THE CHRONICLE August 2026 (No. 291)

The Chronicle of the U.S. Classic Postal Issues

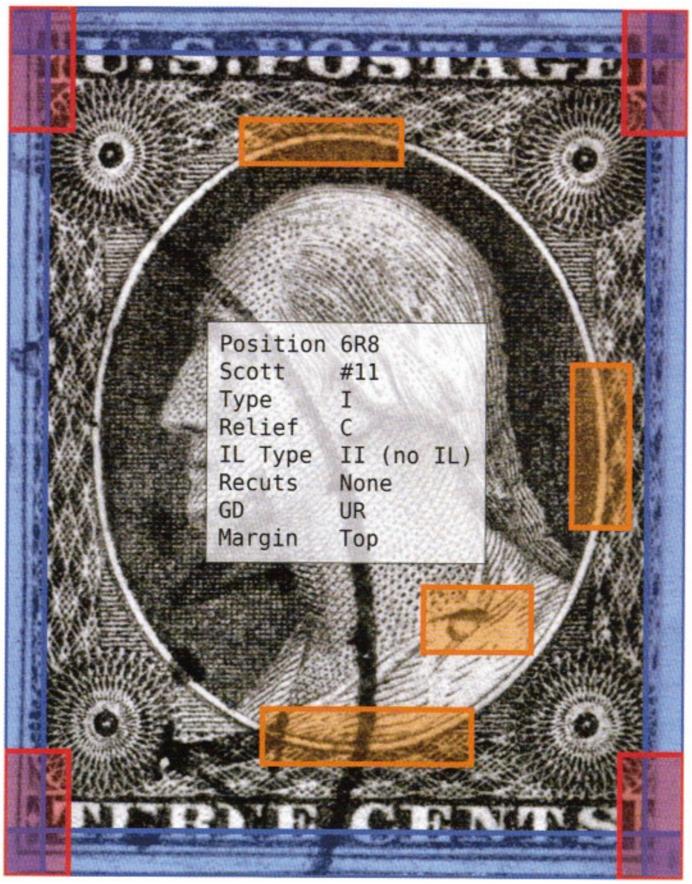

Artificial Intelligence comes to philately: Image from PlateAI, a free software application created using Anthropic's Claude. Upload a scan of an imperforate 3¢ 1851-57 stamp and PlateAI tells you the likely plate position. Shaded areas show the diagnostic regions analyzed. Developer David Fussichen, RA 6028, provides details in our 1851 section.

August 2026 Volume 78, No. 3 Whole No. 291

THE EDITOR'S PAGE

MICHAEL LAURENCE

IN THIS ISSUE: BACK TO OUR ROOTS

The *Chronicle* originated almost 80 years ago as a study-group newsletter—an occasional set of mimeographed pages distributed to platers of the 3¢ imperforate Washington stamps. *Chronicle* 1, 13 pages printed on one side only and neatly stapled in the upper left corner, bore a publication date of July 25, 1948. Its title was *3¢ 51-57 Chronicle.*

I wonder how the 20-odd initial members of this group would react to the *Chronicle* you now hold in your hands. An article in our 1851 section (page 271), featured on the cover of this issue, tells how member David Fussichen used Claude, the artificial intelligence chatbot created by the Anthropic firm, to develop a website that plates the imperforate 3¢ 1851 Washington stamps. How Fussichen accomplished this feat makes an engaging and informative story, and the plating software itself is remarkable. The software is available free to anyone who wants to use it. Take a scan of any 3¢ imperforate 1851 stamp in your collection, upload it into Plate AI (which is what the device is called) and you almost immediately are shown the most likely plate position—along with images of known plated stamps from that position to enable comparison. As the author makes clear, the possibilities suggested by this software are staggering. We've come a long way since 1948.

At the recent Boston international stamp show, Steven Walske, our Western Mails section editor, won not one but two of our Society's trophies (the Chase Cup and the McDonald Award) for his recent contributions to the literature of our hobby. Not one to rest on his laurels, Walske contributes to this issue (page 235) a well-researched essay showing the difficulties of correspondence between early settlers in the Oregon Territory and friends and family back east. Walske weaves the players, their journal entries and some of the covers they left us, into a story that brings to life the challenges of intercoastal letter correspondence in pioneer days. This is postal history at its best.

Also highly recommended in this issue is Ardy Callender's article in our Foreign Mail section (page 309), exploring the beginnings of New York Foreign Mail markings. Callender uses an understanding of the architecture of the New York Post Office, gleaned from contemporary illustrations and newspaper accounts, to reconstruct how mail was handled and marked by New York's Foreign Mail division at the dawn of the stamp era.

In our 1847 section, Steven Pacetti makes good use of our Society's on-line database of 1847 cover information. Within that searchable listing of almost 16,000 covers, Pacetti found 27 covers on black-bordered envelopes signifying a death announcement. "1847 Stamps on Mourning Envelopes" (page 266) provides a chronological listing of the covers, illustrates a few of the most striking, and offers some helpful generalizations about these solemn artifacts, whose use began during the 1847 era.

Cover information of a different variety is provided in two other articles in this issue. In our 1861 section (page 288), German member Burkhard Krumm provides a census of 10¢ 1861 Type I covers to foreign destinations, 99 covers in all. When added to covers to domestic destinations (listed and discussed by Krumm in *Chronicle* 282 and modestly updated here) the total recorded population of these covers is 177. And in our Bank Note

(*IN THIS ISSUE concluded on page 243*)

THE 1851 PERIOD

ARTIFICIAL INTELLIGENCE COMES TO PLATING THE 3¢ U.S. IMPERFORATE STAMP OF 1851–57

DAVID FUSSICHEN, RA 6028

Introduction: new tool for an old puzzle

The 3¢ Washington stamp of 1851-57 has fascinated collectors for well over a century. Produced from 13 engraved plates with 200 subtly unique positions each, the issue contains 2,600 collectible varieties. Identifying the plate position of each stamp, a process known as plating, is one of the keys to understanding where your stamp fits in the broader story of the issue, and thus where it fits in your collection. But plating this stamp comes with a very steep learning curve. Since Dr. Carroll Chase first studied the issue systematically in the 1920s, collectors have developed better references and better tools for the task. The latest tool is PlateAI, a free machine-learning tool available at stampplating.com.

With PlateAI, a collector can plate a 3¢ Washington stamp in a minute or two, with accuracy matching that of experienced platers.

Applying machine learning to plating is not a new idea, but until recently it would have required specialized expertise, substantial computing budgets and a team of engineers. The technical barriers fell dramatically in the last year. Working with Anthropic's Claude on the technical side, and with guidance and encouragement from several experienced platers, I was able to turn a longstanding idea into a working tool. This article discusses both the tool itself and the way it was built. The story is interesting on its own, and it might serve as a template for other AI-enabled stamp identification applications.

What is plating and why this stamp?

The 3¢ Washington of 1851-57 was the first high-volume postage stamp of the United States. According to Chase, about 362 million imperforate stamps were printed over the seven-year life of the issue, which makes surviving examples both plentiful and affordable compared to other classic stamps. They also varied remarkably. Because of the relatively primitive technology of the era, no two of the 2,600 positions are identical—and that's just the underlying stamp itself. Add cancellations, color and usages, and the collecting possibilities multiply toward the infinite.

These stamps connect directly to antebellum America, a time when steamboats plied the Mississippi, when railroads, steamships, and stagecoaches connected east to west, and when transatlantic travel became routine. The stamps were issued in a country still half slave and half free, and they were demonetized in the opening months of the Civil War. Every 3¢ Washington stamp in every collection is a direct link to that era. Plating the stamp ties it even more closely to a specific time and place.

The printer, Toppan, Carpenter, Casilear & Co. of Philadelphia, engraved the stamp image on a single die. From that die, they created a transfer roll with three impressions, known as reliefs: A, B, and C. Reliefs are distinguished by subtle but consistent differences in the medallion oval surrounding Washington's head, and by the presence of a "gash on the shoulder" on the less-common C reliefs. These differences originated in clean-up work performed by engravers directly on the roll.

The transfer roll was then used to lay down 100 positions on each of two panes (Left and Right), giving 200 stamps per plate. Across the seven years of production, the printer made nine physical plates, three of which were reworked one or more times to produce additional plate states. Counting the states as if each were a separate plate, the issue is conventionally described as coming from 13 plates: 0, 1E (Early), 1i (Intermediate), 1L (Late), 2E, 2L, 3, 4, 5E, 5L, 6, 7 and 8.

Plates 0, 1E, 1i, 2E and 5E produced the "Orange Brown" stamps catalogued as Scott 10 and 10A. The remaining plates produced the "Dull Red" stamps catalogued as Scott 11 and 11A. The Dull Red label covers a wide range of colors that aren't Orange Brown, and color alone is not a reliable way to distinguish the two groups. Plating is the only reliable way. Many collectors hold 11 or 11A stamps believing them to be the rarer 10 or 10A, and many dealers try to sell them the same way. A collector who can plate avoids this mistake.

Beyond Orange Brown and Dull Red, many collectors enjoy the wide range of colors found in the 3¢ Washington. Because each plate and time period produced its own range of colors, plating and color study are intertwined. To work seriously with color, stamps must be plated.

The transfer roll was positioned on each plate using small guide dots impressed into the plate. The guide dots appear in similar but not identical locations near the corners of about half of the 2,600 positions, and their placement is a useful diagnostic.

Once the impressions were laid down, an engraver hand-retouched every one of the 2,600 positions. That retouching is the main reason each position is unique. All 2,600 had at least three, and usually all four, of the outer frame lines recut. The issue is categorized into two types based on the inner frame lines. Type I (Scott 10 and 11) did not have the inner lines recut. Type II (Scott 10A and 11A) did.

Many positions across the 13 plates carry additional recutting beyond the frame lines, most often near the corners in the triangles or diamond blocks. Each plate has a typical recutting style that often allows a plater to identify the plate from which a stamp comes. Even so, the most experienced platers can rarely differentiate between plates 1E and 1i, plates 2L and 3, or plates 6 and 7 by eye alone.

To plate a stamp, then, is to inspect the combination of features: the relief, the guide dots, the recuts, along with secondary clues such as spacing between adjoining stamps or even clues from the cancellation. The goal is to narrow 2,600 candidates toward one. Once a plater has the shortest list possible, the work becomes a side-by-side comparison of the patient stamp, as the unidentified copy is known, against reference images of the candidate positions. Different platers favor different workflows, but the underlying process is the same.

Platers, and PlateAI, use a digital version of an old technique Chase described as "foreshortening." Holding the stamp at a low angle to the eye lets the plater see small bends and curves in the recut framelines that are not obvious head-on. The digital equivalent, called a "compression scan," compresses the long edge of a frameline to ten percent of its width or height, making bends and curves easy to see.

Readers who want to go deeper into reliefs, plate layout, or the production history of the stamp should consult *The 1851 Issue of United States Stamps: A Sesquicentennial Retrospective* by Richard C. Celler and Elliot Omiya (published by the Classics Society in 2006), which remains the standard reference. Most of what is summarized above is treated there in much greater depth.

Plating is slow work and it takes years to develop fluency. A plated 3¢ Washington stamp pins a specific piece of paper to a specific position on a specific plate. Combined with cancellation and color, the plate identification narrows the date range when the stamp was used. It becomes more than a stamp. It is an artifact with a traceable origin.

The shoulders we stand on

Plating the 3¢ 1851 stamp has always been a collective effort. Every generation has built on the work of the one before. The work began with Chase. His foundational study, first published in 1929 and revised in 1942 as *The 3¢ Stamp of the United States 1851–57 Issue*, established the vocabulary, the methodology, and the basic plating data still in use today. Almost everything that has been since written about this stamp references Chase.

Additional contributions came from William K. McDaniel, whose articles in the *Chronicle* from 1973 through 2006 systematically corrected and extended Chase's recut analysis. Celler and Omiya documented relief characteristics in the *Sesquicentennial Retrospective*, the standard reference just cited.

Articles and books remain academic without reference material. To plate a stamp, a collector needs a complete set of references covering all 2,600 positions. Before flatbed scanners and high-resolution monitors, platers needed either the stamps themselves or printed high-resolution photographs of Chase's plate reconstruction. Today, plating work is done with scans enlarged on monitors rather than with magnifying glasses.

Chase himself assembled the first complete plating. His black-and-white photographic record now resides at the Smithsonian's National Postal Museum. Collecting an example of each position became a challenging objective for subsequent collectors.

J. Bryan O'Doherty (RA 4569) assembled and published scans from four complete platings. This formed the backbone of PlateAI's training data. For convenience in discussion, the four platings have been designated as follows: **Chase** = Chase's original black-and-white plating; **LAO** = Robin Lund's plating, supplemented by partial platings from O'Doherty and Wilbur F. Amonette—collectively "Lund-Amonette-O'Doherty" or "LAO;" **Celler** = Richard C. Celler's plating, via Robert J. Lampert; and **Amonette** = Amonette's plating which, after many years, resurfaced in a December 2025 Robert A. Siegel auction and was acquired by O'Doherty.

PlateAI was trained against scans of all four platings, plus another 10,000 plated stamp images from less comprehensive collections. Per-plate coverage and the data pipeline behind it are described in detail below, in the section on training data.

O'Doherty's StampPlating.com, launched in 2018 and fully described by O'Doherty in *Chronicle* 260, deserves particular mention. It was the first project to bring the entire body of plating work into the digital age, and it remains the most comprehensive online reference for the 3¢ 1851. A great deal of what PlateAI does is what the Plating Wizard on StampPlating.com does—narrowing the field of candidate positions through query and comparison—only with machine learning automating the narrowing. PlateAI resides at stampplating.com because Bryan and I share an interest in seeing this work continue.

I came to plating recently and from the side. I founded and ran a data consulting business for 20 years, but stepped away from hands-on technical work years ago. A parallel interest in classic United States stamps had been building for some time, and the 3¢

Washington stamp gradually became my focus. Plating at first looked impenetrable. The vocabulary was unfamiliar, and the work was tedious and frustrating. That changed as I got help from two people. Robert J. Lampert (RA 3316), an experienced plater and the owner of Celler's complete plating, patiently encouraged me and helped me learn to plate these stamps. O'Doherty's StampPlating.com made the reference material searchable and visual rather than buried in print. With those two resources, my learning curve was flattened enough that I did not give up and move on to other philatelic interests.

Using PlateAI

PlateAI runs in any modern web browser. There is nothing to install. It's free and available at *stampplating.com*, which can be accessed directly or through the Classics Society website at *USPCS.org*. Once open, drag a stamp scan to the browser window or use the upload button to load your stamp scan. The tool walks you through the rest.

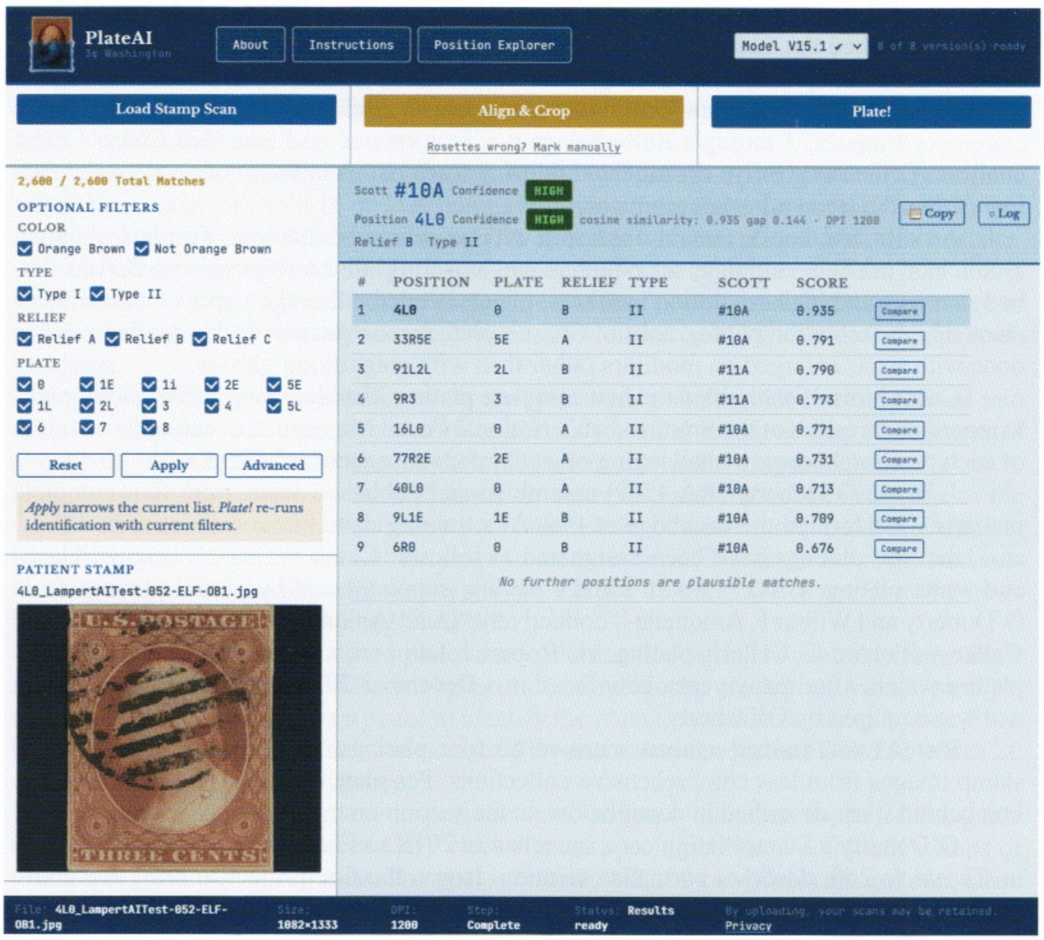

Figure 1. Screenshot of the PlateAI interface after a scan of a 3¢ Washington stamp has been loaded and identified. The top candidate is position 4L0 (Plate 0, Relief B, Type II, Scott #10A), with a cosine similarity of 0.935 and a gap of 0.144 over the second candidate — a very high confidence result. The candidate table, ordered by similarity, shows the most likely positions. The Compare button on any row opens a side-by-side view of that candidate against plated reference images. The patient stamp appears at lower left, cropped and aligned to match the reference library.

The tool handles Scott numbers 10, 10A, 11, 11A (the imperforate stamps of 1851-56) and 25, 25A (the perforated stamps of 1857). It does not handle Type III or IV stamps (Scott numbers 26 and 26A); see the Limits section below for what happens if one of those is uploaded.

For best accuracy, use a 1200 DPI flatbed scan; 2400 DPI also works, but the extra resolution provides little benefit. Lower resolutions and even cell-phone photographs are accepted but accuracy degrades with image quality.

Once a scan is loaded, PlateAI aligns it to a standard orientation by detecting the centers of the four rosettes. Automatic detection works on most stamps. When it works, four colored crosshairs appear on the rosette centers and a yellow box connects them. The user clicks Align & Crop and an aligned, cropped version of the stamp appears in the left panel. When too many rosettes are obscured by cancels for automatic detection, the user clicks "Mark Rosettes Manually' and clicks the four rosette centers in clockwise order, starting at the top left. Either way, the result is the same: a stamp aligned, cropped and ready to be plated.

After alignment, the user clicks the Plate! button. PlateAI returns a ranked list of the ten most likely plate positions, drawn from the 2,600 candidates. A screenshot of this display is shown in Figure 1. The patient stamp is shown at lower left, and likely positions are shown in the candidate table in the main field. Each row in this table shows the position identifier, the plate number, the relief type, the Scott number and a similarity score between zero and one. The higher the score, the stronger the match. In the example in Figure 1, the score of .935 indicates a high degree of confidence that the patient stamp is Position 4L0 (Position 4 from the left pane of Plate 0, Relief B, Type II, Scott #10A).

Two confidence badges appear above the results: one for the position identification and one for the Scott number. Both report HIGH or LOW, based on the score of the top candidate and the gap between it and the second candidate. HIGH confidence means the top position is very likely correct. LOW does not mean the top candidate is wrong; even on LOW results, the top candidate is the correct position over 80 percent of the time on clean scans. LOW just means further verification is warranted.

The Scott number prediction draws on all ten candidates rather than just the top one, and it is often HIGH even when position confidence is LOW.

PlateAI will correctly plate a 1200 DPI scan of a four-margin stamp about 97 percent of the time, meaning the top predicted position is a correct match. When it is not the top match, it is very often the second or third match. It would be unusual if the correct match were not in the top ten.

In every case, and regardless of the confidence scores or labels, the user should use the additional tools to confirm the position prior to finalizing.

The Compare module

The first of the additional tools is the **Compare Module**. Clicking Compare on any candidate opens a side-by-side window showing the patient stamp alongside reference images from four vetted complete platings: Chase, Amonette, Celler and LAO. All images, both patient and references, are cropped and aligned identically. Four tabs offer different views.

Corners shows the four corner regions of each stamp, where the diamond blocks and triangles appear, along with many recuts and guide dots. **L/R Framelines** shows the left and right frameline strips, displayed as ten percent compression scans alongside full uncompressed strips. Because the reference stamps have been aligned (as has the patient stamp), the compression scans in the Compare module are clean and easy to use, and there are four references to compare against instead of one. Figure 2 shows the L/R Framelines tab of the Compare module, showing the salient segments of the Position 4L0 patient stamp against reference scans from the four complete platings in PlateAI's library. The top row shows 10 percent compression scans and the lower row shows the same frameline seg-

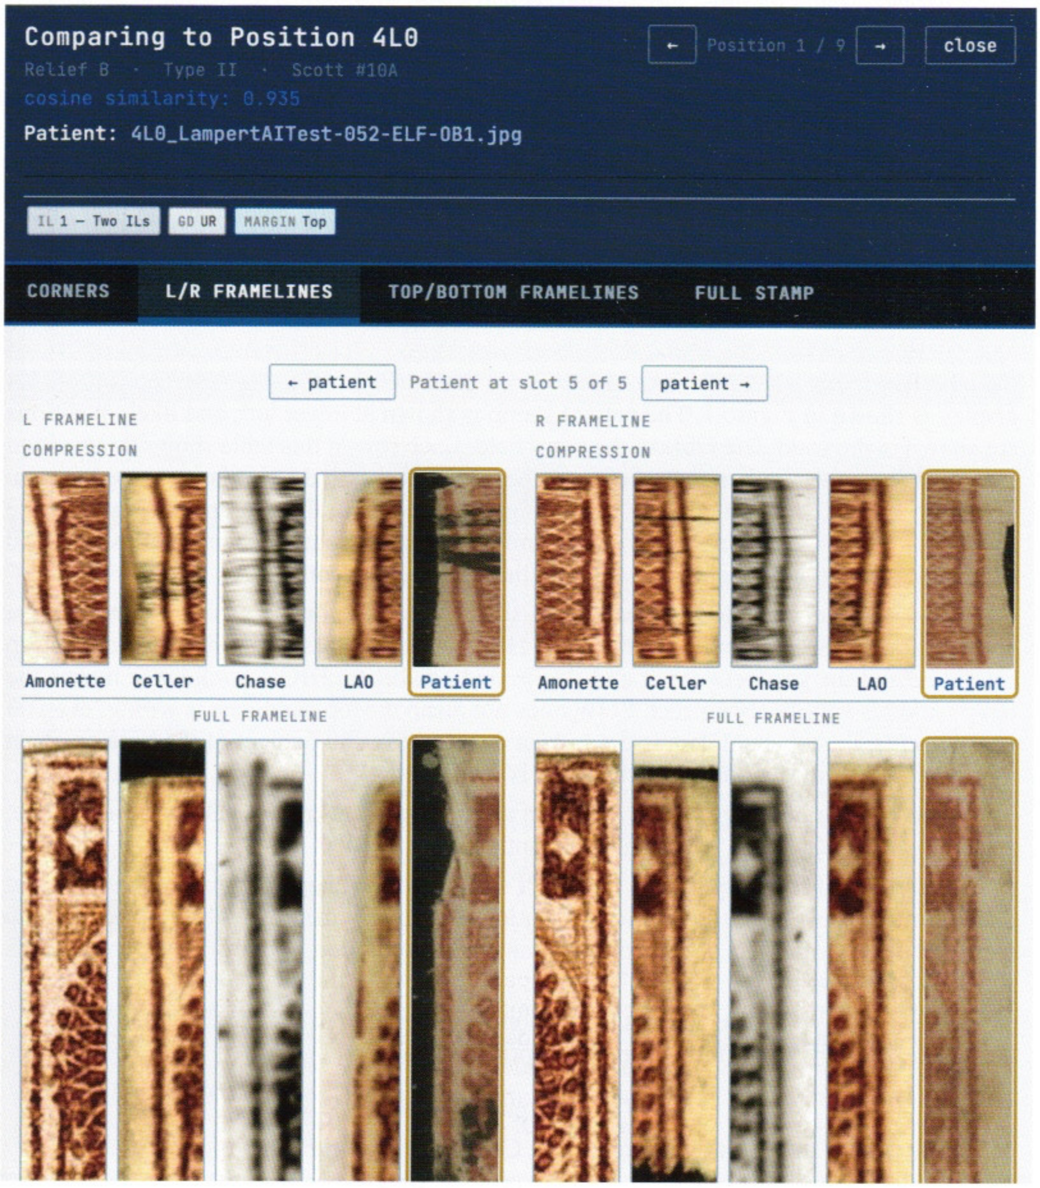

Figure 2. The L/R Framelines tab of the Compare module, showing frameline segments from the patient stamp (outlined in brown) matched against reference scans from the four complete platings in PlateAI's library. The top row shows 10% compression scans, condensing the image for easier analysis. The bottom row shows the same framelines at their natural height. The Patient image sits at the end of each row, highlighted, and can be moved to any position in the row for direct adjacent comparison. The Corners, Top/Bottom Framelines, and Full Stamp tabs at the top offer similar views of other features of the same stamps.

ments at their natural height. The patient column sits at the end of each row, framed in yellow-brown, and can be moved to any position in the row for direct adjacent comparison. The Corners, Top/Bottom Framelines, and Full Stamp tabs (on the menu bar above the images) offer three other views of the same top candidates. **Top/Bottom Framelines** shows the same for the horizontal framelines and **Full Stamp** shows the full patient stamp on the left, and one of the references on the right. The user can cycle through references from each of the four complete platings. A screenshot of the Full Stamp display is shown in Figure 3,

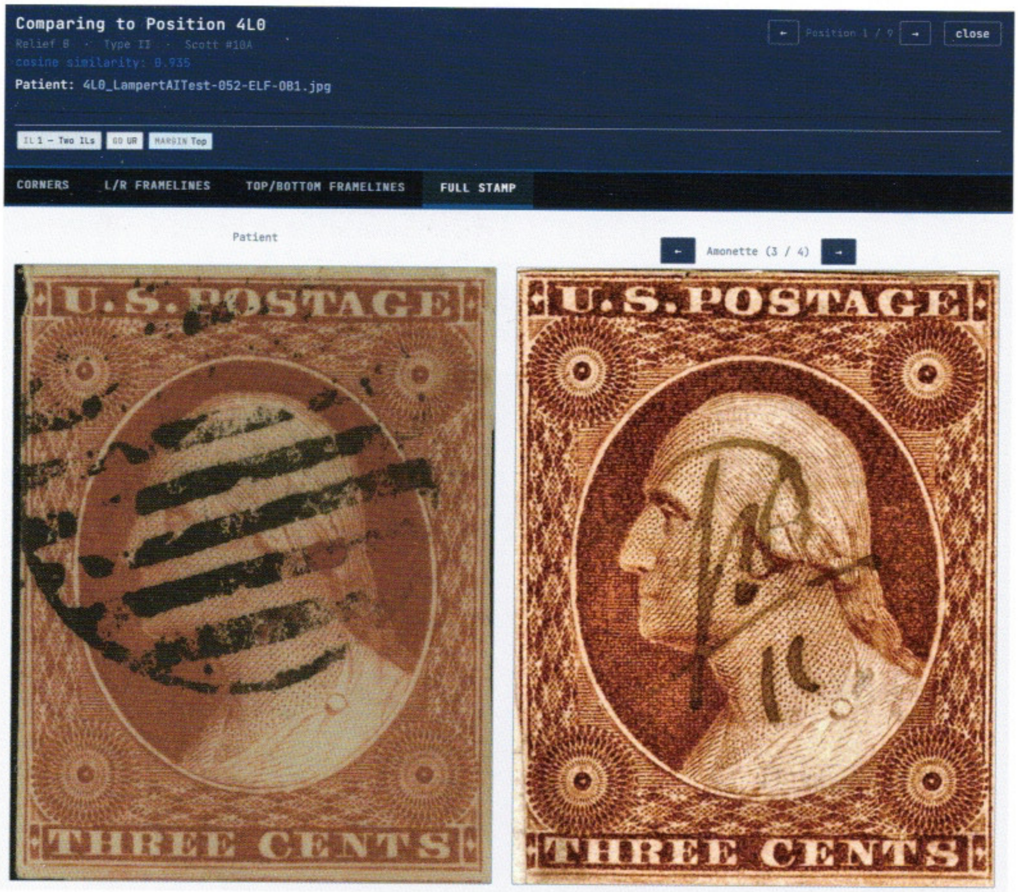

Figure 3. The Full Stamp tab of the Compare module, showing position 4L0 against the Amonette reference copy. Patient stamps and reference stamps are cropped and aligned identically, which makes direct side-by-side comparison straightforward. The Amonette collection is one of four complete platings in PlateAI's reference library; arrows above and below the reference column cycle through the Chase, Celler, Amonette, and Lund-Amonette-O'Doherty (LAO) platings.

here comparing the Position 4L0 patient against the Amonette reference.

Because 1E and 1i stamps present the most challenging positions for humans to differentiate (but interestingly, not for the model), the Compare module has an additional feature to bring in 1E references when examining 1i positions, and vice versa. For all of the tabs, a user may advance to the next potential position. He can reposition the patient stamp, zoom in, or download the cropped and aligned patient and reference scans.

The Attribute Filter

A second tool is the **Attribute Filter**. The left panel offers basic filters for color

(Orange Brown or Not Orange Brown), Type (I or II), Relief (A, B, or C), and Plate. An Advanced module extends these options with recuts, guide dot locations, plate anomalies, and margin position. An example is shown in the screenshot in Figure 4. In the setting illustrated, the search has been limited to positions with two inner lines (IL 1), a one-line recut in the upper-left triangle (Recut 11), a lower-right guide dot, and a non-marginal position on the plate. These selections, combined with Type II, Relief B, and Not Orange Brown on the basic filter panel, narrow the 2,600 candidates to 87 positions, which can then be ranked by the model or inspected directly.

Figure 4. The Advanced Plating Filters module, used to narrow candidate positions by detailed plating features. In the instance shown here, the search is limited to positions with two inner lines (IL 1), a one-line recut in the upper-left triangle (recut 11), a lower-right guide dot, and a non-marginal position on the plate. These selections, combined with Type II, Relief B, and Not Orange Brown on the basic filter panel, narrow the 2,600 candidates to 87 positions, which can then be ranked by the model or inspected directly.

If you want to pare down the top ten by eliminating positions you know are not correct, you can select and apply the criteria you have chosen. If the correct position is not shown, hit Plate! again with the options selected, and the tool will return a new top ten list restricted to the user's selection.

The Position Explorer

A third tool is the **Position Explorer.** This is a reference browser, independent of PlateAI's identification capability. The user picks a plate, applies any combination of attribute filters, and inspects the eligible positions directly. No stamp needs to be loaded. The screenshot in Figure 5 shows the Position Explorer, presenting the 200 positions of Plate 5L with Relief B positions highlighted. Match counts at the top show how many positions pass the current filters, both for the selected plate and across all 2,600. A patient stamp can be loaded but is not required: the explorer is equally useful as a study tool with no stamp at all. In fact, the highlighting pattern shown in Figure 5 reveals a philatelic feature: Some positions on Plate 5L were re-entered with a different relief than the original, causing a

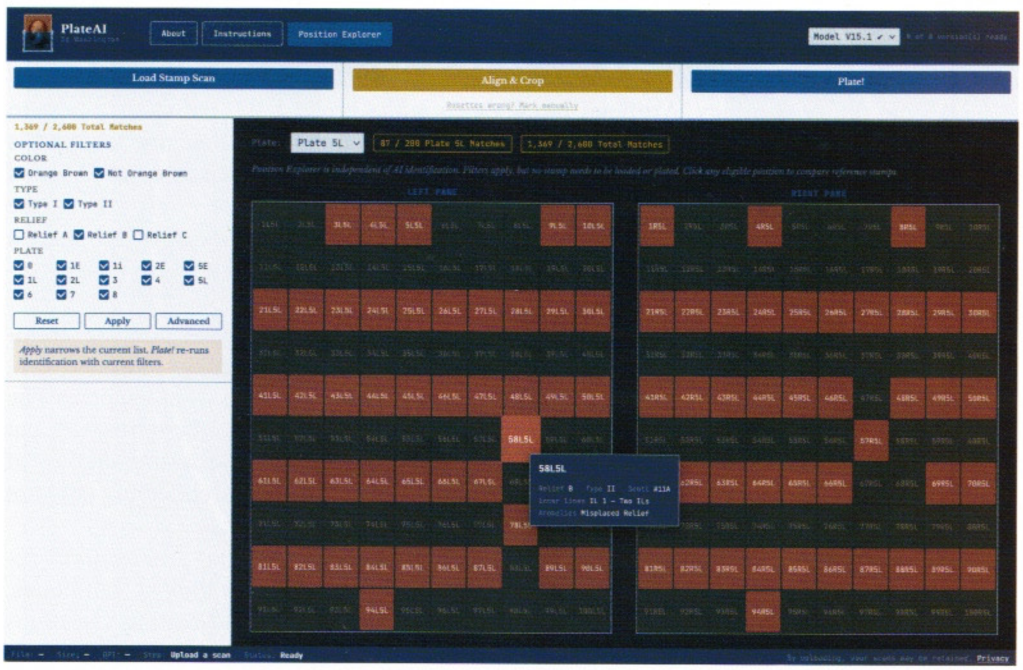

Figure 5. The Position Explorer, here showing the 200 positions of Plate 5L with Relief B positions highlighted. The Explorer works independently of AI identification — no stamp needs to be loaded. Match counts at the top show the count for the selected plate and across all 2,600 positions. The highlighting pattern on Plate 5L reveals a philatelic feature: Some positions on Plate 5L were re-entered with a different relief than the original (5E), causing a handful of B-relief impressions to land in rows normally reserved for A or C reliefs. The Explorer makes plate-level patterns like this visible at a glance.

handful of B-relief impressions to land in rows normally reserved for A or C reliefs. The Explorer makes this kind of plate-level pattern visible at a glance.

Two smaller features round out the workflow. **Copy** puts the results on the clipboard. **Log** starts a session log that captures every plating result. While logging, the user can mark his final choice on each stamp, which is captured in the log alongside the AI's predictions, useful when plating many stamps in a session. The CSV button downloads the log as a spreadsheet.

Early users have consistently observed that the tool is fun to use. Plating traditionally has been slow and demanding work, and the learning curve kept many collectors from attempting it. PlateAI eliminates most of the drudgery. A beginner gets results approaching what an experienced plater gets, in a fraction of the time. Experienced platers among the early users report a different experience: PlateAI lets them devote their attention to cases that are genuinely interesting (the close calls, the outliers, the stamps that don't fit easily) rather than on the routine identifications that formerly consumed most of their time. They also report that they are correcting misplated stamps that have been sitting unnoticed in their collections for years.

How PlateAI works

At its core, modern machine learning is a very simple idea applied at very large scale. A model is a mathematical function with millions of adjustable parameters. The user shows it many examples (in this case, scans of plated stamps with positions known) and uses an automatic procedure to nudge the parameters so that the model's outputs become correct, or at least more correct than they were before. The model is not programmed with rules about what makes a position distinguishable. It learns those distinctions by itself, from the examples it sees.

A note on terminology. Artificial intelligence is the broad field that studies how to make machines do things that traditionally required human judgment. Machine learning is one approach within that field—one that learns from examples rather than from hand-coded rules. Neural networks are one kind of machine-learning model, loosely inspired by the way biological brains process information. Convolutional neural networks, or CNNs, are a kind of neural network specialized for working with images. Like most modern image-recognition systems, PlateAI uses CNNs.

PlateAI's first version was a proof of concept. It used the simplest possible architecture: feed the entire stamp image database into a standard image-recognition network, ask it to learn which of 2,600 positions each image came from, and see what happens. The result was about one percent accuracy. That sounds terrible, and in absolute terms it is, but a random guess across 2,600 positions would produce accuracy around four-hundredths of a percent. One percent meant the model had learned something. The architecture was not the right one for the problem, but the approach was viable.

The second version, built within days of the first, made the central architectural decision that PlateAI still uses. Platers do not look at whole stamps. They look at framelines, diamond blocks, triangles, and other diagnostic regions where the hand-cutting variations are most visible. The model needed to do the same. The stamp was divided into multiple overlapping regions (settling on ten for the most recent version): one for each of the four framelines, one for each of the four corners, and a compression scan each for the left and right framelines. These are indicated by the colored frames overlain on the black and white enlargement in Figure 6. A separate small network, called a branch, was trained to analyze each region. The branches were combined into a single output. The first multi-branch version reached about 30 percent accuracy. The architecture was right; everything since then has been refinement. Note: The current version does not use the horizontal compression scans, but they are available to users in the compare module. The current version also does not use the four internal rectangles (orange in Figure 6) that mark the relief regions. These are available for future refinements in the software.

Refinement has been incremental rather than dramatic, with improvements in the model itself and in the training data. Each training run takes several hours on rented GPU hardware—graphics processing units, the specialized chips that make training a neural network fast enough to be practical. Cloud compute can get expensive fast; providers are happy to sell more hardware to compensate for inefficient code. The entire out-of-pocket

FIGURE 6: DIAGNOSTIC REGIONS ANALYZED BY PLATE AI

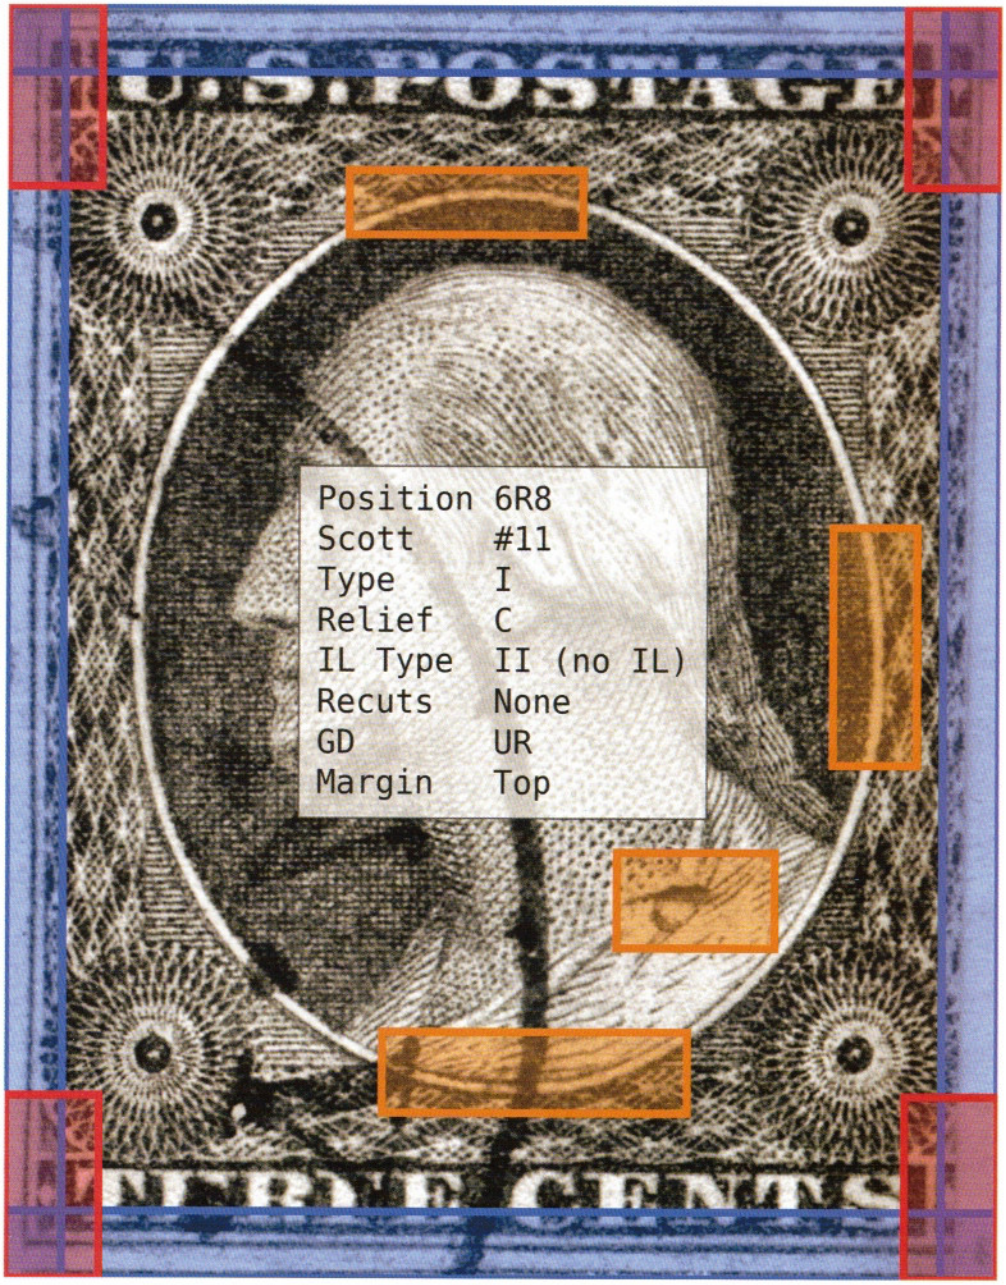

Figure 6. A plated imperforate 3¢ Washington stamp (Position 6R8) showing the regions of the stamp that are examined and assessed by PlateAI. On every stamp PlateAI's multi-branch convolutional neural network analyzes the four framelines (here outlined and shaded in blue) and the four corners (outlined and shaded in red). The orange rectangles indicate relief regions, present in the model's architecture but unused in the current version of the tool. The central panel summarizes the stamp's attributes as a human plater would record them. (IL = inner line; GD = guide dot).

budget for this project, including GPU compute and a Claude subscription, has been under $500, not counting my time. Keeping that number low required attention to efficiency rather than just throwing more computer power at the problem.

The working model was in usable shape within about two weeks of starting the project, and the basic user interface was built in a couple of days after that. The rest has been steady, incremental improvement to both. The current version, v15.1, sits at about 97 percent accuracy on a test panel of 455 stamps (discussed further below).

What the model is computing, when it produces a similarity score, is something called cosine similarity. The model converts each stamp into a list of numbers, let's call it a fingerprint, that captures the visual essence of the stamp's diagnostic regions. To compare two stamps, the model compares their fingerprints. Cosine similarity is a specific way of measuring how alike two such fingerprints are. It returns a single number between negative one and one, with one meaning identical. (If you plate a stamp in PlateAI and get a cosine similarity of 1.000, it likely means that very stamp is already part of the training data.) The score the user sees in PlateAI's results is the cosine similarity between the patient stamp's fingerprint and the various reference stamps' fingerprints for that position.

The model was trained using a technique called triplet loss. Instead of teaching the model to assign each stamp to a specific position label (the approach the proof of concept used) triplet loss teaches the model what the same position looks like and what different positions look like. The training procedure shows the model three stamps at a time: an anchor, another stamp from the same position, and a stamp from a different position. The model is told to make the anchor's fingerprint look more like the same-position stamp's fingerprint and less like the different-position stamp's fingerprint. Repeat this many millions of times across the training data and the model learns, on its own, what features distinguish positions. Nobody told it to look at frame lines, or guide dots, or recuts. It found those things itself.

One useful consequence of training the model this way is that it learns what to ignore. Cancellations vary wildly. A stamp's reference image and the user's patient stamp will never have the same cancel, so during training, cancellations behave like random noise that has nothing to do with the position. The model gradually learns to treat them that way. By the time training is finished, a heavily cancelled stamp can still be plated correctly, because the cancellation is essentially invisible to the features the model relies on. The same is true for paper imperfections, ink spread, and other accidental marks. The model attends to what is consistent across examples of the same position and ignores what is not.

PlateAI trains on grayscale images rather than on full color scans. Thus the black-and-white stamp image in Figure 6, Grayscale has one color channel instead of three for color, which cuts training time by about two-thirds at a small cost in maximum possible accuracy. I initially expected color information to help the model isolate and ignore the mostly black or blue cancellations. But in practice, the grayscale model handles cancellations just fine. A future version trained on color scans might gain a small amount of accuracy, but it has not yet been worth the additional compute.

PlateAI produces some results that surprise human platers. The model discriminates positions on plates 1E and 1i as well as it does with any other two positions. That's not so with a human. A human plater can often narrow down a position to the same spot on either plate 1E or 1i, but making that final determination requires extensive experience, lots of time, and lots of references.

One thing the model does not do is identify the relief. The training-data overlays in Figure 6 include regions for relief analysis, but the current production model is blind to them. An earlier version did include a relief classifier, and, surprisingly, it was less accurate with that addition than without. The relief classifier constantly overfit: it learned the specific training stamps rather than the general pattern, so it could not generalize to unseen stamps, while at the same time reporting that it was highly accurate. Curiously, though the model is blind to the portions of the stamp a human uses to identify reliefs, it clusters C-relief stamps anyway. Plate a C-relief stamp in PlateAI and many of the top ten will be C-reliefs. Something in the diagnostic regions distinguishes C-reliefs from the others. What that something is remains unknown.

The model's accuracy depends on the consistency of its inputs. If reference stamps were cropped and aligned variably, the model would have to learn that variation as well as the position differences, which would dilute its ability to distinguish positions. PlateAI's reference library is therefore aligned identically using the four rosettes as anchors. But there is a trade-off: cropping means loss of any information that lies outside the framelines, including most of the left-side guide dots, which typically appear several millimeters outside the framelines.

Claude: The AI assistant

A year ago, building a tool like PlateAI would have required a team of data scientists and a six-figure budget. What changed in the last year is that a single person with the right knowledge can now do this kind of work in collaboration with an AI assistant. As noted, I built PlateAI working with Anthropic's Claude. Claude is a large language model, the same family of AI system that produced ChatGPT and the now-familiar conversational assistants. Most users interact with these tools to draft emails, summarize documents, or answer questions. Claude can do all of that, but it can also write working software code, hold an extended technical conversation across many sessions, and act as the engineering side of a project while a human supplies the domain knowledge and the judgment.

Encouraged by friends outside of the stamp collecting world ("Just give Claude a try to build an app—it's fun!") I decided to see what it could do with plating stamps. The first message I sent to Claude on March 6, 2026, was this: "I would like to build an app that plates United States 1851-57 3¢ Washington stamps: Scott numbers 10, 10A, 11, 11A, 25, 25A. I have the domain expertise to do it manually, and I have access to 4-5 examples of each of the 2,600 positions. Is this something that is feasible?"

Claude's response began: "This is a fascinating philatelic project! Let me make sure I understand the scope before diving in. The 1851-57 3¢ Washington stamps are among the most complex to plate—with the 150 positions across the two panes (L and R) of the sheet, subtle constant plate varieties, and the differences between the imperforate (Scott 10/11) and perforated (Scott 25) issues."

That last sentence is dead wrong. The plates have 200 positions, not 150. Within 30 seconds of the project starting, Claude had confidently invented a fact about the very stamp the project was about to be built around. This pattern—Claude being articulate, helpful, fast, and occasionally wrong about specifics—would recur throughout the project.

What followed the opening exchange was a clarifying back-and-forth about scope, approach, image sources, and whether to use rule-based logic, AI image analysis, or both (the answer was both). Everything since has been refinement.

I barely read any of the code Claude wrote. My job was to direct the project, judge the results, and recognize the symptoms when something was broken. Claude would then trace and fix it. The working relationship was the kind a manager has with a senior engineer, except that the engineer never gets tired, turns work around in seconds rather than days, and is willing to throw away half a day's work and start over without ego.

Claude has limits. It forgets context over long sessions and has to be reminded of decisions made earlier. It sometimes confuses versions of files or loses track of which change went where. It is fast and competent at coding but has very little built-in knowledge of this stamp. Almost everything useful Claude knew about the 3¢ Washington came either from my direct explanations or from Claude reading StampPlating.com. And as the opening exchange showed, Claude will confidently invent stamp facts if not constrained. It frequently conflated relief and type. It constantly forgot that plates 1E and 1i contain both Type I and Type II stamps. It kept referring to non-existent variations in rosettes. It once suggested designing for top and side "ornaments" (almost certainly conflating 1¢ facts with 3¢ facts). At one point it claimed the 3¢ was green. And it made historical claims that sounded plausible but weren't true. Catching those errors before they made it into code or documentation was a constant background task.

What Claude does well is conversation — extended, technical conversation in which I could think out loud, propose an approach, get pushback, refine the idea, and have working code at the end of an hour. Claude switches contexts almost instantly. It reads code and design files faster than a human can. It accepts pushback gracefully when I had a reason for a particular decision and does not get attached to its own suggestions. It is very fast at fixing bugs I identify by their symptoms. The pace of work was rapid: weeks' worth of traditional software development was done in hours and sometimes minutes. And despite the shortcomings, my friends were right—it is fun to build apps with an AI assistant like Claude!

The supporting infrastructure was deliberately mainstream. PlateAI is built with standard HTML, CSS, and Python. The model was trained on Google Colab, a cloud notebook environment that rents access to GPUs by the hour, because training a multi-branch CNN requires more processing power than my Mac can provide. The code lives in a private GitHub repository. The production server runs on DigitalOcean, a standard cloud host. Claude walked me through every technical task that was new, including the exact syntax for command lines or the precise options to apply in a web-based interface as I set up the infrastructure.

The wider implication of the build process matters more than PlateAI itself. The model architecture, the workflow, the training-data discipline, the deployment pattern: all transfer directly to other multi-position issues. The 1857 perforated Type IV 3¢ Washingtons (Scott 26A) could be added to PlateAI with minimal architectural changes if at least three complete platings of those issues were available as training data. The 1¢ Franklins, the 10¢ Washingtons, the 1847 stamps and the Postmaster Provisionals could each have dedicated AI plating tools. Beyond classic U.S. philately, tools to identify Washington-Franklin varieties would likely be a simpler task. The barriers to any of these projects are no longer software development or data science expertise. They are, first, access to sufficient reference images for training, and second, domain expertise in the specific issue.

The training data, and the work it took to make it usable

The first week of the PlateAI project was spent acquiring, organizing and understanding training data. That work continued through the whole project, but the initial week was exclusively focused on the data, before I had any idea how the machine learning would work.

In the data world, the tooling to turn raw inputs into a usable form is called a data pipeline. Raw stamp scans come from many sources, with different resolutions, different formats, different naming conventions, different image conditions. Some are full panes containing a hundred stamps each; some are individual stamps; some come in pairs or blocks of stamps, some have positions known and labeled, some do not. None of this is ready to feed into a neural network. The pipeline is what turns it into something usable.

PlateAI's pipeline is a small library of utilities Claude and I built side by side with the model itself, image extraction tools to produce individual stamp images efficiently from scanned panes or sheets of stamps organized in multiple ways: 1. A utility that crops and aligns individual stamp images in the same way the "align and crop" feature in PlateAI does it, only in batch to process thousands of stamp images at a time and to flag the ones that need manual intervention. 2. A utility that allows a person to quickly handle misaligned stamps. 3. A utility that runs quality checks, looking for dimensional outliers, misplated or duplicate training images. 4. Various utilities to move and rename files as they move through the data pipeline and to run statistics on the training data.

These utilities make it possible to process thousands of stamp images. Without them, the project could not have succeeded because the effort to handle the data would have required months of tedious, error-prone manual work. These tools were built incrementally as the project unfolded, and the effort has paid for itself many times over and continues to pay off every time a new set of scans arrives.

The training library now stands at 19,000 stamp scans, drawn from the four complete platings described earlier and from smaller partial collections contributed by other collectors. The average position has between four and ten reference scans. The scarcest plate (1E) has the fewest, while the most common plates (1L, 2L, 3) have the most. Additional scans for the Orange Brown plates and plates 6, 7, and 8 could incrementally improve the accuracy. Plates 1L, 2L, 3, 4, and 5L are likely covered about as well as they can be covered.

One of the unexpected benefits of building the pipeline this way is what it surfaced about the existing reference material. Once the model became reasonably accurate, it could be turned around and run against its own training data, looking for stamps it strongly believed were labeled incorrectly. More than two hundred such cases turned up—stamps that had been catalogued as one position by their original platers but that the model identified as a different position with high confidence. Included in the 200 are about 75 identified misplatings in the four reference collections. Those positions were corrected in the training data. They were removed from the reference photos available to PlateAI users, but at present, replacement reference images have not been added. No position has fewer than three references to compare to.

Each potential misplated stamp was reviewed by a human, and the great majority of them turned out to be genuine misplates that had sat in private collections, sometimes for decades, attributed to the wrong position. The model is not always right when it disagrees with a human plater. But it is right often enough that running this kind of check has become a regular part of the pipeline. Each correction, in turn, becomes cleaner training data for the next version of the model.

The library is not closed. New scans, new partial platings, and new complete platings continue to come in. Each addition strengthens the model in the positions it covers.

Testing and performance

PlateAI was tested against a panel of 455 stamps that were never used for training. Each stamp in the panel is a high-quality, four-margin scan with a known position, drawn proportionally across the 13 plates and the three reliefs. This is the model's report card.

Rather than throwing arbitrary scans at the panel, the testing framework starts from the ideal case—a complete stamp with four clean margins—and then simulates what happens when the conditions degrade. It masks one frameline at a time, then two adjacent framelines, then a corner. This systematically simulates the model's performance on less-than-perfect stamps.

Table 1 shows the results. The model is highly accurate on four-margin stamps, degrades gracefully when a single margin is missing, degrades more sharply when two adjacent framelines are gone, and holds up well when only a corner is lost. Even in the worst tested condition, the right answer is in the top ten more than 95 percent of the time. Across all conditions tested, the right answer was in the top 15 for every single stamp in the panel.

The Type I/Type II split in Table 1 reveals that the model handles Type II stamps better than Type I. That's not surprising given the fact that Type II stamps have more opportunity for variation with recut inner lines.

TABLE 1: TEST PANEL ACCURACY BY SCOTT NUMBER AND MARGIN CONDITION

Margin condition	Scott variety	N	Top-1	Top-3	Top-5	Top-10
Four-margin	10 (OB Type I)	19	89.5%	94.7%	94.7%	100.0%
	10A (OB Type II)	150	96.7%	100.0%	100.0%	100.0%
	11 (Not OB Type I)	141	94.3%	98.6%	99.3%	100.0%
	11A (Not OB Type II)	144	100.0%	100.0%	100.0%	100.0%
	Overall	454	96.7%	99.3%	99.6%	100.0%
Three-margin: (single frameline masked)	10 (OB Type I)	76	81.6%	93.4%	93.4%	97.4%
	10A (OB Type II)	600	94.0%	98.3%	99.2%	99.7%
	11 (Not OB Type I)	564	81.7%	93.3%	95.6%	98.4%
	11A (Not OB Type II)	576	98.1%	99.3%	99.5%	99.5%
	Overall	1816	91.0%	96.9%	97.9%	99.1%
Two-margin: (two adjacent framelines masked)	10 (OB Type I)	76	64.5%	88.2%	92.1%	94.7%
	10A (OB Type II)	600	88.7%	96.3%	97.8%	99.0%
	11 (Not OB Type I)	564	67.6%	83.0%	88.5%	94.5%
	11A (Not OB Type II)	576	92.2%	97.7%	98.4%	99.1%
	Overall	1816	82.2%	92.3%	94.9%	97.5%

Model accuracy on a 455-stamp held-out test panel (ie: images not included in the training data). Three- and two-margin stamps were simulated by masking one or two adjacent framelines. Small N for Scott 10 reflects the scarcity of Type I positions on the Orange Brown plates.

What the test data implies for the user is straightforward. When PlateAI is confident, the top candidate is very likely correct. When it is unsure, the right answer is almost always still on the list. If the top candidate does not match what the user expected, a few seconds of attribute filtering or a side-by-side comparison against the next few candidates will surface the right one. Either way, the right discipline is to verify with the compare module on any plating that matters. Confidence labels are guides, not guarantees.

For context on what these numbers mean against human performance: I plate correctly about 95 percent of the time without aid. Amonette plated correctly about 99 percent of the time. Chase and Celler approached 100 percent. These figures come from running PlateAI against stamps each plater had identified, with the disagreements adjudicated by a human. The numbers in Table 1 indicate where PlateAI sits along that scale.

Limits and open questions

PlateAI works on the imperforate 3¢ stamps of 1851-56 (Scott 10, 10A, 11, 11A) and the Type I and II perforated stamps of 1857 (Scott 25, 25A). It does not work on Type III or IV stamps (Scott 26, 26A). PlateAI will appear to work on Type III and Type IV stamps and may even claim high confidence that a Type III/IV stamp matches a position from a Type I or II plate. A future enhancement to the application could catch these cases more gracefully, but at present, PlateAI may confuse a user who unknowingly uploads a 26 or 26A.

The tool is calibrated for 1200 or 2400 DPI flatbed scans. Lower resolutions and cell phone photographs are accepted and will produce results, but accuracy degrades with image quality. A user working from a phone photograph or low-resolution scan (even 600 DPI) should treat any single result with skepticism and verify against the compare module rather than acting on the top candidate alone.

PlateAI and other AI tools that are sure to be built will open new research possibilities. For the 3¢ Washington, it is a small next step to build heatmaps of the regions of the stamps the model distinguishes—potentially finding new varieties or adding deeper understanding to the production process. Why does PlateAI cluster C reliefs without looking at the ovals or the "gash on the shoulder?" How does it distinguish 1E and 1i so well? What does it see that we don't?

Thanks

Robert J. Lampert and J. Bryan O'Doherty were central figures in this project's progress. Bob's encouragement and his patient instruction in plating made the project possible. Bryan's encouragement and his StampPlating.com were both formative. The site is a central place to learn about the 3¢, and its plating wizard is an example of what a web-enabled stamp identification application can do. Gary Schrader, Elvin Fritz and Charles Temple were early users whose feedback on the tool, and broader perspective on plating, shaped the project as it developed.

The 3¢ Washington stamp has been studied for 175 years. Each generation of collectors has added something—a corrected recut, a recovered plating chart, a digital reference, and now a machine learning model. PlateAI is one more contribution in that long line. ■