Special Issue: Innovation and the Evolution of Human Behavior

Quantifying Technological Innovation

GORDON RUGG
School of Computing and Mathematics, Keele University, Keele, ST5 5BG Staffordshire, UNITED KINGDOM; g.rugg@cs.keele.ac.uk

ABSTRACT

Many aspects of technological innovation can be represented using graph theory. This makes possible both quantitative and qualitative analysis of technologies and of technological innovation. This article describes how graph theory can be applied in this way, with worked examples from lithic technology. The results are analyzed quantitatively, to assess complexity using several measures, and qualitatively, to assess potential interactions between technologies. Innovation can be quantified as the minimum edit distance between the graph for a new technology and the previous situation. The resulting model also is linked to ergonomic measurement and to foraging theory.

The “Innovation and Evolution” workshop was held at the Centre for the Archaeology of Human Origins, University of Southampton, United Kingdom; workshop papers guest edited by Hannah Fluck (University of Southampton; and, Landscape, Planning and Heritage, Hampshire County Council), Katharine MacDonald (Faculty of Archaeology, University of Leiden), and Natalie Uomini (School of Archaeology, Classics and Egyptology, University of Liverpool). This is article #5 of 7.

INTRODUCTION

At an everyday level, “complexity” is a useful concept, and there is general agreement that the archaeological record shows an increase in complexity from, say, an Oldowan core tool to a modern computer. Under closer examination, however, complexity itself starts to become a complex concept, and one which produces some apparent paradoxes. For instance, a clockwork watch consists of many separate parts, but does this mean that it is more complex than a digital watch which consists of only a few separate parts, but whose manufacture requires highly advanced technology? Measuring complexity is clearly a key issue in any attempt to study innovation empirically, and it is equally clear that measuring complexity is a non-trivial problem. This article discusses issues involved in defining complexity; distinguishes between different types of complexity; gives a worked example of operationalizing one type of complexity; and, relates this discussion to the problem of interpretation of the archaeological record.

MEASURING COMPLEXITY

There are various possible measures of complexity, each useful for different purposes. This article describes four main measures, as follows:

- Process complexity
- Cognitive complexity
- Complexity of use
- Complexity of manufacture

The first three of these measures are well described by other authors, and so are described here relatively briefly, to give general context. This article focuses primarily on complexity of manufacture, which has received less attention previously; it also examines the issues of surface complexity and deep structure complexity.

SURFACE COMPLEXITY AND DEEP STRUCTURE COMPLEXITY

One important point from the outset is that complex appearances do not necessarily require complex causes. We can distinguish between the surface complexity of an item (i.e., how complex its outward appearance is) and the deep structure (i.e., causal) complexity of the same item (i.e., how complex the underlying cause was which led to that outward appearance). This is very similar to, but not identical to, the concept from computer science of algorithmic complexity, i.e., the number of instructions needed to produce a given output.

A classic example of this is a fractal image. The complexity of an image can be measured in various ways. For example, one widely used metric in computer science is the compressibility of the image. There are numerous algorithms for tackling this problem. One popular approach is to treat an image as being like a mosaic laid out in regular rows and columns, where the algorithm in essence records how many consecutive squares are of the same color as it works through each row in turn, rather than recording the color of each square individually. For instance, the sequence “red, red, red, red, red, blue, blue, red, red” could be stored more efficiently as “5 red, 2 blue, 2 red.” The more the image can be compressed, the lower its in-

PaleoAnthropology 2011: 154–165. ISSN 1545-0031
doi:10.4207/PA.2011.ART49

formation content. Fractal images typically have very high information content using this approach and can reasonably be described as complex. However, each fractal image is typically derived from a very short equation, which can reasonably be described as simple. The fractal therefore has considerable surface complexity but little deep structure complexity. The concept of algorithmic complexity tends to be used in relation to deterministic systems, where the algorithm produces all the outputs being modeled. This can lead to obvious problems when trying to model chaotic systems, so the term "deep structure complexity" is used here as a way of sidestepping the various theoretical issues which might be raised by simply adopting "algorithmic complexity" unchanged.

An archaeological example is the pattern in a pattern-welded sword blade. Although the shape of the pattern is complex, the complexity has a simple cause—rods of metal are twisted together and then hammer-welded flat. The only extra step in the sword-making process which is required to add this complex pattern is the twisting of the rods; without this single step, the blade would be visually much simpler. This example shows a high degree of surface complexity produced by a process involving comparatively little deep structure complexity. In terms of technological innovation, this example shows comparatively little change in technological complexity. In terms of cultural innovation, however, pattern-welded blades show much more surface complexity than an "ordinary" blade, and would initially be rare, matching two classic requirements for prestige items (conspicuousness and rarity); they might therefore be expected to spread rapidly as a cultural innovation.

PROCESS COMPLEXITY

Process complexity can be illustrated by the example of a bag made out of string consisting of twisted grass. In terms of materials, the bag simply consists of string, which in turn consists of grass, and it can be described as low in complexity. In terms of process, however, the bag might involve tying a series of knots for the main body of the bag, and a series of different knots for the mouth of the bag, with perhaps a drawstring as well, and is likely to require considerable time and skill to manufacture; in this respect, the bag is high in complexity.

There are obvious advantages in being able to represent the processes involved in making an artifact, as well as the materials and tools involved. There are numerous ways in which this can be done. This section briefly describes an example. This issue is covered in detail by Haidle (2007, 2009). In brief, Haidle (ibid.) analyzed the processes involved in making Paleolithic spears from the archaeological record. Although the spears each consisted only of a single piece of wood, and were therefore simple in terms of materials, they showed evidence of a non-trivial series of processes in their manufacture, including heating and straightening. The time taken to do this was also non-trivial. Much of that time was spent waiting for slow processes to happen, so the time on task (i.e., time when the maker was actively doing something to the spear) was limited, but the elapsed time (i.e., time from the start of the manufacture to the end of the manufacture) was considerably longer. Haidle (ibid.) estimated that the total elapsed time taken to manufacture each spear was in the order of several days. There is also the consideration that a tool might be modified during its lifetime; for instance, by sharpening or other forms of reworking (Bousman 1993). This can be represented in various ways, including flowcharts (e.g., Bousman op.cit.)

COGNITIVE COMPLEXITY

When multiple processes are involved, and when there are substantial periods of "waiting time" involved, then there are obvious advantages in having a plan for the overall process. Ling, McGrew, and colleagues at the Department of Biological Anthropology at the University of Cambridge have studied the cognitive complexity and planning sequences required for various Paleolithic manufacturing technologies. They have also studied the cognitive and physiological issues involved in using a given artifact, as described in the following section (Ling pers. comm., 2007).

Even an apparently simple artifact such as a self-wood spear (i.e., a spear made out of a single piece of wood—essentially a sharpened stick) can require considerable cognitive complexity for its production. The issue of elapsed time versus time on task has already been raised—if the elapsed time required is in the order of several days or more, as in the case of the spears examined by Haidle (2007, submitted), then this implies advance planning considerably beyond anything observed in present-day great apes. If seasoned wood is involved, as is the case in traditional manufacture of longbows, then the seasoning time is in the order of years, even before manufacture begins. There is also the issue of sequencing of tasks, where a long sequence of actions may be required as preparations for a later task—a classic example is Levallois flake technique, where dozens of preliminary flake detachments are required before the intended flake is released.

The processes involved in manufacture can be represented diagrammatically via, e.g., planning diagrams and *chaînes opératoires*; a good example is provided by Bleed (2001). The resulting diagrams can look highly complex. It is important, however, to distinguish between the types of cognition which can be used to reach a given goal. One main type involves sequential symbolic reasoning, often linguistically mediated— e.g., "you do this, then you do that, then you do this third thing." Another main type involves parallel processing, often combined with pattern matching, where the learner sees how something is done, and can imitate it accurately, even if they are unable to put it into words (e.g., "it should look like this"). The two types are useful for different types of activity, and what is simple to handle via one type may be difficult or impossible via another. A classic example is recognizing a familiar face, which is normally performed easily using parallel processing and pattern matching; trying to describe that face unambiguously in words with sufficient detail to distinguish it from other faces is extremely difficult.

In addition, highly practiced skills may become compiled, i.e., so habitualized that they no longer require conscious thought. Compiled skills are typically performed much faster and more accurately than non-compiled skills, which can be very useful in the case of motor skills (sometimes known as “muscle memory”). A good overview is provided by Ericsson et al. (2006). An example of this phenomenon in the ethnographic literature is Stout (2002), who describes a classic case of what appears to be compiled skill in stone tool production. The phenomenon is also described in other literatures, often with names such as “unconscious skills” or “tacit skills” and with varying levels of insight into the underlying neurological processes involved (e.g., Baars 1997). The issue of how concepts are mentally organized has received considerable attention in the psychological literature. One early landmark was the introduction of the concept of the “schema” as a way of modeling mental prototypes by Bartlett (1932). An interesting feature identified by Bartlett, and also found by a substantial body of subsequent work, was the tendency for individual cognitive actions to converge towards a prototypical version of that action; for instance, for re-tellings of a story to converge on a simplified, prototypical version of that story. Another landmark publication was Miller’s (1956) classic article on human cognitive limitations, where Miller outlined the concept that human cognition is typically limited to about seven “chunks” of information. A full description of this literature, and of the closely related Artificial Intelligence literature on knowledge representation, is outside the scope of the present article. The issue of convergence on prototypicality has obvious implications for the extensive literature on lithic classification. One recurrent question within that literature is to what extent lithic classification reflects planned templates in the mind of the knapper, versus constraints imposed by the nature or the availability of the lithic materials, or statistical templates in the mind of the lithic researcher (e.g., McPherron, 2006).

A consequence of skill compilation is that if an activity consists of a series of highly compiled skills, then it may look cognitively complex, but may only require a small number of cognitive steps linked by highly habitualized actions which involve little cognitive load. This complicates any attempt to assess the scale of a particular innovation—an apparently large innovation may turn out to be the result of a very small increase in complexity, or vice versa.

A further complicating issue is identifying where the complexity resides within a given task. Some tasks involve materials which require a considerable number of activities in reaction to the material; for instance, knapping low-quality flint typically requires numerous sub-tasks to work around each new flaw that is discovered in the flint. The sub-tasks were not planned in advance, and are reactions to complexity in the flint, but they are numerous, and the knapper needs to retain throughout them the overall aim of the knapping. This issue is well recognized within the lithic research community and elsewhere (e.g., Clarke, 1999). Rather than trying to decide on a single “true” solution to this issue, it makes more sense to use a more extensive set of answers, and to accept that each gives a useful but different insight.

COMPLEXITY OF USE

An issue closely related to complexity of planning is complexity of use of an artifact. A thrusting spear, for instance, requires a set of activities to bring the user sufficiently close to the target. A bow requires simultaneous use of both hands, and a highly practiced grip of bow, bowstring, and arrow. This raises a further issue, namely physiological constraints. Tools such as spears and lithic artifacts need to be grasped in some way in order to be used, which leads into the anatomy of the hand, where different species have different constraints as regards potential tool use.

This introduces a further challenge, namely the well-recognized problem of interpreting the archaeological record. This is particularly problematic for attempts to study processes, since these by definition are actions and therefore ephemeral; those actions leave traces through which the actions can be inferred, but it is easy to make mistaken inferences, and there is always the possibility that a given trace was produced using a different process overlooked by the researcher. This distinction between what might conceivably have happened, what is demonstrably feasible and what actually happened carries through to the topic of the next section, namely complexity of manufacture.

COMPLEXITY OF MANUFACTURE

Various attempts have been made to quantify technological complexity on the basis of the artifacts and materials involved, as opposed to, e.g., the processes involved. For example, percussion techniques can be classified on the basis of how many items are involved, ranging from a score of zero when a nut is hit directly against a tree, to four when hammer and anvil technique are used on a nut, with the anvil being chocked up on wedges to keep it level (Matsuzawa 1996, cited in Marchant and McGrew 2005). Another approach is Oswalt’s concept of the technounit, i.e., the number of components of which an artifact is composed (Oswalt 1976). So, for instance, a self-wood spear would be a single technounit, and a spear with a head held in place by bindings would consist of three technounits (shaft, head, and binding). This concept has been applied both to human tool use (e.g., Bousman 1993) and to non-human primates (e.g., Westergard 1994).

Although both these approaches produce a single, reasonably objective measure of complexity, they are unable to handle various types of complexity which we might want to measure. For instance, two polished flint axeheads might both consist of one technounit, but the first axehead might be made from surface flint, whereas the second is made from flint which has been deliberately mined, and which therefore involves greater complexity in the overall production process.

One way of handling this was proposed by Rugg and McGeorge (1995) who used graph theory to produce diagrams showing the full chain of tools and materials needed to make the tools which were required to make the artifact.

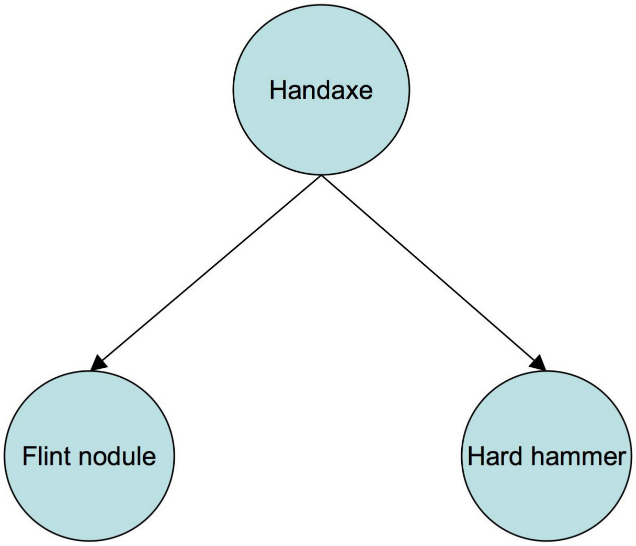

Figure 1. Tool and material requirements for an early handaxe.

For example, an early handaxe can be produced using a hard hammer and a flint nodule. This can be shown diagramatically as follows (Figure 1).

This diagram contains three entities (handaxe, flint nodule and hard hammer) and is two layers deep. The same approach can be extended to handle larger numbers of entities, which is described in more detail below. Although the underlying concept is simple, there are various ways in which it can be applied, which raises the issue of choice of appropriate representation for different aspects of complexity.

REPRESENTATIONS

There are numerous formalisms which can be used to represent aspects of each of the types of complexity described above. Processes, for instance, can be represented using planning diagrams, Gantt charts, or PERT Charts. There is, however, an important distinction between representing complex processes and measuring complexity per se. For example, a Gantt chart shows a representation of which processes are active at which times and a PERT chart shows which tasks feed into which other tasks in what sequence. In both cases, the resulting chart gives a visual indication of the complexity involved, but this is not the same as attempting to quantify it.

These issues can be illustrated by representing the production of a core tool via direct percussion with a hammerstone. One useful formalism from ergonomics is the therblig. Therbligs were invented by Frank and Lillian Gilbreth in the early twentieth century; the name "therblig" is an anagram of "Gilbreth." Therbligs are useful for modeling tasks in ergonomics, particularly workflow analysis, and are still widely used today. A good overview is given by Ferguson, who pulls together the fragmented primary literature (Ferguson 2000). Therbligs are a system of simple pictograms which is used to represent chunks of activity, with each pictogram representing one chunk. Examples of commonly occurring chunks include "pick up" and "visually inspect." A chunk normally involves a set of lower-level activities which can be usefully treated as a single unit for the purpose in question, for instance, the chunk "pick up" subsumes various lower-level activities of arm extension, grasping, etc. Different purposes will have different implications for choice of chunks, for instance, a study of hand positions in knapping would need to treat each hand position as a separate therblig, rather than lumping them all together as "hold." Whether the Therblig chunks map directly onto cognitive chunking in the sense described in the expertise literature (e.g., Ericsson et al. 2006; Simon 1967) is an interesting question, and one which would probably repay detailed investigation.

Each task being modeled may in principle require development of its own therblig system, but in practice the number of chunks required is usually tractably small, and there are many therbligs which occur across a wide range of activities and which therefore do not need to be developed from scratch. For manufacture of early lithic core tools, the therbligs would include "change hand position," "hit with hammerstone," "visually inspect," and "turn core over." Therbligs are usually represented as simple sketch diagrams—for instance, a stylized eye to represent "visually inspect"—to enable faster note-taking, but for clarity the worked examples below will show each therblig in words, within a square bracket.

For a hypothetical Oldowan core tool with minimal retouch on two surfaces, the therblig representation would look as follows:

[visually inspect] [hit] [hit] [hit] [turn core over]
[visually inspect] [hit] [hit] [hit]

When the task is represented in this way, the regularities become obvious, as do the similarities to other human regular structured activities such as song and dance, raising the interesting speculation that an underlying cognitive facility for regular structured activity was co-opted for various purposes.

Adding extra elements to the diagram above illustrates how greater surface complexity can rapidly arise from a comparatively small change in deep structure complexity. We can treat the two lines above as a single process, naming it "{retouch-section}" in braces, to distinguish it from the unitary therbligs in square brackets. If we now add the concept of "change hand position"—i.e., changing position of the hand holding the core so as to expose a new working unretouched area—then we can produce the following diagram.

{retouch-section}
[change hand position]
{retouch-section}

This representation shows that the underlying change in process is small. Only one new entity has been introduced, and one already existing entity, namely {retouch-section}, has been repeated. However, if we unpack this representation into the activities which it represents, then the surface complexity is considerably higher than in the first diagram, as follows.

[visually inspect] [hit] [hit] [hit] [turn over]
[visually inspect] [hit] [hit] [hit]
[rotate]
[visually inspect] [hit] [hit] [hit] [turn over]
[visually inspect] [hit] [hit] [hit]

This example illustrates several useful concepts. One is the concept of chunking—aggregating several entities together into a single concept, which can significantly reduce cognitive load for the person doing the task. A classic example is chess, where chess masters typically conceptualize the arrangement of pieces in a game as a combination of "chunks" of several pieces in a familiar configuration such as a specific, named defense (de Groot 1965; Gobet and Charness 2006). An excellent introduction to this and to related areas of expert behavior is Ericsson et al. (2006). Another useful concept is the distinction between deep structure and surface complexity. A third is the distinction between representation and measurement. The therblig notation above, for instance, makes it easy to measure the number of individual actions performed and the number of higher-level chunks into which they can be organized.

Therbligs are useful for modeling physical activity. The following section focuses on a different type of representation which is useful for representing and measuring complexity, namely graph theory, illustrating its use for modeling and measuring complexity of manufacture for an artifact.

USING GRAPH THEORY TO MODEL COMPLEXITY OF MANUFACTURE FOR AN ARTIFACT

The approach described here is based on the concept of fabricatory depth (Rugg and McGeorge 1995), which applied graph theory to technology with particular reference to materials and tools used in the total production chain. Graph theory is a branch of mathematics originated by Euler (1741), which is now widely used in numerous fields. A classic introduction is provided by Ore (1996). The core concept of graph theory involves points which are joined by lines, as in Figure 2. The points are variously known as vertices or as nodes, and the lines as edges or arcs. In Figure 2, the top node (A) is joined by one arc to node B and by another arc to node C. Node B is not joined directly to node C.

One widely used distinction in graph theory is between undirected and directed graphs. In an undirected graph, the arc between two nodes has the same meaning regardless of the direction in which it is traveled, for instance, the arc between A and B might mean "is a relative of," in which case "A is a relative of B" means the same as "B is a relative of A." In a directed graph, the arc has different meanings depending on the direction in which it is traversed, for instance, "A is a parent of B" has a different meaning from "B is a parent of A." Examples of commonly used directed arcs in practical applications include "is a type of" and "consists of." When the label for an arc consists of more than one word, a common convention is to join the words with underscores.

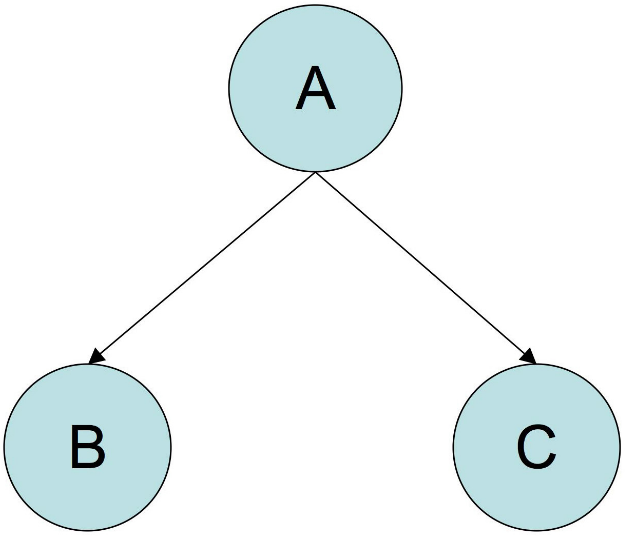

Figure 2. A simple graph.

The parenthood metaphor is widely used in directed graphs, regardless of what the graph is actually representing. In Figure 2, node A would be described as the parent of nodes B and C. Another widely used metaphor in graph theory is the concept of a "leaf node." A leaf node is one which cannot be further subdivided, whether for practical or logical reasons. For instance, a graph showing the components of a machine might show an individual bolt as a leaf node, since the bolt does not consist of any lower-level sub-components. The stage of reaching leaf level is known by various names, such as "instantiation" and "bottoming out."

Graph theory includes metrics for the number of layers involved in a directed graph and the number of different nodes involved, making it possible to measure the complexity of a graph without any subjective assessment being involved. It also is possible to measure the minimum edit distance between two graphs, i.e., the minimum number of changes required to transform one graph into another, which makes it possible to quantify a technological innovation when the previous technology and the innovation are represented as two graphs.

It is often convenient to label the arcs for clarity. This can be done in various ways, as described in Figure 3.

The hypothetical example in Figure 3 uses solid lines to indicate the links involving entities which were definitely used to make an artifact (the flint nodule and the hard hammer) and a dashed line to indicate an entity which is postulated but not proven (a piece of leather to protect the knapper from cuts while knapping). For simplicity, the tools and materials for producing the leather have been omitted from this diagram. Various conventions can be adopted for line types; for instance, different thicknesses of line, or different proportions of dots and dashes in lines, to indicate plausibility of links, or weight of evidence for links. It is also possible to add numeric values to the arcs, as in Figure 4.

The numbers on the arcs can be used to represent a wide variety of measurement, for instance, the example

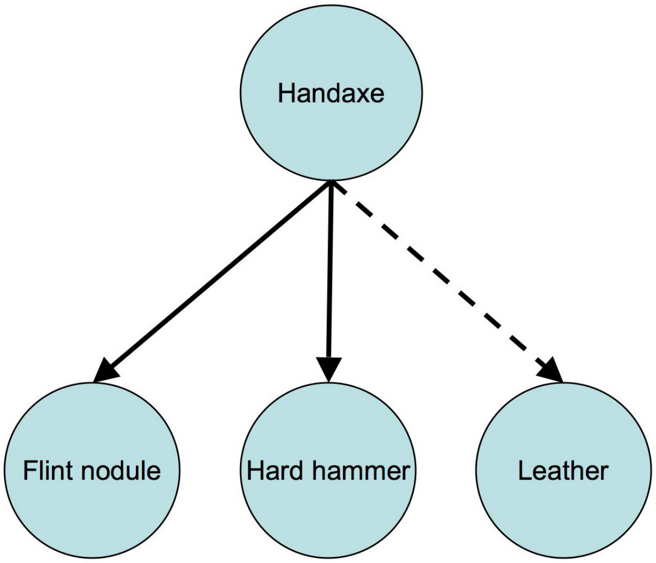

Figure 3. Arc labeling to show attested and postulated entities.

in Figure 4 could illustrate the elapsed time required in a particular ethnographic example to obtain a flint nodule, a hard hammer, and to prepare a piece of untanned leather.

The rest of this section describes the application of graph theory in the manner described above, using worked examples. The convention used here is that the topmost node in the graph represents the artifact whose production is being modeled, and the bottom level of nodes represents the "leaf level" entities being used, i.e., unmodified entities such as stones and antlers naturally shed by deer. (This is different from raw materials, since the term "raw material" is often used to describe something which has already undergone pre-processing, such as copper ore which has been dug out of a mine and ground into powder form.) This approach makes it possible to measure the depth and the width of the graph, i.e., the number of entities involved at a given level, although the latter measure needs to be used with caution in the more complex graphs. The examples are forms of axehead, showing changes in technology through time. As a comparison, they include the handaxe, which was probably not generally used as an axe, but which is used as a convenient example of a related tool which shows minimal fabricatory depth.

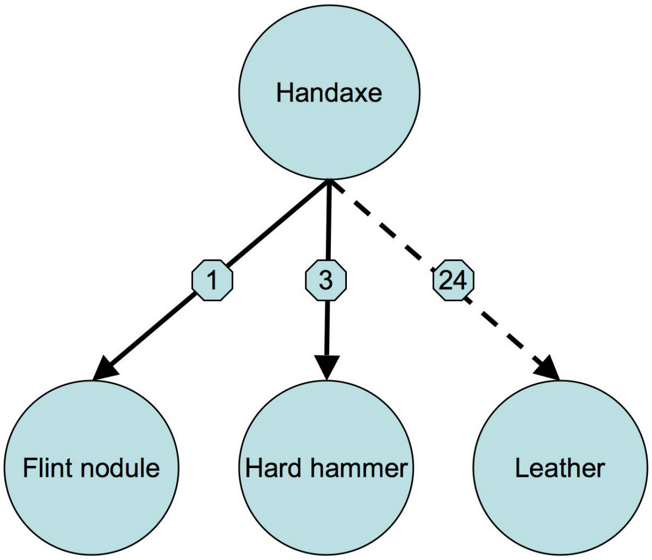

Figure 4. Arcs with numeric labels.

EXAMPLE 1: MANUFACTURING COMPLEXITY IN A HANDAXE

Figure 5 shows a graph for a simple form of handaxe, produced using only hard hammer technique on a flint nodule. The depth of this graph is 2 layers, the maximum width is 2, and the number of entities involved is 3.

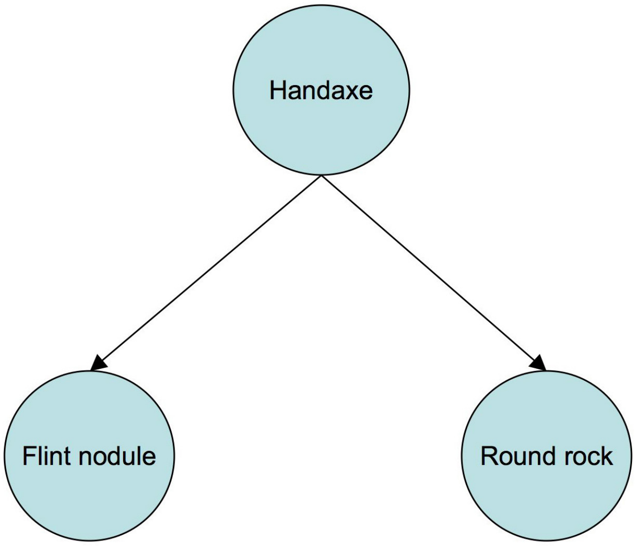

Figure 5. A simple graph showing handaxe manufacture.

The reality of lithic manufacture can be much more complex, however, and can involve activities such as using both soft and hard hammers. The next example demonstrates this for a handaxe produced using both types of hammer.

In Figure 6, a hard hammer is used both directly on the flint nodule and also to modify an antler into a soft hammer, as opposed to the earlier production method of using only a hard hammer. This graph has a maximum depth of 3, a maximum width of 3, and contains 5 entities. Although these two graphs show increasing complexity across time in terms of tools and materials being used, the increased complexity using this metric is small. The situation is likely to be very different if an Oldowan tool is compared with, say, an Acheulian handaxe in terms of complexity of process, illustrating how the use of different measures of complexity can bring out a richer picture of what is going on. If there is a noticeable increase in one form of complexity, but not of other forms, this can generate potentially testable hypoth-

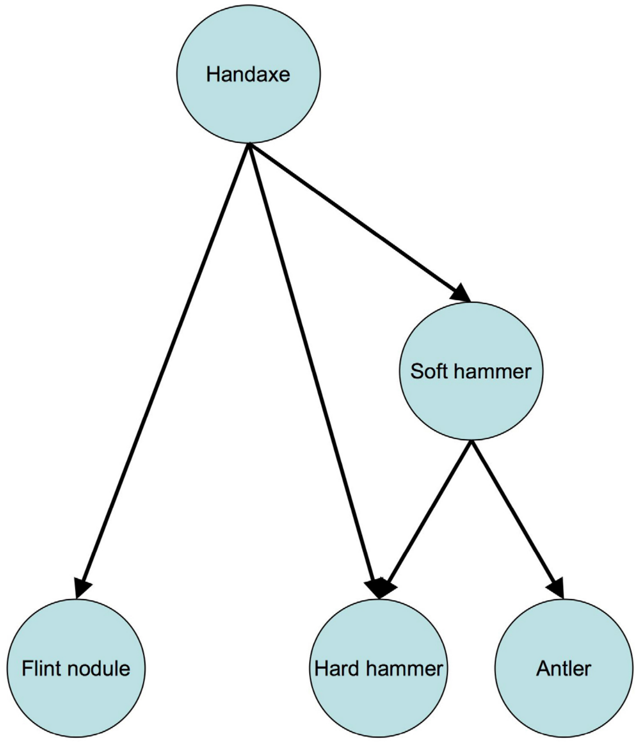

Figure 6. Example 2: A more sophisticated form of handaxe production.

eses about what was happening at a given point in prehistory, for instance, whether a given change is more readily attributable to an increase in cognitive capacity rather than physical dexterity or availability of materials. The relative importance of these factors and of the interactions between them is a well-recognized topic, and is a recurrent theme in the literature on lithic technology (e.g., Roux and Bril 2005).

The situation became more complex with the innovation of polished flint axeheads. These need to be produced from high quality flint to reduce the risk of breakage due to end shock when the axehead strikes a tree, and the high quality flint often was obtained by mining. The next example illustrates this point.

EXAMPLE 2: MINING FLINT

Figure 7 shows the main items required for flint mining. It differs from the previous examples in that the flint nodule is the intended end product (of the mining process), so the flint nodule is shown at the top of the graph. The axe in this example is a tool used in the process of mining the flint nodule, and is shown lower in the graph. This figure is illustrative rather than exhaustive; individual mines will differ in the specific items used and will need to be represented individually. This issue is discussed in more depth below.

Figure 7 is deliberately incomplete, as indicated by the dashed lines, partly because of space constraints, partly because some entities can be produced in more than one way (e.g., baskets and lamps) and would therefore need to be represented differently depending on the production method which was used in a given case. The Figure 7 graph shows a maximum depth of 3 (although some items could be extended to a greater depth) and a maximum width of 6, with a total of 9 entities shown. The omitted parts would roughly double the number of entities involved—antler and hard hammer for the pick, handle, axehead, and binding for the axe, honeysuckle or leather and cutting tool for the rope, ox shoulderblade for the shovel, withies and cutting tool for the basket, chalk and fuel for the lamp. This graph raises several questions, which are discussed below.

One apparently trivial question involves the space required to show a graph. Graphs of this sort often are too big to represent conveniently on a single sheet of A4 or on a standard computer screen. The usual solutions involve breaking the graph into smaller components which can be shown separately, or in the case of online representations, allowing the user to scroll around the graph. Both of these can limit the usefulness of the graph, by preventing the user from seeing it as a whole. There are practical and theoretical advantages in using lower-technology solutions, such as cards stuck to large sheets of paper, so that the user can visualize the entire graph as a whole. Although the graphs can be large, they do terminate, and the systematic, exhaustive process on which they are based is good for detecting key components and tools which might otherwise be overlooked. This issue is closely connected to the concept of craft skills, i.e., skills typically viewed as too low level to merit inclusion in formal education and training (e.g., Brown and McIntyre 1992). In both cases, focusing on the lowest levels of description and categorization can lead to significant new insights.

A second issue is that this graph shows part of a loop, indicated by the dashed circle. The flint nodule will be turned into a polished flint axehead, but mining the nodule requires a ladder, and the ladder is made from a tree trunk worked to shape with a polished flint axe. Loops like this are fairly common. They can become an issue if part of the loop involves a rare entity which cannot easily be replaced; in such cases, loss of that entity could lead to the breakdown of the entire process in which it is used (examples include artifacts where an essential component is obtained via trade and is not obtainable locally; disruption of the trade network would then have serious implications for producing that artifact). In this example, the axehead used to make the ladder could be made using flint found on the surface, so the loop could be reconstituted if it broke down. This issue is an important one for anyone attempting to produce a software model of this approach—causal loops have implications for the software design, which needs to be able to handle the circular referencing involved.

A third issue involves the distinction between how something could have been made, and how it actually was made in a specific case. For example, a lamp could be made out of a stone with a naturally occurring hollow in it, or out of a stone hollowed out using a burin. The material for the basket could be cut using a simple struck flake or with a more sophisticated implement such as a pressure-flaked

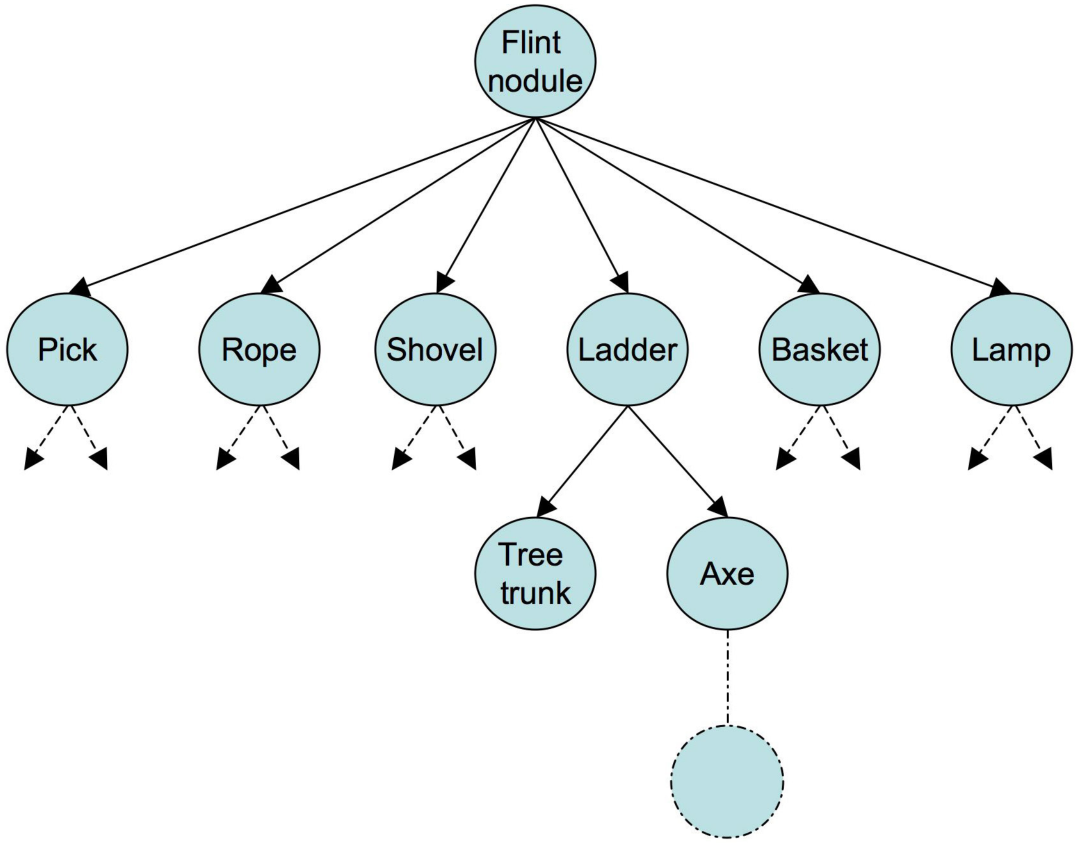

Figure 7. Items required for mining flint.

tool. Options of this sort can be represented using graph coloring. Graph coloring involves representing different types of link with different colors or line types, such as dotted lines, as described earlier. For instance, the graph could use solid lines to show entities known to have been used in a particular case, dashed lines to show entities which were probably used, and dotted lines to show entities which could have been used; it could also show raw entities differently from manufactured tools, or the use of the same tool in different parts of a process (including loops).

When this method is used to represent the manufacture of a copper axehead, there is a considerable increase in the number of tools and materials involved—at least 43, many of which are used more than once in different parts of the overall manufacturing process. A list of tools and materials is included in Appendix 1. The complete graph is too large to be conveniently reproduced here (four sides of A4)—the first layer of tools and materials, for instance, contains copper ore, bellows, tuyere, furnace, crucible, air pipes, tongs, a stone hammer for sprue removal, and three types of abrasive finisher. Most of these lead on to further sub-graphs, for instance, the copper ore, unless surface-mined or collected as native copper, would involve a similar graph to that for mining flint, plus extra nodes and arcs for, e.g., fire and water used to crack hard rocks. The maximum depth of the graph is about five levels, depending on the specific manufacturing process being modeled (e.g., whether a leather bucket or a wooden bucket was used to carry water).

It is possible to plot the various fabricatory metrics against time. Figure 8 shows the number of separate tools and materials used in manufacturing the items described in this chapter, in chronological sequence. The figure could be extended back to the Lower Paleolithic, but would remain constant at a value of 2 tools/materials throughout that time. Plotting values for different tools (e.g., scrapers) and different metrics (e.g., number of levels in the graph for manufacturing complexity) would obviously produce different results, but it is likely that the overall shape of the histogram would remain similar. This metric is useful as an indicator of amount of effort and the number of materials involved. Other metrics are useful for other purposes, for instance, the number of levels in the graph might give some interesting insights into the cognitive complexity involved in producing a given artifact.

The values for the polished flint axehead and copper axehead are conservative; even for these values, it is clear that there is a steep rise in the complexity of manufacture from the Mesolithic through to the Chalcolithic. The energy costs in terms of time for manufacture also rise steeply through this time, from about one or two hours for a tranchet axehead to about two hundred and fifty hours for a copper axehead; the time for a polished flint axehead varies considerably depending on whether surface flint or mined

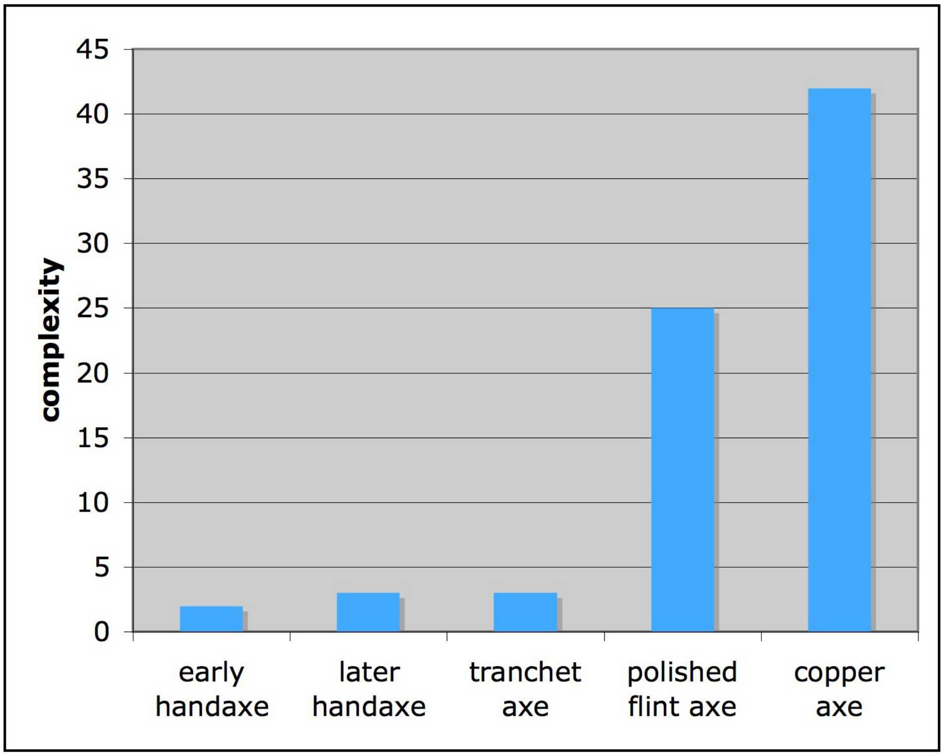

Figure 8. Number of tools and materials involved in producing forms of handaxes and axeheads through time.

flint are used. Another change during that period involves the time on task for manufacture. A tranchet axehead can be produced by one person working alone, and the manufacture can be interrupted and resumed at any point in between other tasks. The copper axehead requires charcoal for fuel, and making charcoal requires several days of constant attention to the burn. Similarly, the melting and casting of the axehead require constant attention, with no time available for other tasks. Figure 9 shows conservative figures for time required to make each of the four types of artifact. It assumes a figure of about an hour to produce an early handaxe, about the same for a later handaxe, and about one or two hours to produce a tranchet axehead. Although the precise figures will vary from case to case, these figures are an order of magnitude lower than the figure for the polished flint axe, and two orders of magnitude lower than the figure for the copper axe.

OTHER ISSUES

The examples above demonstrate how graph theory can be used to represent and quantify complexity. This in turn makes it possible to generate testable hypotheses arising from the representation, as illustrated below.

EXAPTATION

The mining process for polished flint axes clearly involves more entities and more levels than the process of making the axehead itself; the situation is similar with copper axeheads. This leads into the issue of exaptation. Exaptation occurs when an existing situation makes something else possible. In the case of flint mining, for example, most of the technology required for the mining already exists in agricultural societies engaged in large scale earth moving, so no extra inventions are required for them to start mining (ladders and lamps are not required for farming, but are likely to have been invented by this stage). Perhaps significantly, very early farming does not necessarily require major earth moving, but can be carried out via slash and burn agriculture, combined with the use of the digging stick. For a society with no history of defensive or monumental earthworks, mining would require the development of several new entities.

This implies that exaptations facilitating the development of mining were not a necessary consequence of farming *per se*, but were a consequence of the construction of large defensive and ceremonial structures which used the same technologies as heavier-duty farming. A further implication is that because large defensive structures require both large workforces to construct them, and a significant threat to guard against, there was neither means nor need to construct them until the population growth associated with widespread farming occurred. This produces the testable, though speculative, hypothesis that mining would be considerably more likely to occur after warfare and large-scale earthworks were common. This is consistent with the archaeological record in Britain, where there are a few cases of pre-Neolithic mining, but the earliest large scale evidence for mining comes from the Neolithic.

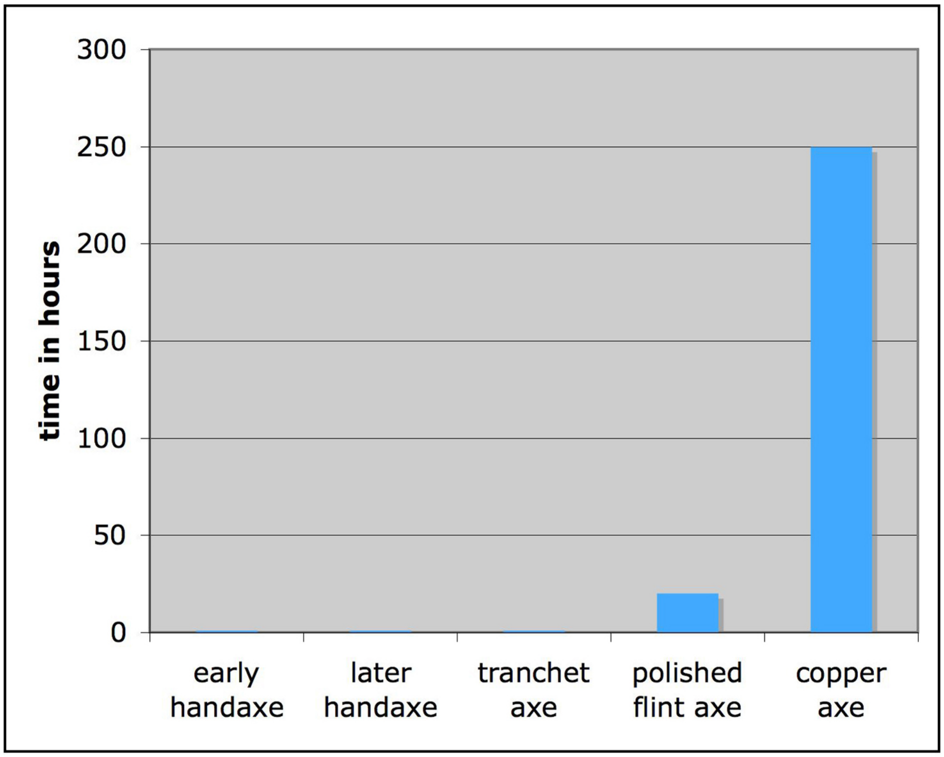

Figure 9. Approximate number of hours required to produce a handaxe, a polished flint axehead, a tranchet axehead, and a copper axehead, assuming use of surface flint for flint artifacts.

LINKS TO OTHER FORMS OF MEASUREMENT

The arcs in a graph can be used to represent ergonomic measures, such as typical energy costs for an activity (e.g., calories required to carry a flint nodule of a given weight for a given distance or to polish an axehead for an hour). This makes it possible to link this representation to standard ergonomics, to game theory (e.g., Maynard Smith [1982]) and also to foraging theory, which assesses the costs and payoffs of different strategies for finding and using resources. This is a well established concept in archaeology and anthropology (e.g., Bousman 1993). So, for example, it is possible to measure the amount of extra effort required to mine flint (with due reference to economies of scale, unit costs, etc.) compared to the effort involved in carrying home flint found during a hunting expedition and then to measure the energy costs of having to replace an artifact which broke because it was made of low-quality flint. Although this could be done without graph theory, using graph theory offers some advantages both in terms of clarity of representation, and also in terms of ability to represent optional routes to achieve a particular goal (for instance, the relative costs of producing an axehead using mined flint versus flint found on the surface).

More speculatively, and moving away from graph theory per se, the energy costs of an artifact could also be compared to the prestige associated with ownership of that artifact. The degree of polishing on some flint axeheads seems to be a good example of this, as the increased energy costs are significant, and this does not seem to be explained by a functional advantage. The energy cost of manufacturing a handaxe or tranchet axe is about two hours; that of a polished flint axe using surface-found flint is about twenty hours, depending on the amount of polish; that of a copper axe is at least two hundred and fifty hours. Using this approach, one can measure the increasing functionality produced by different artifacts (e.g., time taken to cut down a given size of tree, or typical use time until the artifact breaks) against the energy cost of production, to identify the point where diminishing returns occurs; this point can be used as an indicator of where prestige value becomes more significant than functional value.

There are other potential insights from applying graph theory in a variety of ways. For instance, it is possible to check for sub-graphs shared between two or more tasks. An example of this is the use of "pinging" in craft skills, where the craft worker hits the material being worked and listens to the sound it makes; usually a clear, ringing tone indicates that the material is not flawed. This is used in domains as different as flint working (where the flint nodule is pinged), bronze working (where scrap bronze of dubious provenance may be pinged to determine whether it is good enough for re-use), and modern dentistry (where the dentist may ping teeth with a dental probe to listen for cracks in the teeth). Other applications are left to the reader's ingenuity.

CONCLUSION

Graph theory provides a simple but powerful way of representing a technology both quantitatively and qualitatively. This makes it possible to measure the degree of innovation involved in a more complex version of an earlier technology and assess the extent to which a society is pre-adapted for a given new technology. It also enables understanding a culture's technology in context; for instance, in terms of dependencies on raw materials or in terms of activities for which a culture will be pre-adapted. It has clear applications in modeling early societies, but can also be used to model aspects of modern technologies. The same approach can be applied to intangible entities such as the stages of a production process (as opposed to the materials and tools used in it), or societal structures, or to epistemological structures such as layers of explanation (elucidatory depth, as described in Rugg and McGeorge 1995).

At a practical level, once a given entity has been modeled using this approach, that entity's model can be incorporated into future graphs as a "white box," so that other researchers can combine previously produced graphs into large, complex composite graphs fairly easily. This should make it possible to model complex activities such as metal working, where there are advantages in using several separate graphs as components. The resulting models can be used both to generate and to test hypotheses, such as the hypothesis about links between mining and agriculture described in this article. Using this representation also has other advantages, for example, facilitating the use of foraging theory as part of describing a given technology or society. It also offers a simple, easily standardized way of tackling the complex problem of modeling complexity itself.

ACKNOWLEDGMENTS

I am deeply indebted to Miriam Haidle and Victoria Ling for their feedback on earlier drafts of this paper and suggestions on ways of improving it. I am also deeply indebted to the anonymous reviewers and to the editors for their insightful and constructive criticisms.

REFERENCES

Baars, B.J. 1997. *In the Theatre of Consciousness*. Oxford: Oxford University Press.

Bartlett, F.C. 1932. *Remembering: A study in experimental and social psychology*. Cambridge: Cambridge University Press.

Bleed, P. 2001. Trees or chains, links or branches: conceptual alternatives for consideration of stone tool production and other sequential activities. *Journal of Archaeological Method and Theory* 1, 101–127.

Bousman, C.B. 1993. Hunter-gatherer adaptations, economic risk and tool design. *Lithic Technology* 18(1+2), 59–86.

Brown, S., McIntyre, D. 1992. *Making Sense of Teaching (Developing Teachers & Teaching).* Maidenhead: Open University Press.

Clark, A. 1999. Control and intervention in complex adaptive systems: from biology to biogen. www.philosophy.ed.ac.uk/people/clark/pubs/control2.pdf. Accessed 15.11.2009.

de Groot, A.D. 1965. *Thought and Choice in Chess*. Mouton, the Hague.

Euler, L. 1741. (facsimile of original paper). Solutio problematis ad geometriam situs pertinentis. *Commentarii academiae scientiarum Petropolitanae* 8, 128-140. http://math.dartmouth.edu/~euler/docs/originals/E053.pdf. Accessed 28.3.2008.

Ericsson, K.A., Charness, N., Feltovich, P.J., and Hoffman, R.R. (eds.). 2006. *The Cambridge Handbook of Expertise and Expert Performance*. Cambridge: Cambridge University Press.

Ferguson, D. 2000. Therbligs: the key to simplifying work. http://gilbrethnetwork.tripod.com/therbligs.html. Accessed 28.3.2008.

Gober, F. and Charness, N. 2006. Expertise in chess. In: Ericsson, K.A., Charness, N., Feltovich, P.J., and Hoffman, R.R. (eds.), *The Cambridge Handbook of Expertise and Expert Performance*. Cambridge: Cambridge University Press.

Haidle, M.N. 2009. How to think a simple spear? In: de Beaune, S.A., Coolidge, F., and Wynn, T. (eds.), *Cognitive Archaeology and Human Evolution*. Cambridge: Cambridge University Press.

Haidle, M.N. 2007. From brainwave to tradition: how to detect innovations in tool behaviour. Paper presented at the Innovation and Evolution Workshop, Centre for the Archaeology of Human Origins, University of Southampton, April 27-28, 2007.

McPherron, S.P. 2006. What typology can tell us about Acheulian handaxe production. In: Goren-Inbar, N., and Sharon, G. (eds.), *Axe Age: Acheulian Tool-Making from Quarry to Discard*. London: Equinox.

Matsuzawa, T. 1996. Chimpanzee intelligence in nature and captivity; isomorphism of symbol use and tool use. In: McGrew, W.C., Marchant, L.F., and Nishida, T. (eds.), *Great Ape Societies.* Cambridge: Cambridge University Press.

Marchant, L.F. and McGrew, W.C. 2005. Percussive technology: chimpanzee baobab smashing and the evolutionary modelling of hominin knapping. In: Roux, V. and Bril, B. (eds.), *Stone Knapping: the necessary conditions for a uniquely hominin behavior*. Cambridge: McDonald Institute Monographs.

Maynard Smith, J. 1982. *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.

Miller, G.A. 1956. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review* 63, 81–93.

Ore, O. 1996. *Graphs and Their Uses*. Washington: The Mathematical Association of America.

Oswalt, W.H. 1976. *An Anthropological Analysis of Food-Getting Technology*. John Wiley, NewYork. Cited in: Bousman, C.B. 1993. Hunter-gatherer adaptations, economic risk and tool design. *Lithic Technology* 18(1+2), 59–86.

Roux, V. and Bril, B. (eds.). 2005. *Stone Knapping: the necessary conditions for a uniquely hominin behaviour*. Cambridge: McDonald Institute Monographs.

Rugg, G. and McGeorge, P. 1995. Laddering. *Expert Systems*

12(4), 339–346.
Simon, H.A. 1979. *Models of Thought*. New Haven, CT: Yale University Press.
Stout, D. 2002. Skill and cognition in stone tool production: an ethnnographic case study from Irian Jaya. *Current Anthropology* 43: 693–722.
Westergard, G.C. 1994. The subsistence technology of capuchins. *International Journal of Primatology* 15(6), 899–906.

APPENDIX 1

Items required for manufacturing copper axe, based on data from:
Date and author unknown. *Bronze Age living history in the Netherlands*. http://1501bc.com/index_en.html (Accessed 31.3.2008)

mould
clay
sand
stylus
charcoal
basket
spade
knife (flake)
flint nodule
hammerstone
reeds
ox shoulderblade
turf
firewood
fire
hammer
stone head
cobble
binding
rawhide
leather
branch (for handles, etc.)
notch-grinding stone
thread
plant fiber
water
bellows
awl
soft hammer
antler
antler tine
tree (for big bits of wood)
tuyere
high temperature clay
furnace
crucible
air pipes
scraper
tongs
small stone hammer for finishing
hard stone grindstone
sandstone grindstone
clay and sand polisher

TOTAL: 43 items