

7 Things to Consider When Buying a Computer for CT Image Processing



Kyle Stamps 

I don’t want to date myself, but I started playing with computers in the ‘80s, using “floppy disks” that were actually floppy back then. I have worked as an IT specialist for Rigaku for fifteen years, configuring computers for different X-ray analytical instruments and individual customers’ needs. So selecting the right computer is my specialty.

Image analysis for X-ray micro-CT ([computed tomography](#)) is one of the most demanding computational tasks people perform for scientific research. Selecting the right computer is important, but it also can be challenging.

But don’t worry. I am here to help. There are seven things to consider when buying a computer for CT image processing:

1. Which image analysis software to use

2. Range of the image size you need to process
3. Processor (CPU)
4. System memory (RAM)
5. Graphics card (GPU)
6. Data storage
7. Price

With items 1 and 2 in mind, you can select the CPU, RAM, GPU, and data storage. First, I will discuss why the analysis software and image size matter and then dive into each component. I will give you an idea of the price range at the end.

Why can't I use a regular computer for CT image processing?

How can I choose the right computer?

Processor (CPU)

System memory (RAM)

Graphics card (GPU)

Data storage

Price

What should I do now?

Why can't I use a regular computer for CT image processing?

Well, you could. But average mid-level computers, say Dell OptiPlex, are often insufficient for processing GB-size image files.

There are two main reasons: First, you need a relatively large system memory, commonly referred to as RAM (random access memory), at least 32 GB; in most cases 128 GB. But most mid-level computers' memory slots are limited to 16 GB or 32 GB.

Second, most analysis software requires a graphics card or GPU (graphics processing unit) dedicated to the CT data analysis and separate from the onboard GPU. The computer needs a power supply with adequate capacity and the correct connectors for a GPU supported explicitly by the analysis software you plan to use. But again, many mid-level computers don't support these GPUs.

So, you don't need a supercomputer, but you need something better than average.

How can I choose the right computer?



Recommended steps

The “right” computer configuration depends on each use case. Before you start, you need to know:

1. Which image analysis software to use
2. Range of the image size you need to process

Each software has its own system requirements. Most software runs on 64-bit Windows 10, and some also run on Linux. The system memory size and sometimes the GPU memory size depend on the image file size to process. The RAM needs to be 2X, 4X, etc., depending on the analysis software’s requirements.

Once you have these two pieces of information, you can start configuring a computer. It might be tempting to find one person with some level of expertise and go with their recommendation, but I would recommend you refer to multiple sources and go through the following steps to be safe:

1. Check the analysis software’s recommended system requirements. (Never go with the minimum requirements.)
2. Give these requirements to a computer manufacturer representative and have them configure a computer.
3. Give that computer configuration to the software manufacturer representative to get their stamp of approval.

Here are links to recommended system requirements pages of major desktop applications for CT image analysis:

- [Dragonfly](#)
- [VGSTUDIO](#)
- [AVIZO](#)
- [GeoDict](#)

- [3D Slicer](#)

Using an existing computer

You might also consider using an existing computer instead of buying a new one. I would not recommend using a super-old computer, because it is more likely that you will face compatibility issues. But you can safely go back five to six years.

If you are considering an existing computer, compare its configuration against the analysis software's requirements. Make sure that all the drivers are up to date.

If there are things you need to upgrade, check if these items are compatible with your computer before buying them. Computers have a limited number of slots for hard disk drives, a limited maximum RAM size, and compatibility limitations for GPUs due to the power supply or the connector types. We will look at more details on these individual items later.

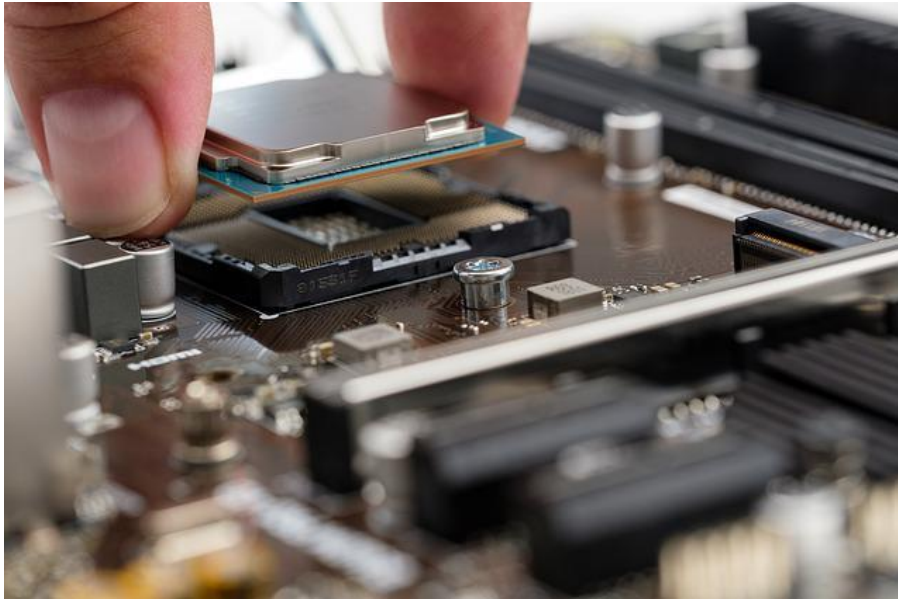
Using a laptop computer

One of the questions people often ask me is whether or not they can use a laptop for CT image analysis. You can probably get a laptop with a decent GPU and 32 GB of RAM, sufficient to analyze up to a few GB size of images. However, RAM expandability and heat generation are always a concern. We will talk more about the latter in a bit.

Workstation laptops can take more RAM, but they are more expensive than a tower computer for the same RAM capacity. There are also additional downsides. For the most part, a laptop doesn't provide performance equal to a tower workstation because it always has a mobile variant or power constraint version of CPU and GPU to compensate for the limited heat dissipation capacity.

While a tower workstation comes with about a 1000 W power supply, a laptop comes with only 180 W at max. This power shortage negatively affects the CPU speed and limits the available GPU choices. A GPU in a full-power tower workstation alone easily pulls 200 - 300 W. You can see how limited a 180 W laptop can be.

Processor (CPU)



Let's take a closer look at the processors, commonly referred to as [CPU \(central processing unit\)](#). A CPU does most of the calculations as its name suggests. You can see it as the thinking brain of a computer. The clock speed and core count are the main factors that affect the speed of a CPU. Most modern software takes advantage of multi-cores, and its speed scales with the number. So, more cores mean faster, technically. I say "technically" because this doesn't necessarily mean you want to go for a 48-core CPU.

When you add more cores, the processors heat up, and the clock speed needs to go down to prevent overheating. You might think a 48-core CPU is great, but each core might need to be slow for this CPU to operate without overheating. CPUs automatically monitor the temperature and power draw and adjust the clock speed as necessary. This balancing act is called CPU throttling.

So, if CPUs are running at an adjusted low clock speed, what does a boost rate of a CPU mean? A CPU's boost rate is the maximum achievable clock speed when heat dissipation is not an issue. In reality, a CPU can run at its maximum boost rate for a while, but it needs to slow down as soon as processors heat up.

You need lots of power and good cooling for a CPU to perform well. As I mentioned above, this is a major reason why a laptop's performance is limited, however much you pay for it.

The takeaway is this: Increasing the number of cores requires sacrificing single-core performance. You don't get much benefit beyond eight cores with today's technology. In other words, eight is the sweet spot.

System memory (RAM)



[RAM](#) is working memory and provides analysis software a temporary space to store and access data quickly. CT images are often large. Let's say we open a 5 GB image file and apply a denoising filter to it. This simple operation alone requires at least 10 GB, double the original file size. To run more complex operations, you need a space four or five times larger than the original file size.

What happens if the software runs out of space? Some programs check the available memory space and warn users that they are running out of space. Then you can close other software applications to free up space or close image files you are not processing right now. If you still don't have enough memory space, you might need to consider reducing the file size or increasing the RAM size.

If there is no warning mechanism in the application, the operating system will automatically start allocating space on the hard drive when running out of RAM space. This usually makes a process so slow that it often looks as if the application is dead and not responding. You want to avoid this by paying attention to your RAM and image file sizes.

Watch a [Mini Tutorial about changing file sizes](#).

You can see why a large RAM is often required for CT image processing. While it is nice to have a large RAM space, it can be expensive, and all computers have a limit of maximum RAM size. The 64

GB and 128 GB memories are often selected for CT image processing. If you don't need to analyze large files and can get by with 32 GB, that is a good size, providing the best value for the price.

The RAM capacity of most consumer computers, including gaming ones, is limited to 128 GB. To increase it to 256 GB or 512 GB, you would need to buy a more expensive computer and more expensive RAM.

RAM can be expensive. However, we are talking about analyzing CT images for scientific research. So you might not have a choice to compromise the integrity of the data or accuracy of the analysis to work with a less expensive computer. In that case, assess the size of images you need to process. Then check the RAM requirements of the analysis software you want to use and go with that size.

Graphics card (GPU)



A [GPU](#) does a lot of calculations like a CPU, but it has many smaller and more specialized cores for image processing. CT reconstruction, 3D rendering, and deep learning calculations can significantly benefit from GPUs.

The way analysis software talks to a GPU is different from software to software. So it is crucial to select a specific GPU supported by the software manufacturer. Some software products also specify

the memory size of the GPU in relation to the image file size. So, always check the manufacturer's recommendations.

If the GPU memory size is not specified, 8 GB is a good size to aim for. When multiple GPUs are supported, go with [NVIDIA](#). NVIDIA has been around for a longer time than other GPU manufacturers, and I've had better luck with their products than others.

There are also gaming GPUs and workstation GPUs. They are very similar. The only difference is gaming cards have a limiting floating point for calculation while workstation cards don't. Gaming cards are good enough for CT image analysis as far as the specific model is supported by the software manufacturer.

Not all computers work with all GPUs. GPUs have specific connectors and draw lots of power. So your computer needs to have the right connection and have enough power supply to support the GPU.

You might not pay much attention to the power supply specification when selecting a computer, but it matters in this case. Most computers are shipped with a minimum required power supply, and you can't simply change it to accommodate a GPU.

Software is particular about which GPU it works with because there is no universal architecture for how software talks to GPUs. For CPUs, a set of rules called x86 architecture was established in the 70s, so you don't need to be picky about CPUs. That is not the case for GPUs. There has been an effort to establish a universal architecture. OpenGL (Silicon Graphics), OpenCL (Apple), and DirectX (Microsoft) are attempting to do this, to name a few.

Meanwhile, NVIDIA GPU design is significantly different from AMD GPU. Thus, NVIDIA GPUs' "native language" is CUDA, and it makes more sense to talk to them in CUDA. Some software supports [OpenGL](#), theoretically eliminating the compatibility issue between software and GPU. However, I still have some concerns and wouldn't use a GPU that the software manufacturer does not explicitly support. Things can get lost in translation.

I said the same thing about RAM, but it is more important for a GPU and worth repeating. Select a GPU supported by the software manufacturer and your computer. You might need to choose a

different computer depending on the GPU you want. If you need to use one not listed as a supported GPU for whatever reason, consult with the software manufacturer.

Data storage



Data storage is another essential component to address because CT images and their analysis result files tend to be large. You can quickly generate 100 GB of data from one experiment. I will leave the general recommendations about CT data storage to another [blog article by Dr. Shawn Zhang from digiM Solutions](#). I will talk about good practices regarding hard drives here.

Data management strategy

After using a CT scanner for a couple of years, you will most likely run out of space on your data drive. You don't want to panic when that happens. The good practice is to plan ahead and categorize data into "still working on" and "ready to archive" groups.

You can keep the "still working on" data on the data drive of your image processing computer to access it easily. You can move the "ready to archive" data to long-term storage to free up space on the data drive.

We will look at good practices for managing these two data groups later, but let's talk about the boot and data drives and the RAID formats first.

Most modern computers come with the boot drive (C drive) and the data drive (D drive). I would recommend an SSD (solid-state drive) for the boot drive for speed, and a mechanical drive for the data drive for large storage size at a low cost. It is common to use 1-2 TB and 8-12 TB drives for the boot and data drives, respectively

The boot drive is relatively static, meaning that you don't make changes to it often. It is also more reliable if it is an SSD. So, a good practice is to create an image backup regularly so that you can recover it in case of failure. If you are also using the computer to control a CT scanner, make sure to update the image when you run any hardware calibrations for it to save the latest settings.

You make changes to the data drive more often, so I recommend using a RAID format for redundancy to protect your data from a disk failure.

We all know making a backup is important. The phrase "[Two is One and One is None](#)" applies well to data. Now, RAID is not a backup. This is a common misconception.

The redundancy some RAID formats provide protects your data only against disk failure. For example, if you accidentally delete or overwrite a file, that change is mirrored on the duplicate disk, so it's gone. You still need a backup even when you apply a RAID format for redundancy.

RAID format

What is a RAID format? RAID stands for Redundant Array of Inexpensive Disks. In other words, RAID is a way to put multiple inexpensive disks together to achieve the speed or reliability of a more expensive disk such as SSDs.

Achievable speed and reliability depend on the type of RAID.

- **RAID 0 (striping)**

RAID 0 requires two drives and is used for speed. It breaks data into two blocks (called "stripes") and writes one block on one drive and the other block on the other drive. RAID 0 can write approximately twice as fast. However, if one drive fails, you lose everything.

This format does not protect your data against a disk failure. RAID 0 became somewhat obsolete after SSDs became available at a reasonable price.

- **RAID 1 (mirroring)**

RAID 1 requires two drives and is used for reliability. It writes the same data on two drives at the same time. The second drive is “mirroring” the first drive. It is commonly used because it is a cheap way to provide redundancy and protect data against disk failure.

RAID 1 is slightly slower than no RAID because the operating system has to do some extra calculations, but the difference is negligible when writing large files (~5%). The difference becomes more significant when writing many small files (~30%).

It means that there is a read/write speed difference between single-page TIFF and multi-page TIFF files, even when their total sizes are the same. The TIFF file size is limited to 4 GB. So for a TIFF stack under 4 GB, there is a speed benefit of saving them as a multi-page TIFF file.

- **RAID 5/6 (parity set)**

RAID 5 and RAID 6 require three and four disks, respectively. These formats are used for reliability. One drive is called a parity drive and is used for a parity set. So you get a three-drive worth of space from four drives.

Parity is a way to detect errors in a storage system by adding checksums. If one of the drives fails, the parity drive can recover the data on it. RAID 5 and RAID 6 can recover from one and two drive failures, respectively.

The read speed is good, but the write speed goes down because the CPU needs to perform a parity calculation unless you have dedicated hardware.

RAID 5/6 is good for more permanent storage and archiving when you have a lot of files but don't care about the speed.

- **RAID 10 (mirroring + striping)**

RAID 10 requires four drives and is a faster and bigger version of RAID 1. RAID 10 writes on two drives as RAID 0 does for speed, and the entire set of stripes is mirrored on two additional drives for redundancy.

The speed is an obvious benefit. Treating a large space (double the single drive size) as one drive also makes file management easier. The downside is that you need four drives.

Now you know what RAID formats are available, let's think about good practices for "still working on" and "ready to archive" data.

Data storage configuration for everyday use - "Still working on" data

If you don't need a large space and can spend extra money, using RAID 1 formatted SSDs would give you the best experience. If you need a large space, more than 4 TB, and can't afford a larger SSD, using RAID 1 or RAID 10 formatted mechanical drives is a good practice.

Long term storage - "Ready to archive" data

Let me just say this first: Hard drives fail. So copying all data from your data drive to an external hard drive and calling it a day is not a good practice. I would recommend either cloud storage for easy or no maintenance or RAID 5 formatted drives on the network for easy access. If you are going to use external hard drives, at least make two copies.

Note that a cloud is not some sort of magical storage space that never fails. Cloud storage is just someone else's computer. It's just that you pay someone else to take care of formatting, redundancy, maintenance, etc. However, when you use a [cloud storage service specialized for CT data](#), there are other benefits such as easy archiving.

Before I close this section, let's look at the Z file system or ZFS. ZFS is similar in concept to RAID in the sense that you use multiple disks to make a single array. The difference is the way it stores data.

There are two ways a drive fails: It mechanically fails, or some bits flip from 0 to 1 or vice versa. In addition to protecting data from mechanical failures by creating redundancy, ZFS scrubs the data to identify and correct flipped bits by putting parity amongst the data. Reading/writing can be faster despite the additional calculation required for scrubbing because it uses multiple disks and a dedicated computer.

If Starship Enterprise in *Star Trek* were to choose a file system, I think they would choose ZFS. It is that secure, yet it is not commonly used because it is not simple to implement. ZFS requires a full-

blown computer, which can be part of the archive rack mount computer with multiple drives on the network. You probably need an IT person to set it up, although some [prepackaged solutions](#) are available.

Price

How much does a computer for CT image processing cost? A workstation equipped with sufficient CPU, RAM, GPU, and data storage for CT image processing costs about \$5K. For example, it will get you a 10-core CPU, 128 GB of RAM, and NVIDIA GPU with 8 GB of memory.

You can certainly configure your computer for higher performance. If you upgrade it to the Zeon Gold and 512 GB of RAM, that would be over \$10K.

You might be on a budget, but you probably don't want to go with a super-cheap computer. Anything under \$3K would make me concerned.

What should I do now?

When you are ready to pick a computer, determine your budget, maximum image file size, and the analysis software you want to use. Then follow the three steps I described at the beginning to find the right computer for you.

1. Check the analysis software's recommended system requirements. (Never go with the minimum requirements.)
2. Give these requirements to a computer manufacturer representative and have them configure a computer.
3. Give that computer configuration to the software manufacturer representative to get their stamp of approval.

You might plan to purchase a CT scanner and use the computer that comes with the scanner for data analysis. In that case, give the file size and analysis software requirements to the CT manufacturer and ask them to provide a computer that satisfies the software manufacturer's requirements.

If you are not sure where to start, our team of CT and IT experts can help you choose the right computer for your research needs. You can talk to one of our CT experts by clicking the "Talk to an expert" button at the right top of the page or send us a message at imaging@rigaku.com.



[Kyle Stamps](#) 

My name is Kyle and I'm Digital Marketing Developer residing in Houston, TX. I have 15 years of development experience with JavaScript, CSS, and PHP with a focus on HubSpot, APIs, micro-services, and data management. In my spare time I enjoy traveling, cooking and good coffee.