
VTK Classes for DICOM Data

Release 0.7.3

David G. Gobbi^{1,2}

January 16, 2016

¹Department of Radiology, University of Calgary, Calgary, Canada

²Calgary Image Processing and Analysis Centre, Foothills Medical Centre, Calgary, Canada
david.gobbi@gmail.com

Abstract

This document describes a library of C++ classes for using DICOM files with VTK. These classes provide easy and complete access to the pixel data and to the meta data, and provide correct orientation and geometrical representation of the images within VTK. Multi-dimensional image volumes are supported, up to a maximum of five dimensions.

Contents

1	Introduction	3
2	Meta Data and Pixel Data	3
2.1	Tag	3
2.2	VR (Value Representation)	4
2.3	Value	4
2.4	Accessing meta data	8
2.5	Iterating over data elements	9
2.6	Dictionaries	9
2.7	Private meta data	10
2.8	Enhanced multi-frame data	11
3	Reading DICOM Files into VTK	12
3.1	Sorting files into a series	12
3.2	Reading a series of files as a VTK image	13
3.3	Enhanced multi-frame and multi-stack files	14
3.4	DICOM image orientation	15
3.5	Obtaining per-slice meta data	16
3.6	Obtaining per-component meta data	16
3.7	Real-world values and lookup tables	18
4	Writing DICOM Files	19
4.1	Conforming to the standard	19
4.2	Writing a VTK image as DICOM	20
4.3	Customizing the generators	21
4.4	Writing a raw pixel buffer to a DICOM file	21
5	Locating DICOM Files	22
5.1	Reading from a DICOM file set	22
5.2	Searching for files that match a query	23
6	Future Work	23
	Appendix A: Enabling the vtkDICOM module in VTK	24
	Appendix B: Building a stand-alone vtkDICOM library	24

1 Introduction

If you are reading this, then you have some knowledge of what DICOM is, and have a need to use DICOM with VTK. You might have been unsatisfied with the `vtkDICOMImageReader` that comes with VTK, since there are many DICOM files that it cannot read, or you might have the need to write DICOM files, which VTK by itself cannot do. You might have tried the VTK reader and writer that come with GDCM, but found that they make it difficult to access the DICOM meta data from within your VTK programs. Read onwards to see if `vtk-dicom` might satisfy your needs.

The purpose of `vtk-dicom` is to provide a set of core classes to enable DICOM data access through VTK. It is currently limited to DICOM files, and provides no network functionality and cannot directly communicate with a PACS. Nor does it provide any of its own GUI components. The sole intent of these classes is to read and write DICOM files with full fidelity and to provide easy access to all of the data contained in the files.

2 Meta Data and Pixel Data

This document makes frequent references to “meta data,” even though this term is not used in the DICOM standard. The meta data is any information in a data set (for our purposes, a DICOM file) apart from the pixels themselves. To better support this distinction, the `vtkDICOMMetaData` class holds only the meta data, while the pixel values are stored in a `vtkImageData` object.

All attributes of a DICOM data set are stored as an ordered series of data elements, where each data element consists of a 32-bit key called a *tag* along with an attached *value* consisting of zero or more bytes (always an even number of bytes). The structure of a data element is shown in Fig. 1. In the `vtkDICOMMetaData` class, an efficient mapping of *tags* to *values* is provided via a hash table.

If you are familiar with the structure of DICOM data elements, then you can safely skip everything the next two sections except for the code samples.

2.1 Tag

Each attribute in the DICOM meta data is identified by a 32-bit tag that is written in hexadecimal as a four-digit *group number* followed by a four-digit *element number*. For example, (0008,103E) is the tag for the DICOM “Series Description” data element. Each tag that uses an even group number has a strict definition within the DICOM public standard, while tags that use an odd group number are defined within the internal standards of one of the many medical device manufacturers. The latter are often called *private tags* because the manufacturers are not required to reveal their definition to the public. All non-private tags, however, are described in parts 3 and 6 of the DICOM standard [1].




Figure 1: A DICOM DataElement.

```

// Construct a tag with default value (0x0000,0x0000).
vtkDICOMTag tag;

// Two ways to create a tag.
tag = vtkDICOMTag(0x0028, 0x0030);
tag = vtkDICOMTag(DC::PixelSpacing);

// Get the group or element from a tag.
unsigned short group = tag.GetGroup();
unsigned short element = tag.GetElement();

```

2.2 VR (Value Representation)

Each DICOM data element has a two-character VR code that specifies the type and format of the associated value. Table 1 provides a full list of the VRs that exist within the DICOM standard. VRs specify either a text or binary format, and most text VR formats use backslash as a value separator. For example an element with VR of CS and contents of “ORIGINAL\PRIMARY” is said to contain two string values, “ORIGINAL” and “PRIMARY”. Notable exceptions to this backslash-as-separator rule are the ST, LT, and UT VRs; these are intended to be used to store one or more paragraphs of free-form text in which a backslash is interpreted as a backslash and not a separator. Within these VRs, line breaks between the paragraphs should be indicated with CR NL (a carriage return followed by a newline), never by NL on its own. In comparison, the SH and LO text VRs are meant for short single-line blocks of text and do not allow CR or NL at all. Furthermore, leading and trailing spaces in single-line text VRs are not significant and should be removed before the value is used in an application.

```

// Construct a VR with an invalid initial value.
vtkDICOMVR vr;

// Two ways to create a VR.
vr = vtkDICOMVR("CS");
vr = vtkDICOMVR(vtkDICOMVR::CS);

// Get information about the VR.
const char *name = vr.GetText();
bool isText = (vr.GetType() == VTK_CHAR);

```

2.3 Value

The value within a DICOM data element can actually be an array of values. The allowed number of items in the array is called the *value multiplicity* or VM. For most attribute values defined in the DICOM standard, the VM is one. As stated in the previous section, text values are separated by backslashes, e.g. “1.000\2.000” while binary values such as FL (float) and FD (double) are stored as an array of their native type.

The vtkDICOMValue class provides automatic conversion between numeric values stored as text, and the native C++ numeric types. So if a number is stored in text as 512 you can retrieve it as an int, unsigned int, or double. And likewise, if a value is stored as short or float, you can retrieve it as a text string. If the conversion is impossible, then the result will be an empty string (if attempting to convert to a string) or zero (if attempting to convert to a C++ numerical type).

Table 1: List of DICOM Value Representations (VRs).

VR	VR Name	Format	Example	Type	Length
AE	Application Entity	(see notes 1,2,3)	MyLabNode01	char	16 bytes max
AS	Age String	nnnY (or D,W,M)	038Y	char	4 bytes fixed
AT	Attribute Tag	<i>binary</i>		<i>tag</i>	4 bytes fixed
CS	Code String	(see note 5)	ORIGINAL	char	16 bytes max
DA	Date	YYYYMMDD	20130604	char	8 bytes fixed
DS	Decimal String	+/-, 0-9, E/e and .	10.2	char	16 bytes max
DT	Date Time	(see note 6)	20130604160259	char	26 bytes max
FL	Float	<i>binary</i>		float	4 bytes fixed
FD	Float (Double)	<i>binary</i>		double	8 bytes fixed
IS	Integer String	+/-, 0-9	15	char	12 bytes max
LO	Long String	(see notes 1,2,4)		char	64 bytes max
LT	Long Text	(see note 1)		char	10240 bytes max
OB	Other Bytes	<i>binary</i>		unsigned char	unlimited
OD	Other Doubles	<i>binary</i>		double	2 ³² –8 bytes max
OF	Other Floats	<i>binary</i>		float	2 ³² –4 bytes max
OL	Other Long	<i>binary</i>		int	2 ³² –4 bytes max
OW	Other Words	<i>binary</i>		short	2 ³² –2 bytes max
PN	Person Name	(see notes 1,2,4)	Doe^John	char	64 bytes max
SH	Short String	(see notes 1,2,4)		char	16 bytes max
SL	Signed Long	<i>binary</i>		int	4 bytes fixed
SQ	Sequence	<i>binary</i>		<i>item</i>	unlimited
SS	Signed Short	<i>binary</i>		short	2 bytes fixed
ST	Short Text	(see note 1)		char	1024 bytes max
TM	Time	HHMMSS[.FF...]	160259.569834	char	16 bytes max
UC	Unlimited Characters	(see notes 1,2,4)		char	2 ³² –2 bytes max
UI	Unique Identifier	0-9, . (see note 7)	1.234.2345	char	64 bytes max
UL	Unsigned Long	<i>binary</i>		unsigned int	4 bytes fixed
UN	Unknown	<i>binary</i>			unlimited
UR	URI or URL	(see note 8)		char	2 ³² –2 bytes max
US	Unsigned Short	<i>binary</i>		unsigned short	2 bytes fixed
UT	Unlimited Text	(see note 1)		char	2 ³² –2 bytes max

1. The default character repertoire for DICOM includes all printing characters in ASCII but no control characters except for CR, NL, FF, and ESC. A new line should always be indicated by CR NL, rather than NL alone. Other character repertoires may only be used if the Specific Character Set (0008,0005) attribute is set for the DICOM data set, and even then the use of the new repertoire is limited to the SH, LO, PN, ST, LT, and UT VRs.
2. Backslash (\) is a value separator, and is not permitted in text VRs that allow multiple data values.
3. This VR does not permit any of the control characters CR, NL, FF, or ESC.
4. This VR does not permit CR, NL, or FF, but it does permit ESC.
5. This VR only permits A-Z, 0-9, underscore, and space. Leading/trailing spaces should be ignored if present.
6. The full DT format is YYYYMMDDHHMMSS.FFFFFFFF+ZZXX where .FFFFFFF optionally gives a fraction of a second (up to 6 digits), and +ZZXX (or -ZZXX) optionally gives an offset in hours and minutes from UTC.
7. If the number of bytes in an element containing Unique Identifiers would be odd, it will be padded to even with a null byte.
8. This VR is restricted to characters allowed in URLs (see RFC 3986). Trailing space must be ignored if present.

Table 2: The default DICOM character repertoire.

00–0F	10–1F	20–2F	30–3F	40–4F	50–5F	60–6F	70–7F
		SP	0	@	P	`	p
		!	1	A	Q	a	q
		"	2	B	R	b	r
		#	3	C	S	c	s
		\$	4	D	T	d	t
		%	5	E	U	e	u
		&	6	F	V	f	v
		'	7	G	W	g	w
		(8	H	X	h	x
)	9	I	Y	i	y
NL		*	:	J	Z	j	z
	ESC	+	;	K	[k	{
FF		,	<	L	\	l	
CR		-	=	M]	m	}
		.	>	N	^	n	~
		/	?	O	_	o	

```

// Create an empty, typeless value.
vtkDICOMValue v;

// Get a value from the meta data.
v = metaData->GetAttributeValue(DC::PixelSpacing);

// Result is valid only if the attribute was present in the meta data.
if (!v.IsValid())
{
    // Attribute not found!
    // You can also use metaData->HasAttribute(DC::PixelSpacing).
}

// Verify that the correct number of items are present.
// The NumberOfValues is zero if v is invalid.
if (v.GetNumberOfValues() == 2)
{
    // The Get methods use assert() for range checking, so be careful!
    // Always check GetNumberOfValues() before calling these methods.
    double xs = v.GetDouble(0);
    double ys = v.GetDouble(1);
}

// Get an array of values from the meta data.
v = metaData->GetAttributeValue(DC::ImagePositionPatient);
double position[3] = { 0.0, 0.0, 0.0 };
if (v.GetNumberOfValues() == 3)
{
    v.GetValues(position, 3);
}

```

```

// Get values that have a VM of 1 (preferred method).
std::string name = metaData->GetAttributeValue(DC::Modality).AsString();
int rows = metaData->GetAttributeValue(DC::Rows).AsInt();

// Get pointer to a native C array (less safe, but very efficient).
// For text VRs, a null-terminated string is always returned. The text
// might be padded with trailing spaces to make an even number of bytes!
// The only exception is UIDs, which are never padded with spaces.
vtkDICOMValue v = metaData->GetAttributeValue(DC::StudyInstanceUID);
const char *uid = v.GetCharData();
if (uid != 0)
{
    // Use the pointer only if the returned value isn't NULL!
    // Also remember that the pointer becomes invalid when the value destructs.
    // Maybe you should use the safer v.AsString() method instead!
}

// Get a value that holds a tag pointer to another value.
vtkDICOMTag tag = metaData->GetAttributeValue(DC::FrameIncrementPointer).AsTag();
if (tag != vtkDICOMTag(0,0)) // make sure tag isn't zero
{
    v = metaData->GetAttributeValue(tag);
    // do something with the referenced value
}

```

Values also provide powerful matching methods, that allow you to check:

1. if a text value matches a wildcard pattern containing "*" or "?".
2. if a date or time falls within a range specified with "-".
3. if a value contains a given code string (VR of CS).
4. if a value is equal to a given numerical value.
5. if a UID is one of a series of specified UIDs.

The `vtkDICOMValue Matches()` method is better than comparing with "==" because it automatically ignores any value padding and handles values with multiplicity greater than one. Furthermore, person names will be compared in a case-insensitive "compatibility mode".

```

// Check if a series of code strings contains a value.
v = vtkDICOMValue(vtkDICOMVR::CS, "ORIGINAL\\PRIMARY");
if (v.Matches("PRIMARY")) { /* true! */ }

// Check against a range of dates (inclusive).
v = vtkDICOMValue(vtkDICOMVR::DA, "20050115");
if (v.Matches("200501")) { /* true! */ }
if (v.Matches("20050101-20050631")) { /* true! */ }

// Check against a wildcard pattern.
v = vtkDICOMValue(vtkDICOMVR::LO, "0Ax-T1w-SPGR");
if (v.Matches("*T1*")) { /* true! */ }

```

2.4 Accessing meta data

The meta data can be loaded from a DICOM file in two ways: either via the `vtkDICOMReader` class (which also reads the image data), or via the `vtkDICOMParser` class (which reads only the meta data). Both of these classes are designed to read a full *series* of DICOM files. To support this, the `vtkDICOMMetaData` object stores data for all files in the series, and its `GetAttributeValue()` method accepts a file index as a parameter. The index is in the range $[0, n - 1]$ for a series of n files.

```
// Get meta data from a vtkDICOMReader
reader->UpdateInformation();
vtkDICOMMetaData *meta = reader->GetMetaData();

// Check if an attribute exists, and get it for the first file.
if (meta->HasAttribute(DC::EchoTime))
{
    int fileIndex = 0;
    double t = meta->GetAttributeValue(fileIndex, DC::EchoTime).AsDouble();
}

// Another way to get an attribute for the first file (here, we assume
// that it will be the same for all files in the series).
std::string str = meta->GetAttributeValue(DC::SeriesDescription).AsString();
```

One interesting property of DICOM meta data is that it can be nested, similar to the way that directories on a file system can have subdirectories. This nesting is used, for example, to store meta data for each frame in the new enhanced multi-frame DICOM files. In order to make it easy to access nested attributes, the `vtkDICOMTagPath` class describes the full path to a nested attribute.

```
// Get an attribute for frame 3 of a multi-frame file.
int frameIdx = 3;
double echoTime = meta->GetAttributeValue(
    vtkDICOMTagPath(DC::PerFrameFunctionalGroupSequence, frameIdx,
                    DC::CardiacSynchronizationSequence, 0,
                    DC::NominalCardiacTriggerDelayTime)).AsDouble();
```

This is rather verbose, so a more convenient method for accessing per-frame data is provided for enhanced multi-frame files. You can give the frame index after the file index, in which case the `GetAttributeValue()` method will perform a search for the attribute without requiring a full path.

```
// Get an attribute for frame 3 of an enhanced multi-frame file.
int fileIdx = 0;
int frameIdx = 3;
vtkDICOMValue vw = meta->GetAttributeValue(fileIdx, frameIdx, DC::WindowWidth);
vtkDICOMValue vc = meta->GetAttributeValue(fileIdx, frameIdx, DC::WindowCenter);
if (vw.IsValid() && vc.IsValid())
{
    // set the window for the image
}
```

The `vtkDICOMMetaDataAdapter` class can also be used to access enhanced multi-frame files as if each frame was a separate file.

2.5 Iterating over data elements

The `vtkDICOMMetaData` object also provides iterator-style access to the data elements. This is useful, for instance, when you want to iterate through all of the elements in the meta data in sequential order. It is also useful if you want to check which attributes vary between files in the series.

```
// Iterate through all data elements in the meta data.
for (vtkDICOMDataElementIterator iter = meta->Begin(); iter != meta->End(); ++iter)
{
    vtkDICOMTag tag = iter->GetTag();
    std::cout << "tag: " << tag << std::endl;
    // Crucial step: check for values that vary across the series.
    if (iter->IsPerInstance())
    {
        int n = iter->GetNumberOfInstances();
        for (int i = 0; i < n; i++)
        {
            std::cout << "instance " << i << ": " << iter->GetValue(i) << std::endl;
        }
    }
    else
    {
        // Not PerInstance: value is the same for all files in series.
        std::cout << "all instances: " << iter->GetValue() << std::endl;
    }
}

// Get the iterator to a specific element (hash table lookup).
vtkDICOMDataElementIterator iter = meta->Find(DC::ImageOrientationPatient);
if (iter != meta->End())
{
    // do something
}
```

You might be surprised by the `PerInstance` check, but it is necessary due to the fact that `vtkDICOMMetaData` holds the meta data for an entire series of DICOM files. Most attributes are the same across the series, but a few attributes vary from one file to the next. These per-file attributes are identified when the file is read by `vtkDICOMReader`.

2.6 Dictionaries

When iterating through the data elements in the meta data, as described in the previous section, it can be useful to get information about the meaning of the data elements that are encountered. Complete information can only be provided by the DICOM standards documents themselves, but the `vtkDICOMDictionary` can at least provide a summary of what kind of data to expect for a given attribute.

```
// do a dictionary lookup on a tag
vtkDICOMDictEntry entry;
entry = vtkDICOMDictionary::FindDictEntry(vtkDICOMTag(0x0008,0x0020));
// check if entry was found in dictionary
if (entry.IsValid())
{
    std::cout << entry.GetName() << std::endl; // prints "StudyDate"
```

```

std::cout << entry.GetVR() << std::endl; // prints "DA"
std::cout << entry.GetVM() << std::endl; // prints "1"
}

// do a dictionary lookup by name
entry = vtkDICOMDictionary::FindDictEntry("StudyDate");
if (entry.IsValid())
{
    std::cout << entry.GetTag() << std::endl; // prints "0008,0020"
}

```

The `vtkDICOMDictionary` class provides information for attributes that are described in the DICOM standard, as well as information for private attributes defined by medical device manufacturers. Every DICOM file is likely to have a mix of standard attributes and private attributes. Fortunately, it is easy to tell the difference between the two: private attributes always use a tag with with an odd group number, while the DICOM standard only uses even group numbers. The lookup of private tags requires the name of the private dictionary.

```

// do a private dictionary lookup in GE dictionary GEMS_ACQU_01
vtkDICOMDictEntry entry;
entry = vtkDICOMDictionary::FindDictEntry("CellSpacing", "GEMS_ACQU_01");
if (entry.IsValid())
{
    std::cout << entry.GetTag() << std::endl; // prints "0019,0004"
    std::cout << entry.GetName() << std::endl; // prints "CellSpacing"
    std::cout << entry.GetVR() << std::endl; // prints "DS"
    std::cout << entry.GetVM() << std::endl; // prints "1"
}

```

2.7 Private meta data

Because private tags are not registered with any central authority, there is no guarantee that they are unique. Instead, each private group within a DICOM file contains 240 blocks (each with 256 elements) that can be individually reserved for elements belonging to a specific private dictionary. The details of how this is done are described in Part 5, Section 7.8 of the DICOM standard.

The result of this is that private tags are of the form (gggg,xxee) where “xx” is a hexadecimal value between 10 and ff that identifies the block that was used to store the attribute. The tricky thing is that this value can vary from one DICOM file to the next, though it is usually consistent within a single series. Some people are surprised by this, because the first block (i.e. 10) is the only block that is used in most files.

To ensure that you are looking for private attributes in the correct location (i.e. within the correct block), you must *resolve* each private tag before using it.

```

// start with the tag in its "dictionary" form, with xx=00
vtkDICOMTag ptag = vtkDICOMTag(0019,0004);
// resolve the private tag (find out what block was reserved)
ptag = meta->ResolvePrivateTag(ptag, "GEMS_ACQU_01");
if (ptag == vtkDICOMTag(0xffff,0xffff))
{
    // the special tag value ffff,ffff indicates that the tag could not be
    // resolved: no private block was reserved for dictionary GEMS_ACQU_01
}

```

```

    }
else
{
    // ptag will now be 0019,xx04 where "xx" is usually 10 (first block)
    double spacing = meta->GetAttributeValue(ptag).AsDouble();
}

```

2.8 Enhanced multi-frame data

DICOM is an evolving standard, and one product of that evolution was the development of enhanced multi-frame data sets that contain an entire series of images within a single file. Unfortunately, the meta data for these files is arranged in a very different manner from old-style “legacy” DICOM files. Many of the attributes of an enhanced file are nested within two sequences: the `PerFrameFunctionalGroupsSequence` (5200,9230) holds meta data that varies from frame to frame (e.g. from slice to slice), while the `SharedFunctionalGroupsSequence` (5200,9229) holds meta data that is the same for all frames.

```

// Get position for frame 3 (4th slice) of a multi-frame file.
int frameIdx = 3;
vtkDICOMValue p = meta->GetAttributeValue(
    vtkDICOMTagPath(DC::PerFrameFunctionalGroupSequence, frameIdx,
        DC::PlanePositionSequence, 0,
        DC::ImagePositionPatient));
// Get orientation for all slices (if orientation varies between slices,
// this will be in the PerFrameFunctionalGroupsSequence instead!)
vtkDICOMValue o = meta->GetAttributeValue(
    vtkDICOMTagPath(DC::SharedFunctionalGroupSequence, 0,
        DC::PlaneOrientationSequence, 0,
        DC::ImageOrientationPatient));

// That was messy. Here is a cleaner way of getting per-frame data that
// searches through the Shared and PerFrame sequences so you don't have to.
int fileIdx = 0; // only one file, but it has multiple frames
p = meta->GetAttributeValue(fileIdx, frameIdx, DC::ImagePositionPatient);
o = meta->GetAttributeValue(fileIdx, frameIdx, DC::ImageOrientationPatient);

```

Even with the “clean” frame-based meta data retrieval, you need to know ahead of time whether you are dealing with an enhanced file (and need to supply a frame index) or whether you are dealing with a legacy file (and need to supply a file index). To work around this, you can use the `vtkDICOMMetaDataAdapter` class, which makes enhanced files *look* like legacy files, and offers a simple pass-through for legacy files.

```

// A third way of dealing with enhanced data is to use the adapter class.
vtkDICOMMetaDataAdapter adapter(meta);
int n = adapter->GetNumberOfInstances(); // returns the number of frames
int i = 3; // The frame we want info for (better be less than n)
p = adapter->GetAttributeValue(i, DC::ImagePositionPatient);
o = adapter->GetAttributeValue(i, DC::ImageOrientationPatient);

```

Unfortunately, this is not the end of the story. Some tags were renamed and given slightly different meanings for the enhanced files. For example, if you need to know the echo time for an MR image slice, you should check for both `EchoTime` (the legacy tag) and `EffectiveEchoTime` (the equivalent enhanced tag).

3 Reading DICOM Files into VTK

3.1 Sorting files into a series

The first difficulty that one often runs into when trying to load DICOM files into VTK is that the files themselves do not have descriptive names. With a directory full of DICOM images it can be difficult to know which ones to load. A typical DICOM folder listing looks something like this:

IM-0001-0001.dcm	IM-0001-0005.dcm	IM-0001-0009.dcm	IM-0001-0013.dcm
IM-0001-0002.dcm	IM-0001-0006.dcm	IM-0001-0010.dcm	IM-0001-0014.dcm
IM-0001-0003.dcm	IM-0001-0007.dcm	IM-0001-0011.dcm	IM-0001-0015.dcm
IM-0001-0004.dcm	IM-0001-0008.dcm	IM-0001-0012.dcm	IM-0001-0016.dcm

These files might be 16 slices of a 3D image, or the first three files might be test images while the remaining 13 files are slices of a 3D volume. Or they might be something else entirely. So the first thing to do with a batch of DICOM files is to find out how they fit together, and the `vtkDICOMFileSorter` performs this task. Given a set of DICOM files, it will discover which files belong to the same DICOM series. Each series will generally be one of the following:

1. A stack of slices.
2. A series of sequential frames.
3. A combination of the above (multi-dimensional).

When we read DICOM images into VTK, we want to load just one series of images (or sometimes just a portion of a series) as a VTK data set. The `vtkDICOMFileSorter` makes this easy.

```
// Instantiate a DICOM sorter.
vtkSmartPointer<vtkDICOMFileSorter> sorter =
    vtkSmartPointer<vtkDICOMFileSorter>::New();

// Provide an array containing a list of filenames.
sorter->SetInputFileNames(filenames);

// Update the sorter (i.e. perform the sort).
sorter->Update();

// Get the first series.
int i = sorter->GetNumberOfSeries();
if (i > 0)
{
    vtkStringArray *sortedFiles = sorter->GetFileNamesForSeries(0);
    // do something with the files
}
```

In addition, the sorter can discover which series belong to the same study. That is, it can tell us which series were collected during the same imaging session. One thing the sorter does *not* do is sort the images in the series according to slice location. It only sorts the images according to the Instance Number embedded in each image, where the Instance Number gives the logical viewing order prescribed by the medical device that generated the images. It is up to the `vtkDICOMReader` to check the slice positions for the files and sort them by location before generating an image volume or time series.

```
// Sort the input filenames by series and study.
sorter->SetInputFileNames(filenames);
sorter->Update();

// Iterate through all of the studies that are present.
int n = sorter->GetNumberOfStudies();
for (int i = 0; i < n; i++)
{
    // Iterate through all of the series in this study.
    int j1 = sorter->GetFirstSeriesForStudy(i);
    int j2 = sorter->GetLastSeriesForStudy(i);
    for (int j = j1; j <= j2; j++)
    {
        vtkStringArray *sortedFiles = sorter->GetFileNamesForSeries(j);
        // do something with the files
    }
}
```

3.2 Reading a series of files as a VTK image

The design goal for the `vtkDICOMReader` is to convert a series of DICOM files into a geometrically accurate `vtkImageData` object. That is, the pixel spacing and the center-to-center distance between adjacent slices in the image data is as specified in the DICOM meta data. Furthermore, a 4×4 matrix is provided that can be used to position and orient the image. In order to achieve this, the reader checks the Image Position and Image Orientation that are recorded in the meta data for each slice. Then, if and only if the image positions line up along the normals of the slice planes, the reader sorts the slices according to location. The `vtkImageData` z spacing is then set to the average center-to-center distance between adjacent slices.

In the absence of Image Position information in the meta data, or if the slices do not form a rectilinear volume, then the slices are sorted only according to the Instance Number in the meta data. There is also a reader method called `SortingOff()` than can be called to disable sorting entirely, so that the order of the slices in the `vtkImageData` will reflect the order of the list of files provided to the reader.

In addition to sorting slices by location, the reader attempts to detect multi-dimensional data sets. It recognizes up to 5 dimensions: x , y , z , t , and a vector dimension. This is best illustrated by example. If an MR raw-data DICOM series provides real and imaginary pixel data at each slice location, then the `vtkImageData` produced by the reader will have two components (real and imaginary). We interpret this loosely as an image with a vector dimension of 2.

When a time dimension is present, things become interesting. The default behavior of the reader is to store adjacent time points in adjacent `vtkImageData` slices. This works well when the images are to be displayed slice-by-slice. It is, however, inappropriate if the `vtkImageData` is to be displayed as a multi-planar reformat or as a volume. For this reason, the `vtkDICOMReader` has a method called `TimeAsVectorOn()` that will cause the reader to treat each voxel as a time vector. In other words, if the DICOM data has 10 individual time slots, then the `vtkImageData` will have 10 components per voxel (or 30 components in the case of RGB data). By selecting a specific component or range of components when displaying the data, one can display a specific point in time.

Five dimensions come into play when the DICOM series has frames that are at the same location and within the same time slot. Going back to the (real,imaginary) example, if such a series of images is read after `TimeAsVectorOn()` is called, then the `vtkImageData` will have 20 components per voxel if there are 10 time

slots. The 20 components can be thought of as 10 component blocks with 2 components per block. A filter like `vtkImageExtractComponents` can be used to extract a block of components that corresponds to a particular time slot.

If the behavior described in the preceding paragraphs is not desirable, then one can use the `SetDesiredTimeIndex(int)` method to read just one time slot, and use a set of N readers to read the N time slots as N separate VTK data sets.

```
// Instantiate the reader
vtkSmartPointer<vtkDICOMReader> reader =
    vtkSmartPointer<vtkDICOMReader>::New();

// Provide a vtkStringArray containing a list of filenames.
reader->SetFileNames (filenames);

// Read the meta data via UpdateInformation()
reader->UpdateInformation();

int numberOfTimeSlots = reader->GetTimeDimension();
if (numberOfTimeSlots > 1)
{
    // for example, read only the final time slot
    reader->SetDesiredTimeIndex (numberOfTimeSlots-1);
}

// Update the reader
reader->Update();
```

3.3 Enhanced multi-frame and multi-stack files

The DICOM standard allows for multiple slices (frames) per file, or even multiple stacks of slices per file. In the case of multi-frame files, each frame is assigned a position and a time slot and the frames are sorted according to the slice sorting method described in the previous section.

In multi-stack files there are, as one might expect, more than one rectilinear (or perhaps non-rectilinear) volume. If sorting has been turned off with the `SortingOff()` method, then all the frames in the file are read sequentially into `vtkImageData` slices. If sorting is on, however, then the reader is only able to read one stack at a time. The method `SetDesiredStackID()` allows one of the stacks to be chosen by name.

```
// Instantiate the reader
vtkSmartPointer<vtkDICOMReader> reader =
    vtkSmartPointer<vtkDICOMReader>::New();

// Provide a multi-frame, multi-stack file
reader->SetFileName (filename);

// Read the meta data, get a list of stacks
reader->UpdateInformation();
vtkStringArray *stackNames = reader->GetStackIDs();

// Specify a stack, here we assume we know the name:
reader->SetDesiredStackID ("1");
```

Sometimes one will find a DICOM series that contains slices that are implicitly arranged into separate stacks, even though there is no information in the meta data that explicitly assigns each slice to a stack. In this situation, the `vtkDICOMReader` will synthesize the stack information by grouping blocks of slices together if that form rectilinear volumes, and it will name the blocks with decimal strings “0”, “1”, etc. This is particularly useful for the localizer series that is present in most MRI studies, since the localizer contains separate blocks of axial, coronal, and sagittal images.

3.4 DICOM image orientation

When people use the term “image orientation” with respect to medical images, they usually mean one or more of the items listed below:

1. The order in which the pixels and slices are stored in the computer’s memory.
2. The orientation of the image slices in a real-world coordinate system, for example the patient coordinate system as defined by the medical imaging equipment that generated the images.
3. The orientation of the patient with respect to the viewer when the images are viewed on a workstation.

The first of these, the way the data is stored in memory, should merely be an implementation detail, but unfortunately the `vtkImageViewer` class insists that the pixels must be arranged in memory such that the pixel at the bottom-left corner of the image is the pixel at the lowest address in memory. This is in conflict with DICOM, which stores the top-left pixel as the first pixel in the file. To provide compatibility with the `vtkImageViewer`, the default behavior of the `vtkDICOMReader` is to flip the image in memory while it is loading it from the file. This behavior can be turned off by calling `reader->SetMemoryRowOrderToFileNative()`.

The second and third items in the list can be referred to as the real-world orientation, and the display orientation, respectively. Neither of these can be considered an implementation detail, as both of them are crucial to the user experience. Also, it is important not to confuse one with the other. An application can handle the real-world orientation incorrectly but still display the images to the user in the correct orientation, but such an application would certainly have a serious flaw.

The real-world orientation is provided by the `GetPatientMatrix()` method of the `vtkDICOMReader`. This method returns a `vtkMatrix4x4` object that describes the coordinate transformation from the data coordinates of the `vtkImageData` that stores the image, to the real-world Patient Coordinate System defined in the DICOM standard [1]. The matrix is used to correctly place the image in the VTK world coordinate system.

The `PatientMatrix` is constructed from the `ImagePositionPatient` and `ImageOrientationPatient` attributes in the series of DICOM files that are provided to the reader. Note that unless `SetMemoryRowOrderToFileNative()` has been called on the reader, the orientation of the matrix will be flipped with respect to `ImageOrientationPatient` in order to account for the fact that the image rows were flipped in memory.

```
reader->SetMemoryRowOrderToFileNative(); // keep native row order
reader->Update(); // update the reader
vtkMatrix4x4 *matrix = reader->GetPatientMatrix();

// Create an image actor and specify the orientation.
vtkSmartPointer<vtkImageActor> actor =
    vtkSmartPointer<vtkImageActor>::New();
actor->GetMapper()->SetInputConnection(reader->GetOutputPort());
actor->SetUserMatrix(matrix);
```

Setting the actor's `UserMatrix` will ensure that the real-world orientation of the image is correctly handled, as far as the VTK display pipeline is concerned. It does not, however, set the display orientation, which is the responsibility of the application. The display orientation can be set via manipulation of the VTK camera, or via certain methods in the `vtkInteractorStyleImage` class. The details of how this is done are outside of the scope of this document.

3.5 Obtaining per-slice meta data

A direct consequence of the sorting that is performed by the `vtkDICOMReader` is that the slices in the `vtkImageData` are not guaranteed to be in the same order as the files that were given to the reader's `SetFileNames()` method. Fortunately, the reader provides a look-up table that can be used to find out which file provided which slice. Hence, given a VTK slice number, the look-up table can provide the corresponding file, and then the meta data can be inspected for that particular file. Similarly, for multi-frame DICOM files, the reader provides a look-up table that gives the DICOM frame that provided each slice. The `vtkMetaData` object provides a `GetAttributeValue()` method that takes both a file index and a frame index, along with the tag of the attribute to be inspected.

```
// Read the files and get the meta data.
reader->SetFileNames(fileNameArray);
reader->Update();
vtkDICOMMetaData *meta = reader->GetMetaData();

// Get the arrays that map slice to file and frame.
vtkIntArray *fileMap = reader->GetFileIndexArray();
vtkIntArray *frameMap = reader->GetFrameIndexArray();

// Get the file and frame for a particular slice.
int sliceIndex = 6;
int fileIndex = fileMap->GetComponent(sliceIndex, 0);
int frameIndex = frameMap->GetComponent(sliceIndex, 0);

// Get the position for that slice.
vtkDICOMValue pv = meta->GetAttributeValue(fileIdx, frameIdx,
    DC::ImagePositionPatient);
double position[3] = { 0.0, 0.0, 0.0 };
if (pv.IsValid() && pv.GetNumberOfValues() == 3)
{
    pv.GetValue(position, 3);
}
```

As a caveat, for multi-frame files, the example given above assumes that the meta data contains a per-frame `ImagePositionPatient` attribute. This is the case for multi-frame CT and MRI files, but not for multi-frame nuclear medicine files. Whenever retrieving meta data from a DICOM image, it is wise to consult the DICOM standard to see how the attributes are defined for the various modality-specific IODs (information object descriptions).

3.6 Obtaining per-component meta data

In the example in the previous section, the `GetComponent()` method was called with two arguments, but the second argument was set to zero. If the `vtkDICOMReader` assigned a vector dimension to the data, then the

the `vtkImageData` will have multiple scalar values in each voxel. For instance, the first component in each voxel may have come from a file that provided the real component of a complex-valued image, while the second component came from a file that provided the imaginary component. In this case, one would do the following to retrieve the meta data from the “imaginary” file:

```
// Read the files and get the meta data.
reader->SetFileNames(fileNameArray);
reader->Update();
vtkDICOMMetaData *metaData = reader->GetMetaData();

// Get the arrays that map slice to file and frame.
vtkIntArray *fileMap = reader->GetFileIndexArray();
vtkIntArray *frameMap = reader->GetFrameIndexArray();

// Get the file and frame for a particular slice and component.
int sliceIndex = 6;
int vectorIndex = 1; // 2nd component is the imaginary component
int fileIndex = fileMap->GetComponent(sliceIndex, vectorIndex);
int frameIndex = frameMap->GetComponent(sliceIndex, vectorIndex);

// Get an attribute from the meta data.
vtkDICOMValue v = metaData->GetAttributeValue(fileIdx, frameIdx, tag);
```

If the data has a time dimension and the reader’s `TimeAsVectorOn()` method was called, then the components of each voxel can correspond both to a specific time slot, and to a specific vector component. To make the situation even more complicated, each pixel in the DICOM files might be an RGB pixel and therefore have three components as given by the `SamplesPerPixel` attribute in the meta data.

The number of components in the `FileIndexArray` and `FrameIndexArray` is equal to the vector dimension, and if `TimeAsVectorOn()` was called, then the vector dimension will include the time dimension. The `FileIndexArray` and `FrameIndexArray` do not have components that correspond to the individual R,G,B components in RGB images, since the R, G, and B components will always have the same meta data because they always come from the same file and frame.

The following strategy is recommended for accessing per-component meta data in multi-dimensional images:

```
// Get the arrays that map slice to file and frame.
vtkIntArray *fileMap = reader->GetFileIndexArray();
vtkIntArray *frameMap = reader->GetFrameIndexArray();

// Get the image data and meta data.
vtkImageData *image = reader->GetOutput();
vtkDICOMMetaData *meta = reader->GetMetaData();

// Get the number of components in the data.
int numComponents = image->GetNumberOfScalarComponents();

// Get the full vector dimension for the DICOM data.
int vectorDimension = fileMap->GetNumberOfComponents();

// Compute the samples per pixel in original files.
int samplesPerPixel = numComponents/vectorDimension;
```

```

// Check for time dimension
int timeDimension = reader->GetTimeDimension();
if (timeDimension == 0)
{
    timeDimension = 1;
}

// Get all attributes for a specific time.
int vectorIndex = timeIndex*vectorDimension/timeDimension;
int vectorEndIndex = (timeIndex + 1)*vectorDimension/timeDimension;

for (int i = vectorIndex; i < vectorEndIndex; i++)
{
    int fileIndex = fileMap->GetComponent(sliceIndex, i);
    int frameIndex = frameMap->GetComponent(sliceIndex, i);
    vtkDICOMValue v = meta->GetAttributeValue(fileIdx, frameIdx, tag);
    // print or display the value
}

// Extract an image at the desired time slot (e.g. for display).
int componentIndex = timeIndex*vectorDimension/timeDimension*samplesPerPixel;
vtkSmartPointer<vtkImageExtractComponents> extractor =
    vtkSmartPointer<vtkImageExtractComponents>::New();
extractor->SetInputConnection(reader->GetOutputPort());
if (samplesPerPixel == 1)
{
    extractor->SetComponents(componentIndex);
}
else if (samplesPerPixel == 2) // rare/nonexistent in DICOM images
{
    extractor->SetComponents(componentIndex, componentIndex + 1);
}
else
{
    extractor->SetComponents(componentIndex, componentIndex + 1, componentIndex + 2);
}
extractor->Update();

```

3.7 Real-world values and lookup tables

The `vtkDICOMReader` always provides the *stored value* of each pixel in the image, that is, it provides the pixel value that was stored in the file. In some circumstances, these stored pixel values can be displayed directly on the screen (for 8-bit grayscale ultrasound data, for example) or can be displayed via a user-defined window-level operation (as is common for MRI). In the general case, however, the DICOM standard defines mappings of these stored values to either *real-world values* that have real physical units, or *presentation values* that are intended for radiographic display.

For computing real-world values, the `vtkDICOMReader` provides `RescaleSlope` and `RescaleIntercept` parameters that, for some image volumes, provide the desired linear mapping values to convert the stored values into real-world values. However, these parameters cannot always be used because (1) they are not available for MRI, and (2) the mapping is not guaranteed to be the same for all slices in a volume, in fact PET images almost always use a different mapping for each slice.

For these reasons, the `vtkDICOMApplyRescale` filter is the recommended method for computing the real-world values. This filter will use the `RealWorldValueMapping` attributes that are stored in the meta data for

each file, or if these attributes are not present, it will apply the `RescaleSlope` and `RescaleIntercept` for each individual slice. The output image volume will use floating-point (double or float) to store the new pixel values. If no mapping information is present in the data set, then an identity mapping is used.

Presentation values for are usually achieved by first applying the `RescaleSlope` and `RescaleIntercept`, and the applying `WindowCenter` and `WindowWidth`. This is, for example, the way that CT images should be displayed under most circumstances. In some cases, however, the DICOM file will provide a palette lookup table for mapping the values to the display. The use of the palette can either be mandatory, in the case of `PALETTE_COLOR` images, or can be optional, in the case of the supplemental palettes that are used for quantitative MR and CT imaging. The `vtkDICOMApplyPalette` filter will extract the palette from the meta data and apply it to the image. If there is no palette in the meta data, then the image data will pass through the filter unchanged.

4 Writing DICOM Files

4.1 Conforming to the standard

The DICOM standard is expansive, with approximately 1500 pages dedicated to how to put together DICOM data sets and write them as DICOM files. Some parts of conformance are simple, for example making sure that all values stored in the file match with the value representations described in Table 1. Other parts are more complex, such as ensuring that a DICOM file that contains an MR image also contains all of the meta data that is required for an MR image.

A crucial facet of conformance is that every DICOM image must have a unique identifier or UID. No two images anywhere in the world are permitted by the standard to have the same UID, unless they are in fact the same image from the same source. The default method that the `vtkDICOMWriter` uses to generate unique IDs is to append a special 128-bit random number called a UUID to the special UID prefix “2.25.” Through the use of such a large random number, the probability that exactly the same number will be generated twice is vanishingly small. For example, `2.25.1267999611695479980069765583325507254` is such a UID.

If you have registered a UID prefix for your organization, then you can use it in the `vtkDICOM` library as shown in the following example. If your prefix is too long to allow for a 128-bit suffix, then a 96-bit suffix will be used instead.

```
// Set the UID prefix for all generated DICOM files (default 2.25.)
vtkDICOMUtilities::SetUIDPrefix("1.2.XXX.YYYYY.");

// Set a UID that identifies your software to the world.
vtkDICOMUtilities::SetImplementationClassUID("1.2.XXX.YYYYY.ZZZZZZZZZZZZZZZZZZZZZ");

// Set the name and version number of your software.
vtkDICOMUtilities::SetImplementationVersionName("SOFTWARE VER X_Y_Z");
```

In addition to having a unique ID, every DICOM data set must conform to one of the classes identified in the DICOM standard. A DICOM “class” is described in the standard by an IOD (Information Object Description) that states which attributes are to be present for that class. A CT image, for example, must belong to one of the CT classes, and must contain information such as the energy of the X-rays that were used to make the image.

4.2 Writing a VTK image as DICOM

The `vtkDICOMWriter` takes a `vtkImageData` object as input, and writes a series of DICOM image files to disk. Since the required meta data for an image varies from one modality to another, the writer delegates the creation of the meta data to another class called a `vtkDICOMGenerator`. A short example of how this is done is as follows:

```
// Create a generator for MR images.
vtkSmartPointer<vtkDICOMMRGenerator> generator =
  vtkSmartPointer<vtkDICOMMRGenerator>::New();

// Create a meta data object with some desired attributes.
vtkSmartPointer<vtkDICOMMetaData> meta =
  vtkSmartPointer<vtkDICOMMetaData>::New();
meta->SetAttributeValue(DC::PatientName, "Doe^John");
meta->SetAttributeValue(DC::ScanningSequence, "GR"); // Gradient Recalled
meta->SetAttributeValue(DC::SequenceVatiant, "SP"); // Spoiled
meta->SetAttributeValue(DC::ScanOptions, "");
meta->SetAttributeValue(DC::MRAcquisitionType, "2D");

// Plug the generator and meta data into the writer.
vtkSmartPointer<vtkDICOMWriter> writer =
  vtkSmartPointer<vtkDICOMWriter>::New();
writer->SetInputConnection(lastFilter->GetOutputPort());
writer->SetMetaData(meta);
writer->SetGenerator(generator);

// Set the output filename format as a printf-style string.
writer->SetFilePattern("%s/IM-0001-%04.4d.dcm");

// Set the directory to write the files into.
writer->SetFilePrefix("/the/output/directory");

// Write the file.
writer->Write();
```

The `vtkDICOMMRGenerator` assists with conformance by generating all the data set attributes that are required by the MR IOD. It will also scan through the `vtkMetaData` object that is provided to the writer, and use any of its attributes as long as 1) they are defined in the MR IOD, and 2) they are deemed to be valid for the image that is being written. A partial list of attributes that are never taken from the input meta data is as follows:

1. SOPInstanceUID (this is always re-generated to ensure its uniqueness)
2. SeriesInstanceUID (ditto)
3. ImageType (this is set to DERIVED\SECONDARY\OTHER by default)
4. PixelSpacing (this is set from the VTK image information)
5. Rows and Columns (ditto)
6. ImagePositionPatient and ImageOrientationPatient (these are set from the PatientMatrix)

The generator always creates a new SOPInstanceUID for each file and a new SeriesInstanceUID for each series. There is no way to set these UIDs manually. The ImageType is set to DERIVED by default, because an image cannot be considered to be ORIGINAL if it was modified in any way after its original acquisition. Finally, all information related to the pixel values or the slice geometry is generated from the vtkImageData information and from the PatientMatrix.

The vtkDICOMWriter allows several parameters, including ImageType, to be set when writing the file. These are demonstrated in the following example.

```
// Plug the generator and meta data into the writer.
vtkSmartPointer<vtkDICOMWriter> writer =
    vtkSmartPointer<vtkDICOMWriter>::New();
writer->SetInputConnection(lastFilter->GetOutputPort());
writer->SetMetaData(meta);
writer->SetGenerator(generator);

// Set the output filename format as a printf-style string.
writer->SetFilePattern("%s/IM-0001-%04.4d.dcm");
// Set the directory to write the files into
writer->SetFilePrefix("/the/output/directory");

// Set the image type to Multi-planar Reformat.
// (forward slashes will be converted to backward slashes)
writer->SetImageType("DERIVED/SECONDARY/MPR");
writer->SetSeriesDescription("Sagittal Multi-planar Reformat");

// Set the 4x4 matrix that gives the position and orientation.
writer->SetPatientMatrix(patientMatrix);
```

4.3 Customizing the generators

At the present time, the vtkDICOMWriter has only three generators available: the vtkDICOMMRGenerator (for MR), the vtkDICOMCTGenerator (for CT), and the vtkDICOMSCGenerator (for Secondary Capture, e.g. screenshots). Writing a new generator class is the recommended method for adding support for a new modality to the vtkDICOM library, though that is beyond the scope of this document.

4.4 Writing a raw pixel buffer to a DICOM file

In addition to the vtkDICOMWriter, there is a class called the vtkDICOMCompiler that can write meta data and image data directly to a file without it being processed by a vtkDICOMGenerator. It can be used to efficiently perform such actions as changing the transfer syntax of the data or tweaking the meta data. By design, the vtkDICOMCompiler will take, as input, a meta data object that describes a series of images, and it will then write the files in the series one-by-one.

```
vtkSmartPointer<vtkDICOMCompiler> compiler =
    vtkSmartPointer<vtkDICOMCompiler>::New();
compiler->SetMetaData(meta);

int n = meta->GetNumberOfInstances();
for (int i = 0; i < n; i++)
{
```

```

char outputFile[256];
sprintf(outputFile, "IM-0001-%04.4d.dcm", i+1);
compiler->SetFileName(outputFile);
compiler->SetIndex(i);
compiler->WriteHeader();
compiler->WritePixelData(rawPixelBufferForFile);
}

```

5 Locating DICOM Files

5.1 Reading from a DICOM file set

When DICOM files are stored on a CD (still a common practice in many hospitals), they are stored as a *file set* that consists of the images files with obscure capitalized file names like IMG00345 plus an index file that is always named DICOMDIR. A file set consists of one or more image series from one or more patients, and the filenames usually give no indication of how the files are meant to fit together.

The `vtkDICOMDirectory` class solves this problem by reading the DICOMDIR index file and telling you which files belong to which patient or which study and series. It has an interface that is very similar to the previously-introduced `vtkDICOMFileSorter` class, but it can sort the files using only the DICOMDIR file, without having to read the individual DICOM files. This is very important if the files are stored on a slow media such as a CD.

```

vtkSmartPointer<vtkDICOMDirectory> dicomdir =
    vtkSmartPointer<vtkDICOMDirectory>::New();
dicomdir->SetDirectoryName("E:");
dicomdir->Update();

// Iterate through all of the studies that are present.
int n = dicomdir->GetNumberOfStudies();
for (int i = 0; i < n; i++)
{
    // Get information related to the patient study
    vtkDICOMItem patient = dicomdir->GetPatientRecordForStudy(i);
    vtkDICOMItem study = dicomdir->GetStudyRecord(i);
    // Iterate through all of the series in this study.
    int j1 = dicomdir->GetFirstSeriesForStudy(i);
    int j2 = dicomdir->GetLastSeriesForStudy(i);
    for (int j = j1; j <= j2; j++)
    {
        vtkDICOMItem series = dicomdir->GetSeriesRecord(j);
        vtkStringArray *sortedFiles = dicomdir->GetFileNamesForSeries(j);
        // do something with the files
        reader->SetFileNames(sortedFiles);
        reader->Update();
    }
}

```

The `PatientRecord`, `StudyRecord`, and `SeriesRecord` in the above example contain attributes for the images that were stored in the DICOMDIR index file. At the bare minimum, these will provide the `PatientName`, `PatientID`, `StudyDate`, `StudyTime`, `StudyUID`, and `SeriesUID`.

Note that the `vtkDICOMDirectory` will work even if provided with a directory that does not contain a `DICOMDIR` file. In the absence of a `DICOMDIR`, it will recursively scan the directory for all DICOM files that it contains. Also, in addition to the `SetDirectoryName()` method, there is also a `SetInputFileNames()` method that allows a specific list of files to be provided.

5.2 Searching for files that match a query

The `vtkDICOMDirectory` class has another trick up its sleeve. It can search for files that have certain attributes, for example it can search a filesystem for all scans of a specific patient. This is similar in purpose to querying a PACS system, except that you are instead providing one or more disk directories where the files might exist. These directories are searched recursively for all files that match the query.

```
// Make a list of the directories to search.
vtkSmartPointer<vtkStringArray> dicompath =
    vtkSmartPointer<vtkStringArray>::New();
dicompath->InsertNextValue("/Volumes/Images1");
dicompath->InsertNextValue("/Volumes/Images2");

// Make a list of attributes to match, using the utf-8 character set.
vtkDICOMItem query;
query.SetAttributeValue(DC::SpecificCharacterSet, "ISO_IR 192");
query.SetAttributeValue(DC::PatientName, "Doe^John");
query.SetAttributeValue(DC::StudyDate, "2012-2015");
query.SetAttributeValue(DC::Modality, "MR");
query.SetAttributeValue(DC::ImageType, "PRIMARY");
query.SetAttributeValue(DC::SeriesDescription, "*T1w*");

vtkSmartPointer<vtkDICOMDirectory> dicomdir =
    vtkSmartPointer<vtkDICOMDirectory>::New();

// The SetInputFileNames method takes directories, too!
dicomdir->SetInputFileNames(dicompath);
// Optionally restrict the search to files ending with ".dcm"
dicomdir->SetFilePattern("*.dcm");
dicomdir->SetFindQuery(query);
dicomdir->Update();
```

All text attributes except for dates accept the wildcards ‘*’ and ‘?’. For dates, you can use a hyphen to specify a range of dates. Matching of all text is done by first converting the text to utf-8 for compability, and for patient names the matching is case insensitive. The query “PRIMARY” will match the multi-valued attribute value “ORIGINAL\PRIMARY”, since only one value has to match.

If the query contains any non-ASCII text, you must set the `SpecificCharacterSet` attribute to whichever character set your program uses internally. This does not have to be the same as the character set used in the files you are searching for, because all matching is done only after converting the text to utf-8.

6 Future Work

One major feature that will be added to the `vtkDICOMWriter` is the ability to write additional kinds of DICOM images, in particular enhanced multi-frame images and parametric maps with floating-point pixel

data. Additional readers will also be added for various non-image data sets, such as spatial transformations and segmentation objects.

It is also worth noting two features that are not scheduled at this time: support for the DICOM networking protocol, and inclusion of DICOM-specific image display classes for VTK. The former is something that might be added in the future, while the latter is beyond the scope of this project and, if released, will be part of a separate project.

References

- [1] Medical Imaging & Technology Alliance. DICOM Base Standard – 2015c. Published by the Association of Electrical and Medical Imaging Equipment Manufacturers (NEMA), 2015.

Appendix A: Enabling the vtkDICOM module in VTK

Starting with VTK 7.0, vtkDICOM is provided as a remote module for VTK. This means that, if you are building VTK from source, you can have cmake automatically download and build vtkDICOM for you. This is done by setting the `Module_vtkDICOM` option for the cmake cache:

```
$ cmake -D Module_vtkDICOM:BOOL=ON ../vtk-source
$ make
```

When the “make” is complete, the vtkDICOM library will be present with all of the other VTK libraries. In addition, if you have enabled wrapping, then all of the vtkDICOM classes will be wrapped.

For versions of VTK prior to 7.0, vtkDICOM must be built as a stand-alone library according to the instructions in Appendix B.

Appendix B: Building a stand-alone vtkDICOM library

The vtk-dicom project is on github, at the following url:

```
http://github.com/dgobbi/vtk-dicom
```

For people who are not familiar with git, the easiest way to access the source code is to click on the “Download ZIP” link on the right side of the web page. The following procedure can be used to build the vtkDICOM library, assuming that you have already built VTK with CMake. This example is for building the package on Linux or OS X in the bash shell.

```
$ # specify the directory where you built VTK, this is an example
$ export VTK_DIR=/Volumes/HD2/vtk-release-build/
$ unzip vtk-dicom-master.zip
$ mkdir vtk-dicom-master-build
$ cd vtk-dicom-master-build
$ cmake -D CMAKE_BUILD_TYPE:STRING=Release ../vtk-dicom-master
$ make
```


Optionally, the following cmake variables can be set with `ccmake` or `cmake-gui`:

<code>BUILD_EXAMPLES</code>	<code>ON</code>
<code>BUILD_PROGRAMS</code>	<code>ON</code>
<code>BUILD_SHARED_LIBS</code>	<code>OFF</code>
<code>BUILD_TESTING</code>	<code>ON</code>
<code>USE_DCMTK</code>	<code>OFF</code>
<code>USE_GDCM</code>	<code>OFF</code>
<code>VTK_DIR</code>	<code>/Volumes/HD2/vtk-release-build/</code>

The `USE_DCMTK` and `USE_GDCM` variables allow you to add the image decompression capabilities of either of these packages (do not specify both!) to the `vtkDICOMReader`. If you do not specify one of these packages, then the `vtkDICOMReader` will only be able to read uncompressed files.

If you want to use DCMTK, then download `dcmtk-3.6.1` from <http://www.dcmkt.org/> and build it using the the instructions provided on the website. If possible, use the same compiler that you use for `vtkDICOM`.

If you want to use GDCM, it can be found at <http://sourceforge.net/projects/gdcm/>. Pay particular attention to the GDCM build instructions, especially the “make install” step, or the `vtkDICOM` library will not be able to link to the GDCM libraries. Additional information on GDCM can be found at http://gdcm.sourceforge.net/wiki/index.php/Main_Page.

The easiest way to use the `vtkDICOM` library in your own project is to add the following command block to the main `CMakeLists.txt` file in your project:

```
find_package(DICOM QUIET)
if(DICOM_FOUND)
    include(${DICOM_USE_FILE})
endif()
set(VTK_DICOM_LIBRARIES vtkDICOM)
```

It is not recommended to try to use `vtkDICOM` (or VTK itself) within projects that are not built with `cmake`.