to estimate the value of $p(\mathbf{y} | \mathbf{x}_s)$. Since we did not have access to the true conditional distribution $p(\mathbf{y} | \mathbf{x}_s)$, we approximated the expected conditional log-likelihood by learning a probabilistic model $q(\mathbf{y} | \mathbf{x}_s; \lambda)$ parameterized by λ . However, learning a model for every sampling pattern $\mathbf{s}$ would be infeasible even for moderately high dimensions. To address this issue, we drew upon the works of [60, 61], where the authors showed that at optimality, a single model $q_{\text{val}}(\mathbf{y} | \mathbf{x}_s; \lambda)$ trained with independently generated sampling patterns that are drawn independent of the data $\mathbf{x}, \mathbf{y}$, is equivalent to a conditional distribution of $\mathbf{y}$ for each sampling pattern. This approach contrasts with the methods that explicitly model $\mathbf{x}$ [62] and have been used in other applications [63]. As such, we trained a model q_{val} by minimizing the following loss function:

$$\mathcal{L}(\lambda) = -\mathbb{E}_{\mathbf{x}, \mathbf{y}} \mathbb{E}_{\mathbf{s} \sim \pi} \log q_{\text{val}}(\mathbf{y} | \mathbf{x}_s; \lambda),$$

where π is a distribution over the sampling pattern that is independent of the data $\mathbf{x}, \mathbf{y}$. In LDLR, the distribution π was a one-dimensional distribution, and the function q_{val}

Algorithm 1 Estimating the conditional likelihood $q_{\text{val}}(\mathbf{y} \mid \mathbf{x}_{\mathbf{s}})$

- 1: **Input:** Training data set $\mathcal{D}_{\text{tr}} = \{(\mathbf{x}^i, \mathbf{y}^i)\}_{i=1}^{N_{\text{tr}}}$, model $q_{\text{val}}(\mathbf{y} \mid \mathbf{x}; \lambda)$ with initial parameters λ , mini-batch size M , acceleration factor α , and prior distribution π over sampling patterns
- 2: **Return:** Trained model $q_{\text{val}}(\mathbf{y} \mid \mathbf{x}_{\mathbf{s}}; \lambda^*)$
- 3: **while** not converged **do**
- 4: Sample a mini-batch of training points of size M
- 5: Draw a sampling pattern $\mathbf{s} \sim \pi$, such that $\frac{r \times c}{\|\mathbf{s}\|_0} = \alpha$
- 6: Update the model parameters

$$\lambda_{t+1} = \lambda_t + \frac{\gamma}{M} \sum_{i=1}^M \nabla_{\lambda} \log q_{\text{val}}(\mathbf{y}^i \mid \mathbf{x}_{\mathbf{s}}; \lambda_t)$$

- 7: **end while**
 - 8: Return the trained model $q_{\text{val}}(\mathbf{y} \mid \mathbf{x}_{\mathbf{s}}; \lambda^*)$
-

was parameterized using the KSPACEENCODER architecture (Section 4.3). The under-sampled data $\mathbf{x}_{\mathbf{s}}$ was created by masking the fully-sampled matrix $\mathbf{x}$ with a mask $\mathbf{s} \in \{0, 1\}^{r \times c}$. This masking ensured that the same model can be used as the input’s dimensions are fixed. This process is summarized in Algorithm 1. Finally, the trained q_{val} was used to define a scoring function $V : \{0, 1\}^{r \times c} \rightarrow \mathbb{R}$, which for each sampling pattern $\mathbf{s}$ estimated the mutual information between that subset of the k -space up to a constant 3. Specifically,

$$V(\mathbf{s}) = \mathbb{E}_{\mathbf{x}} \mathbb{E}_{\mathbf{y} \mid \mathbf{x}} \log q_{\text{val}}(\mathbf{y} \mid \mathbf{x}_{\mathbf{s}}; \lambda). \quad (4)$$

The higher the score achieved by a sampling pattern, the higher its diagnostic signal.

4.4.2 Identifying the “Disease Signature”

Once the score function $V : \{0, 1\}^{r \times c} \rightarrow \mathbb{R}$ was trained, it was straightforward to identify the “disease signature”. Specifically, for a given sampling rate α , LDLR found the sampling pattern $\mathbf{s}^*$ (the disease signature) that had the maximum mutual information with the response variable $\mathbf{y}$ by optimizing:

$$\mathbf{s}^* = \arg \max_{\mathbf{s} \in \{0, 1\}^{r \times c}} I(\mathbf{y} \mid \mathbf{x}_{\mathbf{s}}) = \arg \max_{\mathbf{s} \in \{0, 1\}^{r \times c}} V(\mathbf{s}), \quad (5)$$

where $\alpha = \frac{r \times c}{\|\mathbf{s}\|_0}$ and $\mathbf{s}$ is the binary mask matrix of dimensions $r \times c$. In practice, $\mathbf{s}^*$ was approximated by a Monte Carlo search within the space of all sampling patterns. N candidate sampling patterns were drawn from the prior distribution π . Each drawn pattern was scored by the scoring function V , and the pattern with the highest score was selected as $\mathbf{s}^*$. The details of the full algorithm are provided in Algorithm 2.

Algorithm 2 Identifying the disease signature $\mathbf{s}^*$

- 1: **Input:** Validation data set $\mathcal{D}_{\text{val}} = \{(\mathbf{x}^i, \mathbf{y}^i)\}_{i=1}^{N_{\text{val}}}$, model $q_{\text{val}}(\mathbf{y} \mid \mathbf{x}; \lambda^*)$ with parameters λ , acceleration factor α , number of candidate sampling patterns to generate N , and prior distribution π over the sampling patterns
- 2: **Return:** Sampling pattern $\mathbf{s}^*$
- 3: **for** $j \in \{1, \dots, N\}$ **do**
- 4: Sample $\mathbf{s}_j \sim \pi$ such that $\frac{r \times c}{\|\mathbf{s}\|_0} = \alpha$
- 5: Estimate the mutual information score in 4 as follows

$$\hat{V}(\mathbf{s}_j) = \frac{1}{N_{\text{val}}} \sum_{i=1}^{N_{\text{val}}} \log q_{\text{val}}(\mathbf{y}^i \mid \mathbf{x}_{\mathbf{s}_j}^i; \lambda^*) \quad (6)$$

- 6: **end for**
 - 7: Let $\mathbf{s}^* = \arg \max_{j \in \{1, \dots, N\}} \hat{V}(\mathbf{s}_j)$
-

4.4.3 Learning the Representations and Detecting the Disease

Given the disease signature $\mathbf{s}^*$, the low-dimensional distributed representations of the associated k -space data and the likelihood of the presence/absence of the disease were obtained by training the KSPACEENCODER model along with the classifier in an end-to-end manner as described in Algorithm 1. During the training of this system, for every mini-batch, we randomly sampled a different sampling pattern from the distribution π . Through our experiments, we found that this is, in fact, the key to training robust representations and a reliable classifier. We also explored retraining a classifier using data $\mathbf{x}_{\mathbf{s}^*}$, obtained from a fixed classification optimized sampling pattern $\mathbf{s}^*$. However, the former approach worked better.

4.4.4 Executing Training of the LDLR Pipeline

We trained each classifier on 1 NVIDIA V100 GPU for 50 epochs with the RESNET50 model [37] as the downstream convolutional network. We used a batch size of 256 and optimized using Adam with the learning rate and weight decay selected over a grid with values in $\{10^{-3}, 10^{-4}, 10^{-5}\}$. We optimized using the weighted cross-entropy loss and selected models with the lowest loss computed over the validation set. We reported metrics with a 95% confidence interval computed using 5 seeds 0, 1, 2, 3, 4.

References

- [1] Rosenkrantz, A. *et al.* Prostate Cancer: Multiparametric MRI for Index Lesion Localization—A Multiple-Reader Study. *AJR Am J Roentgenol* **199**, 830–837 (2012).
- [2] Garcia-Reyes, K. *et al.* Detection of prostate cancer with multiparametric MRI (mpMRI): effect of dedicated reader education on accuracy and confidence of index and anterior cancer diagnosis. *Abdominal imaging* **40**, 134–142 (2015).

- [3] Rastinehad, A. R. *et al.* Improving detection of clinically significant prostate cancer: magnetic resonance imaging/transrectal ultrasound fusion guided prostate biopsy. *The Journal of urology* **191**, 1749–1754 (2014). 921–924
- [4] Brown, L. C. *et al.* Multiparametric MRI to improve detection of prostate cancer compared with transrectal ultrasound-guided prostate biopsy alone: the PROMIS study. *Health technology assessment (Winchester, England)* **22**, 1–176 (2018). 925–927
- [5] Thompson, J. *et al.* The Diagnostic Performance of Multiparametric Magnetic Resonance Imaging to Detect Significant Prostate Cancer. *J Urol.* 2016;195(5):1428-1435. *Journal of Urology* **195**, 1428–1435 (2016). 928–931
- [6] Wysock, J. *et al.* Predictive Value of Negative 3T Multiparametric Magnetic Resonance Imaging of the Prostate on 12-core Biopsy Results. *BJU Int.* **118**, 515–520 (2016). 932–935
- [7] MacRedmond, R. *et al.* Screening for lung cancer using low dose CT scanning: results of 2 year follow up. *Thorax* **61**, 54–56 (2006). 936–937
- [8] Stracci, F., Zorzi, M. & Grazzini, G. Colorectal cancer screening: tests, strategies, and perspectives. *Frontiers in public health* **2**, 210 (2014). 938–941
- [9] Marmot, M. G. *et al.* The benefits and harms of breast cancer screening: an independent review. *British journal of cancer* **108**, 2205–2240 (2013). 942–944
- [10] Ilic, D. *et al.* Prostate cancer screening with prostate-specific antigen (PSA) test: a systematic review and meta-analysis. *Bmj* **362** (2018). 945–947
- [11] Winawer, S. J. *et al.* Colorectal cancer screening: clinical guidelines and rationale. *Gastroenterology* **112**, 594–642 (1997). 948–949
- [12] Hao, J. *et al.* Early detection of dementia through retinal imaging and trustworthy AI. *npj Digital Medicine* **7**, 294 (2024). URL <https://doi.org/10.1038/s41746-024-01292-5>. 950–953
- [13] Stang, A. & Jöckel, K.-H. The impact of cancer screening on all-cause mortality: what is the best we can expect? *Deutsches Ärzteblatt International* **115**, 481 (2018). 954–957
- [14] Kansal, A. *et al.* Cost-Effectiveness of a Multicancer Early Detection Test in the US. *The American Journal of Managed Care* **30**, e352–e358 (2024). URL <https://www.ajmc.com/view/cost-effectiveness-of-a-multicancer-early-detection-test-in-the-us><https://www.ajmc.com/view/cost-effectiveness-of-a-multicancer-early-detection-test-in-the-us>. 958–963
- [15] Elsink, K. *et al.* Cost and impact of early diagnosis in primary immunodeficiency disease: A literature review. *Clinical Immunology* **213**, 108359 (2020). URL 964–966

- 967 <https://www.sciencedirect.com/science/article/pii/S152166161930186X>.
- 968
- 969 [16] Ratushnyak, S., Hoogendoorn, M. & van Baal, P. H. M. Cost-Effectiveness of
 970 Cancer Screening: Health and Costs in Life Years Gained. *American Journal of*
 971 *Preventive Medicine* **57**, 792–799 (2019). URL [https://doi.org/10.1016/j.amepre.](https://doi.org/10.1016/j.amepre.2019.07.027)
 972 [2019.07.027](https://doi.org/10.1016/j.amepre.2019.07.027).
- 973
- 974 [17] Slatkoff, S., Gamboa, S., Zolotor, A. J., Mounsey, A. L. & Jones, K. PSA testing:
 975 When it’s useful, when it’s not. *The Journal of family practice* **60**, 357 (2011).
- 976
- 977 [18] Elmore, J. G. *et al.* Ten-year risk of false positive screening mammograms and
 978 clinical breast examinations. *New England Journal of Medicine* **338**, 1089–1096
 979 (1998).
- 980
- 981 [19] Kilpeläinen, T. P. *et al.* False-positive screening results in the Finnish prostate
 982 cancer screening trial. *British journal of cancer* **102**, 469–474 (2010).
- 983
- 984 [20] Fenton, J. *et al.* Prostate-Specific Antigen-Based Screening for Prostate Cancer:
 985 Evidence Report and Systematic Review for the US Preventive Services Task
 986 Force. *JAMA* **319**, 1914–1931 (2018).
- 987
- 988 [21] Lafata, J. E. *et al.* The economic impact of false-positive cancer screens. *Cancer*
 989 *Epidemiology and Prevention Biomarkers* **13**, 2126–2132 (2004).
- 990
- 991 [22] Brodersen, J. & Siersma, V. D. Long-term psychosocial consequences of false-
 992 positive screening mammography. *The Annals of Family Medicine* **11**, 106–115
 993 (2013).
- 994
- 995 [23] Taksler, G. B., Keating, N. L. & Rothberg, M. B. Implications of false-positive
 996 results for future cancer screenings. *Cancer* **124**, 2390–2398 (2018).
- 997
- 998 [24] Eldred-Evans, D. *et al.* Population-Based Prostate Cancer Screening With Mag-
 999 netic Resonance Imaging or Ultrasonography: The IP1-PROSTAGRAM Study.
 1000 *Jama Oncology* **7**, 395 – 402 (2021).
- 1001
- 1002 [25] Wallis, C. J. D., Haider, M. A. & Nam, R. K. Role of mpMRI of the prostate in
 1003 screening for prostate cancer. *Translational andrology and urology* **6**, 464 (2017).
- 1004
- 1005 [26] Marino, M. A., Helbich, T., Baltzer, P. & Pinker-Domenig, K. Multiparametric
 1006 MRI of the breast: A review. *Journal of Magnetic Resonance Imaging* **47**, 301–315
 1007 (2018).
- 1008
- 1009 [27] Biederer, J. *et al.* Screening for lung cancer: Does MRI have a role? *European*
 1010 *journal of radiology* **86**, 353–360 (2017).
- 1011
- 1012 [28] Sali, L. & Regge, D. CT colonography for population screening of colorectal
 cancer: hints from European trials. *The British Journal of Radiology* **89**, 20160517
 (2016).

- [29] Alkabbany, I., Ali, A. M., Mohamed, M., Elshazly, S. M. & Farag, A. An AI-Based Colonic Polyp Classifier for Colorectal Cancer Screening Using Low-Dose Abdominal CT. *Sensors* **22**, 9761 (2022). 1013
1014
1015
1016
- [30] Sathianathen, N. J. *et al.* Negative Predictive Value of Multiparametric Magnetic Resonance Imaging in the Detection of Clinically Significant Prostate Cancer in the Prostate Imaging Reporting and Data System Era: A Systematic Review and Meta-analysis. *European urology* **78**, 402–414 (2020). 1017
1018
1019
1020
1021
- [31] Hugosson, J. *et al.* Prostate Cancer Screening with PSA and MRI Followed by Targeted Biopsy Only. *New England Journal of Medicine* **387**, 2126–2137 (2022). 1022
1023
URL <https://doi.org/10.1056/NEJMoa2209454>. 1024
1025
- [32] Zbontar, J. *et al.* fastMRI: An Open Dataset and Benchmarks for Accelerated {MRI}. *CoRR* **abs/1811.0** (2018). 1026
1027
- [33] Zhao, R. *et al.* fastMRI+: Clinical Pathology Annotations for Knee and Brain Fully Sampled Multi-Coil MRI Data. *arXiv preprint arXiv:2109.03812* (2021). 1028
1029
1030
- [34] Sriram, A. *et al.* *End-to-end variational networks for accelerated MRI reconstruction*, 64–73 (Springer, 2020). 1031
1032
1033
- [35] Muckley, M. J. *et al.* Results of the 2020 fastmri challenge for machine learning mr image reconstruction. *IEEE Transactions on Medical Imaging* **40**, 2306–2317 (2021). 1034
1035
1036
1037
- [36] LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015). 1038
URL <https://doi.org/10.1038/nature14539>. 1039
1040
- [37] He, K., Zhang, X., Ren, S. & Sun, J. *Deep Residual Learning for Image Recognition*, 770–778 (2016). 1041
1042
1043
- [38] Long, J., Shelhamer, E. & Darrell, T. *Fully convolutional networks for semantic segmentation*, 3431–3440 (IEEE Computer Society, Los Alamitos, CA, USA, 2015). URL <https://doi.ieeecomputersociety.org/10.1109/CVPR.2015.7298965>. 1044
1045
1046
- [39] He, K., Gkioxari, G., Dollár, P. & Girshick, R. *Mask R-CNN*, 2980–2988 (2017). 1047
1048
- [40] Weinreb, J. C. *et al.* PI-RADS prostate imaging–reporting and data system: 2015, version 2. *European urology* **69**, 16–40 (2016). 1049
1050
1051
- [41] Tygert, M. & Zbontar, J. Simulating single-coil MRI from the responses of multiple coils. *Communications in Applied Mathematics and Computational Science* **15**, 115–127 (2020). 1052
1053
1054
1055
- [42] Roemer, P. B., Edelstein, W. A., Hayes, C. E., Souza, S. P. & Mueller, O. M. The NMR phased array. *Magnetic Resonance in Medicine* **16** (1990). URL <https://api.semanticscholar.org/CorpusID:9482029>. 1056
1057
1058

1059 [43] Larsson, E. G., Erdogmus, D., Yan, R., Principe, J. C. & Fitzsimmons, J. R. SNR-
1060 optimality of sum-of-squares reconstruction for phased-array magnetic resonance
1061 imaging. *Journal of Magnetic Resonance* **163**, 121–123 (2003).
1062

1063 [44] He, K., Zhang, X., Ren, S. & Sun, J. Leibe, B., Matas, J., Sebe, N. & Welling,
1064 M. (eds) *Identity Mappings in Deep Residual Networks*. (eds Leibe, B., Matas,
1065 J., Sebe, N. & Welling, M.) *Computer Vision – ECCV 2016*, 630–645 (Springer
1066 International Publishing, Cham, 2016).
1067

1068 [45] van der Maaten, L. & Hinton, G. Visualizing data using t-SNE. *Journal of Machine*
1069 *Learning Research* **9**, 2579–2605 (2008).
1070

1071 [46] Bahadir, C. D., Dalca, A. V. & Sabuncu, M. R. *Learning-based Optimization of*
1072 *the Under-sampling Pattern in MRI*, 780 – 792 (2019).
1073

1074 [47] Knoll, F. *et al.* Accelerated Knee Imaging Using a Deep Learning Based Recon-
1075 struction. *Proceedings of the International Society of Magnetic Resonance in*
1076 *Medicine (ISMRM)* (2017).
1077

1078 [48] Zhang, J. *et al.* Extending LOUPE for K-space Under-sampling Pattern Opti-
1079 mization in Multi-coil MRI. *International Workshop on Machine Learning for*
1080 *Medical Image Reconstruction* 91 – 101 (2020).
1081

1082 [49] Weiss, T. *et al.* Joint Learning of Cartesian under Sampling Andre Construction
1083 for Accelerated MRI, 8653–8657 (2020).
1084

1085 [50] Knoll, F. *et al.* Advancing machine learning for MR image reconstruction with an
1086 open competition: Overview of the 2019 fastMRI challenge. *Magnetic resonance*
1087 *in medicine* **84**, 3054–3070 (2020).
1088

1089 [51] Hammernik, K. *et al.* Learning a Variational Network for Reconstruction of
1090 Accelerated MRI Data. *Magn Reson Med* **79**, 3055–3071 (2018).
1091

1092 [52] Zhu, B., Liu, J. Z., Cauley, S. F., Rosen, B. R. & Rosen, M. S. Image
1093 reconstruction by domain-transform manifold learning. *Nature* **555**, 487–492
1094 (2018).
1095

1096 [53] Lee, D., Yoo, J. & Ye, J. C. *Deep residual learning for compressed sensing MRI*,
1097 15–18 (IEEE, 2017).
1098

1099 [54] Hyun, C. M., Kim, H. P., Lee, S. M., Lee, S. & Seo, J. K. Deep learning for
1100 undersampled MRI reconstruction. *Physics in Medicine & Biology* **63**, 135007
1101 (2018).
1102

1103 [55] Cole, E. K., Pauly, J. M., Vasanawala, S. S. & Ong, F. Unsupervised MRI recon-
1104 struction with generative adversarial networks. *arXiv preprint arXiv:2008.13065*
(2020).

[56]	Knoll, F. <i>et al.</i> Assessment of the generalization of learned image reconstruction and the potential for transfer learning. <i>Magn Reson Med</i> 81 , 116–128 (2019).	1105 1106 1107
[57]	Mathieu, M., Henaff, M. & LeCun, Y. Fast Training of Convolutional Networks through FFTs. <i>CoRR</i> abs/1312.5 (2014).	1108 1109 1110
[58]	Rippel, O., Snoek, J. & Adams, R. P. <i>Spectral Representations for Convolutional Neural Networks</i> , NIPS’15, 2449–2457 (MIT Press, Cambridge, MA, USA, 2015).	1111 1112 1113
[59]	Cover, T. M. <i>Elements of information theory</i> (John Wiley & Sons, 1999).	1114 1115
[60]	Covert, I., Lundberg, S. & Lee, S.-I. Explaining by removing: A unified framework for model explanation. <i>arXiv preprint arXiv:2011.14878</i> (2020).	1116 1117 1118
[61]	Jethani, N., Sudarshan, M., Aphinyanaphongs, Y. & Ranganath, R. <i>Have We Learned to Explain?: How Interpretability Methods Can Learn to Encode Predictions in their Interpretations.</i> , 1459–1467 (PMLR, 2021).	1119 1120 1121
[62]	Sudarshan, M., Tansey, W. & Ranganath, R. Deep direct likelihood knockoffs. <i>Advances in neural information processing systems</i> 33 , 5036–5046 (2020).	1122 1123 1124
[63]	Jethani, N., Sudarshan, M., Covert, I., Lee, S.-i. & Ranganath, R. FastSHAP: Real-Time Shapley Value Estimation. <i>ICLR 2022</i> URL https://par.nsf.gov/biblio/10350825 .	1125 1126 1127 1128 1129 1130 1131 1132 1133 1134 1135 1136 1137 1138 1139 1140 1141 1142 1143 1144 1145 1146 1147 1148 1149 1150




Fig. 1: Part A: Overview of the MRI (top) and CT (bottom) pipeline showing the three components required for an image-based diagnosis. **Part B:** Inference pipeline of the Learning Data to Learn Representations (LDLR) framework, which identifies a subset of k -space data (red lines overlying the k -space matrix) with the highest disease-specific information (“disease signature”), learns its low-dimensional representation, and uses it as input to a classifier to detect the disease. **Part C:** The three training steps of LDLR. *Step 1:* Train a score function to estimate mutual information between the output variable y (presence/absence of the disease) and any subset of the k -space data. *Step 2:* Compute the score associated with a set of pre-selected candidate subsets and select the subset with the highest score (“disease signature”). *Step 3:* Train a DL model to represent disease signature in low-dimensional space and use the learned representation to detect the disease using a classifier. **Part D:** The architecture of the kSPACEENCODER model that takes subset of the k -space data as input and generates a low-dimensional representation encoding information for to the downstream task.





Fig. 2: (a) The ROC curves of a classifier trained using the low-dimensional representations generated by the KSPACEENCODER of the raw k -space data (blue curve) and the state-of-the-art image-based model trained on images reconstructed from fully-sampled k -space data (red curve). Both models were trained to infer the presence or absence of multiple pathologies across multiple anatomies. (b) Qualitative comparison of the disease information present in the low-dimensional representations learned by the KSPACEENCODER. The figure shows two-dimensional projections of the representations obtained using the t-SNE algorithm. Each point in the plot corresponds to one example, and its color represents the class (blue implies that the example does not have a disease, and red implies that the example has a disease). *Top row:* t-SNE projections of the image data. *Bottom row:* t-SNE projections of the representations learned by KSPACEENCODER. We clearly see that both sets of representations carry similar discriminatory signals (inferred from the cluster patterns).







Fig. 3: (a) Test set ROC curves of multiple LDLR models trained to detect disease at different sampling rates for the five diseases. LDLR’s performance is compared against $\text{BASELINE}_{\text{OPTIMG}}$, an image-based DL classifier trained using images reconstructed from the fully-sampled k -space data. (b) The 2-D t-SNE projections of the LDLR learned representations for different sampling rates across the five diseases. Each point in the plot is one example, and its color represents the class (blue: disease is not present, and red: disease is present). The clusters remain robust across different sampling rates, indicating that the disease-specific discriminative information remains intact.



(a)



(b)

Fig. 4: (a) Images reconstructed using a state-of-the-art reconstruction method from the subset of the k -space identified by LDLR at different sampling rates. Anatomical features that are clearly visible in the image reconstructed from fully-sampled k -space data become progressively more difficult to discern as we decrease the sampling rates $\alpha \in \{12.5\%, 10\%, 8\%, 5\%\}$. **(b)** Comparison of the performance of LDLR and BASELINE_{LOUPE:VARNET} as a function of sampling factor for each detection task. The plots validate the superiority of LDLR in detecting diseases over methods optimized for image reconstruction across all sampling factors.

1335
 1336
 1337
 1338
 1339
 1340
 1341
 1342
 1343
 1344
 1345
 1346
 1347
 1348
 1349
 1350
 1351
 1352
 1353
 1354
 1355
 1356
 1357
 1358
 1359
 1360
 1361
 1362
 1363
 1364
 1365
 1366
 1367
 1368
 1369
 1370
 1371
 1372
 1373
 1374
 1375
 1376
 1377
 1378
 1379
 1380




Fig. 5: Comparison of the k -space sampling patterns identified by LDLR (which optimizes for disease detection performance) and LOUPE (which optimizes for the image reconstruction performance). LDLR-identified patterns (top row) consist of a mix of low- and high-frequency data, whereas LOUPE-identified patterns (bottom row) consist of predominantly low-frequency data. The prostate and brain patterns have an 8% sampling rate, while the knee pattern has a 5% sampling rate.