English AI Library Edition

AI Deciphers Ancient Scripts

Ancient Records Revived by AI

**Kim Kyung-jin,
Attorney at Law**

kimkj.com · AI Library · 2026

AI Deciphers Ancient Scripts

Ancient Records Revived by AI

Kim Kyung-jin, Attorney at Law

A scroll burned into charcoal, pieces of wood buried in mud, and clay tablets broken and scattered had long been objects no one could read. This book follows twelve events that broke that silence. It moves from virtual unwrapping, which read the Herculaneum scrolls buried in volcanic ash without opening them, to Leviticus from En-Gedi revived by X-rays, the examination of the Isaiah Scroll that identified the number of scribes by handwriting, digital glue that joins tens of thousands of fragments, the reconstruction of the Epic of Gilgamesh, the translation of Akkadian account books, the three-dimensional virtual binding of oracle bone script, the reading of cursive writing on Silla wooden slips, computational linguistics applied to Linear A, the Indus script, and rongorongo, language models that fill in broken Greek inscriptions, and multispectral archives that remain even when the original disappears. Each chapter is organized in the order of a case file, the investigation by a scientific inquiry team, and the truth revealed. I explained terms so that readers can follow the book without background knowledge, and I added verified references to each subchapter. The appendix gathers public data portals that readers can open at home. This book is based on 251 materials collected in the NotebookLM note "AI-Based Research Collection on the Decipherment of Ancient Languages and Ancient Texts (2024-2026)."

Public Research Archive

This book was organized and produced with artificial intelligence. It collects and synthesizes materials in major languages from around the world. NotebookLM does not limit the language of questions based on the language in which the materials were written. Readers can ask questions in their own language in the same research archive used to prepare this book. This is true even when some materials are in Korean, English, or other languages. If readers open the settings in NotebookLM and set the output language to the language they want, answers and generated results will appear in that language.

The public NotebookLM archive used to prepare this book (AI-Based Research Collection on the Decipherment of Ancient Languages and Ancient Texts 2024-2026, 251 materials):
https://notebook.google.com/notebook/cf32a6a1-66cc-473d-99d8-ed5a51596eb3

The original Korean edition of this book is in the AI Library on kimkj.com:
https://kimkj.com/ai-library/?mod=document&uid=12641

Table of Contents

Click each chapter title to move to the text. To read from the beginning, select Chapter 1.

Chapter 1. The Philosopher’s Library Swallowed by Volcanic Ash: The Carbonized Scrolls of Herculaneum

1.1. Case File

The Day the Volcano Covered the Library

In 79 CE, Mount Vesuvius erupted in southern Italy. The volcano released hot gas and volcanic ash all at once. At the foot of the mountain, by the sea, stood a city called Herculaneum. It was a resort city where wealthy Romans spent their summers. This city too was buried underground in a single day[1].

Outside the city stood a large villa. Today, scholars call it the Villa of the Papyri. Its owner appears to have been Lucius Calpurnius Piso Caesoninus. He was the father-in-law of Julius Caesar. The Greek philosopher Philodemus visited this house. Philodemus was a scholar of the Epicurean school. On one side of the villa was a library where scrolls were kept. Because so many of his writings were found there, scholars view it as a library connected to him[1].

Papyrus was the paper of the ancient Mediterranean world. It was made from a plant that grew along the Nile River in Egypt. The inner part of the stem was sliced thin and laid out in horizontal strips. A second layer was then placed on top in the vertical direction. When water was sprinkled on it and it was pressed, the plant sap acted as glue. Once dried, it became a sheet of paper. This left a crosswise grain on the surface of the papyrus. That grain later became an important clue[7].

That day, hot gas rushed into the library. The scrolls were baked without open flame. The papyrus that had served as paper turned into black charcoal. This process of burning into charcoal is called carbonization. The scrolls shrank, and their surfaces swelled[2]. More than 20 meters of volcanic deposits piled up over them. Volcanic deposits are layers formed when volcanic ash and stone powder accumulate and harden[1].

Because they were buried deep underground, air could not reach them. When air cannot reach plant fiber, it does not rot easily. The disaster acted like a lid that protected the artifacts. In that way, the scrolls endured for 2,000 years[2].

Black Sticks from Underground Tunnels

In 1738, the royal court of Naples began digging in this area. In 1750, a worker digging a well came across a colorful mosaic floor. That is how the location of the villa was revealed. Karl Weber took charge of the excavation. He was a Swiss military engineer[1].

Weber dug straight down into the hardened volcanic deposits. This kind of shaft is called a vertical shaft. He then dug narrow tunnels like a maze. Inside the tunnels, it was dark and the air was bad. There was also constant danger that the walls might collapse. Even under these conditions, Weber carefully drew a floor plan of the villa. Thanks to that map, we know what was found in each room[1].

Starting in 1752, workers began taking out objects that looked like black sticks. At first, they thought they were lumps of charcoal. They turned out to be papyrus scrolls. About 1,100 scrolls are known to have come from this villa. It is known as the only case in which a library from the ancient world has survived almost whole[1].

This library contained an unusually large number of writings by Philodemus. Finished works, copies, and drafts were stored together. Scholars believe he wrote and organized texts in this house[1]. The scrolls that were opened with great difficulty contained writings on philosophy, rhetoric, and ethics.

Rhetoric is the study of how to persuade others through speech and writing. No one knew what was inside the scrolls that could not be opened[2].

Scrolls Broken by Forced Unrolling

Scholars wanted to read the writing inside them. The problem was that the scrolls crumbled when touched[2]. In 1753, Father Antonio Piaggio built a machine. It used threads and weights to pull and open a scroll little by little[3]. This machine saved several scrolls, but it also tore inner layers[2].

In the 19th century, the British chemist Humphry Davy tried another method. He used chlorine gas, iodine vapor, and acid on the scrolls. The idea was to separate the layers with chemicals. The result was poor. The chemicals softened the papyrus and broke it apart[3].

Most attempts to open the scrolls physically damaged the artifacts. Fragments with letters on them sometimes scattered across the floor and disappeared. In the end, scholars stopped trying to open them. Even today, hundreds of unopened scrolls remain[2]. These artifacts are divided among holding institutions in Naples, Paris, and Oxford[4].

The Carbon Wall That Blocked X-Rays

Scientists looked for a way to read the scrolls without opening them. The first tool that came to mind was computed tomography. This technology is commonly called CT. CT measures differences in density inside an object without cutting it. Dense materials block more X-rays. So they appear bright in cross-sectional images[5].

Ink mixed with metal components glows white in CT images. In medieval Europe, ink containing iron was widely used. The burned scroll from En-Gedi in Israel was similar. This artifact is an ancient Hebrew manuscript from Israel. Because the ink appears to have contained metal, it left traces in the images. So in 2016, it was read without being opened[5].

The scribes of Herculaneum faced different conditions. A scribe was a person who copied writing by hand. They made ink by mixing soot from burned wood with water and resin. The main component of this ink is carbon. But the papyrus baked by volcanic heat also became carbon. The paper and the writing had become the same substance[6].

When density is the same, X-rays pass through both areas in the same way. No boundary line appears in the cross-sectional image. Scholars called this wall the carbon ink problem. In 2015, Italian and French researchers tried a different method. It was a technique called phase-contrast imaging. It uses the fact that X-rays bend slightly when they pass across the boundary of a material. Even if density is the same, the signal changes at the boundary surface. Using this method, the researchers captured the outlines of a few letters inside a rolled scroll. But it was not enough to read sentences[6]. For a long time, people believed this wall could not be crossed[2].

References

[1] Herculaneum Society. The Villa of the Papyri. University of Oxford. https://www.herculaneum.ox.ac.uk/index.php/the-villa-of-the-papyri/ (accessed September 3, 2026)

[2] Vesuvius Challenge. FAQ. scrollprize.org. https://scrollprize.org/faq (accessed September 3, 2026)

[3] Villa of the Papyri. Wikipedia. https://en.wikipedia.org/wiki/Villa_of_the_Papyri (accessed September 3, 2026)

[4] University of Kentucky. (2026, 6월 25일). The day the Herculaneum scrolls began speaking again. UKNow. https://uknow.uky.edu/research/day-herculaneum-scrolls-began-speaking-again (accessed September 3, 2026)

[5] Seales, W. B., Parker, C. S., Segal, M., Tov, E., Shor, P., & Porath, Y. (2016). From damage to discovery via virtual unwrapping: Reading the scroll from En-Gedi. Science Advances, 2(9), e1601247. https://doi.org/10.1126/sciadv.1601247

[6] Mocella, V., Brun, E., Ferrero, C., & Delattre, D. (2015). Revealing letters in rolled Herculaneum papyri by X-ray phase-contrast imaging. Nature Communications, 6, 5895. https://doi.org/10.1038/ncomms6895

[7] Papyrus. Encyclopaedia Britannica. https://www.britannica.com/topic/papyrus-writing-material (accessed September 3, 2026)

1.2. Forensic Science Team's Pursuit 1: Peeling a Virtual Onion

The Idea of Reading Without Opening

Professor Brent Seales of the University of Kentucky in the United States chose another path. The idea was to open the scroll inside a computer instead of opening the artifact. The physical object would remain untouched, and only 3D images would be obtained. Then the layers would be separated one by one on the screen. It was a method of moving inward as if peeling an onion[1].

The research team named this method virtual unwrapping. In 2016, this method succeeded for the first time. The object was the En-Gedi scroll discussed in the previous section. This artifact was found in 1970 at the site of an ancient synagogue. From the outside, it looked like a blackened lump. When the research team unfolded it on the screen, Hebrew writing appeared. It was the beginning of Leviticus in the Old Testament[1].

Virtual unwrapping has five stages. They are obtaining 3D images, separating layers, assigning ink values, flattening, and stitching. In the En-Gedi manuscript, this sequence worked as it was. That was because the ink was already visible in the images[1].

The Herculaneum scrolls were in worse condition than the En-Gedi manuscript. Their cross sections were crushed by volcanic heat. Their round shape had collapsed, and layer stuck to layer. In addition, the ink did not appear in the images. Of the five stages, the second and third became obstacles at the same time[2].

So with the Herculaneum scrolls, the order was changed. First, the layers are separated and flattened. The search for ink is then carried out separately. Section 1.3 tells the story of that ink search[2].

3D Images Taken by a Particle Accelerator

The first stage is to image the entire inside of the scroll[1]. Hospital CT usually distinguishes details down to around 1 millimeter. Papyrus layers are far thinner than that. A single layer is thinner than a human hair. With hospital equipment, the layers look blurred together[3].

The research team went to a particle accelerator. They first visited Diamond Light Source in the United Kingdom[3]. Later, they also used the European Synchrotron Radiation Facility in Grenoble, France[4]. These facilities produce extremely bright X-rays that travel in a straight path. The brighter the light, the clearer the images that can be obtained in a short time[3].

The research team placed the scroll upright in a container and fixed it firmly. Even a slight movement during imaging would blur the image. The data obtained in this way had a resolution of about 8 micrometers per voxel. A voxel is a three-dimensional pixel that also has a depth value. It can be understood as turning one dot in a drawing program into a small cube[3].

The amount of data was enormous. Imaging one scroll produces about 300 terabytes. That is enough to fill the storage space of thousands of smartphones. Handling this data itself became a new challenge[4].

The facility in Grenoble, France, is called a fourth-generation light source. New equipment began operating there in 2020. Twenty-one European countries spent 150 million euros to improve its performance. The facility was originally built to study materials and life processes. At first, no one knew it would be used to read ancient scrolls[4].

Volumetric Segmentation That Separates Layers

After the 3D images are obtained, the next stage is separating the layers. This work is called volumetric segmentation. It means separating each layer in a 3D image. The computer follows the path of the paper surface in each cross-sectional image. When those lines are connected, a 3D model of one sheet of paper is created. This model is called a mesh[1].

At first, people marked each boundary with a mouse[2]. The length of paper contained in one scroll is more than 1 meter[4]. It was too long and too dense to trace by hand. In addition, the layers had been pressed together, making the boundaries unclear[2].

The pace increased when contest participants released automation tools. A program made by Julian Schilliger of Switzerland is a leading example. This program follows patterns in neighboring cross-sectional images and connects the layers. It reduced work that had taken people several days to just a few hours[2].

This work was too large for a single laboratory to do alone. The organizers released the imaging data openly on the internet. Participants published the source code of the programs they made. One person added new functions on top of tools built by others. The readings and code can still be downloaded by anyone[5].

Flattening and Stitching

Once the layers are separated, the paper model is still curled. In that state, the letters look bent. So a calculation is performed to flatten the model. This stage is called flattening[1].

The computer turns the vertices of the model into small beads. It assumes that the beads are connected to one another by springs. Then it repeatedly calculates how to lay this net down on a floor. It minimizes how much the springs stretch or shrink. This prevents the shapes of the letters from becoming badly distorted[1].

The final stage is stitching. A scroll cannot be flattened all at once. The computer divides it into several regions and flattens them separately. The flattened regions must then be fitted back into one sheet. At this point, the distinctive grid-like fiber pattern of papyrus is used. If the fiber patterns of two regions fit together, it means they were originally neighbors. When this comparison is complete, virtual unwrapping is finished[1].

The five stages do not end in a single pass. If a mistake occurs at any stage, all later stages become misaligned. If the layers are separated incorrectly, writing from another sheet becomes mixed in. So the researchers process the same scroll repeatedly. They check the result with their eyes, change the settings, and run it again[5].

References

[1] Seales, W. B., Parker, C. S., Segal, M., Tov, E., Shor, P., & Porath, Y. (2016). From damage to discovery via virtual unwrapping: Reading the scroll from En-Gedi. Science Advances, 2(9), e1601247. https://doi.org/10.1126/sciadv.1601247

[2] Vesuvius Challenge. (2024, 2월 5일). Vesuvius Challenge 2023 Grand Prize awarded: we can read the scrolls! scrollprize.org. https://scrollprize.org/grandprize (accessed September 3, 2026)

[3] Vesuvius Challenge. FAQ. scrollprize.org. https://scrollprize.org/faq (accessed September 3, 2026)

[4] University of Kentucky. (2026, 6월 25일). The day the Herculaneum scrolls began speaking again. UKNow. https://uknow.uky.edu/research/day-herculaneum-scrolls-began-speaking-again (accessed September 3, 2026)

[5] Vesuvius Challenge. (2026, 6월 25일). An entire Herculaneum scroll has been read for the first time. scrollprize.org. https://scrollprize.org/firstscroll (accessed September 3, 2026)

1.3. Forensic Science Team's Pursuit 2: Finding the Ink's Microscopic Fingerprints

Carbon Written on Carbon

Even after virtual unwrapping was finished, the screen was still blank. The paper had been flattened, but the letters were not visible[1]. The cause was the carbon problem discussed earlier. The papyrus was carbon, and the ink was carbon too. X-rays could not distinguish between them[2].

Professor Seales's research team formed a hypothesis. They believed the ink must have left some trace[1]. Ink touches paper in liquid form. The liquid seeps between the fibers. Then, as it dries, it pulls slightly on the surface[3].

This means the shape of the paper surface changes in tiny ways. Even if the density is the same, the texture is different. The research team decided to find these surface changes[3]. They were too small to see with the eye, so the task had to be given to machines[1].

The tool used here was machine learning. This is not a method in which people write every rule by hand. Instead, many examples with correct answers are shown, and the machine finds the rules on its own. It is the same principle as recognizing cats after seeing thousands of cat photos. In the search for ink, the machine is shown samples marking places with ink and places without ink[1].

A Clue Called Crackle

In March 2023, a contest called the Vesuvius Challenge was launched[1]. Professor Seales started it together with Nat Friedman and Daniel Gross[4]. The organizers released the imaging data online for free. They offered prize money and gathered programmers from around the world[1].

For the first few months, no one could find the ink. That summer, a participant named Casey Handmer found a clue. He simply stared for a long time at the data from the unfolded surface. Then he noticed an unusual pattern. It was a trace of many fine cracks[1].

Think of a dry rice paddy in a drought. Dry mud cracks like a net. Similar fine cracks appeared where ink was present. The participants called this pattern crackle. Handmer received the first ink prize for this discovery[1].

Handmer posted the discovery in an online chat room at once. This sharing changed the course of the competition. Other participants soon used the pattern as training material[1].

Making the Machine Look Through a Small Window

An American college student, Luke Farritor, moved first. He taught a machine learning model the crackle pattern. The model answered with probabilities about where ink was present. As training progressed, its answers became more accurate[1].

The problem was that there was little training data. When data is scarce, a machine memorizes the answers. This is called overfitting. It is like a student who memorizes only the answers without understanding the test questions[1].

The participants increased the amount of data by rotating and flipping it. Even when an image is rotated, the nature of the fine crack pattern remains the same. Thanks to this, the machine learned only the features of the pattern itself[1].

They also adjusted the size of the input window. At first, they showed a wide area at once. When the view is wide, the whole shape of a letter comes into sight. Then the machine began to invent letters. It drew plausible strokes in places with no ink. This is called hallucination in artificial intelligence[1].

The researchers reduced the window to 64 by 64 pixels. At that size, the shape of a letter does not appear. The machine must judge only by the texture of the fine cracks. Fabrication decreased sharply. Instead, real brush marks appeared on the screen as they were[1].

The model used for this computation was made by connecting two kinds of models. One was a three-dimensional convolutional neural network. It is a structure that scans patterns inside a small cube and extracts features. The other was a transformer. It is a structure that checks how related distant signals are to each other. When the two structures are used together, they do not miss faint patterns inside thin layers[1].

The screen produced by the machine is not accepted as the answer at once. Palaeography experts read the screen separately. They check whether the arrangement of Greek words fits the grammar. They also check whether the following sentences make sense. If the machine reading and the human reading do not match, that target is calculated again[1].

Using Only the Surface Relief

After that, the method for finding ink kept improving. The participant Youssef Nader used a pretraining method. This means first showing the machine a large amount of scroll data with no answer key. The machine learns the appearance of papyrus fibers by itself. Then it is taught only how to find ink with a small answer key[1].

In March 2026, a study appeared that dealt with the relief itself. It was written by Giorgio Angelotti, Federica Nicolardi, Paul Henderson, and Brent Seales. The research team measured the surface of already opened papyrus samples with optical equipment. They trained a machine on data that contained only height differences. They did not include brightness or density information[3].

Even so, the machine identified the places with ink. This means ink can be found using only the height differences on the surface. The researchers also calculated how densely the surface must be measured. This number becomes the resolution standard for the next scans. The paper was published that year in the journal Scientific Reports[3].

Even so, ink detection is not yet a finished technology. Sections where the ink remains faint still appear blurred. Where layers are pressed together, the scan data itself is hazy. Two different sheets may also stick together as one sheet. A model that works well on one scroll may fail on another. So people check the results with their eyes and correct them along the way. The research is still moving forward one step at a time[5].

References

[1] Vesuvius Challenge. (2024, 2월 5일). Vesuvius Challenge 2023 Grand Prize awarded: we can read the scrolls! scrollprize.org. https://scrollprize.org/grandprize (accessed September 3, 2026)

[2] Mocella, V., Brun, E., Ferrero, C., & Delattre, D. (2015). Revealing letters in rolled Herculaneum papyri by X-ray phase-contrast imaging. Nature Communications, 6, 5895. https://doi.org/10.1038/ncomms6895

[3] Angelotti, G., Nicolardi, F., Henderson, P., & Seales, W. B. (2026). Ink detection from surface topography of the Herculaneum papyri. Scientific Reports. https://doi.org/10.1038/s41598-026-58467-1

[4] University of Kentucky. (2026, 6월 25일). The day the Herculaneum scrolls began speaking again. UKNow. https://uknow.uky.edu/research/day-herculaneum-scrolls-began-speaking-again (accessed September 3, 2026)

[5] Vesuvius Challenge. (2026, 7월 10일). Open problems: why reading every Herculaneum scroll is still a challenge. scrollprize.org. https://scrollprize.org/2026_open_problems (accessed September 3, 2026)

1.4. The Truth Revealed

The First Word Recovered in 2023

The first result was a single word. The announcement came on October 12, 2023. The main figure was Luke Farritor, an American college student. He was twenty-one years old at the time. The word he read was porphyras. It is a Greek word meaning purple. It is not a word commonly used in ancient texts. Farritor received 40,000 dollars as the first letters prize. Youssef Nader, who placed second, received 10,000 dollars. It was the first case of reading a word from an unopened scroll[7].

The results of the competition were announced on February 5, 2024. The grand prize was shared by three people. They were Youssef Nader of Egypt, Luke Farritor of the United States, and Julian Schilliger of Switzerland. The prize money was 700,000 dollars[1].

The artifact they read was a scroll held by the Institut de France in Paris. Its name is PHerc. Paris. 4. The three restored more than 2,000 Greek characters from this scroll. It was text from fifteen columns. A column is a vertical section of writing on a scroll. This was about 5 percent of the whole scroll[1].

The judging standards were strict. The participants had to submit four passages of 140 characters each. To pass, they had to read at least 85 percent of the characters in each passage. Palaeography experts reviewed the submissions anonymously[1].

The content was a discussion about pleasure. It asks whether rare things are necessarily more pleasant than common things. The text includes a passage that uses food as an example. Its point is that something is not believed to be more pleasant merely because it is rare. The sentence clearly shows Epicurean thought[1].

The Scroll Whose Title Was Recovered

The next achievement was a title. The first title prize was awarded on May 5, 2025. The winners were Marcel Roth and Micha Novak. The prize money was 60,000 dollars[2].

The target artifact was PHerc. 172 at the Bodleian Library in Oxford. The letters shown on the screen read Philodemus' On Vices. It was read as Book 1. Federica Nicolardi's research team at the University of Naples Federico II confirmed it. This was the first case of reading the title of a rolled scroll without opening it[2].

More than 70 columns were later read from this scroll. It is a text about flaws in character, such as greed, arrogance, and flattery. Now the flow of the argument can be followed from the beginning[3].

The First Whole Book Read from Beginning to End

Bigger news came on June 25, 2026. It was the announcement that one whole scroll had been read from beginning to end. The artifact is PHerc. 1667 at the National Library of Naples. The competition participants call it Scroll 4[4].

The scanning was done at the European Synchrotron Radiation Facility in France. After that, it went through virtual unwrapping and ink detection in order. The restored text was about 1.4 meters long. In columns, it was about twenty-two. This was the first time a Herculaneum scroll had been read in full[4].

The content is a Greek treatise on ethics and human character[3]. The text mentions a nephew of the Stoic philosopher Chrysippus. His name is Aristocreon[4]. The text is thought to have been written from the late 3rd century BCE to around the 2nd century BCE. The author has not yet been confirmed[3].

The research results were also released as a paper. Twenty-seven people were listed as authors, including Giorgio Angelotti, Stephen Parsons, and Federica Nicolardi[5]. The data, code, and transcriptions were also released[4].

The research team also disclosed its verification process. They rescanned PHerc. Paris. 4, which had been read in 2023, at a higher resolution. The transcription obtained from the new data matched the earlier result exactly. This is evidence that the machine had not invented the text. It also means the same method can be applied to other scrolls[4].

The 600 Scrolls Still Remaining

The same announcement contained other news as well. The title of a scroll called PHerc. 139 was read. It was Book 8 of Philodemus' On Gods. Until then, only Book 1 of this work had been known. For the first time, it became clear that this work continued for more than eight books[3].

More than 600 unread scrolls still remain[3]. The competition has set a Grand Prize for 2027. The total prize money is 1 million dollars. The first-place winner will receive 800,000 dollars. The targets are thirteen designated scrolls. More than 70 percent of the remaining letters in each column must be read. The deadline is June 25, 2027[6].

There are still many blank spaces in this story. Even after the letters are read, the work of interpreting their meaning remains separate. It also takes a long time to refine one column of transcription for scholarship. Some scrolls still have unidentified authors[3].

Even so, the direction is now clear. A book that no one could open for 2,000 years unfolded on a screen. A library thought to have disappeared in fire is slowly returning[4].

This work was not done by one person. There were physicists who handled the scanning and programmers who separated the layers. There were also palaeographers who read the ancient Greek. Students and office workers who joined the competition also helped. The book opened after 2,000 years because these people made it happen together. The remaining scrolls may be opened in the same way[4].

References

[1] Vesuvius Challenge. (2024, 2월 5일). Vesuvius Challenge 2023 Grand Prize awarded: we can read the scrolls! scrollprize.org. https://scrollprize.org/grandprize (accessed September 3, 2026)

[2] Vesuvius Challenge. (2025, 5월 5일). $60,000 First Title Prize awarded! scrollprize.substack.com. https://scrollprize.substack.com/p/60000-first-title-prize-awarded (accessed September 3, 2026)

[3] University of Kentucky. (2026, 6월 25일). The day the Herculaneum scrolls began speaking again. UKNow. https://uknow.uky.edu/research/day-herculaneum-scrolls-began-speaking-again (accessed September 3, 2026)

[4] Vesuvius Challenge. (2026, 6월 25일). An entire Herculaneum scroll has been read for the first time. scrollprize.org. https://scrollprize.org/firstscroll (accessed September 3, 2026)

[5] Angelotti, G., Parsons, S., Nicolardi, F., Nader, Y., Johnson, S., Josey, D., et al. (2026). Complete virtual unwrapping and reading of a rolled Herculaneum papyrus. arXiv:2606.29085. https://arxiv.org/abs/2606.29085 (accessed September 3, 2026)

[6] Vesuvius Challenge. Open Prizes. scrollprize.org. https://scrollprize.org/prizes (accessed September 3, 2026)

[7] Vesuvius Challenge. (2023, 10월 12일). First word discovered in unopened Herculaneum scroll by 21yo computer science student. scrollprize.org. https://scrollprize.org/firstletters (accessed September 3, 2026)

Chapter 2. Words Blooming from a Burned Ark: The Sealed Leviticus of En-Gedi

2.1. Case File

An Oasis West of the Dead Sea

Ein Gedi is an oasis on the western shore of the Dead Sea. It is also written as En-Gedi. The area around it is a dry desert, but spring water rises here. Where there was water, people gathered. For thousands of years, people built villages and lived by this water. The name of this land also appears in the Old Testament book of 1 Samuel. It is the story of David hiding in the wilderness of Ein Gedi to escape King Saul.

According to archaeological research, people lived on this hill from the late 8th century BCE. Then the village collapsed in a great fire around 600 CE[1]. This is land where more than 1,300 years are stacked in layers under the soil.

The area around the Dead Sea has little rain and dry air. This climate helps old objects survive for a long time. Wood, leather, and cloth decay less than in other places. This is also why the Dead Sea Scrolls were found in this region.

A major excavation began here in 1970. The excavation was led by Dan Barag and Ehud Netzer. Both were at the Institute of Archaeology of the Hebrew University of Jerusalem. Yosef Porath unearthed this scroll on May 5, 1970[1]. What the excavation team found was the site of an old Jewish synagogue. A synagogue is a building where a Jewish community gathered to pray and read the Bible. The floor was covered with a mosaic made of small stones of different colors.

Inside a synagogue, there is a space for the ark. The ark is a chest that holds sacred biblical scrolls. A Jewish community regards this chest as a holy place in the synagogue. The excavation team cleared burned debris from that ark. A black lump came out of the debris[1].

The fact that it came from the ark was itself an important clue. A synagogue ark does not hold just any document. It holds biblical scrolls that the community read and observed. This gave scholars a basis for guessing what the black lump was.

An old biblical manuscript is a precious source in itself. Before the printing press, people copied texts by hand. Each time a text was copied, letters could change. When an old manuscript is found, scholars can check how the text was transmitted. That is why scholars devoted themselves to the black lump from the ark.

The Black Lump Left by Fire

Fire has the power to destroy records. Paper and leather both become ash in fire. But in Ein Gedi, the situation was a little different. The scroll, pressed under debris, was slowly baked with too little oxygen. It could not scatter into ash and hardened like charcoal. This change, in which something hardens like charcoal, is called carbonization. It happens when something burns in a fire with too little air.

The principle is the same as the process by which wood becomes charcoal in a charcoal kiln. If wood is simply burned, only ash remains. If it is covered with soil to block the air and then baked, black charcoal remains. The carbonized scroll was made in the same way.

The carbonized lump was buried under ash and soil from the collapsed building. Because air could not reach it, decay also stopped. Leather is normally a material that decays easily. If it had been left on the ground, it would have disappeared long ago. The fire that damaged the artifact was also the force that preserved it.

At first glance, the lump taken out by the excavation team looked like a piece of charcoal. But on the broken side, rounded rolled layers remained. Archaeologists judged it to be a rolled leather scroll[1]. The artifact was named after the land where it was found. That is how the name Ein Gedi Scroll came about.

This scroll is parchment made from prepared animal skin[1]. The scrolls from Herculaneum in Italy were made of papyrus. Papyrus is a plant-based writing material made by pressing together plant stalks. The two artifacts differ from the base material onward. This difference later becomes the dividing line in how they are read.

Parchment takes much labor to make. Hair and fat are removed from animal skin, and the skin is washed. It is stretched tightly on a frame and dried. Text is written on the thin sheet prepared in this way. It is tougher and lasts longer than paper, but it was expensive.

The state of discovery was also unusual. Old scrolls usually come out as small broken fragments. The Dead Sea Scrolls found in the Qumran caves from 1947 onward also include many fragments. The Ein Gedi artifact came out while preserving its rolled shape. This meant that several more layers remained inside.

An Artifact That Could Not Be Touched

To read the text, the scroll had to be opened. But this artifact could not be opened. Heat had shrunk the leather, causing the layers to stick to one another. The outside was hard like stone, and the inside was easy to crumble. It was in a condition where a touch of the fingertip could turn it into powder[1].

There were not no methods. One way was to cut between the layers with a knife. Another was to apply chemicals to soften the leather. Both methods would damage the artifact. Once a scroll was broken, it could not be restored.

Looking only at the outside could not even show whether there were letters. The carbonized surface was entirely black and had no pattern. The ink was black and the leather was black, so the eye could not tell them apart. Even the idea that text existed inside was a guess based on the shape of the layers.

People who handle cultural heritage always keep this point in mind. An artifact is one of a kind, and it cannot bear even one failure. So conservation science has a principle of examining without touching. This kind of examination is called non-invasive investigation. Invasive means going inside, and non-invasive means the opposite. The Ein Gedi Scroll was an artifact that allowed only non-invasive investigation.

45 Years of Silence

The scroll went into an artifact storage facility in Israel. Later, the Israel Antiquities Authority took charge of and preserved it[1]. Waiting inside a box was the only preservation method available at the time. A storage facility is a space that keeps temperature and humidity stable. It blocks air, light, and insects so that artifacts do not deteriorate further. Even for an unreadable artifact, preservation continues.

Pnina Shor, who was in charge of preserving the Dead Sea Scrolls, took over its care. Later, her name also appeared on the research that read this scroll[1].

The waiting continued from 1970 to 2015[1][2]. For 45 years, no one saw the letters inside it. The surface was so black that it was hard even to know whether letters existed. Some scholars said it might be an artifact that would never be read.

There was a reason for the wait. The technology of 1970 could not handle this lump. Imaging devices that could capture the inside of an object in fine slices were not common. Even if such images were taken, there were no computers capable of calculating that much data. The artifact was left as it was, waiting for technology to grow.

During that time, similar concerns continued around the world. The Herculaneum papyri in Italy were blocked by the same wall. Old attempts to open them physically only broke the artifacts. The two artifacts lived in different places and used different materials, but the problem was the same. A way was needed to read the inside without touching it.

What broke the silence was not a new manual technique. It was imaging that looked inside without cutting the object, and computer calculation. Seales and his team had long studied ways to read old texts without touching them. This research was supported by the U.S. National Science Foundation and Google[2]. They needed an artifact to test methods built up over many years, and the Ein Gedi Scroll became that target.

In 2015, a team led by Brent Seales at the University of Kentucky took on this work. The research results were published on September 21, 2016, in the journal Science Advances[1][2].

Now one question remained. How did they read the inner letters without touching the scroll? The answer lies in the sequence of imaging devices and calculation.

References

[1] Seales, W. B., Parker, C. S., Segal, M., Tov, E., Shor, P., & Porath, Y. (2016). From damage to discovery via virtual unwrapping: Reading the scroll from En-Gedi. Science Advances, 2(9), e1601247. https://doi.org/10.1126/sciadv.1601247

[2] University of Kentucky Research. (2016, 9월 22일). Virtually Unwrapping the En-Gedi Scroll. research.uky.edu. https://research.uky.edu/news/virtually-unwrapping-en-gedi-scroll (accessed September 3, 2026)

2.2. The Forensic Science Team's Pursuit: X-ray Digital Purification

Imaging That Sees Inside Without Cutting

The first tool chosen by the researchers was micro-CT. CT means computed tomography. It works on the same principle as the imaging used in hospitals to look inside the body. X-rays are sent from many angles to measure density inside an object. A computer gathers those values and creates a three-dimensional map. Micro-CT does this at a very small scale.

The device used to image the Ein Gedi Scroll was a Bruker SkyScan 1176. The imaging was done by Merkel Technologies in Israel. The company donated the imaging data to the Israel Antiquities Authority. The imaging scale was 18 micrometers per unit[1]. One micrometer is one thousandth of 1 millimeter. It is a size equal to a fraction of the thickness of a human hair.

There is a reason for this scale. To separate the layers, the object must be measured much more finely than a single layer of leather. If the scale is larger than the thickness of a layer, two layers blur into one lump. The imaging scale is a condition that sets the limit of readable layers.

The data produced by the imaging is not a single photograph. It is a three-dimensional mass densely packed with very small cube-shaped particles. These particles are called voxels. If a pixel in a photograph is a flat particle, a voxel is a three-dimensional particle. Each particle contains a number that shows the density at that location.

At this stage, the artifact was moved into the computer without being touched. From this point on, the object being handled was not leather but numbers. Numbers do not crumble no matter how much they are handled. If the calculation is wrong, it can be started again from the beginning. This is the strength of non-invasive investigation.

Bright Points Left by Ink Mixed with Metal

Obtaining three-dimensional data does not make letters visible. For letters to appear, the density of the ink and the base material must differ. This condition failed in the Herculaneum papyri. Their ink was made from charcoal powder, and the burned papyrus is also a mass of carbon. Their densities are similar, so X-rays do not separate them well.

The Ein Gedi Scroll had better conditions. The researchers confirmed in the scan data that the density where letters were placed was higher than the surrounding area. So they thought this ink likely contained metals such as iron or lead. However, because no outer surface was exposed, they could not directly analyze its composition. The researchers stated this limitation plainly in the paper[1].

Metal absorbs many X-rays. Areas with high absorption appear bright on the screen. Burned leather absorbs fewer X-rays and appears dark. This is how the brightness difference between letters and the base material is formed. This difference became the first clue that allowed the computer to find the letters. This is also why the Ein Gedi Scroll was read before the Herculaneum papyri.

However, the composition of this ink has not yet been confirmed. What the researchers presented was only an observation of high density and an inference based on it[1]. What the ancient scribe mixed into the ink can be known only by obtaining a sample. Treating what is unconfirmed as confirmed fact is against the principle of this field.

The Four Steps of Virtual Unwrapping and Removing Stains

Finding bright points does not make the text readable. The letters are scattered on rounded rolled layers. These layers must be opened inside the computer. This work is called virtual unwrapping. Seales and his team divided virtual unwrapping into four steps[1].

The first step is to find the surface of the layers. The computer slices the three-dimensional data thinly and identifies the coordinates of the leather layers. The points gathered in this way are connected to make a thin sheet-like surface. If the layers are stuck to one another, the process loses its path at this stage. If two layers are read as one joined layer, the letters also become mixed.

The second step is to apply brightness values to that surface. The imaging values at the coordinates through which the surface passes are read and transferred onto the surface. The ink with high density then appears as bright strokes.

The third step is to flatten the crumpled surface. When a three-dimensional object is pressed into a plane, its shape stretches or distorts. It is the same problem as the polar regions becoming enlarged when a globe is unfolded into a paper map. The calculation unfolds the surface in a way that reduces this distortion as much as possible.

The fourth step is to stitch together the separately unfolded surfaces. The boundaries of each surface are aligned and combined into one large image[1].

Among the four steps, the first is difficult. If the layers are traced incorrectly even once, all three later steps go wrong. An incorrectly unfolded image may even create letters that were not there. This is why researchers check this stage by hand.

After these four steps, the black cylinder becomes a single plane. The researchers unfolded five layers of leather in this way. The area of the unfolded surface was estimated at about 94 square centimeters[1]. That is a little smaller than one page of a notebook.

Problems still remain in the unfolded image. The leather thickness differs from place to place, and wrinkles remain. Noise from the imaging process is also mixed in. So some places are too bright and others are dark. When letters and stains are mixed, reading becomes difficult.

The researchers removed these stains by calculation. First, they measured the average brightness of surrounding pixels and made a brightness map of the base. This map contains only broad brightness flows. Then they divided the original image by this base map. After this division, broad brightness differences disappear. Only the ink signals concentrated in narrow strokes remain clearly. Hebrew letters appeared in the image that passed through this process[1].

Tasks That Continue in 2026

Virtual unwrapping was later used on other artifacts as well. On June 25, 2026, the Vesuvius Challenge announced news. It said that one Herculaneum scroll had been read from beginning to end. The artifact number is PHerc. 1667. The imaging was done at the European Synchrotron Radiation Facility in Grenoble, France. The unfolded Greek text consists of twenty-two columns about 1.4 meters long. It was identified as a work on ethics from the 2nd century BCE[2].

A synchrotron is a device that bends fast-moving electrons with a magnetic field. When electrons bend, very strong X-rays come out. Images taken with this light reveal even small differences in density. This gives a clarity that is hard to obtain with laboratory imaging machines. It is a tool one step above the device used at En Gedi.

This achievement came from improved imaging methods. In the new imaging method, carbon ink was seen directly inside the material. It detected ink that had not been visible because the density difference was small. It created through imaging technology the brightness difference that had appeared naturally at En Gedi. It was reported that the same method could be used on more than 600 remaining scrolls[3].

Work still remains. A summary released on July 10, 2026, lists six remaining problems. In compressed layered sections, X-rays scatter and make the image cloudy. The computation that finds layers creates holes or merges two layers. Voxel data does not contain information about which layer continues into which layer. Human-made training labels are not accurate at the grain level. A model learned from one scroll does not work well on another scroll. It is also unclear at what depth of the layer the ink lies[4].

Possible solutions were presented as well. If correct labels are lacking, there is a path that lets the data itself learn the rules. Self-supervised learning models are considered one candidate. New neural network structures are also being used in the computation that finds layers[4].

A new clue also appeared. A study released on March 29, 2026, relied on the fact that ink changes the shape of the surface. The researchers precisely measured the surface curves of already unfolded papyrus and taught them to machine learning. As a result, they could identify where letters were located using only the curves. The study also examined what imaging resolution is needed to read those curves[5]. It is a study that found again, in another property, the clue that En Gedi had obtained from density.

References

[1] Seales, W. B., Parker, C. S., Segal, M., Tov, E., Shor, P., & Porath, Y. (2016). From damage to discovery via virtual unwrapping: Reading the scroll from En-Gedi. Science Advances, 2(9), e1601247. https://doi.org/10.1126/sciadv.1601247

[2] Vesuvius Challenge. (2026, 6월 25일). An entire Herculaneum scroll has been read for the first time. scrollprize.org. https://scrollprize.org/firstscroll (accessed September 3, 2026)

[3] ASNT. (2026, 8월 13일). X-Ray Imaging and AI Recover First Complete Text from Sealed Ancient Scroll. Materials Evaluation. https://www.asnt.org/me/26/8/x-ray-ai-recover-text-sealed-ancient-scroll (accessed September 3, 2026)

[4] Vesuvius Challenge. (2026, 7월 10일). Open Problems: Why Reading Every Herculaneum Scroll Is Still a Challenge. scrollprize.org. https://scrollprize.org/2026_open_problems (accessed September 3, 2026)

[5] Angelotti, G., Nicolardi, F., Henderson, P., & Seales, W. B. (2026, 3월 29일). Ink Detection from Surface Topography of the Herculaneum Papyri. arXiv:2603.27698. https://arxiv.org/abs/2603.27698 (accessed September 3, 2026)

2.3. The Truth Revealed

Leviticus Appears on the Screen

The letters that appeared on the computer screen were Hebrew letters. The researchers unfolded five layers of leather, and two columns of writing emerged inside them. The text was Leviticus, the third book of the Old Testament. The part read was Leviticus chapters 1 and 2[1].

The combined image contained thirty-five lines of writing from one sheet of the scroll. Of these, eighteen lines still had letters. The remaining seventeen lines were restored by scholars. Restoration is not a computer inventing text. It is the work of matching the remaining strokes with the known text and filling in the missing parts. Each line contained thirty-three to thirty-four letters and spaces[1].

It is important that the paper separated the surviving lines from the restored lines. This lets readers distinguish what has been confirmed as fact. If surviving text and filled-in text are mixed in a presentation, later researchers cannot verify it. Because of this, other scholars could examine the same material again.

Leviticus 1 deals with the rules for burnt offerings. A burnt offering is a sacrifice in which the offering is completely burned on the altar. Chapter 2 deals with the rules for grain offerings. It explains how to mix fine flour with oil and place incense on it. From a burned Torah ark came the rules for a sacrifice offered by burning.

The restored text also includes the first verse of Leviticus[1]. It is the sentence saying that the Lord called Moses from the Tent of Meeting and spoke to him. The survival of the beginning of the book means that this scroll was the start of a copy of Leviticus.

The phrase two columns also has meaning. Ancient scrolls were written in long vertical blocks. When one block was filled, the writing moved to the next block. The appearance of two columns means that two such blocks appeared in sequence.

What was read was not only letters. Even the thickness of strokes made by a scribe pressing a brush remained on the screen. This is evidence that it was a manuscript copied by hand.

Filling a Thousand-Year Gap

The researchers also measured the age of the scroll. Radiocarbon dating gave a date in the 3rd to 4th century CE[1]. This method measures changes in carbon remaining in an artifact to determine its age. This result fills a long-standing gap in the study of ancient biblical manuscripts.

Let us first look at what the gap is. The biblical manuscripts among the Dead Sea Scrolls found in the Qumran caves date from the 3rd century BCE to the 1st century CE. By contrast, the standard Hebrew

Bible manuscripts that survive in complete form date to around 1000 CE. There is a gap of nearly one thousand years between the two groups. Manuscripts showing how the text was transmitted during this period were rare.

When the gap is large, doubt also grows. There is no way to confirm whether medieval manuscripts copied the ancient text as it was. Some scholars believed that medieval scribes may have refined the text.

The En Gedi scroll came from the middle of this gap. And the text that was read was exactly the same as the consonantal text of the medieval Hebrew Bible[1]. Hebrew writes only consonant letters and does not write vowels. A text written only with those consonants is called a consonantal text. This medieval text is called the Masoretic Text. It is the standard text preserved by Jewish scribes under fixed rules. Even the paragraph divisions were the same[2]. This result shows that the text was transmitted unchanged over a long period.

Jewish scribes are known to have followed strict rules when copying the Bible. Counting letters for checking was one of those rules. The En Gedi scroll became physical evidence that those rules were actually followed. A single manuscript added one piece of evidence to a long debate.

This manuscript carries one more record. Among the Pentateuch manuscripts found inside a Torah ark so far, it is the earliest dated[1]. The Pentateuch refers to the first five books of the Old Testament together.

What Is Settled and What Remains

What has been revealed and what is still under discussion must be separated.

What has been revealed is as follows. The text that was read is Leviticus chapters 1 and 2. The radiocarbon date points to the 3rd to 4th century CE. The consonantal text is the same as the medieval standard text. The density of the ink is higher than that of the leather, so it appears to have contained metallic components[1].

The remaining issues are also clear. The components of the ink could not be analyzed directly[1]. This was because no exposed surface was available from which a sample could be taken. There is also a view that dates the scroll differently based on the handwriting style. The palaeographer Ada Yardeni dated it to the late 1st century or early 2nd century CE. A palaeographer is a scholar who estimates dates by comparing old handwriting styles. Drew Longacre responded that radiocarbon dating is more reliable[2]. The two views have not yet been settled.

The storage period is another unresolved question. The scroll was made in the 3rd to 4th century CE. The synagogue burned around 600 CE[1]. This means the manuscript remained in the Torah ark for more than two hundred years. It is not known why the community kept an old manuscript for so long. There is only the assumption that they valued an inherited manuscript.

The part that was read is also only one piece of the scroll. What was unfolded was the outer five layers[1]. Deeper inside, the layers are more severely compressed. The problem of separating layers in compressed sections still remained in 2026. It is the problem that researchers working on the Herculaneum scrolls are still trying to solve[3].

To read the inside, two things must improve together. Imaging resolution is needed that can divide the layers more finely. The computation that separates fused layers must also improve. If these two conditions are met, the En Gedi scroll can be placed on the imaging stage again. This is possible because the artifact still remains intact.

Methods that help dating itself are also increasing. In 2025, a model was announced that uses machine learning to learn the handwriting styles of ancient Hebrew manuscripts and estimate their dates. This model learned both radiocarbon dates and handwriting data. It measures even curves in strokes that are hard for the human eye to catch[4]. As such tools accumulate, the debate over the date of the En Gedi scroll may also be reviewed again.

The path opened by virtual unwrapping does not stop with biblical manuscripts. The same method was used on old letters that remained folded and sealed. Imaging and computation are also being applied to fragile books and documents. The principle of reading without touching is becoming a basic rule for handling materials.

The image produced by the computer is not the final answer by itself. In the 2016 study, the computer's work went as far as unfolding the layers and making an image. It was people who read the letters in that image and attached meaning to them. For that reason, palaeographers must examine each letter again. What this case left behind is not only one text. It created a way to handle artifacts that break when touched. The order is to obtain numbers through imaging, unfold them through computation, and have people review them. This order is used again for several artifacts discussed later.

The Leviticus read at En Gedi also went through that review and became part of the scholarly record. What broke 45 years of silence was work done by machines and people together. The imaging device measured density, and the computation unfolded the layers. Then people read the letters and attached meaning to them. The next chapter continues with the story of examining the shapes of letters themselves.

References

[1] Seales, W. B., Parker, C. S., Segal, M., Tov, E., Shor, P., & Porath, Y. (2016). From damage to discovery via virtual unwrapping: Reading the scroll from En-Gedi. Science Advances, 2(9), e1601247. https://doi.org/10.1126/sciadv.1601247

[2] En-Gedi Scroll. (2026, 9월 3일). Wikipedia. https://en.wikipedia.org/wiki/En-Gedi_Scroll (accessed September 3, 2026)

[3] Vesuvius Challenge. (2026, 7월 10일). Open Problems: Why Reading Every Herculaneum Scroll Is Still a Challenge. scrollprize.org. https://scrollprize.org/2026_open_problems (accessed September 3, 2026)

[4] Popović, M., Dhali, M. A., Schomaker, L., van der Plicht, J., Rasmussen, K. L., La Nasa, J., Degano, I., Colombini, M. P., & Tigchelaar, E. (2025). Dating ancient manuscripts using radiocarbon and AI-based writing style analysis. PLOS ONE, 20(6), e0323185. https://doi.org/10.1371/journal.pone.0323185

Chapter 3. Two Shadows Hidden in the Isaiah Scroll: Handwriting Analysis

3.1. Case File

A Leather Scroll from a Cave

In 1947, old leather scrolls were found in caves in the Qumran area west of the Dead Sea. Bedouin goat herders found jars and scrolls inside cliff caves. These artifacts are called the Dead Sea Scrolls. Excavations later continued in several caves. The Dead Sea Scrolls collected so far are organized as about 930 manuscripts. A book copied by hand is called a manuscript. If counted as individual leaves, the leather and papyrus fragments number in the thousands[1]. They are thought to date from between the 3rd century BCE and the 1st century CE[2].

The Dead Sea is a very salty lake, and the land around it is a dry desert. There is little rain, and the air is dry. The temperature inside the caves also does not change much. This is why the leather and papyrus survived without rotting. Few collections of ancient documents of this size have survived.

Among these artifacts, one scroll drew special attention from scholars. It is the Great Isaiah Scroll. Its scholarly symbol is 1QIsaa. This means the first manuscript of the Book of Isaiah from Qumran Cave 1. It is a manuscript written in Hebrew on animal skin in the 2nd century BCE. It is about 7.3 meters long. It contains almost the entire text of the Book of Isaiah in the Old Testament. The writing is divided into 54 long vertical sections[2]. In scholarly terms, these sections are called columns. Among the biblical manuscripts found in the Dead Sea caves, it is the best preserved.

The scroll was made by sewing several sheets of leather together with thread. Several columns were written on one sheet, and then the writing continued onto the next sheet. After it was written, the scroll was rolled up for storage. To read it, one side was unrolled while the other side was rolled up. This is where the name scroll comes from.

The value of this manuscript comes from its age. The Dead Sea Scrolls include the oldest surviving Hebrew Bible manuscripts[3]. The Great Isaiah Scroll is one of them. Scholars gained material with which to compare how the biblical text was transmitted over many centuries. Yet nothing was known about the person who made the scroll.

A Scholarly Divide over Who Wrote It

Reading the scroll and finding out who wrote it are different tasks. The letters can be read, but this manuscript does not give the name of the person who copied it. A person who copied writing is called a scribe. This scroll has no scribe's signature and no date[2]. So scholars had to estimate the number of people by looking only at the shape of the handwriting.

Scholars stood at two crossroads on this question. One side argued that one scribe wrote it alone from beginning to end. Their reason was that the handwriting across the whole scroll looked consistent. The other side argued that two or more people divided the work. Their reason was that the shapes of the strokes changed little by little after the middle.

Both views had one difficult problem. Even handwriting by one person is not exactly the same from beginning to end. When the hand becomes tired, strokes become loose. When the pen is changed, the thickness changes. Without knowing the size of this variation, it is impossible to decide whether the writer changed. Both claims relied on visual impressions. There was no weight on the scale to decide which side was right.

Qumran is thought to have had a community that made and guarded scrolls. People placed scrolls in jars and hid them in caves. They seem to have tried to protect the documents as war approached. The names of the people who hid them also did not survive. Finding the hands that made the scroll is also a way to approach those people.

This dispute is not only about handwriting. If several people divided the work of making a manuscript, it means there was an organization in which scribes worked together. It connects to the question of whose hands carried the biblical text down to the present[2]. This is why researchers of the Dead Sea Scrolls held on to this question for so long.

The Limits of Visual Analysis

The study of old handwriting is called palaeography. It is a method in which a scholar examines letter forms with a magnifying glass. The scholar measures the height, slant, and stroke thickness of letters and copies them into drawings. Scholars may also place the handwriting of several manuscripts in chronological order to trace patterns of change. For a long time, this was the only tool for handwriting analysis.

This method has a weakness. It depends heavily on the scholar's eyes and experience. Scholars often gave different evaluations after looking at the same letters[2]. Human eyes cannot measure very small numerical differences. The force used to press a reed pen or the thickness of spread ink is hard to judge by sight. Memory also has limits. It is not easy to place two letters several meters apart side by side in one's mind and compare them. In the end, the disagreement repeated itself.

Another wall was the educational method of that age. A person training to become a scribe learned by copying the teacher's handwriting exactly. Scribes were taught not to show their own individuality. As a result, the handwriting of scribes from the same school looked almost identical on the surface. For human scholars, this was a thick fog. They could not tell whether the tiny differences between the front and back of the scroll came from a change in one person's condition or from a change of person[2].

If the analysis is wrong, the judgments that follow are also wrong. If there was one scribe, one manuscript contains the time of one person. If there were two, it means there was an organization that divided the work. The size and character of the place where the manuscript was made would change. Many conclusions depended on one handwriting analysis.

The situation is similar to imagining children in a classroom taking dictation from the same handwriting model. Notebooks written by children who all tried to write neatly are hard to tell apart on the surface. Still, each child bends strokes in slightly different ways. The problem was that there was no ruler to measure that difference.

The situation changed when manuscript photographs were released digitally. The real scroll is weak against light and moisture, so it cannot be opened often. With photographs, scholars can examine it again and again without touching the artifact. The Israel Antiquities Authority created a digital collection of the Dead Sea Scrolls with Google. They photographed the artifacts by shining light of several wavelengths with equipment called MegaVision. They captured even near-infrared light, which human eyes cannot see, and recovered faded letters. Their goal was image quality close to looking directly at the real object[1].

As high-resolution photographs that computers could read accumulated, a path opened for new tools. A photograph is a collection of numbers. Each small point that makes up a screen is called a pixel. Each pixel has a brightness value. This means the thickness and angle of strokes can be measured as numbers. The work once done by eye was ready to become calculation. From here, the artificial intelligence detective steps in.

References

[1] Israel Antiquities Authority. (n.d.). The Digital Library. The Leon Levy Dead Sea Scrolls Digital Library. https://www.deadseascrolls.org.il/about-the-project/the-digital-library (accessed September 3, 2026)

[2] Popović, M., Dhali, M. A., & Schomaker, L. (2021). Artificial intelligence based writer identification generates new evidence for the unknown scribes of the Dead Sea Scrolls exemplified by the Great Isaiah Scroll (1QIsaa). PLOS ONE, 16(4), e0249769. https://doi.org/10.1371/journal.pone.0249769

[3] University of Groningen. (2021, 4월 21일). Cracking the code of the Dead Sea Scrolls. rug.nl. https://www.rug.nl/news/2021/04/cracking-the-code-of-the-dead-sea-scrolls?lang=en (accessed September 3, 2026)

3.2. The Forensic Team's Trail: Biometric Fingerprints Left in Handwriting

The Investigation Team Is Formed

A research team at the University of Groningen in the Netherlands took on this case. The palaeographer Mladen Popović led the team. Artificial intelligence researchers Lambert Schomaker and Maruf Dhali joined him[1]. In 2017, they first published a method for measuring the handwriting of the Dead Sea Scrolls by computer[2]. They refined that method and aimed it at the Great Isaiah Scroll. The results came out in 2021[1, 5].

The team's goal was clear. It was to turn human impression-based evaluation into numbers. First, they measured the size of the variation that appears within handwriting by the same person. Then they asked whether the difference between the front and back of the scroll was larger than that variation. If the difference was larger, it meant another person had written it.

The feature calculation method created in 2017 was used to compare manuscripts with one another[2]. The research team applied that method again, this time focusing on one manuscript, the Great Isaiah Scroll. Finding where the scribe changed within a single manuscript is harder than comparing manuscripts. The differences in handwriting are much smaller.

This method has one advantage. If the calculation procedure is written down, anyone can get the same answer from the same photographs. There is no need to rely on a scholar's eye. If the result is wrong, the steps can be traced to find where it went wrong. Even if scholars argue again over the result, the grounds for the dispute remain as numbers.

The First Step: Removing Stains

The first task is to leave only the ink in the photograph. This process is called binarization. It separates every pixel into one of two groups: ink or background. It divides the photograph into black writing and a light background.

This step is needed for a reason. The scroll had been in a cave for about 2,000 years. Stains, mold marks, and fine cracks formed on the leather surface. To a computer, these marks and ink strokes are both dark pixels. If the background remains as it is, the machine learns the texture of the leather instead of the writing. Then it distinguishes types of leather instead of scribes, and the analysis goes wrong[1].

The research team made a separate neural network for this task. Its name is BiNet. A neural network is a calculation structure modeled on the connections among nerve cells in the human brain. BiNet learns the criteria by looking at samples where ink and background were marked by hand. After learning, BiNet removes stains from new photographs and leaves only the strokes.

There is a rule that must be kept here. The original stroke thickness and edge shape must not be changed. That edge is where the scribe's habits are found. If strokes are smoothed while removing noise, the evidence disappears. BiNet was designed to preserve ink traces as they are[1]. This technology first appeared as a public preprint in 2019[3]. A paper that reorganized its design and performance appeared in the 2026 proceedings of an international conference on document recognition[4].

A Ruler for Measuring the Habits of the Whole Handwriting

Two kinds of features are extracted from the cleaned letters. The first concerns texture, and the second concerns shape.

The texture feature is called Hinge. A hinge is the joint of a door. At each place where a stroke bends, it measures two directional angles. It then makes a table of how often each pair of angles appears. This table records as numbers the habit of wrist movement when writing. It can be calculated without recognizing each letter one by one[1].

The shape feature is called Fraglets. The word means short fragments cut from strokes. Strokes are cut up and gathered into a list. The research team made a standard list from fragments taken from many manuscripts. When new handwriting is entered, the method counts which item in the list each fragment resembles. The basic stroke forms learned when a person learned to write appear here[1].

These two features have something in common. They can be calculated without knowing the meaning of the letters. Even a computer that cannot read Hebrew can measure the angles and curved shapes of strokes. Even when damaged letters are mixed in, values can be extracted from the remaining strokes. This is a useful property when handling ancient documents.

The two features look at different things. Hinge looks at the movement of the hand, and Fraglets looks at the learned letter form. The research team used them separately and also combined them. This was to check whether different measures gave the same answer[1].

Traces Left in One Aleph

The two features described above measure the handwriting as a whole. Now it is time to enlarge one letter and examine it closely.

The research team selected one Hebrew character for detailed examination. It was aleph. Aleph is the first letter of Hebrew, and it appears most often in the scroll. It has the shape of one diagonal stroke with two short strokes attached.

There was a reason for choosing aleph. The more samples there are, the stronger the statistics become. Only a frequent letter can be collected evenly from the front and back sections. There are 5,011 alephs in the entire scroll. The research team selected 758 of them by computer and used them for analysis[1]. They chose only those whose strokes remained complete and unbroken.

The selected alephs were divided into two groups. One group came from the front of the scroll, and the other from the back. The images in each group were overlaid in the same position to make an average shape. It is like stacking many sheets of translucent paper and looking down from above. Strokes that overlap often become darker, and strokes that differ from person to person become lighter. If the two average images are subtracted from each other, only the differences remain. The colors show which strokes contain those differences[1].

Points Divided into Two Groups

The features turned into numbers are long lists containing hundreds of values. People cannot compare such lists by eye. So the research team reduced each list to two values and plotted them as points on a plane. Handwriting with similar values gathers close together, and different handwriting moves farther away. The image is made by repeating calculations that move the positions of the points little by little.

Among the 54 columns, all columns with usable photographs were placed on this plane. Columns 16 and 46 were excluded because they did not have usable photographs. When the points were plotted, two groups appeared. The front columns formed one group, and the back columns formed another[1]. There was a visible gap between the two groups.

The research team tested whether this split was accidental. They used statistical methods such as the chi-square test and the t-test. A chi-square test calculates whether a difference is due to chance. A t-test calculates the difference between the averages of two groups. Both methods calculate the probability that a pattern like this could appear by chance. If that probability is very small, it is not treated as chance. The variation that occurs within one person's handwriting alone did not produce a difference this large[1].

It is also important that several measures gave the same answer. Whether measured with Hinge or with Fraglets, the boundary appeared in the same place. The average image of aleph also split at the same place. The boundary lay between column 27 and column 28[1]. If different methods pointed to one place, that boundary was not caused by blemishes in the photograph. The machine drew a line that the human eye had missed.

The research team also considered the opposite possibility. The result could have come from differences in photographic brightness or in the condition of the leather before and after the boundary. But even after the background was removed through binarization, the boundary remained[1]. This means the cause lay not in the leather but in the writing itself.

References

[1] Popović, M., Dhali, M. A., & Schomaker, L. (2021). Artificial intelligence based writer identification generates new evidence for the unknown scribes of the Dead Sea Scrolls exemplified by the Great Isaiah Scroll (1QIsaa). PLOS ONE, 16(4), e0249769. https://doi.org/10.1371/journal.pone.0249769

[2] Dhali, M. A., He, S., Popović, M., Tigchelaar, E., & Schomaker, L. (2017). A digital palaeographic approach towards writer identification in the Dead Sea Scrolls. Proceedings of the 6th International Conference on Pattern Recognition Applications and Methods, 693-702. https://doi.org/10.5220/0006249706930702

[3] Dhali, M. A., de Wit, J. W., & Schomaker, L. (2019). BiNet: Degraded-manuscript binarization in diverse document textures and layouts using deep encoder-decoder networks. arXiv:1911.07930. https://arxiv.org/abs/1911.07930 (accessed September 3, 2026)

[4] Dhali, M. A., de Wit, J. W., & Schomaker, L. (2026). BiNet: A deep encoder-decoder network for binarizing degraded ancient manuscripts. In Document Analysis and Recognition - ICDAR 2025 Workshops (Lecture Notes in Computer Science, pp. 166-193). Springer Nature Switzerland. https://doi.org/10.1007/978-3-032-09371-4_11

[5] University of Groningen. (2021, 4월 21일). Cracking the code of the Dead Sea Scrolls. rug.nl. https://www.rug.nl/news/2021/04/cracking-the-code-of-the-dead-sea-scrolls?lang=en (accessed September 3, 2026)

3.3. The Truth Revealed

The Hand That Changed at Column 27

The analysis supported the theory of multiple scribes. The Great Isaiah Scroll was written by two people. The first scribe wrote from column 1 to column 27. The second scribe continued from column 28 to column 54[1].

This boundary also fits other evidence. The lower part of column 27 is blank for the height of three lines. Column 28 begins on a newly joined sheet of leather. It is also the place where a new chapter of the Book of Isaiah begins[1]. It was a fitting point to hand over the work. The statistical boundary found by the computer and the physical boundary of the scroll were in the same place.

There had already been a guess that the scribe changed after column 27. It came from the place where the leather sheets were joined and the text was divided. But it did not become a scholarly consensus. Artificial intelligence analysis gave numbers to that guess[2].

The meaning of this discovery is large. Numerical evidence showed that one biblical manuscript was produced by several people working in sequence[1]. The two scribes left neither names nor signatures. What they left was only the habit of a moving hand. A machine 2,000 years later read that habit. Popovic said that their names may never be known, but that through their handwriting it felt as if we were shaking hands with them[2].

Two People Trained in the Same School

The handwriting of the two scribes still resembles each other. Even after machine measurement, the differences were slight. The research team read this resemblance itself as another piece of evidence. It means that the writing styles of the two people were close to each other[1]. The team thought the two scribes had trained in the same place or came from the same tradition[2].

Here we can picture an ancient scribal school. Trainees learned an approved script stroke by stroke. A good scribe had to suppress personal style. Because of that discipline, the writing in one scroll remained even.

Still, the habits of the body could not be erased. The angle at which a reed pen is held differs from person to person. The speed of the hand and the pressure applied also differ. These differences remain in the angle where strokes bend and in the shape of the stroke edges. The human eye cannot gather these tiny differences and judge them. A computer gathers thousands of values and calculates averages and variation. In this way, the individual traces hidden under discipline were revealed[1].

This result shows that manuscript production was collaborative work. Making one long scroll takes many months. There is no reason one person had to hold on to the work until the end. This is where the possibility arises that the work was divided and handed over[1]. The biblical text has come down to the present through such hands passing it on to other hands.

What Machines Still Cannot Do

This achievement does not mean that artificial intelligence can do everything. The next tool made by the same research team shows its limits.

In 2025, the research team introduced a dating model called Enoch. Enoch is a model that estimates when a manuscript was written by looking at the shape of the letters. Radiocarbon dating is a method that measures age from carbon left in old objects. Enoch learned from 24 samples dated by that method. The team also published results from applying it to 135 manuscripts of unknown date. The average prediction error ranged from 27.9 years to 30.7 years. When palaeographers reviewed these predictions, 79 percent were judged valid[3].

Debate arose over the remaining 21 percent. The palaeographer Christopher Rollston objected by citing one case. Enoch dated manuscript 4Q114, which contains the Book of Daniel, to between 230 BCE and 160 BCE. Rollston thought Daniel chapters 7 through 12 referred to one event. That event was the desecration of the Jerusalem Temple by Antiochus IV in 167 BCE. The text could not have been written before the event. But most of the range given by Enoch falls before 167 BCE. Rollston pointed out that Enoch's date was too early. He said Enoch should be one tool among several in the toolbox[4].

This debate tells us something. The machine sees only the shape of letters. It cannot read the historical events written in the text or the meaning of the sentences. So the machine's answer is not a final decision but a candidate. After narrowing the candidates, historians must compare the context and decide. It was also palaeographers who decided which letters should be measured. The decision to choose aleph came from long experience with ancient handwriting.

The number of training materials is also an obstacle. Enoch learned from 24 samples whose dates were known[3]. It is hard to increase the number of samples for ancient artifacts. This is because dating requires taking a small piece from the artifact and burning it. A model trained on a small sample becomes unstable outside the range it has learned. The situation is similar in scribe identification. Frequently appearing letters provide enough samples, but rare letters may appear only a few times. Where samples are few, the machine's judgment loses strength.

For that reason, the research team does not treat artificial intelligence as the only judge. The places where leather sheets were joined, the places where the text was divided, and the views scholars have built up are considered together. A conclusion stands when several pieces of evidence point to one place. The answer in the Great Isaiah Scroll came in that way. The machine merely answered, with numbers, a question that humans had asked.

The Next Target of the Investigation

The study of the Great Isaiah Scroll led to a larger plan. On June 23, 2026, the European Research Council decided to award Popovic an Advanced Grant. The research funding is 2.5 million euros, and the period is five years[5]. The new project is called Tracing Scribes and Scrolls[6]. It means following the traces of scribes and scrolls.

On June 30, 2026, the Israel Antiquities Authority announced that it would take part in this project. The University of Pisa, the University of Naples Federico II, the University of Southern Denmark, and the University of Leuven are also participating. The Egyptian museums in Berlin and Turin are also listed[6].

This time, the question moves from who wrote it to where it was made. The research team will examine about 250 samples from the Dead Sea Scrolls. It will measure the chemical composition of leather, papyrus, and ink. It will compare those components with the components of ancient Egyptian papyrus. If the origin of the materials can be known, the place where the manuscript was made can be narrowed down. Artificial intelligence will organize this chemical data and find hidden rules in the handwriting[7]. Popovic said that combining laboratory analysis with the study of ancient handwriting can bring us closer to questions that could not be reached before[6].

This method is not used only for the Great Isaiah Scroll. Among the Dead Sea Scrolls, many manuscripts still remain whose copyists are unknown. If manuscripts written by the same hand are grouped together, a work list for one scribe can be made. As such lists accumulate, we can begin to draw how the scribes were connected to one another[1]. The study that found the two people became the starting line.

Another achievement left by this study is the method itself. The sequence was organized: remove blemishes from photographs, turn strokes into numbers, and test them with statistics. The same sequence can be applied to ancient documents in other languages. The studies of clay tablets and wooden slips that appear later in this book follow a similar path.

The investigation that began with the Great Isaiah Scroll is not over yet. The tool that found the shadows of two people is now setting out to find the place where the scroll was made.

Reading ancient documents does not end with recognizing letters. We must also ask who wrote those letters and where they wrote them. The Great Isaiah Scroll is the first case in which the question of how many people wrote a scroll was answered with numbers. The scroll said nothing for 2,000 years. Humans now read its answer in the tremor of the handwriting.

References

[1] Popović, M., Dhali, M. A., & Schomaker, L. (2021). Artificial intelligence based writer identification generates new evidence for the unknown scribes of the Dead Sea Scrolls exemplified by the Great Isaiah Scroll (1QIsaa). PLOS ONE, 16(4), e0249769. https://doi.org/10.1371/journal.pone.0249769

[2] University of Groningen. (2021, 4월 21일). Cracking the code of the Dead Sea Scrolls. rug.nl. https://www.rug.nl/news/2021/04/cracking-the-code-of-the-dead-sea-scrolls?lang=en (accessed September 3, 2026)

[3] Popović, M., Dhali, M. A., Schomaker, L., van der Plicht, J., Rasmussen, K. L., La Nasa, J., Degano, I., Colombini, M. P., & Tigchelaar, E. (2025). Dating ancient manuscripts using radiocarbon and AI-based writing style analysis. PLOS ONE, 20(6), e0323185. https://doi.org/10.1371/journal.pone.0323185

[4] Steinmeyer, N. (2025, 6월 20일). Can AI date the Dead Sea Scrolls? Bible History Daily, Biblical Archaeology Society. https://www.biblicalarchaeology.org/daily/biblical-artifacts/dead-sea-scrolls/can-ai-date-the-dead-sea-scrolls/ (accessed September 3, 2026)

[5] University of Groningen. (2026, 6월 23일). Mladen Popović awarded ERC Advanced Grant to trace the origins of the Dead Sea Scrolls. rug.nl. https://www.rug.nl/about-ug/latest-news/news/archief2026/nieuwsberichten/mladen-popovi-awarded-erc-advanced-grant-to-trace-the-origins-of-the-dead-sea-scrolls?lang=en (accessed September 3, 2026)

[6] The Jerusalem Post. (2026, 6월 30일). New AI-powered research project aims to uncover the origins of the Dead Sea Scrolls. jpost.com. https://www.jpost.com/archaeology/article-900938 (accessed September 3, 2026)

[7] Archaeology News Online Magazine. (2026, 6월). Where the Dead Sea Scrolls come from: AI and chemical analysis trace origins. archaeologymag.com. https://archaeologymag.com/2026/06/dead-sea-scrolls-origins-studied-with-ai-and-chemical-analysis/ (accessed September 3, 2026)

Chapter 4. Digital Glue for Solving a Puzzle of Tens of Thousands of Clay and Parchment Pieces

4.1. Case File

A Jar Found While Looking for a Goat

In the spring of 1947, something happened on the cliffs west of the Dead Sea. Bedouin boys were looking for a runaway goat. They saw a narrow cave entrance among the rocks. Inside the cave stood tall clay jars. Inside the jars were scrolls with writing on leather[1].

These scrolls are the Dead Sea Scrolls. The Dead Sea Scrolls are Jewish texts written from around the 3rd century BCE to the 1st century CE. Over the next nine years, ten more nearby caves were found. The scrolls came from eleven caves in all. More than 900 texts came from these caves. They were written in Hebrew, Aramaic, and Greek[1]. Most were written on parchment made by stretching sheep or goat skin thin. Some copies were written on papyrus, a paper made from plant stems[7].

The scrolls included some of the oldest copies of the Hebrew Bible. Texts on community rules and calendars were also found[1]. Caves in the dry wilderness protected these letters for 2,000 years.

Near the caves is the site of an old village called Qumran. Many scholars think a Jewish community that lived there left the scrolls behind. They copied the Bible and wrote down their own rules. Scholars believe they hid the scrolls in the caves as the Roman army approached.

The Fragment Pile in Cave 4

The condition of the artifacts differed greatly from cave to cave. The scrolls in Cave 1 survived fairly well inside jars. Cave 4 was different. The scrolls in this cave were piled on the floor without jars. From this single cave came thousands of fragments from about 500 scrolls[1].

Moisture inside the cave slowly damaged the leather. Insects and animals gnawed at it. Over many years, the scrolls broke into small pieces. What remained was a pile of fragments the size of fingernails.

The Israel Antiquities Authority preserves more than 25,000 fragments[2]. Most fragments do not say which scroll they came from. It is common for only two or three letters to remain on one fragment. Scholars laid them out on large glass plates and matched them by eye.

There were only a few clues for matching them. Scholars checked whether the letter shapes looked alike. They checked whether the leather colors were similar. They tried to see whether broken edges fit together. With this method, only two or three pieces could be compared at once. When the number of fragments grows into the tens of thousands, the number of possible pairs grows out of control.

Solving the puzzle by hand took a long time. Human eyes cannot compare the edges of tens of thousands of fragments at once. The fragment edges were already worn and shrunken. Even after scholars worked on them for decades, many fragments still had not found their place.

It was also a problem that only a small group of scholars had the right to read the documents first. Outside researchers had difficulty seeing even the photographs[1]. In September 1991, the Huntington Library in the United States opened this lock. The library made its photo films available without conditions[3]. Once the photographs were opened, efforts began to match fragments by computer.

The scattering of artifacts across several countries was another obstacle. Soon after discovery, some scrolls were sold through an antiquities dealer in Bethlehem[1]. Institutions in countries that funded excavations also received shares of fragments. As a result, fragments from one scroll ended up in different cities. To see whether two fragments joined, a person had to take a plane.

Fake Fragments on the Market

As the price of Dead Sea Scroll fragments rose, another problem appeared. Around 2002, objects described as Dead Sea Scroll fragments began to appear on the antiquities market. About 70 pieces with Bible verses appeared[4]. Scholars called them “post-2002 fragments.” These objects had no records showing which cave they came from or when they were found[6].

The Museum of the Bible in Washington, D.C., also bought and displayed 16 such fragments. Several scholars suspected that the letter forms looked awkward. The museum asked an art forgery investigation company to carry out a detailed examination. The investigation was led by investigator Colette Loll[4].

The results came out in March 2020. All 16 pieces were modern fakes[4][5]. The team examined the surfaces with a 3D microscope. The material was not the thin parchment of real Dead Sea Scrolls. At least 15 pieces were thick leather with a coarse grain. The ink had pooled inside cracks that were already present[4]. This was evidence that the writing was added after the leather had cracked.

This case shook the whole field of biblical archaeology. Professor Christopher Rollston of George Washington University says fakes are still coming onto the market. He often meets people who believe their pieces are real, but he says almost all of them are fake. This shows the difference between artifacts with excavation records and artifacts bought on the market[6].

Matching broken fragments and filtering out fakes have the same root. Both require reading tiny physical traces that human eyes cannot catch. Artificial intelligence has entered this work. In June 2026, the European Research Council assigned new funding. It gave 2.5 million euros to a five-year study on identifying the origins of the Dead Sea Scrolls. The study is called “Tracking the Scribes and the Scrolls.” It combines material analysis, artificial intelligence, and the study of ancient handwriting[7].

This chapter follows two cases. One is reconnecting scattered fragments. The other is filtering out fakes. The two tasks ask different questions about the same materials. The first question is whether these two pieces were originally one. The second question is whether this piece is truly ancient.

References

[1] Israel Antiquities Authority. (n.d.). Discovery and publication. The Leon Levy Dead Sea Scrolls Digital Library. https://www.deadseascrolls.org.il/learn-about-the-scrolls/discovery-and-publication (accessed September 3, 2026)

[2] Israel Antiquities Authority. (n.d.). The digital library. The Leon Levy Dead Sea Scrolls Digital Library. https://www.deadseascrolls.org.il/about-the-project/the-digital-library (accessed September 3, 2026)

[3] Deseret News. (1991, 9월 22일). Research library will give scholars access to Dead Sea Scroll photos. Deseret News. https://www.deseret.com/1991/9/22/18942605/research-library-will-give-scholars-access-to-dead-sea-scroll-photos/ (accessed September 3, 2026)

[4] Greshko, M. (2020, 3월 13일). "Dead Sea Scrolls" at the Museum of the Bible are all forgeries. National Geographic. https://www.nationalgeographic.com/history/article/museum-of-the-bible-dead-sea-scrolls-forgeries (accessed September 3, 2026)

[5] Solly, M. (2020, 3월 16일). All of the Museum of the Bible's Dead Sea Scrolls are fake, report finds. Smithsonian Magazine. https://www.smithsonianmag.com/smart-news/all-museum-bibles-dead-sea-scrolls-are-fake-report-finds-180974425/ (accessed September 3, 2026)

[6] Govier, G. (2026, 4월 8일). New Dead Sea Scrolls exhibit is the real deal. Christianity Today. https://www.christianitytoday.com/2026/04/new-dead-sea-scrolls-exhibit-is-the-real-deal/ (accessed September 3, 2026)

[7] The Jerusalem Post. (2026, 6월 30일). New AI-powered research project aims to uncover the origins of the Dead Sea Scrolls. The Jerusalem Post. https://www.jpost.com/archaeology/article-900938 (accessed September 3, 2026)

4.2. Forensic Tracking 1: Geometric Edge Affinity and Physical Fingerprints

Turning Torn Edges into Numbers

The first clue for matching broken fragments is the shape of the edges. When paper is torn by hand, an uneven sawtooth line appears. The lines of two neighboring fragments fit into each other. People compare these lines by eye. Computers turn these lines into numbers.

A computer follows an edge and measures its curvature. When the degrees of curvature are arranged in order, they become a wave-like sequence of numbers. If the sequences of two fragments fit well when one is flipped and overlaid on the other, they become a candidate match. This method is called edge affinity calculation.

In Korean, this score can be called an edge interlocking score. It is a number that shows how well the edges of two fragments fit together.

The limits of looking only at edges are clear. The edges of ancient fragments are already worn. Many places have shrunk as they dried and lost their original shape. So researchers also look at clues beyond the edges.

Fingerprints Left by the Material Itself

Edges alone are not enough. The edges of fragments are often already worn. So researchers look at the material itself.

Papyrus is made by cutting plant stems into thin strips and layering them crosswise. This leaves stem patterns like grain on the surface. In two fragments that were originally one sheet, this grain continues across the join[1].

A team led by Roy Avitbol at the University of Haifa in Israel read this grain. The team cut the areas near fragment edges into small squares. A neural network looked at two squares and scored whether the grain continued. The team gathered these scores and judged matching pairs across whole fragments[1].

The test target was a set of Dead Sea Scroll fragments in poor condition. Because the correct answers were unknown, many wrong candidates also appeared. Still, the method removed 98 percent of pairs that could not match[1]. The amount people had to examine fell to one fiftieth.

The meaning of this result is clear. The computer did not confirm the pairs. The computer removed pairs that did not need to be checked. In fragment matching, this filtering is usually the role of artificial intelligence.

A team led by Antoine Pirrone in Bordeaux, France, used a Siamese neural network for this. A Siamese neural network takes two images side by side and judges whether they have the same origin. The team named the model Papy-S-Net. The model learned from about 500 papyrus fragments. It chose matches by looking only at paper grain, not letters. In real data, it found 79 percent of the matching pairs[2].

Parchment has similar fingerprints. Leather preserves marks from pores where hair once grew. Fragments from the same animal skin continue the same pattern of marks. Holes made by insects boring through a scroll are also clues. When a rolled scroll is opened, insect holes repeat in lines. The inner layers of a scroll are shorter, and the outer layers are longer. So the spacing between holes becomes slightly wider toward the outside. Measuring this spacing can suggest which layer of the scroll a fragment came from.

To handle these physical fingerprints, the photographs must first be cleaned up. Photographs of the Dead Sea Scrolls held by the Israel Antiquities Authority include rulers, color charts, and number plates. The background is black, so it can be confused with ink. A team led by Bronson Brown-deVost created a four-step processing method that removes these distractions first. The team also released a labeled dataset[3].

Releasing a labeled dataset is important. Other researchers can test their own methods on the same data. They can also compare which of two methods is better using numbers. The study of ancient documents is becoming a field where researchers verify one another's work in this way.

Technology Spreading to Clay Tablets, Bones, and Paper

The same idea moved to other artifacts. At Ludwig Maximilian University of Munich in Germany, there is Professor Enrique Jimenez. He leads the electronic Babylonian Library (eBL) project. The fragment matching tool of this project is the Fragmentarium[4].

Mesopotamians pressed cuneiform signs into wet clay with reed pens. These clay tablets break easily. The research team photographed tablets at the British Museum and the Iraq Museum. Since 2018, it has handled about 22,000 fragments as data[5]. The computer compares strings transcribed from the writing left on fragments. For this, the team borrowed methods used in biology to compare gene sequences. This was because ancient scribes used different spellings from region to region[4].

The results appeared in numbers. In five years, this project found about 1,500 new joins. Over the previous 150 years, scholars had found about 5,000 joins[4].

Clay tablets have one more clue that parchment does not have. It is the three-dimensional shape of the broken surface. Clay tablets were usually dried in the sun, and only some were fired. As they hardened, their interiors became firm unevenly. So each broken surface developed different curves. If these surfaces are measured in 3D, matching pairs can be found.

In China, the same technology is used for oracle bone fragments. Oracle bones are records that Shang dynasty people carved on turtle plastrons and ox bones while divining the future. A team led by Zhang Chongsheng at Henan University created an oracle bone joining dataset. Its joining algorithm achieved a top-10 accuracy of 98.39 percent[6]. A team led by Zhang Zan also trained a deep rejoining model with photographs of oracle bone fragments. This team also released its dataset[7].

The phrase top-10 accuracy needs explanation. The machine does not give only one answer. It provides the ten most likely candidates in order. The probability that the correct answer is among those ten is 98.39 percent[6]. The final judgment is made by a person.

Work to gather data in one place also took place. Anyang Normal University operates an oracle bone data platform called Yinqiwenyuan. As of April 2023, it contained more than 230,000 oracle bone photographs. It also gathered 152 works and more than 3,000 papers related to oracle bone script[8].

Buddhist scripture fragments from the Mogao Caves in Dunhuang are also scattered across countries. In 2026, Chen Mingkun's team introduced a new method. This method learns by itself even without human-labeled answers. This learning approach is called self-supervised learning. The team gathered complete manuscripts and fragments with known matches. It also added isolated fragments whose matches were unknown. The computer created its own practice problems from this data. The team designed a new coding method that also captured edge direction. Then it used a Siamese neural network to score how similar two fragments were[9].

For the Dead Sea Scrolls, European and Israeli researchers built an online workbench together. Its name is Scripta Qumranica Electronica. On this workbench, researchers can move, rotate, and arrange fragment photographs on the screen. They can share their arrangements with other researchers[10]. Fragment matching has changed from work inside one person's head into work that several people can review together.

References

[1] Abitbol, R., Shimshoni, I., & Ben-Dov, J. (2021). Machine learning based assembly of fragments of ancient papyrus. Journal on Computing and Cultural Heritage, 14(3), 1-21. https://doi.org/10.1145/3460961

[2] Pirrone, A., Beurton-Aimar, M., & Journet, N. (2019). Papy-S-Net: A Siamese network to match papyrus fragments. Proceedings of the 5th International Workshop on Historical Document Imaging and Processing, 78-83. https://doi.org/10.1145/3352631.3352646

[3] Brown-deVost, B., Kurar-Barakat, B., & Dershowitz, N. (2024). Segmenting Dead Sea Scroll fragments for a scientific image set. arXiv. https://arxiv.org/abs/2406.15692

[4] Jiménez, E. (2023, 5월 25일). AI is helping us read ancient Mesopotamian literature. The Conversation. https://theconversation.com/ai-is-helping-us-read-ancient-mesopotamian-literature-205091 (accessed September 3, 2026)

[5] Phys.org. (2023, 2월 2일). Researcher uses AI to make texts that are thousands of years old readable. Phys.org. https://phys.org/news/2023-02-ai-texts-thousands-years-readable.html (accessed September 3, 2026)

[6] Zhang, C., Zong, R., Cao, S., Men, Y., & Mo, B. (2020). AI-powered oracle bone inscriptions recognition and fragments rejoining. Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, 5309-5311. https://doi.org/10.24963/ijcai.2020/779

[7] Zhang, Z., Zhang, H., Guo, A., Jiao, Q., Liu, Y., Li, B., & Gao, F. (2025). Deep rejoining model and dataset of oracle bone fragment images. npj Heritage Science, 13(1), 66. https://doi.org/10.1038/s40494-025-01651-9

[8] 王者. (2023, 4월 8일). "殷契文渊"数据平台为甲骨文研究提供助力. 人民日报. https://news.hangzhou.com.cn/gnxw/content/2023-04/08/content_8509027.htm (accessed September 3, 2026)

[9] Chen, M., Xiang, Y., & Zheng, Y. (2026). A self-supervised learning approach for pairwise matching of ancient Dunhuang fragments. Journal on Computing and Cultural Heritage, 19(1), 1-19. https://doi.org/10.1145/3776642

[10] Scripta Qumranica Electronica. (n.d.). Scripta Qumranica Electronica. University of Haifa. https://megillot.haifa.ac.il/index.php/en/research/scripta-qumranica-electronica (accessed September 3, 2026)

4.3. Forensic Tracking 2: Forgery Detective AI

How 16 Fakes Were Caught

All 16 Dead Sea Scroll fragments purchased by the Museum of the Bible were fake[1][2]. The investigation team did not look only at the letters with the naked eye. It examined the materials and ink with the methods of physics and chemistry. Artificial intelligence was not used in this authentication. People detected the forgery with microscopes and material analysis[1].

The Museum of the Bible's authentication did not end in one step. In October 2018, Germany's Federal Institute for Materials Research and Testing first examined five fragments. These five did not match the features of ancient artifacts, so they were removed from display[5]. The results for all the remaining fragments came out in 2020[1][2].

The first piece of evidence was the base material. The real Dead Sea Scrolls were written on thinly prepared parchment. At least 15 of the Museum of the Bible fragments were leather. This leather was thicker than parchment and had a rougher grain[1].

The second piece of evidence was where the ink lay. The investigation team looked at the surface with a 3D microscope. Ink had pooled inside cracks. On torn edges, traces showed that ink had run downward[1].

Why these traces were decisive becomes clear when we think about the order of events. A real manuscript was first written on flat parchment. Then the leather dried and cracked over a long period.

So in a real manuscript, letter strokes break at the cracks. In a fake, the order is reversed. The forger obtained old leather that was already cracked and put ink on it. The ink flowed into the open cracks[1].

The third piece of evidence was a substance spread on the surface. The investigation team found traces showing that the fragments had been soaked in an amber-colored liquid. This liquid appears to have been animal-skin glue. The forger was trying to imitate the glossy feel of the surface of the real Dead Sea Scrolls[1].

The three kinds of evidence pointed to the same answer from different directions. One piece alone could be a coincidence. When three overlap, it is not a coincidence. This is why artifact authentication gathers different kinds of evidence together.

Traces Left by the Hand

A person who makes a fake looks at a photograph of the real thing and copies the letters. This act of copying leaves traces.

When people write familiar letters, the hand flows without stopping. It is different when they copy unfamiliar letter shapes. They stop while drawing a stroke and check where the next stroke should go. The speed of the brush tip drops, and the line trembles slightly. When enlarged, this difference appears in the outer shape of the stroke.

A computer can detect these differences as numbers. In June 2026, a research team led by Hassan Ugail at the University of Bradford in the United Kingdom published a related study. The team built an artificial intelligence system to judge whether old drawings were genuine. This model looks at only five numbers extracted from a drawing. It measures how much patterns repeat and how widely brightness values are spread. It also measures the contrast between bright and dark areas. It adds how even the texture is and how much lines branch like small twigs[3].

The team trained the model on genuine drawings by ten artists. The dates of the works ranged from around 1510 to 1917. The data came from six museums that provide open data. The Metropolitan Museum of Art and the Ashmolean Museum were among them[3].

The paper reported the results of 900 classification tests. The balanced accuracy was 87.6 percent. The rate at which genuine works were accepted as genuine was 77.8 percent. The rate at which works by others were wrongly accepted as that artist's work was 2.6 percent[3].

This study has one more result worth noting. Large deep learning models that are widely used actually performed poorly on this task. The large models often rejected genuine works as fake. When data was scarce, a small model designed for the problem worked better[3].

This result is useful for ancient manuscript research as well. Few ancient manuscripts survive. There is not enough data to feed a large model well. This means a design that tells the model in advance what to look at can be better.

A Machine That Measures Handwriting Style

The same technology is used to study genuine artifacts as well. A research team led by Mladen Popović at the University of Groningen in the Netherlands analyzed the handwriting styles of the Dead Sea Scrolls with artificial intelligence. The team used a neural network to extract only the traces where ink had passed. Then it converted the curvature and slant of letters into numbers[4].

The team trained the model on 24 manuscripts whose radiocarbon dates had already been measured. The model's name is Enoch. Enoch estimated the ages of 135 manuscripts whose dates were unknown.

The average error was about 28 to 31 years. Several manuscripts turned out to be older than previously known[4].

This technology also helps with authentication. A real scribe's handwriting has consistent statistical features in each period. A fake falls outside this distribution. The human eye feels that two writings look alike, but the machine presents the difference as numbers.

However, this help is only another reference point. The Enoch study needs caution. This model was trained to estimate dates. It was not trained to distinguish the genuine from the fake. An unusual date estimate does not immediately mean that an item is fake. A machine's result has meaning only within the scope of the task it learned.

What Machines Cannot Replace

Machine judgment has limits. The final authentication of the Museum of the Bible fragments was not made by artificial intelligence alone. Microscopic observation, chemical testing, and philological knowledge were used together[1][2].

The weight of provenance records remains the same. A provenance record is a record of where an artifact came from. Professor Rollston says we should distinguish artifacts with excavation records from artifacts bought on the market. A scroll found through excavation has a record of which cave it came from and when. A fragment that appears on the market has no such record[5].

Machines read subtle signals that people miss. People judge what those signals mean. The work of detecting forgeries is complete only when both are present.

References

[1] Greshko, M. (2020, 3월 13일). "Dead Sea Scrolls" at the Museum of the Bible are all forgeries. National Geographic. https://www.nationalgeographic.com/history/article/museum-of-the-bible-dead-sea-scrolls-forgeries (accessed September 3, 2026)

[2] Solly, M. (2020, 3월 16일). All of the Museum of the Bible's Dead Sea Scrolls are fake, report finds. Smithsonian Magazine. https://www.smithsonianmag.com/smart-news/all-museum-bibles-dead-sea-scrolls-are-fake-report-finds-180974425/ (accessed September 3, 2026)

[3] Ugail, H., Ritch-Frel, J., Matuzava, I., & Stork, D. G. (2026). Verification of historical sketches via one-class learning on compact feature representations. PLOS One, 21(6), e0344796. https://doi.org/10.1371/journal.pone.0344796

[4] Popović, M., Dhali, M. A., Schomaker, L., van der Plicht, J., Lund Rasmussen, K., La Nasa, J., Degano, I., Perla Colombini, M., & Tigchelaar, E. (2025). Dating ancient manuscripts using radiocarbon and AI-based writing style analysis. PLOS One, 20(6), e0323185. https://doi.org/10.1371/journal.pone.0323185

[5] Museum of the Bible. (2018, 10월 22일). Museum of the Bible releases research findings on fragments in its Dead Sea Scrolls collection. Museum of the Bible. https://www.museumofthebible.org/newsroom/museum-of-the-bible-releases-research-findings-on-fragments-in-its-dead-sea-scrolls-collection (accessed September 3, 2026)

4.4. The Truth Revealed

The Line Between Confirmed Joins and Candidate Joins

Even if a computer says two fragments fit together, the work is not over. The machine only gives a probability. Researchers strictly distinguish confirmed joins from candidate joins.

A confirmed join is accepted only when several kinds of evidence match together. First, the shapes of the broken edges must interlock without gaps. Next, the grain of the material must continue. In papyrus, the stem pattern must continue[1]. In parchment, the hair-follicle marks must continue. Then the letter strokes that cross the boundary must connect naturally. Finally, the joined text must fit the grammar of that language.

A candidate join is a case in which only some of these conditions fit. A result that estimates the position of fragments from the spacing of wormholes belongs here. A result that groups fragments into one scroll group because the ink composition is the same is also a candidate. Such results become the starting point of research, but not its conclusion. The sequence is that people check the candidates filtered by the computer[1].

It is dangerous not to keep this distinction. If two fragments are joined simply because the probability is high, a sentence that never existed can be created. If that sentence is cited, the error spreads into the next study. That is why records keep confirmed joins and candidates separate.

The numbers produced by a machine also assume this distinction. A high filtering rate does not confirm a join[1]. The same is true of the probability that the correct answer is included among the top candidates[2]. In either case, people make the final judgment.

Texts Brought Back to Life

Actual texts have been brought back to life in this way. Fragmentarium has handled about 22,000 clay tablet fragments since 2018. This program found several new joins and new manuscripts. In February 2023, the research team opened Fragmentarium to the public. It also released a digital edition of the Epic of Gilgamesh[3].

Among the revived texts is the Epic of Gilgamesh. The Epic of Gilgamesh is a Mesopotamian hero story. In November 2022, the program recognized one fragment. This fragment came from a copy made around 130 BCE. It is late among the Gilgamesh manuscripts identified so far. This means that people were still copying this story at that time[3].

Cuneiform spelling differed by region. This was because scribes changed signs to match their own dialects. When people searched by eye, this difference was a major barrier. The research team crossed that barrier by borrowing a method used to compare genetic sequences. It became possible to compare strings that differed slightly[4].

The data base for the Dead Sea Scrolls has also expanded. The Israel Antiquities Authority photographed and released more than 25,000 fragments in high resolution[5]. Researchers move and arrange fragments on these photographs[6]. A processing method has also been released to automatically remove obstacles such as rulers and color charts[7].

The Next Door Opened in 2026

The “Tracking the Scribes and Scrolls” study, which began in June 2026, opens the next door. The question this study asks goes one step beyond fitting fragments together. It asks where and by whom the Dead Sea Scrolls were written[8].

The research team examines about 250 samples of parchment, papyrus, and ink. A sample is a small specimen used for testing. The team combines this with artificial intelligence handwriting analysis and the study of ancient handwriting. The Israel Antiquities Authority, the University of Pisa, and the University of Naples are working together. The University of Southern Denmark and KU Leuven in Belgium are also participating. The team also compares the materials with Egyptian papyri held by museums in Berlin and Turin[8].

If samples have the same composition, the likelihood grows that they came from the same workshop. This information is also used to fit fragments together. If items from the same material source are grouped first, the number of pairs to compare decreases.

Remaining Fragments and Remaining Work

Even though achievements have accumulated, many fragments still have not found their proper place. In the storage boxes of the Antiquities Authority, fragments remain whose original scrolls are unknown[8]. On the clay tablet side, many joins have not yet been found[4].

Differences in equipment remain a problem as well. To measure an artifact with a 3D scanner, the artifact must be moved in front of the equipment. Many museums find it difficult to send artifacts outside. So some artifacts have only photographic data and no 3D data. For these artifacts, methods that compare broken edges cannot be used.

Technology has also left tasks to solve. Fragment photographs were taken with different equipment in different countries. If lighting and backgrounds differ, the same algorithm gives different answers. So researchers created open datasets with correct answers attached. They did this so results could be compared by the same standard[7]. Dunhuang fragment research also moved toward learning without people attaching the correct answers[9].

This technology does not push people out. Machines reduce the number of pairs to compare, and people judge meaning. Only people can read whether a revived sentence fits the speech of that period. Because the two roles do not overlap, they can work together.

Access has changed as well. In the past, only a small number of scholars saw the original photographs. Now clay tablet data and Dead Sea Scroll photographs are open on the internet[4][5]. Students can also move fragments around on a screen[6].

Digital glue does not physically attach fragments. The artifacts remain in their boxes. The scroll is rejoined only on the computer screen. So if a judgment is wrong, it can always be reversed. Texts scattered for 2,000 years are slowly returning to sentences in this way.

References

[1] Abitbol, R., Shimshoni, I., & Ben-Dov, J. (2021). Machine learning based assembly of fragments of ancient papyrus. Journal on Computing and Cultural Heritage, 14(3), 1-21. https://doi.org/10.1145/3460961

[2] Zhang, C., Zong, R., Cao, S., Men, Y., & Mo, B. (2020). AI-powered oracle bone inscriptions recognition and fragments rejoining. Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, 5309-5311. https://doi.org/10.24963/ijcai.2020/779

[3] Phys.org. (2023, 2월 2일). Researcher uses AI to make texts that are thousands of years old readable. Phys.org. https://phys.org/news/2023-02-ai-texts-thousands-years-readable.html (accessed September 3, 2026)

[4] Jiménez, E. (2023, 5월 25일). AI is helping us read ancient Mesopotamian literature. The Conversation. https://theconversation.com/ai-is-helping-us-read-ancient-mesopotamian-literature-205091 (accessed September 3, 2026)

[5] Israel Antiquities Authority. (n.d.). The digital library. The Leon Levy Dead Sea Scrolls Digital Library. https://www.deadseascrolls.org.il/about-the-project/the-digital-library (accessed September 3, 2026)

[6] Scripta Qumranica Electronica. (n.d.). Scripta Qumranica Electronica. University of Haifa. https://megillot.haifa.ac.il/index.php/en/research/scripta-qumranica-electronica (accessed September 3, 2026)

[7] Brown-deVost, B., Kurar-Barakat, B., & Dershowitz, N. (2024). Segmenting Dead Sea Scroll fragments for a scientific image set. arXiv. https://arxiv.org/abs/2406.15692

[8] The Jerusalem Post. (2026, 6월 30일). New AI-powered research project aims to uncover the origins of the Dead Sea Scrolls. The Jerusalem Post. https://www.jpost.com/archaeology/article-900938 (accessed September 3, 2026)

[9] Chen, M., Xiang, Y., & Zheng, Y. (2026). A self-supervised learning approach for pairwise matching of ancient Dunhuang fragments. Journal on Computing and Cultural Heritage, 19(1), 1-19. https://doi.org/10.1145/3776642

Chapter 5. Fitting the Broken Edges of Clay Tablets: Rebuilding the Epic of Gilgamesh

5.1. Case File

A Story That Crossed 4,000 Years

The Epic of Gilgamesh is among the oldest works of narrative literature that survive today. Earlier Sumerian songs have also come down to us. The main character is Gilgamesh, the king who ruled the Mesopotamian city of Uruk[1]. He goes with his friend Enkidu to a distant cedar forest. The two fight Humbaba, the monster who guards that forest[2]. Later, after Enkidu dies, Gilgamesh sets out to find a way not to die[1].

The material that carried this story was not paper, but clay tablets. Ancient scribes cut reeds at an angle to make pointed sticks. These sticks are called styluses. When a stylus is pressed into wet clay, it leaves wedge-shaped marks. These are characters made by pressing, not by drawing lines. When the clay dries, the marks harden as they are. That is why this writing is called cuneiform. Students training to become scribes copied this epic as writing practice. Thanks to that, manuscript fragments have been found at many sites across the Middle East. Fragments of this epic were also found at Amarna in Egypt. This means the story had spread far from Mesopotamia[1].

The surviving versions fall into two broad groups. One is the Old Babylonian version, written in the early second millennium BCE. The other is the Standard Babylonian version, arranged later. The standard version is made up of twelve clay tablets. Eleven of them form the main story. The twelfth tablet is seen as an appendix added later. The translations we read today are based on this standard version. Gaps are filled with the Old Babylonian version and other fragments. The person who refined the standard version is known as a scribe named Sin-leqi-unninni. He wove scattered old songs into one long work. He decided what each tablet would contain and matched the number of lines. Yet even this standard version has many gaps. No manuscript survives intact from beginning to end[1].

The story still feels familiar to readers today. Gilgamesh was a powerful king, but he tormented his people. The gods created Enkidu and sent him as someone who could stand against him. The two fight, then become friends. Together they defeat Humbaba and also bring down the Bull of Heaven. As the price, Enkidu falls ill and dies. After seeing his friend die, Gilgamesh realizes that he too will die. Tablet XI contains the scene in which he visits a man who survived a great flood and hears his story[1].

This story was forgotten for more than 2,000 years. The people who could read cuneiform had disappeared. In the 19th century, scholars deciphered cuneiform again, and the epic returned to the world. It began when a researcher searching through boxes of clay tablets brought from Nineveh recognized the flood story. For the next 150 years, scholars gathered scattered manuscripts and filled the text little by little[1].

Nineveh and the Fire of 612 BCE

Most manuscripts of the standard version came from one ancient library. It was the collection that Ashurbanipal, king of Assyria in the 7th century BCE, built in the palace at Nineveh. This library gathered works on divination, medicine, dictionaries, and literature. The king had clay tablets collected from across the empire. He sent letters to his officials and ordered them to find books that were missing. Ashurbanipal was a king who could read cuneiform himself[3].

Around 612 BCE, Nineveh was swept by a great fire. The palace collapsed, and the library burned with it. This fire did two things to the artifacts at once. Clay tablets that had only been dried in the shade were fired hard in the flames[4]. The baked tablets gained a durable body, like bricks from a kiln. If papyrus or leather had been in the same fire, it would have turned to ash and disappeared.

But the weight of the collapsing buildings shattered those hard tablets. The British Museum now holds more than 30,000 tablets and fragments excavated from Nineveh[3][4]. These are artifacts that excavation teams collected from mounds of earth and loaded onto ships in the mid-19th century[3]. A single story came down to us scattered into dozens of fragments.

The fragments vary in size. Some are as large as a palm. Others are as small as a thumbnail. In many cases, only a few lines of writing remain on a fragment. No one tells us where those few lines belong among the twelve tablets of the epic. Reading the writing and finding its position are different tasks.

When one fragment finds its place, the meaning can change. Only when the front and back connect can we know who is speaking and which scene it is. Writing that looked like a list of obscure words when isolated becomes a sentence. A single join can revive an entire scene of the story.

350,000 Lines Cannot Be Matched by Hand

For a long time, matching broken fragments depended on human hands and memory. A scholar picks up a fragment in a storage room. Then the scholar recalls photographs and hand copies of other fragments and looks for a match. If two fragments are judged to have come from the same tablet, that is called a join.

This method runs into the limits of human memory. In the 150 years since Assyriology began, about 5,000 joins have been made this way[5]. The number of remaining fragments is far larger. The British Museum alone holds more than 30,000 tablets from Nineveh[3]. If other collections are added, the number grows much higher. Many of them have not been read since antiquity[5]. If all surviving cuneiform tablets in the world's museums are counted, the number is estimated at about 500,000[6].

There is another reason the work is hard. The fragments are scattered across several countries. The British Museum, Berlin, Yale University, and the Penn Museum in Philadelphia each hold their own collections[7]. To test whether two fragments fit together in person, one may have to cross two continents.

Cuneiform itself also takes a great deal of work. Different scribes wrote the same word in slightly different ways. Spelling shifts by region and period[5]. Even when a person compares the writing on two fragments by eye, it is hard to notice that they belong to the same work.

A single cuneiform sign can also be read with several sounds. The same mark has different values depending on context. Searching for matches among hundreds of thousands of lines of such writing cannot be completed by human eyes alone.

That is why Professor Enrique Jimenez at Ludwig Maximilian University of Munich in Germany chose another path. In 2018, he began a project to restore Babylonian literature as fully as possible. The work brings readings of clay tablets scattered around the world into one place[5]. The project is called the electronic Babylonian Library, or eBL. The fragment-matching workspace created by eBL is Fragmentarium. By 2024, this workspace held readings of about 25,000 clay tablets. Counted by lines, it contained more than 350,000 lines[7].

The urgency comes from time. Even if clay tablets were fired, they do not last forever. Salt and moisture swell the surface and erase the writing. Hundreds of thousands of tablets are still waiting to be read in museum storage rooms[5]. No one knows how many stories they contain.

The number of people who can read cuneiform is also a problem. There are not many Assyriologists in the world. Only some of the surviving clay tablets have been translated[6]. There is a limit to how many tablets one person can read in a lifetime. This is why researchers sought help from machines.

References

[1] George, A. R. (2003). The Babylonian Gilgamesh Epic: Introduction, Critical Edition and Cuneiform Texts. Oxford University Press.

[2] Al-Rawi, F. N. H., & George, A. R. (2014). Back to the Cedar Forest: The beginning and end of Tablet V of the Standard Babylonian Epic of Gilgameš. Journal of Cuneiform Studies, 66, 69-90. https://doi.org/10.5615/jcunestud.66.2014.0069

[3] The Ashurbanipal Library Project. (n.d.). What is the Library? Oracc, University of Pennsylvania. https://oracc.museum.upenn.edu/asbp/whatisthelibrary/index.html (accessed September 3, 2026)

[4] British Museum. (n.d.). What was Ashurbanipal's Library? britishmuseum.org. https://www.britishmuseum.org/research/projects/what-was-ashurbanipals-library (accessed September 3, 2026)

[5] Jiménez, E. (2023, 5월 25일). AI is helping us read ancient Mesopotamian literature. The Conversation. https://theconversation.com/ai-is-helping-us-read-ancient-mesopotamian-literature-205091 (accessed September 3, 2026)

[6] Wang, Z. (2026). TabletCraft: Bridging a 4,000-year cultural gap with bidirectional Akkadian NMT and cuneiform rendering. arXiv:2608.02609. https://arxiv.org/abs/2608.02609 (accessed September 3, 2026)

[7] Cobanoglu, Y., Laasonen, J., Simonjetz, F., Khait, I., Cohen, S., Földi, Z., Hätinen, A., Heinrich, A., Mitto, T., Rozzi, G., Sáenz, L., & Jiménez, E. (2024). Transliterated cuneiform tablets of the electronic Babylonian Library platform. Journal of Open Humanities Data, 10, 19. https://doi.org/10.5334/johd.148

5.2. Forensic Tracking: Clay-Edge Digital Fusion and Fragmentarium

The First Step: Turning Signs into the Alphabet

Fragmentarium begins with the characters. Researchers look at photographs of clay tablets and write the cuneiform in Roman letters[1]. This work is called transliteration. In Korean terms, it is similar to writing the sounds of Chinese characters in Hangul. Each cuneiform sign is converted into alphabet letters according to its sound value. Places that are broken and unreadable are marked with gap symbols.

The transliterations created this way become strings that computers can handle. eBL accumulated these transliterations from May 2018 to August 2023. As a result, about 25,000 clay tablets and more than 350,000 lines of text became searchable data. The catalog of tablet records managed by eBL contains more than 260,000 entries. It includes collections from the British Museum, the Yale Babylonian Collection, and the Penn Museum[1].

Transliteration is done by people. Machines still cannot extract the writing from photographs perfectly. The eBL team continued this transliteration work for more than five years. The result was released as public data that anyone can download[1]. To make a machine work, people must first create the data.

String Matching That Resembles Genetic Testing

The core algorithm came from an unexpected field. It is a string alignment method used by biologists to compare genetic sequences[2]. A gene is a long chain made of the four letters A, T, G, and C. When comparing the genes of two organisms, researchers align them in the same positions even if some letters differ.

The same problem appears in cuneiform. Scribes wrote the same word in different ways. Professor Jimenez saw this variation as resembling genetic mutations. So he adapted a biological alignment algorithm for cuneiform[2].

What the computer actually counts are short pieces of text. Pieces cut from a transliteration a few characters at a time are called n-grams. The system counts how many n-grams two tablets share. Pairs that share longer pieces receive higher scores. High-scoring pairs are presented to people as candidates.

Fragmentarium opens as a web interface. Researchers anywhere in the world can log in and see the same data. A fragment in London and a fragment in Munich can be placed side by side on one screen and compared[8]. There is no longer any need to pack a bag and board a plane.

This method can withstand spelling variation. Even if one word is written differently, a score remains if the surrounding pieces overlap. eBL has placed this matching program in a public repository. Other researchers can download and use the same tool[6].

Computers also help with the act of reading characters. A single cuneiform sign is often read with several sounds. The model trained by eBL reached 98 percent accuracy in choosing the right reading among these possible values[2].

An example makes clear why this is needed. The same wedge mark may be read as one sound in one sentence. In another sentence, it may be read as a sign representing an entire word. If it is read incorrectly, the matching score becomes distorted. Correcting the reading directly raises the accuracy of joins.

The results appeared in numbers. Over 150 years, Assyriology made about 5,000 joins by hand. In five years, eBL found about 1,500 more joins[2]. In the process of creating the transliteration corpus alone, about 1,250 joins were revealed[1]. Thousands of fragments that do not physically touch but belong to the same works were also identified[2].

Language Models That Fill Broken Places

Even after fragments are joined, the text still has gaps here and there. These gaps are called lacunae. From this point, language models step in.

In 2023, researchers at Tel Aviv University and Ariel University released an Akkadian translation model. This model translates into English whether the input is cuneiform signs or alphabetic transliteration. It handles both kinds of input. It did not handle translation back from English into Akkadian. The training data consisted of already published readings of clay tablets[3].

In 2026, a tool appeared that also handled the opposite direction. It is open software called TabletCraft. This tool translates both from Akkadian to English and from English to Akkadian. The training data contains 116,000 sentence pairs. The cuneiform sign conversion table contains more than 14,000 correspondences. The developer said tools are needed to read the remaining 500,000 clay tablets[4].

Such models look at the context before and after a blank and offer likely words as probabilities. Babylonian poetry has a rule of placing two lines with the same structure side by side. Models use this form as a clue too. What the model provides is not a fixed answer, but a list of candidates.

A 2026 Tool That Reads Handwriting Style Too

On May 18, 2026, the University of Würzburg in Germany unveiled a new tool. It was created together with the Mainz Academy of Sciences and Literature and TU Dortmund University. Its name is Paleograficum. This tool learns the shape of each cuneiform sign from photographs of clay tablets. Then it finds signs impressed in the same way across the entire collection[5].

This tool handles 70,000 photographs and more than 5 million signs. In the past, comparing the handwriting on five fragments took three days. With this tool, it takes five minutes. Just as each person has different handwriting, each scribe pressed the wedges at different angles and intervals. Those differences can be used as clues to pick out fragments from the same tablet[5].

The Hittite cuneiform handled by this tool has 375 signs. So far, about 30,000 Hittite clay tablet fragments are known. Professor Daniel Schwemer, who led the research, said this tool could save thousands of hours[5].

The researchers' next goal is to identify individual scribes. This means finding all clay tablets written by the same person. If that becomes possible, scholars can map where one scribe worked and where he moved. The working lives of people from 3,500 years ago would appear in the data[5].

Another way to fit broken clay edges together

Research has long continued on matching fragments by shape rather than by letters. A broken clay tablet leaves broken surfaces. These surfaces are called fracture surfaces. If the fracture surfaces of two fragments interlock, they came from the same tablet.

A leading project on this path is the Virtual Cuneiform Tablet Reconstruction Project. Sandra Woolley of Keele University in the United Kingdom leads the project. Tim Collins and Erlend Gehlken work with her. The researchers scan fragments with a 3D scanner. Then they have a computer judge whether the broken surfaces of two fragments interlock. Old fragments have worn edges. The calculation includes this wear as well[9].

One result of this method came from a flood story. It involved fragments of Tablet 3 of a flood epic called Atrahasis. One fragment was in London, and the other was in Geneva. The two places are about 1,000 kilometers apart. They could not be placed together in person, so the join had not been confirmed. When the 3D data were overlaid, the two fragments connected[9]. This method was published in a paper in 2014[10].

Research using 3D data continues today. A paper was released in March 2026. It deals with a neural network that classifies information about artifacts using 3D point cloud data from clay tablets. A point cloud is data that records the surface of an object as the coordinates of countless points. This kind of data is produced when a clay tablet is scanned with a 3D scanner. The curves and thickness of the surface remain as numbers. This study performed better than earlier transformer-based models[7].

For now, however, the main weapon actually used by the Fragmentarium is not 3D shape. It is photographs and transliterations, that is, letters and images[1][2]. Matching fracture surfaces takes a great deal of work because each fragment must be scanned one by one. The two paths do not replace each other. They move together.

The work of detecting signs in photographs has also advanced quickly. In June 2026, eBL researchers published a new paper. It presents a method that automatically finds the positions of cuneiform signs in photographs of clay tablets. The target set included more than 87,000 fragment photographs. The method detected 2.9 million signs. Its performance improved greatly over earlier methods[11].

References

[1] Cobanoglu, Y., Laasonen, J., Simonjetz, F., Khait, I., Cohen, S., Földi, Z., Hätinen, A., Heinrich, A., Mitto, T., Rozzi, G., Sáenz, L., & Jiménez, E. (2024). Transliterated cuneiform tablets of the electronic Babylonian Library platform. Journal of Open Humanities Data, 10, 19. https://doi.org/10.5334/johd.148

[2] Jiménez, E. (2023, 5월 25일). AI is helping us read ancient Mesopotamian literature. The Conversation. https://theconversation.com/ai-is-helping-us-read-ancient-mesopotamian-literature-205091 (accessed September 3, 2026)

[3] Gutherz, G., Gordin, S., Sáenz, L., Levy, O., & Berant, J. (2023). Translating Akkadian to English with neural machine translation. PNAS Nexus, 2(5), pgad096. https://doi.org/10.1093/pnasnexus/pgad096

[4] Wang, Z. (2026). TabletCraft: Bridging a 4,000-year cultural gap with bidirectional Akkadian NMT and cuneiform rendering. arXiv:2608.02609. https://arxiv.org/abs/2608.02609 (accessed September 3, 2026)

[5] Julius-Maximilians-Universität Würzburg. (2026, 5월 18일). AI reads cuneiform: A milestone for Ancient Near Eastern Studies. EurekAlert!. https://www.eurekalert.org/news-releases/1128630 (accessed September 3, 2026)

[6] Electronic Babylonian Literature. (n.d.). ngram-matcher [소프트웨어]. GitHub. https://github.com/ElectronicBabylonianLiterature/ngram-matcher (accessed September 3, 2026)

[7] Hagelskjær, F. (2026). A novel network for classification of cuneiform tablet metadata. arXiv:2603.03892. https://arxiv.org/abs/2603.03892 (accessed September 3, 2026)

[8] electronic Babylonian Library. (n.d.). Fragmentarium. ebl.lmu.de. https://www.ebl.lmu.de/fragmentarium (accessed September 3, 2026)

[9] International Association for Assyriology. (2022, 10월 22일). In the spotlight: The "Virtual Cuneiform Tablet Reconstruction Project". iaassyriology.com. https://iaassyriology.com/in-the-spotlight-the-vctr-project/ (accessed September 3, 2026)

[10] Collins, T., Woolley, S. I., Hernandez Munoz, L., Lewis, A., Ch'ng, E., & Gehlken, E. (2014). Computer-assisted reconstruction of virtual fragmented cuneiform tablets. In Proceedings of the 2014 International Conference on Virtual Systems and Multimedia (pp. 70-77). IEEE. https://doi.org/10.1109/VSMM.2014.7136691

[11] Che, W., Garcés Arias, E., Niaz, A., Bender, A., & Jiménez, E. (2026). Automated sign detection across the Electronic Babylonian Library: A large-scale dataset and end-to-end cuneiform OCR pipeline. arXiv:2606.22608. https://arxiv.org/abs/2606.22608 (accessed September 3, 2026)

5.3. Truths Revealed

Lines Returned to the Epic of Gilgamesh

The Fragmentarium newly identified 20 fragments of the Epic of Gilgamesh. Thanks to that, more than 100 lines of text returned to their proper places. Several lost scenes also came back to life. One is the scene where Enkidu tells Gilgamesh not to kill Humbaba. Another newly connected scene shows the two men going to Nippur after killing Humbaba. Nippur was the religious center that served the god Enlil[1].

A case in 2025 showed how fast this method can be. Ammar Fadhil of the University of Baghdad and Professor Jimenez revived a long poem praising Babylon. The poem had been forgotten for more than 1,000 years. Using the eBL platform, the researchers found 30 copies belonging to the same poem. Done by hand, this work would have taken decades. Once the scattered copies were gathered, most of the 250-line poem could be connected again[3].

This poem shows a day in the city of Babylon. It includes the river, the gates, and the work of the people. It also contains new records about duties performed by priestesses. The fact that 30 copies were found means that the poem was widely read. It was likely a text that students copied in school[3].

The time required for this work shows the power of the method. To find 30 copies by hand, one would have to search the catalogs of museums around the world one by one. On the eBL platform, the same work ends with a single search[3].

This method is not used only for the Epic of Gilgamesh. The same method is applied to Babylonian divination texts, medical texts, dictionaries, and prayers. The joins found by eBL span many kinds of writing[5].

The Cedar Forest First Solved by Human Hands

Machines did not solve everything. A widely known new discovery in the Epic of Gilgamesh came from human hands.

In 2011, the Sulaymaniyah Museum in Iraq purchased a clay tablet. Farouk Al-Rawi and Andrew George read this tablet. They published the results in 2014. The tablet contained the beginning and the end of Tablet 5 of the Standard Babylonian version. About 20 new lines of text came from it[4].

The new text portrays the Cedar Forest through sound. Monkeys cry noisily, and cicadas sing in chorus. All kinds of birds chirp. These sounds delighted Humbaba, the guardian of the forest, every day. The text also includes a scene in which the two men see the forest for the first time and stop in wonder. There is also a passage where Gilgamesh regrets cutting down trees after killing Humbaba[4]. There are no forests with monkeys on the plains of Mesopotamia. So scholars still ask where this description came from.

The route by which this clay tablet appeared was not clean. The museum bought it from an antiquities dealer. This was meant to prevent looted artifacts from being dispersed after the war in Iraq. There is no record of which site it came from[4]. Artifacts like this offer fewer clues for finding matching fragments.

This discovery came without artificial intelligence. But how often discoveries of this kind can be made in the future depends on tools. Before the Sulaymaniyah tablet became known to the world, chance and the human eye were needed. eBL aims to turn that chance into search[5].

The Story of the Round Ark

The flood story is similar. If Tablet 5 of the Epic of Gilgamesh contains the Cedar Forest, Tablet 11 contains the great flood[6]. The clay tablet discussed below is not a copy of the Epic of Gilgamesh. It is an older tablet that contains the same flood tradition[8].

Dr. Irving Finkel of the British Museum examined a clay tablet in 2009. A private collector had brought it to the museum. He judged that it was written around 1750 BCE. It is about 11.5 centimeters wide and 6 centimeters tall. It contains 60 lines[8]. The content is a story in which a god instructs a person to build a boat. It also says to take animals aboard two by two[7].

What stands out is the shape of the boat. This boat is not a long box. It is a round boat. It has the same form as a round basket boat used on rivers, called a coracle[7]. The clay tablet mentions rope and a wooden frame. It also gives the amount of bitumen to apply so the boat would not leak. Bitumen is a black, sticky natural oil. Dr. Finkel published this reading in a book in 2014[8].

This clay tablet was not read by artificial intelligence. A person read it line by line. The role of artificial intelligence lies elsewhere. It is to gather copies of flood stories scattered around the world into one place and find overlapping lines among them[5].

What Machines Do and What People Do

This field treats results in three layers.

The first layer is physical joining. This is when the broken surfaces of two fragments actually interlock. In that case, the two fragments are confirmed to have come from one clay tablet.

The second layer is textual joining. This is when two fragments do not touch directly, but are shown to be different parts of the same work. Most of what the Fragmentarium finds belongs here[5]. Thousands of fragments belonging to the same works have been identified in this way[2].

These two layers have different characters. A physical join becomes fact once it is confirmed by sight. A textual join depends on probability, so it can change when new fragments appear. eBL records the two layers separately[5].

The third layer is interpretation. This means asking what the restored text means. A machine does not know where the memory of monkey sounds in the Cedar Forest came from. Nor can a machine answer how a round boat became the square ark of the Bible. These questions belong to scholars of ancient texts and historians.

Candidates produced by machines become official records only after human review. A scholar opens a list of candidates and looks at photographs of two fragments side by side. After checking the shapes of the signs, the spacing between lines, and even the thickness of the tablet, the scholar confirms the join[5]. The machine's work goes only as far as selecting a few worth looking at from countless possible pairs[2].

A review article published in June 2026 draws the same line. Today's artificial intelligence cannot decipher cuneiform on its own. It is a tool that helps human research[9]. The Paleograficum at the University of Würzburg is the same. It only reduces a handwriting comparison that once took three days to five minutes. A person decides what to compare and judges the result[10].

What has changed is speed. Joins that numbered 5,000 over 150 years increased by 1,500 in five years[2]. The eternal life that Gilgamesh searched for was eventually gained by his story. Sentences that waited for 2,600 years, broken in the earth, are now returning one by one to their proper places.

Much work remains. Of the world's 500,000 clay tablets, only some have been read[11]. The eBL transliteration data include 25,000 tablets[5]. The 30,000 Hittite fragments have only just been organized as photographs[10]. The 3D work of matching broken surfaces is only beginning. The next generation of students will inherit this list. By then, the blanks in the Epic of Gilgamesh will be fewer than they are now. The fragments left by the fire of Nineveh have not yet all been read. The tools for reading them increase every year.

References

[1] Ofgang, E. (2024, 8월 12일). Piecing together an ancient epic was slow work. Until AI got involved. The New York Times. https://edition.pagesuite.com/tribune/article_popover.aspx?guid=235bb893-7f2d-41cd-9857-139c7d1e8355 (accessed September 3, 2026)

[2] Jiménez, E. (2023, 5월 25일). AI is helping us read ancient Mesopotamian literature. The Conversation. https://theconversation.com/ai-is-helping-us-read-ancient-mesopotamian-literature-205091 (accessed September 3, 2026)

[3] Fadhil, A. A., & Jiménez, E. (2025). Literary texts from the Sippar Library V: A hymn in praise of Babylon and the Babylonians. Iraq, 1-58. https://doi.org/10.1017/irq.2024.23

[4] Al-Rawi, F. N. H., & George, A. R. (2014). Back to the Cedar Forest: The beginning and end of Tablet V of the Standard Babylonian Epic of Gilgameš. Journal of Cuneiform Studies, 66, 69-90. https://doi.org/10.5615/jcunestud.66.2014.0069

[5] Cobanoglu, Y., Laasonen, J., Simonjetz, F., Khait, I., Cohen, S., Földi, Z., Hätinen, A., Heinrich, A., Mitto, T., Rozzi, G., Sáenz, L., & Jiménez, E. (2024). Transliterated cuneiform tablets of the electronic Babylonian Library platform. Journal of Open Humanities Data, 10, 19. https://doi.org/10.5334/johd.148

[6] George, A. R. (2003). The Babylonian Gilgamesh Epic: Introduction, Critical Edition and Cuneiform Texts. Oxford University Press.

[7] PBS NewsHour. (2014, 1월 24일). New evidence suggests Noah's ark may have been round. pbs.org. https://www.pbs.org/newshour/world/new-evidence-suggests-noahs-ark-may-have-been-round (accessed September 3, 2026)

[8] Finkel, I. (2014). The Ark Before Noah: Decoding the Story of the Flood. Hodder & Stoughton.

[9] Centre for the Study of Manuscript Cultures. (2026, 6월 14일). Can artificial intelligence decipher cuneiform tablets? Universität Hamburg. https://www.csmc.uni-hamburg.de/publications/mesopotamia/2026-06-14.html (accessed September 3, 2026)

[10] Julius-Maximilians-Universität Würzburg. (2026, 5월 18일). AI reads cuneiform: A milestone for Ancient Near Eastern Studies. EurekAlert!. https://www.eurekalert.org/news-releases/1128630 (accessed September 3, 2026)

[11] Wang, Z. (2026). TabletCraft: Bridging a 4,000-year cultural gap with bidirectional Akkadian NMT and cuneiform rendering. arXiv:2608.02609. https://arxiv.org/abs/2608.02609 (accessed September 3, 2026)

Chapter 6. Ancient Mesopotamian Receipts and Merchants' Whispers: AI Translation

6.1. Case File

Receipts Carved in Clay

Ancient Mesopotamians wrote on wet clay tablets. They cut the end of a reed at an angle to make a stylus. A stylus is a writing tool that leaves marks without ink. They pressed this stylus into clay and lifted it right away[1]. The pressed marks became triangular shapes like wedges. This writing is called wedge-shaped writing. Scholars also use the name cuneiform. Cuneiform means wedge-shaped.

Clay hardens like stone when it dries. Unlike paper or leather, it does not rot over time. Even in a fire, it does not turn to ash. Instead, heat bakes the clay hard like brick. When cities burned in war, clay tablets became stronger[3]. Thanks to that, account books from 4,000 years ago have survived to this day.

Many clay tablets are about the size of a palm. They are small enough to carry in a pocket. A scribe kneaded wet clay and shaped it into a tablet. A scribe was a person who wrote and read for others. After writing, the tablet was dried in the sun. It was rarely fired in a kiln on purpose. When museums fire tablets today, they do it for preservation[2]. If something needed correction, the surface was rubbed smooth and erased.

About 500,000 clay tablets are estimated to have been excavated so far[4]. More tablets still lie buried underground. The writing on them is not only the boasting of kings. There are receipts saying how many sacks of grain were received. There are contracts lending silver and setting interest. There are also letters sent by wives to husbands who had gone far away[6].

One receipt is short. The item name and quantity come first. Then come the names of the giver and the receiver. At the end, the date and witness names are written. The order is similar to a receipt from a shop today. One such document has little content by itself. But when tens of thousands gather, the story changes. They show in which year grain prices rose. They also show what each city bought and sold.

Most Remain Untranslated

The problem is that these texts are barely being read. Only some clay tablets have been translated into modern languages[5]. The rest are sleeping in museum boxes.

Not just anyone could write. To become a scribe, a person studied for a long time at school. Students wrote signs again and again on clay tablets. A tablet with mistakes was soaked, kneaded, and shaped again. Surviving practice tablets are also excavated as artifacts. They show classrooms from that time.

Few people knew how to write. Most people had scribes write for them. When they received a letter, a scribe read it aloud. This is why letters had a polite fixed form. The person speaking and the person writing were different.

Even reading one clay tablet requires several steps. First, each sign is identified. Next, the sound of the sign is decided. Then the words are separated and their meanings are found. Finally, the sentence is assembled and translated. If the format is familiar, the work ends quickly. If the document is unfamiliar, it takes a long time.

The reason is that there are too few people who can read them. The study of cuneiform is called Assyriology. It takes decades to train one scholar. They must learn ancient languages and ancient history together. So there are very few people in the world who can read these signs[5].

The number of newly found clay tablets exceeds the speed at which scholars can read them. Since the number of people does not grow, the boxes only pile up. In June 2026, the University of York in the United Kingdom began a new research project. The project is called "Towards Universal Access to Cuneiform Culture." It aims to translate cuneiform materials into Arabic as well so that anyone can read them. The project will run for three years, until 2029[7].

Work is also underway to gather materials in one place. The Cuneiform Digital Library Initiative is one such project. It collects information on clay tablets scattered around the world and releases it online. It places photographs, transliterations, and translations side by side. A transliteration writes the sounds of cuneiform signs in the alphabet. So far, more than 400,000 artifacts have been listed[4]. Only when data is gathered can machines have something to learn from.

Clay tablets are now scattered across many countries. Many are in museums in Iraq and Turkey. Museums in London, Paris, and Berlin also hold large shares. They are divided among universities in the United States and Japan as well. If the front and back of one document are in different countries, it is hard to match them. Only when photographs are posted together online can they be compared.

Account Books Pushed Aside by Royal Records

There is an old reason why everyday documents were pushed back. Excavation teams in the 19th and early 20th centuries dug palaces and temples first. Large stone monuments and gold objects drew people's eyes. Scholars also read royal inscriptions and law codes first. They spent their time first on translating myths and epics.

Meanwhile, merchants' receipts and letters were pushed down the order. Grain ration lists were treated as records of minor errands. Seen one by one, they seemed to contain little. Scholars put off the work of copying these dull documents by hand. In this way, the voices of ordinary people stayed buried for a long time.

Something was lost because of this bias. It became harder to know how prices rose and fell. What women did became hidden. The names of children and servants remained only in lists. Royal inscriptions record boasts. Receipts do not boast. They record only the fact of a transaction. For that reason, receipts can become more honest sources.

One Wedge Has Many Meanings

Even for machines, cuneiform is not easy to read. One sign can play several roles at once[8]. In one place, it becomes a syllable sign that represents a sound. In another place, it represents the meaning of a whole word. In another place, it makes no sound. It is used only as a marker that tells whether the following word is wood or a god. This kind of marker is called a determinative.

Cuneiform was not used to write only one language. Sumerian and Akkadian shared the same signs. Later, Hittite and Elamite also borrowed these signs. The same sign can have different sounds in each language. The reader must first decide the language of the document.

The Sumerian sign gish, meaning wood, is a good example. When it is placed before the name of a piece of furniture, it marks the material as wood. When it stands alone, it is the word wood itself. In another place, it represents only the sound gish. The reader must look at the surrounding sentences and choose its role.

The way numbers were written was also different from today. Mesopotamians grouped numbers by 60. The 60 minutes and 60 seconds of the clock come from this. Quantities on receipts were also written in this system. Grain and silver used different units. Even the same sign has a different value when it refers to grain and when it refers to silver. Machines must learn these unit differences as well.

Writing words without spaces is another obstacle. Ancient scribes did not leave spaces between words. Computers have trouble finding where one word ends and the next begins[8]. Broken and worn clay tablets are another problem. When wedge marks become shallow, even human eyes struggle to tell them apart. The frequent changes in Akkadian word forms also make the work hard. Many endings attach to one root. The final sound changes depending on whether the word is a subject or an object. It changes again by tense and person. The base form to look up in a dictionary must be found first.

Clay tablets break for many reasons. They are pressed underground and soaked by water. Corners can break off during excavation. Tablets broken through the old antiquities trade have been scattered across several countries. After they are taken from the ground, salt can seep out and cover the signs[2]. Scholars call the empty places where signs have disappeared lacunae.

Because of these many layers of walls, hundreds of thousands of account books have remained silent.

References

[1] Truman State University. (2026). How it works. Truman State University Cuneiforms. https://exhibits.truman.edu/cuneiform/how-it-works/ (accessed September 3, 2026)

[2] Guinan, A., Oller, G., & Ormsby, D. (1976). Nippur rebaked: The conservation of cuneiform tablets. Expedition, 18(3). https://www.penn.museum/sites/expedition/nippur-rebaked/ (accessed September 3, 2026)

[3] Bouscaren, D. (2023). Assyrian women of letters. Archaeology, 76(6). https://archaeology.org/issues/november-december-2023/features/assyrian-women-of-letters/ (accessed September 3, 2026)

[4] Cuneiform Digital Library Initiative. (2026). About the CDLI. cdli.earth. https://cdli.earth/about (accessed September 3, 2026)

[5] Gutherz, G., Gordin, S., Sáenz, L., Levy, O., & Berant, J. (2023). Translating Akkadian to English with neural machine translation. PNAS Nexus, 2(5), pgad096. https://doi.org/10.1093/pnasnexus/pgad096

[6] Michel, C. (2020). Women of Assur and Kanesh: Texts from the Archives of Assyrian Merchants (Writings from the Ancient World 42). SBL Press.

[7] Cuneiform Digital Library Initiative. (2026). Research Programme "Towards Universal Access to Cuneiform Heritage" 2026-2029. cdli.earth. https://cdli.earth/postings/239 (accessed September 3, 2026)

[8] Gordin, S., Gutherz, G., Elazary, A., Romach, A., Jiménez, E., Berant, J., & Cohen, Y. (2020). Reading Akkadian cuneiform using natural language processing. PLoS ONE, 15(10), e0240511. https://doi.org/10.1371/journal.pone.0240511

6.2. The Forensic Trail: Akkadian AI Translation

Opening Two Translation Paths

In 2023, researchers at Tel Aviv University opened the way first. Gai Gutherz and his colleagues built a translator that turns Akkadian into English[1]. Akkadian was used for a long time in Mesopotamia. It was written in cuneiform, and no one speaks it today.

The researchers opened two translation paths. The first path is called C2E. It takes cuneiform signs converted into computer characters and returns English. It does not read clay tablet photographs directly. Since it does not pass through a scholar's transliteration, the process is shorter. The second path is called T2E. It takes text in which a scholar has written the sounds of cuneiform signs in the alphabet. This alphabetic conversion is called transliteration. At the transliteration stage, the meaning of each sign is partly organized. So T2E translation is more accurate[1].

Making a transliteration is not easy either. A scholar looks at a sign and decides its sound. If there are several sounds, the scholar chooses by context. This choosing is judgment. Once that judgment is made, the text becomes easier for a machine to handle. This is why T2E scores higher.

The inside of the translator is divided in two. The front part turns the original text into clusters of numbers. The back part looks at those numbers and chooses English words one by one. The front part is called the encoder, and the back part is called the decoder. The encoder stores the meanings and order of words together. The decoder looks at the words already chosen and decides the next word.

Research that reads signs directly from photographs continues separately. In June 2026, researchers at the Electronic Babylonian Library released their results. A machine scanned photographs of more than 87,000 broken clay tablets. It found nearly 2.9 million sign marks. Detection performance rose by up to 37 percent compared with earlier methods. It is still weak with badly broken tablets and disordered line layouts[6].

A Machine Trained on 8,000 Documents

A machine learns translation by looking at example sentences. The researchers gathered documents from an open dataset called ORACC. ORACC is an online collection that organizes and releases cuneiform texts. The number of documents gathered from it was 8,056. In sentence count, there were more than 50,000[1].

The collected documents were of many kinds. About 3,000 were inscriptions recording the achievements of kings. About 2,000 were administrative letters. Court records and omen reports were included as well[1]. Each document had an English translation made by a scholar beside it. The machine looked at these pairs tens of thousands of times and found rules on its own.

The way a machine learns is similar to studying for an exam. First, it looks at the question and the answer together. Then it hides the answer and tries to solve the question by itself. If it is wrong, it traces where it went wrong. It repeats this tracing millions of times. What differs from humans is speed and persistence.

A machine does not handle whole words as they are. It breaks words into smaller pieces and handles those pieces. These pieces are called tokens. Even rare words can be handled if they are broken into pieces. This method fits Akkadian, which has many endings attached to words.

Translation quality is measured with a score called BLEU. It counts how much a machine translation overlaps with a human translation. It scores the ratio of overlapping word groups. The score is shown from 0 to 100. A score above 30 is often seen as meaning the sense comes through. But it is not an absolute standard. Interpretation changes depending on document type and sentence length.

The researchers chose a structure called a convolutional neural network. It learns by grouping nearby information together. They chose it because training is fast and stable when data is limited. In validation, C2E received 36.52 points. T2E was higher, at 37.47 points. Older automatic search methods reached only 27.09 and 23.51 points. The system performed well on short documents with fixed formats[1].

The way the data is divided also matters. The portion for learning and the portion for testing are set apart. If the documents for testing are shown in advance, the score is meaningless. It is like knowing the exam questions in advance. The researchers kept this division and calculated the scores.

The same original text can be translated differently by different people. So quality is not judged by score alone. Human evaluation is used together with the score. Scholars are shown the translations and asked to grade them. They are asked separately whether the meaning comes through. They are also asked separately whether there are errors. In T2E, about 44 percent of translations were judged usable by scholars. In C2E, the figure was about 38 percent[1]. Sometimes scores and human evaluations diverge. In those cases, human judgment comes first.

Akkadian translation is a low-resource task. Low-resource means there are few example sentences to learn from. English-French translators learn from tens of millions of sentences. Akkadian has less than one-thousandth that number of paired sentences. When data is scarce, a large model can get more confused.

In 2025, a shared task was held in which several research teams competed. The task was to find the base forms of words and restore missing signs. The data came from the Electronic Babylonian Library and Archibab. Measuring performance on the same data showed right away which method was better[2].

Finding base forms is groundwork for translation. In scholarship, this task is called lemmatization. It means returning words expanded by endings to their dictionary forms. Only after this groundwork is done can word-by-word search and statistics become possible.

Catching Plausible Lies

Machine translation has one dangerous habit. It invents plausible sentences even when there is no basis for them. This habit is called hallucination. The researchers examined hallucination in two types. One type misreads the meaning even after seeing the source text. The other ignores the source text and creates a different story[1].

Hallucination often occurs with broken clay tablets. The machine fills the missing spaces with imagination. So a scholar must always review the translation. The sentence produced by the machine is a first draft, not a final judgment.

There are several ways to reduce hallucination. Broken parts are clearly marked as blanks. The machine is allowed to answer that it does not know that passage. The same sentence may be translated several times. If the answer changes each time, it is not trusted. Scholars use these marks to decide where to look again.

Turning wedge-shaped signs into alphabetic form and separating words comes before translation. The same line of research also tried giving this task to machines. A neural network called bidirectional LSTM performed better. Its accuracy reached as high as 97 percent[3].

A Door Open in Both Directions

Until now, translators moved in only one direction. They only turned ancient texts into modern English. In May 2026, Wang Zhaohui of the University of Southern California opened the opposite direction. The tool he made is called TabletCraft[4].

TabletCraft also translates English into Akkadian. It draws the result in cuneiform signs and even creates an image of a clay tablet. For training, it used 116,000 pairs of bilingual sentences. Its sign mapping table contained more than 14,000 entries. In testing, it scored 49.1 from Akkadian to English. From English to Akkadian, it scored 48.5[4].

This tool was used in museum experience devices and university classes. Its creator also stated a clear limit. Machine-made Akkadian must not be cited as if it were a real ancient record. The texts the tool handles well are royal inscriptions from the Neo-Assyrian period[5].

Two-way translation has real use. Students can turn their own sentences into cuneiform signs. When they draw the signs themselves, they can see how the writing is built. In museums, visitors can carve their own names. Ancient writing changes from something to look at into a tool to handle.

The studies discussed above released their data and programs together[1][5]. When they are made public, other teams can repeat the same tests. If something is wrong, it is found quickly. Few people work on ancient languages. Openness strongly shapes the speed of progress.

Methods for using translators are also taking shape. A scholar displays the machine's first draft on the screen. The source text and translation are placed side by side and compared. Doubtful words are marked. As the marks accumulate, they become material for the next round of training.

References

[1] Gutherz, G., Gordin, S., Sáenz, L., Levy, O., & Berant, J. (2023). Translating Akkadian to English with neural machine translation. PNAS Nexus, 2(5), pgad096. https://doi.org/10.1093/pnasnexus/pgad096

[2] Gordin, S., Sahala, A., Spencer, S., & Klein, S. (2025). EvaCun 2025 shared task: Lemmatization and token prediction in Akkadian and Sumerian using LLMs. In Proceedings of the Second Workshop on Ancient Language Processing (pp. 242-250). Association for Computational Linguistics. https://doi.org/10.18653/v1/2025.alp-1.33

[3] Gordin, S., Gutherz, G., Elazary, A., Romach, A., Jiménez, E., Berant, J., & Cohen, Y. (2020). Reading Akkadian cuneiform using natural language processing. PLoS ONE, 15(10), e0240511. https://doi.org/10.1371/journal.pone.0240511

[4] Wang, Z. (2026, 5월 14일). TabletCraft: Bridging a 4,000-year cultural gap with bidirectional Akkadian NMT and cuneiform rendering. arXiv:2608.02609 (2026년 5월 14일 제출, 2026년 8월 공개). https://arxiv.org/abs/2608.02609 (accessed September 3, 2026)

[5] Buyukyildirim, K. (2026, 8월 11일). TabletCraft opens a two-way AI translation channel for archaeologists. Arkeonews. https://arkeonews.net/tabletcraft-opens-a-two-way-ai-translation-channel-for-archaeologists/ (accessed September 3, 2026)

[6] Che, W., Garces Arias, E., Niaz, A., Bender, A., & Jimenez, E. (2026, 6월 21일). Automated sign detection across the Electronic Babylonian Library: A large-scale dataset and end-to-end cuneiform OCR pipeline. arXiv:2606.22608. https://arxiv.org/abs/2606.22608 (accessed September 3, 2026)

6.3. Truths Revealed

Letters from the Merchants of Kanesh

Near Kayseri in central Turkey, there is a mound called Kultepe. Four thousand years ago, this place was called Kanesh. Assyrian merchants kept a trading base there. The merchants' base community was called a karum. Karum originally meant a quay by a river. A place where goods were unloaded became a merchants' district.

The karum of Kanesh was a self-governing organization of Assyrian merchants. They made rules among themselves and settled disputes. They paid taxes to the local king and received protection. Goods were carried on the backs of donkeys. The route from Ashur to Kanesh was a long journey that took several weeks.

Excavations continuing since 1948 have yielded a flood of clay tablets. About 23,500 documents have been found so far. This is the number reported by Fikri Kulakoglu, who leads the excavation[1]. Their contents are not royal boasting. There are loan documents and cargo invoices. A cargo invoice is a list of goods that were shipped. There are court records and family letters.

The letters preserve demands and complaints just as they were. They ask when money will be sent. They scold someone because goods are late. They add greetings telling the recipient to take care. Some tell a son to study. The worries of people from 4,000 years ago remain in the sentences.

These documents revealed the structure of an ancient economy. Merchants brought textiles and tin from Ashur. In Anatolia, they exchanged them for silver and gold. The documents also recorded how taxes and tolls were paid. They even preserved how disputes were judged.

Tin is a metal needed to make bronze. Bronze is made by mixing copper and tin. Tin was hard to obtain in Anatolia. It had to be brought from far to the east[8]. This shortage was the force that created long trade routes.

The role of women who remained in Ashur was also revealed. The wives and daughters of merchants wove textiles at home and sent them out[2][8]. They controlled output and bargained over prices. In letters to their husbands, they also discussed market prices. Cecile Michel collected and translated into English about 300 such documents. Michel works at the French National Centre for Scientific Research and the University of Hamburg[2].

Reading an Envelope Without Opening It

Ancient merchants placed important clay tablets inside clay envelopes and sealed them. To read the inside, the envelope had to be broken. But the envelope also bore seal impressions. If the envelope was broken, that information disappeared.

The envelope was a device to prevent forgery. The contract terms were written on the inner clay tablet. Clay was then wrapped over it to make an envelope. The same contents were written again on the outside of the envelope. Witnesses rolled their seals over it. If the outside and inside differed, the alteration would be exposed.

A way was needed to read the inside while preserving the envelope. In June 2026, the Kayseri Archaeology Museum showed an answer. It read one clay tablet inside an envelope without opening it. The work was carried out with the French National Centre for Scientific Research and the University of Hamburg. They used a method like CT scanning, which looks inside an object without cutting it. The contents inside were a promise involving a wheat and barley transaction. Kulakoglu said the scan showed even a single barley grain[1]. The images of letters obtained this way become input data for translators. Machines gain something to learn only when there are more documents to read.

A Translation Contest with $50,000 at Stake

A translation contest was also held. An organization called the Deep Past Initiative prepared the contest[3]. It was held on Kaggle, an artificial intelligence competition site[4]. The task was to translate commercial documents from Kanesh into English[3]. The contest began on December 16, 2025. The entry deadline was March 23, 2026. The prize was $50,000. A financial company paid the operating costs[7].

The contest data was released so anyone could download it[4]. Public data remains even after the contest ends. The next researcher can start again from that data. One contest leaves behind one dataset.

Among the Kultepe documents, many still have not been translated[3]. Excavation continues every year even now[1]. When new clay tablets are found, they are added to the list. If the number of readers does not grow, only the list grows longer. This is why help from machines is needed.

The same situation does not end at one site. Similar documents are piled up in museums around the world. There are boxes there that still have not been read. Documents with fixed forms are common everywhere. Grain ledgers and tax registers are examples. The method that worked at Kanesh can also be used in other archives.

The contest announcement put forward real people who appear in the documents. They were the names of merchants such as Ashur-nada and Ashur-idi. The names Ishtar-lamassi and Ennam-Ashur were listed as well[3]. The names of one merchant family from 4,000 years ago became the subject of a contest.

The contest method has strengths. Developers around the world, not just a few laboratories, tackle the same problem. Because performance is measured with the same yardstick, it becomes clear which method is better. Prize money has the power to bring people together.

This group of documents suits machine translation well. The forms are fixed, and the sentences are short. The same expressions repeat across many documents. Machines are good at finding repeated rules. There are also many documents that scholars have already translated. Those translations become teaching material for the machine.

What Machines Cannot Do

Achievements have accumulated, but the remaining barriers are also clear. The first barrier is a shortage of data. Translators for modern languages grow by consuming hundreds of millions of sentences. Akkadian has only tens of thousands of paired sentences[5]. When data is scarce, rare place names and personal names are easy to get wrong.

The second barrier is the many meanings of signs. The same sign can be hard to classify as a sound, a word, or a marker[6]. When the surrounding context is short, the machine also loses its way.

The third barrier is how machines judge. A machine calculates which word is likely to come next. It cannot judge whether the sentence is historically correct. That is why scholarly verification is needed as the final step[5]. Translation becomes trustworthy when machine and human work as a pair.

A division of work is taking root. The machine scans tens of thousands of tablets and produces a rough meaning. The scholar selects important documents from among them. Only the selected documents are then reread carefully by hand. This greatly increases the number of documents that can be read.

Translations need tags. They should record which passages were produced by the machine. Places corrected by people should also be marked. Only then can later readers retrace the evidence. Historical research begins with retracing evidence.

Names are a difficult part of translation. Personal names and place names cannot be translated by meaning. Their sounds must be transcribed as they are. But the same name may be written differently from one document to another. A machine can easily turn an unfamiliar name into the wrong word. So scholars make separate name lists and provide them to the machine.

Even so, the direction is clear. In the past, one scholar read hundreds of tablets over a lifetime. Now the machine first creates a draft translation. The scholar spends time correcting that draft. Receipts from 4,000 years ago that slept in boxes are waiting their turn. The voices of ordinary people, not kings, are beginning to be heard.

The merchants of Kanesh did not think their letters would survive for so long. They merely wrote down the accounts of that day. Those accounts now show the economy of an age. What remains is not boasting, but records.

References

[1] Türkiye Today. (2026, 6월 13일). 4,000-year-old clay tablet sealed in mud envelope goes on display in central Türkiye. Türkiye Today. https://www.turkiyetoday.com/culture/4000-year-old-clay-tablet-sealed-in-mud-envelope-goes-on-display-in-central-turkiye-3221841 (accessed September 3, 2026)

[2] Michel, C. (2020). Women of Assur and Kanesh: Texts from the Archives of Assyrian Merchants (Writings from the Ancient World 42). SBL Press.

[3] Deep Past Initiative. (2026). Deep Past Challenge: Introduction. deeppast.org. https://www.deeppast.org/challenge/intro/ (accessed September 3, 2026)

[4] Kaggle. (2026). Deep Past Challenge: Translate Akkadian to English. kaggle.com. https://www.kaggle.com/competitions/deep-past-initiative-machine-translation (accessed September 3, 2026)

[5] Gutherz, G., Gordin, S., Sáenz, L., Levy, O., & Berant, J. (2023). Translating Akkadian to English with neural machine translation. PNAS Nexus, 2(5), pgad096. https://doi.org/10.1093/pnasnexus/pgad096

[6] Gordin, S., Gutherz, G., Elazary, A., Romach, A., Jiménez, E., Berant, J., & Cohen, Y. (2020). Reading Akkadian cuneiform using natural language processing. PLoS ONE, 15(10), e0240511. https://doi.org/10.1371/journal.pone.0240511

[7] Deep Past Initiative. (2025, 12월 16일). Deep Past Initiative partners with Kaggle on prize challenge to translate 4,000-year-old tablets. EIN Presswire. https://www.einpresswire.com/article/875885564/deep-past-initiative-partners-with-kaggle-on-prize-challenge-to-translate-4-000-year-old-tablets (accessed September 3, 2026)

[8] Bouscaren, D. (2023). Assyrian women of letters. Archaeology, 76(6). https://archaeology.org/issues/november-december-2023/features/assyrian-women-of-letters/ (accessed September 3, 2026)

Chapter 7. An Empire's Future Asleep in Turtle Plastrons: Oracle Bone Script and 3D Virtual Rejoining

7.1. Case File

Characters from 3,000 Years Ago Found in a Pharmacy

In 1899, Wang Yirong, a scholar of the Qing dynasty, recognized an ancient character. A long-told story is attached to this event. It says that Wang Yirong fell ill and received a medicinal material called dragon bones. Dragon bones were powdered medicine made by grinding old animal bones. People in the past believed these bones were dragon bones. It is said that the bones he received had not yet been ground. On their surface, he saw unfamiliar markings[1].

This story has not been established as fact. Chinese scholars believe there is no record to support it. The palaeographer Li Xueqin rejected the story in a 2007 article. He said the dragon bones sold as medicine had already been broken up, so the characters could not be seen. In reality, an antique dealer bought bones from Anyang and moved them elsewhere. Those bones are thought to have reached Wang Yirong, a scholar of bronze and stone inscriptions, for appraisal[2].

Wang Yirong died the following year and could not trace the source of the bones. His friend Liu E published the first book collecting these characters in 1903. In 1908, the scholar Luo Zhenyu identified where the bones had come from[1]. They were being dug out of fields in Anyang, Henan Province.

That land in Anyang is Yinxu. Yinxu was the site of the last capital of the Shang dynasty. The Shang ruled this place from the 13th century BCE to the 11th century BCE[3]. Formal excavation began in 1928. Before then, bones dug up by farmers were sold as medicine. Some sellers even scraped off the inscribed surface before selling them[1]. Records that had survived 3,000 years nearly disappeared as medicine.

A Kingdom Where the King Asked the Gods

The characters carved on the bones from Yinxu are called oracle bone script. No earlier written records are known in East Asia[3].

When the people of Shang faced important matters, they asked the gods. They asked whether it would rain. They asked whether the harvest would be good. They asked whether they could go to war. They even asked about the day when a queen would give birth. This ritual of asking the will of the gods is called divination. Divination was state business managed directly by the royal court[3].

Fu Hao, a queen of the Shang dynasty, was also a commander who led armies. Her name appears many times in oracle bone records. In 1976, Fu Hao's tomb was excavated intact at Yinxu[4]. The bronze vessels and jade objects inside the tomb showed her status. In this way, oracle bone script tells us about ancient lives together with tombs.

The questions written in oracle bone script were not limited to affairs of state. Some questions worried about the king's toothache. Others asked whether a hunt could proceed. So these bones became records that show a day in the lives of Shang people[3].

Bones Scorched by Fire, Characters Carved by Knife

The people of Shang used two materials for divination. One was the flat plastron of a turtle. The other was the broad shoulder blade of an ox. Together, these two materials are called oracle bones[3].

The preparation process was careful. Scribes smoothed the outer surface of the bone. Then they carved rows of small hollows on the back. They pressed a heated rod into the hollows. The heated bone cracked on the outer surface. The king and the diviner examined the direction and angle of the cracks. They read those cracks as the answer of the gods[3].

When the reading was finished, the scribes took up knives. They carved the process of divination onto the surface of the bone, stroke by stroke. These carved writings are oracle bone script.

The records followed a fixed order. First, they wrote the date of the divination and the name of the person in charge. Next, they wrote the question asked of the gods. Then they wrote the king's judgment of good or bad fortune. Last, they wrote what actually happened. The record filled four fields in sequence. Yet records with all four fields preserved are rare. A bone might break and lose one or two fields. A scribe might also leave a field blank from the beginning[3]. Even so, the order of writing itself was steady.

This regularity later becomes a great help. A fixed format narrows the blanks. If the first field has a date, the next field contains a question. Computers also learn this order and guess missing characters[3].

The way dates were written was fixed as well. The Shang counted days by pairing ten signs with twelve signs. These pairs completed one cycle every sixty days. Even if only two date signs remain on a bone, their order can be checked. This gives a clue for narrowing the date of a broken record[3].

Records Broken into 150,000 Pieces

After the Shang dynasty fell, Yinxu was buried in earth. Over those long years, the bones could not remain whole. The weight of the soil pressed down on the thin plastrons and split them. Water and microbes ate away at the grain of the bone. A large plastron that was once one piece broke into dozens of fragments.

About 150,000 oracle bone fragments are now known to the world[5]. These relics are not gathered in one place. After their discovery in the late 19th century, many flowed into the antique market[1]. As a result, fragments broken from a single bone came to sit in museums in different countries[5]. Some relics are in Anyang and Beijing in China. Others are held by institutions in Europe and Korea[6].

The work of restoring scattered fragments to their original single piece is called rejoining. For a long time, rejoining depended only on human eyes. Scholars examined tens of thousands of black-and-white rubbings under magnifying glasses. A rubbing is a copy made by placing paper over a bone and tapping ink onto it. It is like placing paper over a coin and rubbing it with a pencil to bring out the pattern. Scholars compared the grain of the bone and the curves of broken edges. They also checked one by one whether the ends of severed strokes connected[5].

This method was slow and difficult. Finding one matching pair could take years. Fragments whose partners were not found remained in boxes. Without a matching piece, the sentence also remained cut off. Records with only questions and no results piled up.

The scale of the work makes the difficulty clearer. If there are 150,000 fragments, the number of possible pairs exceeds 10 billion. Even if a person compared one hundred pairs a day, a lifetime would not be enough to see them all. The fragments are also scattered across several countries[5]. In many cases, there was not even a chance to place two fragments side by side.

3,000 Characters Still Have Not Spoken

Even if all the fragments are joined, a problem remains. The characters themselves cannot be read. About 4,500 distinct characters have been identified across oracle bone script. Yet scholars agree on the meanings of only about 1,500 of them. The meanings of the remaining roughly 3,000 characters are still unknown[7].

There are several reasons why reading is difficult. Knife marks have worn down, blurring the outlines. When bones broke, half of a character disappeared. Each scribe held the knife differently, so the same character was carved in different forms. As a result, several characters with the same meaning but different shapes appeared. These are called variant forms[3]. Modern Chinese character dictionaries cannot explain these shapes from 3,000 years ago.

This silence leaves holes in Shang history as well. If we cannot read the names of rituals, we cannot count the types of rituals. If we cannot read place names, we cannot draw the borders of the kingdom. If we cannot read personal names, we cannot know who was responsible for what. The silence of one character leads to the silence of one scene in history.

About 3,000 characters whose meanings are unknown remain, along with tens of thousands of fragments that have lost their partners. Artificial intelligence has entered Yinxu to break this silence.

References

[1] Mark, E. (2016, 2월 26일). Oracle bones. World History Encyclopedia. https://www.worldhistory.org/Oracle_Bones/ (accessed September 3, 2026)

[2] 姚小鸥. (2020, 7월 16일). 谁是甲骨文的最早发现者. 澎湃新闻. https://www.thepaper.cn/newsDetail_forward_8297008 (accessed September 3, 2026)

[3] Chen, Z., Hua, W., Li, J., Zhu, Y., Zhi, X., Liu, Z., Chen, T., Zhang, W., & Zhai, G. (2026). Oracle bone inscriptions information processing: a comprehensive survey. npj Heritage Science, 14, 220. https://doi.org/10.1038/s40494-026-02511-w (accessed September 3, 2026)

[4] Smarthistory. War and sacrifice: The tomb of Fu Hao. Smarthistory. https://smarthistory.org/tomb-of-fu-hao/ (accessed September 3, 2026)

[5] Zhang, C., Wang, B., Chen, K., Zong, R., Mo, B., Men, Y., Almpanidis, G., Chen, S., & Zhang, X. (2022). Data-driven oracle bone rejoining: A dataset and practical self-supervised learning scheme. In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (pp. 4482-4492). ACM. https://doi.org/10.1145/3534678.3539050

[6] 梁文艳. (2026, 5월 26일). 以科技赋能，让甲骨文"活"起来. 中国青年报. https://zqb.cyol.com/pc/content/202605/26/content_426168.html (accessed September 3, 2026)

[7] Liu, Y., Guan, H., Wang, P., Wang, X., Wan, J., Zhang, K., Zheng, H., Liu, X., Kuang, Z., Yang, H., Li, B., Liu, Y., Jin, L., & Bai, X. (2026). AlphaOracle: Oracle bone script decipherment via human-workflow-inspired deep learning. The Innovation, 7(11), 101462. https://doi.org/10.1016/j.xinn.2026.101462

7.2. The Scientific Investigation Team's Pursuit: Oracle Bone-Specific Multimodal Models and 3D Virtual Rejoining

Bringing Scattered Materials onto One Screen

An investigation begins with collecting materials. Oracle bone research was the same. In 2019, researchers at Anyang Normal University and the Chinese Academy of Social Sciences opened a data platform called Yin Qi Wen Yuan. Yin Qi Wen Yuan gathers oracle bone rubbings, character information, and papers in one place and makes them freely available[1].

This platform does more than upload photographs. It photographs oracle bone fragments in high resolution and turns them into 3D models. A 3D model is data that scans an object from several directions and recreates it as a solid form inside a computer. By rotating the screen, researchers can see the front, back, and broken side surfaces of a bone. The team makes these models with millimeter-level precision. The depth of shallow knife marks is also included in this data[1].

The effort to recover oracle bones that had gone overseas was also linked to this platform. The researchers chose to bring in 3D data instead of the physical objects. Working with collections in Germany, Korea, France, and other places, they gathered data for more than 900 items[1]. The actual objects stayed with their original collections, and only their forms returned to Anyang.

In October 2025, a tool called Yin Qi Xing Zhi was added to this platform. This tool recognizes characters, compares rubbings, and finds related literature[1]. A warehouse that had stored data became an assistant that answers questions.

The amount of accumulated data can also be confirmed in numbers. A comprehensive review paper published in 2026 introduces 36 public datasets related to oracle bones. Some are datasets for character recognition, and others are for rejoining research[2].

Using several kinds of data together is called multimodal. In oracle bone research, this means looking at three kinds of data together. One is a black-and-white rubbing made with ink. Another is a color photograph that preserves color and texture. The last is 3D scan data that contains depth information. A rubbing shows strokes clearly but loses depth. 3D data contains depth but can make strokes look faint. When the three are overlaid, they can fill one another's gaps[2].

Following the Order in Which Humans Read Characters

AlphaOracle, released in 2026, is an artificial intelligence system that works on this data. This model was created jointly by Huazhong University of Science and Technology, Anyang Normal University, and others. Its design principle is to follow the workflow of human scholars[3].

A human scholar goes through four stages. First, the scholar studies the rubbing and organizes the strokes. Next, the scholar examines the shape of the character. Then the scholar looks for other sentences where the same character appears. Last, the scholar checks whether the meaning fits with knowledge from ancient texts. AlphaOracle is divided into parts that handle each of these four stages[3].

There is a reason for this structure. Scholars find it hard to trust a machine that simply throws out an answer. When the process is divided into stages, they can see where the answer changed. A scholar can return to that stage and compare views with the machine[3].

The amount of material read by this model exceeds what a person could see in a lifetime. It handled 73,883 rubbings. It also handled 45,364 scan records and 23,755 research papers. The researchers asked 86 experts to try the tool. As a result, the time needed to analyze one character fell by 64 percent. Seventy-nine percent of participants said the tool was useful[3].

Viewing Characters by Breaking Them into Parts

Another research team chose a method that does not memorize characters as whole units. Oracle bone script characters are made of repeated pictorial parts[4]. Parts shaped like hands, feet, and vessels appear across many characters. If the meaning of a part is known, even an unfamiliar character can be guessed.

This idea led to the OB-Radix dataset, published in April 2026. OB-Radix contains 934 distinct characters. It includes 1,853 drawings of the parts that make up those characters. There are 478 types of parts. Each part has an explanation verified by experts[4].

When artificial intelligence meets an unfamiliar character, it first breaks it into parts. Then it examines how those parts are arranged. Whether two parts are placed vertically or side by side can change the meaning. The model studies these relationships as well and proposes a meaning[4].

In June of the same year, a model called OracleAnalyzer also appeared. This model retrained a small language model with about 3 billion parameters on oracle bone script data. The researchers created and added new oracle bone reasoning data and preference data. Preference data is a collection of answers that humans chose as better. As a result, the small model performed better than much larger models[5]. This means that well-chosen data can succeed without relying on size alone.

A Model That Redraws Faded Strokes

Some research first makes characters visible before reading them. OBSD, published in 2024, is such a tool. This tool uses a diffusion model that generates new images[6]. A diffusion model works by gradually making a clear image from a blurred stain.

OBSD takes worn and blurred oracle bone character images as input. It then proposes, as an image, which modern Chinese character that character may have led to. This study was selected as the best paper at an international conference in 2024[6]. Its idea of turning character reading into an image-generation problem was recognized.

Computers Match the Cross Sections of Broken Fragments

Alongside character reading, research on matching fragments also progressed. OB-Rejoin, released in 2022, is a dataset for rejoining research. It contains 998 actual oracle bone rubbings. The same research team also released a joining neural network called S3-Net[7].

This neural network calculates how closely the curves along fragment edges resemble one another. When two edges interlock, the score rises. The machine lists the candidates from the highest score down. The same research team also published an earlier study in 2020. At that time, the correct match appeared within the top ten candidates 98.39 percent of the time[11].

When people solve a puzzle, they look at the shapes of the pieces before they look at the picture. Machines use the same method. The difference is that a machine can compare hundreds of thousands of pieces in a day.

What the machine provides is not the answer. It is a list of candidates. The item at the top of the list can be wrong. So the researchers use whether the correct answer appears within the top ten as their performance measure[7]. There is a real difference between checking only ten items and checking one hundred thousand.

The Center for Oracle Bone Script Studies at Capital Normal University built a joining program with Henan University. Its name is Tiedada, and it has been in use since 2020. This program compares 180,000 rubbings at once to find matches. With this method, the research center organized records for about 1,500 oracle bone fragments. In the past, finding the matching piece for one fragment could take eight to ten years. Now candidates appear in a few days[8].

A Test That Moves to the Sentence Level

Reading one character and understanding a sentence are different tasks. S-OBI, released in late June 2026, is a test set that measures this difference. The researchers drew questions from 95 original rubbings verified by experts. There are 695 questions in all, and the answers are given at the sentence level. They then had several large AI models solve the same questions. The test results showed clear differences among models. They also showed that grasping the flow of an entire sentence remains difficult[9].

This test revealed one more fact. General AI does not solve oracle bone script problems well[9]. Even models trained on texts from across the world falter before sentences from 3,000 years ago. The characters and grammar are ones they have never learned. That is why research continues on retraining models for oracle bone script[5].

Research has also moved beyond broken edges to connect context. One research team created a task that judges whether two sentences originally appeared on the same bone. They released a dataset called OBID-ACR. They also proposed a method that uses both character-shape information and sentence information[10]. This opened a way to connect context even before the fragments are physically joined.

References

[1] 梁文艳. (2026, 5월 26일). 以科技赋能，让甲骨文"活"起来. 中国青年报. https://zqb.cyol.com/pc/content/202605/26/content_426168.html (accessed September 3, 2026)

[2] Chen, Z., Hua, W., Li, J., Zhu, Y., Zhi, X., Liu, Z., Chen, T., Zhang, W., & Zhai, G. (2026). Oracle bone inscriptions information processing: a comprehensive survey. npj Heritage Science, 14, 220. https://doi.org/10.1038/s40494-026-02511-w (accessed September 3, 2026)

[3] Liu, Y., Guan, H., Wang, P., Wang, X., Wan, J., Zhang, K., Zheng, H., Liu, X., Kuang, Z., Yang, H., Li, B., Liu, Y., Jin, L., & Bai, X. (2026). AlphaOracle: Oracle bone script decipherment via human-workflow-inspired deep learning. The Innovation, 7(11), 101462. https://doi.org/10.1016/j.xinn.2026.101462

[4] Zhang, J., Li, R., Pang, H., Xia, D., Zhu, Z., Zhang, Q., Li, C., & Yang, X. (2026, 4월 8일). Specializing large models for oracle bone script interpretation via component-grounded multimodal knowledge augmentation. arXiv:2604.06711. https://arxiv.org/abs/2604.06711 (accessed September 3, 2026)

[5] Song, Z., Wang, Y., Ma, Z., Yu, Z., Wang, T., Zhang, J., Wang, T., & Yu, K. (2026, 6월 24일). OracleAnalyser: Analysing implicit semantics of oracle bone scripts through MLLMs with post-training. arXiv:2606.25906. https://arxiv.org/abs/2606.25906 (accessed September 3, 2026)

[6] Guan, H., Yang, H., Wang, X., Han, S., Liu, Y., Jin, L., Bai, X., & Liu, Y. (2024). Deciphering oracle bone language with diffusion models. In Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). Association for Computational Linguistics. https://aclanthology.org/2024.acl-long.831/ (accessed September 3, 2026)

[7] Zhang, C., Wang, B., Chen, K., Zong, R., Mo, B., Men, Y., Almpanidis, G., Chen, S., & Zhang, X. (2022). Data-driven oracle bone rejoining: A dataset and practical self-supervised learning scheme. In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (pp. 4482-4492). ACM. https://doi.org/10.1145/3534678.3539050

[8] 北京日报. (2026, 8월 28일). 高校团队"唤醒"三千年前文明碎片. 新浪教育. https://edu.sina.com.cn/l/2026-08-28/doc-inipwhrw0020843.shtml (accessed September 3, 2026)

[9] Li, Z., Chen, Z., Chen, T., & Zhai, G. (2026, 6월 30일). Beyond single character: Evaluating MLLMs for sentence-level oracle bone inscription understanding. arXiv:2606.31169. https://arxiv.org/abs/2606.31169 (accessed September 3, 2026)

[10] Zhang, H., Wang, T., Zhang, Z., Li, B., Sun, H., Hou, C., Wang, N., Yu, Y., Jiao, Q., Xiong, J., & Liu, Y. (2026). A multi-modal dataset and method for bone-level association prediction in oracle bone inscriptions. npj Heritage Science, 14, 19. https://doi.org/10.1038/s40494-025-02282-w (accessed September 3, 2026)

[11] Zhang, C., Zong, R., Cao, S., Men, Y., & Mo, B. (2020). AI-powered oracle bone inscriptions recognition and fragments rejoining. In Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI-20), 5309-5311. https://doi.org/10.24963/ijcai.2020/779

7.3. The Truth Revealed

When One Character Changed, the Map Changed

AlphaOracle's first result was a single character. The researchers put into the system one character that had long been debated. It is the character that later developed into the Chinese character lao (劳)[1]. The system organized the shapes of the strokes and the composition of its components. It then found and listed other sentences in which the character appeared. Next, it compared them with knowledge from ancient texts and narrowed the meaning.

The answer proposed by the system was that this character was either a place name or a clan name. A clan means a family group descended from a common ancestor. This result differed from interpretations that had read the character as a modifier in the sentence. The researchers stated that this interpretation had been reviewed by scholars of ancient writing[1].

When the meaning of one character changes, every sentence containing that character changes. If it is a person's name, the sentence tells the story of an individual. If it is a place name, the same sentence becomes a record of local administration. How the Shang dynasty governed that land must also be redrawn. The researchers reported that this result changes how we view Shang administration and society[1].

What deserves attention in this study is not the answer itself. It is that the system presented the evidence with it. AlphaOracle shows which sentences and which components led it to choose that answer[1]. Scholars can follow that evidence and decide whether to agree. The core of this study was to produce answers in a form that humans can review.

Matching That Once Took Years Now Takes Days

There were results in joining work as well. They came from the Tiedada program discussed earlier. This program found more than 300 matches that scholars had long missed[2]. It added that number on top of matches people had searched for by hand for more than a century.

The human role remains. Scholars check the list given by the machine one item at a time. The check ends with comparison against the physical objects.

The Anyang Normal University laboratory also reported its own results. By May 2026, it had newly joined 155 oracle bones. The same laboratory identified nearly 7,000 cases of duplicate records[3]. Sometimes one bone had been published separately in several catalogues. In that case, one bone was counted as if it were two different artifacts. The number of duplicates identified is far larger than the number of newly joined pieces. This means the machine's work is not only new discovery. It is also a check on old records.

When fragments are joined, sentences become longer. Often, only the question remains on one side, while the result is on the other. Only when two fragments meet do the day of divination and what actually happened connect in one line. They become one event record, not several short records.

This difference directly affects historical writing. If we look only at cut-off records, events seem scattered. If we look at joined records, the sequence appears. We can see when the king asked the question and when he received the result.

The 7,000 duplicate records that were identified carry the same weight[3]. If the same bone is counted as two, the number of sources itself becomes distorted. The same sentence may also be cited as two different records. By comparing rubbings in large numbers again, the machine filtered out these overlaps. This is not a dramatic discovery, but it sets the basic data straight.

Bones Scattered Across the Sea Meet on a Screen

Distance was the wall that blocked joining. If fragments from one bone were kept in different countries, they could not be placed side by side. Even if photographs were exchanged, thickness and curvature could not be known. Three-dimensional data lowered this wall.

The roughly 900 overseas items discussed earlier are an example of crossing this wall[3]. The physical objects remain where they have always been. What was brought over was information about the shape of their surfaces.

A researcher in Anyang can rotate those fragments on a screen. A fragment in Beijing and a fragment in Europe can also be placed on the same screen so their broken sections can be compared.

The way materials are opened to the public has also changed. In the past, people had to visit the holding institution to see an artifact. Getting permission alone could take months. Now students can open oracle bone materials from home[3]. When the number of people who handle the data grows, the number of eyes checking it also grows. When the number of eyes grows, the chances of catching wrong joins also grow.

What Machines Still Cannot Do

The achievements also made the limits clear. The S-OBI test released in late June 2026 showed those limits in numbers. Several large AI models solved the same oracle bone sentence questions. Their scores fell more on questions that required understanding the whole sentence than on questions that asked them to recognize one character. The cause also became clear. The models first misread characters that were plainly preserved. They then carried that error into their reasoning, so the answer went wrong[4]. Understanding the flow of a sentence from 3,000 years ago is still largely a human task.

The problem that requires even more caution is fabrication. AI chooses answers by probability. In a field with limited training data, such as oracle bone script, it can invent characters that look plausible but never actually existed. There is also not enough neighboring language material for comparison with the Shang language[5]. This makes it hard to check an answer against a dictionary.

Data imbalance remains as well. The oracle bone materials gathered so far are concentrated in certain periods and certain excavation sites. A machine trained on biased data is likely to make mistakes when it encounters oracle bones from an unfamiliar region. A 2026 review paper also named data quality and balance as future tasks[5].

In numbers, there is still a long way to go. About 3,000 characters still have unknown meanings. AlphaOracle solved one of them with evidence[1]. When we consider the labor required to solve one character, the work ahead is large. Still, once the method is organized, the next character becomes a little faster.

How Humans and Machines Divide the Work

For this reason, today's research method has settled into placing humans and machines side by side. The machine produces a broad set of candidates and organizes the evidence. The human checks those candidates against texts and physical objects.

The reduction in analysis time discussed earlier is also the result of this division of labor. The machine did not replace human judgment. It narrowed in advance the materials humans needed to inspect. This is why experts said the tool was useful[1].

The same method has taken hold in joining work. A program completes in a few days the comparisons that once took years[2]. Scholars take those candidates to the storage room. They place the actual bones together and check whether the broken sections fit. The final judgment is made at human fingertips.

This division of labor comes with one condition. The evidence for the machine's answer must be visible to humans. If the evidence cannot be seen, scholars have no way to check it. An answer that cannot be checked cannot enter historical writing. This is why AlphaOracle divided the process into stages[1].

Three thousand years ago, a Shang king looked at cracks in bone and asked about the future. Today's researcher looks at the broken sections of the same bone and asks about the past. The direction of the question is opposite, but the mark being examined is still the same.

References

[1] Liu, Y., Guan, H., Wang, P., Wang, X., Wan, J., Zhang, K., Zheng, H., Liu, X., Kuang, Z., Yang, H., Li, B., Liu, Y., Jin, L., & Bai, X. (2026). AlphaOracle: Oracle bone script decipherment via human-workflow-inspired deep learning. The Innovation, 7(11), 101462. https://doi.org/10.1016/j.xinn.2026.101462

[2] 北京日报. (2026, 8월 28일). 高校团队"唤醒"三千年前文明碎片. 新浪教育. https://edu.sina.com.cn/l/2026-08-28/doc-inipwhrw0020843.shtml (accessed September 3, 2026)

[3] 梁文艳. (2026, 5월 26일). 以科技赋能，让甲骨文"活"起来. 中国青年报. https://zqb.cyol.com/pc/content/202605/26/content_426168.html (accessed September 3, 2026)

[4] Li, Z., Chen, Z., Chen, T., & Zhai, G. (2026, 7월 1일). Beyond single character: Evaluating MLLMs for sentence-level oracle bone inscription understanding. arXiv:2606.31169. https://arxiv.org/abs/2606.31169 (accessed September 3, 2026)

[5] Chen, Z., Hua, W., Li, J., Zhu, Y., Zhi, X., Liu, Z., Chen, T., Zhang, W., & Zhai, G. (2026). Oracle bone inscriptions information processing: a comprehensive survey. npj Heritage Science, 14, 220. https://doi.org/10.1038/s40494-026-02511-w (accessed September 3, 2026)

Chapter 8. Brushwork Revived from a Mud Pit: Silla Wooden Slips and AI Cursive Reading

8.1. Case File

An Age When People Wrote on Wood

There was a time when paper was precious. People in ancient East Asia shaved wood into thin pieces and wrote on them. These wooden boards with writing are called wooden slips. A wooden slip was usually as wide as two or three fingers placed together. Many were a little longer than a handspan. Some were shaped like small pillars. These wooden slips could be written on all four sides. They are called multi-sided wooden slips because they are wooden boards with writing on several faces.

Wooden slips had many uses. One was as a tag attached to goods. When grain or salted seafood was sent from a local area to the royal capital, a wooden slip was tied to it with a cord. It recorded the name of the sending district, the name of the goods, and the quantity. Another use was as an official document. Some wooden slips record an official's name and rank, received goods, and a date[1]. There are also practice wooden slips on which the same Chinese character was written several times to learn writing.

Wood was harder than paper, so it could be used again. When a wooden slip was no longer needed, its surface was thinly shaved off with a knife. When a new surface appeared, people wrote on it again. The thin wood shavings cut off at this time are also found at archaeological sites. These scraps are called sakseol. The word means wood shavings cut away to erase writing. This can be seen in the 329 wooden slip-related finds from the Gwanbuk-ri site in Buyeo. Of these, 247 were sakseol. This means that a facility for making and revising documents was nearby[1].

Historians value wooden slips for a reason. History books are usually written after events are over. They contain the writer's judgment and the circumstances of the dynasty. Wooden slips are closer to records of the day's work written on that day. They were not decorated texts made to impress later generations. They preserve who sent how many seom of grain and when. Records made at the time an event happened are called primary sources. Wooden slips also contain mistaken characters and corrected places. So researchers place several slips side by side and compare them.

Wooden slips are found both on the Korean Peninsula and in the Japanese archipelago. Both places were societies that made documents in Chinese characters. At the same time, people searched for ways to write their own languages. Silla people mixed the meanings and sounds of Chinese characters to write Korean word order. These writing systems are called idu and hyangchal. Japanese people also borrowed the sounds of Chinese characters to write their own language. Wooden slips are physical evidence that shows when these attempts began. This is because almost no paper documents remain. This is why research continues to compare wooden slips from the two countries side by side.

A Thousand Years Protected by Mud

Wooden slips are usually found in certain places. They are found on pond floors near palaces. They are found in wetlands where water pooled inside fortress walls. They are found in mud layers that filled ditches. Wood that people used and threw away sank into the water and was buried there.

Almost no air enters underwater mud. The organisms that rot wood are mostly microorganisms that live by breathing air. When air is blocked, these microorganisms cannot do their work. Some microorganisms that live without air remain. Even so, the speed at which wood decays becomes much slower. That is why wood kept its shape for more than a thousand years. Wooden slips remained even in soil where leather and paper disappeared without a trace. A trash pit, not a tomb, protected the records.

Instead, other changes occurred. Water seeped into the wood and filled the cell walls. Iron and other minerals in the soil entered with it. The wood became fully soaked with water. This kind of wood is called waterlogged wood, or saturated wood.

Saturated wood must not simply be dried after excavation. As the water leaves, it shrinks quickly and cracks. This is why wooden slips are placed in water containers right away at excavation sites. In a conservation lab, the water inside the wood is slowly replaced with chemicals. This process alone can take months or years. Reading must be finished before then. This is because treatment can change the color of the surface.

Black Writing on Black Wood

Wood penetrated by minerals becomes dark. Many pieces pass through brown and turn completely black. This darkening of wood is called blackening. The problem is the material used for writing. People in the past made ink by mixing soot from burned pine with animal glue. The main component of ink is carbon, and its color is black.

There is black ink writing on blackened wood. No matter how much bright light is shone on it, the two cannot be distinguished. During the long time spent underwater, many ink particles were also washed away. Dark vertical lines in the wood grain mix with brush marks. Soil and decayed matter fill the grooves left by the brush. To the human eye, it is hard even to tell whether letters are present.

When reading is difficult, scholarly views also split. Different scholars may read the same wooden slip as different Chinese characters. If one character changes, the meaning of the sentence changes, and the interpretation of history changes too. So wooden slip researchers have long searched for new eyes. They needed a tool that could catch what human eyes miss.

The number of wooden slips that must be read keeps growing. Silla wooden slips have steadily been found at Seongsansanseong Fortress in Haman, Gyeongnam. In August 2025, two more inscribed wooden slips were newly identified[2]. An inscription means characters written on an artifact. At the Gwanbuk-ri site in Buyeo, Chungnam, large numbers of wooden slips from the Baekje Sabi period were found[1]. At Daemosanseong Fortress in Yangju, Gyeonggi, a wooden slip estimated to be from the 5th century was found. If this is correct, it would be among the earliest known Baekje written materials[3].

The situation is the same in Japan. The Nara National Research Institute for Cultural Properties gathers and discloses information on wooden slips found across Japan. This database is called Mokkan-ko[4]. It does not search only by the characters written on wooden slips. It can also search by use, date, material, and excavated region[5]. So much wood remains to be read.

References

[1] 파이낸셜뉴스. (2026, 2월 5일). 부여 관북리 '백제 피리' 첫 확인..목간 329점 무더기 발견. 파이낸셜뉴스. https://www.fnnews.com/news/202602051448564101 (accessed September 3, 2026)

[2] 경남일보. (2025, 8월 7일). 함안 성산산성서 새로운 목간 2점 명문 확인. 경남일보. https://www.gnnews.co.kr/news/articleView.html?idxno=616028 (accessed September 3, 2026)

[3] 뉴시스. (2025, 11월 20일). 1500년 전 '백제 문자' 깨어났다...양주 대모산성서 5세기 목간 출토. 뉴시스. https://www.newsis.com/view/NISX20251120_0003410107 (accessed September 3, 2026)

[4] 奈良文化財研究所. 木簡庫(MOKKAN-KO). https://mokkanko.nabunken.go.jp/ja/ (accessed September 3, 2026)

[5] カレントアウェアネス・ポータル. (2018). 奈良文化財研究所、木簡に関する情報を検索するシステム「木簡庫」(MOKKAN-KO)を公開. 国立国会図書館. https://current.ndl.go.jp/car/35714 (accessed September 3, 2026)

8.2. Forensic Tracking 1: A Digital Ink Magnifier

Using Light the Eye Cannot See

Human eyes see only a narrow range of light. There is light beyond the red of the rainbow. This light is called infrared. Its wavelength is long, so the eye cannot detect it, but a camera can capture it. The signal that a remote control sends to a television is also infrared.

Infrared is used to read wooden slips because of the properties of the materials. The carbon particles in ink absorb infrared well. Blackened wood grain reflects much infrared light. Two materials that look equally black to the human eye separate in front of an infrared camera. The side that absorbs light appears dark, and the side that reflects it appears bright. Brush marks that seemed to have disappeared rise up in the photograph.

The light mainly used is near-infrared light around a wavelength of 800 nanometers. A nanometer is one billionth of a meter. The red light that humans can see ends around 700 nanometers. This means researchers use light just beyond that range. Photography is done in a room where outside light is blocked. The light is spread evenly so it does not lean to one side. If it is shone from only one side, long shadows form in each groove of the wood grain.

Infrared photography has become a basic procedure in wooden slip research. In excavation reports, infrared photographs are usually placed beside photographs of wooden slips. The Nara National Research Institute for Cultural Properties has long developed equipment that makes wooden slip character images clearer with infrared cameras[1].

From Infrared to Hyperspectral Imaging

A single infrared photograph does not bring every character back to life. The wavelength that shows a character well differs depending on the position of stains and the condition of the wood grain. That is why hyperspectral imaging appeared.

An ordinary photograph stores only three values: red, green, and blue. Hyperspectral imaging is different. It records one scene by dividing it into dozens or hundreds of narrow wavelength bands. You can think of it as finely splitting light into rainbow colors and photographing each part separately. For each point on the screen, brightness values for each wavelength are recorded in a row. So instead of one photograph, it creates a thick bundle of photographs.

By comparing these values, researchers can distinguish places where ink remains from places where soil stains sit. This is because ink and soil change in brightness across wavelengths in different ways. Some characters are easy to see around 900 nanometers. Other characters come alive in other bands. Researchers select wavelength bands and find the combination that is easiest to read. Computers may also automatically compare several bands and narrow down candidates. This reduces by calculation work that people once checked one image at a time with their eyes.

The National Research Institute of Gaya Cultural Heritage first used this method in 2025. It applied hyperspectral imaging to the reading of wooden slips from Seongsansanseong Fortress in Haman. The institute said this method is more precise than existing infrared photography. It explained that even characters hard to recognize with the naked eye or ordinary photography can be recovered[2].

Finding Characters by Their Shapes

Even when a photograph becomes clear, the next problem remains. Researchers must decide which Chinese character a faint brush mark represents. Characters written in cursive script can change greatly in shape. Many characters have only a few strokes left. There are tens of thousands of Chinese characters, so there are limits to searching dictionaries by hand.

The Nara National Research Institute for Cultural Properties and the Historiographical Institute at the University of Tokyo built a system to help with this work. Its name is Mojizo. The user cuts out and uploads an image of the character they want to read. The system searches character images gathered from wooden slips and old documents. It finds and shows shapes that look similar[3]. A person then looks at the candidates and judges. It is not a machine that decides the answer. It is a search tool that brings up similar examples.

This kind of tool changed the order of work. In the past, scholars thought of candidates by relying on memory and character dictionaries. Now the machine first selects similar characters from tens of thousands of cases. The person chooses from the narrowed candidates by considering the context. The time needed for reading has been greatly reduced.

Distinguishing Strokes Drawn by the Machine

Today's generative AI can make blurred images clearer. It connects broken lines and fills empty places. If this technology is used on photographs of wooden slips, the result looks better. There is a trap here.

Strokes filled in by the machine are not ink that actually remains on the wood. They are guesses drawn from the average of tens of thousands of learned handwriting samples. If a wood grain pattern remains in a round shape, the machine may read it as a round stroke. Then a character that was never there can appear convincingly. The source material becomes contaminated, and historical interpretation is distorted.

So there are rules in research settings. The original photograph and the corrected result must be published side by side. Ink marks that can actually be confirmed and pixels created by the machine are marked separately. The machine's answer is treated not as a conclusion but as a list of candidates. A person makes the final reading. When a proposed reading is announced, researchers also state what equipment was used and under what conditions the image was taken.

Preparation is also under way at the government level. On April 16, 2026, the Presidential Advisory Council on Science and Technology approved the implementation plan of the Korea Heritage Service. It is a plan to introduce AI, robots, and sensors into the investigation and restoration of national heritage from 2026 to 2030. This year's investment is 14 billion won. Of that, 4.4 billion won was allocated to the development of intelligent advanced conservation technology[4]. The Gimhae National Museum held an academic symposium on June 23, 2026. The theme was a scientific view of Gaya heritage research. Seven presentations followed on nondestructive analysis, digital restoration, and three-dimensional X-ray microscopy[5]. Methods for looking inside artifacts that could not be seen with the naked eye are spreading beyond wooden slips.

References

[1] 奈良文化財研究所. 赤外線カメラを利用した木簡文字画像鮮明化システムの開発. 学術機関リポジトリデータベース. https://irdb.nii.ac.jp/01144/0004944317 (accessed September 3, 2026)

[2] 경남일보. (2025, 8월 7일). 함안 성산산성서 새로운 목간 2점 명문 확인. 경남일보. https://www.gnnews.co.kr/news/articleView.html?idxno=616028 (accessed September 3, 2026)

[3] 奈良文化財研究所, 東京大学史料編纂所. 木簡・くずし字解読システム MOJIZO. https://aimojizo.nabunken.go.jp/ (accessed September 3, 2026)
[4] 뉴시스. (2026, 4월 16일). "AI, 블록체인으로 국가유산 보존관리"...과기자문회의, 유산청 시행계획 의결. 뉴시스. https://www.newsis.com/view/NISX20260416_0003594216 (accessed September 3, 2026)
[5] 경남일보. (2026, 6월 17일). 국립김해박물관, 23일 '가야학술제전' 심포지엄 개최. 경남일보. https://www.gnnews.co.kr/news/articleView.html?idxno=639314 (accessed September 3, 2026)

8.3. Forensic Tracking 2: KuroNet and Miwo

Connected Writing Cannot Be Divided

In Japan, millions of old documents used until the Edo period remain. These documents were written in a cursive script called kuzushiji. Kuzushiji is a handwriting style used for more than a thousand years from the 8th century. In 1900, Japan changed its writing system. After that, schools stopped teaching this script. So most Japanese people today cannot read even books from 150 years ago[1].

The reason it is hard to read is not only the shapes of the characters. There is a writing method in which several characters are connected without lifting the brush from the paper. The last stroke of the previous character and the first stroke of the next character join into one. Then no boundary remains between characters. A whole line looks like one long continuous line.

Computer character recognition long worked in the same order. First, it cut out each character into a square box. Then it guessed what character the image inside the cut box was. This method works only when there are boundaries. In front of connected cursive writing, it collapsed at the first step. If the beginning was wrong, the rest was wrong too, so there was no way to improve accuracy.

KuroNet, Which Does Not Divide Characters

Researchers at the Center for Open Data in the Humanities in Japan found another path. Tarin Clanuwat, Alex Lamb, and Asanobu Kitamoto built a model together. Its name is KuroNet. It was first presented in 2019 at the International Conference on Document Analysis and Recognition[1]. In 2020, they refined its structure and published the results again in an academic journal[2].

KuroNet does not cut out characters in advance. It takes a whole page of an old document as input. The central structure of the model is U-Net. U-Net gradually reduces an image and extracts features. Then it enlarges it back to its original size. It is a neural network. The reducing side and the enlarging side are symmetrical, so it is drawn in the shape of the letter U. While reducing the image, it learns the thickness and direction of brush strokes. While enlarging it, it returns that information to its original position.

This neural network predicts two things at once across the whole page. One is where the center of each character is. The other is what character is at that position. The result appears as a probability map. It is an image that marks places on the screen where characters are likely to exist. If one follows the places with high probability, the position and name of the character are fixed at the same time. Because there is no cutting stage, it does not break down even with connected cursive writing.

For training, the researchers used data made by Japan's National Institute of Japanese Literature. It is called the Japanese classical Kuzushiji dataset. It contains more than one million character samples[2]. The same research team also released a smaller version so that machine learning researchers could use it more easily. This is Kuzushiji-MNIST, modeled on handwritten digit data. They also released a version with more characters and a version that collected only Chinese characters[3]. Thanks to this, researchers who did not know Japanese classics could enter this problem.

KuroNet handed over its mission in October 2022. The name of the recognition model changed to Ruri[4]. Ruri is a model that refines object detection technology for cursive writing. It reads better than before even when the color or pattern behind the characters is complex. It was trained with an improved version of the same dataset, and it reads characters that earlier models could not read[5].

Miwo in the palm of the hand

The next task was to put a model that had run on a server into a mobile phone. The app that came out of that work is Miwo. It was first released on August 30, 2021[5]. Its Japanese spelling is みを, and its English spelling is Miwo[6].

It is simple to use. One photographs an old document with a mobile phone camera. After a short while, the screen shows the result converted into modern characters. The original and the result can be compared by sliding them over each other. The app also shows what the hentaigana were. Hentaigana are old hiragana forms that are no longer used. There is also a function for exporting the result as a text file. The photographed images are not left on the server. It can be used for free on both Android and iPhone[6].

From version 1.1, released on October 26, 2022, the recognition model changed to Ruri[5]. In the same month, the app received the 2022 Good Design Award[7]. Usage records continued to accumulate. As of September 3, 2026, the number of images recognized since its release was 3,523,373[6]. This means that people are photographing old documents in libraries, temples, and local archives.

Citizens read together

The method in which a machine reads first and a person corrects it is used elsewhere as well. There is an online transcription project called Minna de Honkoku. Transcription means copying old handwriting into modern characters. Its spelling is みんなで翻刻, and it means transcribing together. It began in 2017 at Kyoto University's Study Group on Historical Earthquakes. It is now run jointly by the National Museum of Japanese History and the Earthquake Research Institute of the University of Tokyo[8].

Participation is open. Anyone with an internet browser can transcribe old documents. The screen shows a first draft read by artificial intelligence. The person only needs to correct the wrong characters. The site uses the Kuzushiji recognition service of the Center for Open Data in the Humanities. So far, about 10,800 participants have transcribed more than 48 million characters[8].

The transcribed text does not simply pile up. The correct answers fixed by people become training data again. When training data increases, the model improves. When the model improves, the first draft becomes more accurate, and people need to do less work. Once this cycle begins, it gains speed.

The uses of the documents that have been read are also broad. Earthquake researchers started this project. Old diaries and government documents record the dates and damage of earthquakes, tsunamis, and poor harvests. When records that could not be handled because of cursive writing become data, researchers can count the cycles of past disasters. Old brush writing becomes today's disaster prevention data.

This is the story of Japanese paper documents. KuroNet, Ruri, and Miwo are all models that read Kuzushiji written on paper. No confirmed case has yet been found in which these models automatically read wooden slips from mud. Wooden slips contain few characters. Their base is wood, and much of the ink has disappeared. Correct-answer data for training has not accumulated as much as it has for paper documents. The situation may change if more Korean and Japanese wooden slip data is gathered. For now, this is an expectation, not an achievement.

References

[1] Clanuwat, T., Lamb, A., & Kitamoto, A. (2019). KuroNet: Pre-modern Japanese Kuzushiji character recognition with deep learning. 2019 International Conference on Document Analysis and Recognition (ICDAR), 607-614. https://doi.org/10.1109/ICDAR.2019.00103

[2] Lamb, A., Clanuwat, T., & Kitamoto, A. (2020). KuroNet: Regularized residual U-Nets for end-to-end Kuzushiji character recognition. SN Computer Science, 1, 177. https://doi.org/10.1007/s42979-020-00186-z

[3] Clanuwat, T., Bober-Irizar, M., Kitamoto, A., Lamb, A., Yamamoto, K., & Ha, D. (2018). Deep learning for classical Japanese literature. arXiv. https://arxiv.org/abs/1812.01718

[4] ROIS-DS 人文学オープンデータ共同利用センター. KuroNetくずし字認識サービス(AI OCR). CODH. https://codh.rois.ac.jp/kuronet/ (accessed September 3, 2026)

[5] ROIS-DS 人文学オープンデータ共同利用センター. みを(miwo)とは？ CODH. https://codh.rois.ac.jp/miwo/about/ (accessed September 3, 2026)

[6] ROIS-DS 人文学オープンデータ共同利用センター. みを(miwo)：AIくずし字認識アプリ. CODH. https://codh.rois.ac.jp/miwo/ (accessed September 3, 2026)

[7] 国立情報学研究所. (2022, 10월 26일). 世界初のAIくずし字認識アプリ「みを(miwo)」が2022年度グッドデザイン賞を受賞. 国立情報学研究所. https://www.nii.ac.jp/news/release/2022/1026.html (accessed September 3, 2026)

[8] みんなで翻刻. みんなで翻刻について. https://honkoku.org/ (accessed September 3, 2026)

8.4. The Truth Revealed

Royal Household Management Shown by the Wolji Wooden Slips

Donggung Palace and Wolji Pond in Gyeongju were a detached palace and pond of the royal family of Unified Silla. Wooden slips first appeared in the 1975 excavation of Wolji. As investigations continued afterward, their number increased. The Gyeongju National Museum displays artifacts from this site in the Wolji Hall[1].

The Wolji wooden slips show a day in the life of the royal household. Some wooden slips were labels attached to jars. Label wooden slips have grooves cut on the left and right at the end. A string was wound around those grooves to tie the label to an object. Their use was similar to today's parcel invoice. Wooden slips from Wolji contain the characters nokgap and ong. Nokgap means salted seafood made from deer, and ong means jar. This means that the Silla royal family recorded and managed the type of salted food and the time it was made[2].

These records do not appear in history books. Samguk Sagi recorded kings, wars, and institutions. It did not record what came into the kitchen. Wooden slips fill that gap. One wooden board from the mud becomes a household ledger of the palace.

The Korea Society for the Study of Wooden Documents and the Gyeongju National Museum held an international academic conference on April 3 and 4, 2026. The theme was the space and written life of Gyeongju, the royal capital of Silla. It took place for two days at the Gyeongju Hwabaek Convention Center. Materials from the Wolseong moat, Donggung Palace, and Wolji Pond were placed together. It was an attempt to redraw the royal capital of Silla through wooden slips, stone monuments, and ink inscriptions. Researchers from Korea and Japan presented together[3].

The Punishment Report on the Seongsansanseong Wooden Slips

Haman Seongsansanseong is a mountain fortress built by Silla in the 6th century. Wooden slips have continued to come out of the wetland inside the fortress wall. As of August 2025, 247 ink-inscribed wooden slips had been found there. That is more than half of all ancient wooden slips found in Korea[4].

Two newly deciphered pieces from the 18th excavation were made from pine wood. One is a multi-sided wooden slip with writing on four sides. The other is a two-sided wooden slip[4]. The multi-sided wooden slip contained administrative content on three of its four sides. It is a record of a punishment imposed on a person[5]. It is thought to have been made in the mid-6th century[4].

This reading was obtained through the hyperspectral imaging seen earlier. Characters that could not be recognized by the eye or in ordinary photographs came back to life[4]. A report by a low-ranking Silla official that had been in mud for more than a thousand years returned as sentences. The Seongsansanseong wooden slips also record many names of villages that sent grain. They are materials that show how 6th-century Silla divided and governed its local areas. They also allow us to trace which routes brought taxes to the fortress.

A Pile of Documents from the Baekje Royal Palace

The Gwanbuk-ri site in Buyeo is known as the royal palace site of Baekje's Sabi period. The Buyeo National Research Institute of Cultural Heritage carried out its 16th excavation in 2024 and 2025. There, 329 wooden slip-related items were found. There were 82 wooden slips with writing and 247 shavings[6]. This is the largest number from a single site in Korea, and they are early examples among Sabi-period wooden slips[7].

Some wooden slips include dates. The sexagenary cycle years Gyeongsin and Gyehae appear. They are understood as 540 and 543, respectively. This means they are records from just after Baekje moved its capital from Ungjin to Sabi. Personnel records of officials, ledgers recording state finances, and names of government posts were also confirmed[7]. An actual Baekje wind instrument made of bamboo was also found for the first time[8].

On July 22 and 23, 2026, an international academic conference was held. It was organized by the Buyeo National Research Institute of Cultural Heritage, the Korea Society for the Study of Wooden Documents, and the Baekje Society. The theme was the musical instruments and wooden slips of the Gwanbuk-ri site in Buyeo. It was held in Haeundae-gu, Busan. It began with a keynote lecture by Park Soon-bal, professor emeritus at Chungnam National University. Twelve papers were presented in three sessions. Researchers from Korea, China, and Japan examined the meaning of the Gwanbuk-ri wooden slips together[8].

Caution Before a Stone Monument

Methods that work for wooden slips do not work for all old characters. When the base changes to stone, the situation changes. The Gwanggaeto Stele is one example. This stele, erected in 414, has the history of Goguryeo carved on it. The surface of the stele has been worn down by long exposure to wind and rain. Here, one must read carved grooves, not ink.

For that reason, researchers compare rubbings. Original-stone rubbings are rubbings made before lime was applied to the surface of the stele. Paper was placed directly on the original stone and tapped to make them. About ten have been confirmed so far. In the late 19th century, another type appeared. These were lime rubbings, made after lime was applied to the stele surface to make the characters look clearer. Several copies called ssanggu gamukbon were also made. Paper was placed over the inscription, and lines were drawn around the characters. The blank space was filled with dark ink to make it look like a rubbing. This is not a rubbing made by tapping. It is a copied replica drawn by hand. Characters may have been copied incorrectly during the process[9]. That is why the same character can look different from one rubbing to another.

The center of the debate is the sentence in the Sinmyo year passage. The interpretation of this sentence became tied to the ancient history of Korea-Japan relations and was disputed for a long time. The issue was whether the characters had been changed by human hands[10]. To identify the original characters, early original-stone rubbings and several lime rubbings must be compared. The first task is to set standards for distinguishing each rubbing[9].

Stories circulate that three-dimensional measurement and artificial intelligence have answered this debate. When checked, the studies offered as evidence are often hard to find. The expectation that a method that worked for wooden slips will also work for a stone monument is different from actual research results. The technologies seen earlier make what was invisible visible. Deciding the meaning of what has been made visible comes next.

So the conclusion has two layers. Infrared and hyperspectral imaging revived ink that had disappeared from Silla and Baekje wooden slips. Recognition models without segmentation read indivisible cursive writing in Japanese paper documents. The objects handled by the two technologies differ, but their direction is the same. Machines first point out strokes that human eyes missed. At the same time, without rules for distinguishing strokes drawn by a machine from strokes left in the original, the same technology creates history that did not exist. It was the machine that awakened the brush writing, and it is people who are responsible for that writing.

References

[1] 국립경주박물관. 월지관. 국립경주박물관. https://gyeongju.museum.go.kr/kor/html/sub02/0201.html (accessed September 3, 2026)

[2] 국사편찬위원회. 목간. 우리역사넷. https://contents.history.go.kr/mobile/tt/view.do?levelId=tt_b36 (accessed September 3, 2026)

[3] 경북일보. (2026, 3월 31일). 1500년 전 신라의 목소리 깨어난다...경주서 국제 목간 학술대회. 경북일보. https://www.kyongbuk.co.kr/news/articleView.html?idxno=4068614 (accessed September 3, 2026)

[4] 경남일보. (2025, 8월 7일). 함안 성산산성서 새로운 목간 2점 명문 확인. 경남일보. https://www.gnnews.co.kr/news/articleView.html?idxno=616028 (accessed September 3, 2026)

[5] 경남도민일보. (2025, 8월 7일). 누군가를 처벌한 신라 행정 보고, 함안 성산산성 목간서 확인. 경남도민일보. https://www.idomin.com/news/articleView.html?idxno=943461 (accessed September 3, 2026)

[6] 파이낸셜뉴스. (2026, 2월 5일). 부여 관북리 '백제 피리' 첫 확인..목간 329점 무더기 발견. 파이낸셜뉴스. https://www.fnnews.com/news/202602051448564101 (accessed September 3, 2026)

[7] 경향신문. (2026, 2월 17일). 부여 관북리서 발견된 목간 329점, '백제 피리'만큼이나 중요한 이유. 경향신문. https://www.khan.co.kr/article/202602171427001 (accessed September 3, 2026)

[8] 이데일리. (2026, 7월 21일). 백제 왕궁터 출토 목간 329점·횡적 발굴 성과, 국제학술대회서 조명. 이데일리. https://edaily.co.kr/News/Read?mediaCodeNo=257&newsId=03276726645516816 (accessed September 3, 2026)

[9] 국립중앙박물관. 광개토대왕비 탁본. 국립중앙박물관 큐레이터 추천 소장품. https://www.museum.go.kr/MUSEUM/contents/M0501000000.do?schM=view&relicRecommendId=254462 (accessed September 3, 2026)

[10] 국사편찬위원회. 광개토대왕릉비. 우리역사넷. https://contents.history.go.kr/mobile/kc/view.do?levelId=kc_r100300 (accessed September 3, 2026)

Chapter 9. The Aegean's Lost Riddle: Linear A and the Geometric Grid

9.1. Case File

Three Kinds of Writing from the Palaces of Crete

Crete lies in the southern Aegean Sea. In the Bronze Age, the Minoan civilization grew on this island. The Minoans built ships and sailed across the Mediterranean. They carried olive oil, wine, and grain. The palaces collected these goods and then distributed them to people.

Around 1900, the British archaeologist Arthur Evans excavated the site of the palace at Knossos. He found clay tablets with writing carved on them underground. The writing was not of one kind. Evans divided the scripts he found into three groups: Cretan hieroglyphic, Linear A, and Linear B[1].

Minoan civilization lasted for more than a thousand years. Then, in the late Bronze Age, it went through great change. One cause was a volcanic eruption on Santorini to the north. The eruption is thought to have occurred between the 17th and 16th centuries BCE. The exact date is still debated[7]. Volcanic ash is thought to have reached Crete, and tsunamis are thought to have struck the coast. Around 1450 BCE, Mycenaean Greeks took control of the island[1]. More than a hundred years separate the eruption and Mycenaean rule. It is difficult to tie the two events directly as cause and effect. After that, the language of the Minoans disappeared from the record.

The three scripts were used at different times. Linear A was used from around 1750 BCE to around 1450 BCE. Linear B appeared later. It was the script used after Mycenaean Greeks took control of Crete. The two scripts share many signs with similar shapes[1].

The survival of the clay tablets was an accident. The Minoans did not fire their clay tablets. They carved writing into wet clay, dried it, used it for about a year, and then discarded it. But when the palaces burned, the clay tablets were fired like pottery. Only the tablets baked by fire survived for 3,500 years. The materials that remain today are materials left by fire.

Among the three scripts divided by Evans, Cretan hieroglyphic appeared first. It is a script in which shapes such as human heads and vessels remain clearly visible. Linear A is closer to a form that reduces those pictures into lines. The two scripts were used together during nearly the same 200-year period. One did not appear only after the other disappeared. Some scholars think Linear A was made after Cretan hieroglyphic. But we cannot conclude that the three scripts form one straight line of development. They clearly share signs, but the path connecting them has not yet been revealed[1].

Linear B Was Read

Linear B was read in the 1950s. The British architect Michael Ventris achieved this. He showed that the language written in Linear B was an old form of Greek. The classical scholar John Chadwick helped his research[1].

Ventris began by drawing a table. He placed consonants along the horizontal rows and vowels along the vertical columns[8]. This kind of table is called a syllabic grid. It is a chart with spaces prepared in advance for each syllable. Ventris used statistics to narrow down which sign belonged in which space[1].

He fitted place names such as Knossos and Phaistos into the spaces he had narrowed down. Once the place names matched, the remaining signs began to fall into place. Scholars tried the same method with Linear A. They transferred the sound values of Linear B to similar signs[3]. Sounds emerged. But no one could understand what language those sounds represented[1].

It is strange to have sound values but no meaning. This is because reading letters and understanding language are different tasks. A person who knows Hangul can read an unfamiliar foreign language written in Hangul. The sounds come out, but the meaning is unknown. Scholars facing Linear A are in the same position.

There Is No Stone Written in Two Languages

When scholars read ancient scripts, they usually look for a matching document. This is a document that writes the same content side by side in two languages. Such a document is called a bilingual inscription. Egyptian hieroglyphs were deciphered thanks to the Rosetta Stone. The Rosetta Stone had the same decree carved in three scripts. One of them was ancient Greek, which scholars already knew how to read.

Mesopotamian cuneiform followed a similar path. On the Behistun cliff in Iran, an inscription remained in three languages. At first, none of the three languages could be read. Ancient Persian was solved first. It had few signs, and the words were separated. Scholars used the deciphered ancient Persian as a foothold to read the rest.

No such document has been found on Crete. There is no artifact with a Greek translation written beside Linear A. The writing left by the Minoans is written only in the language of the Minoans. Without a matching counterpart, there is no way to confirm the meaning[1].

There are no markings that could serve as clues either. Champollion, who read Egyptian hieroglyphs, had cartouches. These were long borders surrounding royal names. It was already known that the writing inside these borders represented personal names. Champollion recovered several sound values from them and used those values to solve the rest. Linear A clay tablets have no such borders[1].

There is another problem. We do not know the roots of the language written in Linear A. Scholars call this language Minoan. It has not been shown which language family Minoan belongs to. Some argue that it belongs to the Indo-European family. Others argue that it comes from western Asia. Among these, the Semitic and Lycian lines are often seen as having better evidence. Lycian was an ancient language used in what is now southwestern Türkiye[1]. Neither view has won scholarly agreement.

Linear A is not found only on Crete. It has also been found on several Aegean islands and on the Greek mainland[1]. It has been found at sites buried by the volcano on Santorini. The script spread along the sea routes used by the Minoans. It is evidence of how widely connected the Bronze Age Mediterranean was.

This is why scholars keep working on Linear A. Linear A is one of the old scripts found on European soil. It is the only record the Minoans left in their own language. If Linear A is deciphered, it will reveal what they believed and what they bought and sold. The roots of the Minotaur story preserved in Greek myth also lie here.

1,400 Short Account Books

The amount of surviving material is not large. An inscription means writing carved or drawn on an artifact. About 1,370 Linear A inscriptions have been identified so far. Even if all the signs written in them are added together, the total is only around 7,400[1]. This amount cannot compare with the records of Egypt or Mesopotamia.

The texts are short as well. Most clay tablets are warehouse accounts. They record what came in, how many items there were, and how much went out to whom. They are mostly lines of personal names, object names, and numbers. Sentences that look like real sentences are rare. Without sentences, there are fewer ways to identify grammar.

When counting signs, we must separate the layers. There are 97 signs that write syllables. There are about 50 signs that represent whole object names. There are 4 signs for numbers and 17 signs for fractions. About 150 combinations made by overlapping two signs are also known[6]. The standard sign list that gathers all of these reaches 300 signs[2]. Among them, only a small number appear often. The rest appear only a few times. To entrust statistics to a computer, the same sign must appear many times. Signs that appear only a few times cannot be captured by statistics.

So researchers decided to rebuild the material itself. This means turning hand-copied drawings into a form that computers can handle[2]. In 2026, a method also appeared for rereading incised marks with light[4]. It reveals shallow grooves that the eye had missed[5]. The work of awakening the silence of the Aegean began again with this refining of the data.

References

[1] Braović, M., Krstinić, D., Štula, M., & Ivanda, A. (2024). A systematic review of computational approaches to deciphering Bronze Age Aegean and Cypriot scripts. Computational Linguistics, 50(2), 725-779. https://doi.org/10.1162/coli_a_00514

[2] Salgarella, E., & Castellan, S. (2021). SigLA: The Signs of Linear A. A palaeographical database. In Grapholinguistics in the 21st Century 2020: Proceedings (Grapholinguistics and Its Applications, Vol. 5, pp. 945-962). Fluxus Editions. https://sigla.phis.me/ (accessed September 3, 2026)

[3] Younger, J. G. (n.d.). Linear A texts and inscriptions in phonetic transcription. University of Kansas. http://people.ku.edu/~jyounger/LinearA/ (accessed September 3, 2026)

[4] Giorgi, L., & Lopez, S. (2026). Aegean Bronze Age writing systems through RTI: New potential and perspectives. Digital Applications in Archaeology and Cultural Heritage, 41, e00508. https://www.sciencedirect.com/science/article/abs/pii/S2212054826000159 (accessed September 3, 2026)

[5] GreekReporter. (2026, 3월 14일). New technology reveals hidden secrets of 3,500-year-old Bronze Age Aegean scripts. greekreporter.com. https://greekreporter.com/2026/03/14/new-technology-reveals-hidden-secrets-aegean-bronze-age-scripts/ (accessed September 3, 2026)

[6] Del Freo, M. (2022). Linear A. Mnamon: Ancient writing systems in the Mediterranean. Scuola Normale Superiore. https://doi.org/10.25429/sns.it/lettere/mnamon022 (accessed September 3, 2026)

[7] Pearson, C. L., Brewer, P. W., Brown, D., Heaton, T. J., Hodgins, G. W. L., Jull, A. J. T., Lange, T., & Salzer, M. W. (2018). Annual radiocarbon record indicates 16th century BCE date for the Thera eruption. Science Advances, 4(8), eaar8241. https://doi.org/10.1126/sciadv.aar8241

[8] Stavrakakis, E. (2024). Ventris' Grids. Bulletin of the Institute of Classical Studies, 67(1), 86-103. https://doi.org/10.1093/bics/qbae033

9.2. The Forensic Science Team’s Pursuit: Fitting an Unsupervised Phonological Grid

A Machine That Finds Matches Without a Dictionary

To build a translator, we usually need matching sentences. Korean sentences and English sentences are stacked side by side and shown to a machine. This kind of material is called a parallel corpus. Linear A has no parallel corpus[3].

So researchers chose a method that learns without an answer sheet. It finds rules from statistics alone, without being given answers in advance. This method is called unsupervised learning. The machine counts the order in which signs appear together. Then it compares those statistics with those of another writing system and looks for matches.

In 2019, researchers at the Massachusetts Institute of Technology introduced this method. The model was created by Jiaming Luo, Yuan Cao, and Regina Barzilay. They turned the task of connecting two writing systems into a minimum-cost flow problem. It is a method that places channels between signs and assigns values to the flow through them. The channel with the lowest cost is more likely to be the correct match[1].

The researchers tested this model on Ugaritic. It beat the previous best result by 5 percentage points. They also tested it on Linear B. It correctly transferred 67.3 percent of words that shared roots with Greek[1]. But Linear B had an answer. The researchers already knew that the script wrote Greek.

There was a reason the researchers chose Linear B as a test case. Only a problem with a known answer can measure a model's performance[1]. If the model is applied directly to a problem with no answer, there is no way to know whether the result is correct. So researchers first measure performance on a script that has already been read, then move on to an undeciphered script[3].

What the model produces is not a finished translation. It is a list of candidate sound values that might attach to each sign[1]. A person takes this list and filters it by comparing it with artifacts. The computer's role is to reduce the amount that people must examine[3].

Turning the Shape of Sound into Coordinates

Linear A is different. We do not even know which language it should be compared with[3]. In 2021, the same laboratory introduced a model that handled this condition. Luo, Frederic Hartmann, and Enrico Santus built it together with Barzilay and Cao[2].

This model handles two difficulties at once. The first difficulty is that there are no word boundaries. Ancient writing often has no spaces. The second difficulty is that the related language has not been identified[2].

The researchers turned the shape of sound into numerical coordinates. Using the International Phonetic Alphabet as a base, they placed the properties of consonants and vowels in a coordinate space. Sounds made with the lips are close to each other in the coordinates. They are far from sounds made with the root of the tongue. With these coordinates, the machine makes fewer absurd matches[2].

The researchers tested this model on Gothic and Ugaritic. Then they applied it to Iberian, whose meaning has not yet been solved. For the Iberian script, the sound values of the signs are fairly well known. What remains unsolved is the language those sounds record. The researchers found no evidence that Iberian is related to Basque. This result did not conflict with the mainstream scholarly view. They also showed a procedure for comparing candidates for related languages[2]. The same method can be used directly in Linear A research[3].

Rules Revealed by Blank Spaces

The machine compares more than the shapes of signs. It also calculates the habits by which signs appear with their neighbors. Some signs appear only at the beginning of words. Some signs attach only at the end. In some cases, two signs never appear side by side at all[3].

This kind of statistic is called co-occurrence statistics, meaning how often things appear together. When co-occurrence statistics are gathered, the blank spaces in the table appear. Signs placed in the same horizontal row are likely to share the same consonant. Signs placed in the same vertical column are likely to share the same vowel[8]. The computer repeats much faster the work Ventris once did by hand[3].

The computer finishes this calculation much faster than a person. If there are 300 signs, the number of possible matches reaches tens of thousands. This is too much for a person to count by writing it out on paper. Ventris carried out this work over several years[3].

Up to this point, the task is to draw the structure of sound. Even if the structure is drawn, it does not reveal what the sound means. Filling the grid and understanding the language are separate tasks[3].

Turning Letter Shapes into Data

For a machine to produce statistics, the material must be in a form the machine can read. Calculations cannot be done with only pictures printed in paper books. A data portal called SigLA took on this task. It was created by Ester Salgarella of Cambridge and Simon Castellan of Rennes[4].

SigLA contains 300 standard signs and about 400 hand-copied inscriptions. Within those inscriptions, more than 3,000 individual signs can be found. The data and images are open for anyone to use. The strokes of a single sign can also be separated and compared[4].

A study published in March 2026 broadened the texture of the data by one more layer. Lavinia Giorgi and Sara Lopez applied reflectance transformation imaging to Aegean scripts. This method photographs the same artifact several times while shining light on it from different directions. When these photographs are combined, tiny surface features become visible. Researchers can even examine which stroke the engraving tool made first[7].

As this kind of data accumulates, it becomes useful for other scripts as well. In 2022, researchers at the University of Bologna took another look at the Cypro-Minoan script. The study was carried out by Michele Corazza, Fabio Tamburini, Miguel Valério, and Silvia Ferrara. They compared sign shapes with a neural network that used no prior knowledge. As a result, they concluded that the old division into three groups should be reconsidered. The statistics fit better if the script is seen as one writing system[5].

Results produced by a computer do not become scholarship right away. Archaeologists and philologists examine them again[3]. They check which tablet came from which room. They compare habits of stroke formation to see whether the same scribe wrote the texts[7]. Only after passing this review do the results enter academic data[3].

Research also continues on the three Aegean scripts together. This approach places Cretan hieroglyphs, Linear A, and Cypro-Minoan script side by side. When the three scripts are compared, rules appear that are not visible when only one script is studied. Computers are powerful in this kind of large-scale comparison[3].

What Machines Cannot Do for Us

These methods have not revealed the meaning of Linear A. Three walls still remain.

The first is a counterpart for comparison. Unsupervised alignment works well when there is a relationship between an undeciphered script and a known script[1]. Minoan has no surviving related language[3]. The coordinates can be drawn, but there is no fixed reference point to anchor them.

The second is the amount of data. The 1,370 inscriptions are too few for machine learning[3]. When data is scarce, a large neural network invents convincing falsehoods. This kind of error is called hallucination.

The third is verification. An article published in The Conversation in July 2026 pointed out this problem. It was written by Jane Adkins of Dublin City University. She wrote that matching patterns with statistics cannot create meanings that do not exist. Her point was that a foothold, such as a known language or a bilingual text, is needed to separate chance matches from real rules. Machines only narrow the candidates. Deciding whether they are right remains the work of humans[6].

References

[1] Luo, J., Cao, Y., & Barzilay, R. (2019). Neural decipherment via minimum-cost flow: From Ugaritic to Linear B. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics. https://aclanthology.org/P19-1303/ (accessed September 3, 2026)

[2] Luo, J., Hartmann, F., Santus, E., Barzilay, R., & Cao, Y. (2021). Deciphering undersegmented ancient scripts using phonetic prior. Transactions of the Association for Computational Linguistics, 9, 69-81. https://doi.org/10.1162/tacl_a_00354

[3] Braović, M., Krstinić, D., Štula, M., & Ivanda, A. (2024). A systematic review of computational approaches to deciphering Bronze Age Aegean and Cypriot scripts. Computational Linguistics, 50(2), 725-779. https://doi.org/10.1162/coli_a_00514

[4] Salgarella, E., & Castellan, S. (2021). SigLA: The Signs of Linear A. A palaeographical database. In Grapholinguistics in the 21st Century 2020: Proceedings (Grapholinguistics and Its Applications, Vol. 5, pp. 945-962). Fluxus Editions. https://sigla.phis.me/ (accessed September 3, 2026)

[5] Corazza, M., Tamburini, F., Valério, M., & Ferrara, S. (2022). Unsupervised deep learning supports reclassification of Bronze age cypriot writing system. PLOS ONE, 17(7), e0269544. https://doi.org/10.1371/journal.pone.0269544

[6] Adkins, J. (2026, 7월 27일). Cracking the code: can AI help us decipher ancient languages? The Conversation. https://theconversation.com/cracking-the-code-can-ai-help-us-decipher-ancient-languages-288238 (accessed September 3, 2026)

[7] Giorgi, L., & Lopez, S. (2026). Aegean Bronze Age writing systems through RTI: New potential and perspectives. Digital Applications in Archaeology and Cultural Heritage, 41, e00508. https://www.sciencedirect.com/science/article/abs/pii/S2212054826000159 (accessed September 3, 2026)

[8] Stavrakakis, E. (2024). Ventris' Grids. Bulletin of the Institute of Classical Studies, 67(1), 86-103. https://doi.org/10.1093/bics/qbae033

9.3. Truths Revealed

Numbers and Fractions Have Been Solved

The first part of Linear A to be solved was its way of counting. The Minoans used a decimal system that counted in groups of ten. One vertical line meant 1, and one horizontal line meant 10. A circle meant 100, and a circle with rays meant 1000[3]. This rule could be guessed just by looking at the signs.

The difficult part was the fraction signs. Seventeen fraction signs appear on the clay tablets[8]. If signs written on top of each other are counted, the number comes close to thirty[9]. For a long time, no one knew which sign meant which value.

In 2021, researchers at the University of Bologna addressed this problem. The study was carried out by Michele Corazza, Silvia Ferrara, Barbara Montecchi, Fabio Tamburini, and Miguel Valério. They combined the old method of studying the shapes of signs with computer calculation[1].

The researchers first organized the rules that the signs followed on the clay tablets. They counted which signs appeared together and which signs did not. Then they used computers to examine combinations of values. There were nearly 4 million candidates. They narrowed these candidates through statistical tests and comparisons with fraction systems from other civilizations[1].

The remaining answer was this. In the Minoan fraction system, the smallest value was 1/60. With this system, most values with 60 as the denominator could be written. The sign meaning 1/10 continued unchanged into Linear B. In Linear B, this sign was used as a unit for measuring dry grain. About 250 years separated the records of the two scripts. Even across that gap, the counting method survived[1].

Once the counting system was solved, the structure of the accounts also began to appear. The clay tablets follow a method of pairing commodity signs with quantities. The same format repeats across several sites. This means the palaces of Crete shared the same accounting rules[3]. Even without knowing the meaning, the nature of the documents appears in this way.

Place Names Can Be Read, but Their Meaning Has Not Been Solved

Something has also been gained on the sound side. John Younger of the University of Kansas has published a list that transcribes Linear A inscriptions into sounds. This transcription assigns the same sound values to signs that have the same shapes in Linear B[2].

In this transcription, several words appear to be old place names on Crete. They are names that survived into later periods, such as Knossos and Phaistos[2]. The fact that place names fit means the hypothesis about sound values is not entirely wrong[3].

The sound-value hypothesis also has clear limits. Even if signs look the same in the two scripts, there is no guarantee that their sounds were the same. The people who created Linear B borrowed another people's script to write their own language. In the process of borrowing it, sounds may have shifted[3]. For that reason, the transcription is used only as a provisional reading[2].

But this is as far as it goes. Apart from place names, the meanings of the other words do not connect. They can be read aloud, but we do not know what they mean. The bridge between sound and meaning has not been built[3].

The Commotion of 2026

In the summer of 2026, an announcement claimed that Linear A had been read. The claim was made by artificial intelligence engineer Tom Di Mino. He had not received formal training in linguistics[4].

Di Mino said he wrote Python scripts and searched two data sources. They were GORILA, a collection of hand-copied materials, and SigLA, discussed above. In his first announcement, he said he had assigned sound values to 40 signs and made a list of 408 words[4]. In August, he released a 41-page paper and expanded the list to 508 entries. It included readings of 42 signs and translations of 443 inscriptions. He argued that Minoan belonged to the Semitic language family, which includes Hebrew and Aramaic[5].

The academic community did not accept this claim. No peer-reviewed paper had been published[4]. Reviews were said to be under way at Rutgers University and the University of Cambridge. But no conclusion had been reached yet[5]. Claims that Linear A has been read have appeared several times over the past hundred years. For that reason, researchers approach new claims with caution[6].

Another interpretation appeared in the same August. It concerned an inscription on a stone table from the peak sanctuary on Mount Juktas. It was a libation table used to pour wine or other liquid offerings to a god. The proposal was that two names carved on this stone were connected to later Greek gods. This proposal also remains a hypothesis. Linear A has not been securely deciphered[7].

This story leaves something for artificial intelligence research as a whole. Machines are good at finding rules where rules exist. They produced results with data that must fit together, such as fraction signs[1]. But there are places that rules alone cannot reach. Which sound is tied to which meaning is a matter of convention. When everyone who knows that convention disappears, calculation has only a narrow path for bringing it back[6].

Keeping the Remaining Data Open

The research method has also changed. In the past, images of inscriptions appeared only in expensive catalogues. Only people who could obtain a few books could run statistics. Now the data is open on the internet. Even students can download it and count the signs[3].

When data is open, verification becomes faster. If someone proposes a new reading, others can check it using the same data[6]. During the commotion of the summer of 2026, the verification process worked in just this way[4]. The side making the claim and the side reviewing it were looking at the same data[5].

Because the data is public, the barrier to research has also fallen. Even people outside universities can try sign statistics. When new eyes increase, rules that were missed can sometimes be found. But unverified claims also increase at the same time. Openness and verification must work as a pair[6].

A Script Still Unread

Artificial intelligence has done four things with Linear A. Machines have gathered scattered data in one place. They have turned sign shapes and strokes into numbers so they can be compared[3]. Computer calculation helped human reading narrow the values of the fraction signs. It reduced a problem with millions of candidates to a size that humans could review[1].

What machines have not done is also clear. What kind of language Minoan was remains unknown. A document that could serve as a Rosetta Stone is still underground, or does not exist[3]. Even if the sounds are transcribed, they do not reveal what the words refer to[6].

The case of the Cypro-Minoan script shows this difference well. A neural network produced a result that unsettled the existing classification[10]. But that does not mean it read the script. What the machine organized was the structure of the signs, not the meaning of the language.

Two paths seem to remain. One is the discovery of new artifacts. A single document written in two languages could change the whole field. The other is to refine the data in greater detail. If even one stroke is recorded, later research can begin from there[3].

Linear A is still an unread script. In front of it, the machine has not opened the door. It has drawn the shape of the lock. Opening the door depends on new artifacts and human judgment[6]. The clay tablets of the Aegean are still not speaking in their own words.

References

[1] Corazza, M., Ferrara, S., Montecchi, B., Tamburini, F., & Valério, M. (2021). The mathematical values of fraction signs in the Linear A script: A computational, statistical and typological approach. Journal of Archaeological Science, 125, 105214. https://doi.org/10.1016/j.jas.2020.105214

[2] Younger, J. G. (n.d.). Linear A texts and inscriptions in phonetic transcription. University of Kansas. http://people.ku.edu/~jyounger/LinearA/ (accessed September 3, 2026)

[3] Braović, M., Krstinić, D., Štula, M., & Ivanda, A. (2024). A systematic review of computational approaches to deciphering Bronze Age Aegean and Cypriot scripts. Computational Linguistics, 50(2), 725-779. https://doi.org/10.1162/coli_a_00514

[4] GreekReporter. (2026, 7월 30일). AI engineer claims to have deciphered Linear A. greekreporter.com. https://greekreporter.com/2026/07/30/ai-engineer-claim-decipher-linear-a/ (accessed September 3, 2026)

[5] GreekReporter. (2026, 8월 13일). AI engineer claims to have deciphered Linear A, identifies Minoan language as Semitic. greekreporter.com. https://greekreporter.com/2026/08/13/ai-engineer-claim-linear-a-minoans-semitic-language/ (accessed September 3, 2026)

[6] Adkins, J. (2026, 7월 27일). Cracking the code: can AI help us decipher ancient languages? The Conversation. https://theconversation.com/cracking-the-code-can-ai-help-us-decipher-ancient-languages-288238 (accessed September 3, 2026)

[7] GreekReporter. (2026, 8월 14일). Linear A inscription points to Minoan origins of Greek gods Dionysos and Demeter. greekreporter.com. https://greekreporter.com/2026/08/14/minoan-linear-a-inscription-greek-gods-dionysos-demeter/ (accessed September 3, 2026)

[8] Del Freo, M. (2022). Linear A. Mnamon: Ancient writing systems in the Mediterranean. Scuola Normale Superiore. https://doi.org/10.25429/sns.it/lettere/mnamon022 (accessed September 3, 2026)

[9] Salgarella, E., & Castellan, S. (2021). SigLA: The Signs of Linear A. A palaeographical database. In Grapholinguistics in the 21st Century 2020: Proceedings (Grapholinguistics and Its Applications, Vol. 5, pp. 945-962). Fluxus Editions. https://sigla.phis.me/ (accessed September 3, 2026)

[10] Corazza, M., Tamburini, F., Valerio, M., & Ferrara, S. (2022). Unsupervised deep learning supports reclassification of Bronze age cypriot writing system. PLOS ONE, 17(7), e0269544. https://doi.org/10.1371/journal.pone.0269544

Chapter 10. Hacking the Code of Solar Cycles and Animal Seals: The Sounds of Indus Script and Rongorongo

10.1. Case File

Four Signs Carved on a Seal

The Indus civilization lasted from around 2600 BCE to around 1900 BCE[1]. Its homeland was in present-day Pakistan and northwestern India. Judging by the size and number of its cities, it was one of the leading civilizations of the Bronze Age. The Bronze Age was an ancient period when people used bronze more than iron. The land covered by its cities and villages is known to have exceeded 1 million square kilometers[2].

Large excavations began at Harappa and Mohenjo-daro in the 1920s[1]. The person who oversaw this work was the British archaeologist John Marshall. He was the head of the Archaeological Survey of India. The Harappa site was excavated in 1921 by Daya Ram Sahni. The Mohenjo-daro site was handled in 1922 by Rakhaldas Banerji. Marshall brought these excavations together and announced the existence of a new civilization to the world[3].

Straight brick streets and drainage systems soon emerged from the ground. Many small square seals, smaller than a palm, also came out[1].

The appearance of the cities alone tells us a great deal. The roads ran straight like a grid, and drains connected to each house. Bricks were fired in standard sizes. Weights used for measuring also followed standards. This means it was a society that made plans and followed rules. Yet the writing they left behind cannot be read.

The seals were made of a soft stone called steatite. Steatite is a stone that can be marked even with a fingernail. Animal images are carved on the lower part of the seals. They include a one-horned animal, a humped bull, and an elephant-like figure. Above the image, signs are lined up in a row in an empty space[1].

A seal is an object used for stamping. When pressed into soft clay, its image and signs are transferred as they are. If clay is attached over the string that ties a package and then stamped, it becomes a seal impression. In Indus cities, clay sealings stamped with seals were also found[1].

Scholars call this group of signs the Indus script or Harappan script[1]. The name includes the word script, but it still cannot be read. We know neither its sounds nor its meanings. This has remained true for more than a hundred years since its discovery[2].

The statement that it cannot be read has two meanings. One is that we do not know the sounds of the signs. The other is that we do not know their meanings. At present, both remain unknown.

The problem is length. A single seal usually has only four or five signs[1]. Even the longest inscription on one face does not exceed 17 characters. Only when signs carved across three faces are counted together does the total reach 26 characters[4]. It is not enough to fill even one line of a classroom notebook.

Writing carved on seals or vessels like this is called an inscription. About 3,000 Indus inscriptions have been collected so far[5]. The number is not small, but each one is too short. It is hard to discover sentence structure from short texts alone. There is no chance to see how subjects and predicates are arranged.

Signs on Wooden Tablets Across the Sea

In the middle of the South Pacific lies an island called Rapa Nui. We know it better as Easter Island. It is more than 2,000 kilometers from the nearest inhabited land. It is also known for the large stone human figures called moai. Wooden tablets marked with unknown signs were also found on this island. Today these records are called rongorongo[6].

The signs are not brush-written letters. They were carved into wood with sharp pieces of obsidian or tools made from bone. The carved shapes resemble people, birds, and fish[6]. Human figures with raised hands and feet appear often.

The reading direction is also unfamiliar. The reader starts at the lower left end of the tablet. After reading one line, the tablet is turned 180 degrees. The next line is read while the tablet is upside down. This method of turning the tablet after each line is called reverse boustrophedon[6]. In simple terms, it is a way of reading by flipping. Ancient Greece also had writing that changed direction line by line. In that case, the tablet was not turned. Only the direction of the letters changed.

This method tests researchers once more. They must decide which line is the first line before the order can be fixed[6]. If the order is set incorrectly, the sequence of signs is completely reversed. This is a point where humans must make a judgment before the data can be entered into a computer.

Only a small number of artifacts remain. There are 26 genuine pieces scattered among museums and private collectors around the world. All of them are outside Rapa Nui[6]. Even when all the signs carved on them are added together, they total about 15,000 characters[7]. That is less than the length of a single book.

There is a reason why so few artifacts remain. In the 19th century, the island's population fell sharply, and many wooden tablets disappeared. The surviving pieces were scattered among European museums and monasteries. Four of them are kept by a religious order in Rome[8].

When these wooden tablets were made was also long disputed. Some suspected that they were imitations made after Europeans arrived on the island. In 2024, Silvia Ferrara's research team measured the age of four wooden tablets in Rome[8]. They used a method that determines age by measuring the amount of carbon left in wood. This method is called radiocarbon dating.

Three of the four pieces were made from 19th-century wood. The remaining piece turned out to be from much older wood. It predates the 1720s, when European ships reached the island, by more than 200 years[8]. This increased the possibility that the people of Rapa Nui created the script on their own.

The researchers added one condition. The time when the tree was cut is not the same as the time when the signs were carved. The age of the wood only tells us that the carving was made later than that. Old wood may have been carved and used at a later date. It also remains true that only one tablet produced an early date[8].

Is This Even Writing?

Scholars have long debated both sign systems. The debate did not begin with decipherment. They first asked whether these signs truly recorded spoken language[1][9].

In 2004, Steve Farmer, Richard Sproat, and Michael Witzel published a paper together. They argued that the Indus signs were not writing. They based this claim on the fact that the records are too short and that no long document exists. In other words, they argued that the signs were closer to family marks, ownership marks for goods, or charms[9].

The opposing side proposed finding an answer through statistics. The idea was to count how often signs appear and what rules govern the signs before and after them[5]. Human language has statistical habits, like habits of speech. This approach measures whether such habits also appear in the signs.

The strength of this method is that it can be used without knowing the meaning. The signs only need to be converted into numbers, and their order can be counted. This is also something computers do well.

Three Walls Blocking Decipherment

Both sign systems remain undeciphered. Three walls block the way forward.

There is no comparison text. Egyptian hieroglyphs were solved thanks to the Rosetta Stone. The Rosetta Stone also contained the same content in Greek. The French scholar Champollion placed known writing and unknown writing side by side and matched them. A text that records the same content side by side in two languages is called a bilingual text. No such record has been found at Indus sites or on Rapa Nui[2][6].

We also do not know what language the signs recorded. The actual language recorded by the signs is called the underlying language. For the underlying language of the Indus script, the Dravidian-language theory and the Indo-Aryan-language theory oppose each other. Some argue that it belongs to no known language family[2].

The remaining wall is the amount of data. Today's artificial intelligence grows by consuming vast amounts of text. But here we have only a few thousand Indus seals and 26 wooden tablets[1][6]. When data is lacking, artificial intelligence produces plausible falsehoods. This phenomenon is called hallucination.

These three walls hold one another in place. If we knew the underlying language, we could try to match sound values. If we had a bilingual text, we could identify the underlying language. If there were more data, statistics could provide clues. Because none of the three has been solved, the others remain stuck as well.

On January 5, 2025, the Indian state of Tamil Nadu offered a prize. The announcement was made at the opening ceremony of an international conference in Chennai marking the 100th anniversary of the discovery of the Indus civilization[10]. The prize offers 1 million dollars to anyone who deciphers the Indus script well enough to be recognized by archaeological experts[10][11]. As of September 2026, no winner is known.

References

[1] Rao, R. P. N., Yadav, N., Vahia, M. N., Joglekar, H., Adhikari, R., & Mahadevan, I. (2009). Entropic evidence for linguistic structure in the Indus script. Science, 324(5931), 1165. https://doi.org/10.1126/science.1170391

[2] Khan, M. A., Khurshid, S. K., & Aslam, M. (2025). Modern algorithms for ancient scripts: A review of AI-based techniques in Indus civilization research. Cultural and Historical Heritage: Preservation, Presentation, Digitalization, 11(2), 11-21. https://doi.org/10.55630/KINJ.2025.110201

[3] Grand Valley State University. (n.d.). Archaeologist of the month: John Marshall. gvsu.edu. https://www.gvsu.edu/archaeology/module-spotlight-view.htm?entryId=A748C17B-A353-F70E-D7422D41F7889433 (accessed September 3, 2026)

[4] Farmer, S. (n.d.). Claims concerning the longest Indus "inscription". safarmer.com. https://safarmer.com/indus-longestinscription/ (accessed September 3, 2026)

[5] Yadav, N., Joglekar, H., Rao, R. P. N., Vahia, M. N., Adhikari, R., & Mahadevan, I. (2010). Statistical analysis of the Indus script using n-grams. PLoS ONE, 5(3), e9506. https://doi.org/10.1371/journal.pone.0009506

[6] Australian Museum. (2026, 5월 28일). Kohau Rongorongo: Talking tablet. australian.museum. https://australian.museum/learn/cultures/pasifika-collections/rapa-nui-collections/kohau-rongorongo-talking-tablet/ (accessed September 3, 2026)

[7] Ross, L. (2026, 3월 31일). Reading Rongorongo (Version 2). Zenodo. https://doi.org/10.5281/zenodo.19362491 (accessed September 3, 2026)

[8] Ferrara, S., Tassoni, L., Kromer, B., Wacker, L., Friedrich, M., Tonini, F., Lastilla, L., Ravanelli, R., & Talamo, S. (2024). The invention of writing on Rapa Nui (Easter Island): New radiocarbon dates on the Rongorongo script. Scientific Reports, 14, 2794. https://doi.org/10.1038/s41598-024-53063-7

[9] Farmer, S., Sproat, R., & Witzel, M. (2004). The collapse of the Indus-script thesis: The myth of a literate Harappan civilization. Electronic Journal of Vedic Studies, 11(2), 19-57. https://doi.org/10.11588/ejvs.2004.2.620

[10] The Tribune. (2025, 1월 6일). Stalin announces $1 mn prize for deciphering Indus Valley script. tribuneindia.com. https://www.tribuneindia.com/news/india/stalin-announces-1-mn-prize-for-deciphering-indus-valley-script/ (accessed September 3, 2026)

[11] Smithsonian Magazine. (2025, 1월 30일). Officials are offering $1 million to anyone who can decode this ancient script. smithsonianmag.com. https://www.smithsonianmag.com/smart-news/officials-are-offering-1-million-to-anyone-who-can-decode-this-ancient-script-180985922/ (accessed September 3, 2026)

10.2. The Forensic Science Team’s Pursuit: The Information Entropy Meter

Measuring the Fingerprints of Writing

Indus seals and Rapa Nui wooden tablets never encountered each other. Yet researchers in both places picked up the same tool. That is because the work they could do before rows of unknown signs was similar.

Faced with unreadable signs, researchers chose a different question. Instead of asking what they meant, they asked how they were arranged. Even without knowing the meaning, the order can be counted.

The tool used here is information entropy. Entropy is a value that measures how hard it is to know what will come next. If the value is high, the next item is hard to predict. If the value is low, the next item is almost fixed.

This concept came from communications engineering. It was made as a ruler for measuring how much information is carried when signals are sent through telephone lines. It calculates only the order, regardless of meaning. That is why it can be applied to unreadable writing.

It is easier to understand by thinking of classroom seats. If students sit by drawing lots, no one knows who will sit next to whom. Entropy is high then. If students sit in number order, the person who comes next is fixed. Entropy is low then.

Human language lies between these two extremes. Grammar binds the order, but people freely choose what they want to say. So the entropy of human language lies in the middle range. The research begins by using this middle range as a ruler.

We can see this in Korean as well. After hakgyo-e, meaning "to school," many words can follow. Kanda (goes), gatda (went), and danyeowatda (has been there) are all possible. But not just any word can come next. This state, where choices exist but are narrowed, is a feature of language.

Measuring the Indus Signs

In 2009, a research team including Rajesh Rao and Nisha Yadav applied this ruler to the Indus signs. They calculated the uncertainty of the next sign when the previous sign was known[1]. This value, which measures how predictable the next sign is when the previous sign is known, is called conditional entropy.

They also prepared comparison targets. On one side, they placed human languages such as English, Sanskrit, and Sumerian. On the other side, they placed things that were not human language. These included biological gene sequences, computer program code, and randomly generated strings of signs[1].

The result was that the Indus signs fell close to human languages. They had clearer rules than random sign strings and more freedom than computer code. The research team published this result in Science[1].

The same research team also presented a probability model. This method creates a table of the probability that the next sign will appear once the previous sign is fixed. A model that calculates the next item by looking only at the previous sign is called a Markov model. Looking at the table reveals which signs often come at the beginning of an inscription. It also reveals which signs often come at the end[2].

The same research team examined sign frequency the following year. They gathered about 3,000 inscriptions and selected 1,548 that were in good condition. A small number of signs accounted for most uses, while the remaining signs appeared rarely. This pattern is always seen in human language[3].

This distribution has a name. It is called Zipf's law. It is the rule that a few frequently used words fill most of a text. In Korean, particles such as eun, neun, i, and ga occupy the front ranks. By contrast, there are many words that appear only once. The Indus signs showed the same pattern[3].

Rules also appeared in the order of placement. Twenty-three signs accounted for 80 percent of the endings of inscriptions. By contrast, 82 signs were needed to account for 80 percent of the beginnings[3]. This means the signs that come at the end are far more limited. It is similar to how particles or endings gather at the end of Korean sentences.

The researchers also made a model that predicts the next sign from the previous sign[2]. This model restored worn and hard-to-read signs with about 75 percent probability[3]. Such a result can appear only when rules exist.

The word restore must not be misunderstood. What the model fills in is only the shape of the sign. We still do not know what sound that sign had. It fills one missing square in a puzzle. It does not read the picture.

Objections and Replies

This model also has limits. It predicts the next sign by looking only at the previous one or two signs[2]. Human sentences link words that are much farther apart. If the data is short, those distant links cannot be learned.

A rebuttal soon appeared against this result. Richard Sproat published a critique in 2010. He pointed out that similar entropy values alone cannot prove that the signs are writing. His point was that nonlinguistic sign systems can also produce similar values[4].

In the same year, Rao's team gave its answer in the same journal. They said they had not proved that the Indus signs were writing. They explained that they had only supported, with statistics, the possibility that they were writing[5].

In 2014, Sproat pushed the criticism one step further. He gathered several sets of sign data that were not human language. They included Mesopotamian deity symbols, North American totem poles, and star patterns painted on barn walls. He also included pictures used in weather forecasts. He compared these materials and the Indus signs by the same standard[6].

The decision procedure he made classified the Indus signs as non-linguistic. Sproat wrote that statistics alone were not enough. He argued that either a successful decipherment or separate archaeological evidence that the culture used writing was needed[6]. This debate is still not over.

A new analysis published in April 2026 faces the same question. Ashish Nair redesigned the test using 1,916 Indus inscriptions. He used a computer to create fake data that imitated non-linguistic sign systems. He placed the statistics of the real Indus data and the fake data side by side and compared them[7].

The Indus data did not match any fake dataset exactly. It lay between two fake datasets. No real non-linguistic sign system reproduced the Indus statistics as they were. Nair wrote that judgment on whether this is writing should be postponed. This paper is a preprint posted online before journal review[7].

Eyes that recognize signs

There is work to do before statistics. It is the work of transcribing faint marks carved in stone into signs. For a long time, this work depended on human eyes and hands.

Transcription is harder than it may seem. Seals have worn down over 4,000 years. Even the same sign differs slightly depending on the person who carved it. Scholars may disagree on whether to see a mark as two signs or one sign. If this judgment changes, the statistics that follow also change.

In 2017, Satish Palaniappan and Ronojoy Adhikari gave this work to a computer. They made a processing workflow in which seal photographs were entered and a list of signs came out[8].

In 2025, Vaishnavi Dixit and her team carried this workflow forward. They made one neural network to recognize signs in seal photographs and another neural network to recognize animal images. They divided the data into five parts and tested them in rotation. Both neural networks showed average accuracy of about 85 percent[9].

To reduce this problem, the criteria for judgment must be made public and compared with one another. In 2025, Muzaffar Ali Khan and his team organized that work. They gathered artificial intelligence studies on the Indus script published from 2009 to 2025 and sorted them by type. They also recorded mistakes that often recur[10].

In 2025, Atul Sharma and Shubhajit Roy Chowdhury released a tool. Its name is AI-EPIGRAPHY. When a user enters a sequence of signs, it returns interpretation candidates with high probability. It is a tool that combines frequency analysis and chain probability calculation[11].

In May 2026, an attempt to rebuild the data foundation itself was announced. Beáta Megyesi and her team built a database that brings together rare and unreadable scripts. It placed high-resolution photographs, transcribed signs, annotations, and source information in one framework. The judgment was that gathering scattered materials is a task that comes before decipherment[12].

References

[1] Rao, R. P. N., Yadav, N., Vahia, M. N., Joglekar, H., Adhikari, R., & Mahadevan, I. (2009). Entropic evidence for linguistic structure in the Indus script. Science, 324(5931), 1165. https://doi.org/10.1126/science.1170391

[2] Rao, R. P. N., Yadav, N., Vahia, M. N., Joglekar, H., Adhikari, R., & Mahadevan, I. (2009). A Markov model of the Indus script. Proceedings of the National Academy of Sciences, 106(33), 13685-13690. https://doi.org/10.1073/pnas.0906237106

[3] Yadav, N., Joglekar, H., Rao, R. P. N., Vahia, M. N., Adhikari, R., & Mahadevan, I. (2010). Statistical analysis of the Indus script using n-grams. PLoS ONE, 5(3), e9506. https://doi.org/10.1371/journal.pone.0009506

[4] Sproat, R. (2010). Ancient symbols, computational linguistics, and the reviewing practices of the general science journals. Computational Linguistics, 36(3), 585-594. https://doi.org/10.1162/coli_a_00011

[5] Rao, R. P. N., Yadav, N., Vahia, M. N., Joglekar, H., Adhikari, R., & Mahadevan, I. (2010). Entropy, the Indus script, and language: A reply to R. Sproat. Computational Linguistics, 36(4), 795-805. https://doi.org/10.1162/coli_c_00030

[6] Sproat, R. (2014). A statistical comparison of written language and nonlinguistic symbol systems. Language, 90(2), 457-481. https://doi.org/10.1353/lan.2014.0031

[7] Nair, A. (2026, 4월 20일). How non-linguistic is the Indus sign system? A synthetic-baseline scorecard. arXiv:2604.17828. https://arxiv.org/abs/2604.17828 (accessed September 3, 2026)

[8] Palaniappan, S., & Adhikari, R. (2017). Deep learning the Indus script. arXiv:1702.00523. https://arxiv.org/abs/1702.00523

[9] Dixit, V., Hussain, N., Basak, S., Atturu, D., Mitra, D., & Bhattacharya, U. (2025). Deep learning in archiving Indus script and motif information. Journal of Computer Applications in Archaeology, 8(1), 156-169. https://doi.org/10.5334/jcaa.175

[10] Khan, M. A., Khurshid, S. K., & Aslam, M. (2025). Modern algorithms for ancient scripts: A review of AI-based techniques in Indus civilization research. Cultural and Historical Heritage: Preservation, Presentation, Digitalization, 11(2), 11-21. https://doi.org/10.55630/KINJ.2025.110201

[11] Sharma, A., & Roy Chowdhury, S. (2025). AI-EPIGRAPHY: An interactive tool for computational decipherment of the Indus Valley script. Proceedings of IndiaHCI '25. https://doi.org/10.1145/3768633.3770145

[12] Megyesi, B., Rattenborg, R., Lang, B., Waldispuhl, M., & Heder, M. (2026). Building a corpus and database for rare and undeciphered scripts. Proceedings of the Fourth Workshop on Language Technologies for Historical and Ancient Languages (LT4HALA 2026) @ LREC 2026, 184-196. ELRA Language Resources Association. https://aclanthology.org/2026.lt4hala-1.17/ (accessed September 3, 2026)

10.3. The Truth Revealed

A wooden tablet that counts the moon

Among the wooden tablets of Rapa Nui, one is called Mamari. This tablet has a section that looks different from other parts. Crescent-shaped signs continue in a row[1].

In the mid-20th century, the German scholar Thomas Barthel focused on this section. He saw it as a record of lunar dates[1]. In 1990, Jacques Guy reorganized the material and reached the same conclusion[2].

In 2011, Paul Horley strengthened this interpretation. The calendar on the Mamari tablet consists of thirty nights. This number matches the actual cycle of the moon. It also matches calendars from several Polynesian islands[1].

Horley also compared it with calendars from Tahiti and New Zealand. The names attached to each night were similar from island to island. The Rapa Nui calendar was about three days behind those of other islands. This seems to be a difference that arose while the island was far apart from them[1].

The arrangement of the signs also followed rules. Other signs appear between the crescent signs and divide them into eight groups. The last two crescents indicate nights when the moon is not visible[1].

The same arrangement appears outside the wooden tablet as well. Among the island's rock drawings, there is one where a turtle pattern and crescents are carved together. There too, the last two crescents are placed inside the turtle's shell. This is the same way the tablet sets apart moonless nights[1]. The wooden tablet and the rock shared the same method of counting.

This interpretation is widely accepted in rongorongo studies. Even so, what has been read is the order of the dates. We do not know what the Rapa Nui people called those nights.

The answers from artificial intelligence and their weight

Counting the moon and counting the sun are connected. A solar eclipse occurs when the sun and the moon overlap. So the claim that a calendar has been read soon spreads into the claim that the whole sky has been read.

In March 2026, Logan Ross released a preprint. He ran a language model on a personal computer and analyzed rongorongo. He made a dictionary of 89 signs and claimed that it covered nearly 97 percent of the remaining records. He also argued that the Mamari tablet was a table for predicting solar and lunar eclipses[3].

This claim has not yet been verified by the academic community. Ross's paper is not a paper that has passed journal review. It is a preprint first posted in an online repository[3]. Anyone can post a preprint. Whether the content is correct must be examined later by other researchers.

When dealing with such claims, the method of testing must be examined together. The remaining rongorongo records consist of 26 items, about 15,000 characters[3]. Where there is only this much material, results can easily fit by chance. If sound values are freely assigned to 89 signs, any sentence can be made. So scholars first ask whether another person can obtain the same result by the same method.

A real decipherment must pass a test. It must take another wooden tablet that was not used to create the decipherment and read it. Meaningful sentences must appear there too. Linear B, which recorded Mycenaean Greek, is an example that passed that test[4]. Rongorongo and the Indus script have not yet crossed that threshold.

What machines can and cannot do

What artificial intelligence has achieved so far with these two scripts is clear.

Machines recognize signs from faint marks[5]. Work that would take months for a person to compare by eye can be finished in hours. They also restore worn and hard-to-read signs from the surrounding context[6]. They gather photographs and records of scattered artifacts in one place and make them searchable[7].

The work of gathering data is not very visible. But without this work, there is no calculation that follows. If the brightness and angle of photographs differ from one another, signs cannot be compared. If transcription rules differ from researcher to researcher, the statistics also diverge. The database published in 2026 is an attempt to bring those standards into one form[7].

There are achievements on the statistical side as well. Several studies have confirmed that the Indus signs are not random arrangements[6][8]. The conclusion that there are rules in the arrangement has been reproduced repeatedly. But this is not the same as concluding that they are writing. Saying that there are rules is different from saying that speech was written down.

But what machines cannot do remains larger. The sound value of each Indus sign is still unknown[9]. The sound values of rongorongo signs are also unknown[3]. In both scripts, there is still no word whose meaning has been confirmed. It has not even been decided whether what is written on Indus seals is a person's name or the name of an object[9].

The reason lies in the nature of the data. If there is a parallel text, decipherment becomes much easier. Neither Indus sites nor Rapa Nui Island has such a text[9]. There are examples solved without a parallel text. Linear B was solved that way. But in that case there were clues for guessing the underlying language, and there was enough material[4]. For the Indus script, even the underlying language has not been settled[9]. For rongorongo, too little writing remains[3].

The analysis from April 2026 showed this situation again in numbers. The 1,916 Indus inscriptions matched neither the statistics of human language nor those of non-linguistic signs exactly[10]. This means that a question that has continued for more than a century remains open.

Standing on the side of waiting

Research on unread scripts requires two attitudes. One is the attitude of pushing new tools to the end. The other is the attitude of doubting the results that come out.

Artificial intelligence does the first task well. It counts, compares, and lines up candidates. The second task belongs to people. Among the candidates, we must judge what was plausible in that age. That judgment includes the location of ruins, objects found in graves, and oral traditions that remain.

The interpretation of the Rapa Nui calendar is an example. Counting crescent signs is calculation. Comparing those thirty nights with calendars from several Polynesian islands is historical verification[1]. When the two tasks worked together, the interpretation became persuasive.

The animal images on Indus seals are also such data. We do not know what the one-horned animal meant. The thirty nights on the Mamari tablet are the same. The order has been identified, but we do not know what people did on those nights.

There is value in saying that we do not know. We must draw a line between what has been found and where speculation begins so that the next study can begin. If the line is blurred, plausible stories pass themselves off as facts. In research on unread scripts, drawing this line is as important as producing results.

The two scripts discussed in this chapter never met each other. They are separated by more than 3,000 years, and their homes lie on opposite sides of the earth. Yet the difficulties researchers face are almost the same. Short materials, no parallel text, and an unknown underlying language. This is why the same tools are used in both places.

The 1 million dollars offered by Tamil Nadu also reflects this situation[11]. More than a year has passed since the prize was announced, but no winner has appeared. Calculators have become faster, but the key has not yet appeared.

References

[1] Horley, P. (2011). Lunar calendar in rongorongo texts and rock art of Easter Island. Journal de la Societe des Oceanistes, 132, 17-38. https://doi.org/10.4000/jso.6314

[2] Guy, J. B. M. (1990). The lunar calendar of Tablet Mamari. Journal de la Societe des Oceanistes, 91(2), 135-149. https://doi.org/10.3406/jso.1990.2882

[3] Ross, L. (2026, 3월 31일). Reading Rongorongo (Version 2). Zenodo. https://doi.org/10.5281/zenodo.19362491 (accessed September 3, 2026)

[4] Chadwick, J. (1958). The decipherment of Linear B. Cambridge University Press.

[5] Palaniappan, S., & Adhikari, R. (2017). Deep learning the Indus script. arXiv:1702.00523. https://arxiv.org/abs/1702.00523

[6] Yadav, N., Joglekar, H., Rao, R. P. N., Vahia, M. N., Adhikari, R., & Mahadevan, I. (2010). Statistical analysis of the Indus script using n-grams. PLoS ONE, 5(3), e9506. https://doi.org/10.1371/journal.pone.0009506

[7] Megyesi, B., Rattenborg, R., Lang, B., Waldispuhl, M., & Heder, M. (2026). Building a corpus and database for rare and undeciphered scripts. Proceedings of the Fourth Workshop on Language Technologies for Historical and Ancient Languages (LT4HALA 2026) @ LREC 2026, 184-196. ELRA Language Resources Association. https://aclanthology.org/2026.lt4hala-1.17/ (accessed September 3, 2026)

[8] Rao, R. P. N., Yadav, N., Vahia, M. N., Joglekar, H., Adhikari, R., & Mahadevan, I. (2009). Entropic evidence for linguistic structure in the Indus script. Science, 324(5931), 1165. https://doi.org/10.1126/science.1170391

[9] Khan, M. A., Khurshid, S. K., & Aslam, M. (2025). Modern algorithms for ancient scripts: A review of AI-based techniques in Indus civilization research. Cultural and Historical Heritage: Preservation, Presentation, Digitalization, 11(2), 11-21. https://doi.org/10.55630/KINJ.2025.110201

[10] Nair, A. (2026, 4월 20일). How non-linguistic is the Indus sign system? A synthetic-baseline scorecard. arXiv:2604.17828. https://arxiv.org/abs/2604.17828 (accessed September 3, 2026)

[11] The Tribune. (2025, 1월 6일). Stalin announces $1 mn prize for deciphering Indus Valley script. tribuneindia.com. https://www.tribuneindia.com/news/india/stalin-announces-1-mn-prize-for-deciphering-indus-valley-script/ (accessed September 3, 2026)

Chapter 11. AI Reads Broken Greek Inscriptions Again: Foundation Models and Ithaca

11.1. Case File

Records of a Country Carved in Stone

Ancient Greeks carved important matters into stone. When they made a new law, they carved its contents on a stone pillar and set it up in the square. Tax ledgers were also left in stone. Lists of objects dedicated to temples were also left in stone. These writings carved on stone or bronze are called inscriptions. The study that collects and reads inscriptions is epigraphy. It means the study of writings carved on stone or metal[4].

Inscriptions are usually records left directly by people of that time. They rarely show traces of later copyists changing the text. So they are valuable sources for historians. But they are not always in their original form. Some stones had letters erased or recarved later. There are also cases where the stone was ground down and used again. Inscriptions tell us when Athens formed ties with which city[4].

Inscriptions were made for many people to see. They were set up in squares or in front of temples so passersby could read them. Their use was similar to notices posted in a school hallway. So inscriptions clearly include the names of officials of the day and the date. Historians use these names and dates as clues[4].

Erased Letters, Lacunae

The problem is time. Stone monuments stood before rain and wind for thousands of years. The stone surface wore away little by little, and the grooves of the letters became shallow. Stones buried in the ground were eaten away by moisture and microorganisms. Sometimes a stone broke, leaving only one side[1].

As a result, inscriptions develop blanks where letters have disappeared. Scholars call these blanks lacunae. In Latin, they are lacunae[1]. A lacuna opens a hole in the middle of a sentence. If a person's name disappears, we cannot know who did the act. If a number disappears, we cannot know how much was paid. If a city name disappears, we cannot know which country the agreement was made with.

Scholars have long tried to fill these blanks. They searched for other inscriptions from a similar period with similar content and compared them. Greek inscriptions had fixed sentence patterns. Documents announcing laws arranged sentences in almost the same order. So if the words before and after remain, scholars can guess the missing word[1].

Stone monuments are too heavy to move to a laboratory. So scholars pressed wet paper against the stone. When the paper dried, the grooves of the letters rose up in relief. This paper copy is called a squeeze. Scholars brought the squeeze back and read the letters at their desks. Old printed catalogs often included only transcribed letters instead of photographs. So later readers had to read only the letters without seeing the shape of the stone[4].

Guessing follows rules. First, scholars count how many letters the blank had. They can estimate it by measuring the spacing between the letters left on the stone. Then they consider only words with the right number of letters. Letters filled in this way must be written inside square brackets. A small dot is placed under a letter whose form is damaged and hard to read. This convention is called the Leiden system. It is a rule made to distinguish letters left on the stone from human conjecture[1, 4].

The Dispute over Three-Stroke Sigma

To read an inscription, scholars must also know when it was carved. An old method is to estimate the date from the shape of the letters. This is because letter shapes changed little by little over time[2].

A common example in Athenian inscriptions is sigma. Sigma is one letter of the Greek alphabet. In an early period, sigma was carved with three strokes. Later, sigma with four strokes became widely used. For a long time, scholars used three-stroke sigma as a standard. If such a letter appeared, they saw the inscription as earlier than around 446 BCE[2].

But some scholars began to question this standard. The speed at which letter forms changed differed by region. Some regions held on to old letter forms for a long time. Even within one city, each craftsman had different carving habits[2].

So disputes continued over the dates of treaty inscriptions that Athens made with other cities. One side said they were earlier than around 446 BCE. The other side said they were carved in the 420s BCE[2]. The 420s BCE were the height of the Peloponnesian War. This war began in 431 BCE. If the date shifts by several decades, the outline of history also changes. The point when Athens pressed allied cities by force moves as a whole.

Filling blanks carries risks. A word supplied by one scholar may be read like the original text by the next generation. A single printed edition that omits square brackets can turn conjecture into fact[1]. For this reason, epigraphers always marked their conjectures.

Scholars must also find out where an inscription originally stood. Greek cities each used different dialects. The same word could be spelled differently from city to city[2]. Scholars inferred the findspot by looking at these spelling differences and the names of offices[4]. The findspot is the place where the stone came out of the ground.

Human Hands Cannot Read Them All

The inscriptions to be read keep piling up. The Packard Humanities Institute has collected 178,551 Greek inscriptions[2]. Latin inscription materials from the Roman period also exceed 176,000 items[3]. By contrast, only a small number of scholars in the world can read inscriptions[4].

Even refining the numbers is not easy. When the data were prepared so machines could read them, 78,608 items remained. More than half were removed because they were too short or in too poor a condition. Choosing the data to feed the machine was itself a major task[2].

Filling blanks is difficult even for humans. In one test, trained epigraphers were given inscriptions with erased letters and asked to fill them in. The rate at which the first answer given by humans was correct was 25 percent[2]. This means they were wrong three times out of four. In an earlier experiment, the error rate for letters filled in by humans was 57.3 percent[1].

The situation for papyrus is worse. Papyrus is the name of a plant that grows by the waters of Egypt. Ancient people sliced the inside of that plant thinly. They laid the slices crosswise and lengthwise, beat them, pressed them, and dried them. The recording material made in that way is also called papyrus. Hundreds of thousands of sheets have come out of Egypt's dry sand. Greek papyri that no one has yet read are counted at one million items[5]. These papers contain tax receipts, letters, and contracts. They are records that preserve the daily lives of ancient people.

Reading does not end with opening a box. The paper is torn and the ink is erased. Conservation and the matching of scattered fragments must come first. Only after that can the letters be deciphered. There are not enough people to do all this work[4].

The way the materials were read was also a problem. For a long time, inscription data accumulated only as lists of letters in books. Photographs of the stone and information about the material were not included[4]. This means that the only information that could be put into a machine was the letters.

Excavation still continues. Every year, new stones and new papyri come out of the ground. The materials to be read increase, but the number of readers stays the same[4]. This gap widened year by year.

So many records remained unread in museum boxes. There were stones and papers, but not enough people to revive the sentences on them. Researchers brought in machines to break through this blockage[4].

References

[1] Assael, Y., Sommerschield, T., & Prag, J. (2019). Restoring ancient text using deep learning: a case study on Greek epigraphy. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP) (pp. 6368-6375). Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-1668

[2] Assael, Y., Sommerschield, T., Shillingford, B., Bordbar, M., Pavlopoulos, J., Chatzipanagiotou, M., Androutsopoulos, I., Prag, J., & de Freitas, N. (2022). Restoring and attributing ancient texts using deep neural networks. Nature, 603(7900), 280-283. https://doi.org/10.1038/s41586-022-04448-z

[3] Assael, Y., Sommerschield, T., Cooley, A., Pavlopoulos, J., Shillingford, B., Herms, B., Suresh, P., Maynard, B., Grayston, J., Wulgaert, R., Prag, J., Mullen, A., & Mohamed, S. (2025). Contextualizing ancient texts with generative neural networks. Nature, 645(8079), 141-147. https://doi.org/10.1038/s41586-025-09292-5

[4] Sommerschield, T., Assael, Y., Pavlopoulos, J., Stefanak, V., Senior, A., Dyer, C., Bodel, J., Prag, J., Androutsopoulos, I., & de Freitas, N. (2023). Machine learning for ancient languages: A survey. Computational Linguistics, 49(3), 703-747. https://doi.org/10.1162/coli_a_00481

[5] Reply. (2026, 3월 19일). Austrian Academy of Sciences is developing the Ancient Greek AI "Apollo" with Mistral AI and Reply. reply.com. https://www.reply.com/en/newsroom/news/austrian-academy-of-sciences-is-developing-the-ancient-greek-ai-apollo-with-mistral-ai-and-reply (accessed September 3, 2026)

11.2. The Forensic Team's Pursuit: Ithaca and Apollo, Ancient Text Autocomplete Systems

A Machine That Learns from Blank-Filling Problems

The way AI reads old writing resembles a fill-in-the-blank question in a Korean language exam. A machine is shown many complete inscriptions. Then several letters are hidden, and the machine is asked to guess what should go there. If the answer is wrong, the numbers inside the machine are adjusted little by little[4]. If this process is repeated millions of times, the machine learns the habits of Greek sentences.

A machine trained in this way does not memorize words. It remembers, as probabilities, which words tend to come after which other words[4]. So even in an inscription it has never seen before, it can produce candidates that fit the blank.

From Pythia to Ithaca

In 2019, researchers introduced a model called Pythia. It was made by Yannis Assael, Thea Sommerschield, and Jonathan Prag. Pythia is a neural network that predicts the letters of erased Greek inscriptions[1]. A neural network is a computational structure that learns rules by looking at data. It was modeled on the connections between nerve cells in the human brain.

Pythia reads at both the character level and the word level. It needs two eyes to handle material in which half a word has been broken away. When the error rate for letters filled in by humans was 57.3

percent, Pythia recorded 30.1 percent. The correct answer appeared among Pythia's 20 candidates in 73.5 percent of cases[1].

In 2022, Ithaca appeared. Google DeepMind and several universities created it together. The University of Oxford, Ca' Foscari University of Venice, and the Athens University of Economics and Business participated. The research results were published in the journal Nature in March 2022[2]. Its name comes from Ithaca, the island in the epic Odyssey. It carries the meaning of returning home after long wandering.

How Ithaca Judges

Ithaca uses a structure called a transformer. A transformer is a neural network that examines how strongly words in a sentence relate to one another. Places where letters have been erased are entered as symbols meaning blanks. The model calculates the probability of the letters that should replace those symbols[2].

Ithaca performs three tasks. It fills erased letters, guesses where an inscription originally stood, and estimates when it was carved. It does not learn these three tasks separately. It learns them together. Dialect and spelling habits are clues to place. Word usage and sentence patterns are clues to time. When one neural network handles these clues together, they help one another[2].

Ithaca does not look at photographs of stones. It reads only transcribed letters and makes judgments[2]. This path differs from the old method of dating inscriptions by the shape of strokes.

Its way of giving dates is also different. Ithaca does not point to one exact year. It divides the period from 800 BCE to 800 CE into ten-year units. Then it assigns probabilities to each interval and shows them as a graph[2]. Scholars look at the peaks of the graph and judge which period is most likely.

The data used for training came from the inscription collection of the Packard Humanities Institute. The 78,608 items seen above are that data. Even counting only words that appear more than ten times, there were 35,884 words[2].

When the research team built Ithaca, they had three goals. It should work with scholars, help judgment, and show the judgment process. The goal was not a machine that replaces humans. They aimed for a tool that enters the scholar's way of working[2].

Ithaca does not give only one answer. It shows several candidates for a blank, together with their probabilities. It also marks on the screen which letters led it to that judgment. Scholars can see with their own eyes what the machine relied on for its answer. The final judgment remains with the human. The research team released Ithaca's code and data[2].

The Meaning of Twenty Candidates

The machine does not produce a single answer. It ranks candidates in order of high probability. The rate at which the first candidate is correct is called top-1 accuracy. The rate at which the correct answer appears among the first twenty candidates is top-20 accuracy[1].

This distinction matters. What scholars need is not one correct answer. They need a list of candidates worth reviewing. If the correct answer is among the twenty, the scholar can choose from within that list[1]. Sometimes they also discover in that list a word they had not thought of themselves[2].

Aeneas Crosses into Latin

The same research team expanded its direction to Latin inscriptions. In July 2025, they announced a model called Aeneas in Nature[3]. Its name comes from Aeneas, the hero of the story of Rome's founding[5].

Aeneas learned from more than 176,000 Latin inscriptions. The data combined three inscription databases from Rome, Heidelberg, and Clauss-Slaby. In places where up to ten letters were erased, the correct answer appeared among the 20 candidates in 73 percent of cases. Even when the number of erased letters was unknown, it recorded 58 percent[3]. Answering without knowing the length of the erasure is a difficult problem even for humans.

It also guessed which Roman province an inscription came from. A province was a local territory ruled by Rome. Among 62 choices, it was correct 72 percent of the time[5]. For the carving date, it narrowed the estimate to within an average of 13 years of the range set by historians[3, 5].

Aeneas has a new function. It can make judgments using a photograph of the inscription as well[3]. This means it looks not only at transcribed letters but also at the shape of the stone. It also finds other inscriptions with similar sentences and shows them side by side. It brings up in seconds similar cases that used to take scholars days to find[5].

The researchers tested this tool on a record from the reign of the Roman emperor Augustus. Aeneas gave two possible date ranges for the inscription. One was between 10 BCE and 1 BCE. The other was between 10 CE and 20 CE, and it had the higher probability[5]. The two views that had divided scholars appeared directly on the probability graph.

The researchers had 23 historians try this function. The scholars said they gained new research ideas from nine out of ten inscriptions[3, 5]. The model was released for anyone to use[5].

Apollo, Learning Greek as a Whole

In March 2026, the Austrian Academy of Sciences announced a new model. Its name is Apollo. It is being built together with the AI company Mistral AI and the information technology company Sail Reply[7]. If Ithaca is a tool for handling one inscription, Apollo aims to be a model that learns the whole Greek language[6].

This kind of model is called a foundation model. It means a basic model trained broadly so it can be used for many tasks. It is not a tool made for only one job[4]. The same base is used for restoring inscriptions, searching texts, and reading handwriting[6].

Apollo learns from about 600 million words of Greek texts. Tens of thousands of already published inscriptions and papyri are added to this[6, 7]. The researchers said this is one of the largest Greek-language collections assembled so far[7].

The Austrian National Library holds many papyri. Tens of thousands are stored in that library alone[7]. Documents that have not yet been read are mixed in among them.

The research is led by Anna Dolganov of the Austrian Archaeological Institute. She said there are one million Greek papyri in the world that no one has yet read[7]. Apollo is meant to restore letters and find content in such materials. Reading handwritten text is also one of its goals[6, 7]. The researchers said this work could be finished not in years, but in hours[7].

References

[1] Assael, Y., Sommerschield, T., & Prag, J. (2019). Restoring ancient text using deep learning: a case study on Greek epigraphy. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP) (pp. 6368-6375). Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-1668

[2] Assael, Y., Sommerschield, T., Shillingford, B., Bordbar, M., Pavlopoulos, J., Chatzipanagiotou, M., Androutsopoulos, I., Prag, J., & de Freitas, N. (2022). Restoring and attributing ancient texts using deep neural networks. Nature, 603(7900), 280-283. https://doi.org/10.1038/s41586-022-04448-z

[3] Assael, Y., Sommerschield, T., Cooley, A., Pavlopoulos, J., Shillingford, B., Herms, B., Suresh, P., Maynard, B., Grayston, J., Wulgaert, R., Prag, J., Mullen, A., & Mohamed, S. (2025). Contextualizing ancient texts with generative neural networks. Nature, 645(8079), 141-147. https://doi.org/10.1038/s41586-025-09292-5

[4] Sommerschield, T., Assael, Y., Pavlopoulos, J., Stefanak, V., Senior, A., Dyer, C., Bodel, J., Prag, J., Androutsopoulos, I., & de Freitas, N. (2023). Machine learning for ancient languages: A survey. Computational Linguistics, 49(3), 703-747. https://doi.org/10.1162/coli_a_00481

[5] Google DeepMind. (2025, 7월 23일). Aeneas transforms how historians connect the past. deepmind.google. https://deepmind.google/blog/aeneas-transforms-how-historians-connect-the-past/ (accessed September 3, 2026)

[6] Mistral AI. (2026). Austrian Academy of Sciences. mistral.ai. https://mistral.ai/customers/austrian-academy-of-sciences/ (accessed September 3, 2026)

[7] Reply. (2026, 3월 19일). Austrian Academy of Sciences is developing the Ancient Greek AI "Apollo" with Mistral AI and Reply. reply.com. https://www.reply.com/en/newsroom/news/austrian-academy-of-sciences-is-developing-the-ancient-greek-ai-apollo-with-mistral-ai-and-reply (accessed September 3, 2026)

11.3. The Truth Revealed

62 Percent and 72 Percent

The researchers measured Ithaca's ability in numbers. They gave it inscriptions with letters deliberately hidden and asked it to fill them in. They gave the same test to trained epigraphers[1].

When people solved the test alone, their first answer was correct 25 percent of the time. When Ithaca solved it alone, the figure was 62 percent. That is more than twice as high[1].

The test had a third condition. People solved the problem while looking at Ithaca's list of candidates. Then accuracy rose to 72 percent[1]. It was higher than people alone and also higher than the machine alone.

This number is the key to the study. The machine calculates probabilities from its memory of reading 78,608 inscriptions. People know the historical setting of each inscription. When the machine lists twenty candidates, people filter them with historical knowledge. Accuracy was highest when the two were combined[1, 2].

The researchers also tested its ability to identify place and time. Ithaca learned the Greek world divided into 84 regions. It identified the original home of inscriptions with unknown findspots with 71 percent accuracy. For the date of carving, its estimate was on average 29.3 years away from the range set by historians. Earlier methods based on judging letter shapes by eye could be off by as much as one hundred years[1].

The Answer 421 BCE

Ithaca also gave an answer to an old dispute. It was the dating problem of the Athenian treaty inscriptions discussed earlier. These inscriptions had been regarded as earlier than around 446 BCE because of the three-stroke sigma[1].

The researchers entered these inscriptions into Ithaca. Ithaca judges by looking not at letter shapes, but at the sentences themselves. It checks which words are used together with which other words. It also checks which official titles appear with which expressions[1].

The average date given by Ithaca was 421 BCE. That is one generation later than the old standard. This result pointed in the same direction as other new evidence. It differed by an average of 5 years from the dates recently reassigned by scholars. Dates assigned by the old standard were off by as much as 27 years[1].

The machine did not rewrite history by itself. It gave more weight to one side of a long-running dispute[1, 2]. Scholars were led to reexamine the habit of fixing dates from a single letter form. One tool changed the direction of a century-old debate.

Plausible Sentences Invented by Machines

Despite these results, problems remain. One is that the training data is biased. Greek inscription data contains many Athenian examples. The periods are also concentrated in the Classical age. Inscriptions from peripheral regions with less data are therefore less accurate[2].

There is also the problem that it is hard to know why the machine gave a certain answer. A neural network contains millions of numbers. It can show which letters it focused on more, but it cannot explain in words why it chose that answer[2].

Problems also arise when the blank is long. Filling in a few letters works well, but places where several whole words have disappeared are difficult. In the Aeneas test, when the deleted length was unknown, performance fell from 73 percent to 58 percent[3]. In such cases, the model can create sentences that are grammatical and plausible but never actually existed. Such invented answers are called hallucinations. This means wrong answers that the machine has fabricated as if they were facts. Hallucinations are dangerous when dealing with ancient records. History that never existed can harden into fact[2].

A study released in August 2026 addressed this limit directly. It tested the restoration of damaged ancient documents using several models and several conditions. The researchers concluded that restoration cannot be left entirely to automation. Performance differed greatly depending on which part of a document had to be filled in. Documents with repeated fixed phrases were restored well. Sentences written differently by each person were difficult. Results also differed depending on whether the erased length was known. The researchers said machines should be made as tools that help palaeographers[4].

Researchers set rules to reduce this risk. Letters filled in by the machine are also placed in square brackets. Scholars filter out low-probability candidates. The notation conventions followed by earlier epigraphers are applied to the machine as well[2].

So the answers produced by Ithaca are not historical facts. They are lists of hypotheses to examine. Only when scholars check them against other materials do they become historical sources[2]. A historical source is evidence used to uncover history.

What Happened in Seven Years

The pace of change in this field can be compared in numbers. In 2019, Pythia lowered the human error rate from 57.3 percent to 30.1 percent[5]. In 2022, Ithaca raised first-answer accuracy to 62 percent[1]. In 2025, Aeneas moved into Latin and narrowed the average dating error to 13 years[3]. In 2026, Apollo began learning Greek as a whole[6].

In seven years, the target languages increased and accuracy improved. The materials handled also expanded from letters to photographs[3]. One thing did not change during that time. The final judgment is made by people[2].

Photographs, Context, and the Next Step

Research continues to expand. One paper published in 2026 retrained a widely used large language model on Greek materials. The experiment used Meta's released Llama 3.1 model. In papyrus restoration, the first answer was correct 73.5 percent of the time. In inscription restoration, the figure was 63.7 percent. The correct answer appeared among the top 20 candidates in 86.0 percent and 83.0 percent of cases, respectively. This means that even a model not built specifically for inscriptions can do this work[7].

Research has also emerged on restoring letter images. MESA, released in April 2026, is a method for filling in erased strokes in photographs of inscriptions. It uses well-preserved letters on the same stone as examples and restores the patterns in worn areas[8]. Models that fill in sentences and models that fill in photographs are developing side by side.

Lost tragedies are also truly returning. In 2022, a sheet of papyrus was found at the Philadelphia site in Egypt. It appears to have been written in the 3rd century CE. In 2024, scholars read and published its contents. They were 97 lines of dialogue from Euripides' tragedies Polyidus and Ino. Until then, the two works had been known only through short quotations and plot summaries[9].

This restoration was done by human hands and eyes. It was not done by a machine. There are one million more such sheets of paper left in the world. Human eyes alone cannot read all of this pile. That is exactly where models such as Apollo are aimed[6].

As the tools were released, research methods also changed. Ithaca and Aeneas have made their code and data public[1, 10]. Anyone can try them on the web, even outside a university laboratory. Upload one photograph of an inscription, and a list of candidates, a map, and a date graph appear[10]. In the past, this comparison work could take several years[3].

Much work still remains. One million Greek papyri alone remain unread[6]. Inscription data is uneven by region and period[2]. Long gaps that machines do not handle well also remain[3]. These problems must be reduced one by one with better data and more verification[4].

The machine proposes candidates, and people choose. The method that raised Ithaca from 62 percent to 72 percent is the answer[1]. For silent stones and papers to speak again, both sides must be present.

References

[1] Assael, Y., Sommerschield, T., Shillingford, B., Bordbar, M., Pavlopoulos, J., Chatzipanagiotou, M., Androutsopoulos, I., Prag, J., & de Freitas, N. (2022). Restoring and attributing ancient texts using deep neural networks. Nature, 603(7900), 280-283. https://doi.org/10.1038/s41586-022-04448-z

[2] Sommerschield, T., Assael, Y., Pavlopoulos, J., Stefanak, V., Senior, A., Dyer, C., Bodel, J., Prag, J., Androutsopoulos, I., & de Freitas, N. (2023). Machine learning for ancient languages: A survey. Computational Linguistics, 49(3), 703-747. https://doi.org/10.1162/coli_a_00481

[3] Assael, Y., Sommerschield, T., Cooley, A., Pavlopoulos, J., Shillingford, B., Herms, B., Suresh, P., Maynard, B., Grayston, J., Wulgaert, R., Prag, J., Mullen, A., & Mohamed, S. (2025). Contextualizing ancient texts with generative neural networks. Nature, 645(8079), 141-147. https://doi.org/10.1038/s41586-025-09292-5

[4] Zhang, S., Caraffa, E., Zuffrano, A., Modesti, M., & Colavizza, G. (2026, 8월 28일). Text restoration of ancient documents with language models (arXiv:2608.28170). arXiv. https://arxiv.org/abs/2608.28170 (accessed September 3, 2026)

[5] Assael, Y., Sommerschield, T., & Prag, J. (2019). Restoring ancient text using deep learning: a case study on Greek epigraphy. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP) (pp. 6368-6375). Association for Computational Linguistics. https://doi.org/10.18653/v1/D19-1668

[6] Reply. (2026, 3월 19일). Austrian Academy of Sciences is developing the Ancient Greek AI "Apollo" with Mistral AI and Reply. reply.com. https://www.reply.com/en/newsroom/news/austrian-academy-of-sciences-is-developing-the-ancient-greek-ai-apollo-with-mistral-ai-and-reply (accessed September 3, 2026)

[7] Cullhed, E. (2026). Instruction-tuning pretrained causal language models to restore ancient Greek papyri and inscriptions. Digital Scholarship in the Humanities, 41(1), 69-76. https://doi.org/10.1093/llc/fqaf131

[8] Toulatzis, V., Theodoridou, S., & Fudos, I. (2026, 4월 19일). MESA: A training-free multi-exemplar deep framework for restoring ancient inscription textures (arXiv:2604.17390). arXiv. https://arxiv.org/abs/2604.17390 (accessed September 3, 2026)

[9] O'Grady, E. (2024, 10월 11일). Unearthed papyrus contains lost scenes from Euripides' plays. The Harvard Gazette. https://news.harvard.edu/gazette/story/2024/10/unearthed-papyrus-contains-lost-scenes-from-euripides-plays/ (accessed September 3, 2026)

[10] Google DeepMind. (2025, 7월 23일). Aeneas transforms how historians connect the past. deepmind.google. https://deepmind.google/blog/aeneas-transforms-how-historians-connect-the-past/ (accessed September 3, 2026)

Chapter 12. A Virtual Library That Never Burns: Ultra-High-Resolution Multispectral Archives and the Future

12.1. Case File

How the Library of Alexandria Disappeared

In Alexandria, Egypt, there was the largest library of the ancient world. The Ptolemaic dynasty built it around the 3rd century BCE. The kings tried to gather the books of the world in one place. Books from ships entering the harbor were collected and copied. The copy was returned to the owner, and the original was kept by the library. Old records say there were hundreds of thousands of scrolls [1].

Many people think this library disappeared in one great fire. But researchers believe it collapsed slowly. In 145 BCE, the king expelled foreign scholars. When the scholars left, the hands that cared for the books left with them. In the Roman period, the buildings were damaged several times. Money to buy new books also ran out [1].

Ancient writing was hard to preserve even in the first place. Scribes of that time often copied existing books rather than writing new ones. A papyrus scroll handled by many hands became worn after about a hundred years and was hard to read. In Egypt's dry soil, it sometimes survived much longer. So someone had to copy it again at the right time for the text to continue. If the copying hands stopped, the book disappeared from the world.

The ancient texts we read today are only part of the writing from that age. Far more books survive only by title, while their main texts are gone. Most works written by ancient Greek tragedians disappeared in this way. The books that remain were lucky because copies had been scattered in several places.

Destruction Continues Today

The forces that destroy records did not exist only in the past. War still brings down libraries and archives. UNESCO is identifying and counting damaged cultural facilities in Ukraine one by one. As of July 29, 2026, 545 sites had been confirmed. They include 43 museums, 22 libraries, and 5 archives. Hundreds of religious sites and old buildings were damaged as well [2].

War does not only destroy artifacts. Artifacts can also slip out secretly from storage facilities where management has stopped. An artifact whose record of origin has been erased loses much of its research value. This is because the room where one clay tablet was found, and the tablets found with it, are information in themselves.

Climate also threatens records. Sites on the coast are worn down by waves and collapse little by little. Buried artifacts decay faster when rainfall increases or groundwater rises. In some places, the rate of loss is faster than the rate of excavation.

Not every damaged artifact disappears. If its form remains, there is a way to restore it. That is why measuring it in advance matters before disaster strikes. In areas where war continues, work is underway to record buildings and artifacts in three dimensions.

Each Material Has a Different Lifespan

Materials that hold records break down in different ways. Mesopotamian clay tablets were plates made by pressing letters into mud. The heat from burning palaces fired the tablets hard like pottery. Clay tablets that did not meet fire dissolve back into mud when water enters.

Clay tablets have another limit. Scribes soaked used tablets in water, scraped off the surface, and used them again. So old tablets usually returned to mud. That is why only the last few years before a kingdom's fall often remain. Only the final tablets in the archive were fired hard when the palace burned.

Papyrus, leather, and paper are weaker. These materials are organic matter from plants and animals. In hot and humid places, they rot before lasting even a few hundred years. On wax tablets used by Roman soldiers, the wax itself has disappeared completely. What remains are only shallow scratches where the stylus cut down into the wooden board.

Hard materials are not safe either. Oracle bone script carved on turtle plastrons and ox shoulder bones was also crushed and broken underground. Most oracle bones that remain today are fragments smaller than a palm. Sometimes one sentence is scattered across three fragments stored in different warehouses [3].

The Library Saved by Fire

Records sometimes survive in unexpected ways. The Assyrian king Ashurbanipal kept a large archive in Nineveh. He ruled from around 669 BCE. The archive gathered myths, divination texts, medical texts, and administrative documents [4].

Around 612 BCE, Nineveh burned. If the books had been made of paper, they would all have disappeared then. But the clay tablets were fired by the flames and became even harder. These tablets were excavated several times from the 1840s to the 1930s. Today the British Museum holds more than 30,000 of their fragments [4].

But they did not come out intact. When the city fell, the tablets were deliberately broken and thrown away across the palace. What excavators found was not an orderly shelf but piles of scattered fragments. It was common for the front and back of one work to appear in different rooms [4]. Fitting these fragments together again is a major task in today's computational archaeology. Computational archaeology is the field that studies old artifacts with computers.

Letters Already Brought Back to Life

Digitization also faces barriers. The world's artifacts far outnumber human hands. The fragment piles from Nineveh, discussed above, are one example [4]. It takes a long time to photograph each piece and attach the text that has been read from it. So machines must extend human speed.

There are two ways to save disappearing artifacts. One is conservation, which preserves the object itself for a long time. The other is to read inside the object without touching it and move the result into a computer. The second method has already produced results several times.

The scroll from En-Gedi in Israel had burned into a lump of charcoal. If someone tried to unroll it by hand, it would have crumbled. The research team looked inside it with CT scanning, which sees inside an object without cutting it. The computer mathematically traced the rolled leather layers and flattened them into a plane. Ink mixed with metallic elements appeared on that surface, and the opening of Leviticus was read [5].

A path also opened for the Herculaneum scrolls buried in the ash of Vesuvius. On June 25, 2026, researchers announced that they had read one scroll from beginning to end. The scroll's name is PHerc. 1667. It contained an ethical discussion in Greek, about twenty-two columns long [6]. A new prize was offered for thirteen sealed scrolls designated as a challenge. The deadline is June 25, 2027 [7].

A virtual library is not a warehouse filled with photographs. The artifact's shape, ink composition, and deciphered text must be linked in one place. Only then can old data be recalculated when new methods appear later. This bundle is needed so verification can continue even after the original disappears.

Imaging and reading require money and people. It is also hard to decide which artifacts should be imaged first. Artifacts in dangerous areas are difficult to image at all. So every country sets priorities and works in order.

These cases leave one question. Even if the object is finally destroyed, can the writing inside it remain? The place that seeks that answer is the virtual library.

References

[1] Mark, J. J. (2023, 7월 25일). Library of Alexandria. World History Encyclopedia. https://www.worldhistory.org/Library_of_Alexandria/ (accessed September 3, 2026)

[2] UNESCO. (2026). Damaged cultural sites in Ukraine verified by UNESCO. UNESCO. https://www.unesco.org/en/ukraine-war/damaged-cultural-sites (accessed September 3, 2026)

[3] Ma, Y. R. (2026). Towards Computational Chinese Paleography. arXiv:2601.06753. https://arxiv.org/abs/2601.06753 (accessed September 3, 2026)

[4] The British Museum. What was Ashurbanipal's Library? The British Museum. https://www.britishmuseum.org/research/projects/what-was-ashurbanipals-library (accessed September 3, 2026)

[5] Seales, W. B., Parker, C. S., Segal, M., Tov, E., Shor, P., & Porath, Y. (2016). From damage to discovery via virtual unwrapping: Reading the scroll from En-Gedi. Science Advances, 2(9), e1601247. https://doi.org/10.1126/sciadv.1601247

[6] Vesuvius Challenge. (2026, 6월 25일). An entire Herculaneum scroll has been read for the first time. scrollprize.org. https://scrollprize.org/firstscroll (accessed September 3, 2026)

[7] Vesuvius Challenge. (2026). Open Prizes. scrollprize.org. https://scrollprize.org/prizes (accessed September 3, 2026)

12.2. The Forensic Team's Trail: Knowledge Graphs and Spectral Analysis of Light

A Large Map Connecting Artifacts

In older laboratories, the shape and text of an artifact were handled separately. Conservation scientists measured the shape, and philologists read the text. The two results went into different drawers. So even when the same artifact was studied twice, the answers could not be combined.

Computational archaeology puts both on one map. That map is called a knowledge graph. A knowledge graph is a picture that draws information as points and lines. Points are things, and lines are relationships between things.

On this map, one clay tablet becomes a point. The scribe who wrote the tablet also becomes a point, and the city where the tablet was found becomes a point. Ink composition and three-dimensional shape data each become points as well. Personal names and object names written on the tablet are also extracted as points. The computer learns by finding lines that connect these points.

Once the map is made, the shape of questions changes. In the past, scholars placed one tablet before them and asked what letter it was. Now they can ask where other documents using this ink are. They can also find at once which cities produced tablets written by the same scribe.

In February 2026, a research team at Jagiellonian University in Poland released such a map. The map contains about 3.2 million lines of relationships among people, events, and objects. The team made the machine recommend relationships and explain the reasons. Experts read those explanations and checked whether they were usable [1].

The Layer That Measures Shape

The first layer of the map is the artifact's outer shape. On old steles and clay tablets, the surface wears away and the grooves of letters become shallow. With the eye, it is hard to tell whether grooves remain. So researchers scan the surface with three-dimensional scanners.

The scanner stores the surface as millions of points. The computer connects those points into a triangular mesh. On this mesh, virtual light can be shone from several directions. When the light direction changes, even shallow grooves cast shadows, and letters appear.

Machines also learn the cross-sectional shape of grooves. Marks carved by a human tool have even cross sections and steady directions. Natural cracks do not. A model that learns this difference separates letters from scratches.

Broken cross sections are measured in the same way. The machine compares whether the curved broken edges of two fragments fit together. This calculation replaces the work of holding fragments by hand and trying to match them. It can also test matches with fragments in museums in other countries that cannot be touched.

Rolled scrolls use another kind of calculation. In the En-Gedi scroll, the research team separated the leather layers one by one inside the computer. Then they flattened each layer into a plane and restored the letters [2].

The Layer That Reads Ink with Light

The second layer of the map is ink. Ink responds to light differently depending on its materials. Multispectral imaging is a method that separates and photographs not only visible light but also infrared and ultraviolet light. At some wavelengths, the background becomes dark and the ink appears bright.

Because of this difference, erased writing can also be revived. Ancient people erased texts and wrote again to save costly parchment. A little ink component remains on the erased surface. When images taken at several wavelengths are combined, those traces rise from underneath. A document written twice in this way is called a palimpsest.

Saint Catherine's Monastery on Mount Sinai in Egypt has many such documents. Researchers from the United States and Greece used multispectral imaging on 74 volumes there. The photographed surfaces totaled 6,800 pages. Beneath them, 305 erased texts appeared again. They were written in ten languages, including Greek and Syriac [3].

There is also a method that measures the components themselves. Raman spectroscopy shines light on ink and reads the pattern of light that bounces back. The pattern differs for each substance, so it acts like a fingerprint. Artificial intelligence looks at this pattern and distinguishes whether the ink is soot-based ink or iron-containing ink. Ink mixtures vary slightly by workshop, so they become clues for guessing origin.

Stronger light was used for the Herculaneum scrolls. Researchers imaged the inside of the scrolls with a synchrotron, which produces very powerful X-rays. Burned papyrus and carbon ink have similar densities, so they cannot be distinguished in an ordinary image. A deep learning model learned subtle differences in texture and found the places where letters had been [4].

The Layer That Connects Texts

The third layer of the map is the relationship between documents. Ancient books spread through copying. Each time they were copied, small errors entered one by one. Scholars trace these errors and draw a family tree of manuscripts.

Computers borrow methods from biology for this work. It resembles drawing the evolutionary tree of living things by comparing gene sequences. The computer counts differences between documents and arranges which manuscript came from which. A family tree drawn in this way helps scholars infer the shape of the lost original.

Handwriting is also a clue that connects relationships. Each scribe has different pressure in the brush and different angles of strokes. Artificial intelligence converts these habits into numbers and stores them. When documents with close numbers are grouped together, a network of scribes appears.

Archives Already in Operation

Several archives already operate by gathering these layers. Trismegistos is a project that assigns numbers to documents from the ancient world. It has linked about 960,000 records in more than 100 languages into one system. The problem of different countries calling the same artifact by different names was solved through these numbers [5].

The Pelagios Network is a community that connects scattered materials. It sets common methods for linking places, people, and time. Even when old place names changed by period, it links them to one coordinate [6]. There is also the Electronic Babylonian Library, operated in Munich. It gathers cuneiform fragments from around the world in one database. Scholars search fragments there and recover lost sentences [7].

Common numbers may look dull, but they are powerful. Only when the numbers match can materials from different institutions line up in a single row [5]. Often, the photographs are in a museum, the deciphered text is at a university, and the excavation record is in a research institute. When one number ties these three together, scattered information gathers into the story of one artifact.

The way materials are released has also changed. In the past, photographs and deciphered texts could be seen only after a paper was published. Today, research teams release scan data and programs together. Then researchers in countries without the equipment can calculate again with the same data [8].

Tasks Still Left

On July 10, 2026, the Vesuvius Challenge organized and released the remaining problems. In places that are pressed and crushed, the layers of a scroll do not separate from each other. A computer may follow one layer and then mistakenly jump to the layer beside it. It is still not certain how deep the ink has soaked below the surface. A model trained on one scroll loses strength on another scroll. The data are so large that ordinary computers have difficulty handling them. The next target date set by the researchers is June 2027 [8].

Other fields that study ancient writing meet the same wall. A 2026 paper reviewing Chinese palaeography identified the lack of data as a major challenge. It recommended a direction in which AI helps human judgment rather than replacing it [9].

References

[1] Kutt, K., Sroka, E., Ishchuk, O., & do Valle Miranda, L. (2026). A three-stage neuro-symbolic recommendation pipeline for cultural heritage knowledge graphs. arXiv:2602.19711. https://arxiv.org/abs/2602.19711 (accessed September 3, 2026)

[2] Seales, W. B., Parker, C. S., Segal, M., Tov, E., Shor, P., & Porath, Y. (2016). From damage to discovery via virtual unwrapping: Reading the scroll from En-Gedi. Science Advances, 2(9), e1601247. https://doi.org/10.1126/sciadv.1601247

[3] Early Manuscripts Electronic Library. Sinai Palimpsests Project. EMEL. https://www.emelibrary.org/sinai-palimpsests-project/ (accessed September 3, 2026)

[4] Vesuvius Challenge. (2026, 6월 25일). An entire Herculaneum scroll has been read for the first time. scrollprize.org. https://scrollprize.org/firstscroll (accessed September 3, 2026)

[5] KU Leuven. Trismegistos portal. Trismegistos+. https://tmplus.kuleuven.be/trismegistos-portal (accessed September 3, 2026)

[6] Pelagios Network. Pelagios Network. https://pelagios.org/ (accessed September 3, 2026)

[7] Ludwig-Maximilians-Universität München & Bayerische Akademie der Wissenschaften. electronic Babylonian Library. https://www.ebl.badw.de/ (accessed September 3, 2026)

[8] Vesuvius Challenge. (2026, 7월 10일). Open problems: Why reading every Herculaneum scroll is still a challenge. scrollprize.org. https://scrollprize.org/2026_open_problems (accessed September 3, 2026)

[9] Ma, Y. R. (2026). Towards Computational Chinese Paleography. arXiv:2601.06753. https://arxiv.org/abs/2601.06753 (accessed September 3, 2026)

12.3. Epilogue

Looking Back at Eleven Cases

This book has followed eleven cases. It began with the scrolls of Herculaneum, buried under volcanic ash. A computer virtually unfolded a lump of charcoal that would crumble if touched. Burned papyrus and carbon ink look almost the same under X-rays. The machine learned the tiny differences in surface texture and found the letters [1].

The Leviticus scroll from En Gedi also recovered its first verses through virtual unwrapping. In this case, ink containing metallic elements appeared clearly in X-rays [2]. In the Isaiah Scroll, handwriting was the clue. AI measured small habits in the strokes and showed that there was more than one scribe. To human eyes, those parts had looked like the same handwriting [3].

From the fourth case on, the work became a matter of fitting fragments together. Scattered clay tablets and parchment fragments were rejoined by the shapes of their edges. The same calculations were used to detect forged artifacts. Missing lines of the Epic of Gilgamesh were filled by joining fragments together.

Mesopotamian receipts began to speak through machine translation. Grain accounts and letters from thousands of years ago were rendered in today's language. Oracle bone script on turtle plastrons was joined together with 3D scans and dedicated models [4]. Wooden slips from Silla were also read by character recognition models. Japanese cursive documents were deciphered in the same way [5].

We also saw doors that have not yet opened. Linear A, Indus script, and rongorongo still have no answer. Machines measured regularities in the arrangement of signs and examined whether they were language or not. Reading their meaning remains the task of the next generation.

The final case was autocomplete for ancient inscriptions. Ithaca is a model that fills missing letters in Greek inscriptions. While filling the letters, this model also estimates where and when the inscription was made [6]. Aeneas, which came later, handles Latin inscriptions. Aeneas looks at both transcribed letters and photographs of inscriptions [7].

The eleven cases have something in common. None of them was solved by machines alone. Physicists handled imaging, and computer scientists handled computation. Scholars who read ancient languages checked the results and added meaning. Only when the three fields met around the same artifact did one line come alive.

What Machines Still Cannot Do

Along the way, the weaknesses of machines also became clear. The first wall is the lack of data. Linear A has only about 7,400 surviving signs in all. With such a small amount, it is hard for a machine to learn rules [8]. Today's language models grow by reading billions of words. The surviving data for lost ancient languages fall far short of that scale.

If a machine is forced to produce answers where data are scarce, it invents words that do not exist. The longer the blank, the lower the reliability of the answer. A sentence may fit the context but never have existed in history. Once such a sentence appears in a paper, later research may accept it as fact.

When training data lean to one side, the answer also leans. If inscriptions from one region are numerous, the machine may read inscriptions from other regions in that region's style. The machine often cannot explain by itself why it gave that answer. So recent research is moving toward providing evidence together with the result [9].

The scroll-reading research team also released a list of remaining problems. In pressed and crushed places, the layers do not separate from each other. A model trained on one scroll may lose strength on another scroll. Stating what cannot yet be done creates the next step [10].

Rules are also needed for handling letters produced by machines. Letters read with certainty and letters guessed must be marked differently. Researchers must also record which model read them that way and with which data. Only then can another researcher later run the same calculation again.

The Original Remains the Original

No matter how precise digital restoration becomes, it cannot replace the original. Scholars see scanned data not as the original, but as a digital counterpart. Imaging includes human choices. A person decides what light to use and from what angle to shoot. If those choices change, the resulting image also changes.

Physical objects retain information that scanners cannot capture. The structure of soil grains on the surface of a clay tablet is one example. The fine pattern of holes in a turtle plastron can also be seen only in the original. How deeply ink soaked into papyrus can also be measured only from the physical object [10].

New measurement technologies will continue to appear. They cannot be used on artifacts discarded today. That is why artifacts should not be cut or forcibly opened. The purpose of a virtual library is not to throw away the original. Its purpose is to make reading possible while touching the original less [2]. When an artifact is handled fewer times, the life of the original grows longer.

A Future of Reading Together

AI is not a tool that pushes historians aside. The Ithaca study showed this with numbers. When Ithaca restored missing letters by itself, its accuracy was 62 percent. When historians worked alone, it was 25 percent. When the two worked together, it rose to 72 percent [6].

These numbers tell us two things at once. In this experiment, the machine alone performed better than humans alone. And its performance improved further when it worked with humans. Neither side alone could have reached that point.

Machines produce many candidates quickly. Humans judge whether those candidates could have existed in that era. Scholars call this alternating confirmation between humans and machines human-in-the-loop collaboration. For this loop to turn, people who can read ancient texts must continue to exist. If there is no one to read them, the letters produced by machines become another riddle.

Teaching is as important as research. The skill of reading ancient languages has passed from person to person. If that chain breaks, there will be no one to verify the letters produced by machines. The new generation must learn ancient scripts and computers together.

This loop does not belong only to experts. In Japan, Miwo, a free app that reads cursive script, is available. When a user takes a photo, modern characters appear on the screen. Even someone seeing an old document for the first time can read the first line [5]. Holding up a smartphone in front of a museum display case can become the beginning of research.

There is still much work to do. Sealed scrolls still remain in Herculaneum [11]. In Ukraine, libraries and archives are still being destroyed [12]. The world's old writings are still disappearing faster than they are being saved.

Readers of this book can also do something. Anyone can download and view publicly released scan data [10]. Students also take part in competitions to match fragments [11]. In many places, citizens join in transcribing old documents. Even if each person takes one line, the data accumulate.

Protecting records is not only about old stories. The materials we create today will also become old materials someday. If file formats change, today's documents may no longer be readable. That is why libraries keep moving materials into new formats. Building a virtual library always comes with hands that maintain it.

Still, the direction is set. A scroll buried under volcanic ash 2,000 years ago has been read again [1]. The method gained from one scroll becomes the starting point for research on the next scroll [11]. If data and code are released, researchers in other countries can walk the same path [10]. Even if an artifact someday returns to the soil, the sentences inside it remain. A library that does not burn is being built one line at a time.

References

[1] Vesuvius Challenge. (2026, 6월 25일). An entire Herculaneum scroll has been read for the first time. scrollprize.org. https://scrollprize.org/firstscroll (accessed September 3, 2026)

[2] Seales, W. B., Parker, C. S., Segal, M., Tov, E., Shor, P., & Porath, Y. (2016). From damage to discovery via virtual unwrapping: Reading the scroll from En-Gedi. Science Advances, 2(9), e1601247. https://doi.org/10.1126/sciadv.1601247

[3] Popović, M., Dhali, M. A., & Schomaker, L. (2021). Artificial intelligence based writer identification generates new evidence for the unknown scribes of the Dead Sea Scrolls exemplified by the Great Isaiah Scroll (1QIsaa). PLOS ONE, 16(4), e0249769. https://doi.org/10.1371/journal.pone.0249769

[4] Ma, Y. R. (2026). Towards Computational Chinese Paleography. arXiv:2601.06753. https://arxiv.org/abs/2601.06753 (accessed September 3, 2026)

[5] ROIS-DS Center for Open Data in the Humanities. みを(miwo): AIくずし字認識アプリ. CODH. https://codh.rois.ac.jp/miwo/ (accessed September 3, 2026)

[6] Assael, Y., Sommerschield, T., Shillingford, B., Bordbar, M., Pavlopoulos, J., Chatzipanagiotou, M., Androutsopoulos, I., Prag, J., & de Freitas, N. (2022). Restoring and attributing ancient texts using deep neural networks. Nature, 603(7900), 280-283. https://doi.org/10.1038/s41586-022-04448-z

[7] Assael, Y., Sommerschield, T., Cooley, A. E., Shillingford, B., Pavlopoulos, J., Suresh, P., Herms, B., Grayston, J., Maynard, B., Dietrich, N., Wulgaert, R., Prag, J., Mullen, A., & Mohamed, S. (2025). Contextualizing ancient texts with generative neural networks. Nature. https://doi.org/10.1038/s41586-025-09292-5

[8] Braović, M., Krstinić, D., Štula, M., & Ivanda, A. (2024). A systematic review of computational approaches to deciphering Bronze Age Aegean and Cypriot scripts. Computational Linguistics, 50(2), 725-779. https://doi.org/10.1162/coli_a_00514

[9] Kutt, K., Sroka, E., Ishchuk, O., & do Valle Miranda, L. (2026). A three-stage neuro-symbolic recommendation pipeline for cultural heritage knowledge graphs. arXiv:2602.19711. https://arxiv.org/abs/2602.19711 (accessed September 3, 2026)

[10] Vesuvius Challenge. (2026, 7월 10일). Open problems: Why reading every Herculaneum scroll is still a challenge. scrollprize.org. https://scrollprize.org/2026_open_problems (accessed September 3, 2026)

[11] Vesuvius Challenge. (2026). Open Prizes. scrollprize.org. https://scrollprize.org/prizes (accessed September 3, 2026)

[12] UNESCO. (2026). Damaged cultural sites in Ukraine verified by UNESCO. UNESCO. https://www.unesco.org/en/ukraine-war/damaged-cultural-sites (accessed September 3, 2026)

Appendix. My Own AI Archaeology Travelogue: A Global Exploration Portal for the Public

A.1. Vesuvius Challenge Official Data Archive

What You Can See

The Vesuvius Challenge is an international competition. It is a contest to read burned scrolls with computers. The scrolls came from Herculaneum. It was an ancient Roman city beside Mount Vesuvius. The scrolls were made of papyrus. Papyrus is an ancient paper made by pressing a plant. The official address is scrollprize.org[1]. This site opens the data used in the competition to everyone[2]. No membership registration is needed, and it costs no money[2][5].

Menus are lined up at the top of the first screen. Prizes, Problems, Tutorials & In-depth, and Data come first. Milestones, Community, Press, and Donate follow them. As of September 2026, 45 scrolls and papyrus fragments have been scanned. One of them has been read from beginning to end[1].

On June 25, 2026, the competition team announced news about that one scroll. It is a scroll called PHerc. 1667[3]. It is the first case in which a scroll was read to the end without being opened[3][4]. The unfolded papyrus is about 1.4 meters long. Ancient books divided writing into vertical columns. One such column is called a column. In this scroll, about twenty-two columns written in Greek letters appeared. The content is a text on Stoic ethics. It discusses human nature, impulses, and moral growth. It is thought to be a text from around the 2nd century BCE[3].

The whole scroll has not survived intact. What remains now is the inner core of the scroll. The outer layers broke apart when people tried to open them by hand in the 19th century and the 1980s. Where the surface fell away, the letters also disappeared. So the text that was read has gaps here and there[3].

The scan was done in Grenoble, France. It was done at the European Synchrotron Radiation Facility, ESRF. This is a large research device that creates very bright X-rays. The BM18 beamline was used for this work[4]. A beamline is an experimental passage that brings X-rays to an object and shines them on it.

The Order to Follow at Home

First, go to scrollprize.org and press the Get Started button. This screen guides first-time visitors through three steps[5].

The first step is finding ink. You mark visible ink on the screen. Then the model learns from those marks and predicts the rest[5].

The second step is virtual unwrapping. This means opening the scroll with a computer. You follow one layer of papyrus in CT cross-section images. CT cross-section images are pictures taken as if the inside had been sliced thinly. You flatten that layer and check the result[5].

The third step is learning the whole process. You try scanning, unwrapping, and reading in a small-scale sequence[5].

You should know that the experience sequence and the actual restoration sequence are different. The experience starts with easy ink finding. In actual restoration, scanning comes first. Then the scroll is virtually unwrapped. Reading the ink is the last step[5].

To download the data directly, go to scrollprize.org/data. The data is in the public repository s3://vesuvius-challenge-open-data/. You can also open it in a web browser. Example files that show how to use it are on GitHub[2].

Prizes are also offered to challengers. As of September 2026, the remaining prize money is 2.14 million dollars. The 2027 Grand Prize is worth 1 million dollars. First place receives 800,000 dollars. The deadline is June 25, 2027. Prizes are also offered for scrolls in which no letters have yet been found. A challenger must read ten letters within an area of 4 square centimeters. For this, 50,000 dollars is given for each scroll. There is also a monthly Progress Prize. The best submission chosen that month receives 20,000 dollars[6].

Points to Watch

The data is released under CC BY-NC 4.0 terms. It cannot be used for making money. When using it, you must identify the source. Some older data has separate papers that must be cited. This applies to Scrolls 1 through 4 and Fragments 1 through 6[2]. In that case, cite the 2023 paper by Parsons and others[2][7].

You should also know in advance that the files are large. The full 3D scan data reaches several terabytes[2]. To run deep learning, you need a powerful graphics card. If your home computer is weak, you can start by looking through the data with your eyes.

Conditions apply when you enter for a prize. You must join the competition Discord room. The code must be released under open terms such as MIT. You must not announce the results outside before the award announcement[6].

References

[1] Vesuvius Challenge. (2026). Vesuvius Challenge: Reading the Herculaneum Scrolls with machine learning. scrollprize.org. https://scrollprize.org/ (accessed September 3, 2026)

[2] Vesuvius Challenge. (2026). Data. scrollprize.org. https://scrollprize.org/data (accessed September 3, 2026)

[3] Vesuvius Challenge. (2026, 6월 25일). An entire Herculaneum scroll has been read for the first time. scrollprize.org. https://scrollprize.org/firstscroll (accessed September 3, 2026)

[4] Angelotti, G., Parsons, S., Nicolardi, F., et al. (2026). Complete virtual unwrapping and reading of a rolled Herculaneum papyrus. arXiv:2606.29085. https://arxiv.org/abs/2606.29085 (accessed September 3, 2026)

[5] Vesuvius Challenge. (2026). Get Started. scrollprize.org. https://scrollprize.org/get_started (accessed September 3, 2026)

[6] Vesuvius Challenge. (2026). Open Prizes. scrollprize.org. https://scrollprize.org/prizes (accessed September 3, 2026)

[7] Parsons, S., Parker, C. S., Chapman, C., Hayashida, M., & Seales, W. B. (2023). EduceLab-Scrolls: Verifiable recovery of text from Herculaneum papyri using X-ray CT. arXiv:2304.02084. https://doi.org/10.48550/arXiv.2304.02084

A.2. eBL Electronic Babylonian Literature Portal

The Structure of the Portal

eBL is short for electronic Babylonian Library. In Korean, it means an electronic Babylonian library. Ancient Mesopotamians wrote in cuneiform. Cuneiform is wedge-shaped writing pressed into wet clay tablets. This portal brings together fragments of those clay tablets in one place. The objects it collects are fragments that contain Babylonian literature. It is not a place that contains all cuneiform materials in the world[1].

The portal is based at Ludwig Maximilian University of Munich in Germany. It is led by Professor Enrique Jimenez. The project began in April 2018. It is supported by the Sofja Kovalevskaja Award from the Alexander von Humboldt Foundation[2].

The address is www.ebl.lmu.de. You can enter without logging in, and there is no fee. On the left side of the screen, there are eight items. eBL Project, Library, Corpus, and Signs come first. Akkadian Dictionary and Bibliography come after them. News and Archaeology are at the end[1].

Library is the place where individual fragments are collected. It used to be called Fragmentarium. eBL researchers have found about 1,200 cases of fragments being joined together there[1]. Corpus shows texts by work. The Epic of Gilgamesh is organized into about 3,150 lines. The creation myth Enuma Elish has 1,096 lines[3]. A data paper published in 2024 gives the scale. This portal contains about 25,000 clay tablets. The transliterations exceed 350,000 lines[4]. A transliteration is a text that rewrites cuneiform in Latin letters.

What You Can Try on the Screen

When you press Library, a search box appears. You can search by name, year, and title. There are also word-level search and advanced search. If you do not know what to look for, press the I'm feeling lucky button. Then a random fragment opens[5].

When you open one fragment, the screen is divided into several sections. At the top, the holding institution, size, and findspot are written. The script style and document type also appear. For example, K.5743 is an item from the Kuyunjik collection of the British Museum. Kuyunjik is the mound where the ancient city of Nineveh stood. It is the name used for artifacts from that mound as a group. Its size is 3.81 centimeters by 2.54 centimeters. The script style is Neo-Assyrian[6]. This is the letter form used during the Neo-Assyrian Empire.

In the middle, the transliteration appears with line numbers. At the bottom, a photograph of the clay tablet is attached. It also states which fragment it has been joined with[6].

The site also tells you how to cite the text. Write the fragment number and eBL edition. Then add the address and the access date[1].

News and Points to Watch

You can see newsletters under News. Issue 22, dated May 28, 2026, is the latest one. In this issue, the portal introduced a new screen layout. The number of holding institutions also increased. The National Museum of World Writing Systems in Incheon is one of them. More than 3,100 British Museum photographs were newly uploaded. More than 1,000 Yale Peabody Museum photographs were also uploaded. A function for directly registering new words in the dictionary was also added[7].

There is something to know before using it. The transliterations in Library are not finished texts. They are provisional texts. The researchers are still correcting them. The same is true of the work texts organized in Corpus. They are interim results assembled from fragments collected so far. When new fragments are joined, the number of lines and the content change. So when using them in a school assignment or article, write the access date as well[1].

The transliterations are released under CC BY-NC-SA 4.0 terms. The photographs have different conditions. To reuse the photographs, you must get separate permission[1]. The museum that owns the clay tablet is the side that grants permission[1][4]. You should first read the notice attached to the photograph description at the bottom of the screen[1].

References

[1] electronic Babylonian Library. (2026). About: Library. ebl.lmu.de. https://www.ebl.lmu.de/about/library (accessed September 3, 2026)

[2] Jiménez, E. (2023, 5월 25일). AI is helping us read ancient Mesopotamian literature. The Conversation. https://theconversation.com/ai-is-helping-us-read-ancient-mesopotamian-literature-205091 (accessed September 3, 2026)

[3] electronic Babylonian Library. (2026). Corpus. ebl.lmu.de. https://www.ebl.lmu.de/corpus (accessed September 3, 2026)

[4] Cobanoglu, Y., Laasonen, J., Simonjetz, F., Khait, I., Cohen, S., Földi, Z., Hätinen, A., Heinrich, A., Mitto, T., Rozzi, G., Sáenz, L., & Jiménez, E. (2024). Transliterated cuneiform tablets of the electronic Babylonian Library platform. Journal of Open Humanities Data, 10(1), 19. https://doi.org/10.5334/johd.148

[5] electronic Babylonian Library. (2026). Library. ebl.lmu.de. https://www.ebl.lmu.de/library (accessed September 3, 2026)

[6] electronic Babylonian Library. (2026). K.5743, eBL edition. ebl.lmu.de. https://www.ebl.lmu.de/library/K.5743 (accessed September 3, 2026)

[7] electronic Babylonian Library. (2026, 5월 28일). eBL Newsletter 22. ebl.lmu.de. https://www.ebl.lmu.de/news (accessed September 3, 2026)

A.3. Guide to Using Miwo

What Kind of App Is It?

Cursive script was widely used in old Japanese books and documents. These letters are called kuzushiji. They often appear in Edo-period printed books and old documents[3]. Most Japanese people today cannot read these letters.

Miwo is a smartphone app that photographs this cursive script and converts it into modern letters. It was made by CODH, the Center for Open Data in the Humanities in Japan[1]. It is free[2]. The app name comes from Miotsukushi, the fourteenth volume of The Tale of Genji. A miotsukushi is a marker that shows a safe water route. The name means that the app will serve as a guide in the sea of old documents[3].

The app first appeared on August 30, 2021. On October 26, 2022, it was updated to version 1.1. At that time, a new recognition model called Ruri was added. Ruri detects letters well even when the letter background has patterns or colors[1]. Training used the Japanese classical kuzushiji dataset. This is a collection of old letter images made for machine learning. It contains more than 1 million characters. This app received the 2022 Good Design Award[3]. As of September 2026, the app has recognized more than 3.5 million images[1].

Installation and Order of Use

First, open the app store. Use the App Store for iPhone and iPad. Use Google Play for Android devices. Enter みを or miwo in the search box[1]. In the App Store, the app size is about 52 megabytes. iPhone and iPad require iOS 11.0 or later[2].

When you turn on the app, a photo-taking screen appears. Put the old document on the screen and press the shooting button. You can also load a photo you have already taken. Then press the recognition button, and the artificial intelligence reads the letters[1].

The result screen has several functions. There is one adjustment bar on the screen. If you slide this bar left and right, the two screens overlap. You can compare the original photo and the read-out letters side by side. There is also a function for pressing Japanese kana letters. It tells you which Chinese character that letter originally came from. You can move the read-out text to another app with the copy button[1]. You can also save the result and open it again later[1][3].

You should remember that the internet must be on. The photo is sent to the CODH server for recognition. The sent photo is not kept on the server[1].

Materials It Reads Well and Materials It Does Not Read Well

If you know the limits before using the app, you will be less disappointed. The artificial intelligence was trained mainly on printed books from the Edo period. Many books from that period were woodblock prints. These are books made by carving letters into wooden blocks and printing them on paper. So it reads printed books from that period well. Accuracy is lower for materials from other periods and for handwritten documents[3].

The condition of the paper also changes the result. Stained areas or places eaten by insects are hard to read. The pattern of the paper, lighting, and shadows also interfere[3]. It is better to photograph under bright and even light.

There are also things it cannot read. Cursive script carved on stone monuments or signs is not currently supported[3].

The developers state that recognition results may contain errors. They make clear that artificial intelligence recognition of cursive script will not become perfect. So important materials must be checked again by a person[3].

The work of rewriting old letters into modern letters is called transcription. There is a separate place where many people share this work. It is a citizen-participation transcription site called Minna de Honkoku. In Korean, it would mean Everyone's Transcription. It is run by an institution different from the one that made Miwo. Its nature is the same as the work of people refining the text read by the app. The address is honkoku.org. So far, more than 10,800 people have participated. The transcribed characters have exceeded 48 million. It had its second renewal in August 2025. The new version is in trial operation. The National Museum of Japanese History participates. The Earthquake Research Institute of the University of Tokyo and researchers from Kyoto University also run it together[4].

References

[1] ROIS-DS 人文学オープンデータ共同利用センター. (2026). AIくずし字認識アプリ「みを（miwo）」. CODH. https://codh.rois.ac.jp/miwo/ (accessed September 3, 2026)

[2] Clanuwat, T. (2022, 10월 27일). miwo (Version 1.1) [Mobile app]. App Store. https://apps.apple.com/jp/app/miwo/id1581794085 (accessed September 3, 2026)

[3] ROIS-DS 人文学オープンデータ共同利用センター. (2026). みを（miwo）とは？. CODH. https://codh.rois.ac.jp/miwo/about/ (accessed September 3, 2026)

[4] みんなで翻刻. (2026). みんなで翻刻 Minna de Honkoku. honkoku.org. https://honkoku.org/ (accessed September 3, 2026)

Support This AI Library Book

If this book was useful, you can support future translations, public editions, and open AI knowledge projects through PayPal.

PayPal

https://paypal.me/kimkjkj

Scan the QR code or open the link above.

Kim Kyung-jin, Attorney at Law