IBO

IB DIPLOMA PROGRAMME
PROGRAMME DU DIPLÔME DU BI
PROGRAMA DEL DIPLOMA DEL BI

COMPUTER SCIENCE
CASE STUDY: COMPUTER TECHNOLOGY AND HUMAN EVOLUTION RESEARCH

For use in May 2002 and November 2002

INFORMATIQUE
ÉTUDE DE CAS : INFORMATIQUE ET RECHERCHE SUR L'ÉVOLUTION DE L'HUMANITÉ

Pour utilisation lors des sessions de mai 2002 et novembre 2002

INFORMÁTICA
ESTUDIO DE UN CASO: TECNOLOGÍA INFORMÁTICA E INVESTIGACIÓN DE LA EVOLUCIÓN HUMANA

Para uso en los convocatorias de mayo de 2002 y noviembre de 2002

IB DIPLOMA PROGRAMME
PROGRAMME DU DIPLÔME DU BI
PROGRAMA DEL DIPLOMA DEL BI

COMPUTER SCIENCE
CASE STUDY: COMPUTER TECHNOLOGY AND HUMAN EVOLUTION RESEARCH

For use in May 2002 and November 2002

INSTRUCTIONS TO CANDIDATES

- Case study booklet required for higher level paper 2 and standard level paper 2 computer science examinations.
- This case study booklet consists of 4 sections:
  1. Text;
  2. The Glossary;
  3. Quote from a research paper;
  4. Figures 1-6 with legends.

14 pages

1. Text

A research group at a university institute studying human evolution is interested in researching fossils by digitising them and analysing the resultant files with the help of sophisticated computer software. (This special field of research has acquired the name “Virtual Anthropology” — VA for short — in order to contrast it from conventional physical anthropology, where measurements, *etc.*, are made using the original fossil.) There are several advantages to using the VA approach. The most obvious one: it enables researchers to investigate (“see”) the inside of a fossil. In conventional physical anthropology, one cannot see the inside of a fossil without breaking it. Human fossils are very rare and very precious; they may therefore never be damaged during any investigations. The researchers investigating and/or analysing human and human predecessor fossils are called palaeoanthropologists.

The palaeoanthropologists begin their analyses of fossils by first producing CT-scans of them using a medical CT scanner (Figure 3 and Figure 4). (CT stands for “computed tomography”.) First, the fossil is put into a box which is then placed on a sliding bed (Figure 5) that can be moved through a circular aperture surrounded by a ring assembly — the CT scanner (Figure 4). The ring assembly consists of an x-ray emitting apparatus on a rotating disk and a circular ring of detectors (Figure 4). The x-ray emitting apparatus revolves around the fossil on the bed; after one complete revolution of the assembly, the fossil is slid forward by a fixed, specific distance (often 1 mm), and the x-ray emitting apparatus is again revolved around the specimen. The distance which the fossil is moved forward between revolutions is called a slice (or “slice distance”). This process is repeated until the fossil has been completely scanned.

The x-rays are attenuated by the fossil material, while the air surrounding the fossil does not (within the sensitivity limits of the detectors) and the material of the box attenuates the rays hardly at all. The file (called the “absorption file”) of a single slice is a one-dimensional array that consists of the absorption strengths — one for every beam angle. A complete scan typically consists of 200–400 slices. The research group, however, is interested in knowing the absorption strength at every point of the fossil. This is not possible, since there are an infinite number of points. Instead, each complete scan is divided into small three-dimensional regions. These regions are called “voxels”, short for “volume pixels”. Each voxel has a specific absorption value, measured in Hounsfield units. In a thresholding step, the background is set to 0 HU (Hounsfield units), so the fossil is represented as a three-dimensional array of non zero, positive values, called CT numbers. The absorption file consists only of directions, so various numerical transformation algorithms must be used to compute the CT number at each voxel. All researchers in the field of VA deplore the fact that the best algorithms require many very numerically intensive computations. The (computed) voxel file is conventionally called a “CT-scan”.

By the way, the same method (with slightly different physical parameters for the x-ray tubes, *etc.*) is used to diagnose patients, for example those suffering from brain tumours. In the case of a brain tumour, a radiologist will look at the slice images on a computer visual display unit and can identify the regions of the tumour by the difference in its CT number ranges relative to the surrounding (healthy) brain tissue CT number range.

In fossil research, a typical CT-scan will usually be several hundred megabytes large, as for each voxel the CT number is stored in 12-bit resolution (grey scale) and a typical scan has 512×512 voxels per slice and about 200–400 slices. Almost always, for 12-bit resolution, each CT number is stored as 2 bytes.

For further research analysis, this file is transferred from the scanner set-up (Figure 5), located in the radiology department of a hospital, via a national communication link to the university institute equipped with several workstations. For transmission, the CT-scan is segmented into packets; each packet is sent separately, together with parity bits and check sums (to reduce the possibility of transmission errors). Also, an elaborate handshaking protocol is in place to ensure complete delivery. The file formats of the various stages of the CT-scan computation are different and quite a bit of computer processing power is necessary for the conversion from one format to another.

The researchers at the institute are in competition with other research groups and are thus concerned that unauthorised parties might gain access to this file. They are also worried that some hacker might attempt to corrupt these valuable files by entering the system out of sheer maliciousness. For these reasons, the institute has two LANs separated by a "firewall" (a specialised software package that prevents access by unauthorised persons). One LAN (the "outer one") is connected to the WAN of the university computing centre and also to the Internet. The other LAN (the "inner one") interconnects two workstations (equipped with RISC processors) and three high-performance desktop computers (1000 MHz CPUs addressing 512 MB RAM each) via a hub server, whose firewall not only limits access to the computers, tape drive back-up systems (streamers), twenty 10 GB hard-disk storage drives and several 20 GB ones, but also monitors unauthorised break-in attempts.

The workstations are used by the researchers to manipulate, analyse, edit or otherwise re-compute the CT-scans. The resultant files must be stored and/or transferred to the high performance desktop machines where graphics are produced for publications or storage on CD-ROMs. Major computer-intensive tasks are: filtering, segmentation, feature extraction and re-assembly (see The Glossary for definitions of these terms). Surface rendering and advanced statistical analysis is done on the high performance desktop machines. The filtering tasks are processed in batch mode which can be interrupted by a request for urgent landmark feature identification tasks. A batch job interrupted in this manner will be placed on a stack, the intermediate state of the volatile core memory is stored as a swap file on a large hard disk drive reserved specifically for this purpose.

Research results and processed images are transferred "over the firewall" to other computers at the institute for inclusion in manuscripts, posting on a web site or making displays/posters for presentations at scientific meetings.

Large, animated image files must also be produced, because high-quality surface rendering alone is often inadequate for visual presentation of these complex fossil structures (Figure 1). An animation sequence consists of slowly rotating the virtual fossil, re-computing the surface rendering for each frame and compressing the output image frame. The sequence of (compressed) images, together with the decompression algorithm, is called an animated file. The animated files are burnt onto CD-ROMs, which are then sold via a commercial, highly reputed scientific publisher to other research universities. The CD-ROMs are copyrighted with the museum that has custody over the fossil.

What are the advantages of CT-scanning fossils and analysing them in digitised form? There are several: (1) The most obvious advantage is the possibility to investigate hidden features without damaging the (priceless) original. (2) A second advantage is that researchers can manipulate the fossil in a reversible way: if the result of some process is deemed unsuccessful (or in contradiction to data derived from other fossil material), the file storing the result can simply be deleted. Indeed, one can make numerous attempts — with slight variations — at re-assembly or reconstruction, for example. Such repetitious approaches are impossible using an original. (3) A third advantage is numerical analysis. The volume of any cavity in the fossil can be found by simply counting the number of voxels inside the cavity and multiplying their sum by the volume of one voxel. With this method, VA researchers have found the brain volumes of many fossil crania. (4) Also, one can locate the 3-dimensional position of any point on or in a fossil. The geometry of human and human predecessor fossils can then be analysed with sophisticated statistics (called "geometric morphometrics") in order to find evolutionary relationships in the human fossil record.

2. The Glossary

attenuation:	The reduction in the intensity of a beam due to some material intercepting it. Attenuation occurs both for particle beams (such as electron rays) and for electromagnetic radiation beams (consisting of photons — the particles of light). Note that x-ray beams are a stream of photons, albeit not visible ones. It is technically incorrect to speak of the absorption of a beam: only some of the beam particles are absorbed during its passage through matter.
cranium:	The skull without the lower jaw (mandible). Plural: *crania*.
encrustation:	Bones lie buried in the ground and fossilise there. Many times, the surrounding material (*e.g.*, sand or some other sediment) fossilises together with the bone. The material surrounding the fossil is called encrustation. When a fossil is recovered from its resting place in a sediment (its "matrix") by its discoverer, remnants of the encrustation usually adheres to it. If a fossil specimen is large, cavities are invariably filled with encrustation. (In the latter case, physical removal is impossible.) Removal of encrustation is a difficult task (see *segmentation*).
feature extraction:	The application of an algorithm which isolates parts of a CT-scan with characteristic (geometric) properties. For example, one tries to find rims around openings by having the algorithm identify a continuous, closed sharp edge. Another feature often looked for is a tip (the extreme point on the cusp of a tooth, for example).
filtering:	The application of some algorithm to separate two or more features in a CT-scan. For example, one could separate regions of a CT-scan with strongly fluctuating CT numbers from those where the CT numbers change only gradually by applying a suitable filter algorithm.
fossil:	An artefact (usually found in geological sediments) that is the result of changing organic material into minerals. The most common fossils are former bones and teeth. However, there also exist fossils which are mineral replacement of the brain tissue (so-called "endocasts"). Occasionally, even trace fossils such as footprints fossilise (see, for example, "the Laetoli footprints" in a popular text, such as *Scientific American*). Fossils of humans and human predecessors (the "Australopithecines") are among the rarest fossils. Complete fossilised skeletons of humans and human predecessors have not yet been found (the famous "Lucy" fossil skeleton is ~60 % complete). In fact, there are only some twenty fossils of almost complete crania. The process of transforming a bone into a fossil is called fossilisation. *Fossilisation* is a very slow process; it takes about 25 thousand years for a bone to fossilise. (Thus there are no 10 thousand-year-old fossils, and no 100 thousand-year-old bones.)

Hounsfield:	One of the inventors of the CT-scanning process, for which he was awarded the Nobel Prize in Physiology and Medicine in 1979. The values assigned to each voxel (*i.e.*, the voxel values) on a suitable scale are called CT numbers. The CT numbers are conventionally measured in HU (Hounsfield units). For calibration purposes, air is assigned –1000 HU and water 0 HU; bone is then ~2000 HU, and fossil material can be very high, even above 4000 HU.
interpolation:	Usually, a fossil specimen is found with parts missing (lost during the fossilisation process). Sometimes, one attempts to approximate the geometry of the missing part by designing an interpolation algorithm in an attempt to fill the gap between the parts of the fossil. A different interpolation algorithm is used when a researcher wants to find a (digital) surface that has a finer scale (*i.e.*, smaller voxels) than the original CT-scan.
iso-voxels:	Iso-voxels are cubic; *i.e.*, the dimensions in all three directions are the same. Almost always, as the result of scanning conditions, the dimensions (within one slice) of a pixel (the base of a voxel) are smaller than the thickness of a slice (Figure 6). The voxels are therefore prisms with a height (the thickness of the slice) greater than the dimensions in the slice plane. The voxels in a CT-scan supplied by the scanner are therefore not iso-voxels. Researchers apply interpolation algorithms in order to create iso-voxels. A file of a digitised fossil must consist of iso-voxels before a rendering algorithm can be applied; otherwise the calculated image of the surface will appear distorted.
modelling:	The process of designing mathematical models that allow the researcher to predict some aspect of the fossil — either its appearance, its physical properties (such as strength), its geometric properties, or its relation (either statistical or evolutionary) with other fossils.
re-assembly:	Even almost complete fossils are never found in one piece. A major task of a physical anthropologist is to attempt the re-assembly of the fossil. Presently, various research teams are trying to design re-assembly algorithms to apply to digitised fossils. The advantage of a re-assembly algorithm in VA is its *reversibility*: if, for some reason, the re-assembly is deemed unacceptable, one can simply discard the file. Physical re-assembly is irreversible (unless one is willing to damage the fossil pieces when breaking the reassembled object). Re-assembly is different from *reconstruction*. In reconstruction, one attempts to guess or estimate missing pieces.

segmentation:	The process of separating two different materials (*e.g.*, a fossil from its encrustation). In a VA environment, segmentation is done by applying segmentation algorithms, which are very difficult to design and are computationally very demanding. If the segmentation is done in the laboratory on the real specimen by a laboratory technician, the segmentation process is often called *preparation*. "Cleaning a fossil" is no longer an acceptable phrase, as this term implies that the removed material is dirt — which it is not, from a researcher's point of view. The encrustation is valuable, as it contains the physical and chemical clues that permit determining a fossil's antiquity and identifying the (paleo) environment in which the fossil individual once lived.
surface rendering:	The voxels that are on the boundary between voxels with non zero CT numbers and those with zero CT numbers in a CT-scan represent the surface of a fossil in digital form — *i.e.*, a digital representation of a surface. Visualising these digital representations of surfaces (for projection on a computer screen, for example) is done with surface rendering algorithms.
thresholding:	The process of assigning all CT numbers less (or more) than some specified non zero number the value zero. For example, if the voxels of a digitised fossil have CT numbers in the range 1000 HU to 3500 HU, and the encrustation voxels always have CT numbers greater than 3800 HU, one can segment the digitised fossil from its encrustation by assigning to all voxels with CT numbers greater than 3700 HU the value zero. Thresholding is the simplest segmentation algorithm.

3. Quote from a research paper

Since the discovery of the Tyrolean Iceman in 1991, advanced imaging and post-processing techniques have been successfully applied to anthropological research. Among the specific techniques are spiral computed tomography and 3-dimensional reconstructions, which include stereolithographic and fused deposition modelling of volume data sets. The Iceman's skull was the first to be produced using stereolithography; subsequently, it has been successfully applied to preoperative planning. With the advent of high-end performance graphics workstations and biomedical image processing software packages, 3-dimensional reconstructions have become established as routine tools for analysing volume data sets. These techniques enabled dramatically new insights to be gained in the field of physical anthropology. Computed tomography became the ideal research tools to access the internal structures of various precious fossils without even touching — let alone damaging — them. Among the most precious are specimens from the genus *Australopithecus* (1.8 Myr–3.5 Myr old), as well as representatives of *Homo heidelbergensis* (200 kyr–600 kyr old) and *Homo neanderthalensis* (40 kyr–100 kyr old); such fossils have been CT-scanned during the last five years. The fossils are filled with a stone matrix or other encrustations. During the post-processing routines, highly advanced algorithms were used to remove these encrustations virtually (the concrete fossils remain untouched). Thus it has been possible to visualise the morphological structures that are hidden by the matrix layer. Some specimens have been partially destroyed, but it has been possible for the missing parts to be reconstructed on the computer screen in order to get estimations of brain volume and endocranial morphology, both major fields of interest in physical anthropology. Moreover, the data in computerised form allows new descriptions of morphological structures using geometric morphometrics. Some of the results may change aspects and interpretations in human evolution and approaches to long-standing questions in this field. We subsume the introduction of these new imaging and post-processing techniques into a new field of research: Virtual Anthropology.

From: *New Methods and Techniques in Anthropology*
W. Recheis, G.W. Weber, K. Schäfer, H. Prossinger, R. Knapp, H. Seidler, D. zur Nedden
Collegium Antropologicum, **23**, 495–509 (1999).

4. Figures 1-6 with legends

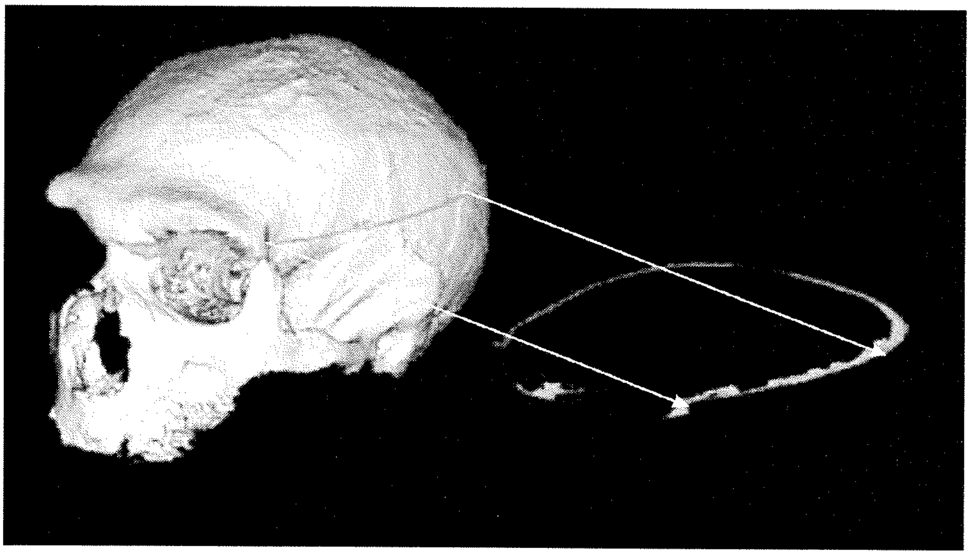

Figure 1

A CT-scan of a Neanderthal fossil cranium (a cranium is the skull without the lower jaw). This cranium is almost complete — making this specimen very rare. The image of one slice of the scan has been translated towards the right (*i.e.*, away from the left-hand side of the cranium). The missing slice can be seen as a fine gap in the rendering of the CT-scan of the cranium. The variation in brightness visible in the *slice* is due to the variation in the attenuation of the x-rays due to the density variation in the fossil material. The brightest regions visible in the slice are *encrustations* (see The Glossary). One task a palaeoanthropologist has to perform is the algorithmic removal of such encrustations in the CT-scan before other analyses can be performed. In this particular CT-scan, most of the encrustation can be removed by *thresholding* (see Figure 2).

The brightness variation (which allows us to perceive the surface as three-dimensional) in the total view on the left is due to a *surface rendering algorithm*. The surface is an interpolation of the surface voxels (which must be iso-voxels in this scan for the surface rendering algorithm to function properly).

Each slice in this CT-scan is 0.4863 mm thick. The total CT-scan file is 77.75 megabytes large.

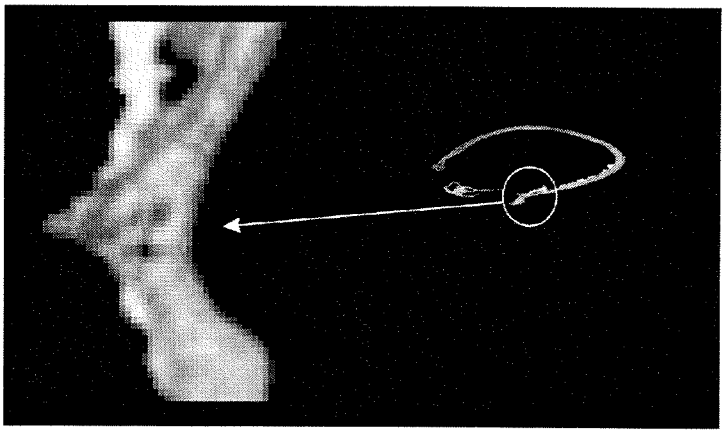

Figure 2

The CT-scan of the slice shown in Figure 1 and a detail enlarged 10×. The individual pixels can be clearly identified. (One can also see a cavity in the nasal region of the cranium!) The brightness of the pixels is a measure of the *attenuation* coefficient of the x-rays passing through the fossil material (fossilised bone and encrustation). The very brightest pixels in this detail are encrustations. Note that the boundary between encrustation and fossilised bone is not sharp everywhere. If a thresholding algorithm cannot be used, then the researcher trying to separate fossilised bone from encrustation must devise a suitable algorithm than can detect the boundary between the two materials. (The slice in this scan has 512 × 512 pixels; the grey scale is 8-bit resolution.)

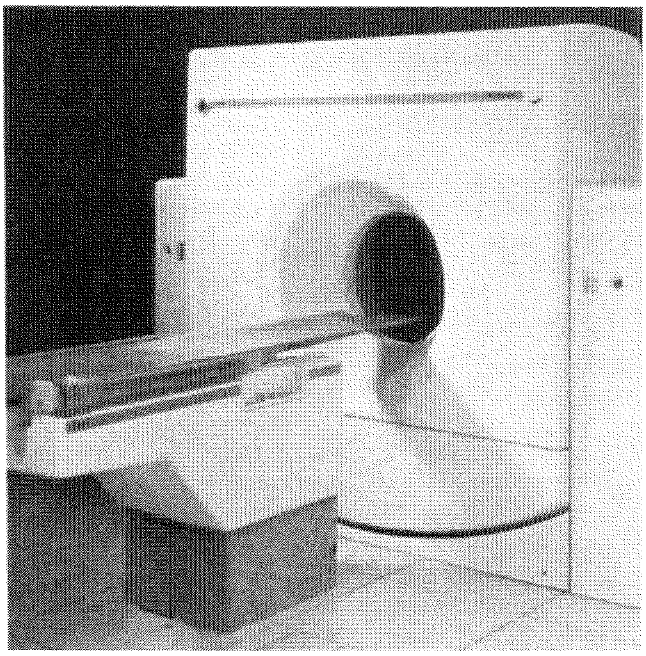

Figure 3

A photograph of a typical medical scanner set-up. The CT scanner is inside the large housing with its circular aperture. A patient (or — in the case of paleoanthropology research — a fossil in a box) rests on a bed that is mounted on a table. The bed is made of a synthetic material that hardly attenuates x-rays. The bed can be moved forward in incremental distances (called slices) into the circular aperture. After the patient/fossil has been moved by one incremental distance (*i.e.*, one slice thickness), the x-ray emitting apparatus (not visible — it is inside the housing) revolves once about the patient and a circular array of detectors (also inside the housing) registers the non-absorbed x-rays during the revolution. After one complete revolution of the x-ray emitting equipment, the bed is then moved forward by one slice and a further slice image is made. The investigator can predetermine how many slices of a patient/fossil should be scanned — it is sometimes not necessary to make a scan of the total object. The CT scanner inside the housing is shown in Figure 4.

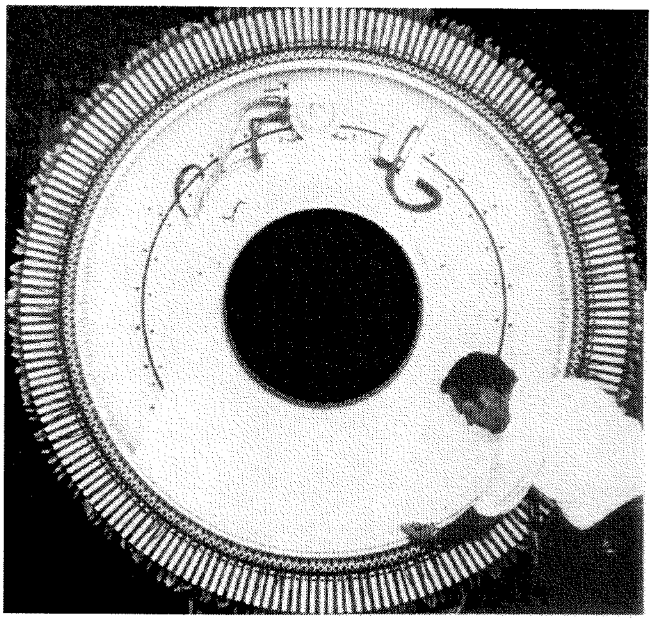

Figure 4

A photograph of a partially disassembled CT scanner while being serviced by a technician. The x-ray emitting apparatus has been removed; during scanning, it would be attached to the plate that is being inspected by the technician. The many detectors arranged on the periphery can be seen. During the scanning operation, the plate with the attached x-ray emitting apparatus revolves around the object to be scanned, and the attenuation by the object's material(s) is registered by the diametrically opposite detectors.

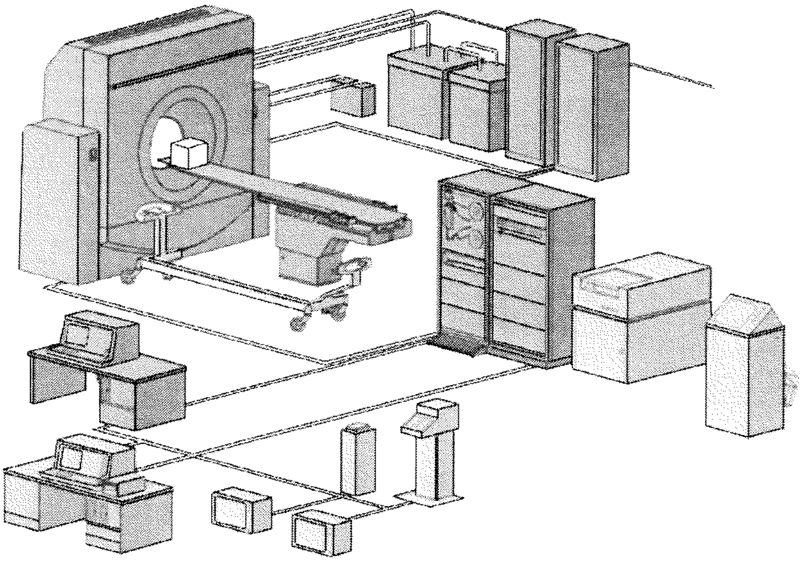

Figure 5

A typical medical scanner set-up together with its peripherals when used to scan a fossil specimen. The CT scanner (see Figure 4) is about to begin scanning the fossil (material) which is inside a white box. The fossil has been placed in a cardboard box so that it can be held fixed inside the box with some soft material, such as tissue paper. The box material and the tissue paper hardly attenuate any x-rays. Keeping the fossil fixed in the box during scanning is essential: while the box (with fossil) is moved into the scanner, one slice distance at a time, the research team does not want the orientation of the fossil to alter between slice images. If the scanner is used for patients, then the patient is lying on the bed on which the box shown here rests. The bed itself can be lifted from the table and placed onto a cart (shown with wheels) so that, after having been scanned, the patient can be rolled back into the hospital wards. The computing peripherals are shown in a lighter grey. Note that the interconnections between scanner and peripherals shown in this figure are not the LAN referred to in the text.

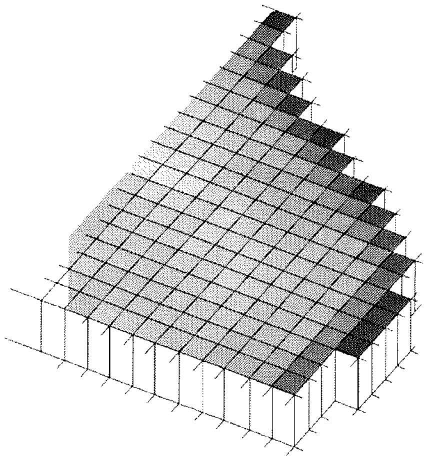

Figure 6 (a)

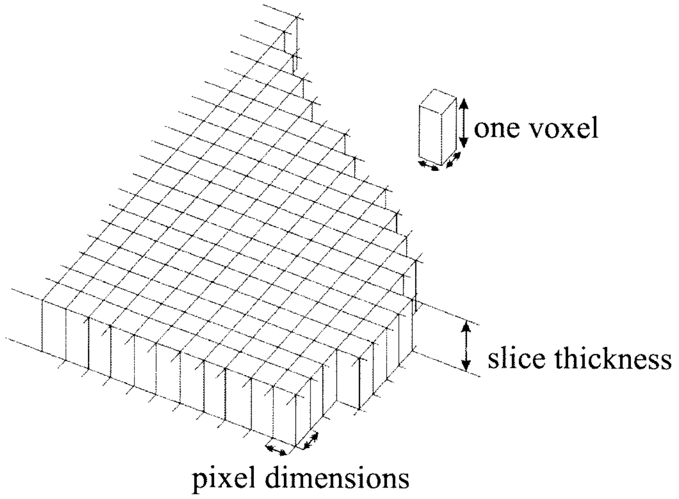

Figure 6 (b)

The relation between slice thickness and voxel height. During scanning, the slice thickness is the distance the fossil is moved forward through the circular aperture. After reconstruction of the CT-scan (file), the slice thickness has become the voxel height. In (a), a small segment of a slice has been extracted from the scan and imaged. The surface of the voxels form the pixel pattern seen in CT-scan slice images (Figure 2). In (b), the geometric relation between pixel dimensions and voxel height are clarified. In most scanners, the resolution is 512 × 512, which means that each slice consists of a square grid, 512 pixels along one side. By applying an interpolation algorithm, researchers can create iso-voxels, in which the height is the same as the dimensions in the plane of the slice (*i.e.*, the pixel dimensions).