



Highdicom: a Python Library for Standardized Encoding of Image Annotations and Machine Learning Model Outputs in Pathology and Radiology

Christopher P. Bridge^{1,2} · Chris Gorman³ · Steven Pieper⁴ · Sean W. Doyle² · Jochen K. Lennerz^{5,6} · Jayashree Kalpathy-Cramer^{1,2,7} · David A. Clunie⁸ · Andriy Y. Fedorov^{7,9} · Markus D. Herrmann^{3,6}

Received: 28 October 2021 / Revised: 20 May 2022 / Accepted: 26 May 2022
© The Author(s) 2022

Abstract

Machine learning (ML) is revolutionizing image-based diagnostics in pathology and radiology. ML models have shown promising results in research settings, but the lack of interoperability between ML systems and enterprise medical imaging systems has been a major barrier for clinical integration and evaluation. The DICOM[®] standard specifies information object definitions (IODs) and services for the representation and communication of digital images and related information, including image-derived annotations and analysis results. However, the complexity of the standard represents an obstacle for its adoption in the ML community and creates a need for software libraries and tools that simplify working with datasets in DICOM format. Here we present the *highdicom* library, which provides a high-level application programming interface (API) for the Python programming language that abstracts low-level details of the standard and enables encoding and decoding of image-derived information in DICOM format in a few lines of Python code. The *highdicom* library leverages NumPy arrays for efficient data representation and ties into the extensive Python ecosystem for image processing and machine learning. Simultaneously, by simplifying creation and parsing of DICOM-compliant files, *highdicom* achieves interoperability with the medical imaging systems that hold the data used to train and run ML models, and ultimately communicate and store model outputs for clinical use. We demonstrate through experiments with slide microscopy and computed tomography imaging, that, by bridging these two ecosystems, *highdicom* enables developers and researchers to train and evaluate state-of-the-art ML models in pathology and radiology while remaining compliant with the DICOM standard and interoperable with clinical systems at all stages. To promote standardization of ML research and streamline the ML model development and deployment process, we made the library available free and open-source at <https://github.com/herrmannlab/highdicom>.

Keywords DICOM · Python · Software · Machine learning · Segmentations · Structured reports

Background

Recent breakthroughs in machine learning (ML) and computational processing capabilities have led to the development of ML models that demonstrate unprecedented

performance on a variety of highly complex computer vision tasks [1]. State-of-the-art convolutional neural network models now regularly achieve near-human or even superhuman performance on a variety of challenging vision tasks and across different imaging modalities, including

✉ Markus D. Herrmann
mdherrmann@mgh.harvard.edu

¹ Martinos Center for Biomedical Imaging, Massachusetts General Hospital, Boston, MA, USA

² MGH & BWH Center for Clinical Data Science, Mass General Brigham, Boston, MA, USA

³ Computational Pathology, Department of Pathology, Massachusetts General Hospital, Boston, MA, USA

⁴ Isomics, Inc., Cambridge, MA, USA

⁵ Center for Integrated Diagnostics, Department of Pathology, Massachusetts General Hospital, Boston, MA, USA

⁶ Department of Pathology, Harvard Medical School, Boston, MA, USA

⁷ Department of Radiology, Harvard Medical School, Boston, MA, USA

⁸ PixelMed Publishing, LLC., Bangor, PA, USA

⁹ Surgical Planning Laboratory, Department of Radiology, Brigham and Women's Hospital, Boston, MA, USA

segmentation and classification of slide microscopy images in pathology [2–4] as well as computed tomography or magnetic resonance images in radiology [5–7]. Over the last couple of years, ML models have evolved technically within the research domain [8] and the vision is that these models will soon be applied widely in clinical practice to support pathologists and radiologists in interpretation of images and ultimately improve diagnostic accuracy and efficiency [6, 9–11]. To realize this vision, healthcare enterprises are now tasked with evaluating model performance in clinical context and integrating the outputs of ML models into clinical workflows. Similarly, development of models may be expedited if image annotations generated by clinical experts and stored within clinical information systems could be directly consumed by ML training and validation pipelines. Unfortunately, this is currently impeded by the lack of standard interfaces for exchange of image annotations and ML model outputs between image analysis, image display, and image management systems.

Digital Imaging and Communication in Medicine (DICOM) is the internationally accepted standard for communication of medical images and related information across a wide range of medical imaging modalities and disciplines. Hospitals around the world have established an extensive enterprise imaging infrastructure, workflows, and software applications based on DICOM [12] and pathology and radiology are converging towards using DICOM for communication of digital images [13–16]. However, existing pathology as well as radiology systems primarily rely on non-standard formats and interfaces for the storage and exchange of image annotations and computational image analysis results, to which we hereafter collectively refer as *annotations*. Similarly, ML models developed by researchers generally receive and return annotations in a variety of customized formats that are incompatible with clinically available image management and display systems and that lack metadata required for interpretation and use of the information in clinical context. Instead, it would be desirable if ML models were developed according to the FAIR guiding principles [17] using standardized metadata to allow for annotations to be findable, accessible, interoperable, and reusable. The DICOM standard provides information object definitions (IODs), such as Segmentations and Structured Reports, for annotations [18, 19], and implementation of these IODs to enable interoperable storage and communication of ML model outputs has been proposed by the Integrating the Healthcare Enterprise (IHE) Radiology Technical Committee [20].

Python is the de facto standard programming language of data science and provides a rich ecosystem for scientific computing, image processing, and machine learning [21–24]. The majority of ML models are developed and deployed in the form of Python programs. The *pydicom* library [25] provides data structures and routines for storing and accessing data of DICOM datasets (parts 5 and 6 of the DICOM standard) as well as reading and writing DICOM files (part 10 of the DICOM standard). However, *pydicom* has no concept of IODs (parts 3 and 16 of the DICOM standard) and as such leaves it to each developer to set all attributes required by an IOD manually and ensure that they follow all relevant constraints when creating new DICOM objects containing annotations. Similarly, parsing the annotation IODs for the information relevant to a particular ML task using the *pydicom* API is challenging due to their highly nested and interdependent structure. Consequently, both tasks are slow, complex, and error-prone and require considerable knowledge of the DICOM standard. We therefore identified a need for a higher-level abstraction layer between the ML model developer and the low-level encoding rules of the DICOM standard. This motivated us to create the open-source *highdicom* library, which provides a high-level application programming interface for creating and reading annotations in DICOM format using the Python programming language. Our goal in releasing this library is to enable ML processes that achieve interoperability between ML models and clinical information systems throughout the entire model development and deployment lifecycle while avoiding the complexity that this currently entails. Furthermore, we aimed to create a library that is applicable across a range of common ML tasks and imaging domains.

In this article, we first describe the design and implementation of the *highdicom* library to meet this unmet need and then assess the library's capabilities in encoding and decoding annotations (either generated by human readers or ML models) in DICOM format. We perform experiments that demonstrate the use of the library during ML model training and inference and show how the library enables the development of ML models that are interoperable with established image management and display systems and thus can be readily integrated into an enterprise medical imaging environment. To this end, we consider a variety of clinically relevant computer vision problems and multiple imaging modalities across different medical disciplines, placing a focus on lung tumor detection in slide microscopy images in pathology and computed tomography images in radiology as an illustrative use case.

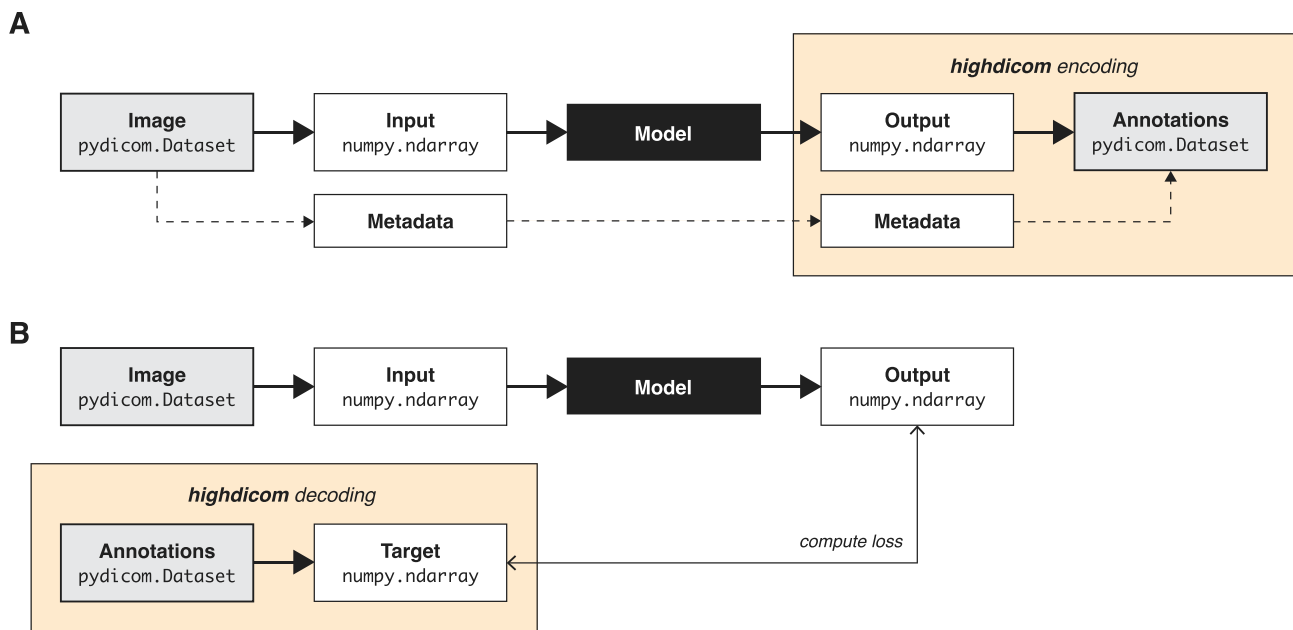


Fig. 1 Intended use of *highdicom* in data pipelines during machine learning model training and inference workflows. **A** Encoding of model outputs upon inference in the postprocessing pipeline. **B** Decoding of image annotations for model training in the preprocessing pipeline

Methods

Design Overview and Application Programming Interface (API)

The software components responsible for transforming the data input and output from ML models, and thereby ensuring interoperability with adjacent systems, are commonly referred to as data pipelines [26, 27]. During inference, pipelines are responsible for retrieving and preprocessing input images into an in-memory format that can be consumed by the model and encoding the model's in-memory outputs into a form suitable for communication and storage. During training, they retrieve and preprocess input images and additionally, if required, decode annotations into an in-memory representation of the target for model training. The *highdicom* library is intended to operate within data pipelines that connect clinical infrastructure using the DICOM standard to popular Python ML frameworks such as *PyTorch* [28] and *Tensorflow* [29], and is focused on annotations rather than the input images themselves. The library's core functionality is twofold: First, encoding model outputs in the form of *NumPy* arrays together with relevant metadata into annotations in the form of *pydicom* objects (Fig. 1A). Second, decoding annotations provided as *pydicom* datasets to obtain targets in the form of *NumPy* arrays (Fig. 1B) by reading and interpreting the included metadata. We chose the n-dimensional *NumPy* array data structure [22] as an in-memory representation of

model outputs and targets because it is interoperable with *pydicom* as well as *PyTorch* and *Tensorflow* and many other well-established Python image processing libraries (e.g., *OpenCV* [30] and *ITK* [31]).

API Overview We designed *highdicom* following the object-oriented programming paradigm and modelled the API according to the DICOM Information Model, which specifies different abstract data types that are referred to as information object definitions (IODs) (Fig. 2). An IOD defines the set of required and optional DICOM attributes that may be included into DICOM objects. We selected various IODs for storage of annotations and implemented each in *highdicom* as a Python class.

Strictly speaking, each Python class implements a DICOM Storage Service-Object Pair (SOP) Class, which is the data structure within the DICOM standard that stores the attributes defined by an IOD. An instance of such a Python class thus represents a DICOM SOP instance and serves as a container for a DICOM dataset, where each instance attribute holds the value of a DICOM data element.

The Python classes are ultimately derived from the *pydicom.Dataset* class from the existing *pydicom* package and therefore inherit low-level behaviors, such as accessing, setting, iterating over data elements, and reading/writing to/from files that many developers are already familiar with. It further allows developers to retain

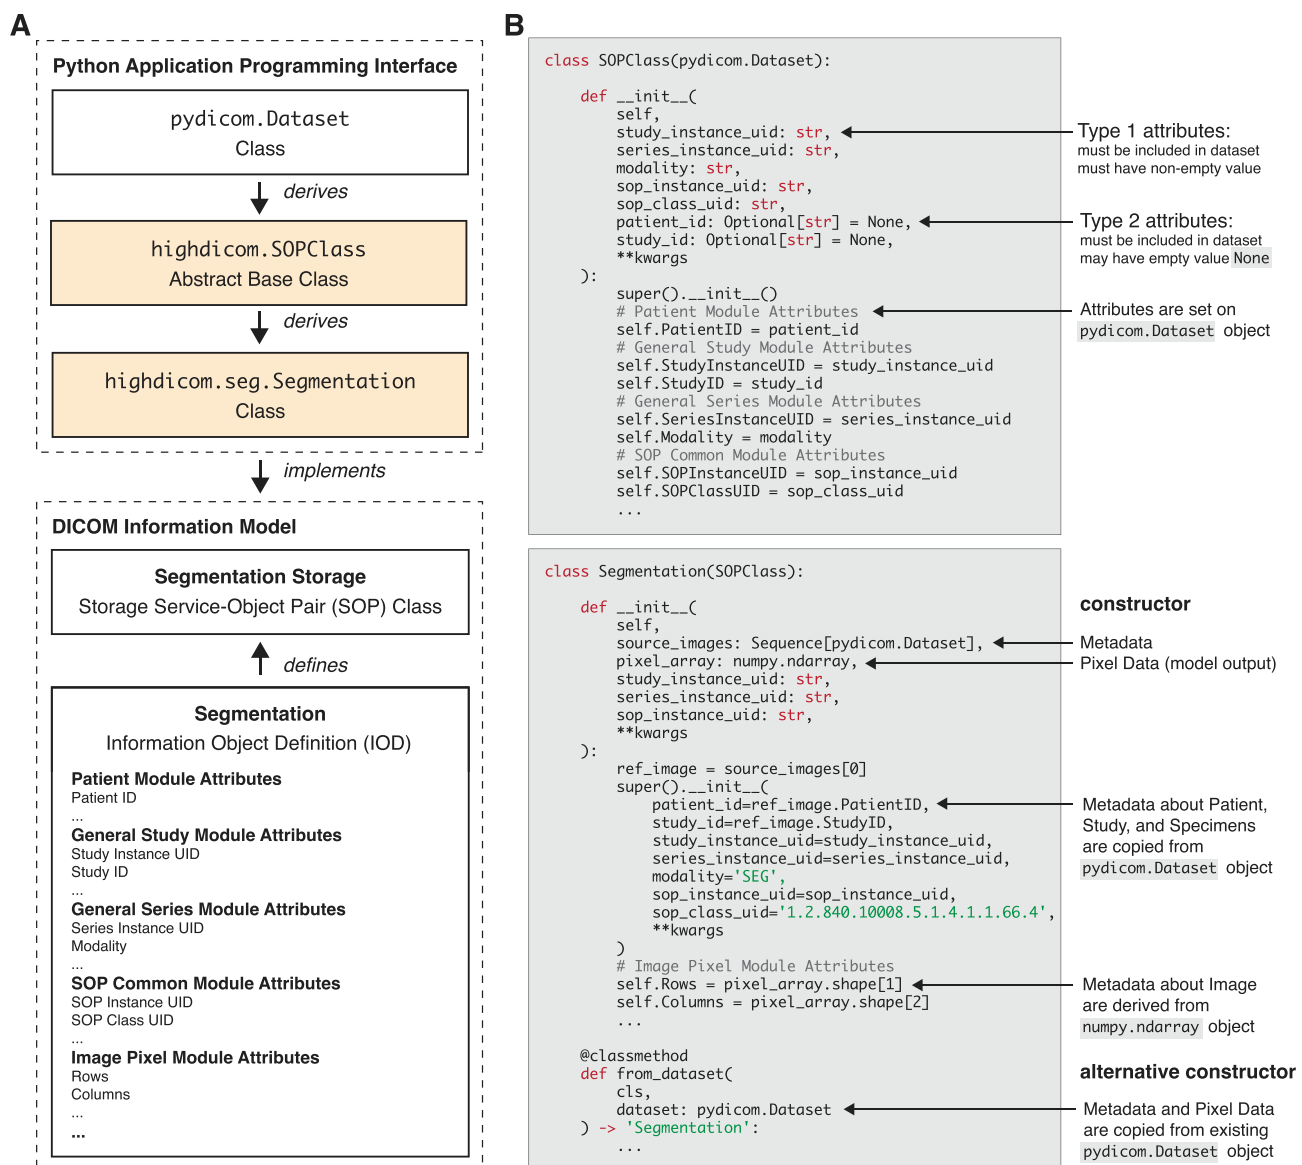


Fig. 2 Implementation of the DICOM Information Model in Python. **A** The *highdicom* Python abstract base class `highdicom.SOPClass` and its relationship to an DICOM information object definition (IOD) and DICOM Storage Service-Object Pair (SOP) Class. **B**

A *highdicom* Python class for a specific DICOM IOD and SOP Class (exemplified by `highdicom.seg.Segmentation` that implements the DICOM Segmentation Storage SOP Class defined by the DICOM Segmentation IOD)

low-level control over all data elements in order to add to or alter information in objects constructed by *highdicom*. Below `pydicom.Dataset` in the class hierarchy, there is a common abstract base class called `highdicom.SOPClass` (Fig. 2A), which abstracts the attributes that are required by all SOP classes. Specific SOP classes are then implemented by dedicated Python classes that are derived from the abstract base class (Fig. 2B). In this way, we aim to provide an idiomatic Python interface that abstracts as much of the low-level DICOM encoding and decoding rules as possible while staying close to the standard DICOM terminology to avoid potential ambiguities.

Encoding of DICOM SOP Instances The process of encoding information in derived objects is implemented in the constructor methods of the corresponding SOP classes (either in the `highdicom.SOPClass` abstract base class or in derived IOD-specific classes). For construction of an SOP instance, the developer provides the image-derived information that is outputted by a model (e.g., pixel data or graphic data) together with descriptive contextual information that the standard requires for the corresponding IOD. Attribute values that are static or can be derived from provided arguments are automatically set upon object construction. For example, relevant metadata about the patient, the study, or

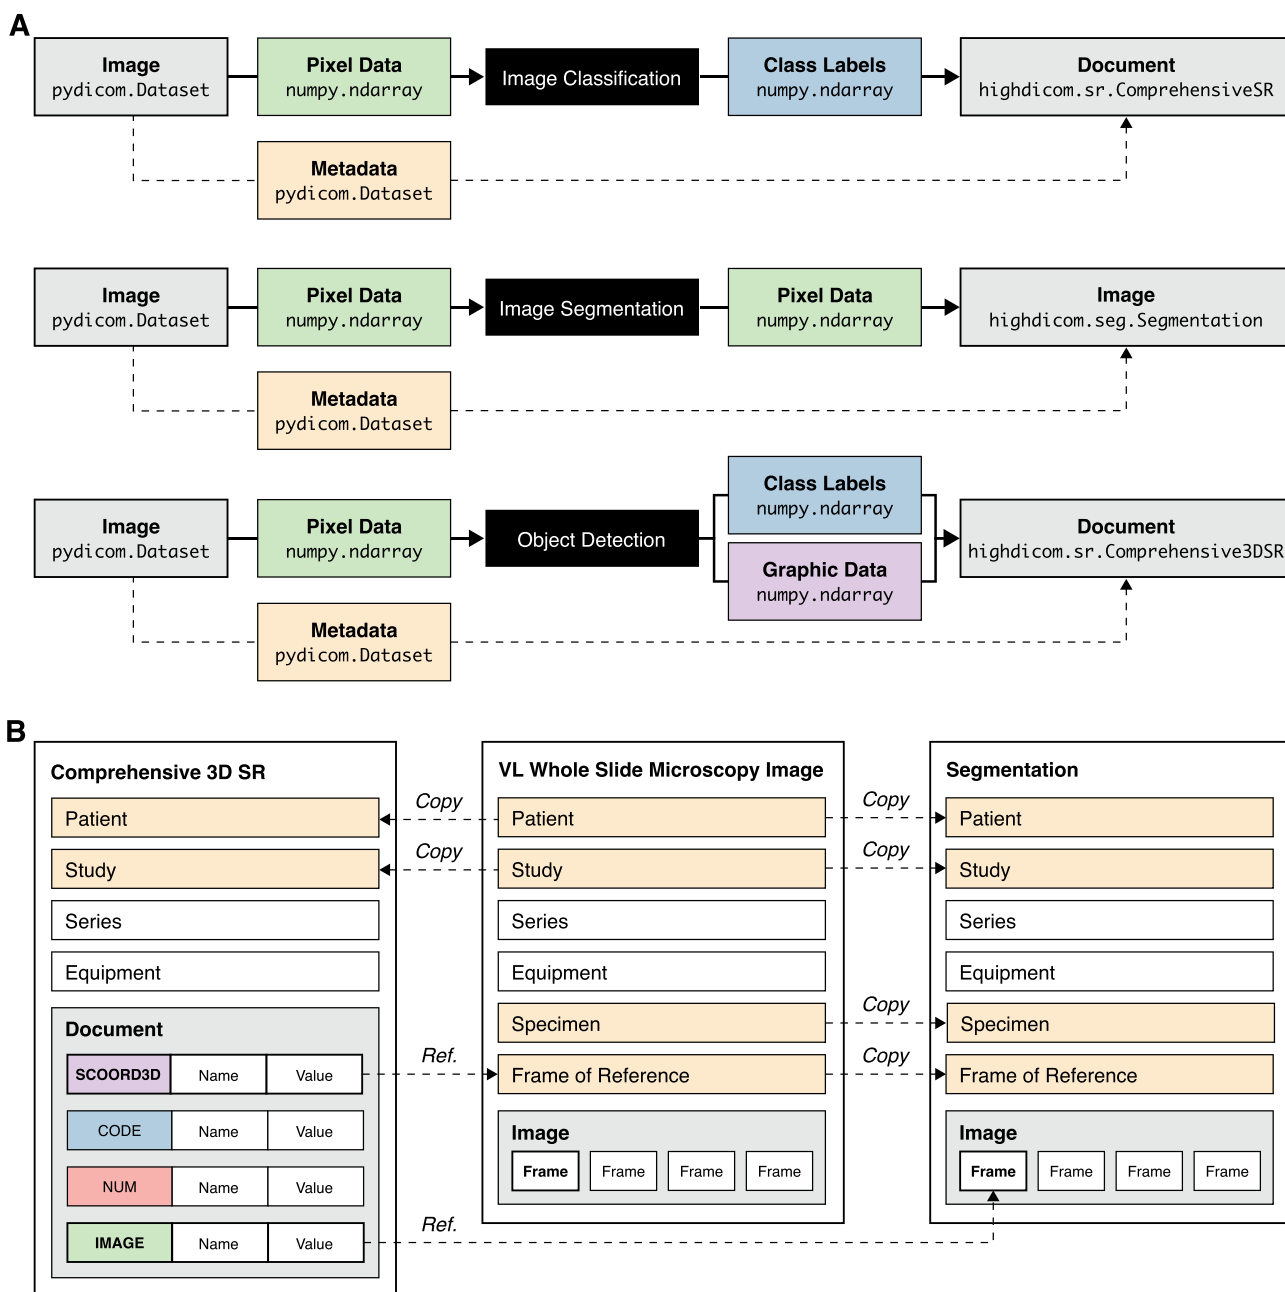


Fig. 3 Encoding of machine learning model outputs in DICOM. **A** Information entities and the Python types used to represent machine learning model inputs (images) and outputs (image-derived information) for three common decision problems. **B** Schematic overview of the content of source image objects (exemplified by a DICOM VL

Whole Slide Microscopy Image) and derived objects (DICOM Comprehensive 3D SR and DICOM Segmentation). Note that descriptive metadata is copied from source to derived objects and derived objects may reference information contained in source images or other derived objects

the specimen are automatically copied from the metadata of provided source images and references to the source images are included in the derived objects (Fig. 3A–B). Furthermore, the constructor automatically validates the content of created SOP instances through runtime checks to ensure that constructed objects are fully compliant with the relevant IOD in the standard.

By design, all required information must be passed to the SOP class constructor when creating the object, and thereafter the object remains immutable through the *highdicom* API (though an experienced developer may use the lower-level interface provided by the *pydicom* API to modify the object if required). This means that the constructor can validate all input parameters at once accounting for all

interdependencies and conditional logic between attributes. It also reflects the intent of the standard in that DICOM objects are immutable following creation.

Decoding of DICOM SOP Instances The *pydicom* library provides a powerful low-level Python interface to developers to access DICOM data elements of a dataset directly, with little abstraction from the details of the data format. While this is appropriate for many image objects, the complexity of the derived objects used for annotations means that accessing the desired information using the *pydicom* API requires a detailed knowledge of the underlying data structures and in our experience results in a verbose, cumbersome, and error-prone process. Therefore, we have endowed *highdicom* SOP classes with additional methods (not in the standard) that provide a means for developers to access, filter, and interpret the content of a DICOM object when preparing image annotations to be used as targets for a training algorithm. In addition, *highdicom* SOP classes implement alternative constructor methods that allow for the creation of *highdicom* SOP instances from existing *pydicom.Dataset* objects, which were read from a file or retrieved over network, and thereby enhance the objects with additional, modality-specific methods and properties for data access.

Data Types and Structures The majority of DICOM metadata attribute values that are passed to and returned from the *highdicom* API upon encoding and decoding of SOP instances have primitive, built-in Python types such as strings (`str`), integers (`int`), and floats (`float`). To further encapsulate closely related metadata of composite DICOM data types (DICOM Sequences or Sequence Items) and to improve code readability and reusability, the *highdicom* API further provides custom Python types, which are implemented in the form of Python classes and are generally derived from either *pydicom.Dataset* or *pydicom.Sequence*. DICOM bulkdata values such as pixel data or vector graphic data are passed to and returned from *highdicom* Python classes as NumPy objects (`numpy.ndarray`).

Storage of Annotations in DICOM Format

Having described the general approach taken by our library, we now begin to discuss the individual IODs that we selected for implementation. The DICOM standard specifies a wide range of IODs for different types of DICOM objects, including images acquired by various modalities (e.g., computed tomography or whole slide microscopy) as well as image-derived information generated by image display, processing, or analysis systems [32]. For implementation in the *highdicom* library, we considered standard IODs that provide mechanisms to store image annotations for common ML tasks across pathology and radiology use cases.

We thereby focused on the following decision problems and their corresponding annotations (Fig. 3A) [1]:

1. *Image classification* — class labels in the form of discrete binary or categorical values and optionally class scores in the form of continuous probabilistic values (Fig. 3A upper panel)
2. *Image segmentation* — class labels at pixel resolution that identify semantically distinct regions of interest (ROIs) within an image in the form of raster graphics (Fig. 3A middle panel)
3. *Object detection* — spatial coordinates for individual ROIs in the form of vector graphics (commonly bounding boxes), combined with class labels and detection scores (Fig. 3A lower panel)

We identified three IODs that together allow for the encoding of annotations for these common use cases: the Segmentation IOD and two Structured Report (SR) IODs. The Segmentation IOD was selected to encode ROIs returned by image segmentation models as raster graphics. The Comprehensive SR and Comprehensive 3D SR IODs were chosen to encode vector graphic ROIs returned by object detection models as well as class labels, scores, and measurements returned by image classification and regression models (Fig. 3A). All three IODs are designed to be agnostic of the imaging modality and able to support use cases across medical disciplines including pathology and radiology.

DICOM Segmentation Images The Segmentation IOD is implemented in *highdicom* as the `highdicom.seg.Segmentation` Python class and allows for the encoding of one or more components, which in DICOM are referred to as *segments*. Each segment may represent a pixel class (category) or an individual instance of a given class as generated by semantic segmentation or instance segmentation models [33], respectively. Segments may further have binary or fractional type, either representing a mask of Boolean values where non-zero pixels encode class membership or a mask of decimal numbers where pixels encode class probability.

In order to encode a DICOM Segmentation image, the developer passes to the constructor a mask as a `numpy.ndarray` (of either Boolean, integer, or floating point data type) along with additional metadata that describe the meaning of each segment within the segmentation (`highdicom.seg.SegmentDescription`) and the algorithm responsible for producing the segmentation (`highdicom.AlgorithmIdentificationSequence`).

To facilitate decoding of DICOM Segmentation images, the `highdicom.seg.Segmentation` class provides methods that allow developers to filter segments by their label, segmented property category or type, or tracking identifiers. It further provides methods to obtain a segmentation

mask as a `numpy.ndarray` for a given set of segments and source image frames. While conceptually straightforward, in practice, several steps are necessary to achieve this correctly: (i) Determining which frames stored in the Segmentation image are relevant to a given set of segments and source image frames based on the multi-frame dimension indexing information, (ii) Sorting the Segmentation image frames according to the query, (iii) Adding in missing pixel values in case of sparse Segmentation images where background image frames were omitted during encoding to save storage space (iv) (Optionally) combining multi binary segments into a multi-class label map.

DICOM Structured Report Documents There are various IODs defined by the standard that utilize structured reporting, but we selected the Comprehensive SR (`highdicom.sr.ComprehensiveSR`) and Comprehensive 3D SR (`highdicom.sr.Comprehensive3DSR`) IODs for implementation in *highdicom* because they provide the most flexible mechanisms for storing annotations. In addition to the IOD definitions, the standard provides SR templates, which serve as schemas that define how the content of an SR document shall be structured and how the information shall be encoded. A template consists of a sequence of *content items*, each defining a name-value pair (or question-answer pair) that encodes a domain-specific property or concept (Fig. 4A). Notably, both concept names and values have a composite data type and are each encoded by one or more DICOM attributes. Concept names are coded using standard medical terminologies and ontologies such as the DICOM Controlled Terminology or the Systematized Nomenclature of Medicine Clinical Terms (SNOMED CT) and thereby get endowed with an explicit, domain-specific meaning [34]. The structure of the corresponding value depends on the *value type*, which defines a set of DICOM attributes that are included in the SR document to represent the assigned value.

Within *highdicom*, the `highdicom.sr.CodedConcept` class is an important data type that encapsulates the DICOM attributes required to code a concept using a standard coding scheme within a single Python object. We further contributed lower-level data types to the underlying *pydicom* library that provide programmatic access to codes included in the DICOM standard, specifically the DICOM Controlled Terminology (DCM), SNOMED-CT (SCT), and Unified Code for Units of Measure (UCUM) coding schemes. These codes that are included in the *pydicom* library are fully compatible with the coded concepts of the *highdicom* library and can generally be used interchangeably throughout the API. Furthermore, for each of the different DICOM content item value types, we have implemented a separate Python class that is derived from `pydicom.Dataset` and encapsulates both the coded concept name and the corresponding value of the given type (Fig. 4B).

Notable content item classes include `highdicom.sr.CodeContentItem`, which may be used to store class labels as coded values, and `highdicom.sr.NumContentItem`, which may be used to store a measurement along with its unit. ROIs may be either encoded by value or by reference and stored within or outside of the SR document content, respectively. In the case of vector graphics (including but not limited to bounding boxes), the graphic data may be stored within the SR document and encoded via DICOM content items of value type `SCORD3D`, which encodes 3D spatial coordinates of geometric objects in the frame of reference (patient or slide coordinate system). This value type is implemented in *highdicom* by the `highdicom.sr.Scoord3DContentItem` Python class (Fig. 4B). In the case of raster graphics, the pixel data of Segmentation images are stored outside of the SR document, but specific segments can be referenced from within the SR document via content items of value type `IMAGE`. (implemented by the `highdicom.sr.ImageContentItem` Python class), which includes DICOM identifiers for the referenced image object and segments contained therein.

The standard provides different SR templates for a variety of common clinical use cases and diagnostics tasks, such as recording X-ray dose exposure or reporting echocardiography findings. We chose to implement the more generic template TID 1500 “Measurement Report” in *highdicom* for encoding annotations, because the template provides standard content items to describe measurements and qualitative evaluations of images as well as individual image ROIs (Fig. 4C) and because it has already been successfully used for standardized communication of quantitative image analysis results [18, 19]. Importantly, sub-templates that can be included in TID 1500 allow for the encoding of annotations of entire images, planar image regions, or volumetric image regions (Fig. 4C). Within the library’s API, these selected templates are implemented by Python classes, which are derived from an abstract base class `highdicom.sr.Template`, which is in turn derived from `pydicom.Sequence` (Fig. 4D). The constructors of these Python classes require the developer to pass the relevant data via named parameters but then handle its inclusion in the template with the correct concept names as well as ensuring all constraints are satisfied.

When decoding SR documents, the high degree of nesting in the document tree and the variable order of content items at each level means that finding a particular content item of interest in the tree potentially requires multiple nested loops. Furthermore, as described above, each content item is a collection of data elements that must first be parsed and interpreted as a unit. The Python classes that implement SR templates and individual SR content items provide methods and properties to facilitate data access. Using the provided methods, measurement groups within a `highdicom.sr.ComprehensiveSR` or

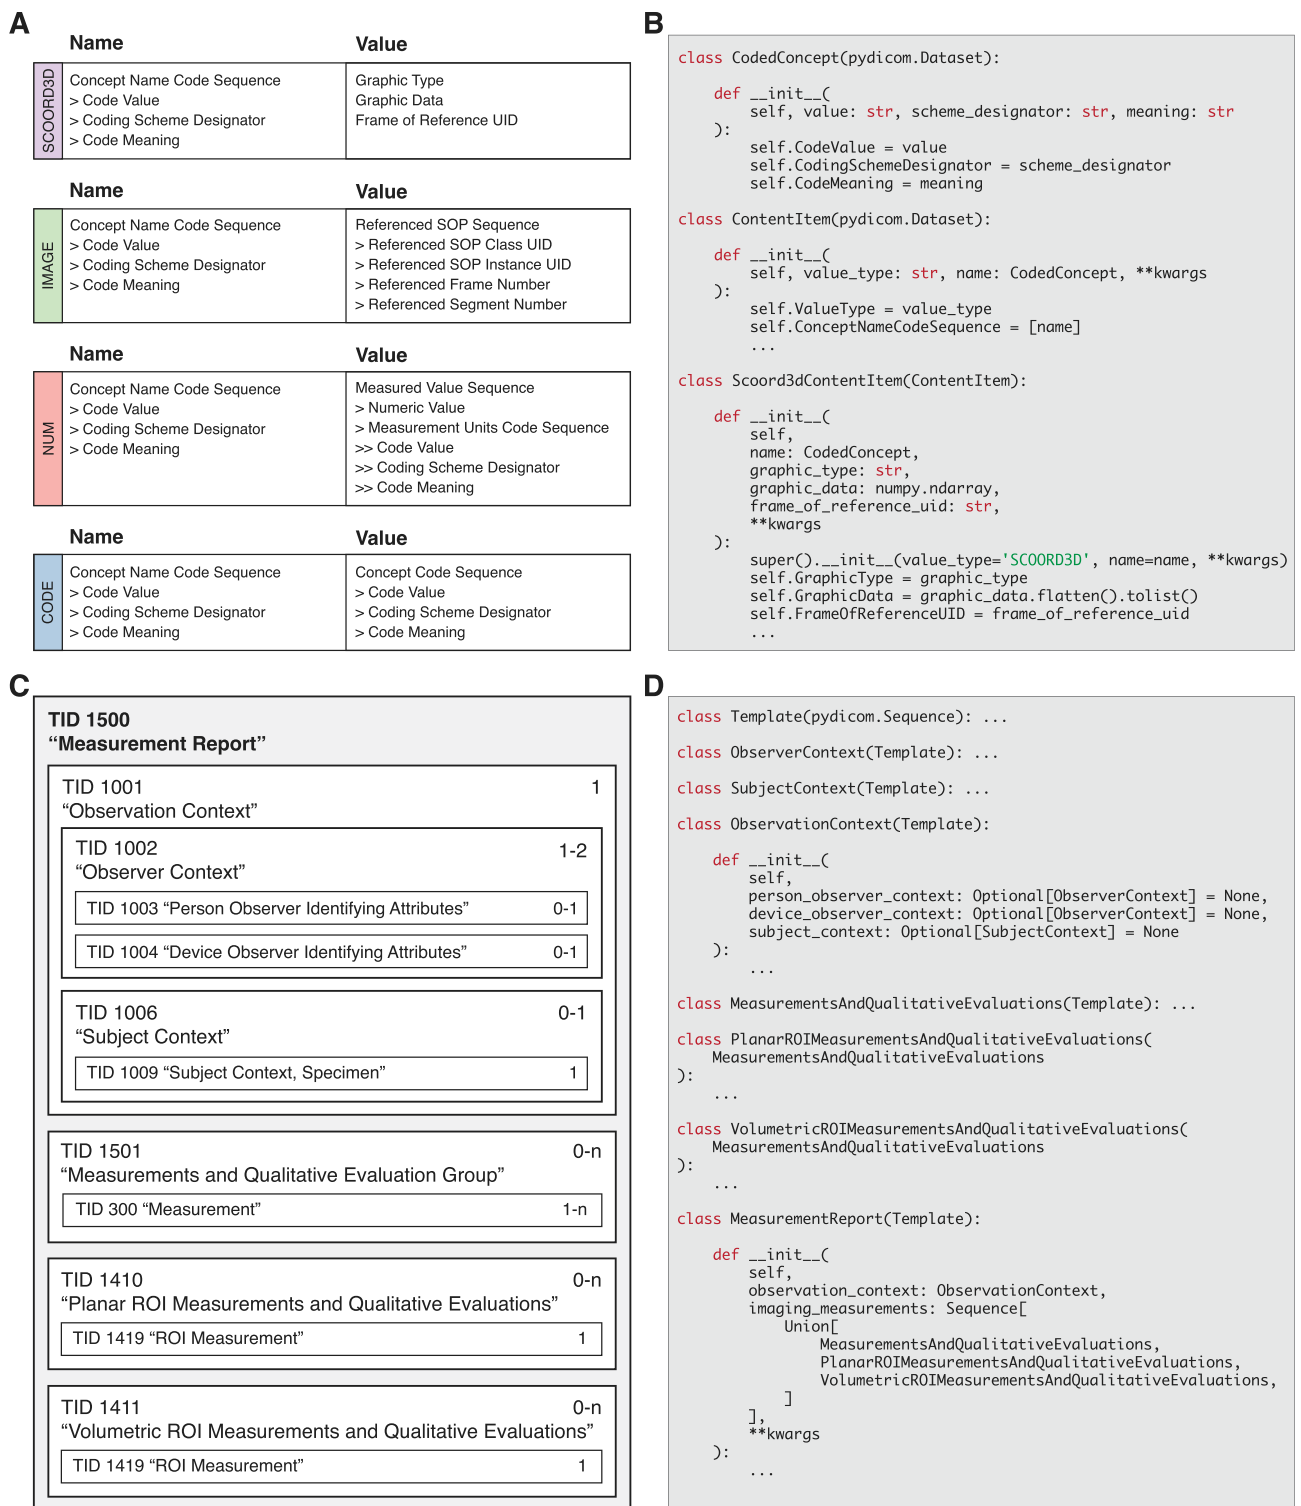


Fig. 4 Encoding of annotations as DICOM Structured Reporting (SR) content items and templates for inclusion into an SR document. **A** SR content items of different values types. **B** Implementation of SR content items in *highdicom* by classes that inherit from *pydicom*.

Dataset. **C** SR template TID 1500 "Measurement Report" and included sub-templates. **D** Implementation of SR templates in *highdicom* by classes that inherit from *pydicom*. Sequence

`highdicom.sr.Comprehensive3DSR` object can be filtered by their finding type, finding site, or tracking identifiers. Individual measurements and qualitative evaluations contained within these groups can similarly be filtered by their concept name. Furthermore, *highdicom* classes representing SR templates and content items provide access to their content items or values, respectively, through Python properties that return the data either as a built-in Python type or a custom *highdicom* type (which will typically match the type of the argument passed to the constructor).

Results

Having laid the foundation through the description of the library's design and implementation, we now proceed to demonstrating the capabilities of the library. We consider a concrete use case of developing machine learning models for lung tumor detection in both pathology and radiology and deploying the models clinically using a common platform and framework that is applicable independent of the medical discipline or imaging modality. In this section, we first describe the steps necessary to encode the annotations in DICOM using *highdicom*, including the description of the detected region of interest, the identified finding, and related measurements and qualitative evaluations. We then show through a series of experiments how *highdicom* can streamline ML model training and inference for this use case.

Highdicom Facilitates Encoding of Image Annotations in DICOM Format

Structured Reporting using Standard Medical Terminologies While the approach of using standardized vocabularies is powerful and important for interoperability, it complicates working with the data. For example, comparing two concepts for equality requires comparison of their code values, coding scheme designators, and coding scheme versions. The `highdicom.sr.CodedConcept` and the lower-level *pydicom* types facilitate the use of coded concepts for structured reporting of annotations in Python at a high level of abstraction (code snippet 1).

```
import highdicom as hd
from pydicom.sr.codedict import codes

tumor = hd.sr.CodedConcept("108369006", "SCT", "Tumor")
neoplasm = codes.SCT.Neoplasm

assert(tumor == neoplasm) # True
assert(tumor.value == neoplasm.value) # True
assert(tumor.scheme_designator == neoplasm.scheme_designator) # True
```

Code snippet 1: Coding of concepts using domain-specific terminologies. "Neoplasm" and "Tumor" are synonyms that map to the same code within the SNOMED-CT coding scheme.

Describing ROI Evaluations and Measurements The coded concept type forms the basis for additional higher-level composite data types for DICOM structured reporting such as SR content items. Code snippet 2 demonstrates example content items for the encoding of a tumor image region of interest, the tumor finding, and an associated tumor measurement.

```
import numpy as np
import highdicom as hd
from pydicom.sr.codedict import codes

# Encode bounding box enclosing the tumor region using 3D spatial coordinates
region_item = hd.sr.ImageRegion3D(
    graphic_type=hd.sr.GraphicTypeValues3D.POLYGON,
    graphic_data=np.array([
        [45., 34., 0.],
        [49., 34., 0.],
        [49., 42., 0.],
        [45., 42., 0.],
        [45., 34., 0.]
    ]),
    frame_of_reference_uid="1.2.3.4"
)

# Encode a qualitative evaluation describing the "morphology" as "adenocarcinoma"
morphology_item = hd.sr.QualitativeEvaluation(
    name=codes.SCT.AssociatedMorphology,
    value=hd.sr.CodedConcept(
        value="8551/3",
        scheme_designator="ICD03",
        meaning="Acinar adenocarcinoma"
    )
)

# Encode a measurement with name "area" and of unit "square millimeter"
area_item = hd.sr.Measurement(
    name=codes.SCT.Area,
    value=12.07,
    unit=codes.UCUM.SquareMillimeter
)
```

Code snippet 2: Encoding of regions of interest, qualitative evaluations and measurements.

This demonstrates using the SCT vocabulary built in to *pydicom* to encode a concept name as "Morphology," and a domain-specific coding scheme, the International Classification of Diseases for Oncology (ICD-O), to specify the exact type of tumor as the concept value. Of note, the area measurement in our example is encoded in a well-defined physical unit, as would be expected for clinical decision-making. The corresponding image region is defined in the same physical space. In DICOM, image regions may be defined by spatial coordinates within either the pixel matrix of an individual image or, as in this example, the frame of reference (the 3D patient- or slide-based physical coordinate system). While the former appears more straightforward, the latter is more general and allows for annotations derived from transformed versions of the original images with arbitrary affine transformations (rotations, scaling, etc.) as well as crops.

Creation of DICOM Annotation Objects The computer vision problem of tumor detection could be solved using either an object detection or image segmentation model. Accordingly, the output of these models and the annotations used to train them can be encoded using the `highdicom.sr.Comprehensive3DSR` (code snippet 3) and `highdicom.seg.Segmentation` (code snippet 4) classes respectively. In either case, this involves describing the finding and the anatomical site of the finding as well as supplying relevant contextual metadata such as the device or person reporting the observation. However, note that it is not necessary to specify patient, study, or specimen information since

highdicom copies this metadata directly from the source images provided as evidence to the constructor.

```
import highdicom as hd
from pydicom.dataset import Dataset
from pydicom.sr.codedict import codes

def load_image() -> Dataset: ...

source_image = load_image()

# Describe the device that is reporting the observation
observation_context = hd.sr.ObservationContext(
    observer_device_context=hd.sr.ObserverContext(
        observer_type=codes.DCN.Device,
        observer_identifying_attributes=hd.sr.DeviceObserverIdentifyingAttributes(
            uid=hd.UID(),
            name="GPU01"
        )
    )
)

# Describe the region of interest finding and associated measurements and evaluations
roi_description = hd.sr.PlanarROIMeasurementsAndQualitativeEvaluations(
    finding_type=codes.SCT.Neoplasm,
    finding_sites=[hd.sr.FindingSite(anatomic_location=codes.SCT.Lung)],
    referenced_region=region_item,
    measurements=[area_item],
    qualitative_evaluations=[morphology_item]
)

# Construct a measurement report from the ROI(s)
measurement_report = hd.sr.MeasurementReport(
    observation_context=observation_context,
    procedure_reported=codes.SCT.ImagingProcedure,
    imaging_measurements=[roi_description]
)

# Construct an SR document containing the measurement report
report_document = hd.sr.Comprehensive3DSR(
    evidence=[source_image],
    content=[measurement_report[0]],
    series_number=1,
    series_instance_uid=hd.UID(),
    sop_instance_uid=hd.UID(),
    instance_number=1,
    manufacturer="NGH Computational Pathology",
    manufacturer_model_name="Object Detection Model",
    software_version="v1",
    device_serial_number="XYZ"
)
```

Code snippet 3: Creation of a DICOM Comprehensive 3D SR document instance to encode predictions of an object detection model using content items constructed in snippet 2.

```
from typing import Tuple
import highdicom as hd
import numpy as np
from pydicom.sr.codedict import codes
from pydicom.dataset import Dataset

def load_image() -> Tuple[Dataset, np.ndarray]: ...

def segment_image(image: np.ndarray) -> np.ndarray: ...

metadata, frames = load_image()
probabilities = segment_image(frames)

segment_description = hd.sr.SegmentDescription(
    segment_numbers=1,
    segment_label="Tumor",
    segmented_property_category=codes.cid7150.MorphologicallyAbnormalStructure,
    segmented_property_type=codes.cid7159.Neoplasm,
    algorithm_type=hd.sr.SegmentAlgorithmTypeValues.AUTOMATIC,
    algorithm_identification=hd.AlgorithmIdentificationSequence(
        name="highdicom-experiments",
        version="v1.0",
        family=codes.cid7162.ArtificialIntelligence
    )
)

segmentation_image = hd.sr.Segmentation(
    source_images=[metadata],
    pixel_array=probabilities,
    segmentation_type=hd.sr.SegmentationTypeValues.FRACTIONAL,
    segment_descriptions=[segment_description],
    series_instance_uid=hd.UID(),
    series_number=1,
    sop_instance_uid=hd.UID(),
    instance_number=1,
    manufacturer="NGH Computational Pathology",
    manufacturer_model_name="Segmentation Model",
    software_versions="v1",
    device_serial_number="Server XYZ"
)
```

Code snippet 4: Creation of a DICOM Segmentation image instance to encode predictions of a semantic image segmentation model.

Highdicom Facilitates Efficient Loading and Decoding of Images and Corresponding Annotations

When it comes to training a model for tumor detection, annotations may be provided in the form of either raster graphics within a Segmentation image or vector graphics within an SR document. In both cases, *highdicom* provides methods that simplify access to, and interpretation of, the relevant content in the annotation SOP instances. If annotations are provided as raster graphics within a Segmentation image, model training may require combining binary bit planes from multiple segments in the Segmentation image to create a single label map, represented as a NumPy array, in which pixels encode tumor identities. If instead annotations are provided as vector graphics within an SR document, the spatial coordinates of image regions will need to be collected from within the document content tree and passed as NumPy arrays to training processes. Snippets 5 and 6 show example usage of the methods that *highdicom* provides for these purposes.

```
import highdicom as hd
import numpy as np
from pydicom import dcmread
from pydicom.sr.codedict import codes

# Create highdicom SR document object from a dataset read from file
dataset = dcmread('annotation.dcm')
sr_document = hd.sr.Comprehensive3DSR.from_dataset(dataset)

# Ensure that the document content is a TID 1500 "Measurement Report"
assert isinstance(sr_document.content, hd.sr.MeasurementReport)

# Find bounding polygons for lung tumor regions of interest
tumor_regions: List[np.ndarray] = [
    group.roi.value
    for group in sr_document.content.get_planar_roi_measurement_groups(
        finding_type=codes.SCT.Neoplasm,
        finding_site=codes.SCT.Lung,
        graphic_type=hd.sr.GraphicTypeValues3D.POLYGON
    )
]
```

Code snippet 5: Example of parsing a DICOM SR document to find lung tumor regions.

```
from typing import List
import highdicom as hd
import numpy as np
from pydicom import dcmread
from pydicom.sr.codedict import codes

# Create highdicom segmentation object from a dataset read from file
dataset = dcmread('annotation.dcm')
seg_image = hd.sr.Segmentation.from_dataset(dataset)

# Find numbers of segments for which the segmented property is "tumor"
tumor_segment_numbers: List[int] = seg_image.get_segment_numbers(
    segmented_property_type=codes.SCT.Neoplasm
)

# Get label map with these segments for a given SOP instance
label_map: np.ndarray = seg_image.get_pixels_by_source_instance(
    source_sop_instance_uids=['1.2.3.4'],
    segment_numbers=tumor_segment_numbers,
    combine_segments=True,
    relabel=True
)
```

Code snippet 6: Example of parsing a DICOM Segmentation image to create a label map of lung tumor regions for a given source frame.

Highdicom Facilitates Decoding and Encoding of Annotations During Model Training and Inference, Respectively

To establish a proof-of-concept standard-based ML workflow and to demonstrate the utility of the *highdicom* library for ML, we performed a set of experiments on the training and evaluation of deep convolutional neural network (CNN) models using publicly available slide microscopy (SM) and computed tomography (CT) image datasets. We emphasize that our intent is to demonstrate a complete ML workflow for pathology and radiology fully based on DICOM, rather than create models with optimal performance or reach state-of-the-art for a particular task.

For pathology, we trained and evaluated models using lung cancer collections of slide microscopy (SM) images from The Cancer Imaging Archive (TCIA) [35] that were acquired as part of The Cancer Genome Atlas (TCGA) Lung Adenocarcinoma (LUAD) or Lung Squamous Cell Carcinoma (LUSC) projects and which we converted into DICOM format as previously described [14, 36]. For radiology, we used the collection of CT images of the Lung Image Database Consortium (LIDC) and Image Database Resource Initiative (IDRI) (LIDC-IDRI) [35, 37, 38], which were already available in DICOM format. We used available measurements and qualitative evaluations for these SM and CT files provided by TCIA as image annotations, which we encoded in DICOM SR documents or DICOM Segmentation images using *highdicom* (see supplementary methods), resulting in training sets for CT lung nodule detection and SM image classification encoded entirely within DICOM format.

We developed proof-of-concept ML models based on published algorithms and implemented data pre- and post-processing pipelines for each model to load model inputs from DICOM SM or CT image instances, annotations from DICOM SR documents or DICOM Segmentation images respectively, and store outputs to DICOM SR instances. For pathology, we implemented a weakly supervised image classification model using multiple instance learning with the objective to classify individual SM image frames of lung tissue sections into slide background, normal lung tissue, lung adenocarcinoma, or lung squamous cell carcinoma similar to prior work described by Coudray et al. [39]. To this end, we used a modified version of a ResNet-101 model [40], which we initialized with parameters from pre-training on ImageNet [41] and further optimized using SM image frames and image annotations from the TCGA collections similar to the algorithms described by Lerousseau et al. [42] and Lu et al. [43]. During training, each training sample was created by selecting one or more frames of an SM image from a given series (i.e., digital slide) together with the corresponding image-level annotations obtained from the SR document using *highdicom*. During inference, the data postprocessing pipeline

collects predicted class probabilities for each image frame, constructs low-resolution probabilistic segmentation mask for each class (with pixels representing class probabilities for individual frames), and finally encodes the constructed masks in a DICOM Segmentation image with FRACTIONAL Segmentation Type and PROBABILITY Fractional Segmentation Type (Fig. 5 upper panel). The postprocessing pipeline further thresholds the individual class probability predictions to generate a binary segmentation mask for each class (normal lung tissue, lung adenocarcinoma, or lung squamous carcinoma), performs a connected component analysis and border following to find the contours of ROIs representing class instances, and encodes each detected ROI together with additional measurements and qualitative evaluations in a DICOM Comprehensive 3D SR document (Fig. 5 lower panel).

For radiology, we implemented an object detection model to detect lung nodules in individual CT slices of the chest. We used an off-the-shelf implementation of the widely used RetinaNet convolutional neural network [44] available with the *torchvision* package¹. Specifically, we used a RetinaNet model with the ResNet-50 backbone [40] and initialized the model with weights from pre-training on the ImageNet dataset [41]. During training, each training sample was created by selecting a random CT image frame (2D axial slice) from a given series. The annotations encoded in DICOM Segmentation images were read using *highdicom* and the bounding box containing each nodule in the slice was calculated on-the-fly from the contained segments and used as a ground truth label for supervised training of the RetinaNet model. The post-processing pipeline for the chest CT model collected predicted bounding boxes and their detection scores outputted by the RetinaNet model for every frame in the CT series and encoded them in a DICOM Comprehensive 3D SR document, with vector graphics used to represent bounding box coordinates and detection scores encoded as a measurement of the region represented by the bounding box (Fig. 6).

Annotations Generated by Highdicom can be Stored in Image Management Systems using DICOMweb Services and Visualized using DICOM-Compliant Display Systems After model training, we selected one pathology and radiology model for further clinical evaluation and deployed it into a production-like environment, consisting of an image management system (IMS) with a DICOMweb interface [45] and DICOM-compliant image display systems. Specifically, a *dcm4chee-arc-light* archive² served as the IMS and we stored SM and CT images in the IMS via DICOMweb RESTful services using the *dicomweb-client* Python library [14].³ Upon inference,

¹ <https://pytorch.org/vision/0.8/index.html>

² <https://github.com/dcm4chee/dcm4chee-arc-light>

³ <https://github.com/mghcomputationalpathology/dicomweb-client>

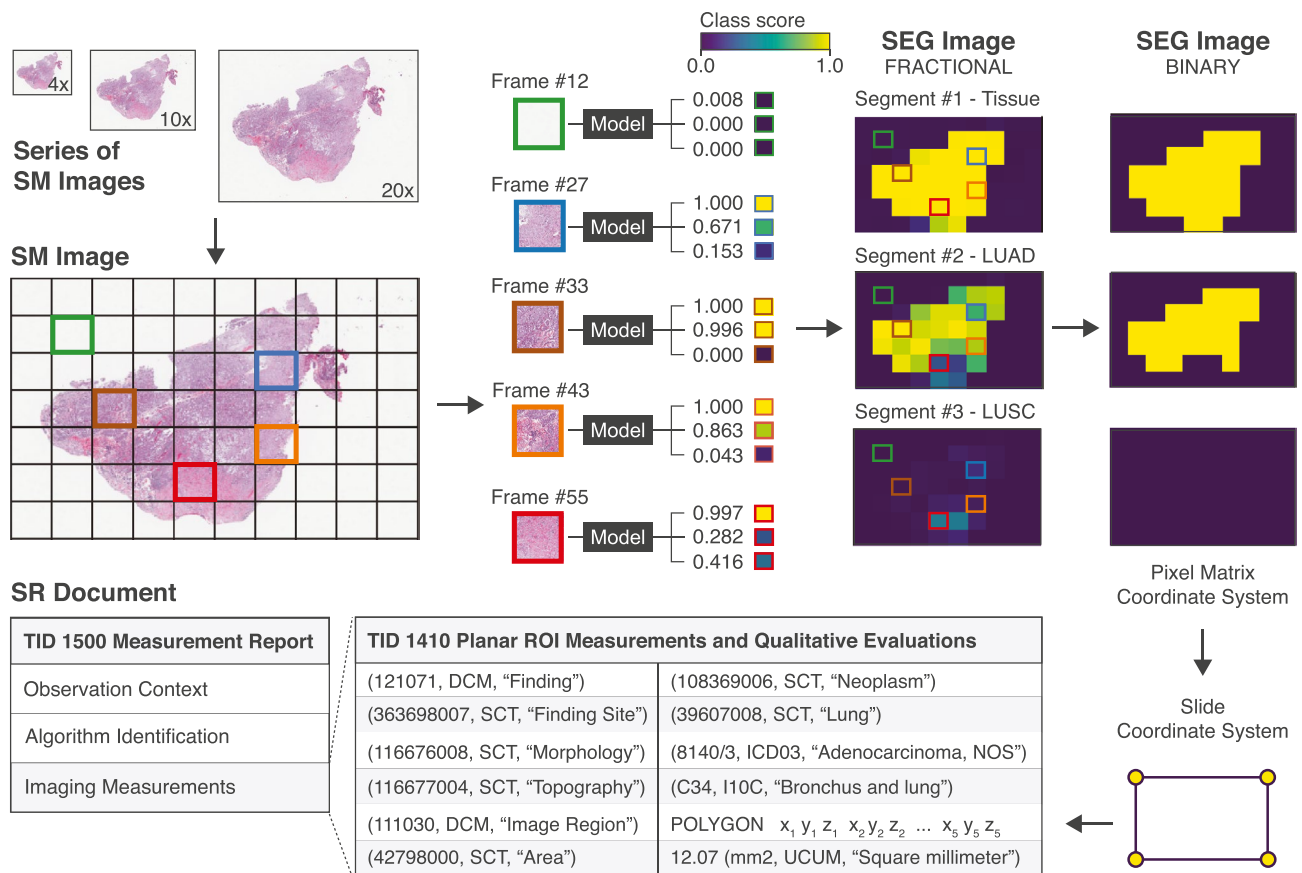


Fig. 5 Schematic overview of output post-processing pipelines of the pathology model, which classifies individual image frames of a multi-frame SM image of a lung tissue section specimen. Outputted scores get transformed into a segmentation mask from which bounding

boxes of the tumor regions are derived. The coordinates of the bounding box vertices are stored as 3D spatial coordinates in the reference slide coordinate system

the data preprocessing pipelines retrieved DICOM SM or CT images from the IMS over network using the *dicomweb-client*, read and interpreted the image metadata and pixel data using *pydicom*, and passed the pixel data as inputs to the model as NumPy arrays. Model outputs received as NumPy arrays were encoded as DICOM SR documents in the data postprocessing pipeline using *highdicom* and stored back in the IMS over network using the *dicomweb-client*.

For radiology, we visualized the ground truth lung nodules using the OHIF⁴ viewer, which retrieved the DICOM Segmentation images over network using the *dicomweb-client* library and displayed each segment as a raster graphic on top of the corresponding CT images (supplementary Fig. S3B). We additionally visualized detected ROIs using the open-source 3D Slicer⁵ software (supplementary Fig. S3B).

For pathology, we visualized detected lung tumor regions using the Slim⁶ viewer, which retrieved the DICOM SR documents over network using the *dicomweb-client* JavaScript library⁷ and displayed the spatial coordinates of each ROI contained in the SR documents as a vector graphic on top of the corresponding SM images (supplementary Fig. S4).

Discussion

The main contributions of this paper are: (i) The demonstration that image annotations can be encoded and exchanged in DICOM format using existing DICOM IODs and services, respectively. (ii) The development of a software library that provides a high-level application programming interface (API) for the Python programming language to facilitate

⁴ <https://github.com/ohif/viewers>

⁵ <https://slicer.org>

⁶ <https://github.com/herrmannlab/slim>

⁷ <https://github.com/dcmjs-org/dicomweb-client>

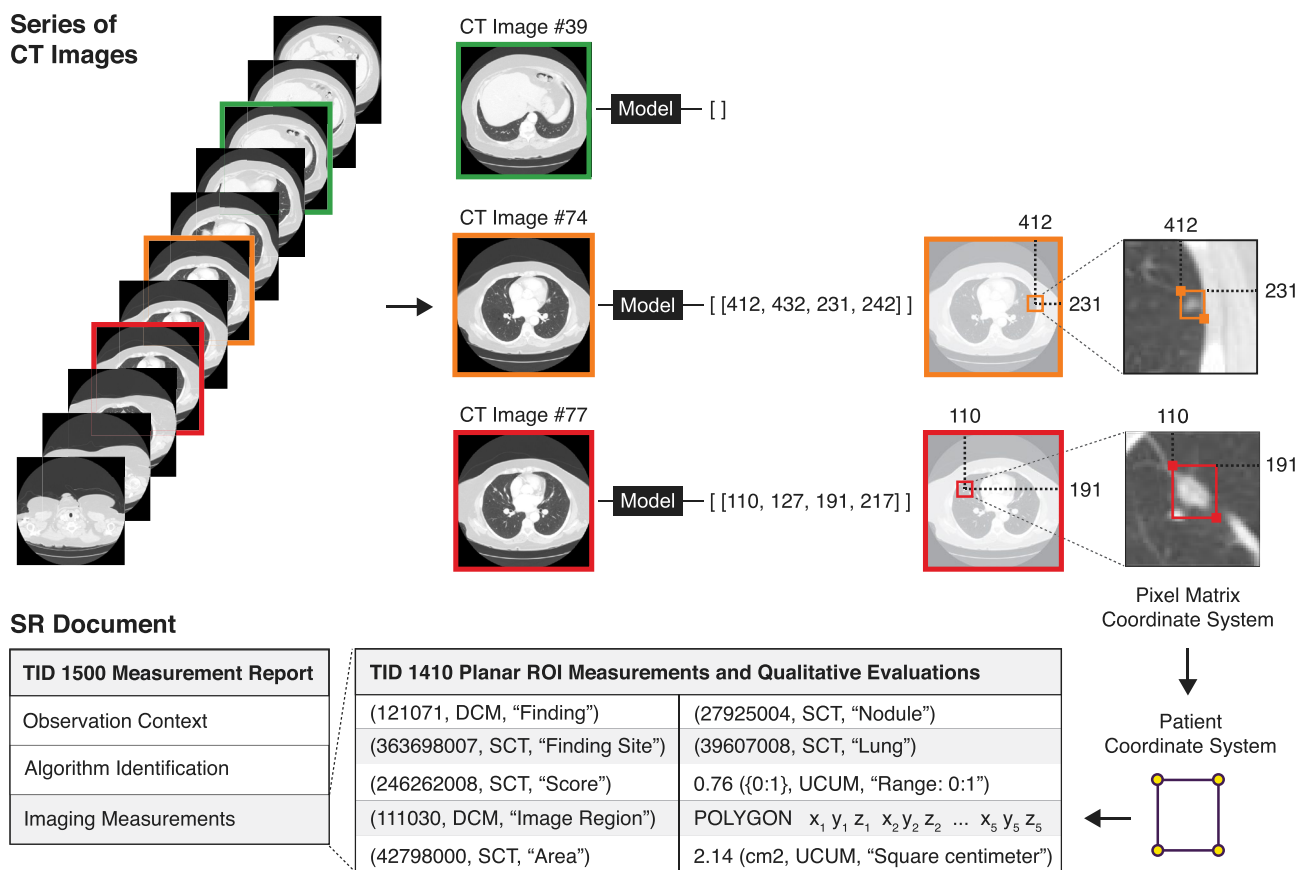


Fig. 6 Schematic overview of output post-processing pipeline of the radiology model, which detects lung nodules in image frames of single-frame CT images of the thorax and outputs bounding boxes of

lung nodule regions. The coordinates of the bounding box vertices are stored as 3D spatial coordinates in the reference patient coordinate system

creation of DICOM objects for storage of image-derived information, including image annotations, as well as accessing and interpreting information stored in DICOM objects. (iii) The establishment of a standard-based workflow for ML model training and inference that is generally applicable across different imaging modalities, computer vision problems, and medical disciplines.

In developing the *highdicom* library and establishing an ML workflow based on DICOM, we made several observations that merit further discussion.

Clinical Use of Machine Learning Model Outputs in Pathology and Radiology Requires Domain-Specific Metadata

Medical images and image annotations must not only contain the actual data, such as the pixel data in case of an image, but require additional metadata that enable interpretation and use of the data. Such metadata can be grouped into information related to data representation, information about the data acquisition process and

equipment, and information related to the clinical context in which the data was acquired, including identifying and descriptive information about the patient, study, and specimens. This contextual information that describes how the data relates to the real world is crucial for unambiguous interpretation of medical images as well as any regions of interest, measurements, or qualitative evaluations derived from them.

To ensure that clinical decisions based on this information are made for the right patient and specimen and in the correct clinical setting, real-world entities need to be uniquely identifiable throughout the digital workflow. As such it is desirable to establish an unambiguous association between the digital information (images and image annotations) on the one hand and clinically relevant real-world entities (patients, specimens, etc.) on the other hand by including clinical identifiers into digital objects. This furthermore facilitates exchange of information between departments and institutions upon transfer and referral of patients. DICOM specifies standard information object definitions and attributes to store and exchange digital images and image-derived

information together with the relevant clinical identifiers as composite objects. The *highdicom* API facilitates access to and creation of such standard DICOM objects using the Python programming language and thereby enables data engineers and scientists to develop ML models and systems that can receive inputs and return outputs that include relevant identifiers for clinical application. Additionally, in many cases including patient information and references to the source images, *highdicom* will find and copy the relevant metadata from the dataset of the source image to reduce the room for human error as far as possible.

In addition to identifiers, DICOM objects contain descriptive metadata about the imaging target (patient or specimen), the imaging modality and procedure, the anatomical location of the imaging or surgical procedure, and in case of pathology the preparation of the specimen. This information can be critical for the interpretation of images or image annotations by ML systems during model training or inference as well as by other systems that use or interpret model outputs. Most importantly this descriptive metadata allows automated systems to decide whether or not a given information object may be appropriate to use in the context of the intended use or select one of several available objects for analysis or display [46]. Descriptive metadata is also useful for performing model validation and error analysis to determine groups of inputs, according to patient demographic information, pre-analytic specimen preparation variables, or image acquisition parameters, upon which models are under-performing. Furthermore, the DICOM standard provides mechanisms for describing the image analysis algorithm (name, version, etc.) as well as the completeness or validity of analysis results at various stages of the clinical decision making process. For example, the DICOM SR IODs include attributes that allow clinical users to verify or, if necessary, complete or correct ML model outputs, to record the verification or modification activity, and to create an audit trail that establishes the relationship between the document containing the verified or modified content and the predecessor document containing the unverified model outputs. These mechanisms are critical for safe clinical application of ML models, since their outputs are generally intended for clinical decision support rather than independent decision making [47] and thus require review by a clinical expert before inclusion into the medical record.

The *highdicom* library enables developers to access relevant descriptive information in received DICOM objects upon preprocessing or include such information into generated DICOM object upon post-processing and thereby make it available to downstream clinical systems. The high-level and well-tested abstractions provided by

highdicom allow developers to achieve this goal with only a few lines of Python code.

Standard Coding Schemes Enable Unambiguous Interpretation of Image Annotations

Subtle differences in the description of imaging findings can lead to drastically different treatment decisions. To ensure that image annotations can be interpreted unambiguously by both clinicians and devices or automated systems that may act upon the information, the terms used to describe and report annotations need to be well-defined. DICOM structured reporting uses codes of established clinical terminologies and ontologies to describe image-derived information rather than using free text. For example, while many words in English and other languages may be used to refer to a “tumor” as the finding type of the ROI, the concept can be unambiguously represented across languages and domains by the SNOMED-CT code “108369006.” The use of structured reports and standardized codes facilitates interpretation of image annotations by both humans and machines and is therefore critical for enabling structural and semantic interoperability between ML models and clinical systems. The standard-based approach further facilitates the re-use of data beyond the scope of the project or use case for which they were initially created. While there are several advantages to using codes, they are cumbersome to work with and increase the complexity of ML programs and are thus in our experience often frowned on by developers. The *highdicom* library provides data structures and methods that abstract the codes and significantly simplify using and operating on coded concepts.

While codes chosen from well-established coding schemes can significantly improve interoperability, the choice of the appropriate code can still pose a significant challenge to both developers and clinical experts. The *highdicom* library does not (and cannot) fully solve this problem. Indeed, in practice, it may be the case that no standard coded concept accurately describes the annotation and a custom coding scheme is required. DICOM allows, and *pydicom* and *highdicom* support, the definition of such custom coding schemes with the convention of a prefix of “99” followed by an identifying text string. Consumers of custom coded concepts should detect this condition and seek out-of-band information for correct interpretation of the annotations. However, for a large range of common clinical use cases, the library (together with the underlying *pydicom* library) exposes value sets defined in the DICOM standard via abstractions, and by depending on these abstractions throughout its API, encourages developers to choose codes from these predefined sets.

Encoding Image Regions in a Well-defined Coordinate System in Three-dimensional Physical Space Allows for Clinically Actionable Measurements

Establishing an unambiguous spatial relationship between ROIs and their corresponding source images for display or computational analysis requires a common frame of reference, which defines the coordinate system to uniquely localize both images or image regions with respect to the imaging target (the specimen in pathology or the patient in radiology) with both position and orientation. Many applications simply specify ROIs relative to the pixel matrix of an image in pixel units. However, this simple approach is problematic for interoperability, because the image pixel grid forms an ill-defined coordinate system and the location (offset, rotation, and scale) of an image with respect to the imaging target changes upon spatial transformation of the image. DICOM specifies a frame of reference for both slide-based and patient-based coordinate systems, which enables accurate and precise localization of a ROI with respect to the patient or the specimen on the slide independent of whether affine transformations have been applied to images. Defining ROIs in physical space in millimeter units further has the advantage that spatial ROI measurements such as diameter or area can be readily taken in this frame of reference without the need to transform coordinates, a process that can be error prone and result in incorrect measurements with potentially serious clinical implications. The *highdicom* library enables developers to work with both 2D pixel matrix and 3D frame of reference coordinates and provides developers methods to readily convert coordinates between the different coordinate systems.

Scaling to Large Numbers of Image Annotations in the Context of Slide Microscopy Imaging in Pathology

As demonstrated in this paper, encoding of ROIs in SR documents works for both pathology and radiology. However, the deeply nested structure of SR documents does not scale well to object detection problems in pathology, where millions of cells or nuclei may be detected per whole slide image. To address this challenge, DICOM Working Group 26 Pathology (WG-26) has developed a supplement for the DICOM standard that proposes the introduction of a *Microscopy Bulk Simple Annotations* IOD and Annotation (ANN) modality specifically designed for the storage and exchange of a large number of image annotations in the form of spatial coordinates [48]. The graphic types used in the ANN objects have been harmonized with those in SRs, and their structure is similar to that of SEG images. This supplement was recently approved and incorporated into the DICOM standard and is

now implemented in *highdicom* as a `highdicom.ann.MicroscopyBulkSimpleAnnotations` SOP class, reusing the existing building blocks of the library for coded concepts and spatial coordinates.

Abstracting the Complexity of the Standard Without Oversimplifying Medical Imaging Use Cases

DICOM is the ubiquitous standard for representation and communication of medical image data and standardizes many aspects of the imaging workflow to enable interoperability in the clinical setting. However, DICOM is often criticized by the biomedical imaging research community for its elaborateness and alternative data formats have emerged in the research setting that are intended to simplify access to and storage of data by researchers that do not want to cope with intricacies of the standard [49]. The first step in an image analysis pipeline is thus often the conversion of DICOM objects into an alternative format that is considered more suitable for research use [50]. While conversion of clinically acquired DICOM objects into another format may work well within the limited scope of a research project, the reverse, i.e., the conversion of a given research output into standard DICOM representation, is generally not possible, since important contextual information is lost along the way [51]. Many of the attributes of DICOM objects that are regarded superfluous by researchers and are readily removed for ease of use are crucial for interoperability with clinical systems and for correct representation and interpretation of the data in clinical practice.

We argue that the discussion regarding the establishment and adoption of standards for clinical deployment of ML models and integration of their outputs into clinical workflows should be guided primarily by the requirements of clinical systems and clinicians for interpretability and clinical decision making, rather than current practices within research communities. The DICOM standard has been evolving over many years through continuous collaboration of an international group of experts and a diverse set of stakeholders based on a considerate and controlled process that takes a variety of use cases as well as legal and regulatory aspects into account. While the comprehensiveness and inclusiveness of the standard has advantages, it has also resulted in significant complexity and demands an implementation that exposes the useful parts of the standards through a layer of abstraction. The *highdicom* library strikes a fine balance, by providing an API that hides as many details of the DICOM standard as possible from model developers, while acknowledging that medical imaging is complex and that efforts aiming for DICOM abstraction should involve technical and domain experts to avoid oversimplification with detrimental effects on interoperability and ultimately patient safety. The result is an API that abstracts the intricate structure of

DICOM datasets, but retains full and direct access to all DICOM attributes and stays close to the terminology of the DICOM data models to avoid any ambiguities.

Bridging the Gap Between Model Development in Research and Model Deployment in Clinical Practice

Researchers, medical device manufacturers, and healthcare providers are generally interested in accelerating the translation of research findings into clinical practice and enable patients to get access to and benefit from diagnostic and therapeutic innovations. However, the incentives for the different stakeholders who participate in the translation process at different time points from model development to deployment are not necessarily well aligned. Currently, the production deployment of an ML model is generally not a major concern to model developers, who primarily operate in a research environment. The developer often does not receive a technical specification against which the model should be developed and is unaware of the environment into which the model should ultimately be deployed for clinical validation. As a consequence, the structure of data outputted by ML models developed in research settings is generally highly customized towards a particular research project and specific use case and lacks identifying or descriptive metadata relevant for clinical application (see above). Furthermore, current ML models store data in a variety of proprietary formats that are incompatible with clinical systems, which generally rely on a DICOM interface for data exchange. Together, these factors impede the deployment of an ML model and its integration into existing clinical workflows for validation or application.

One opportunity for streamlining this process is to rely on DICOM as a common format and interface for data exchange during both model development and deployment. In our experiments, we demonstrated that *highdicom* makes feasible a fully DICOM-based workflow in which all files stored on storage devices are in DICOM format with minimal increase in complexity for the developer. Adapting a model developed in such a workflow for clinical deployment becomes a straightforward task.

A common use for non-DICOM formats is for storage of intermediate results within the input image preprocessing pipeline, such as the results of image registration operations. A limitation of our proposed DICOM-only workflow is that it assumes that model training and inference pre-processing pipelines operate directly on the source images. However, we argue that models developed for eventual clinical deployment must have input preprocessing pipelines that are able to operate efficiently from the raw source data and as such having this constraint in place through model development process simplifies deployment. Furthermore, intermediate

results could also be represented in DICOM format (e.g., using the Spatial Registration IOD for image registration results) and future versions of *highdicom* may provide tools to help with the creation and access of intermediate results in DICOM format.

Common Platforms, Services, and Tools Will Facilitate Enterprise Medical Imaging, Interdisciplinary Research, and Integrated Diagnostics

Standardization of images, image annotations, and model predictions between pathology and radiology opens new avenues for enterprise medical imaging, interdisciplinary quantitative biomedical imaging research, and integrated image-based diagnostics. Despite unique challenges and use cases for image management in digital pathology and radiology, there are opportunities for streamlining the investment into and use of IT infrastructure and platforms across medical disciplines within the enterprise. Given that most hospitals already have an existing medical imaging infrastructure based on DICOM, encoding image annotations in DICOM format may lower the barrier for integration of ML systems into clinical workflows.

Relying on the DICOM standard may further promote interdisciplinary biomedical imaging research by, for example, clearing the way for the use of annotations of slide microscopy images in pathology as ground truth for training ML models for analysis of CT images in radiology or vice versa. Furthermore, leveraging a standard data format and communication interface provides an opportunity to synthesize different imaging modalities and interpret pathology and radiology ML model outputs side-by-side. In this paper, we demonstrate that *highdicom* facilitates the creation and interpretation of image annotations independent of a specific medical imaging modality, discipline, department, or institution. We further show that data can be exchanged and stored using DICOM-compatible image management systems, which already exist in hospitals worldwide and are increasingly being adopted by biomedical imaging research initiatives around the world. For example, the National Cancer Institute's Imaging Data Commons (IDC) in the USA will make large public collections of pathology and radiology images, image annotations, and image analysis results available in DICOM format [52]. The *highdicom* library will allow researchers to leverage these resources and enable them to readily share their results and make them usable by other researchers. We therefore see the potential for *highdicom* to streamline the development and deployment of ML models across departmental boundaries, accelerate the translation of technological innovations from research into clinical practice, and to assist in the realization of AI in healthcare.

Conclusion

The *highdicom* library abstracts the complexity of the DICOM standard, and exposes medical imaging data to ML model developers via a pythonic interface that ties into the scientific Python ecosystem for machine learning and image processing and allows data scientists to think of imaging data at a high level of abstraction without having to worry about the low-level details and rules of the DICOM standard. Focusing on the use case of detecting lung tumors in slide microscopy images of surgical tissue section specimens as well as in computed tomography images of the chest, we examined examples for the interpretation of DICOM-encoded image annotations during model training and encoding of model outputs during model inference. Through a series of experiments, we have demonstrated the utility of the library for the development of ML models and shown that, by relying on the DICOM standard, the library enables interoperability of the developed ML models with commercially available DICOM-compliant information systems and allows for unambiguous interpretation of model outputs in clinical context independent of the specific medical imaging modality or discipline. By facilitating the use of DICOM throughout the model development and deployment process, *highdicom* has the potential to bridge the gap between research and clinical application and thereby streamline clinical integration and validation of ML models.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10278-022-00683-y>.

Acknowledgements The authors thank Dandan Mo and Aidan Stein for assistance with testing the *highdicom* library and deployment of image management and display systems. Further thanks are extended to James Alexander Mays, Lida Harriri, and Anthony John Iafate for discussions and feedback regarding the histomorphologic classification of lung tumors. The authors would also like to express their gratitude to Justin St. Marie, Ira Cooper, and Tom Schultz for support with compute and storage infrastructure and to Bernardo Bizzo for discussions about graphical image annotations and display of model outputs. The results shown here are in part based upon data generated by the TCGA Research Network: <http://cancergenome.nih.gov/>. The authors acknowledge the National Cancer Institute and the Foundation for the National Institutes of Health, and their critical role in the creation of the free publicly available LIDC/IDRI Database used in this study.

Author Contribution M.D.H. initiated and designed the study. C.P.B. and M.D.H. wrote initial version of the manuscript. C.P.B., C.G., and M.D.H. conducted experiments, processed and analyzed data, and created the figures. C.P.B., C.G., S.D., A.Y.F., S.P., D.A.C., and M.D.H. designed, implemented, or tested software. All authors contributed to writing of the manuscript.

Declarations

Conflict of Interest Christopher P. Bridge receives research funding from GE Healthcare, Nvidia Corporation, and Bayer, AG. Steven Pieper is in

part supported by the National Institute of Biomedical Imaging and Bioengineering (NIBIB) (P41 EB015902) and by the National Cancer Institute, National Institutes of Health, under Task Order No. HHSN26110071 under Contract No. HHSN2612015000031. Jayashree Kalpathy-Cramer is in part supported by the National Cancer Institute CI U01CA242879, and receives research funding from GE Healthcare, Genentech, Inc, and Bayer, AG. Andriy Y. Fedorov is in part supported by the National Institute of Biomedical Imaging and Bioengineering: contracts 75N92020C00008 and 75N92020C00021; grant P41 EB028741, and by the National Cancer Institute: Task Order No. HHSN26110071 under Contract No. HHSN2612015000031; grants U24 CA264044 and R01 CA241817. David A. Clunie receives financial compensation as a consultant of Philips Algotec, as a consultant for Essex Leidos CBIIT NCI, as a consultant for Brigham and Women's Hospital NCI Imaging Data Commons (IDC), as a consultant for the University of Leeds Northern Pathology Imaging Co-operative (NPIC), and as a contractor for NEMA as DICOM Editor. Markus D. Herrmann is in part supported by the National Cancer Institute: Task Order No. HHSN26110071 under Contract No. HHSN2612015000031. The content of this publication does not necessarily reflect the views or policies of the Department of Health and Human Services, nor does mention of trade names, commercial products, or organizations imply endorsement by the US Government.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. LeCun, Y., Bengio, Y., and Hinton, G. "Deep learning". *Nature* 521.7553 (2015), pp. 436–444.
2. Campanella, G., Hanna, M. G., Geneslaw, L., Miralflor, A., Werneck Krauss Silva, V., Busam, K. J., Brogi, E., Reuter, V. E., Klimstra, D. S., and Fuchs, T. J. "Clinical-grade computational pathology using weakly supervised deep learning on whole slide images". *Nature Medicine* 25.8 (July 2019), pp. 1301–1309.
3. Lu, M. Y., Chen, T. Y., Williamson, D. F. K., Zhao, M., Shady, M., Lipkova, J., and Mahmood, F. "AI-based pathology predicts origins for cancers of unknown primary". *Nature* 594.7861 (June 2021), pp. 106–110.
4. Laak, J. van der, Litjens, G., and Ciompi, F. "Deep learning in histopathology: the path to the clinic". *Nat Med* 27.5 (May 2021), pp. 775–784.
5. McKinney, S. M., Sieniek, M., Godbole, V., Godwin, J., Antropova, N., Ashrafian, H., Back, T., Chesus, M., Corrado, G. S., Darzi, A., Etemadi, M., Garcia-Vicente, F., Gilbert, F. J., Halling-Brown, M., Hassabis, D., Jansen, S., Karthikesalingam, A., Kelly, C. J., King, D., Ledsam, J. R., Melnick, D., Mostofi, H., Peng, L., Reicher, J. J., Romera-Paredes, B., Sidebottom, R., Suleyman, M., Tse, D., Young, K. C., De Fauw, J., and Shetty, S. "International evaluation of an AI system for breast cancer screening". *Nature* 577.7788 (Jan. 2020), pp. 89–94.

6. Choy, G., Khalilzadeh, O., Michalski, M., Do, S., Samir, A. E., Panykh, O. S., Geis, J. R., Pandharipande, P. V., Brink, J. A., and Dreyer, K. J. "Current Applications and Future Impact of Machine Learning in Radiology". *Radiology* 288.2 (Aug. 2018), pp. 318–328.
7. Ardila, D., Kiraly, A. P., Bharadwaj, S., Choi, B., Reicher, J. J., Peng, L., Tse, D., Etemadi, M., Ye, W., Corrado, G., Naidich, D. P., and Shetty, S. "End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography". *Nat Med* 25.6 (June 2019), pp. 954–961.
8. Hosny, A. and Aerts, H. J. W. L. "Artificial intelligence for global health". *Science* 366.6468 (Nov. 2019), pp. 955–956.
9. Allen, B., Seltzer, S. E., Langlotz, C. P., Dreyer, K. P., Summers, R. M., Petrick, N., Marinac-Dabic, D., Cruz, M., Alkasab, T. K., Hanisch, R. J., Nilsen, W. J., Burleson, J., Lyman, K., and Kandarpa, K. "A Road Map for Translational Research on Artificial Intelligence in Medical Imaging: From the 2018 National Institutes of Health/RSNA/ACR/The Academy Workshop". *J Am Coll Radiol* 16.9 Pt A (Sept. 2019), pp. 1179–1189.
10. Granter, S. R., Beck, A. H., and Papke, D. J. "AlphaGo, Deep Learning, and the Future of the Human Microscopist". *Archives of Pathology & Laboratory Medicine* 141.5 (2017), pp. 619–621.
11. Abels, E., Pantanowitz, L., Aeffner, F., Zarella, M. D., Laak, J. van der, Bui, M. M., Vemuri, V. N., Parwani, A. V., Gibbs, J., Agosto-Arroyo, E., Beck, A. H., and Kozlowski, C. "Computational pathology definitions, best practices, and recommendations for regulatory guidance: a white paper from the Digital Pathology Association". *J. Pathol.* (July 2019).
12. Roth, C. J., Lannum, L. M., and Persons, K. R. "A Foundation for Enterprise Imaging: HIMSS-SIIM Collaborative White Paper". 29.5 (Oct. 2016), pp. 530–538.
13. Clunie, D., Hosseinzadeh, D., Wintell, M., De Mena, D., Lajara, N., Garcia-Rojo, M., Bueno, G., Saligrama, K., Stearrett, A., Toomey, D., Abels, E., Apeldoorn, F. V., Langevin, S., Nichols, S., Schmid, J., Horchner, U., Beckwith, B., Parwani, A., and Pantanowitz, L. "Digital Imaging and Communications in Medicine Whole Slide Imaging Connectathon at Digital Pathology Association Pathology Visions 2017". *Journal of Pathology Informatics* 9 (2018), p. 6.
14. Herrmann, M. D., Clunie, D. A., Fedorov, A., Doyle, S. W., Pieper, S., Klepeis, V., Le, L. P., Mutter, G. L., Milstone, D. S., Schultz, T. J., Kikinis, R., Kotecha, G. K., Hwang, D. H., Andriole, K. P., Iaffrate, A. J., Brink, J. A., Boland, G. W., Dreyer, K. J., Michalski, M., Golden, J. A., Louis, D. N., and Lennerz, J. K. "Implementing the DICOM Standard for Digital Pathology". *Journal of Pathology Informatics* 9 (2018), p. 37.
15. Dash, R., Jones, C., Merrick, R., Haroske, G., Harrison, J., Sayers, C., Haarselhorst, N., Wintell, M., Herrmann, M., and Macary, F. "Integrating the health-care enterprise pathology and laboratory medicine guideline for digital pathology interoperability". *J Pathol Inform* 12.16 (Mar. 2021).
16. IHE PaLM Technical Committee in collaboration with DICOM WG26. *IHE Pathology and Laboratory Medicine Technical Framework Supplement Digital Pathology Workflow – Image Acquisition (DPIA)*. https://www.ihe.net/uploadedFiles/Documents/PaLM/IHE_PaLM_Suppl_DPIA.pdf. Aug. 2020.
17. Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., Blomberg, N., Boiten, J. W., Silva Santos, L. B. da, Bourne, P. E., Bouwman, J., Brookes, A. J., Clark, T., Crosas, M., Dillo, I., Dumon, O., Edmunds, S., Evelo, C. T., Finkers, R., Gonzalez-Beltran, A., Gray, A. J., Groth, P., Goble, C., Grethe, J. S., Heringa, J., Hoen, P. A., Hooft, R., Kuhn, T., Kok, R., Kok, J., Lusher, S. J., Martone, M. E., Mons, A., Packer, A. L., Persson, B., Rocca-Serra, P., Roos, M., Schaik, R. van, Sansone, S. A., Schultes, E., Sengstag, T., Slater, T., Strawn, G., Swertz, M. A., Thompson, M., Lei, J. van der, Mulligen, E. van, Velterop, J., Waagmeester, A., Wittenburg, P., Wolstencroft, K., Zhao, J., and Mons, B. "The FAIR Guiding Principles for scientific data management and stewardship". *Sci Data* 3 (Mar. 2016), p. 160018.
18. Fedorov, A., Clunie, D., Ulrich, E., Bauer, C., Wahle, A., Brown, B., Onken, M., Riesmeier, J., Pieper, S., Kikinis, R., Buatti, J., and Beichel, R. R. "DICOM for quantitative imaging biomarker development: a standards based approach to sharing clinical data and structured PET/CT analysis results in head and neck cancer research". *PeerJ* 4 (2016), e2057.
19. Herz, C., Fillion-Robin, J. C., Onken, M., Riesmeier, J., Lasso, A., Pinter, C., Fichtinger, G., Pieper, S., Clunie, D., Kikinis, R., and Fedorov, A. "dcmqi: An Open Source Library for Standardized Communication of Quantitative Image Analysis Results Using DICOM". *Cancer Research* 77.21 (2017), e87–e90.
20. IHE Radiology Technical Committee. *IHE Radiology Technical Framework Supplement AI Results (AIR)*. https://www.ihe.net/uploadedFiles/Documents/Radiology/IHE_RAD_Suppl_AIR.pdf. June 2020.
21. Virtanen, P. et al. "SciPy 1.0: fundamental algorithms for scientific computing in Python". *Nat Methods* 17.3 (Mar. 2020), pp. 261–272.
22. Harris, C. R., Millman, K. J., Walt, S. J. van der, Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., Kerkwijk, M. H. van, Brett, M., Haldane, A., Del Ró, J. F., Wiebe, M., Peterson, P., Gérard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. "Array programming with NumPy". *Nature* 585.7825 (Sept. 2020), pp. 357–362.
23. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, É. "Scikit-Learn: Machine Learning in Python". *J. Mach. Learn. Res.* 12 (Nov. 2011), pp. 2825–2830.
24. Walt, S. van der, Schönberger, J., Nunez-Iglesias, J., Boulogne, F., Warner, J., Yager, N., Gouillart, E., Yu, T., and contributors, the scikit-image. "scikit-image: image processing in Python". *PeerJ* 2 (2014), e453.
25. Mason, D. "SU-E-T-33: pydicom: an open source DICOM library". *Medical Physics* 38.6 Part 10 (2011), pp. 3493–3493.
26. Hapke, H. and Nelson, C. Building machine learning pipelines. O'Reilly Media, 2020.
27. Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P., and Aroyo, L. M. "“Everyone Wants to Do the Model Work, Not the Data Work”: Data Cascades in High-Stakes AI". *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. Yokohama, Japan: Association for Computing Machinery, 2021.
28. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang L. Bai, J., and Chintala, S. "PyTorch: An Imperative Style, High-Performance Deep Learning Library". *Advances in Neural Information Processing Systems* 32. Curran Associates, Inc., 2019, pp. 8026–8037.
29. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. "Tensorflow: A system for large-scale machine learning". *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. 2016, pp. 265–283.
30. Bradski, G. "The OpenCV Library". *Dr. Dobbs's Journal of Software Tools* (2000).
31. Yoo, T. S., Ackerman, M. J., Lorensen, W. E., Schroeder, W., Chalana, V., Aylward, S., Metaxas, D., and Whitaker, R. "Engineering and algorithm design for an image processing API: a technical report on ITK-the insight toolkit". *Medicine Meets Virtual Reality 02/10*. IOS press, 2002, pp. 586–592.

32. Pianykh, O. S. *Digital imaging and communications in medicine (DICOM): a practical introduction and survival guide*. Vol. 6. Springer, 2008.
33. Hafiz, A. M. and Bhat, G. M. "A survey on instance segmentation: state of the art". *International Journal of Multimedia Information Retrieval* 9.3 (2020), pp. 171–189.
34. Bidgood, W. D. "The SNOMED DICOM microglossary: controlled terminology resource for data interchange in biomedical imaging". *Methods Inf Med* 37.4-5 (Nov. 1998), pp. 404–414.
35. Clark, K., Vendt, B., Smith, K., Freymann, J., Kirby, J., Koppel, P., Moore, S., Phillips, S., Maffitt, D., Pringle, M., Tarbox, L., and Prior, F. "The Cancer Imaging Archive (TCIA): Maintaining and Operating a Public Information Repository". *Journal of Digital Imaging* 26.6 (Dec. 2013), pp. 1045–1057.
36. Clunie, D. A. "Dual-Personality DICOM-TIFF for Whole Slide Images: A Migration Technique for Legacy Software". *Journal of Pathology Informatics* 10 (2019), p. 12.
37. Armato III, S. G., McLennan, G., Bidaut, L., McNitt-Gray, M. F., Meyer, C. R., Reeves, A. P., Zhao, B., Aberle, D. R., Henschke, C. I., Hoffman, E. A., Kazerooni, E. A., MacMahon, H., Beek, E. J. R. van, Yankelevitz, D., Biancardi, A. M., Bland, P. H., Brown, M. S., Engelmann, R. M., Laderach, G. E., Max, D., Pais, R. C., Qing, D. P.-Y., Roberts, R. Y., Smith, A. R., Starkey, A., Batra, P., Caligiuri, P., Farooqi, A., Gladish, G. W., Jude, C. M., Munden, R. F., Petkovska, I., Quint, L. E., Schwartz, L. H., Sundaram, B., Dodd, L. E., Fenimore, C., Gur, D., Petrick, N., Freymann, J., Kirby, J., Hughes, B., Vande Castele, A., Gupte, S., Sallam, M., Heath, M. D., Kuhn, M. H., Dharaiya, E., Burns, R., Fryd, D. S., Salganicoff, M., Anand, V., Shreter, U., Vastagh, S., Croft, B. Y., and Clarke, L. P. "The Lung Image Database Consortium (LIDC) and Image Database Resource Initiative (IDRI): A Completed Reference Database of Lung Nodules on CT Scans". *Medical Physics* 38.2 (2011), pp. 915–931.
38. Armato III, S., McLennan, G., Bidaut, L., McNitt-Gray, M., Meyer, C., Reeves, A., Zhao, B., Aberle, D., Henschke, C., Hoffman, E., Kazerooni, E., MacMahon, H., Beek, E. van, Yankelevitz, D., Biancardi, A., Bland, P., Brown, M., Engelmann, R., Laderach, G., Max, D., Pais, R., Qing, D., Roberts, R., Smith, A., Starkey, A., Batra, P., Caligiuri, P., Farooqi, A., Gladish, G., Jude, C., Munden, R., Petkovska, I., Quint, L., Schwartz, L., Sundaram, B., Dodd, L., Fenimore, C., Gur, D., Petrick, N., Freymann, J., Kirby, J., Hughes, B., Castele, A., Gupte, S., Sallam, M., Heath, M., Kuhn, M., Dharaiya, E., Burns, R., Fryd, D., Salganicoff, M., Anand, V., Shreter, U., Vastagh, S., Croft, B., and Clarke, L. *Data From LIDC-IDRI*. Tech. rep. The Cancer Imaging Archive, 2015.
39. Coudray, N., Ocampo, P. S., Sakellaropoulos, T., Narula, N., Snuderl, M., Fenyó, D., Moreira, A. L., Razavian, N., and Tsirigos, A. "Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning". *Nature Medicine* 24.10 (2018), pp. 1559–1567.
40. He, K., Zhang, X., Ren, S., and Sun, J. "Deep Residual Learning for Image Recognition". *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
41. Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. "ImageNet: A Large-Scale Hierarchical Image Database". *CVPR09*. 2009.
42. Lerousseau, M., Vakalopoulou, M., Classe, M., Adam, J., Battistella, E., Carré, A., Estienne, T., Henry, T., Deutsch, E., and Paragios, N. "Weakly Supervised Multiple Instance Learning Histopathological Tumor Segmentation". *Medical Image Computing and Computer Assisted Intervention – MICCAI 2020*. Springer International Publishing, 2020, pp. 470–479.
43. Lu, M. Y., Williamson, D. F. K., Chen, T. Y., Chen, R. J., Barbieri, M., and Mahmood, F. "Data-efficient and weakly supervised computational pathology on whole-slide images". *Nat Biomed Eng* (Mar. 2021).
44. Lin, T.-Y., Goyal, P., Girshick, R., He, K., and Dollár, P. "Focal loss for dense object detection". *Proceedings of the IEEE international conference on computer vision*. 2017, pp. 2980–2988.
45. DICOM Standards Committee. *DICOM PS3.18 – Web Services*. <http://dicom.nema.org/medical/dicom/current/output/chtml/part18/PS3.18.html>. 2021.
46. Gauriau, R., Bridge, C., Chen, L., Kitamura, F., Tenenholtz, N. A., Kirsch, J. E., Andriole, K. P., Michalski, M. H., and Bizzo, B. C. "Using DICOM Metadata for Radiological Image Series Categorization: a Feasibility Study on Large Clinical Brain MRI Datasets". *J Digit Imaging* 33.3 (June 2020), pp. 747–762.
47. Magrabi, F., Ammenwerth, E., McNair, J. B., De Keizer, N. F., Hypponen, H., Nykänen, P., Rigby, M., Scott, P. J., Vehko, T., Wong, Z. S.-Y., and Georgiou, A. "Artificial Intelligence in Clinical Decision Support: Challenges for Evaluating AI and Practical Implications". *Yearb Med Inform* 28 (01 2019), pp. 128–134.
48. DICOM Standards Committee, Working Group 26 (Pathology). *Supplement 222: Microscopy Bulk Simple Annotations Storage SOP Class*. ftp://medical.nema.org/medical/dicom/sups/LB/sup222_lb_WSIAnnotations.pdf. 2021.
49. Larobina, M. and Murino, L. "Medical image file formats". *J Digit Imaging* 27.2 (Apr. 2014), pp. 200–206.
50. Li, X., Morgan, P. S., Ashburner, J., Smith, J., and Rorden, C. "The first step for neuroimaging data analysis: DICOM to NIfTI conversion". *J. Neurosci. Methods* 264 (Apr. 2016), pp. 47–56.
51. Roberts, M., Driggs, D., Thorpe, M., Gilbey, J., Yeung, M., Ursprung, S., Aviles-Rivero, A. I., Etmann, C., McCague, C., Beer, L., et al. "Common pitfalls and recommendations for using machine learning to detect and prognosticate for COVID-19 using chest radiographs and CT scans". *Nature Machine Intelligence* 3.3 (2021), pp. 199–217.
52. Fedorov, A., Longabaugh, W., Pot, W., Clunie, D., Pieper, S., Aerts Hugo, J., Homeyer, A., Lewis, R., Akbarzadeh, A., Bontempi, D., Clifford, D., Herrmann, M., Höfener, H., Octaviano, I., Osborne, C., Paquette, S., Petts, J., Punzo, D., Reyes, M., Schacherer, D., Tian, M., White, G., Ziegler, E., Shmulevich, I., Pihl, T., Wagner, U., Farahani, K., and R, K. "NCI Imaging Data Commons". *Cancer Research* (2021).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.