TC: A Journal of Biblical Textual Criticism 25 (2020): 129–143

Digital Tools for Paleography in the OCHRE Database Platform

Sarah Yardney, Miller Prosser, and Sandra R. Schloen
University of Chicago

Introduction[1]

When integrated with text editions and images, digital paleography tools can deliver more accurate and informed results and enable specialists to communicate their message clearly and powerfully. In this article we introduce new digital tools for paleography in the Online Cultural and Historical Research Environment (OCHRE). Specifically, we demonstrate how to produce reconstructed readings of Dead Sea Scrolls manuscripts in the Critical Editions for Digital Analysis and Research project (CEDAR), and how to produce dynamic letter charts of Hebrew letters in CEDAR and of alphabetic Ugaritic letters in the Ras Shamra Tablet Inventory (RSTI).[2]

First, however, a few words on why we are presenting tools for paleography in a journal devoted to textual criticism. Textual criticism of the Bible is an ancient endeavor, dating back to the time of Origen, but the discovery of the Dead Sea Scrolls in the middle of the 20th century brought us a vast quantity of new evidence. The manuscripts from Qumran are crucial for text critics, but most are also badly damaged, and the published editions are not always reliable. As a result, it behooves text critics to read the fragments and reconstruct them when necessary.[3] Using real characters from the scribe's own hand is the most methodologically sound way to test reconstructions, and we illustrate how OCHRE's reconstruction tools facilitate this process.[4] When the identification of a fragment is in doubt, letter charts can enable a scholar to determine which scroll the fragment belongs to; OCHRE can also assist here by generating interactive charts taken directly from manuscript images.

1 The authors would like to thank Doren Snoek for his helpful comments on drafts of this article.

2 CEDAR and RSTI are each discrete research projects in the OCHRE database environment.

3 High-resolution digital images of the Scrolls are available through the Leon Levy Dead Sea Scrolls Digital Library, https://www.deadseascrolls.org.il/ [accessed 8 October 2020]. The manuscripts from Qumran are the main focus of our discussion, but these tools could be used to reconstruct broken manuscripts in any language.

4 The Scripta Qumranica Electronica (SQE) will provide a font created from a scribe's handwriting that can be used to type reconstructions over a manuscript image; see Bronson Brown-deVost, review of *Evidence of Editing: Growth and Change of Texts in the Hebrew Bible*, *JHebS* (forthcoming). The font, however, represents an average of the scribe's characters (Bronson Brown-deVost, personal communication) and thus offers little methodological advance over the practice of calculating the average width of letters, proposed in Edward D. Herbert, *Reconstructing Biblical Dead Sea Scrolls: A New Method Applied to the Reconstruction of 4QSam[a]*, STDJ 22 (Leiden: Brill, 1997). Bruce Zuckerman has developed a more rigorous approach, using Photoshop to test characters from elsewhere in a scroll; see, for example, his reconstructions of col. 2 of 11QT[a] in Bernard M. Levinson, "Refining the Reconstruction of Col. 2 of the *Temple Scroll* (11QT[a]): The Turn to Digital Mapping and Historical Syntax," *DSD* 23 (2016): 1–26. The tools in CEDAR produce similar results to Zuckerman's work with the additional functionality of being able to query for all instances of a character for testing (see below).

The CEDAR and RSTI Projects in OCHRE

Based at the University of Chicago, CEDAR is a multi-project digital humanities initiative dedicated to creating new digital tools for text-critical research using the OCHRE platform.[5] The CEDAR-Bible project team includes Sandra Schloen and Miller Prosser of the OCHRE Data Service; Jeffrey Stackert and Simeon Chavel of the University of Chicago; Ronald Hendel of the University of California, Berkeley; Sarah Yardney, postdoctoral research specialist at the University of Chicago Divinity School; and graduate students from the University of Chicago Divinity School and the Department of Near Eastern Languages and Civilizations.

RSTI began as an attempt by Prosser to organize his dissertation research on a subset of tablets from Ras Shamra-Ugarit. With Dennis Pardee's partnership, the project expanded to serve multiple purposes. The RSTI project in OCHRE now documents an inventory of all inscribed objects discovered at Ugarit by the French excavation team, listing findspots, museum number, publication information, and attested language(s).[6] RSTI has also become the central database for managing tens of thousands of photographs of tablets produced by various researchers. Integrated with these various sets of data are detailed text editions and geospatial data about the excavations at Ras Shamra.

The OCHRE database was originally developed to record and manage archaeological data, but it now supports over fifty projects across a wide range of disciplines in the humanities, social sciences, and hard sciences.[7] A distinctive feature of the OCHRE platform is its item-based approach to data management, which makes possible rich integration of diverse data. In particular, OCHRE's item-based structure allows for the granularity and flexibility required for the study of ancient texts, enabling the researcher to record observations about the text at many different scales, from the level of the entire manuscript down to the level of individual characters, and even of individual diacritical marks where these are present. This atomization of the textual data is beneficial for paleographical studies that focus on discrete characters or clusters of characters. Using OCHRE, every linguistically meaningful mark on the manuscript can be a separately addressable entity—an item. OCHRE items of various kinds are organized and linked in many helpful ways, creating an integrated network of related content. For the tasks we describe in this article, the database allows us to integrate high-quality digital images and text transcriptions. In short, for both the textual reconstruction and letter chart features, we associate every letter from our text transcription with the place on the image where the letter is attested.

5 For more information on the text-critical aspects of CEDAR, see Sarah Yardney, Sandra R. Schloen, and Miller Prosser, "New Digital Tools for a New Critical Edition of the Hebrew Bible," *Open Theology* 5, no. 1 (2019): 80–94. A full description of the project is also available at https://cedar.uchicago.edu [accessed 22 July 2020].

6 See http://ochre.lib.uchicago.edu/RSTI/ [accessed 10 September 2020]. This information was originally published in Pierre Bordreuil and Dennis Pardee, *La trouvaille épigraphique de l'Ougarit* (Paris: Éditions Recherche sur les civilisations, 1989). However, because excavations continued after publication of this volume, it quickly became out of date.

7 David Schloen and Sandra Schloen, *OCHRE: An Online Cultural and Historical Research Environment* (Winona Lake, IN: Eisenbrauns, 2012) and David Schloen and Sandra Schloen, "Beyond Gutenberg: Transcending the Document Paradigm in Digital Humanities," *Digital Humanities Quarterly* 8, no. 4 (2014). See also "OCHRE," https://ochre.uchicago.edu [accessed 21 July 2020].

Integrating Texts and Images

Let us consider first the process of integrating text transcriptions and digital images. The structure of a text in OCHRE is important background for the current discussion because the paleography tools in OCHRE leverage the highly itemized nature of the text. A paleographical study in OCHRE begins by constructing in the database a transcription of the manuscript(s) under consideration. Once modeled in the database, a text can be understood as a hierarchical structure with two interconnected branches.[8] At the top of the hierarchy is an item representing the text in question. This database item can be described with metadata, links, notes, events, bibliography, and properties such as genre. This item contains two branches of textual content: the epigraphic and discourse hierarchies. These are each a hierarchy of recursively organized items that record the details of the text. The epigraphic hierarchy records the researcher's interpretation of the observable signs and letters and their spatial layout. These items are recursive in that the epigraphic item that represents line one is the hierarchical parent of an epigraphic item that represents the first letter in line one, which in turn might be the hierarchical parent of the epigraphic item that represents the first vowel mark of the letter. This recursive pattern repeats to describe the text in whatever detail is necessary for the research question being addressed. The discourse hierarchy represents the interpretation of the letters as words, phrases, stanzas, verses, chapters, or any other discourse structure. Discourse items are also recursive. A chapter may contain verses that contain phrases that contain words. In the OCHRE data model, a text is composed of these two highly articulated recursive branches of content. Further, letters in the epigraphic hierarchy are linked to words in the discourse hierarchy, thereby creating relationships between these two branches.

One of OCHRE's strengths is its integration of text and images. Multiple images of a manuscript can be added to OCHRE and linked to text transcriptions;[9] images and transcriptions can then be viewed side by side. To achieve the degree of integration that allows us to propose reconstructions or create letter charts, specific areas of an image are demarcated and linked to parts of the text, be they words, letters, or even diacritical marks. As shown below (Figure 1), the scholar simply draws a polygon around the relevant portion of the image and then chooses which element of the transcription to link with it. This feature is called *hotspotting* and is particularly useful when applied to damaged texts.

8 Here we are describing the underlying arrangement of the data in the database. At any time, the text can be viewed as a normal line-by-line document.

9 Images can be added to OCHRE from a local drive or from a server location accessible through standard web protocols. There is no limit to the number of images that can be linked to a text. Most projects using the OCHRE platform, CEDAR included, store digital images on a project server, making them available to any project member around the world through the OCHRE app or its API (application program interface).

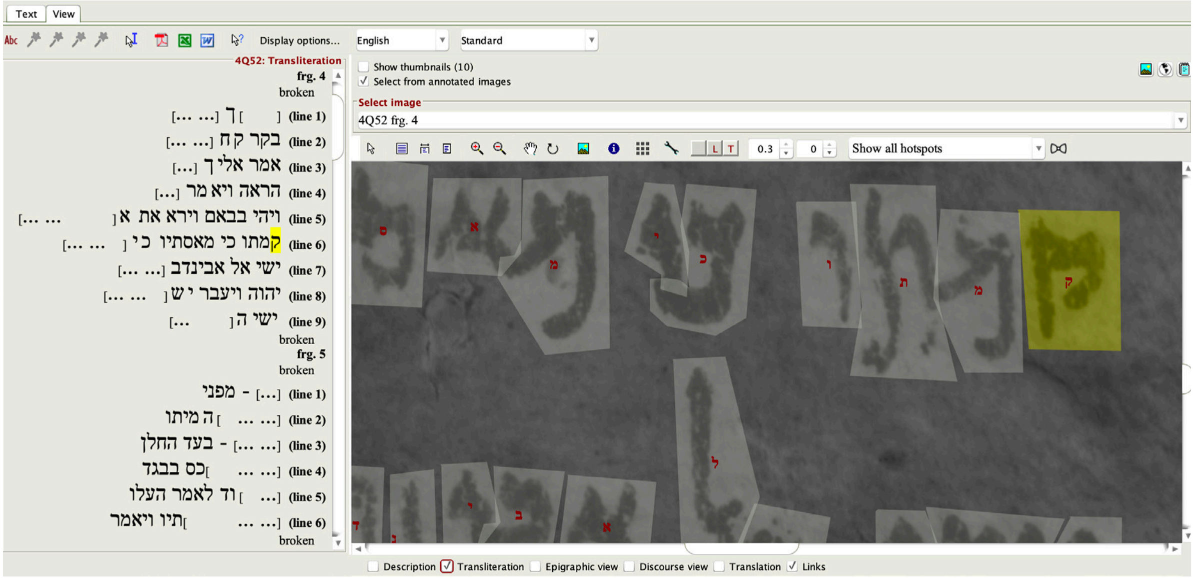

Figure 1. A screenshot of the OCHRE app shows an infrared image of a fragment of 4Q52 (4QSam^{b}) next to a transcription of the manuscript. The gray polygons on the image are the hotspots, with small red characters indicating the letter to which each hotspot is linked. Clicking on a hotspot or on the transcription will highlight the linked characters; here the *qoph* in line 6 is highlighted. The image of 4Q52 is used courtesy of The Leon Levy Dead Sea Scrolls Digital Library, Israel Antiquities Authority; photo: Shai Halevi.

Not only does hotspotting document the scholar's proposed reading of the text, but it also makes the manuscript image effectively text-searchable: a scholar can query the database for all instances of a character in the manuscript, and the database will return images of that character. These thumbnail images can be viewed in a table with or without additional information about the location of the character in the manuscript and the word of which it is a part; see Figure 2 and Figure 3.

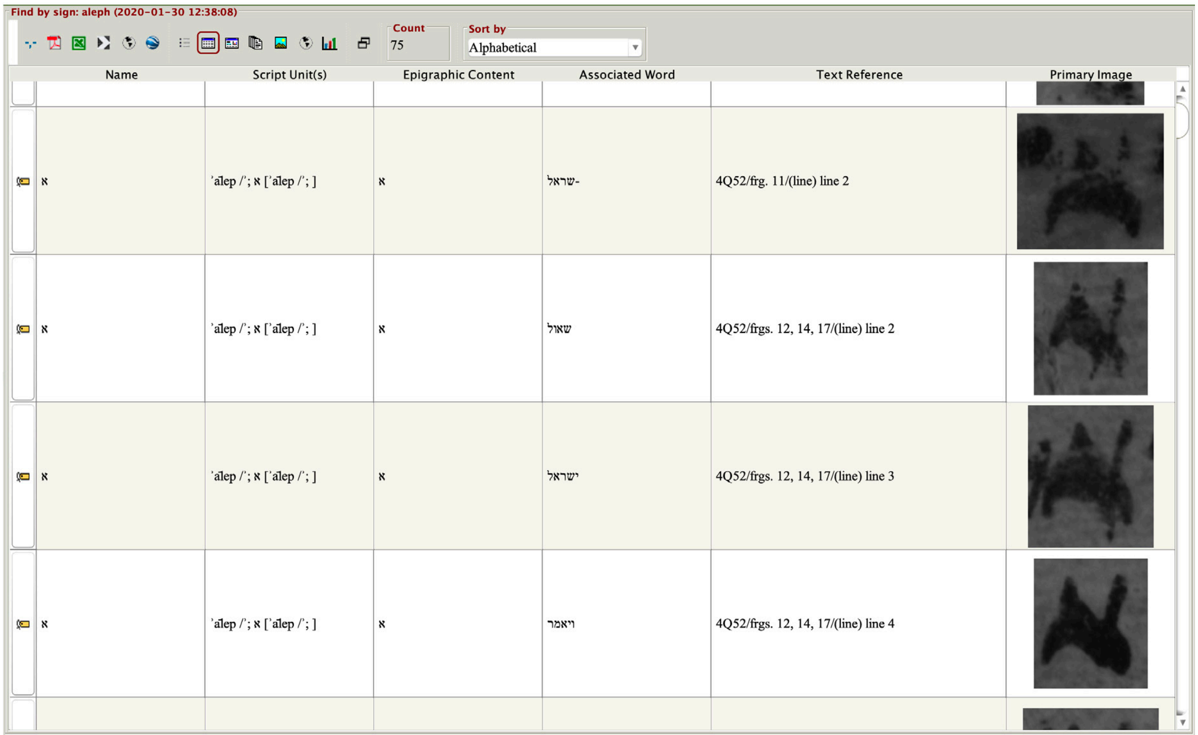

Figure 2. OCHRE can display every instance of a character—here, *aleph*—that has been hotspotted in the images of a manuscript. The "Associated Word" column gives the word in which the *aleph* appears, and the "Text Reference" column gives the location of the *aleph* in the manuscript.

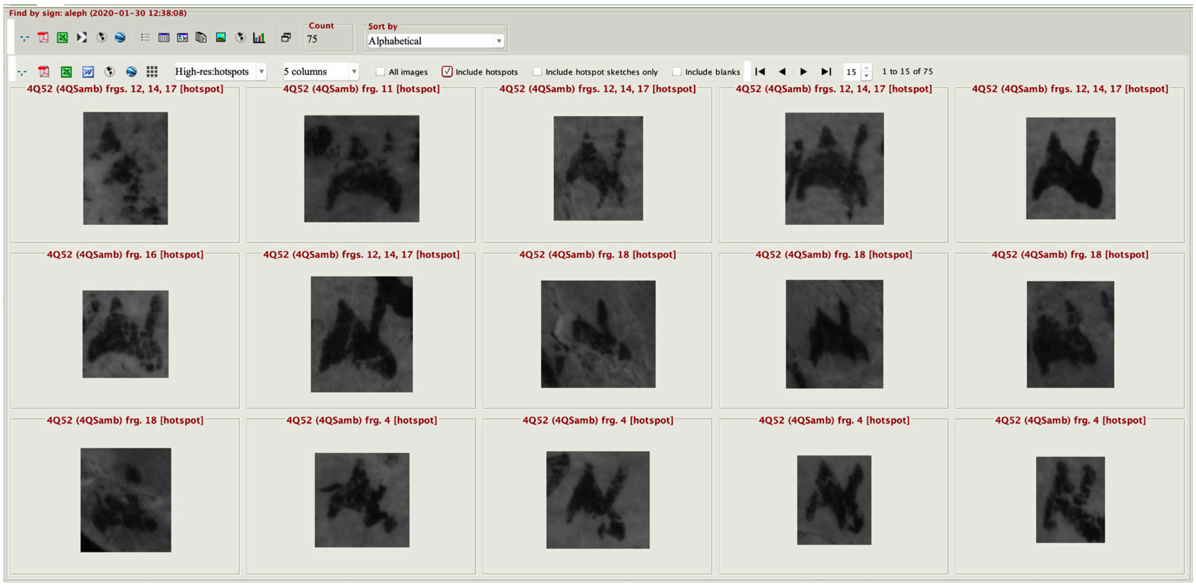

Figure 3. OCHRE can also display every instance of a character in a more condensed gallery format without additional information about its context. Contextual information can, however, be viewed in a popup window by clicking on a hotspot.

Both views are useful for different purposes. The information about context is necessary when choosing a character to use in the reconstruction of a damaged manuscript: a methodologically sound reconstruction will use characters from the same letter environment and from nearby in the manuscript. At other times a scholar will simply want to see every instance of *aleph* in a scribal hand, for example when visually scanning for shapes that could match a very broken character; in this case the gallery view without the contextual information is more useful. These tables are more comprehensive and more accurate than the handwritten character charts produced in the past;[10] we will discuss this topic in greater detail below. But first, we demonstrate the utility of the OCHRE reconstruction tools for testing and proposing textual reconstructions.

The OCHRE Reconstruction Tools

Before a text critic can make an argument about the evidence of a particular manuscript, she must first read the manuscript. In some cases, however, the heavily damaged state of the characters challenges this seemingly basic endeavor. When the reading of a manuscript is in doubt, the most methodologically rigorous way of testing reconstructions is by overlaying characters in the scribe's own handwriting onto an image of the damaged text. Performed on paper, this task is cumbersome and imprecise to the point of being impractical, but the reconstruction tools in OCHRE enable the scholar to do this work with a high degree of accuracy and control.

We will demonstrate the reconstruction tools on a particularly broken fragment of 4Q52 (4QSam^{b}), numbered frg. 2 in *DJD XVII*;[11] see Figure 4. The *DJD* editors claim that this frag-

10 See, for example, Ada Yardeni, *The Book of Hebrew Script: History, Palaeography, Script Styles, Calligraphy & Design* (Jerusalem: Carta, 1997); and Frank Moore Cross, *Leaves From an Epigrapher's Notebook: Collected Papers in Hebrew and West Semitic Palaeography and Epigraphy* (Winona Lake, IN: Eisenbrauns, 2003).

11 Frank Moore Cross, Donald W. Parry, Richard J. Saley, and Eugene Ulrich, eds., *Discoveries in the Judaean Desert XVII: Qumran Cave 4 XII 1–2 Samuel* (Oxford: Clarendon Press, 2005).

ment attests the Hebrew *Vorlage* of the Septuagint plus in 1 Sam 14:41.[12] Given the text-critical significance of this plus, the authenticity of which has been widely contested, a close investigation of the paleographical evidence from the fragment is worthwhile. While we will refrain from presenting an argument on the reading of the fragment in order to maintain focus on the digital tools, we will demonstrate how the tools can be used to construct and support such an argument.

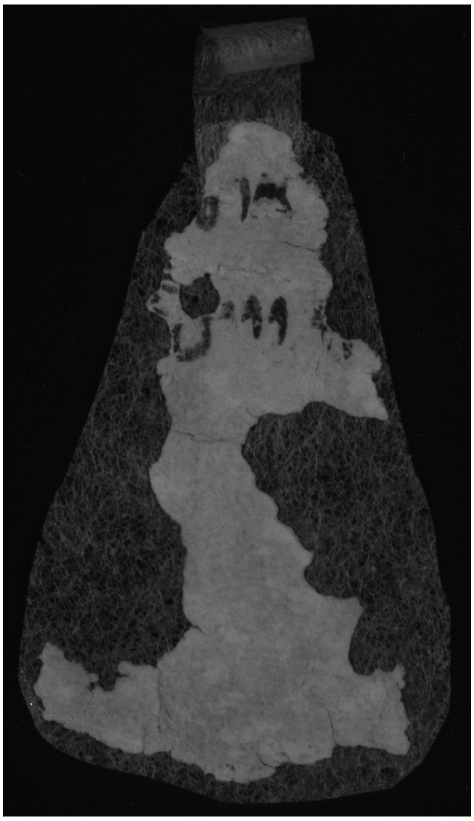

Figure 4. 4Q52 (4QSam^{b}) frg. 2. Image is used courtesy of The Leon Levy Dead Sea Scrolls Digital Library, Israel Antiquities Authority; photo: Shai Halevi.

Let us imagine that a scholar wants to test the reading of 4Q52 frg. 2 found in *DJD XVII* and perhaps propose a different reading. OCHRE can display an image of the fragment along with a library of all the hotspots, cut out from the current image, shown as a scrolling list on the right side of the screen; see Figure 5. A manuscript with extensive textual content is likely to have a good representation of hotspot exemplars. A small damaged fragment such as this one, on the other hand, may have only a few exemplars that may not include all characters in the writing system. The researcher can choose to add hotspots to the library of exemplars in two ways. The "Add" button adds all hotspots from whichever images the researcher specifies. Typically, one would choose images of fragments that are more extensively preserved, in the expectation that one could find a good "fit" for the character in question. Alternatively, the "More" button adds up to 50 instances of a selected character.

The list of hotspots can be filtered by letter value so that only instances of the desired letter are displayed. As shown below (Figure 5), hotspots can be dragged and dropped from the library of cutouts directly onto the image of the damaged manuscript, and moved around using the mouse, to test their fit relative to the visible ink or the available space. Here we have chosen a series of *aleph's* to compare with the first character in line 1. We chose this character for demonstration because the reading is unlikely to cause controversy, but the same process we will show could also be used to test the characters that are more difficult to identify.

[12] *DJD XVII*, 224–25.

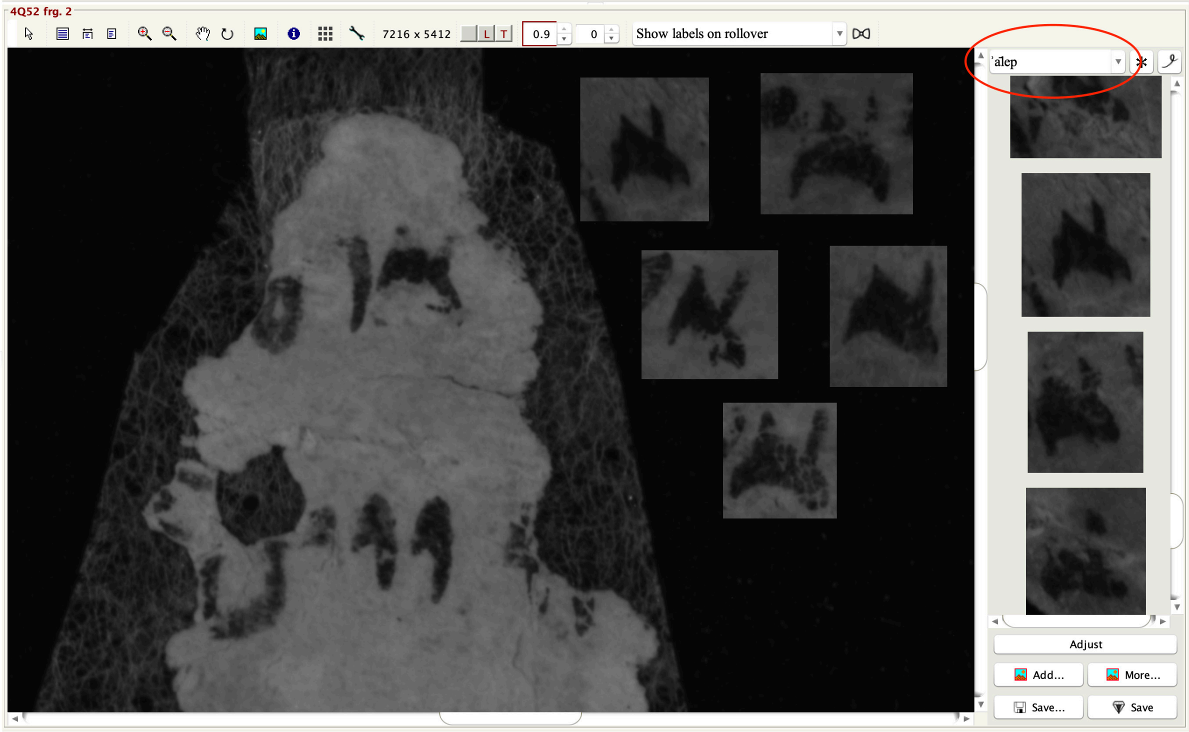

Figure 5. The list of hotspots from images of 4Q52 has been filtered to show only instances of *aleph*; see top right of the screenshot. Several hotspots from the list have been dragged and dropped onto the image of 4Q52 frg. 2. The image of 4Q52 is used courtesy of The Leon Levy Dead Sea Scrolls Digital Library, Israel Antiquities Authority; photo: Shai Halevi.

OCHRE's reconstruction tools allow the researcher to adjust these hotspots in order to compare them more precisely with the image or to highlight a reconstruction.[13] First, however, the character must be extracted from the background. In OCHRE, this is a two-step process. In the first step, OCHRE uses an image processing technique called *masking*. When the mask is applied to the original image, pixels that fall within the polygon are retained, while pixels that fall outside the polygon are dropped. Figure 6 demonstrates the process from left to right: original cutout area, then the negative and positive masked images.

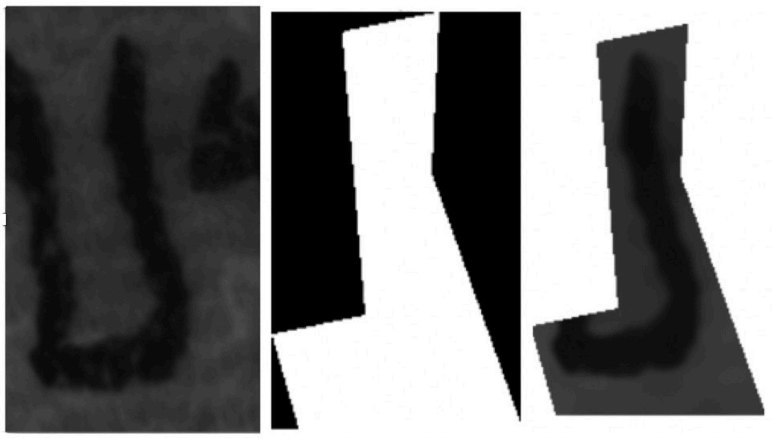

Figure 6. Masking technique produces a useable letter form.

[13] These digital tools, which will be described in greater detail below, were inspired by common image processing strategies that provide powerful options for the manipulation of the hotspot cutout graphics. OCHRE's tools use the Java version of OpenCV (opencv.org), an open source library widely used for computer vision (CV) applications, distributed under the 3-clause BSD license.

It is common for signs on a manuscript to overlap in ways such that a neighboring sign will infringe on the extent of another. Helpfully, OCHRE gives researchers the ability to draw precise polygon-shaped hotspots (instead of only rectangular ones) to outline the meaningful boundary of the sign, carefully avoiding neighboring intrusions. In the following example (Figure 7), polygon hotspots were used to delimit the bounds of three signs in close proximity.

Figure 7. Polygon hotspots are useful for creating accurate coverage of the sign.

In the next step, visual "noise"— minor variations that are not paleographically meaningful (e.g., unevenness in the texture of the writing surface)—is minimized. *Denoising* is a technical process that, in effect, averages out the variations in the image, attempting to minimize noise while retaining the essential features. Experimentation with filters of different sizes shows that, indeed, the denoising process smooths out minor variations in preparation for the next step.

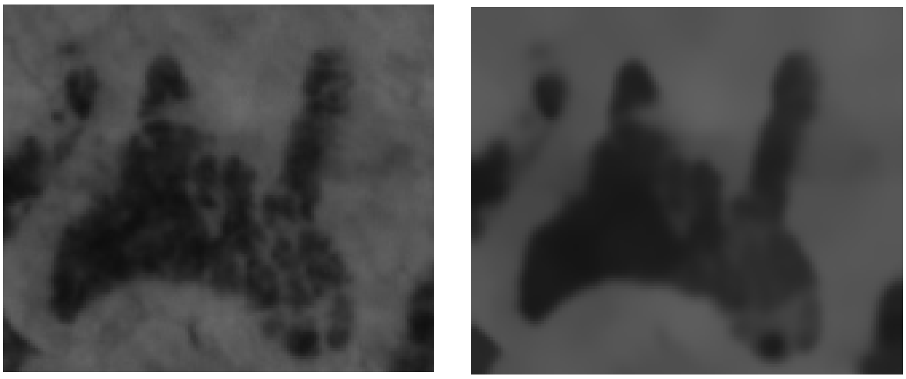

Figure 8. Applying a denoising filter (value of 12) to a hotspot cutout (right) results in a much better exemplar than the original image (left).

After the denoising process, the image background is removed, leaving a high-contrast black exemplar of the letter. The image processing technique used for separating the foreground of an image from the background is called *thresholding*. Notice how the image on the right side of Figure 9 (below) results in a much cleaner thresholded image having first been denoised.

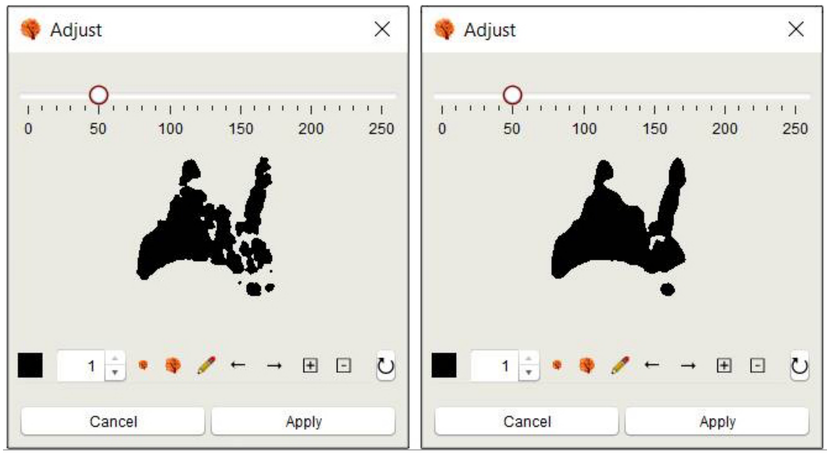

Figure 9. Original image shown thresholded (left) compared to the denoised image thresholded (right).

At this point, the researcher can begin to manually alter the thresholded exemplar. First, in cases where it is difficult to automatically identify the edges of a character, the researcher can use a slider to increase or decrease the threshold at which the computer identifies ink versus background. In the following sequence of images (Figure 10), we adjust the threshold of our *aleph*.

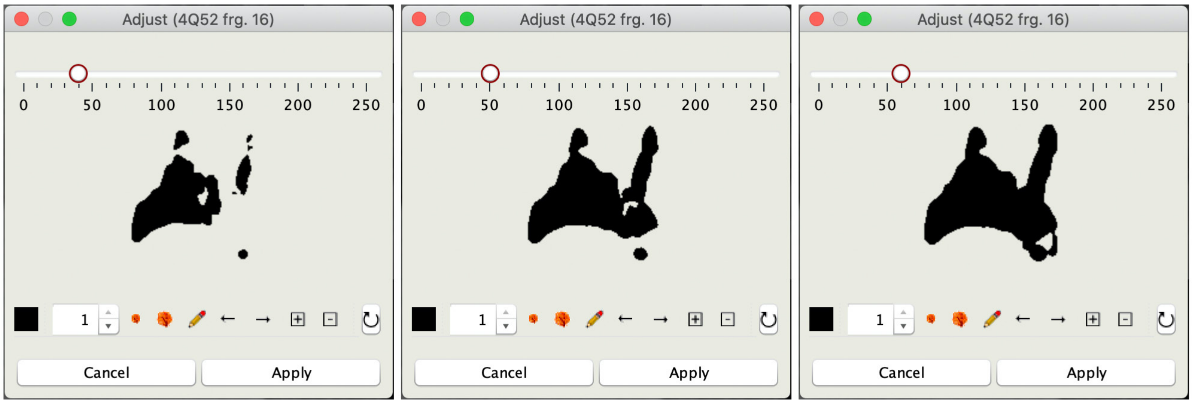

Figure 10. The threshold for the extraction of a character can be adjusted. Here the same *aleph* is shown extracted at three different thresholds.

When the threshold is increased, unwanted pixels may be added from elsewhere in the frame. The eraser tool allows the researcher to refine the extracted character by removing pixels from specific areas. The pencil tool can also be used to add pixels that thresholding does not pick up.

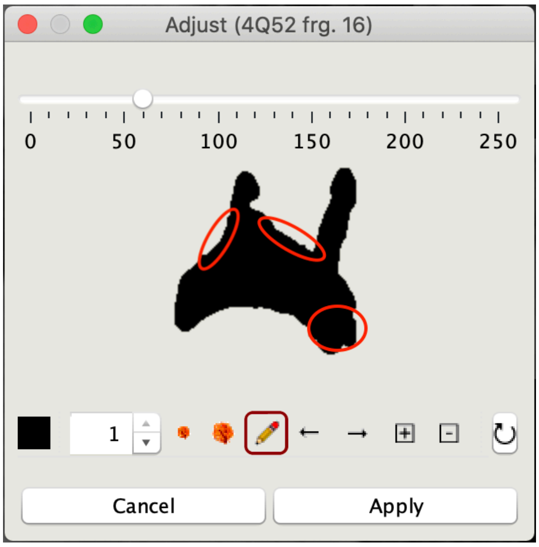

Figure 11. The pencil can add pixels to the extracted character, and erasers in two sizes can be used to remove unwanted pixels. Here the pencil and eraser tools have been used to touch up the circled areas.

The researcher may also wish to adjust the rotation of the hotspot to make it match more closely with the stance or orientation of a partial letter on the image. The adjustment tools include left and right rotation buttons (Figure 12).

Figure 12. Arrow buttons rotate the character right or left by a few degrees with each click. The *aleph* used in this demonstration did not require rotation.

Finally, the exemplar can be colored and its opacity can be adjusted (Figure 13). Bright colors such as the red we used make the exemplar stand out from the original. Multiple colors could be used for purposes such as indicating the confidence of the reading or recording different attempts at reconstruction.

Figure 13. The extracted character has been colored red, and its opacity has been adjusted to 50%.

When the scholar is satisfied with these many adjustments, the exemplar is applied to the damaged manuscript and can be placed directly over the character to be tested (see Figure 14). Again, we intentionally chose to demonstrate these techniques on a character whose identification is not in question; however, the process of selecting a hotspot, adjusting its appearance, and overlaying it directly on the image of the fragment could be used productively to test identifications of the more broken characters. At any stage in the process, the reconstruction can be saved as an image file that can be added to a publication in support of the scholarly argument.[14]

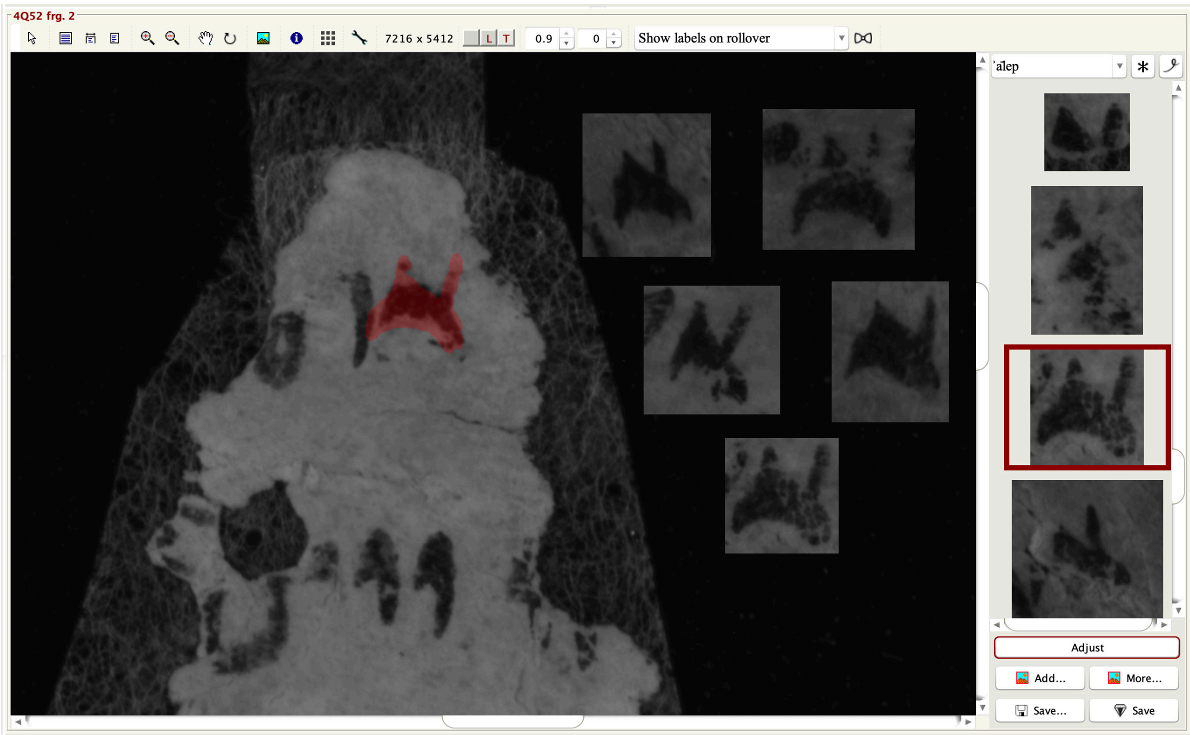

Figure 14. The adjusted hotspot has been placed directly on top of a character for testing. The image of 4Q52 is used courtesy of The Leon Levy Dead Sea Scrolls Digital Library, Israel Antiquities Authority; photo: Shai Halevi.

The ability to layer examples of the scribe's own writing directly over a damaged manuscript is an enormous advance in the study of ancient manuscripts. Rather than attempting to estimate visually or mathematically whether a character would fit a given location, the scholar can test proposed readings in a controlled way. Note the extent to which scholarly judgment is still required: the software does not match characters or suggest reconstructions—and this is a strength of the system. The tools available in OCHRE do not reduce the skill or agency of the researcher in any way. They simply enable her to do her work with greater precision and efficiency.

14 Files are saved in PNG format to respect transparency.

Generating Dynamic Letter Charts

In the sections above, we introduced the OCHRE data model for text and the process of integrating text transcriptions using hotspotting. In this section, we introduce a process for producing digital letter charts for the Dead Sea Scrolls and for Ugaritic tablets. The outcome is different from the reconstruction tools above, but we use the same highly atomized approach to textual data.

While often thought of as a tool for dating manuscripts based on their handwriting, letter charts are also useful for text critics when assessing whether a fragment belongs to a particular manuscript and thus whether it can be used to reconstruct a particular passage. In a pre-digital context, a researcher would carefully sketch and ink letter forms, arranging them in columns for comparison. These hand-drawn charts are helpful but limited. In many cases, it is not clear to the reader if a given example is taken directly from a text or is a generalized form. Also, it is typically too cumbersome for the researcher to note the exact source of each sign. Thanks to the possibilities opened up by digital technology, paleography now has an opportunity to evolve. We will demonstrate the creation of two kinds of dynamic letter charts. By dynamic, we mean that the contents of the letter charts are generated on-the-fly from the database and can be altered using different query criteria. We also mean that the end result is interactive: clicking on a letter form reveals the character in context on the original image and allows the researcher to navigate to the specific character in the text transcription.

The first kind of letter chart uses the same polygonal hotspots introduced above. In addition to searching for hotspots within images of one manuscript, as described above, the researcher can also search for hotspots across images of multiple manuscripts, which can be displayed in the same kinds of tables shown in Figure 2 and Figure 3 above. The researcher can query for all instances of a character, yielding the maximum amount of paleographic evidence from the selected manuscripts; or, if this quantity of data proves unwieldy, the researcher can restrict the query to only those characters marked as of interest.[15] This property can be assigned to characters during the transcription process or after the transcription has been completed. A table such as the one in Figure 15 would help a researcher determine whether a very broken fragment has been properly identified or perhaps belongs to a different scroll.

15 Because OCHRE is a semi-structured XML graph database, the primary query language is XQuery. Thankfully, OCHRE users do not need to know how to compose an XQuery. OCHRE includes a user-friendly graphical user interface that allows the user to make a series of selections using pick lists, all of which generates a complex XQuery expression to extract data from the backend database.

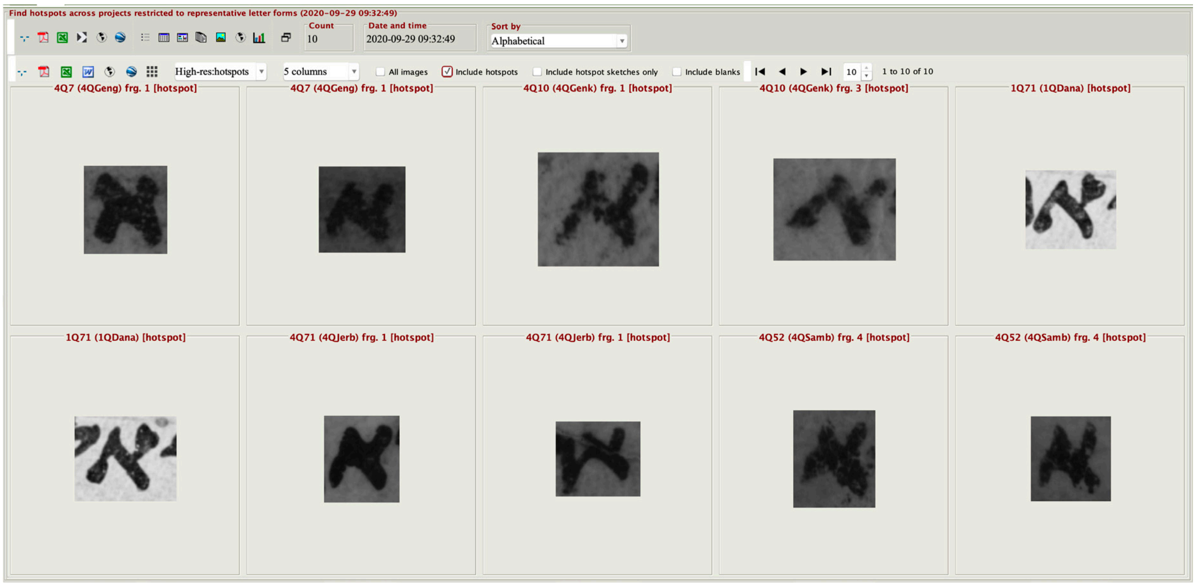

Figure 15. This table displays examples of aleph from five different Dead Sea Scrolls: 4Q7 (4QGen[b]), 4Q10 (4QGen[k]), 1Q71 (1QDan[a]), 4Q71 (4Jer[b]), and 4Q52 (4QSam[b]). These particular characters have been tagged as being of interest, allowing the researcher to restrict the number of results returned in a query.

By clicking on any of the letter forms, the researcher is able to inspect the image source (see Figure 16). Another click brings up the researcher's transcription of the manuscript. This feature removes the barrier that typically exists between a letter chart and the source of any specific form: because of their format, letter charts frequently do not list where forms are found, and when they do, images of the manuscript are often located in a separate volume or are not available at all.

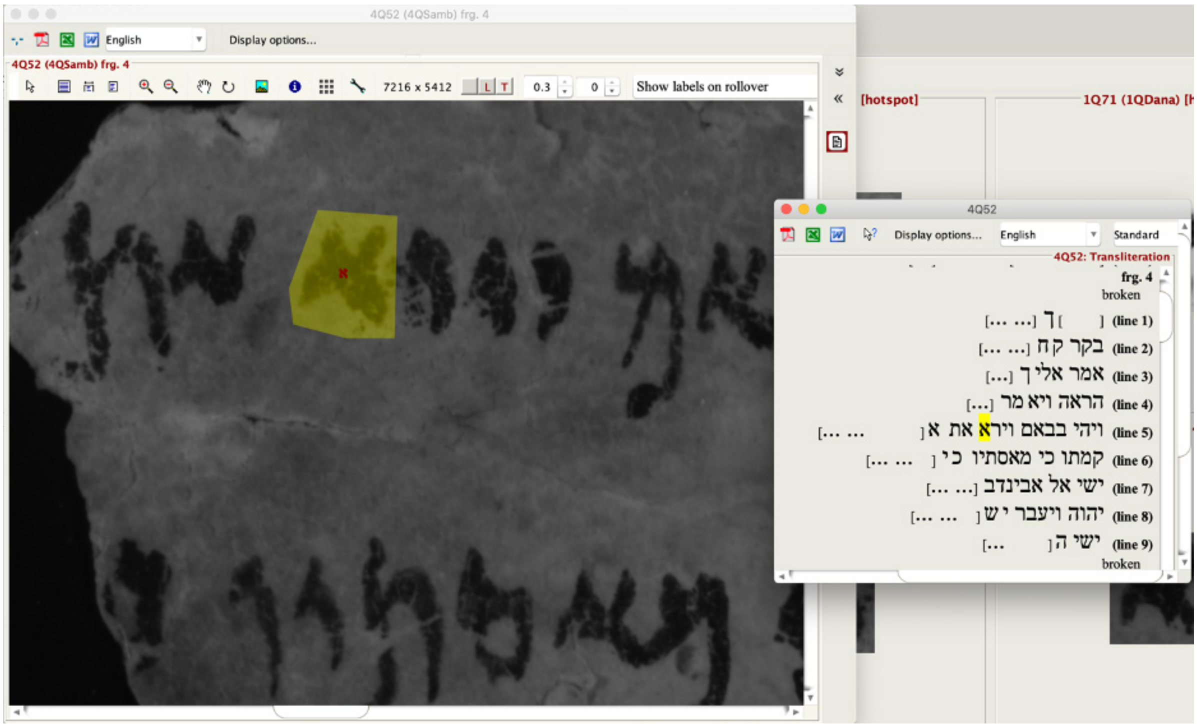

Figure 16. Clicking on a letter form in the table opens a popup window with a high-resolution image of the manuscript so the character can be viewed in context. Clicking on the text icon in the popup window brings up the transcription of the manuscript. The image of 4Q52 is used courtesy of The Leon Levy Dead Sea Scrolls Digital Library, Israel Antiquities Authority; photo: Shai Halevi.

The second kind of letter chart uses a different kind of hotspotting, originally designed for cuneiform. Instead of drawing a polygon shape, the researcher traces the shape or outline of a character. The end result is a hotspot that consists of a series of linear elements. The series of lines are saved as a single hotspot and then associated with a letter from the text transliteration. This style of hotspot is a natural fit for cuneiform (see Figure 17) and will be demonstrated here on a cuneiform corpus taken from the RSTI project in OCHRE, but it could also be productively used for biblical manuscripts.

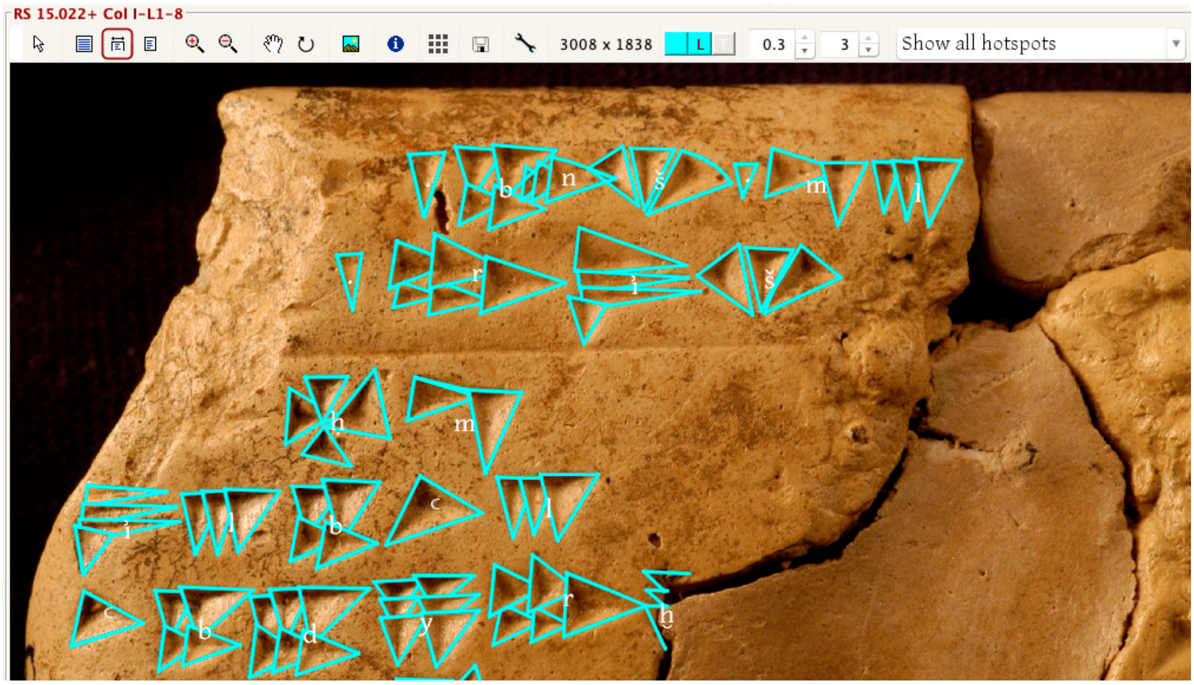

Figure 17. RS 15.0222+, from the Ugaritic tablets found at Ras Shamra, has been hotspotted with linear hotspots, shown in bright blue to make them more visible. In several cases, multiple wedges are associated with a single transliterated letter.

In cuneiform studies, there are various approaches to representing a two-dimensional drawing of a three-dimensional cuneiform wedge. In early publications of the Ugaritic texts, wedges are represented only as a head and tail (left example in Figure 18 below). Some researchers attempt to represent the internal structure of the wedge (middle example below). Others simply represent the outline of each wedge (right example below). In RSTI, wedges are drawn as outlines, but the choice is one of personal preference. OCHRE could support any of these approaches.

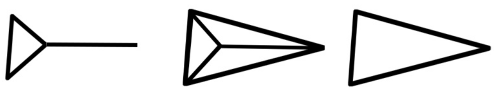

Figure 18. Three approaches to representing cuneiform wedges.

Because vector graphics scale more effectively than raster graphics, these linear hotspots can be viewed in a more traditionally formatted chart. Figure 19 represents a portion of a dynamic letter chart displaying only those characters the researcher has marked as representative. The columns in this chart have been configured to represent textual genres, but they could also represent date of composition, tablet findspot, or any other properties the researcher wishes to assign. This flexibility is enabled by OCHRE's model of data integration: any information the researcher records about the text is linked to the epigraphic analysis of the text and to the associated hotspots.

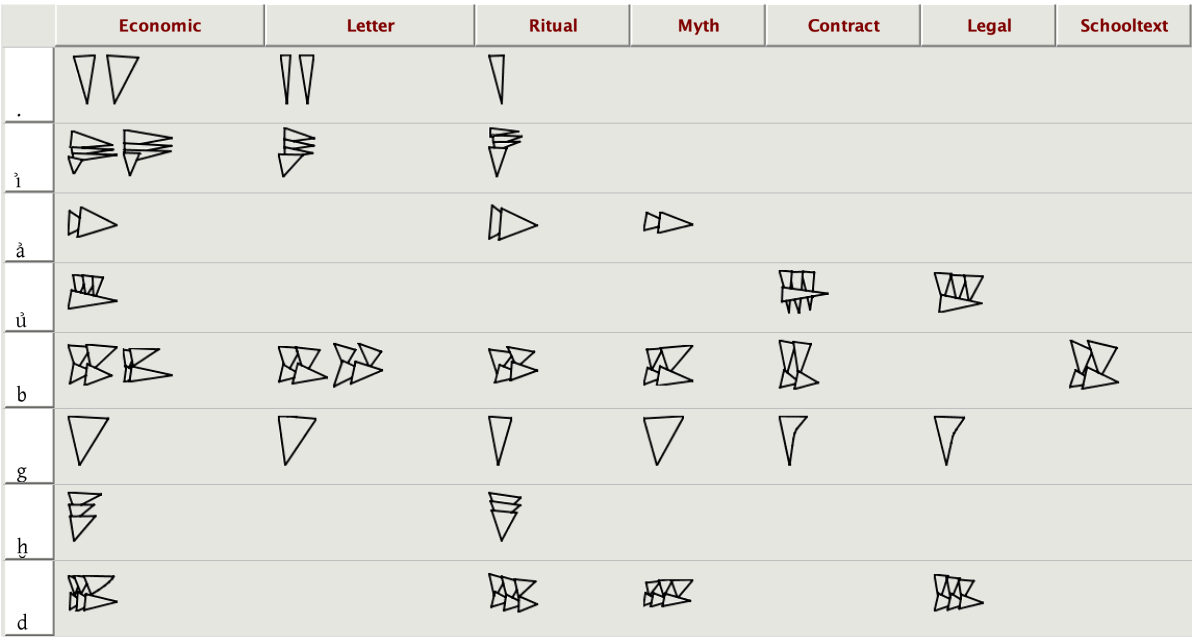

Figure 19. Dynamic letter chart of Ugaritic sign forms.

As in the Hebrew examples above, the Ugaritic characters can also be viewed as a gallery. Figure 20 below shows the results of a query of all hotspots from a specific text, with the hotspot outlines displayed against the source image. Figure 21 presents the hotspot outlines independent from the source image.

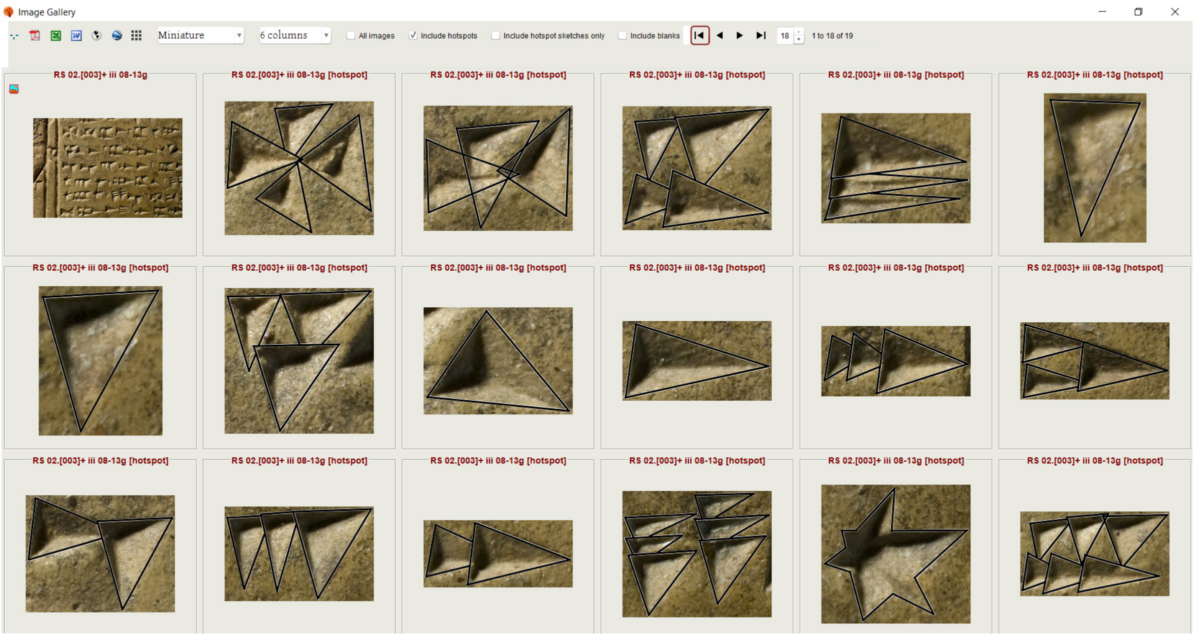

Figure 20. Gallery view of Ugaritic hotspots.

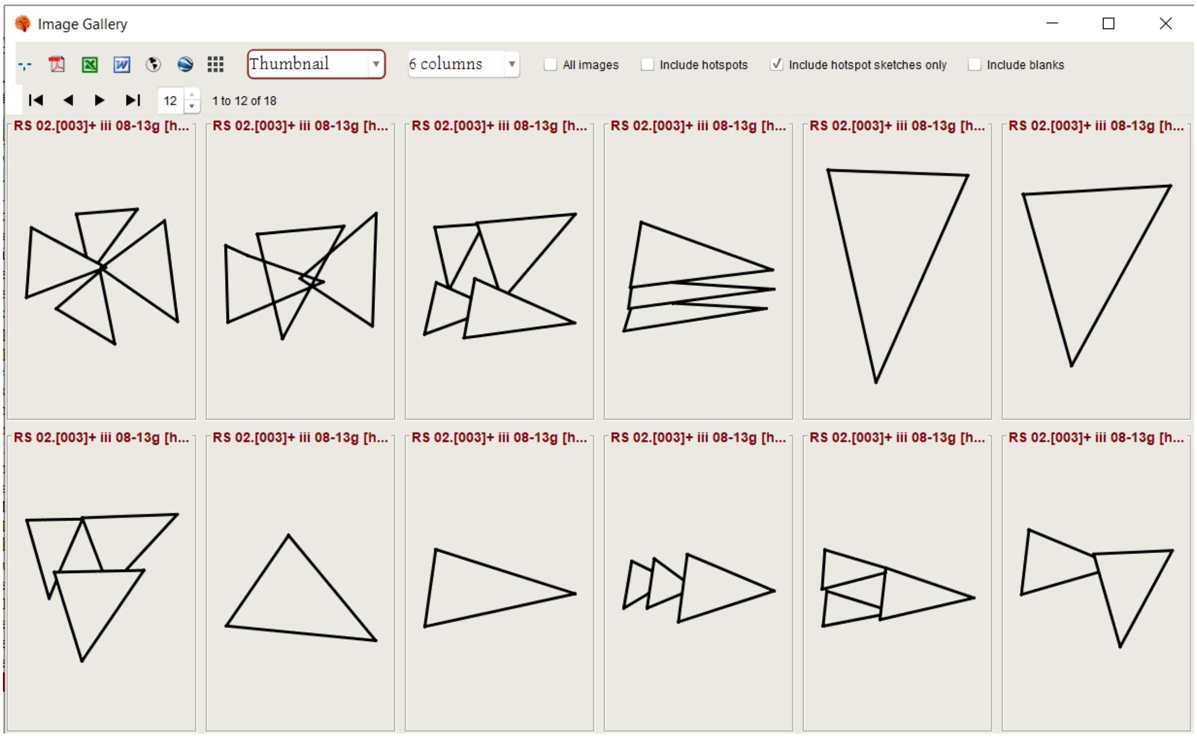

Figure 21. Ugaritic hotspots displayed as linear letter forms.

Conclusion

The digital tools for paleography that have been built in the OCHRE database platform provide an accurate, controlled, and methodologically rigorous way for scholars to reconstruct damaged manuscripts.[16] The biblical Dead Sea Scrolls, with their (sometimes extensive) physical damage and high text-critical value, make these tools relevant for scholars studying the history and development of the biblical text. Textual critics should be reading the Dead Sea Scrolls from (images of) the fragments themselves and performing their own reconstructions when necessary.

We wish to emphasize that this paper is not meant to be a technical manual for how to perform paleographic analysis in OCHRE. Instead, we hope to have demonstrated the value of two computational concepts: atomization and integration. It is only because the CEDAR and RSTI projects use OCHRE's highly granular textual data model that the reconstruction and letter chart tools are possible. In OCHRE, rather than storing a text edition as a single document, each letter is a separately addressable database item. These items can be linked to polygons or linear hotspots on a digital image, creating a uniquely and powerfully integrated dataset. The high degree of integration allows for the live and dynamic browsing or searching of textual content linked to the digital image.

[16] These tools are available on the OCHRE platform. For more information about OCHRE, contact the OCHRE Data Service at ochre@uchicago.edu.

Revision history
18 February 2021: published online
19 February 2021: updated Figure 16 with the correct image