

Chapter 10

Chemical and Isotopic Studies of Ancient Metals

A. Mark Pollard and Peter Bray

Introduction to Chemical and Isotopic Analysis of Metals

It is not the intention of this chapter to give a thorough introduction to the many and varied methods which have been used to analyze archaeological metals. Some of these can be found in the selected bibliography at the end of this chapter. The aim of this section is to guide and advise the student about the different sorts of analyses which can be done, and which ones might be appropriate under what circumstances.

When embarking upon an analytical study of ancient metalwork, the researcher must start with a single fundamental question: ‘what is the purpose of the analysis?’ or, perhaps more precisely, ‘to what purpose will the analytical data be put?’ Answers to this question might range from the apparently trivial ‘I want to know what it is made of’ to the ‘I need to know from which particular ore deposit this metal was ultimately derived’. Another key question, and not unrelated, is ‘how destructive can my analysis be?’ Increasingly, museums and collection managers are asking for completely nondestructive (preferably ‘noninvasive’) analyses, which might make sense from a curatorial point of view but not necessarily from an academic perspective. It calls into question the reason why we are collecting these things at all—simply to preserve them as ‘objects’ in perpetuity, or as an archive of information which can help us understand the human past? In many ways, this is a false dichotomy. On the one hand, analytical methods are tending toward being more-or-less nondestructive (but not noninvasive!), and so it might be seen as prudent to restrict sampling access for as long as possible. On the other hand, one might argue that if the information is worth having, then it is worth a small sacrifice. We are back to the quality of the question again.

A. M. Pollard (✉)

Research Laboratory for Archaeology and the History of Art, University of Oxford,
Dyson Perrins Building, South Parks Rd, OX1 3QY Oxford, UK

e-mail: mark.pollard@rlaha.ox.ac.uk

P. Bray

e-mail: peter.bray@rlaha.ox.ac.uk

In addition to the rise of 'nondestructive' methods, museums are beginning to prefer analyses that bring the instrument to the collection rather than the objects to the instrument. Handheld instruments that can be taken into the field or to a museum have varying strengths and weaknesses in comparison to standard laboratory-based equipment, or large international facilities such as synchrotrons or neutron sources, which require the object to be taken to the facility. As many countries are increasingly taking control of the export of antiquities even for research purposes, a student of the metalwork may have to visit the country to look at recently excavated material, and either take the portable analytical tools with them, or make arrangements to have the work done with local partners.

In terms of the choice of analytical tools, let us be quite clear—there is no universal panacea for the analysis (chemical or isotopic) of metals. No one technique has all the answers, although some may now seem to approach it. No one technique is 'best'. Nor can the analysis be carried out from a recipe book. Each case is different. Different forms of ores and mineral deposits mean that, in some studies, one particular element or suite of elements may be the critical distinguishing factor, but in others, it may be irrelevant. In some cases, isotopic measurements may be diagnostic, while in others they may be completely ambiguous. What is 'right' depends on the nature of the question. Of course, as professionals we would always claim that the choice of analytical tool is made calmly and logically, to reflect the nature of the question. At one level this is true, but in many cases the choice is severely influenced by levels of availability and access, cost, and what is actually working at the time. Most graduate students know this from firsthand experience!

One issue to deal with is cost. All analyses cost money, even if no money changes hands. The true cost of some analyses can be enormous—consider the full cost of studying a metal object using a multimillion dollar synchrotron radiation source, or even just the marginal cost of doing so (i.e., not contributing to the cost of building the machine itself, but paying for the running costs). Somebody is paying for the capital investment and infrastructure that support these facilities. Archaeologists, and particularly (but by no means exclusively) commercial archaeologists (salvage or rescue archaeology), have a habit of considering some technical approaches as 'too expensive' to consider. Often there is no money in the budget to pay for any post-excavation analysis, and thus a lot of important work never gets done. However, there is no such thing as 'too expensive'. There is only a rational cost/benefit analysis, which balances the cost of obtaining the information against the intellectual value of that information. A cheap analysis which concludes that Roman nails contain iron, or that, in another context, pottery is made of clay, is poor value. An analysis (expensive or not) that rewrites our views on the role of copper alloys in the Early Bronze Age is good value. Again, it is the quality of the question that counts. There is a clear onus on the analyst to be able to explain to managers and finance officers what their work can deliver intellectually, above and beyond an appendix to a site report.

There is also a need to take into account the different levels of analysis that are possible, and what might be necessary for the question asked. A key differentiation is between destructive and nondestructive analysis, as described in more detail below. Although these categories appear to be mutually exclusive, the distinction between

them is not always clear-cut. A second distinction to be made is between *qualitative*, *semiquantitative*, and *quantitative* analyses. Crudely speaking, *qualitative* is simply asking whether a particular element is present or not. This might be sufficient to identify a metal as gold rather than something else, or an alloy as bronze rather than brass, and may be all that is required. *Semiquantitative* can mean either an approximate analysis or that the amount of a particular element is categorized into broad groups, such as ‘major’ (a lot), ‘minor’ (a bit), or ‘trace’ (just enough to detect), or perhaps into bands such as ‘1–5 %’.¹ Again, this might be sufficient to answer the question ‘is the major alloying element tin or zinc?’ A *quantitative* analysis (sometimes referred to somewhat tautologically as ‘fully quantitative’), on the other hand, is an attempt to quantify to some specified level of precision all of the components present in an alloy above a certain level.² Ideally, we should aim at quantitative analysis of as many elements as possible, validated by accompanying data on internationally accepted standard materials. In the end, quantitative data will have lasting value, and will mean that future analysts will not need to go back to the object and take more samples. Mostly, however, we compromise.

Nondestructive Chemical Analysis

Nondestructive methods are increasingly important in the study of archaeological and museum objects. First and foremost, among these techniques is X-ray fluorescence (XRF), which is now increasingly available as a handheld portable system. XRF uses a source of X-ray (usually an X-ray tube, but occasionally a sealed radioactive source) to irradiate the sample. The incoming (‘primary’) X-rays interact with the atoms that make up the sample and cause them to emit secondary (‘characteristic’) X-rays of a very specific energy. In the majority of XRF systems (the so-called energy dispersive systems), these secondary X-rays are captured in a solid-state detector, which both measures their energy and counts their number. By knowing the energy of the secondary X-ray, we can identify which atom (i.e., element) it has come from, and from the number, we can estimate how much of that particular element is present in the sample. If the analysis is done in air (i.e., without the need to enclose the sample in an evacuated chamber), then the whole device can be made portable and the machine can be taken to the object, although the sensitivity and precision of the analysis is somewhat compromised. Some portable devices now come with ‘skirts’ that fit over the object and allow at least a partial vacuum, which improves the performance somewhat.

¹ It is important to note whether results are expressed in atomic or weight %. Atomic % denotes the percentage of the total number of atoms in a sample that are of a particular element. Weight % denotes what percentage of the mass of the sample each element contributes. If an archaeometallurgy paper does not explicitly state which system has been used, with caution, it can be assumed the results are in weight %.

² Bear in mind that, with at least 93 naturally occurring elements in the Periodic Table, no analysis containing fewer elements can truly claim to be ‘fully quantitative’!

The advantages of XRF are many. It can give a rapid analysis, typically done in a few minutes or so. For metals, it can analyze most elements of interest, but if an air path is used (as described above), it cannot detect any element below calcium (Ca, atomic number 20) in the Periodic Table. Even under vacuum, the lightest element which can be seen by XRF is sodium (Na, atomic number 11). Thus, if using an air path, modern alloys containing magnesium or aluminum will show no sign of these components. Of the metals known in antiquity, however, this is only a problem for iron, where the main alloying elements (e.g., carbon and phosphorus) are too light to be seen by XRF, even under vacuum. Nevertheless, XRF is particularly well suited to metal alloy identification and quantification, as the majority of alloys used in antiquity (gold–silver/electrum and the bronzes and brasses) are primarily made up of heavy metals which show up well under XRF.

The disadvantages of XRF are also numerous. It is not a particularly *sensitive* method—sensitivity here is taken to mean the smallest amount of a particular element that can be measured accurately. For most elements in XRF, this is typically around 0.1 wt%, but the exact value will depend on the element being detected and nature of the other elements present. This is more than adequate if the nature of the question is ‘what alloy am I dealing with?’ Or, indeed, ‘what are the major impurities present in this metal?’, but it is nowhere near as sensitive as many of the techniques that involve sampling. More significantly, it is also a surface-sensitive technique due to absorption of the secondary X-rays by the metal matrix as they leave the sample. The degree to which these secondary X-rays are absorbed depends on the energy of the secondary X-ray and the nature of the material(s) through which it needs to pass in order to get to the detector. Typically, information is only obtainable from the top few tenths of a millimeter below the surface. Thus, if the surface is corroded, dirty, or unrepresentative of the bulk composition for any reason (see below), then the analysis given may be completely misleading. Sometimes this problem can be reduced by judicious mechanical cleaning of the surface (or preferably an edge) to reveal bright, uncorroded metal, but this may be curatorially unacceptable. This may be a particularly serious problem with handheld instruments, because often the primary X-ray beam diameter is a centimeter wide (to compensate for the relatively low intensity of the source); thus, the area needing to be cleaned may be too big for most museum curators. Other nondestructive analytical tools are available to the archaeometallurgist, provided he/she is prepared to transport the object to the facility. These include synchrotron XRF (sXRF), which is identical to other forms of XRF apart from the source of the X-rays, which in this case is a huge circular electron accelerator called a synchrotron. Accelerated to almost the speed of light, these electrons give off electromagnetic radiation of all frequencies from high-energy X-rays to long-wavelength radio waves. By harnessing the X-ray emission, the synchrotron can provide a highly intense collimated beam of monochromatic X-rays, which provide an excellent source of primary X-rays for XRF. Because they are so intense, it is normal to carry out sXRF in air, so that large objects of any shape can be accommodated in front of the beam, although the secondary X-rays are still attenuated by passage through air as described above.

In a similar vein, other forms of particle accelerator can be used to provide the primary beam for analytical use. Perhaps the most common is proton-induced X-ray emission (PIXE). In this case, a beam of protons is accelerated to high energies in a linear accelerator and extracted as a fine beam, which then strikes the object being analyzed. The protons interact with the atoms in the sample, causing them to emit secondary X-rays, as in XRF. Because the beam is external to the accelerator, objects of any shape can be analyzed in air, with the same advantages and disadvantages as above. The analytical sensitivity of PIXE is generally better than XRF, because proton impact causes a lower background X-ray signal than is the case with primary X-rays or electrons, so sensitivities on the order of a few parts per million (ppm) are possible.

Recent instrumental developments in PIXE have resulted in highly focused proton beams (down to a few microns) that allow chemical analysis and elemental mapping on the microscopic scale. This technique is called microPIXE, or μ -PIXE. Variants of PIXE include particle- (or proton-)induced gamma-ray emission (PIGE) in which the elements in the sample are identified by the gamma rays they emit as a result of proton beam irradiation (rather than by secondary X-rays). This method can be used to detect some of the lighter elements that XRF or other methods cannot detect.

(Mildly) Destructive Chemical Analysis

Where at all possible, it is usually better from an analytical point of view to remove a small sample from the metal object, polish it, and mount it for analysis. This also allows physical examination of the artifact (e.g., metallography, hardness, etc., as described in Chapter x) before any chemical analysis is carried out. Sampling also allows for spatially resolved chemical analysis—i.e., how does the composition of the sample vary from the surface to the interior of the artifact? When combined with information on the phase structure provided by optical microscopy, chemical analyses of this kind are an extremely valuable entry point into understanding the manufacture and use—the ‘biography’—of the object.

Once a sample has been taken, there are a very large number of choices of analytical instrumentation available. The two most common ones used today are based on either electron microscopy (EM) or inductively coupled plasma spectrometry (ICP).

Electron Microscopy

EM has been the workhorse of the chemical study of metals for more than 40 years. It is important to realize that although the microscopy side of the instrument is dependent upon a beam of electrons, the actual chemical analysis is still primarily dependent upon the detection of X-rays. EM has many advantages. It is widely available and can produce high-resolution images of the physical structures present, as well

as micro-point chemical analyses and two-dimensional (2D) chemical maps of the exposed surface. As a technique, it has had many pseudonyms over the years, based upon the particular configuration available—electron microprobe analysis (EMPA), scanning electron microscopy (SEM), scanning electron microprobe with energy dispersive spectrometry (SEM-EDS), SEM with wavelength dispersive spectrometry (SEM-WDS), and probably more.

As an analytical tool (as opposed to a high-powered imaging device), the key aspect to understand about EM is how the chemical analysis is actually achieved, which in effect means: how are the secondary X-rays detected? In an EDS system, the process is almost identical to that described above for XRF, except that the primary beam in this case is made up of electrons rather than X-rays. These primary electrons have the advantage over X-rays of being steerable and focusable using electrostatic devices, and having more controllable energies. The electrons strike the sample and interact with the atoms to produce characteristic X-rays, which are collected and counted using solid-state detectors as before. Because such an analysis usually takes place in a high-vacuum chamber on a prepared and polished flat sample, many of the problems inherent with XRF are reduced. The sensitivity to lighter elements is much better than XRF, even to the extent that some systems can measure carbon in iron. Furthermore, the problem of surface sensitivity can be controlled by cleaning and/or preparing the sample in such a way that it provides a cross section of the artifact. One big advantage of energy dispersive analysis combined with SEM is that the speed of the analysis is such that 2D chemical maps of the prepared surface can be readily produced, showing just how and where the elemental inhomogeneities are distributed within the metallographic structures (i.e., are certain elements within the grains or between the grains, near the surface or near the core, etc.).

Less widely available (not only because of higher cost but also because of the expertise required to produce good data), but regarded by many as the 'gold standard' for metals analysis, is wavelength dispersive X-ray detection (WDS) in EM. This method differs from EDS only in the way that the X-rays are detected. However, it tends to make the machine much more suited to analytical rather than imaging applications, making it rather more specialized and less flexible than the 'all purpose' machines that use (energy dispersive) ED detection. In practice, many of the larger and more expensive machines have multiple wavelength detectors and an ED detector, giving multiple functionality. In WDS machines, the X-rays emitted from the sample are considered to be waves in the X-ray region of the electromagnetic spectrum, rather than particles of a particular energy. The nature of the parent atom is defined by the characteristic wavelength of the X-rays produced, and the number of atoms present is defined by the intensity of this wavelength. This is a graphic illustration of the quantum mechanical principle of particle-wave duality! The wavelength of the X-rays is measured using a crystal as a diffraction grating, which resolves the X-rays into their component wavelengths (in the same way as water droplets produce a rainbow), and the intensity of each wavelength is measured via a solid-state detector.

Without going into too much detail, WDS analysis in EM gives better sensitivity than EDS (typically detecting elements down to 0.001–0.01 %, rather than 0.1 % by EDS), but it is slower, which makes 2D mapping more cumbersome. In both

forms of detection, however, because the primary irradiation is by electrons rather than X-rays, the beam can be focused down to a few microns in diameter, giving a spatial resolution for the chemical analysis in the order of a few tens of microns. This is necessary when looking at phenomena such as age embrittlement in alloys, where microscopic phases are precipitated at grain boundaries, or when attempting to characterize small inclusions in slag or metal that can tell us about the ore sources used and the metalworking processes applied.

It is worth pointing out that much more can now be achieved through EM without sampling than was possible just a few years ago. Environmental chambers, operating at near-atmospheric pressure, are available which can accommodate large objects (up to 30 cm in diameter and 8 cm high). Such equipment can produce images and analyses without the need for applying a conducting coat of gold or carbon to prevent charging. Software developments include ‘hypermapping’, where the data are collected and stored as a series of superimposed elemental maps. This allows the analyst to go back to the data at any time and generate a point analysis or a 2D map without needing to reanalyze the sample. The advantage of this is that it is feasible to collect all the data that might be conceivably needed now and in the future, and thus limit the analysis of an object to a one-off event. The downside of this method is the same as for nondestructive XRF: surface-only analysis and lowered sensitivity.

Inductively Coupled Plasma Spectroscopy/Spectrometry

ICP in various formats is now the method-of-choice for chemical analysis in a wide range of applications across research and industry. In its current form, it represents the refinement of a long line of optical spectroscopy techniques for chemical analysis, going back to optical emission spectroscopy that was developed in the 1930s. The principle is simple—when an atom is ‘excited’ (i.e., given a lot of energy), it reorganizes itself, but almost instantaneously de-excites back to its original state with the emission of a pulse of electromagnetic radiation, which is often in the visible part of the spectrum. In other words, it gives off light. The wavelength of this light is characteristic of the atom from which it came, and the intensity (amount) of light is proportional to the number of atoms of that particular element in the sample. Measurement of the wavelength and intensity of the light given off from a sample therefore forms the basis of a quantitative analytical tool of great sensitivity. In ICP, the sample is introduced into an extremely hot plasma at around 10,000 °C, which causes the atoms to emit characteristic frequencies. This emitted light is resolved into its component wavelengths using a diffraction grating, and the intensity of each line of interest measured. In this form, the instrument is known as an inductively coupled plasma optical emission spectrometer (ICP-OES), or sometimes as an ICP atomic emission spectrometer (ICP-AES). The use of ICP as an ion source for mass spectrometry (inductively coupled plasma mass spectrometer (ICP-MS)) is discussed in more detail below.

ICP is an extremely sensitive means of measuring elemental composition. Depending on the element to be measured and the matrix it is in, it can usually detect many elements down to levels of ‘parts per billion’ (ppb, equal to 1 atom in 10^9) from a very small sample (which varies depending on the particular needs of the experiment). The sample can be introduced into the spectrometer as a liquid (i.e., dissolved in acid), but more recently attention has switched to a device which uses a pulsed laser to vaporize (ablate) a small volume from a solid sample into a gas stream which then enters the plasma torch. This is known as laser ablation ICP (LA-ICP-OES). Laser ablation has several advantages over solution analysis, not least of which is simplicity of sample preparation and the potential to produce spatially resolved chemical analyses of complex samples (including 2D elemental maps, line scans, etc., as described above). The disadvantages include somewhat lower levels of sensitivity and a much more complex procedure for producing quantitative data. However, the fine scale of the ablation technique (the crater produced by the laser can be typically 10–15 microns in diameter and depth) means that for all intents and purposes, the analysis can be regarded as ‘nondestructive’. This assumes, of course, that the object is small enough (perhaps up to 10 cm across) to fit into the laser ablation chamber (and, of course, that clean metal can be obtained to give a non-biased sample). It is therefore now possible to analyze small metal objects without cutting a sample, in which up to 20 or 30 elements are quantified, covering the concentration range of the major and minor alloying elements, down to sub-ppm levels (or 1 in 10^6) of trace elements. As discussed below, in some circumstances isotopic ratios for particular elements can also be measured if a mass spectrometric detector is used.

The elemental analytical capability of most ICP methods is comparable to or better than the data previously produced by neutron activation analysis (NAA), which is historically the preferred means of analyzing trace and ultra-trace elements in archaeological and geological sciences. NAA is not discussed in detail here because it is rapidly becoming obsolete as a result of increasing difficulties in obtaining neutron irradiation facilities. However, it is described in most standard texts on analytical and archaeological chemistry and is an important method to understand due to its prevalence in earlier research. Suffice it to say that the archaeological science literature (including archaeometallurgical literature) contains a good deal of high-quality NAA data, which raises the question of how we should use this legacy data—a matter discussed in more detail below.

Analyzing Metals with Surface Treatments

One important feature of the analysis of ancient metals (which is much less common in other branches of inorganic archaeological chemistry) is the fact that the surfaces of ancient metals are usually different in chemical composition from the bulk of the metal. This provides a challenge to the analyst. For example, a simple handheld XRF scan of an archaeological metal surface will almost certainly produce an analysis that is not representative of the entire artifact. This may not be a major problem if the purpose of the analysis is to simply categorize the object into a class of metal—arsenical copper, brass (copper–zinc), bronze (copper–tin), leaded bronze, etc., or to

decide if the object contains gold or silver (essentially, a qualitative analysis). If the purpose of the analysis is to be more precise than this (e.g., to talk in terms finer than percentage points), then handheld XRF on unprepared surfaces is not particularly useful.

Metal surfaces can differ from bulk composition for two main reasons, described in detail below. One is that they were made that way, either deliberately or accidentally. The second is that the long-term interaction of the metal with the depositional environment has altered the surface, sometimes but not always resulting in a mineralized surface known as corrosion or patina.

Deliberate surface treatments are well known from antiquity, and come in a wide variety of forms. These include the deliberate tinning of a bronze surface to give a silvery appearance, and the surface enrichment of gold objects to remove base metals to give a richer gold appearance. These come under the category of deliberately applied surface finishes, and they may be the result of either chemical or physical treatment, or both. Some treatments are highly sophisticated and still poorly understood, such as the use of niello (a black mixture of sulfides of copper and silver, used as a decorative inlay), or shakudo (a gold and copper alloy which is chemically treated to give dark blue–purple patina on decorative Japanese metalwork), or even the enigmatic Corinthian Bronze ('aes'), which may have had either a black or golden appearance.

This area of research opens up the knotty question of 'how was an object or statue expected to look in antiquity?' We are today most familiar with bronze statues looking either green from copper corrosion or bronze-colored if 'cleaned'. However, it is highly likely that in certain periods of antiquity at least some statues were intended to look very different—they may have been gilded, or surface-treated to give black–purple colors, or even painted. The careful detection through chemical analysis of any remnants of surface finishes is crucial to understanding how these objects were intended to look, and what might have been their social function. It also leads into the fascinating area of deception, either for sheer fakery (e.g., making gold coinage look finer than it actually is), or into the mystical world of alchemy, where base metals can be given the appearance of gold. There is much serious archaeological research to be done on the metallurgical underpinnings and consequences of alchemy.

Accidental surface treatments include various segregation phenomena, such as 'tin sweat' on bronzes, in which the casting conditions are such that the tin-rich phases preferentially freeze at the surface of the mould, giving a silvery appearance. It is sometimes difficult to decide whether a particular effect was deliberately intended, or purely accidental.

The long-term interaction of a metal with its burial environment (or, in rarer or more recent cases, with its atmospheric environment if it has never been buried) is essentially the story of the attack of metal by water, and is often electrochemically mediated. For many years, aesthetics dictated that metal corrosion products should be scrubbed or stripped away to leave the bare metal, in search of the 'original appearance' of the object. Thankfully, this barbaric approach has diminished. In fact, if one ever wants to provoke an 'old-school' metal conservator into a frenzy, it is worth simply injecting into the conversation that the most interesting bit of a metal object is the corrosion products, which historically at least used to be thrown away. Most

modern conservators would probably not entirely agree with this statement, but would concede that the corrosion products are part of the object's biography. In theory at least, the metal corrosion contains encoded within it the whole environmental history of the object, including evidence for manufacturing, use, discard, and deposition. Not only does this have implications for the materiality of the object, but it could also have implications for authenticity studies. Decoding this history is, however, another story.

In the light of all of this, it should be clear that when approaching a metal archaeological object for the purposes of chemical or isotopic analysis, the analyst must *expect* to find surface anomalies that may be vital to the biography of the object, or which will at least, if not identified and countered, render any analyses quantitatively unreliable. Each object is in this sense at least unique. The best advice to the analyst is to start with a careful optical microscopic examination, and then work upward!

Isotopic Analysis of Metallic Objects

In addition to the chemical composition of an object, another key attribute that can be measured is the isotopic composition of some, or all, of the elements in that object. Isotopes are different versions of the same element, which differ only in their mass. The chemical identity of an element is determined solely by the number of protons in the nucleus, though the number of neutrons can be variable. Thus copper (proton number 29) has two naturally occurring isotopes, indicated as ^{63}Cu and ^{65}Cu , meaning that one isotope has 29 protons plus 34 neutrons in the nucleus (i.e., 63 particles in the nucleus in total) and the other has 29 protons plus 36 neutrons (65 in total). The 'natural abundance' of these two isotopes is roughly 75 % and 25 %, respectively, meaning that the average atomic weight of naturally occurring copper is approximately 63.546.

Chemically speaking, the isotopes of an element behave identically, but sometimes they take part in chemical reactions or physical transformations (e.g., evaporation) at slightly different rates. This is because the difference in mass of the isotopes means that bond energies between one isotope and, for example, oxygen will be slightly different from that of another isotope and oxygen. Thus, in the course of various chemical reactions or transport phenomena, the ratio of one isotope to another of the same element may change. In the case of chemical reactions, this is termed *fractionation*, and means that the environmental or geological history of certain metals can be reconstructed by measurements of isotopic ratios.

Some stable isotopes are formed as the result of the radioactive decay of another element. For example, strontium has four stable naturally occurring isotopes: ^{84}Sr (natural abundance 0.56 % of the total Sr), ^{86}Sr (9.86 %), ^{87}Sr (7.0 %), and ^{88}Sr (82.58 %). Of these, ^{87}Sr is produced by the decay of the radioactive alkali metal ^{87}Rb , and is therefore termed *radiogenic*. Thus, the isotope ratio $^{87}\text{Sr}/^{86}\text{Sr}$ in a rock is related to the original isotopic composition of the rock, but will also change with time as the ^{87}Rb originally present turns into ^{87}Sr .

Almost all metals of interest to the archaeometallurgist have more than one naturally occurring isotope. In fact, elements with only one isotope are exceptional in nature, and the ones of most relevance here are gold, which occurs naturally only as ^{197}Au , and arsenic (^{75}As). Many, such as Cu (with the two isotopes listed above) and silver (^{107}Ag and ^{109}Ag) have a couple, but a significant proportion has several stable isotopes, including iron (with 4), lead (4), Ni (5), Zn (5), and, champion of them all, tin (10). There is thus plenty of scope for measuring isotopic ratios among the metals of interest to archaeologists.

Foremost and by far, the most intensively studied is lead. It has four stable isotopes, ^{204}Pb , ^{206}Pb , ^{207}Pb , and ^{208}Pb , the latter three of which are the stable end members of the three long natural radioactive decay chains found in nature (starting with uranium (^{235}U and ^{238}U) and thorium (^{232}Th)). The approximate natural abundances of these four isotopes of lead are 1.4, 24.1, 22.1, and 52.4 %, respectively, but the precise abundances in any particular mineral or geological deposit depend on the geological age of that deposit, and the initial concentrations of uranium and thorium. Because of the large range of possible starting conditions and the differing geological ages of mineral deposits, there is large variability in the measured abundance of geological lead isotope ratios—much larger than in any other metal of interest here. For reasons of measurement discussed below, it is conventional in the study of heavy metal isotopes to deal with ratios of one isotope to another rather than absolute values of abundance. Thus, in archaeological discussions of lead isotope data, it is conventional to deal with three sets of ratios ($^{208}\text{Pb}/^{206}\text{Pb}$, $^{207}\text{Pb}/^{206}\text{Pb}$, and $^{206}\text{Pb}/^{204}\text{Pb}$) which can be plotted together as a pair of diagrams.

Measurement of Isotopic Ratios in Metals

Measurement of the isotopic ratios of metals requires some form of mass spectrometry, in which the numbers of constituent atoms in the sample are directly measured after being separated according to weight. Conventionally, this has been achieved using thermal ionization mass spectrometry (TIMS), in which the sample to be measured is chemically deposited from a solution onto a fine wire, mounted in a mass spectrometer, and heated to evaporate and ionize the sample. The ions are then extracted into the spectrometer electrostatically, and ions of different mass are steered into a bank of detectors—one for each mass of interest (this is called a ‘multi-collector’ or ‘MC’ instrument). Rather than measure each isotope independently, however, it is better to report the results as ratios of one isotope to another, as the data can be directly recorded as the ratio of the electric current flowing through each of the two detectors. This means that any fluctuation in ionization efficiency at the source is seen as an equal fluctuation in both detectors. The potential inaccuracy is therefore canceled out, making the ratio more precise than a single measurement. The main drawback of TIMS is the time taken to prepare the sample. It needs to be dissolved into high-purity acids, often concentrated by passing through an exchange column, and then deposited on the wire. The number of samples that can be analyzed in this way is relatively slow.

The preeminent position of TIMS for isotopic studies was vastly altered in the late 1990s, when it was realized that an ICP torch is also an extremely efficient source of ions. Rather than using optical techniques to measure the concentration of atoms in the sample (as in ICP-OES, described above), it was realized that if the plasma containing the ionized sample could be fed into a mass spectrometer, then a new instrument was possible—ICP-MS. The technical trick is to interface an extremely hot gas from the plasma source with a mass spectrometer under vacuum in such a way that the mass spectrometer does not melt and the vacuum in the mass spectrometer is not destroyed. This was achieved in the 1980s, and such instruments became commercially available through the 1990s. Early systems of this type used a low-resolution quadrupole mass spectrometer as the detector, because it is fast, cheap, and efficient. It is very effective for trace element abundance analysis, but the isotopic ratio measurement precision of a quadrupole is typically 100 times poorer than can be achieved by TIMS, and so the instrument saw limited use as an isotopic measurement tool. The next generation of ICP instruments used much more sophisticated mass spectrometers—either bigger systems with much higher resolution, or (more effective still) feeding the output from one mass spectrometer into a second mass spectrometer to give much improved resolution—combined with multi-collector detectors. Such instruments, termed ‘HR-MC-ICP-MS’ (high-resolution multi-collector ICP-MS) or ‘ICP-MS-MS’ (ICP with two mass spectrometers), are capable of measuring isotopic ratios with precisions at least as good as TIMS (if not better), but with much simpler sample preparation. TIMS instruments are quoted as giving 95 % confidence levels in precision (i.e., measurement reproducibility) of ± 0.05 % for the $^{207}\text{Pb}/^{206}\text{Pb}$ ratio, and ± 0.01 % for $^{208}\text{Pb}/^{206}\text{Pb}$ and $^{206}\text{Pb}/^{204}\text{Pb}$: a modern high-resolution multi-collector will do at least as well as this, if not an order of magnitude better.

Initially, ICP-MS instruments were designed for solution input, but as described above, the addition of a laser ablation ‘front end’ for spatially analyzing solid samples has just about revolutionized the practice of isotopic ratio mass spectrometry for heavy metals such as lead. As long as a solid sample can fit into the ablation chamber, such analyses can be carried out without sampling and with virtually no visible damage. That is not to say such measurements are cheap or easy, but it does mean that large numbers of measurements can now be made quickly and at high precision, reopening the potential for large-scale studies of ancient metals.

What Can Chemical and Isotopic Data Tell Us?

There is a long history of the chemical analysis of ancient metalwork, going back at least to the late eighteenth century in Europe, with Martin Heinrich Klaproth’s investigation in 1795 of the composition of ancient copper coins and Pearson’s 1796 paper analyzing Bronze Age tin-bronze. Until the mid-twentieth century, the question was primarily one of ‘what is it made from?’—that is, which metals and alloys were used by ‘the ancients’, and what was the sequence of their development? Interestingly,

from this simple question, systematic patterns of behavior appeared to emerge over much of the Old World, although at different times. The large-scale chemical analysis of metal artifacts from known archaeological contexts during the twentieth century allowed scholars to propose a model for the 'development of metallurgy in the Old World' from an early use of 'native' metals (i.e., metals like copper, gold, silver, or iron which can occur naturally in the metallic state), to the smelting of 'pure' copper, followed by arsenical copper alloys, followed by tin bronzes, followed by leaded tin bronzes, and then eventually an Iron Age. This developmental sequence often provided a framework for archaeometallurgical research, which, as the rest of this volume demonstrates, has increasingly been superseded by detailed regional studies.

Since the 1950s, the dominant driver in the chemical analysis of metals has been the quest for provenance—i.e., can metal artifacts be traced back to a particular ore deposit? Given that the dominant model of technological evolution at this time was one of diffusion, it was plainly logical to use the artifacts themselves to see if tracing metal back to ore source could identify the source whence the knowledge of metals diffused. It quickly became apparent that matching the trace elements in a particular copper alloy object with those of its parent ore deposit was a tall order. Indeed, there are many potential sources of copper ores in the ancient world, and broadly speaking, ores from similar genetic and geochemical environments are likely to have similar patterns of trace elements. Nevertheless, from the tens of thousands of analyses that have been done on European Bronze Age objects, it is certainly true to say that particular combinations of trace element patterns are discernible within the data, which many think are inevitably linked with particular regional ore sources. Given that typological approaches to European Bronze Age copper alloy artifacts have defined particular assemblages of objects as 'industries', then it is clear that analytical and typological analyses were moving in the same direction. Hence the obsession with 'groups' of trace elements which has permeated much of the twentieth-century thinking on European Bronze Age metallurgy.

The general failure of this approach has been well documented. It is certainly true that particular types of copper ore deposits are likely to give rise to specific combinations of trace element impurities in the smelted copper (or, in some cases, to the absence of specific impurities). These impurity patterns may well have regional or even specific locality significance, but only in the simplest (and rarest) of cases—that is, when the metal is 'primary' (i.e., made directly from smelted copper from a single ore source) and not mixed with copper from other sources or recycled and remelted. In other words, if a metal artifact has received the minimum of manipulation, consistent with that required to convert an ore to metal, then this approach may be suitable. In all other cases, which represent the vast majority of ancient metalwork that has come down to us, a more sophisticated model is required.

A fundamental problem with this standard model of metallic provenance is that the dominant factor which is thought to affect the trace element pattern in the metal artifact is the original trace element composition of the primary ore. Although all authors engaged in this work have recognized that issues such as recycling and remelting are also certain to perturb this standard model, in practice most authors then proceed to interpret the data by ignoring these factors. Another factor to be




Fig. 10.1 Impurity content of copper alloy produced as a function of smelting temperature for charges containing 10% As, 2% Sb, and 2% Ni

considered is that there is not necessarily a one-to-one relationship between the trace element pattern in the ore and the same pattern in the smelted metal. Not only do some trace elements preferentially concentrate in the slag (if there is any) as opposed to the liquid metal (called ‘partitioning’), but there is even variation in this partitioning behavior as a function of temperature and ‘redox’ conditions (i.e., the degree of oxidation or reduction) in the furnace. It has been shown from laboratory experimentation, for example, that if a copper ore contains nickel, arsenic, and antimony, then all the antimony (Sb) in the charge is transferred to the metal at all temperatures between 700 and 1,100 °C. In contrast, the nickel (Ni) will not appear in the smelted metal at all unless the temperature is above 950 °C, while only low percentages of the original arsenic (As) content will be transferred at temperatures below 950 °C (see Fig. 10.1). Thus, if the ‘impurity pattern’ of interest consisted of the presence or absence of As, Sb, and Ni (all important elements in early copper-base metallurgy), then the impurity pattern for metal smelted from the same charge changes dramatically if the furnace temperature goes above 950–1000 °C. Under the standard model of provenance, such an observation would be taken as indicating a switch in ore source. This is not to say that all conclusions about provenance based upon changes in trace element pattern are wrong; it does, however, council caution if considering only one possible explanation for such a change.

Partly in response to these problems with the many large chemical analysis programs carried out in the 1960s and later, archaeometallurgists turned with glee to the new technique of lead isotope analysis (LIA). Much has been written about the value and limitations of LIA in archaeology as it has the potential to be an invaluable tool for provenance studies. LIA is based upon the fact that different ore bodies will contain varying initial amounts of thorium, uranium, and primogenic lead. Over varying amounts of time (depending of course on when the ore body was formed) radiogenic decay produces distinct lead isotope signatures that are dependent on the starting parameters of the ore body. This method proved to have importance not only for ancient metals that contain lead (i.e., copper alloys, silver alloys, as well as lead and lead/tin alloys) but also for glass, ceramics, and even human bone. The fact that LIA has yet to fully achieve this potential (at least as far as ancient metals are concerned) is probably also due to a lack of sophistication in the way that the provenance model has been applied, in much the same way as chemical provenancing discussed above.

One key issue that has yet to be fully addressed is the degree to which different ore deposits within a particular region can be expected to have distinctly different lead isotope signals. Clearly, this primarily depends upon the geological nature of the deposits, but it is worth recalling that much of the 1990s was taken up with a polarized debate between two key groups of LIA practitioners. One group stated that the various metalliferous outcrops on the Aegean islands could largely be distinguished from each other, while the other group (using much the same data) felt that they could not and preferred to use a single 'Aegean field'. This debate, as much about philosophy as about ore geology and mathematics, is not dissimilar to the debates in evolutionary biology between 'lumpers' and 'splitters'.

The most definitive recent statement using British data has come from the work of Brenda Rohl and Stuart Needham who have shown that the lead isotope fields for the four main mineralized regions of England and Wales effectively overlap and are therefore indistinguishable. They coined the term EWLIO—the England and Wales Lead Isotope Outline—to signify the lack of resolution between these sources. Moreover, the ore fields of the neighboring countries (Ireland, Western France, Germany, and Belgium) are also shown to be considerably overlapping with EWLIO. In other words, lead isotope data on their own are clearly not sufficient to distinguish between ore sources in northwestern Europe. Their research did, however, indicate a possible way forward by combining trace element patterns with lead isotope signatures and archaeological typologies, in a manner which has subsequently been built on elsewhere, as described below.

One interesting outcome of this discussion of the use of lead isotope data as a technique for provenancing has been the investigation of the conditions under which anthropogenic processing might affect the value of the isotopic ratio. It is well known elsewhere in isotope systematics that for the light elements (hydrogen, carbon, nitrogen, etc.), fractionation (changing of the isotopic ratio) is an inevitable consequence of many natural processes—indeed, it is this very fact that makes carbon isotopes such a powerful tool in tracing carbon cycling through the ecosystem. The question is: do heavy metal isotopes (Pb, Cu, Zn, etc.) undergo similar fractionation, especially as a result of anthropogenic processes such as smelting? Orthodoxy would say

no. It is, however, possible to show that under certain conditions, such as nonequilibrium evaporation from a liquid phase, it is theoretically plausible that fractionation would occur. Such conditions might easily occur during the processing of metals in antiquity. All that is required is that the vapor phase above a liquid metal is constantly removed during evaporation. As light isotopes preferentially enter the vapor phase, the remaining liquid is gradually enriched in the heavier isotope. It is not difficult to conceive of some stages of metal processing which approximate to these conditions.

Unfortunately, at the time these processes were being investigated (1990s), measurement techniques were too imprecise to detect what, for all practical metallurgical processes, were likely to be small effects. Certainly, the attempts of the time failed to measure significant fractionation as a result of anthropogenic processing of lead or tin. Zinc, the major component of brass, and an extremely volatile element, showed much more promise. Unfortunately, none of these experiments have been repeated using the much more sensitive analytical equipment available today. We know relatively little about the natural variation in copper, tin, and zinc isotopes between different types of deposit. More interestingly, perhaps, it is possible that anthropogenically induced fractionation in these isotopes might offer an opportunity to observe smelting and melting processes. It would not be surprising, for example, if measurement of zinc isotopes allowed a distinction to be made between brass made by the direct process (adding zinc metal to copper metal) and that made by the calamine process (vaporizing zinc to form zinc-oxide and then adding this to copper). Similarly, does the tin isotope ratio in a bronze change systematically in proportion to the length of time the metal is molten? Does it therefore change measurably each time a bronze object is recycled? There is clearly a need for a series of careful laboratory experiments, leading to a program of experimental archaeology. Working with Legacy Data

The world of archaeometallurgy is blessed with a plethora of data on the chemical composition of ancient copper alloys (and, to a much lesser extent, other alloys including iron), and also with a corpus of lead isotope data from ancient mine sites and from copper, lead, and silver objects from the ancient world. In theory, this dataset is of immense value, and provides a basis upon which all scholars should be able to build their research. Apart from the usual problems of non-publication of some key data, the frequent lack of sufficient descriptive or contextual detail, and the occasional erroneous values, these data (particularly the chemical data) suffer from one major problem: they were measured using the best techniques of their time, which were (usually) not as good as those we have now. It would, however, be wrong to assume that all data collected in the past are inferior to that obtained now. For example, gravimetric determinations of major elements in copper or iron alloys carried out during the nineteenth and early twentieth centuries will almost always bear comparison with—indeed, may be marginally better than—instrumental measurements carried out today. Obviously, there was no way these pioneers could have determined the minute levels of trace elements that can be done today, but in terms of major element analysis, these data should be good enough to use today.

What is perhaps more of a problem is data collected through the major analytical campaigns of the 1930s and later, using now obsolete instrumental methods of

analysis, such as optical emission spectrometry (OES) and subsequently atomic absorption spectrometry (AAS). OES in particular has well-known limitations, such as relatively poor reproducibility and a tendency to underestimate some elements in the alloy when they get to high concentrations. The key question for modern researchers is to what extent could or should we use this large database of ‘legacy data’? Some limitations are obvious. Many historical analyses lack data on certain key elements, which can only be remedied by reanalyzing the original sample, if it can be traced, and if it can be resampled. Archives of old samples taken from known objects, where they can be tracked down and properly identified, are priceless to modern researchers and must be curated at all costs—the size of the samples taken from precious objects as late as the 1980s would make most modern museum curators’ eyes water from pain!

Even if data on all the elements of interest are present, it is still difficult to imagine how to combine old and new datasets in a meaningful way—particularly if there are known deficiencies in the data for some elements. One approach is to say ‘don’t do it’. Indeed, major analytical projects on British metal objects from the 1970s onward were designed to create an entirely new dataset that excluded previous work. This is inherently cautious, but is probably a reflection of the tendency of most analysts only to trust their own data. However, is this simply wasting a vast archive of potential information, much of which, for financial or ethical reasons, we are unlikely to ever get again? This is the dilemma posed by legacy data.

In many ways, the existing scientific archaeometallurgical archive falls between two stools as it is often considered to lack sufficient scientific and archaeological information. The use of outdated analytical techniques and poor publication standards mean that modern archaeologists often have a stark list of numbers with little scientific context. The standards of modern chemistry require that analytical data are published alongside a number of measures of data quality. These include for each element the *limits of detection* (LoD, or *minimum detectable level*, MDL) of the analytical instrument, the *precision* of the measurements (how closely an analysis can be reproduced), and the *accuracy* of the data (how far the analysis deviates from the known true value of an internationally agreed laboratory standard). The large legacy dataset rarely provides this information and new research has to essentially trust the scientific proficiency and integrity of previous generations.

The large chemical legacy dataset has also been tarnished by the archaeological conclusions that were initially drawn from them. It was common for the chemical results to be interpreted solely by chemists (or mathematicians) who were relatively unfamiliar with the archaeological context of the metal objects. Baldly statistical interpretations of data often led to models of metal provenance and trade that clashed directly with the typologies, theories, and traditions of other archaeological specialists. An archetypal example of this was the *Studien zu den Anfängen der Metallurgie* (SAM) project, which relied on mathematically derived decision trees to assign provenance groups that were roundly attacked in the 1960s and 1970s. From a modern viewpoint, various statistical interpretative models of metallurgical data have been rightly ignored, but this has been at the expense of ignoring important raw data. Popular academic opinion has tended to confuse the notes with the tune; as Eric

Morecambe famously never said ‘we (probably) have all the right analyses, but not necessarily in the right order’.

Discussions of the future of chemical and isotopic archaeometallurgy commonly stress the two foci of research described in the first section of this chapter. First, gathering new scientific data using modern techniques and standards is essential. This must of course be undertaken within a framework of a well-designed set of scientific and archaeological questions. This will hopefully ensure that scarce funds are efficiently applied and that in the future more money will be available. The second research strand is the continuing development and application of new interpretative methods. It is hard to overplay the revolutionary impact of metallography and cheap quantitative chemical analysis on archaeometallurgy in previous centuries. The more widespread application of, for example, tin isotopic work and synchrotron radiation may have similar era-defining effects in the future.

The collation, confirmation, and reinterpretation of the legacy dataset form a potential third facet of archaeometallurgy’s future. A purely pragmatic viewpoint would stress that financial and conservational constraints demand that we focus on existing datasets. While often true, there is no need to see the use of legacy data as a kind of scientific crisis cannibalism—a choice of last resort. Of course, a new interpretative and theoretical approach would need to be developed for working with legacy data. In the traditional model of interpreting analytical data, the individual artifact’s chemical signature is given primacy. The uncertainties of legacy data require that we now develop a more robust, ‘fuzzy’, methodology. Rather than setting arbitrary but firm thresholds within our numerical data to define the difference between metals from different sources, a much more reasonable approach is to look at averages of archaeological groups. Coherent sets of artifacts can be proposed using the traditional combination of typological or chronological schemes and geological insights. It is important to remember that this is also a very powerful way of interpreting *new* scientific data collected by recent methods. As touched upon above, the precise level of individual chemical elements within an archaeological object is affected by conditions in the smelt, losses through oxidation under melting, heterogeneity caused by differing levels of solubility (e.g., lead will not dissolve into copper), and variability within ore deposits. Even where the quality of the scientific data available is extremely good, the usefulness of going beyond the peculiarities of individual objects and exploring broad archaeological trends is increasingly being recognized.

Archaeometallurgists are becoming ever more sophisticated in using scientific data to engage with wider archaeological problems than merely provenance and ‘compositional industries’. The last 20 years have seen several new interpretative concepts come to the fore such as the artifact biography, material agency, and the interaction of human choice with underlying material properties. Creating a history for an individual artifact obviously requires the integration of a wide range of archaeological datasets. Therefore, some archaeometallurgists stress the importance of applying a wide range of new analytical techniques on a few key artifacts. However, the large database of legacy data is actually an ideal tool to answer many of the modern questions of archaeomaterials, if we use an appropriate scale of analysis.

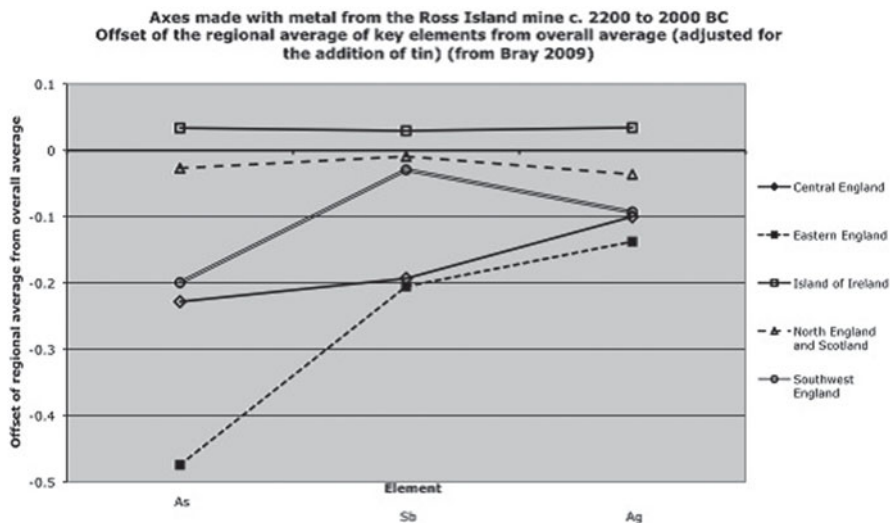


Fig. 10.2 The average level of key diagnostic elements can be a powerful way of interpreting chemical composition data; here showing the slow flow of metal away from an Early Bronze Age source in Ireland. As metal is remelted its composition alters in predictable ways. Eastern England is farthest from the source and, on average, uses copper that has been remelted the greatest number of times

Approximately half of all known Western European Early Bronze Age copper-alloy objects have been analyzed to find their chemical composition. This wealth of data allows averages to skip over individual weaknesses to create a strong regional picture that can be grounded in wider archaeological frameworks. Systematic *averagetrends* of the loss of elements such as arsenic and antimony under heating reveal regional patterns of recycling, curation, alloying, smithing, and exchange. A clear example is the relationship between Ireland and Scotland in the late second millennium BC. The simple fact that arsenic is lost from a melt due to oxidation explains clear trends in the average composition of copper axes from the two regions. Scottish axes consistently show lower average levels of arsenic at this time. This can be simply explained by arguing that Ireland was a center of extractive metallurgy at this time, while melting and casting Irish axes into new forms produced local Scottish axe types. Beyond Scotland there are further losses of arsenic and antimony caused by the reworking of objects and slow movement of the metal away from the ore source at Ross Island, County Kerry (see Fig. 10.2). Focusing on individual objects without creating a broader regional picture misses these important archaeological trends. A preliminary re-interpretation of the archival data from the western European Early Bronze Age using these ideas has been published by Bray and Pollard (in press).

Though only in its preliminary stages, the correct use of legacy data must stand alongside new analytical programs and techniques as a third strand for the future of scientific archaeometallurgy. Work on new data mining techniques and theoretical frameworks is necessary to combat the concern that old datasets are neither scientific

nor archaeological enough. The view must become not ‘they are wrong, what can replace them?’ but instead ‘how wrong are they, how can we incorporate them?’ They are simply too large, useful, and hard-won to ignore.

Summary

Although each case is unique, it is essential to anchor the chemical and isotopic study of metal around a systematic question and methodology. Due to the perennial difficulties of securing access, time, and money, it is important that any scientific intervention is carefully planned. Figure 3 summarizes the wide range of factors that typically affect such planning.

When archaeometallurgists ask themselves ‘What are we trying to learn from the analysis?’, the traditional response is one of the following:

‘*What is it made from?*’ can simply be answered from a qualitative analysis (perhaps by XRF) of the surface without too much cleaning—but is this the best approach? If it is a rare object, and/or something that is unlikely to be available again for analysis, then would it not be better to do a ‘proper job’ and produce a quantitative analysis which will have more lasting value (this assumes that any such analyses will then be properly published!).

‘*Where does it come from?*’ is the provenance question, and requires the usual set of considerations to be taken into account, including:

- Are we dealing with a single object, or a coherent group of objects?
- What do we know about possible ore sources?
- What do we know about mixing/recycling of ores and metals?
- Can we sample the object and get original metal?
- Do we need isotopes or trace elements, or both?
- Do we have access to sufficient comparative data to answer this question?
- How narrowly defined (geographically) does the answer have to be to be useful archaeologically—do we need a particular mine, or a region, or a geological unit?

How was it made is the technology question. If the object(s) can be sampled, then traditionally this is approached through metallography and physical testing (hardness, etc.), but if the samples come from a metalworking site then there may be additional evidence provided by the associated debris (slag, furnace remains, etc.). Again, we may not require a quantitative chemical analysis to answer this particular question, but it must always be worth considering doing it anyway, if only so that the object need not be ‘disturbed’ again. It is very likely that the sample taken for metallography can be used for chemical analysis, so if it is not analyzed at the time (or even if it is), it is vital that the sample taken is properly curated—preferably with the original object.

These three questions are likely to remain as important concerns. However, recent developments in archaeometallurgy have widened the range of our influences and

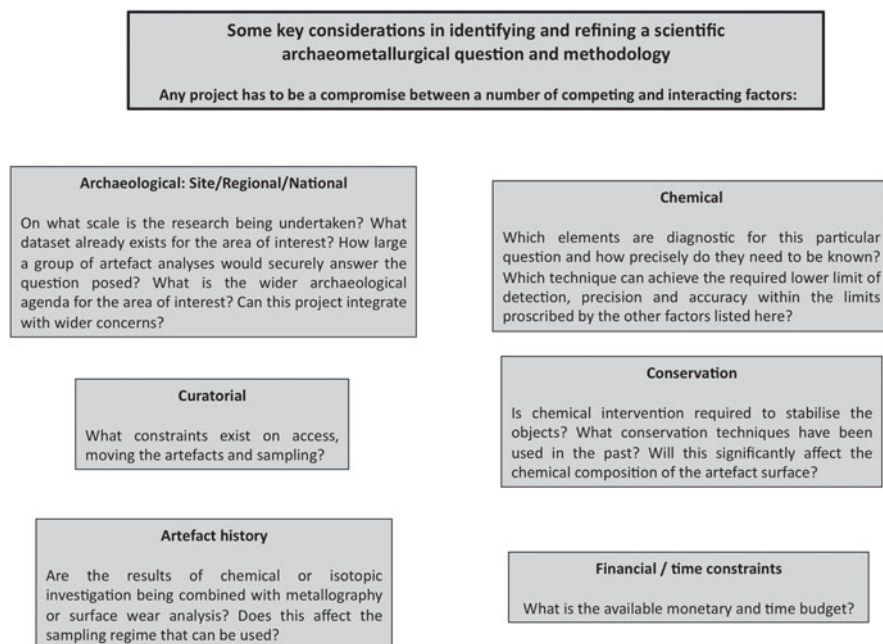


Fig. 10.3 Key considerations when choosing and refining a scientific archaeometallurgical question and methodology

sources (Fig. 10.3). Integration with other specialists must now be the primary concern of archaeologists of every flavor. When posing archaeometallurgical questions, we must consider the wider archaeological questions of that site, region, or time period. This can lead to our datasets being used to tackle social questions that at first seem to have very little connection to the chemistry of metalwork. In addition, it must be remembered that a chemistry-based approach to metallurgy need not include fresh analyses. Free and nondestructive legacy data can lead to radical new archaeological theories being produced from data originally intended for another purpose. In conclusion, we are beginning to recognize that a rigorous, scientific approach to composition, alteration, manufacturing technology, and source can tell us about the objects we excavate and the societies behind them in *equal* measure.

Bibliography

- Bray, P. (2009). Exploring the social basis of technology: Reanalysing regional archaeometric studies of the first copper and tin-bronze use in Great Britain and Ireland. Unpublished D.Phil. Thesis, University of Oxford.
- Bray, P., & Pollard, A. M. (2012). A new interpretative approach to the chemistry of copper-alloy objects: source, recycling and technology. *Antiquity* 86, 853–867.

- Brothwell, D. R., & Pollard, A. M. (Eds.). (2001). *Handbook of Archaeological Sciences*. Chichester: John Wiley and Sons.
- Junghans, S., Sangmeister, E., & Schröder, M. (1960). *Metallanalysen kuperzeitlicher und früh-bronzezeitlicher Bodenfunde aus Europa. Studien zu den Anfängen der Metallurgie 1*. Berlin: Gebr. Mann.
- Junghans, S., Sangmeister, E., & Schröder, M. (1968). *Kupfer und Bronze in der frühen Metallzeit Europas. Studien zu den Anfängen der Metallurgie 2* Berlin: Gebr. Mann
- Pernicka, E. (1998). Whither metal analysis in archaeology? In C. Mordant, M. Peront, & V. Rychner (Eds.), *L'Atelier du Bronzier en Europe du xxe au viiie siècle avant notre ère; vol 1: 259–267*. Paris: Edition du CTHS.
- Pollard, A. M., & Heron, C. (2008). *Archaeological Chemistry*. Cambridge, Royal Society of Chemistry (2nd revised ed).
- Pollard, A. M., Thomas, R. G., Ware, D. P., & Williams, P. A. (1991). Experimental smelting of secondary copper minerals: implications for Early Bronze Age metallurgy in Britain. In E. Pernicka, & G.A. Wagner (Eds.) *Archaeometry '90:127–136*. Basel: Birkhauser Verlag
- Pollard, M., Batt, C., Stern, B., & Young, S. M. M. (2007). *Analytical Chemistry in Archaeology*. Cambridge: Cambridge University Press.
- Rehren, Th., & Pernicka, E. (2008). Coins, artefacts and isotopes – archaeometallurgy and archaeometry. *Archaeometry*, 50(2), 232–248.
- Rohl, B., & Needham, S. (1998). The Circulation of Metal in the British Bronze Age: the Application of Lead Isotope Analysis. Occasional Paper 102, British Museum, London
- Rowlands, M. J. (1971). The archaeological interpretation of prehistoric metalworking. *World Archaeology*, 3, 210–224.
- Shortland, A. J., Freestone, I. C., & Rehren, Th. (Eds.). (2009). *From Mine to Microscope – Advances in the Study of Ancient Technology*. Oxford, Oxbow Books.
- Tylecote, R. F. (1976). *A History of Metallurgy*. Metals Society, London (2nd ed., Institute of Materials, London, 1992).
- Tylecote, R. F. (1987). *The Early History of Metallurgy in Europe*. London: Longman.
- Young, S. M. M., Pollard, A. M., Budd, P. D., & Ixer, R. A. (Eds.). (1999). *Metals in Antiquity. BAR International Series 792*. Oxford: Archaeopress.