Chapter 4
A Theoretical Framework for Stigmergetic Reconstruction of Ancient Text

Eugene Ch′ng, Andrew Lewis, Rolf Erlend Gehlken and Sandra I. Woolley

Abstract Cuneiform script, an intellectual breakthrough 5,000 years ago, made recording information possible. Cuneiform is mankind's first ever script, recorded and communicated using clay tablets for thousands of years across the entirety of the Ancient Near East. Remnants of the medium are now stored worldwide in many of collections and time required for the joining of the fragments using traditional methodologies means that the information recorded within these fragments will not be known in our lifetime. The research narrated in this chapter opens up a novel method for reconstructing the fragments, using nature-inspired approaches and new mobile digitising technology. It covers groundwork done to date for supporting a full-scale stigmergy reconstruction of cuneiform tablets and provides hypothetical scenarios within a theoretical framework for testing 'in the wild'.

E. Ch′ng (✉) · A. Lewis
IBM Visual and Spatial Technology Centre, University of Birmingham,
Edgbaston B15 2TT, UK
e-mail: e.chng@bham.ac.uk

A. Lewis
e-mail: AXL148@bham.ac.uk

S. I. Woolley
Electronic Electrical and Computer Engineering, University of Birmingham,
Edgbaston B15 2TT, UK
e-mail: s.i.woolley@bham.ac.uk

R. E. Gehlken
Institut für Archaeologische Wissenschaften, Goethe-Universität,
Gruneburgplatz 1 60629 Frankfurt, AM, Germany
e-mail: gehlken@mailer.uni-marburg.de

E. Ch′ng
Centre for Creative Content and Digital Innovation, Universiti Malaya,
50603 Kuala Lumpur, Malaysia

E. Ch′ng et al. (eds.), *Visual Heritage in the Digital Age*,
Springer Series on Cultural Computing, DOI: 10.1007/978-1-4471-5535-5_4,

Keywords Complexity · Stigmergy · Stigmergetic · Stimuli-response · Cooperative · Collaboration · Cuneiform · Reconstruction · Fragments

4.1 Cuneiform Script: Medium of Communication in the Ancient Near East

Cuneiform script is, according to our knowledge, mankind's first ever script (3,300 BC). This intellectual breakthrough made the recording of textual information possible. The writer uses a reed stylus, the end of which is cut in the shape of a right-angled triangle for impressing the script into the surface of small pieces of clay (so called "cuneiform tablets"). Due to the shape of the stylus a single tablet imprint looks like a small wedge (Latin *cuneus*, therefore cuneiform).

Characters ("letters") are composed of a combination of wedges. Cuneiform was the medium of communication throughout the entirety of the Ancient Near East (an area that stretched from modern Egypt to Iran) and was even used at the beginning of the Christian era.

Current collections of thousands of inscribed fragments from complete tablets are distributed in museums worldwide. Pieces of any one tablet may be located in different museums due to illicit and unprofessional excavations. The largest collections are found in the British Museum, the Louvre (Paris), the Vorderasiatische Museum (Berlin) and the Iraq Museum in Baghdad. Smaller sets of collections in Britain are housed in The Ashmolean Museum in Oxford and also in The Birmingham Museum and Art Gallery. The University of Heidelberg houses a collection of about 2,000 unpublished fragments that belong to a single temple archive from the time of Nebuchadnezzar II (604–562 BC). In addition to those in collections, cuneiform tablets continue to be discovered. For example, in 2009, archaeologists found a cache of tablets dating back to the Iron Age (1200–600 BC) in a 2,700 year old Turkish temple in south eastern Turkey at Tell Tayinat (Bettam 2009).

The majority of the tablets in the collections are catalogued using photographs and hand-drawn copies. In some cases, large collections of tablets numbering thousands remain uncatalogued even after a century of storage. The ancient text on the fragments reflects a wide range of activities encompassing religious, literary, scientific, and administrative documentation including encyclopaedias, dictionaries, political texts, letters and school tablets. However, the fragmentary nature of so many of these tablets means that we can only obtain a very limited understanding of their potential. Hence, the urgent reconstruction of the tablets is a desideratum. Then researchers will be able to gain valuable insights into the social, political, scientific and historical aspects of the ancient cultures represented by the cuneiform tablets. There is also a time pressure on the need for bringing these fragmentary records together. As time progresses, fragments will inevitably gradually disintegrate and the text will be lost forever. Hence, at one level it is critical that these fragments are appropriately recorded to ensure the potential for future research. At another, the full interpretation of these objects will only be facilitated once the fragments are rejoined so that the full texts can be read.

This research proposes a nature-inspired framework for facilitating crowd-sourced, collaborative reconstruction of cuneiform tablets that will allow both the academic community and the public to participate in.

4.2 Traditional Methods for Reconstructing Cuneiform Tablets

The task of reconstructing cuneiform tablets presents a problem for scholars. It is projected that joining these fragments manually by individuals or even groups of researchers locally would take decades to complete. Under the present system of reconstruction, each academic accesses 2D photographic or lithographic representations of tablets from a series of printed catalogues. If a potential match between fragments is discovered, the curator of the museum housing the fragments would be notified and a manual reconstruction would follow. Such methods are often time consuming. Furthermore, 2D pictorial references present a real problem to scholars as 2D images do not fully present the surface information of physical pieces. Even if the reconstruction is conducted locally, the projected amount of time spent could be enormous. For example, the 2,000 unpublished fragments which belong to a single temple archive from the time of Nebuchadnezzar could take nearly two decades to reconstruct; a simple calculation projected that it will take one person 18.255 years working everyday on location, even discounting joining problems that may arise during the reconstruction. The number *c(n, k)* of all combinations of the *k*th class of n elements without repetition is

$$c(n,k) = \binom{n}{k} \tag{4.1}$$

Starting with 2000 pieces ($n = 2000$) of cuneiform fragments, and trying to join every piece to every other piece ($k = 2$) we therefore have the possibilities

$$c(2000,2) = \frac{2000!}{2!(2000-2)!} = 1999000 \tag{4.2}$$

The average speed for an expert is 30 possibilities per hour (at times faster, for example: one will not try to join a letter fragment and an economic text, or a thick and a thin piece. At times slower: there may be problems such as parts missing between the pieces, or tablets are encrusted by salt and not immediately readable). So $1999000/30 = 66633.33$ hours are likely to be needed. A person may not be able to work more than 10 hours per day, thus about 6663 days that is 18 years, are needed $(66633.33/(10\times365) = 18.26)$.

A more rapid solution might be possible through the use of a digital approach and particularly through the use of digitised 3D representations of the objects rather than the physical ones. Previous works on digital cuneiform databases exist, such as the Cuneiform Digital Forensic Project (CDFP) (Arvanitis et al. 2002; Woolley et al. 2001), and The Cuneiform Digital Palaeography Project (CDP 2004) at The University of Birmingham. These earlier data acquisition experiments provide an insight into important technical issues such as graphical resolution and data volume. The resolution required for fine cuneiform details meant that the digitised data volumes could be extremely large. The experience also provided an insight into the need for clear and robust object tagging for any digital archive of this type (Woolley et al. 2002). In addition, the multidisciplinary aspect highlighted the importance of intuitive data representation to both expert and non-expert users across the disciplines. A third project involving the digitisation and visualisation of cuneiform tablets is The Stanford Cuneiform Tablet Visualisation Project (Anderson and Levoy 2002). Together, these related research projects have laid an early foundation which the present research project will build upon. For example, one problem is the stringent forensic resolution requirement of 0.025 mm resolution over the entire surface of a cuneiform tablet set by the CDFP. Such high resolution scan of a single tablet typically uses up a storage space of a single layer DVD (4.7 GB). The size of the digital data makes interactive visualisation impossible, not to say manipulate. Interactive 3D (i3D) reconstruction requires efficiency in processing speed and visualisation; this requires a much smaller storage capacity (or better digital representation). The requirement for efficient storage and processing becomes even higher when collaboration between multiple users is necessary. For collaboration, a synchronous client–server architecture-based tool is needed. Users wishing to collaborate may access an interface on the client-side via the network that displays 3D cuneiform fragments retrieved from the digital archive. The advent of research in Human–Computer Interaction, interactive 3D (i3D) computer graphics and high-speed networks should have brought advances to the community but a survey of literature did not find evidence for 3D collaboration platforms from progress in technology application, or any useful outcome in this area of research. Although concurrent versioning systems and computer supported cooperative work (Eseryel et al. 2002) and cloud-based groupware such as Google Docs are available, they are office tools rather than i3D platforms supporting cooperative reconstructions. However, lessons can be learned from issues related to psychological ownership and perceived quality of work (Blau and Caspi 2009; Dekeyser 2004).

A nature-inspired solution proposed in the present research for the reconstruction of cuneiform fragments using current technology can be structured into three main components. The first component prepares the medium digitally for facilitating the nature-inspired reconstruction, it addresses issues related to the digitisation and cataloguing of 3D cuneiform information in a structured database that allows fast and efficient transmissions over the network. The second seeks to construct an algorithm to filter fragments based on accessible information (size, form, ratio, and orientation) for fitting 3D fragments. The third component formulates a nature-inspired, crowd-sourced and agent-assisted reconstruction that advocates indirect cooperation via a virtual environment that supports i3D interface for the direct manipulation of real-time 3D data.

With regard to the first research component, a survey of the design of structured databases for storing cuneiform tablets that could facilitate the representation to be efficiently sent over the networks, and at the same time provide 3D manipulation, did not find any large-scale projects dealing with 3D data. The Cuneiform Digital Palaeography Project (CDP 2004) is perhaps the most structured that deals with 2D pictorial data. Some work has been done on scanning and coding (Hahn et al. 2006), and visualisation of a small collection of cuneiform tablets (Anderson and Levoy 2002; Woolley et al. 2002), but assisted reconstruction of cuneiform tablets from 3D data has not yet been attempted. A review of literature in other application areas yielded a number of projects that are of relevance to the problems that we may encounter in the reconstruction of cuneiform tablets (e.g. Kampel and Sablatnig 2004; Koller and Levoy 2004; Papaioannou et al. 2002; Zhu et al. 2008). Given the broad geographical grouping of cuneiform fragments, the traditional method of matching can be very slow and even a partial automation of the process would be beneficial. The CDP has been initiated to provide computerised catalogues of cuneiform signs, there has been no attempt to reconstruct fragmented cuneiform tablets using a computerised system.

When investigating the online, collaborative aspects of fragment reconstruction, there appears to be very little evidence of research that is directly applicable to the challenge presented by the fragments from cuneiform tablets. Some insights can be gained from related research (Farella et al. 2002; Bulmer 2002; Löffler 2002), however, the 2D jigsaw puzzling in Bulmer project is significantly different from the 3D mosaic problem associated with cuneiform fragments. The latter two papers also provide insights into considerations of bandwidth, concurrency and security when using distributed system for the visualisation of 3D resources.

4.3 Methodology: Nature-Inspired Approach for the Reconstruction of Cuneiform Tablets

Nature has a way of finding the best solution to difficult problems (Bonabeau and Théraulaz 2008), and those solutions are frequently the product of simple local interactions (i.e. cooperation) amongst entities. Cooperation is important in any large-scale mosaic-based reconstruction, particularly when the rare expertise of scholars plays an important role in the identification and joining of fragments belonging to the same text and period. There are approximately 600 cuneiform-related scholars in the world, including students. For a number of reasons, scholars scattered worldwide do not necessarily interact, communicate, or cooperate well with a common purpose. It seems that groups are frequently disjointed. How do we then facilitate common activities that will promote worldwide collaboration? What if scholars in the field do not want to communicate, can they collaborate without direct communication? Are there proxies that will allow the information worked on by an individual provide clues that will help other individuals reconstruct fragments without direct collaboration? It turns out that insects have been involved in such activities for thousands of years. The word stigmergy describes such activities.

4.3.1 Stigmergy: Coordination Without Direct Communication

Stigmergy was first coined in 1959 by Pierre Paul Grassé (Grassé 1959), who used the word to describe the effect of pre-existing environmental states on the actions of termites building a mound. Grassé noted (in French) that “The coordination of tasks and the regulation of constructions do not depend directly on workers, but the constructions themselves. The worker does not direct his work, he is guided by it. It is the stimulation of this particular type that we give the name of stigmergy,” tanslated by Holland and Melhuish (1999, p.174). Studies in stigmergy are originally associated with social insects, which rely on the stimulus–response mechanism for coordinated behaviour. Stigmergetic interactions can occur in both social species (Bonabeau et al. 1997; Camazine 1991; Goss et al. 1990; Smith 1978) and those that are solitary (Peters 1970). It also occurs in highly intelligent species, such as human beings (Helbing et al. 1997a, b).

The observed collective behaviour of insects in the construction of mounds and nest structures (Bruinsma 1979; Grassé 1959, 1984; Smith 1978) is intriguing. At the individual level, it would appear that each insect works in solitude yet, at the collective level, an apparent work of coordination is observed. In the eye of an observer, it appears that an invisible hand is guiding such work. The mechanism coordinating the construction of structures depends on various concepts well covered in *A Brief History of stigmergy* (Theraulaz and Bonabeau 1999), and is reiterated here for the purposes of formulating a methodology for fragment reconstruction. One of the two concepts (Rabaud 1937) is *interaction*, in which an individual’s behaviour acts as stimulus for modifying the behaviour of another. The other is *interattraction*, in which any animal within a social species is attracted by any other animal of the same species. Furthermore, the behaviour state of insects is influenced by a “critical number of specific stimuli from its nestmates” called a group effect. Each individual is a source of stimuli for other individuals, this stimuli-response “opens the way for an indirect coordination of individual activities” (Theraulaz and Bonabeau 1999). If the presence of an individual is a stimulus in itself, and so also the physical and chemical structures that these insects collate and leave behind, then there is not the need for direct coordination.

Two examples will make the concept clear (Fig. 4.1). The wasp *Paralastor* sp. (Smith 1978) constructs nest based on a stimulus–response mechanism where the completion of a stage of construction provides a stimulus for the construction of the next stage, until the nest is completed. In pheromone-based stigmery, termites *Termitoidae* (Grassé 1959) use soil pellets covered with pheromone for building pillars, first depositing at an uncoordinated random location. Chance deposition of a critical pile of pellets provides a positive pheromone feedback for more activities at the location, therefore increasing the likelihood of coordinated construction of pillars. This process is termed quantitative stigmergy (Theraulaz and Bonabeau 1999).

The use of stigmergetic algorithms for problem solving is not new. Stigmergy has mainly been applied to optimisation problems, e.g. (Bonabeau and Théraulaz 2008; Holland and Melhuish 1999). However, there is no literature relating to the use of this approach within the context of large multi-threaded problem solving.

Stigmergy can possibly form the cornerstone of a nature-inspired and crowd-sourced reconstruction of cuneiform tablets, and it is easy to see the potential that stigmergy communication will have in a large, multi-threaded problem solving environment. A stigmergy reconstruction system does not need a single, central intelligence algorithm for the reconstruction of fragments. Individual reconstruction agents, an autonomous software unit that mimics insects in this context, can cooperate without direct communication, within a system at radically different speeds—so that humans and non-human agents can work together to solve a problem without direct communication. Later, we will formulate a methodology based on stigmergy for the reconstruction of cuneiform tablets.

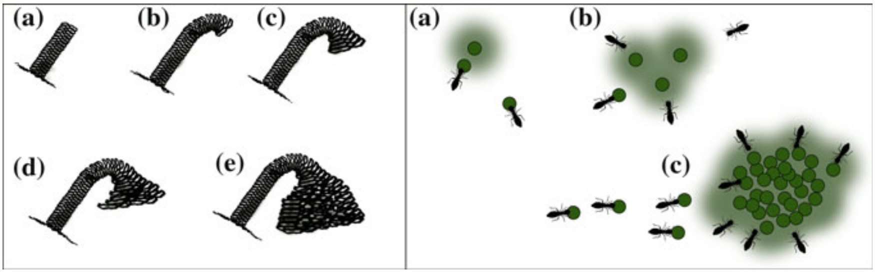

Fig. 4.1 Two examples of stigmergy in collective insect behaviour. The wasp *Paralastor* sp. (*left*) and pheromone-based stigmergy of termites *Termitoidae* (*right*). In the wasp nest construction mechanism, each stage of the construction process (**a**–**e**) provides a stimulus for another stage of construction. For the pillar building termites, a chance deposition of pheromone infused soil pellets **a** provides positive feedback for more activities in that location **b**, with subsequent positive feedback loops **c** until the pillar is completed

4.3.2 A Client–Server Framework Supporting Stigmergetic Interaction

A major requirement for nature-inspired reconstruction is a sufficiently powerful distributed multiprocessing agent framework. Given that a large number of agents may be operating at any given time, it would be a mistake to limit the processing engine to a single processor or workstation. Consequently, a framework has been developed to support multiple artificial agents, alongside a database of fragment geometry, and a socket.IO-based server that allows for real-time Web-based communication between different subsections of the reconstruction system (Fig. 4.2). The framework server receives and broadcasts messages from agents and users via JavaScript Object Notation (JSON), utilising the socket.IO specification to ensure that connected systems are kept in sync. The server holds a virtual representation of a 2D plane, which is used to sort and manipulate fragments. Agents may jump fragments from one position to another within this virtual space (or ‘grid’), move a fragment on-screen, or rotate a tile in space. Any interaction or manipulation of fragments is ultimately processed and recorded by the server. In this way the messaging system and the server simulate what might be called the virtual law, an equivalent of the physical laws of the agent’s universe. The server controls the passage of time and information between agents, and any rules implemented at the server level are immutable. Server rules affect every agent within the system in a way that cannot be avoided by any agent.

A simple example of a server-side or “universal” rule is the implementation of “caution” within the system. All agents within the system will avoid interacting with fragments that have recently been manipulated by a more dominant agent (e.g. a human). The sense of dominance and the reaction of an agent after it is

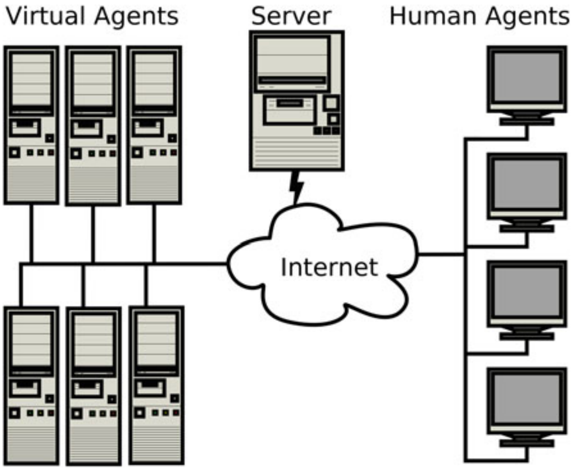

Fig. 4.2 A human agent cooperative client–server architecture

"frightened away" from a fragment can be controlled at the client level, but the overwhelming instinct to avoid the fragment is compulsory.

The socket server also acts as a rudimentary Web-server, transmitting the front end interface and any fragment model data to clients. Combining the file and socket servers onto a single system is both convenient and necessary for the proper operation of the system, as some Web browsers and brands of antivirus software wrongly classify the normal activities of the application framework as a potential cross-site scripting attack when resources were hosted on multiple servers.

All clients within the system are built using an agent framework. The framework is based in Python, and support for parallel processing as standard. Threading and FIFO queue systems have been implemented within the system to facilitate synchronous and asynchronous support for socket communications.

4.3.3 Agents as Stigmergetic Virtual Insects

A basic 'agent' class (a software template) is provided that supports a Life-Timer (a timer within agents that decides the longevity of its lifetime) and also a throttle to control execution speed. New agents can be automatically spawned on an appropriate local or distributed processing node, and can be given the power to spawn new agents if they require it. The life-timer of an agent can be adjusted at any point, extending the life of successful agents, and reducing the life expectancy of agents that make poor choices.

Agents can be created and added to the framework at any point, and may interact using one of two methods. The *Network Method* presents a location independent method of access, where the normal bounds of an environment are removed. An agent accessing the framework in *Network Mode* can access information about any fragment, rotate, or move any fragment to any position in a single tick. Agents within the *Network Space* have a complete overview of the fragments, abstracted from their physical location. They rely on metadata to make decisions about what to do with an individual fragment.

In contrast to this, the second method of access within the framework is the *Grid mode* (Fig. 4.3). Agents accessing fragments using the Grid mode are bound by the order and location of fragments within a 2D grid, and their knowledge of the fragment database is limited to fragments within adjacent grid positions. Movement of fragments is also limited to adjacent grid squares, effectively throttling the maximum speed at which fragments can be moved. Agents within the Grid make their decisions based on the best local solution to a problem, with the intention that multiple, simple agents can generate complex local structures and hierarchies that network-based agents (including human users) can exploit to improve the matching process.

Possibly the simplest example of a network layer agent is a *Lay-Flat Agent*. The Lay-Flat Agent visits all fragments in turn, and rotates them so that they lay flat on the surface, for example, with the text surface pointing face upwards. Once this

Fig. 4.3 Agents A1 and A2 are *Grid Agents*, and can only see fragments F1 and F2 on the grid. Agent A3 is a *Network Agent*, and can see all fragments including F3, which is not visible on the grid

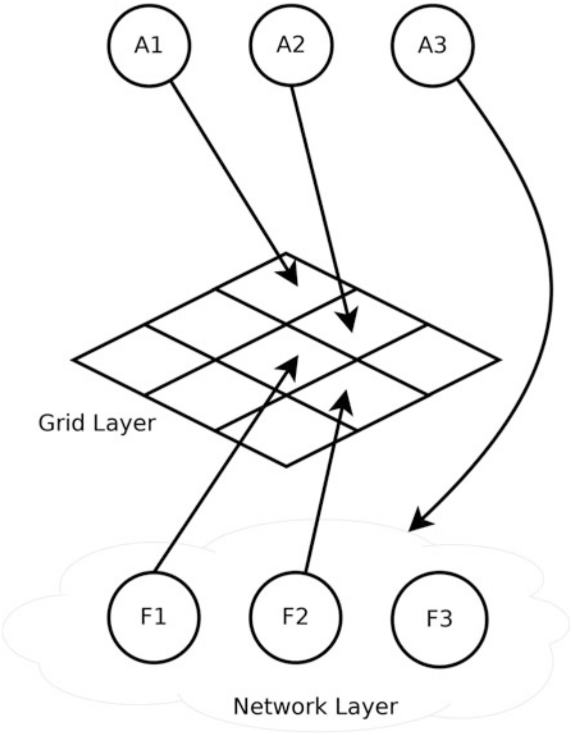

agent has visited a fragment once, it leaves a non-degrading tag in the fragment's metadata and does not revisit the fragment unless the tag is removed.

A slightly more complex example of a network agent is a *Clustering Agent*, which takes fragments with similar metadata aspects (such as dig-site, text subject, or interaction history) and attempts to group them together in a particular area of the grid. The clustering agent will not interact with a fragment unless it has been left idle for some time, with no other human or agent interaction occurring. It is worth noting that this agent operates at the network level, but has a direct effect in the Grid space.

A Grid Agent operates in a manner akin to termites or ants. The actions of Grid Agents are defined by relatively simple rules that affect the environment of the grid, which in turn guide the actions of other grid agents.

A grid agent may be programmed with the simple rule set that swaps the position of fragments so that the total size of juxtaposing fragments is as close to the average size of a complete cuneiform tablet as is possible. If the size is close to the ideal, the agent's life-timer or 'energy level' will be extended. Once the agent has swapped the fragments and gained energy, it leaves a tag on the fragments and hops to an adjacent grid location in search of previously unmodified fragments. If there are no fragments to change, the agent will continue to hop from location to location, searching for fragments until its life-timer reaches zero. Conversely, if the agent is repeatedly successful, its life-timer reaches a pre-determined level, and the agent will duplicate itself, and divide its life-timer value between both agents.

The action of the swapping agent described above is very simple, in some ways akin to Cellular Automata (Wolfram 1986), but the alterations that the agent causes within the grid can lead to a very complex interaction with other agents. The simple act of grouping fragments by size can improve or worsen the ability of other agents to match fragments together, affecting their life-timer and therefore their prevalence within the system.

4.3.4 Special Agents

There are also a number of special cases where an agent may be created with unusual characteristics. In the case of a *Human Agent*, it is not beneficial to use the life-timer of the user in a normal manner. Users will typically work in environments with the potential of being affected by real world distractions, which could result in a premature end to an online session if the user is away for too long. So a human interacting with the system will remain visible as an agent until the user disconnects from the server by closing their session. Also, the association of a particular user with successful matches (assessed through matching of the user's tag in fragment metadata) could increase the reputation of the user within the system and could, for example, allow them to override the actions of less experienced users in the event of a real-time fragment locking issue.

Another instance where a special class of agent may be needed is the initial seeding and maintenance of the system. A *Super-Agent* can be used to act as a caretaker to the framework, setting up the initial conditions of the Grid and Network layers when the system is initialised or restarted. A similar agent may be used to track and record the growth and die-off of particular agent classes, and attempt to respawn or cull them if their numbers become radically unbalanced as the result of a short-term issue within the framework (such as deliberate human intervention). The *Caretaker Agent* could also contribute to the overall entropy of the system, shuffling idle fragments or affecting the behaviour of an agent contrary to its normal operation.

Finally, a *Glue Agent* could be used to hold the position of two fragments relative to each other, effectively bonding them into a single fragment. Similar to User and Caretaker Agents, the Glue Agent should have an infinite life-timer, hunt for fragments marked as matching, and bond them if their proximity tags have reached a particular level of potency. Upon bonding, the Glue Agent tags the fragments with a unique identifier, locks their relative positions, and co-locates them on the Grid layer. The Glue Agent then spawns a duplicate of itself to continue searching for more fragments. A Glue agent may bond more than two fragments together, and will release a fragment if the unique identifier which applies to the fragment is removed. If a Glue Agent has less than two connections, it sets its life-timer to 0 and then dies. This prevents superfluous agents from cluttering the system if incorrect bonds are formed.

4.3.5 An Example of Agent Creation

The following example creates a new agent that prints out its life expectancy, and then dies after three ticks. Note that in a multiprocessing environment, the text output would be suppressed, but in a single instance, the text will be output to the interpreter. The start() and stop() functions of the Agent class start and stop the execution of the agent. An agent that has been stopped will not age (its lifetime will not reduce) until it is started again.

```
from agent import Agent
class My_Agent(Agent):
      def main(self):
         print "I have ", self.LIFE, " life left"

a = My_Agent(life=3)
a.start()
```

4.3.6 An Example of Communication with the Server

All agents created with the agent class can communicate with the Node.js server from within the main() function by using self.ul.send() and self.ul.recv() functions. It is also possible to access the server without an Agent, using the Uplink class as shown below.

```
from uplink import Uplink
ul = Uplink(hn = "localhost", pn = 8000, pt = 10)
ul.connect()

for i in range(100):
   ul.send("rotate","fragment.js",1.2,1.2,1.2)
   incoming = ul.recv()
   if incoming:
      print incoming
```

The ul.send() and ul.recv() functions are tailored to the transmission of positional data, with the send() function accepting an action, fragment id, and x, y, z data. Differently formatted messages can also be transmitted and received using the ul.sendraw() and ul.recvraw() functions.

4.4 Methods Supporting the Nature-Inspired Framework

Designing a nature-inspired framework for research of this particular nature requires supporting methodology. These methods are photogrammetry, light weight 3D scanning of cuneiform fragments, a Web-based user interface for supporting an environment for hosting human agents, and 3D printing.

4.4.1 Photogrammetry

One of the most useful clues to the reconstruction of any object is a good understanding of its finished appearance. In order to understand the general shape of cuneiform tablets, a photogrammetric analysis of approximately 8,000 complete tablets from the Cuneiform Digital Library Initiative (CDLI) database has been made. This detailed analysis of images was a necessary step towards the reconstruction of cuneiform tablets, since many text records (both within the CDLI database and in the wider community) lack data on the physical size and shape of complete tablets and fragments.

By exploiting a priori knowledge of the CDLI image scanning system, it was possible to extract the scanning size in DPI (dots per inch) of images within the database by using the Exchangeable Image Format (EXIF) data within the stored images. A simple script was used to perform luminance threshold analysis of the image data, to extrapolate the dimensions of each tablet automatically. When checked against known good data from the CDLI database, the output of the photogrammetric analysis program was found to be accurate within approximately one millimetre of the recorded values. The study has revealed some interesting physical features and characteristics that could help increase the speed and efficacy of cuneiform fragment matching algorithms (Lewis and Ch'ng 2012).

4.5 3D Scanning

The impartial and accurate recording of physical features of cuneiform tablets is not a new problem. The issues associated with the recording and visualisation of tablets has been known since the 1800s. In 1864, the Quarterly Review (Murray 1864) in an article noted that:

> It is a boon of enormous value to be able in any instance to eliminate that fruitful source of error, the fallibility of the observer. Photography is never imaginative, and is never in any danger of arranging its records by the light of a preconceived theory.

To the modern recorder with a 3D scanner, the most 'fruitful source of error' (Murray 1864) is the subsampling and smoothing of data. From an archaeological standpoint, it makes sense to capture as much detail as possible when scanning a cuneiform fragment, for even though the process of 3D scanning is not normally destructive, it is possible that a fragment may not be available for examination in the future. In the absence of the actual fragment, unedited scan data could be of unexpected importance to future researchers. However, the accurate recording of small 3D features is still on the edge of our capability, and expensive equipment is frequently needed to scan objects at high detail. At the present time, the initial cost of scanning equipment, rather than the cost of reproduction, presents itself as a modern day analogue for the photographic dilemma faced by the nineteenth century curators of the British Museum. Some objects lend themselves to 3D scanning, while others can present a considerable problem to certain types of 3D scanner. Reflective, absorptive, and dispersive surfaces present a particular problem for 3D laser scanners. The shortcomings of laser scanning technology which are well reported by Hahn (Hahn et al. 2007) and Chen (Chen 2008) have also been anecdotally expressed by those familiar with 3D laser scanning in the field. There are many cases where coherent laser light is affected by an object in such a way that it cannot be detected by the optical sensors, and scans of dark or reflective objects are frequently corrupt and unusable.

Cuneiform fragments require multiple scans to reproduce them in their entirety. Each scan is taken from a different angle so that all surfaces of the tablet can be recovered in 3D. It is not unusual for the surface of a tablet to have a firing pattern that varies in tone from black to light grey or orange. This deviation in surface colour can leave significant holes in captured data, and make matching together individual scans very difficult. When unusable scan data is encountered, a human operator must decide whether to leave the hole in the data, attempt to rescan the area with different settings or let the computer fill in the hole using specialised algorithms. The traditional method used to overcome the problem of missing data in industrial models involves the application of a non-reactive powder or water soluble spray to mask difficult areas. Within the context of the heritage community, the application of such a chemical masking agent could irreparably harm a unique artefact. In order to scan difficult objects, an alternative to traditional laser scanning must be used.

A popular avenue for the capture of cuneiform fragments is Reflectance Transformation Imaging (RTI). RTI scanners usually employ an SLR camera in a hemispherical rig that is populated by fixed lights with a known position. Triggering the lights in sequence and photographing the effect on the fragment generates an interactive lighting model that can be adjusted to view the surface of a fragment. Although RTI was primarily envisioned as a 2D process, RTI images can be used to generate a 3D surface map. The principal issues associated with using RTI for 3D model generation are described by Macdonald (Macdonald 2011), and it is shown that while the dome scanner presents an excellent method for visualisation of fragments with variable lighting conditions, the quality of the 3D model generated from RTI images is less accurate than other methods.

Autodesk’s 123D-Catch is also a popular contender for 3D scanning without lasers, but practical experimentation with small objects has shown that the time taken to capture a fragment and the low level of detail captured by the system prohibit the widespread use of this capture method for cuneiform fragments. Other examinations of the 123D-Catch system have shown that the resulting models are of lower accuracy than other methods, and are prone to distortion in some circumstances (Chandler and Fryer 2011).

4.5.1 Web-Based Collaborative Virtual Environment

To fulfil the needs of end users, a front end system has been developed that comprises a browser-based interface for the manipulation of multiple fragments (Fig. 4.4). The browser-based front end runs primarily in JavaScript, and interacts with a database and with a node.js server. Although the hardware requirements of the interface are higher than most websites, the interface, when tested, performs as expected on a computer or laptop of moderate power, with a normal connection to the Internet. Data transactions are minimised by local caching, and positional information is transferred by pulling information from the server when required.

A user may interact with the system by selecting fragments from the database using a catalogue. Since the database is hosted in SQL, it is possible to sort and filter the data in a variety of ways. Chosen fragments appear along the bottom of the in-browser 3D visualisation, so that they can be added to the work area one at a

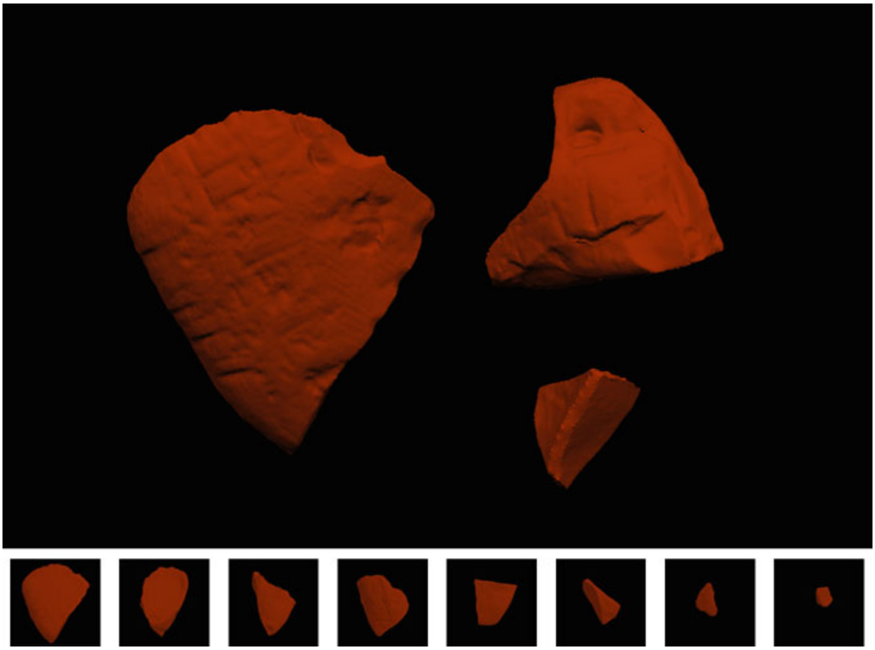

Fig. 4.4 The user interfaces supporting the Web-based collaborative virtual environment

time. As a user manipulates the fragments on-screen, the position and rotation of the pieces are transmitted to the server, and the position of on-screen fragments is checked against the server-side record every 100 ms. This means that multiple users can access the same fragment or collection of fragments, and the result of the manipulation can be viewed in real-time. In addition to standard information about position and rotation, each fragment within the framework supports the storage of metadata. The Web interface can read and manipulate the metadata as needed, supporting a diverse range of functionality including ‘scent’ or ‘pheromone’ trails that record user interactions, the last proximity of other fragments, potential matches and tags applied by different users.

As an example, an individual user may adjust the orientation and position of two fragments within the Web interface, concluding that the two fragments fit together. The user does not tag the fragments with any metadata, but keeps them close together on the screen. The proximity of the two fragments to each other creates a ‘scent tag’, a virtual insect pheromone trail that links the two fragments temporarily, but slowly degrades over time.

A second user may access one of the fragments at a later date and see from the ‘scent tag’ that it has been placed next to another fragment that seems to be part of the same tablet. The second user may then also choose to keep those two fragments close together on-screen, and possibly even tag them as a potential match. These actions will both increase the potency of the ‘scent tag’ linking these two fragments together. Once a critical level is reached, a specialist human or Artificial Intelligence (AI) agent will become aware of the potential match, and may bond the two fragments together permanently.

The metadata format used by the system is sufficiently flexible that new metadata fields can be added without affecting the operation of other agents, so that the evolution of the system agents and the web interface is possible without disruption to the system as a whole. No preference is given in the interface for cooperative or competitive work strategies, and a locking strategy is employed that takes the importance of an individual user and the time of the last interaction into account.

4.6 3D Printing

Manual reconstruction of cuneiform fragments allows for the direct manipulation of fragments in a 3D space, and it is reported by experts in cuneiform reconstruction that the tactility of a physical fit is an important feedback system in manual cuneiform reconstruction. The haptic feedback experienced when reassembling fragments of stone or clay often provides a strong indication of whether two fragments fit together.

In a virtual reconstruction system, however, the user is abstracted from the direct interactions of the physical world by a keyboard or touch screen, and manipulation of fragments by remote control is clumsy in comparison to direct manual reconstruction.

Modern domestic 3D printing technology makes it possible to reproduce cuneiform tablets using plastic or resin, and restore the haptic feedback of a real object to digitally stored fragments. The physicality of the printed object acts as a real world interface to the digital objects, communicating the effectiveness of a match between two fragments in a non-verbal way.

An additional benefit of 3D printing technology is the ability to place physical copies of fragments with scholars who may previously have had access to photographic resources. Also, fragments may be rescaled or manipulated before printing to allow for more accessible study or demonstration of a particular text.

Broadly speaking, modern 3D printers have the potential to provide each and every student of cuneiform studies (and every member of the general public) with complete, real world access to the entire body of virtual cuneiform source materials. 3D-printed parts are not precious or fragile, and so the process of sorting and matching fragments could be carried out by untrained hands prone to hazardous mistakes. Even museum visitors could be encouraged to try their hand at matching fragments together.

Until recently, powder deposition and UV resin printers provided the only practicable option for the reproduction of cuneiform tablets. However, the introduction of inexpensive plastic deposition printers such as the Makerbot Replicator or the UP! Personal Portable 3D Printer has provided a much more cost effective avenue for object printing. Despite the lower resolution of these systems, The VISTA-CR project and other projects like the Cornell Creative Machines Lab have both shown that the reproduction of 3D cuneiform fragments is possible in a variety of materials (principally PLA and ABS plastic) with sub-100 micron accuracy.

The quality of reproduction using deposition printing is generally good, but care must be taken when attempting to match fragments together printed using different types of plastic. The colour and type of plastic used can have a marked effect on the level of material shrinkage encountered as the fragment cools. For PLA plastic the effect is negligible, but for ABS, the final print may shrink by as much as 3 %.

4.7 Hypothetical Scenarios of Stigmergetic Interaction

This section attempts to formalise various hypothetical scenarios where users and agents may coordinate using the stigmergetic framework presented in this article.

4.7.1 Single User

A reconstruction system with only a single human user presents itself as an elaborate aide-memoire. The facility to tag and statically position fragments within the system endows the user with an augmented capacity for information organisation and recollection. Fragments are abstracted from the real world, and can be viewed and manipulated digitally without the special handling and security requirements associated with the original artefacts. The system becomes a workspace for linear examination and reconstruction that parallels traditional methods.

4.7.2 Two Users and Above

With the introduction of multiple human users to the system, additional variables could increase the complexity of the interactions dramatically. Even one additional human user adds permutations to the scenarios in the reconstruction process. For two users working with the same fragment, it is likely that the following aspects of the working relationship will affect the outcome of the reconstruction process.

Method of Encounter
An encounter within the system may be either accidental or deliberate. In the case of a deliberate encounter, the parties involved will likely have shared goals, such as the reconstruction of a particular fragment or collection. It is reasonable to assume that users working in the way will be more likely to work collaboratively, and may have developed a shared strategy to facilitate the achievement of their goals. In accidental encounters, users may potentially have antagonistic goals, and have dissimilar systems. In these cases, a series of activities is more likely to be the result of stigmergy, since no other form of communication will exist.

Strategy
A user may be either cooperative or antagonistic to the goals of another user. Antagonistic behaviour may not be deliberate, but may be the product of competing goals or strategies. As an example of this, consider the case of a tablet that is broken into three pieces. Two users already have a piece of the complete tablet, and both users attempt to manipulate the third piece (which fits between the other two pieces) so that it will match with the fragment that they already have. One or both of the users may behave antagonistically to reach their goal, constantly moving the piece back when the other user reorients it to suit their fragment's orientation. Alternatively, they may work cooperatively, by first allowing one user to move the disputed piece into alignment, and then moving their own piece to match the alignment of the other two pieces. If the users work cooperatively, both will achieve their goal, and their reputation will increase as a result of the successful match. Antagonism will result in either the match not being made, or a third party stepping in and taking credit for the match.

Timescale
Stigmergic communication does not require real-time interaction, but neither does it preclude it. Real-time non-verbal communications can operate independently of stigmergy, and subtle visual cues like shaking a fragment on the screen to draw

another user's attention to it, or moving a fragment away from a user can be used to enrich the matching process. Alternatively, users may rely solely on stigmergic communication to work with fragments when actions have a gap of hours, days or even months between them.

Geography

There are obvious real-time communication opportunities for physically and electronically co-located users, but there are also subtler reasons that geographical location is important to user–user interaction. The physical and social environment of the users has the potential to alter the interaction style of the user at a very basic level. Language and real time gestures common in one location may be completely different in another location, and at this level, it is easy to see how non-real-time stigmergic communication can be more useful than direct communication.

Dominance

As users interact with the system they gain reputation, which increases their effective dominance over other users. In situations where a deadlock may occur, the system will resolve in favour of the user with the highest reputation. The tag of a dominant user lasts for longer than that of a weaker user, and can have a greater effect on the actions of artificial agents. For these reasons, the outcome of an encounter between two users could be altered significantly if their dominance or equipotentiality were altered.

In the case of *n* users interacting with multiple fragments, the importance of dominance and cooperation increase. The interaction between multiple users could be considered similar to that of only two users, with the two most dominant individuals (or groups of individuals) being the main users in the system. The same variables discussed above would apply to the subgroupings of users, with the overall effectiveness of the group being decided by their individual effectiveness and ability to work cooperatively as a group.

Human Users and Artificial Agents

With the possible exception of some of the special agents described in the Sect. 4.3.4, a human agent will always exhibit dominance over an artificial agent within the system. Aside from this lack of dominance, it is hoped that the interaction between human and artificial agents should be very much like human–human interactions. Given that the primary method of communication within the system is stigmergetic, the normal issues of complex language parsing and real-time interaction faced by interactive artificial agents do not apply. Artificial agents within the system behave as simple creatures, they act as supporting actors in the environment, and a human would no more desire to communicate with them than he would with a termite or a wasp.

4.8 Conclusion

In this chapter, we formulated a novel theoretical framework for reconstructing ancient text. The long-term reconstruction of cuneiform tablets has been a difficult task for scholars due to various reasons—the distribution of uncategorised fragments in wide geographical locations as a result of the division of the finds after excavation between the country where the excavation was carried out and the countries the archaeologists came from (in Iraq until 1969), excavations carried out at the same ruin by different nations, illicit and unprofessional excavations, and other circumstances. Other difficulties are reflected in the way the fragments are traditionally catalogued as photographs and hand-drawn copies; others, however, have been left in an unstructured state in the archives even after a century of storage. A complete reconstruction of the tablets is important, as it will reveal centuries of social, political, literary, and scientific knowledge acquired since 3,300 BC.

A computational approach to reconstructing cuneiform tablets is greatly needed here. If n fragments are given, there are $0.5\ n(n-1)$ match possibilities (compare Gehlken 1990, pp. 7–8). A Brute force computation calculating 3D facts match for all possibilities is too expensive and might take far longer than manual human matching. A computational approach involving subdomains of Artificial Intelligence techniques is a better approach. However, popular search heuristics such as Genetic Algorithms and biologically inspired adaptive and learning algorithms such as Artificial Neural Networks are not going to be the most efficient as there are massive permutations in the fragmentation of the pieces. A methodology partly involving crowd sourcing from human and artificial agents is needed. The concept of stigmergy therefore, naturally falls into this category. The nature-inspired approach based on how social insects coordinate without direct communication is intriguing, and perhaps strikingly similar to a community of scholars that may not communicate very much yet have the same goal. These insects, apart from the fact that they only leave information in their local environment, either through pheromones or through their action on the physical materials and structures they built, do not have other means of direct communication.

Designing a framework for supporting stigmergetic interactions is not straightforward. For stigmergy to work as intended, all fragments need to be digitised as 3D textured models, this can be accomplished through our development of a pipeline of 3D scanning processes for this research. Our process pipeline for digitising fragments can be easily constructed with off-the-shelf software and inexpensive hardware. This will make it possible for scholars and curators to capture 3D models on and off locations for uploading to our Web-based virtual environment. Since networked virtual environment has no spatial or temporal limitations, and 3D models hosted within such an environment will make it possible for unlimited access to collections. This necessarily brings researchers from various geographical locations together into a single space. As stigmergy does not require synchronised time; stigmergetic actions can be sequential and still work, and therefore time is of no importance, although continual actions will greatly speed up the reconstruction. Our client–server framework supporting both human and artificial agents and Web-based user interface has made this a possibility. Furthermore, virtually reconstructed pieces can be confirmed via the 3D printing process we presented as a supporting method. Finally, a good practice in piecing together the fragments is to start with as much information as possible. Our work on photogrammetric analysis of approximately 8,000 complete tablets from the CDLI has given us an initial understanding of the general shapes and ratios and nature of cuneiform tablets. The supporting technical foundation for a truly nature-inspired reconstruction of cuneiform tablets has been prepared. We hypothesised that such a method will greatly speed up the reconstruction of the tablets, bringing the long-sought knowledge hidden in the scattered fragments to the public. Finally, we proposed various hypothetical scenarios on how human and artificial agents will likely coordinate within the stigmergetic framework we developed. These will allow us a structured approach for facilitating better coordination amongst the agents. We are but a short distance from obtaining ancient knowledge.

The success of the project has direct relevance to the digital preservation of cuneiform tablets for future access and cooperative reconstruction globally. The significant contribution that this project will bring to the academic community is the restoration of past knowledge and insights into the hidden past. Only 5–10 % of the ancient Greek literature has survived; saving and reading the burnt papyri from Herculaneum (Villa of the Papyri) by modern technology might increase this percentage considerably. Our task, comparable to the task of saving burnt papyri from Herculaneum is to recover cuneiform tablets by computationally joining the fragments from the ancient corpus of Sumerian, Babylonian, Assyrian and Hittite literature.

Acknowledgments The authors gratefully acknowledge the support of The Leverhulme Trust research grant number F000 94 BP, and the multidisciplinary support of the University of Birmingham's Digital Humanities Hub.

References

Anderson, S., & Levoy, M. (2002). Unwrapping and visualizing cuneiform tablets. *IEEE Computer Graphics and Applications, 22*(6), 82–88.

Arvanitis, T. N., Davis, T., Livingstone, A., Pinilla-Dutoit, J., & Woolley, S. I. (2002). The digital classification of ancient near eastern cuneiform data. *British Archaeological Review, BAR International Series. 1075.*

Bettam, S. (2009). Archaeologists find cache of cuneiform tablets tablets in 2,700-year old Turkish temple. http://www.artsci.utoronto.ca/main/news-archives/archaeologists-find-cache-of-cuneiform-tablets-tablets-in-2-700-year-old-turkish-temple-aug.-7-09-by-sean-bettam-excavations-led-by-a-university-of-toronto-archaeologist-at-the-site-of-a-recently-discovered-temple-in-southeastern-turkey-have-uncovered. Retrieved March 12, 2013.

Blau, I., and Caspi, A. (2009). What type of collaboration helps? Psychological ownership, perceived learning and outcome quality of collaboration. *Proceedings of the Chais*

Conference on Instructional Technologies Research 2009: Learning in the Technological Era. pp. 48–55.

Bonabeau, E., & Théraulaz, G. (2008). Swarm smarts. *Special Editions, 18*(1), 40–47.

Bonabeau, E., Theraulaz, G., Deneubourg, J.-L., Aron, S., & Camazine, S. (1997). Self-organization in social insects. *Trends in Ecology and Evolution, 12*(5), 188–193.

Bruinsma, O. H. (1979). *An analysis of building behaviour of the termite Macrotermes subhyalinus (Rambur)*. Wageningen: Landbouwhogeschool te Wageningen.

Bulmer, L. (2002). A study of collaboration using jigsaw puzzles CVE'02: *Proceedings of the 4th international conference on Collaborative virtual environments* (pp. 157–158).

Camazine, S. (1991). Self-organizing pattern formation on the combs of honey bee colonies. *Behavioral Ecology and Sociobiology, 28*(1), 61–76.

CDP (2004). The cuneiform digital palaeography project. Birmingham: The University of Birmingham. Retrieved March 12, 2013, from http://www.cdp.bham.ac.uk/

Chandler J., and Fryer, J. (2011). Accuracy of AutoDesk 123D Catch. Retrieved March 12, 2013, from http://homepages.lboro.ac.uk/cvjhc/otherfiles/accuracy%20of%20123dcatch.htm

Chen, T., Seidel, H. (2008). Modulated phase-shifting for 3D scanning. *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition.*

Dekeyser, S. (2004). Towards a new approach to tightly coupled document collaboration. *9th Australasian Document Computing Symposium, Melbourne, Australia.*

Eseryel, D., Ganesan, R., & Edmonds, G. S. (2002). Review of computer-supported collaborative work systems. *Education Technology and Society*, *5*(2), 1–8.

Farella, E., Brunelli, D., Bonfigli, M.E., Benini, L., Gaiani, M., & Ricco, B. (2002). Using palmtop computers and immersive virtual reality for cooperative archaeological analysis: The Appian way case study: *Proceeding of the 8th International Conference on Virtual Systems and Multimedia VSMM.*

Gehlken, E. (1990). Uruk. Spätbabylonische Wirtschaftstexte aus dem Eanna-Archiv. Teil I. Ausgrabungen in. Uruk-Warka, Endberichte, Band 5.

Goss S., Beckers, R., Deneubourg, J.L., Aron, S. & Pasteels, J.M. (1990). How trail laying and trail following can solve foraging problems for ant colonies. In R. N. Hughes (Ed.), *Behavioral mechanisms of food selection. NATO ASI Series* (Vol. 20. pp. 661–678).

Grassé, P. P. (1959). La théorie de la stigmergie: essai d'interprétation du comportement des termites constructeurs. *Insectes Sociaux, 6*, 41–81.

Grassé, P.P. (1984). Fondation des Sociétés–Construction. Termitologia, tome II.

Hahn, D. V., Duncan, D. D., Baldwin, K. C., Cohen, J. D., & Purnomo, B. (2006). Digital hammurabi: Design and development of a 3d scanner for cuneiform tablets. *Proceedings of SPIE, 6056*, 130–141.

Hahn, D. V., Baldwin, K. C., & Duncan, D. D. (2007). Non-laser-based scanner for three-dimensional digitization of historical artifacts. *Applied Optics, 46*(15), 2838–2850.

Helbing, D., Keltsch, J., & Molnár, P. (1997a). Modelling the evolution of human trail systems. *Nature*, *388*(6637), 47–50.

Helbing, D., Schweitzer, F., Keltsch, J., & Molnár, P. (1997b). Active walker model for the formation of human and animal trail systems. *Physical Review E. 56*(3), 2527–2539. http://link.aps.org/doi/10.1103/PhysRevE.56.2527. Retrieved March 12, 2013

Holland, O., & Melhuish, C. (1999). Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life, 5*(2), 173–202.

Kampel, M., and Sablatnig, R. (2004). 3D puzzling of archeological fragments. *Proceedings of 9th Computer Vision Winter Workshop*. pp. 31–40.

Koller, D., & Levoy, M. (2004). *Computer-Aided Reconstruction and New Matches in the Formae Urbis Romae*. Formae Urbis Romae: Nuove scoperte.

Lewis, A., & Ch'ng, E. (2012). A photogrammetric analysis of cuneiform tablets for the purpose of digital reconstruction. *International Journal of Cultural Heritage in the Digital Era, EuroMED Supplement, 1*(1), 49–53.

Löffler, J. (2002). Adaptive visualization of distributed 3D documents using image streaming techniques. *Proceedings of the Sixth Eurographics Workshop on Multimedia*. pp. 9–19.

Macdonald, L. W. (2011). Visualising an egyptian artefact in 3D: Comparing RTI with laser scanning. *Proceedings of the 2011 International Conference on Electronic Visualisation and the Arts.* pp. 155–162.

Murray, J. (1864). Photography. *London Quarterly Review, 116*(1), 497–498.

Papaioannou, G., Karabassi, E., & Theoharis, T. (2002). Reconstruction of three-dimensional objects through matching of their parts. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 24*(1), 114–124.

Peters, P.J. (1970). Orb web construction: Interaction of spider (Araneus diadematus Cl.) and thread configuration. *Animal Behaviour. 18 Part 3*(0), 478–484.

Rabaud, E. (1937). *Phénomène social et sociétés animales*. Paris: Felix Alcan.

Smith, A. P. (1978). An investigation of the mechanisms underlying nest construction in the mud wasp Paralastor sp. (Hymenoptera: Eumenidae). *Animal Behaviour. 26 Part 1*(0), 232–240.

Theraulaz, G., & Bonabeau, E. (1999). A brief history of stigmergy. *Artificial Life, 5*(2), 97–116.

Wolfram, S. (1986). *Theory and applications of cellular automata*. Singapore: World Scientific.

Woolley, S. I., Davis, T. R., Flowers, N. J., Pinilla-Dutoit, J., Livingstone, A., & Arvanitis, T. N. (2002). Communicating cuneiform: The evolution of a multimedia cuneiform database. *The Journal of Visible Language, Special Edition on Research in Communication Design., 36*(3), 308–324.

Woolley, S. I., Flowers, N. J., Arvanitis, T. N., Livingstone, A., Davis, T. R., & Ellison, J. (2001). 3D capture, representation and manipulation of cuneiform tablets. *Proceedings of SPIE (Three Dimensional Image Capture and Applications IV)., 4298*(0277), 103–110.

Zhu, L., Zhou, Z., & Hu, D. (2008). Globally consistent reconstruction of ripped-up documents. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 30*(1), 1–13.