Marking Time: Increased Scope and Accuracy for Sketch Classification of the Clock Drawing Test

by

Jonathan Chien

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2015

Author . . . .
Department of Electrical Engineering and Computer Science
May 22, 2015

Certified by . . . .
Randall Davis
Professor
Thesis Supervisor

Accepted by . . . .
Albert R. Meyer
Chairman, Master of Engineering Thesis Committee

Marking Time: Increased Scope and Accuracy for Sketch Classification of the Clock Drawing Test

by

Jonathan Chien

Submitted to the Department of Electrical Engineering and Computer Science
on May 22, 2015, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

In this thesis, I designed and implemented improvements to an automatic classifier for the digitized clock drawing test, a diagnostic tool for assessing cognitive impairment, which asks the patient to draw an analog clock face using a digital pen. The classifier handles both the grouping of strokes into clock components and the subsequent labeling of those groups. Despite the domain-specificity, classification is a challenging problem because the subject often has cognitive or motor impairments. It is thus important for the classifier to be able to handle a wide range of input with distorted, overwritten, or missing components. I improve the robustness of the classifier, particularly for messy clinical data, by incorporating intrinsic stroke properties, developing additional symbol recognizers, and creating a global context evaluator. I describe in this thesis properties for isolated symbol recognition, features for symbol recognition and match scoring, as well as common sense rules based on a symbol's local and global context in a drawing. I combine these elements into a new system that locally maximizes a global label assignment based on match quality and context. I demonstrate that this system accurately recognizes a wide variety of clinical input, improving overall classification performance.

Thesis Supervisor: Randall Davis
Title: Professor

Acknowledgments

I would like to acknowledge Professor Randall Davis for his ideas, guidance, and feedback over the past year. Not only did I learn a tremendous amount technically, but I also learned a lot about communication and academic research, which I hope to continue applying in future pursuits.

I'd also like to thank the rest of my research group, my friends, and my family for their support and encouragement throughout my time at MIT.

Contents

Chapter 1

Introduction and Motivation

This chapter gives an introduction to the clock drawing test, the ClockSketch program, and the problem of sketch classification.

1.1 The Digital Clock Drawing Test

The clock drawing test is a widely used cognitive screen for a broad range of cognitive disorders. The test asks patients to draw a complete analog clock at a specific time, often 11:10. After this task is complete, the patient is shown a correctly drawn reference clock, and is asked to copy the correct clock drawing. The first part of the test is called the Command clock and the second part the Copy clock. Cognitively impaired patients often omit, distort, or misplace clock components in their drawings. By evaluating any errors that are in each clock, as well as those present in the command clock but not the copy clock, a clinician can estimate the degree of a patient's mental impairment and sometimes identify the specific elements such as memory or spatial reasoning that are impaired. While the test is simple and somewhat subjective, it has been shown to exhibit high sensitivity, specificity, and correlation with other cognitive tests.[Shulman, 2000]

While the clock drawing test has traditionally been administered with pen and paper, a digital version has been created using a digital pen and specialized paper. The digital pen and paper provide a time sequence of the pen's position when it is in contact with the paper. This provides precise timing data, something not present in the traditional clock drawing test, but potentially useful both to improve classification of a patient's pen strokes and in diagnosis. For example, if someone scribbles over a digit or tries to overwrite the original digit with another digit, this information may be lost in the final drawing but can be visualized separately by relying on timing data. The timing between various strokes and how that differs between the Command and Copy portions can also be helpful in determining a patient's cognitive deficiencies.

1.2 Sketch Recognition and ClockSketch

Sketching is a natural, efficient, and unconstrained method of inputting information. Digitizing sketches has historically been time-consuming and error-prone, often requiring the help of domain-specific software for sketch analysis. With the increasing prevalence of phones and computers with touch and stylus interfaces, sketch recognition research has grown in popularity and importance. Given the inherent tradeoff between drawing freedom and labeling accuracy, most sketch recognition system have targeted domain-specific problems such as note-taking[Ispas et al., 2011], physics[Lee et al., 2008], or non-Roman writing[Ma and Liu, 2008], constraining the possible component labels. We aim to tackle the sketch recognition problem in the domain of the clock drawing test.

We do this as part of the ClockSketch program, a system designed to aid clinicians in analyzing and evaluating digitized sketches from the clock drawing test. In addition to allowing clinicians to zoom in on or highlight specific components of the completed drawing, ClockSketch can actually replay sketches, unraveling layers of overwriting. Clinicians can also manually group and label each stroke of a patient's drawing into various clock components. Accurately identifying the components is important to the clinician's ability to make a diagnosis. However, manually labeling every single stroke is tedious and time consuming. In response, work has been done to enable ClockSketch to automate the grouping and classification of strokes into various clock components. Accurate grouping and labeling of the strokes is a necessary prerequisite to eventually being able to automatically suggest a diagnosis from higher-order features in a patient's drawing. Even an imperfect system for automatic grouping and labeling can save clinicians significant amounts of time, requiring them to manually correct only a small subset of difficult to label strokes rather than every single stroke. Examiners can then focus their time on the difficult cases, and on large-scale analysis with the labeled data.

The classification problem can be thought of as two separate problems, grouping and labeling. The goal of grouping is to identify strokes that belong to the same component (e.g. the 1 and the 2 in the numeral 12) and to separate strokes that belong to different components. The goal of labeling is to determine which component each of the groups represents. The goal is to enable the ClockSketch classifier to achieve high performance on both of these problems, even on inputs with messily drawn, repeated, or missing components, which can occur frequently with cognitively impaired patients.

Some portions of the clock drawing classification problem have been well studied. For example, digit recognition has been a staple benchmark of machine learning models, with many models achieving greater than 99.5% performance on the MNIST data set of handwritten digits. However, these digits are size-normalized and mostly well-written as they do not come from cognitively impaired patients. Proper classification of the clock drawing requires solving the additional problem of grouping and must also deal with messily drawn components as well as poorly studied challenges such as overwrites. As such, achieving robust performance with the ClockSketch program is of interest due to the unique challenges that it must handle.

Chapter 2

Definitions and Data

This chapter defines relevant terminology used in the rest of the document. It explains how the data from the clock drawings is represented and provides an overview of the data sources used.

2.1 Terminology

2.1.1 Data Representation

Point : a recording of the x and y positions of the digitizing pen at a given timestamp. The pen records 80 points a second. Since points are recorded based on time and not distance, points cannot share the same timestamp but can share the same position.

Stroke : a Stroke is a time-series of points from when the pen touches the paper to when it leaves the paper.

Group : a collection of one or more strokes during any phase of the classification process.

Label : a classification assigned to a collection of one or more strokes such as numerals "1" through "12", "hour hand", or "clockface".

Symbol : a collection of one or more strokes specifically intended to represent a single clock component such as a hand. As such, every Symbol has an associated Label. At the end of classification, all strokes in a drawing will be grouped into labeled symbols.

Clock Drawing : A collection of one or more symbols that corresponds to a single phase of the clock drawing test.

2.1.2 Labels

Clockface : The outline enclosing the hands and other clock components. Generally a complete circular stroke. Sometimes, clockfaces are more oval than circular but they are rarely not round or missing.

Numeral : A number from 1 to 12 drawn around the clockface border.

Hand : The minute or hour hand, usually emanating from the center of the clockface, including any arrowheads if drawn.

Center Dot : A dot, often filled in, usually drawn at the base of the hands at the center of the clockface.

Crossout : An overwritten or scribbled over symbol. Depending on the original intended symbol, a crossout can be either a Crossout Digit or Crossout Hand. Includes the scribble stroke if it exists.

Noise : Stray unintentional marks or strokes that do not fit into any category that contribute to the clock drawing.

Tick Mark : A mark that generally intersects or is near the clockface border. Often placed either in between numerals, or intersecting the clockface at the same angular position as the numerals.

2.1.3 Properties of Strokes, Symbols, or Drawings

Bounding box : The smallest rectangle that includes every point of a stroke, symbol, or drawing.

Center of mass : The point defined by the mean x and y coordinate of all the points of a stroke, symbol, or drawing.

Radius : The mean distance from the center of mass to each point in a stroke, symbol, or drawing.

Average Span : The mean of the height and width of the bounding box of a stroke, symbol, or drawing.

Inner : Strokes and symbols that are part of a clockface, hand, or center dot are considered inner, as they are usually drawn around the center of the clockface.

Outer : Strokes and symbols that are part of a numeral or tick mark are considered outer, as they are usually drawn around the outline of the clockface.

Temporally sampled : A method of sampling points where the difference in timestamps between consecutive points is constant. This is how the digital pen currently samples its points.

Spatially sampled : A method of sampling points where the distance between consecutive points of a stroke is constant. Since this is not how the digital pen collects the data, the points must be resampled with interpolated timestamps and positions to obtain a spatially sampled stroke.

Ouyang score : The match distance from a symbol to a specific numeral as determined by the Ouyang recognition algorithm, which is described in 3.3.3.

R^2 : 1 - the ratio of the sum of squared residuals from a fitting algorithm divided by the variance of the points of a stroke, symbol, or drawing. This value indicates the goodness of fit of a particular shape.

2.1.4 Other Terms

Command : The first phase of two in the clock drawing test where the patient is asked to draw a clock at 11:10.

Copy : The second phase of the clock drawing test where the patient is given a reference image of a properly drawn clock at 11:10 and asked to copy the drawing.

Angular extent : The length of an arc in degrees. The angular extent of a perfect circle or ellipse is 360 degrees. Used when trying to fit a clockface as repair strokes are often drawn to increase the overall angular extent of the clockface.

2.2 Data Sources

The seven data sources selected for system training and performance evaluation are described below. The first five sets are the ones that the current ClockSketch classifier was developed, trained, and tested on. The two additional clinical data sets provide rarer strokes such as tick marks and harder classification challenges due to the presence of significantly distorted and messy drawings.

YDU-51 healthy training set : 51 tests drawn by healthy patients. The clock faces in this set generally contain all 12 numerals and have no major errors present. The numerals from this data set were used to train the Ouyang nearest-neighbors numeral classifier.

YDU-100 healthy test set : 99 different tests drawn by healthy patients from the same testing site as the YDU-51 set. These tests are used to ensure that the Ouyang numeral classifier and overall system perform accurately on well-drawn clocks that the Ouyang numeral classifier was not trained on.

VIN-96 clinical test set : 96 tests that were chosen effectively randomly. Was used to act as a representative sample of the standard clinical use case for evaluating the performance of the current ClockSketch classifier on clinical data.

drawn effectively at random from our corpus of clock tests, intended to represent a sampling of the average clinical input to the classifier.

EGE/ORU-112 clinical test set : 112 tests that were also chosen effectively randomly but from a different testing location. Was used to ensure the current classifier performed well on a wide range of clinical tests.

EMD-20 clinical test set : 19 tests that were chosen as a sample of some of the more difficult clock labeling challenges.

CIN-170 clinical test set : 170 tests that also pose some particularly difficult clock labeling challenges. Several clocks have particularly interesting anamolies, such as the lack of a clockface or digits, crossouts of an almost completed drawing, or rarer strokes such as spokes.

TICK-74 clinical test set : 74 additional tests that also come from the CIN test site that all contain tick marks. Even disregarding the presence of tick marks, these clocks are pose difficult classification problems. This data set ensures that there is a large breadth of tick mark strokes to evaluate the performance of the new classifier on tick marks.

Each example clock in this text is drawn from one of these data sets. A full list of the names of the clocks in each data set can be found in Appendix A.

Chapter 3

Background and Previous Work

This chapter summarizes relevant research for sketch grouping and labeling in the context of challenges that are relevant to the clock drawing test. It also summarizes the current strategy of the ClockSketch classifier and reports its performance.

we summarize relevant research for sketch grouping and labeling in the context of challenges that are relatively specific to the clock drawing test. We provide a summary of the previous ClockSketch classifier's algorithm and report its performance.

3.1 Segmentation and Grouping

While there have been a wide variety of approaches to grouping strokes, none are completely adequate for the segmentation challenges in the clock drawing test.

[Dickmann et al., 2010] tackles the problem of grouping in the context of two mock drawings, landscapes and drawings consisting solely of hatchings and arrows. He evaluates two types of classifiers to make a decision of whether to start a new grouping between each stroke based on spatiotemporal features. The features he looks at are duration between end of previous stroke and start of next stroke, minimum distance between the two strokes, duration for a given stroke, and velocity for a given stroke. He evaluates support vector machines (SVMs) and echo state networks, which are a type of neural network that retains memory over a sequence of inputs. However, his work assumes symbols consist only of consecutive strokes, which does not apply to a large number of symbols in our data set as cross outs and augmentations are not uncommon.

In the domain of recognizing circuit drawings, [Sezgin and Davis, 2008] build a dynamic Bayesian network, which allows for interspersed strokes between different symbols. However, building this Bayesian network requires having priors for each symbol type as well as hidden variables for each transition. While they achieve 87-95% accuracy in recognizing circuits with five distinct components, given the approximately 20 different symbol types in the clock drawing test, a Bayesian or Markov model would be overly complex and difficult to adequately train for our task.

[Yin and Sun, 2005] develop a domain-independent system that combines the segmentation and recognition problem into one by framing it as a dynamic problem minimization problem. They formalize a system to break the sketch down into candidate sets of primitive components such as lines or arcs with relationships such as intersection or tangency. The sketch is matched against a variety of training templates to minimize the match distance between its components. The most complex sketch they examine contains 13 line or arc components, which is significantly less complex than a standard clock drawing. The ClockSketch classifier must also identify intermediate components such as numerals or hands, which are not handled by their work.

Domain specific research has also primarily focused on simple spatial or temporal rules that achieve adequate performance for data in the domain. Segmenting handwritten text has been heavily researched but since text is generally written in an orderly directional sequence with clear spatial differentiation, work such as [Zhou et al., 2009]'s in segmenting Japanese is able to achieve high performance with just spatiotemporal features. In their work, overlapping text can always be merged into one line component. Such simple spatial rules would not apply to our task, where overlapping strokes can represent different clock components or a crossed-out symbol.

3.2 Labeling

There has been significant research on both numeral and primitive shape (e.g. arcs, arrows, and circles) recognition, which are similar to subproblems encountered by the ClockSketch classifier.

3.2.1 Numerals

MNIST is a data set of 70,000 handwritten digits that are already separated and scaled to 20x20 images. State-of-the-art machine learning algorithms using deep neural nets have achieved "near human" performance of $>$99.7%. [Ciresan et al., 2012] However, isolated digit recognition is an easier challenge for a few reasons.

1. Numeral classification for the clock drawing test must deal with significantly distorted numerals drawn by cognitively impaired patients. In some cases, a numeral is completely unintelligible without the help of context from surrounding digits.

2. Clock numerals also include two-digit numerals, making proper stroke grouping a requirement for good numeral labeling performance. Our analysis in Chapter 5 shows that the bulk of errors in numeral labeling are the result of poor segmentation rather than mistakes in isolated numeral recognition.

3. Multiple numerals, or many numerals and a scribble can be written in the same position. In this case, the ClockSketch classifier must be able to filter the overwritten symbol and determine what it is independently, requiring the consideration of temporal data.

3.2.2 Hands, Clockface, and Center Dot

While there has not been significant work outside of the ClockSketch classifier targeting the clock drawing domain, there has been research on recognizing shapes that are similar to various clock components. [Yu and Cai, 2003] describe properties, such as direction or curvature, to recognize primitive components such as curves or line segments. They provide methods to segment more complex strokes into combinations of primitive components, as well as recognizers that can detect combinations of primitive components and build them into a basic shape such as a rectangle or helix. The hands, clockface, and center dot clock components are analogous to arrows, ellipses, and dots respectively. For example, the arrow primitive shape can be modeled as a specific combination of line segments. While grouping rather than stroke segmentation is the major challenge in our domain, the stroke properties they use are helpful in labeling the inner clock components.

3.3 Previous ClockSketch Program

The previous ClockSketch classifier was developed by [Ma, 2014]. We provide an overview of each of its steps.

3.3.1 Noise and Clockface Detection

The previous ClockSketch program first classifies any stroke that spans less than 375 milliseconds, covers less than 1 square millimeter, and has an average span less than 1 millimeter as noise. Noise strokes are removed from all further consideration and are not included in calculations of any drawing properties (e.g. average stroke length).

It then looks for the longest stroke in the drawing and classifies it as the clockface if the stroke's radius is at least 1/4 of the drawing's average span. If a clockface is found, it considers adding the second longest stroke in the drawing to the clockface group provided the stroke's radius is at least 1/8 the drawing's average span, and the stroke's distance to the clockface center is larger than the average stroke distance to the clockface center. If the drawing's longest stroke does not meet the clockface criterion, it concludes that there is no clockface and uses the drawing's center of mass as the clockface center for later calculations. This method is designed to correctly identify clockfaces that are drawn with one or two large major strokes, which is the majority of cases.

After noise and clockface strokes are identified, all remaining strokes are assumed to be part of a numeral, hand, or center dot. The hand and center dot are considered inner symbols (i.e. near the center of the clock), while the numerals are considered outer symbols (i.e. near the edge of the clock).

3.3.2 Inner and Outer Stroke Separation

To distinguish between the inner and outer strokes, the classifier takes advantage of two common properties of the drawing: people generally draw the numerals consecutively rather than alternating drawing numerals and inner symbols, and the numerals are typically near the edge of the clock. To cluster the strokes, the classifier calculates each stroke's distance from the clockface center and timestamp from the start of the drawing. The clustering algorithm scales these two values in a 2:5 distance:time ratio and runs k-means clustering with 4 clusters on the resulting pair of values. It labels a cluster as consisting of outer strokes if the cluster's average distance to the clockface center is at least $3/4$ of the average stroke's distance (the value is less than 1 as there are generally more outer than inner strokes). Otherwise, it labels the cluster as consisting of inner strokes. The algorithm's parameters were empirically chosen based on performance on the previous data sets.

3.3.3 Numeral Segmentation and Labeling

Now that the outer strokes have been determined, the next step is to further segment the strokes into groups that each represent only one numeral. The angular position of each stroke is first calculated using the center of the stroke's bounding box and the clockface center. The strokes' angular positions are examined in order, and if the difference between consecutive positions is at least 0.22 radians, the strokes are identified as different numerals. Otherwise, they are grouped together as the same numeral. With the various numerals segmented into groups, the final step is to determine which numeral each group is. This step uses the Ouyang symbol recognizer, an overview of which follows:

1. The input symbol is resampled so that the points are spatially instead of temporally spaced.

2. The now spatially-sampled symbol is rescaled so that the vertical and horizontal positions of the points relative to the center of mass each have unit standard deviation.

3. The angles between successive points are calculated, and five 24x24 feature images are now extracted from the transformed symbol, spanning 2.5 standard deviations in both dimensions. Four of the feature images correspond to reference angles of 0, 45, 90, and 13 degrees 5, measuring how much a stroke is horizontal, vertical, or diagonal at a given point. The features values range from 0.0, indicating the stroke and reference angles differ by at least 45 degrees, to 1.0, indicating the angles are identical. The fifth feature represents each stroke endpoint.

4. The five feature images are Gaussian blurred and downsampled with a max filter to minimize effects from spatial variations.

5. Finally, a template matching algorithm that accounts for small spatial differences is used to calculate the total distance between the input's feature images and a reference template. The Ouyang score, or minimum distance between the input and any of the recognizer's reference templates is returned.

The YDU-51 data set is used to train twelve Ouyang recognizers for each numeral from 1 to 12. This gives each Ouyang recognizer a set of 101 templates to find the best match among, capturing a variety of drawing styles. Since the YDU-51 data set comes from healthy adults, the numerals are all relatively well drawn.

Twelve Ouyang scores or best-match distances are calculated for each numeral group using the Ouyang recognizers. Note that the label of a numeral can often be determined by its position and order (e.g. a symbol on the positive x-axis relative to the clockface center that is between a 2 and 4 is probably a 3). To make use of this information, the Ouyang scores are scaled by a constant ranging from 1.0 for the numeral that would generally belong in that angular position (e.g. 9 on the negative x-axis) to 2.0 for the numeral that would generally be angularly opposite. Finally, each group is labeled as the numeral that produces the minimum scaled-Ouyang score.

3.3.4 Hands and Center Dot Labeling

With outer strokes labeled, the remaining unlabeled strokes are the inner strokes. The previous classifier uses the following algorithm:

1. If both dimensions of the stroke's bounding box are less than 1/3 of the average stroke length and the stroke's farthest point from the center is closer than the average stroke's distance from the center, the stroke is determined to be a center dot.

2. Otherwise, the angular position of the stroke's farthest point from the center is calculated. If the angular position is between -90 and 90 degrees, i.e. on the right side of the clockface center, the stroke is labeled as a minute hand. Otherwise, it is labeled as an hour hand.

3.3.5 Crossout Detection and Error Correction

After every stroke is grouped and labeled, three repair steps are run to identify crossouts as well as common segmentation and labeling mistakes.

Identifying inner strokes that were accidentally clustered into the outer strokes

To detect inner strokes grouped as outer strokes, the classifier uses both a spatial and a temporal criteria. If a group of outer strokes meets either criteria, it is judged to potentially include an inner stroke. The spatial criteria catches groups whose average span is greater than the average stroke's distance to the clock center. Since the outer strokes were angularly separated, a large average span cannot be caused by two angularly distant numerals in the same group. The criteria is primarily designed to catch instances where the center dot is included as an outer stroke. The current clustering approach is prone to this error, as many individuals do not draw the center dot with the rest of the inner strokes. The temporal criteria catches groups that contain temporally consecutive strokes whose timestamps differ by more than 3 seconds. This criteria is designed to catch cases where the arrowhead of a hand is included as an outer stroke. If a group meets either criteria, all strokes in the group that are closer to the clock center than 0.6 times the average stroke's distance from the clock center are removed and added to the inner group.

Identifying joined and crossed out numerals

When there are strokes in a symbol that were not drawn consecutively or were drawn with more than a 0.3 second pause in between, the symbol is considered to potentially include two joined-together components or a cross out. The classifier splits the symbol into multiple sub-symbols using this temporal criterion. If a sub-symbol's convex hull has at least 20% overlap with a previous sub-symbol's convex hull, the classifier relabels the previous sub-symbol as a crossed out numeral. The classifier then runs a scribble detection algorithm on the overlapping sub-symbol to determine whether it should also be part of the crossed out numeral, or labeled as a separate numeral. An example of this process is shown in Figure 3-1.

The scribble detection algorithm first checks the sub-symbol's Ouyang score. It must be at least 3 standard deviations above the mean Ouyang score for actual numerals, indicating a poor match for a numeral. Next, it determines whether the sub-symbol can be judged to be a scribble used to cross out the previous sub-symbol. This determination is done by looking at pairs of points 7 points apart. If the direction between the two pairs of points changes by more than 120 degrees in magnitude, it is characterized as a sharp turn. If this occurs more than 35 times in a group of strokes, the group is determined to be a scribble, and the sub-symbol is also labeled as part of the crossed out numeral.

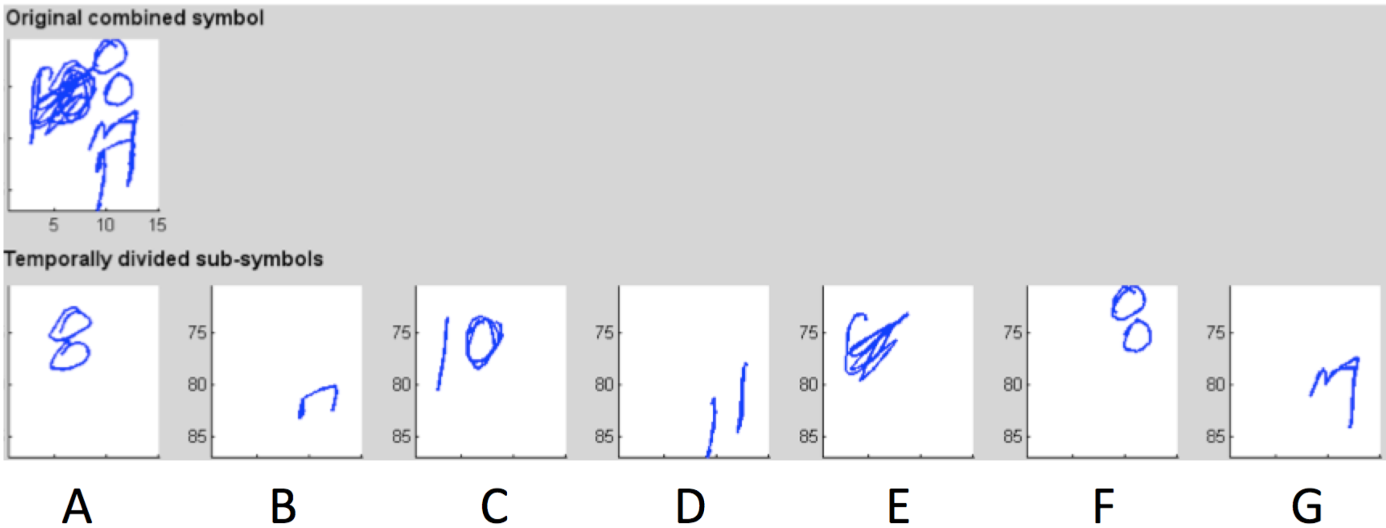

Figure 3-1: TCD3. The bottom images show seven sub-symbols generated from the top image by splitting between strokes that are non-consecutive or drawn with at least a 0.3 second pause. Sub-symbol E has more than 20% overlap with A and C, so A and C are labeled as cross outs. Sub-symbol E is detected as a scribble, so it is included with A and C as a cross out. Sub-symbol G overlaps with B and D, so B and D are labeled as cross outs. Sub-symbol G is not detected as a scribble, so it gets labeled as a 7 through the Ouyang recognizer. Sub-symbol F does not overlap with any other sub-symbol, so it is labeled as an 8 through the Ouyang recognizer. The original symbol ends up as fou

Identifying split numerals

This final error correction step is run to identify any numerals that were split into two symbols (e.g. an 11 being classified as two 1s) or any false positives from the previous error correction that detected joined numerals. The algorithm orders the numerals by angular position and looks for adjacent symbols that both contain only one stroke. It then combines the two strokes into one symbol and compares the best-match Ouyang score for the combined strokes to that of the old Ouyang scores. If the new Ouyang score (normalized by each numeral's mean and standard deviation) is better than the average of the old scores, then the strokes are combined into one symbol. Otherwise, they are restored to two separate numerals.

3.3.6 Results

Table 3-1 shows accuracy results for the previous ClockSketch classifier. The table presents the classifier's combined accuracy on the old data sets that it was empirically trained and evaluated on. Table 3-2 presents the classifier's combined accuracy on two new clinical data sets.

Table 3.1: Previous Classifier Performance on Old Data Sets

Old Data Sets	**Accuracy**	**# Strokes**	**# Containing Clocks**
All strokes	91.4%	18803	754
Numerals	96.4%	12967	752
Hands & Center Dot	87.3%	3871	753
Crossed out digits	49.0%	145	57
Noise	78.1%	576	232
Tick marks	0.0%	149	14
Clockfaces	90.8%	862	749
Other strokes	0.0%	233	53

Table 3.2: Previous Classifier Performance on New Clinical Data Sets

New Data Sets	**Accuracy**	**# Strokes**	**# Containing Clocks**
All strokes	72.4%	13548	488
Numerals	83.0%	8039	473
Hands & Center Dot	78.4%	2214	450
Crossed out digits	38.3%	175	76
Noise	76.7%	1092	291
Tick marks	0.0%	1033	117
Clockfaces	86.0%	584	486
Other strokes	0.0%	411	91

We see that the classifier performs reasonably on the old data sets, with an overall accuracy of over 90% primarily due to its good performance in classifying numerals. However, when evaluated on the new data sets, the performance significantly drops. This is partially because we chose clocks containing tick marks for one of the new data sets, and the classifier does not handle tick marks. The new data sets also feature a higher proportion of drawings with significant distortions, which accounts for some of the drop in performance for other symbols.

Figure 3-2 presents heat maps depicting errors of the classifier for both old and new data sets. Even though classification accuracy is highest on numerals on the old data set and second highest on the new data set, numerals still represent the largest source of error in absolute terms as they comprise the majority of strokes. We see that numerals are mostly misclassified as other numerals, and unsurprisingly due to their proximity, crossed out digits, tick marks, and clockface strokes are also misclassified as numerals. Hands seem to be equally misclassified as the other hand, digits, center dot, and noise, probably owing to the fact that their arrowheads are near the digits, while stray strokes near the base may be misclassified as another inner stroke or noise.

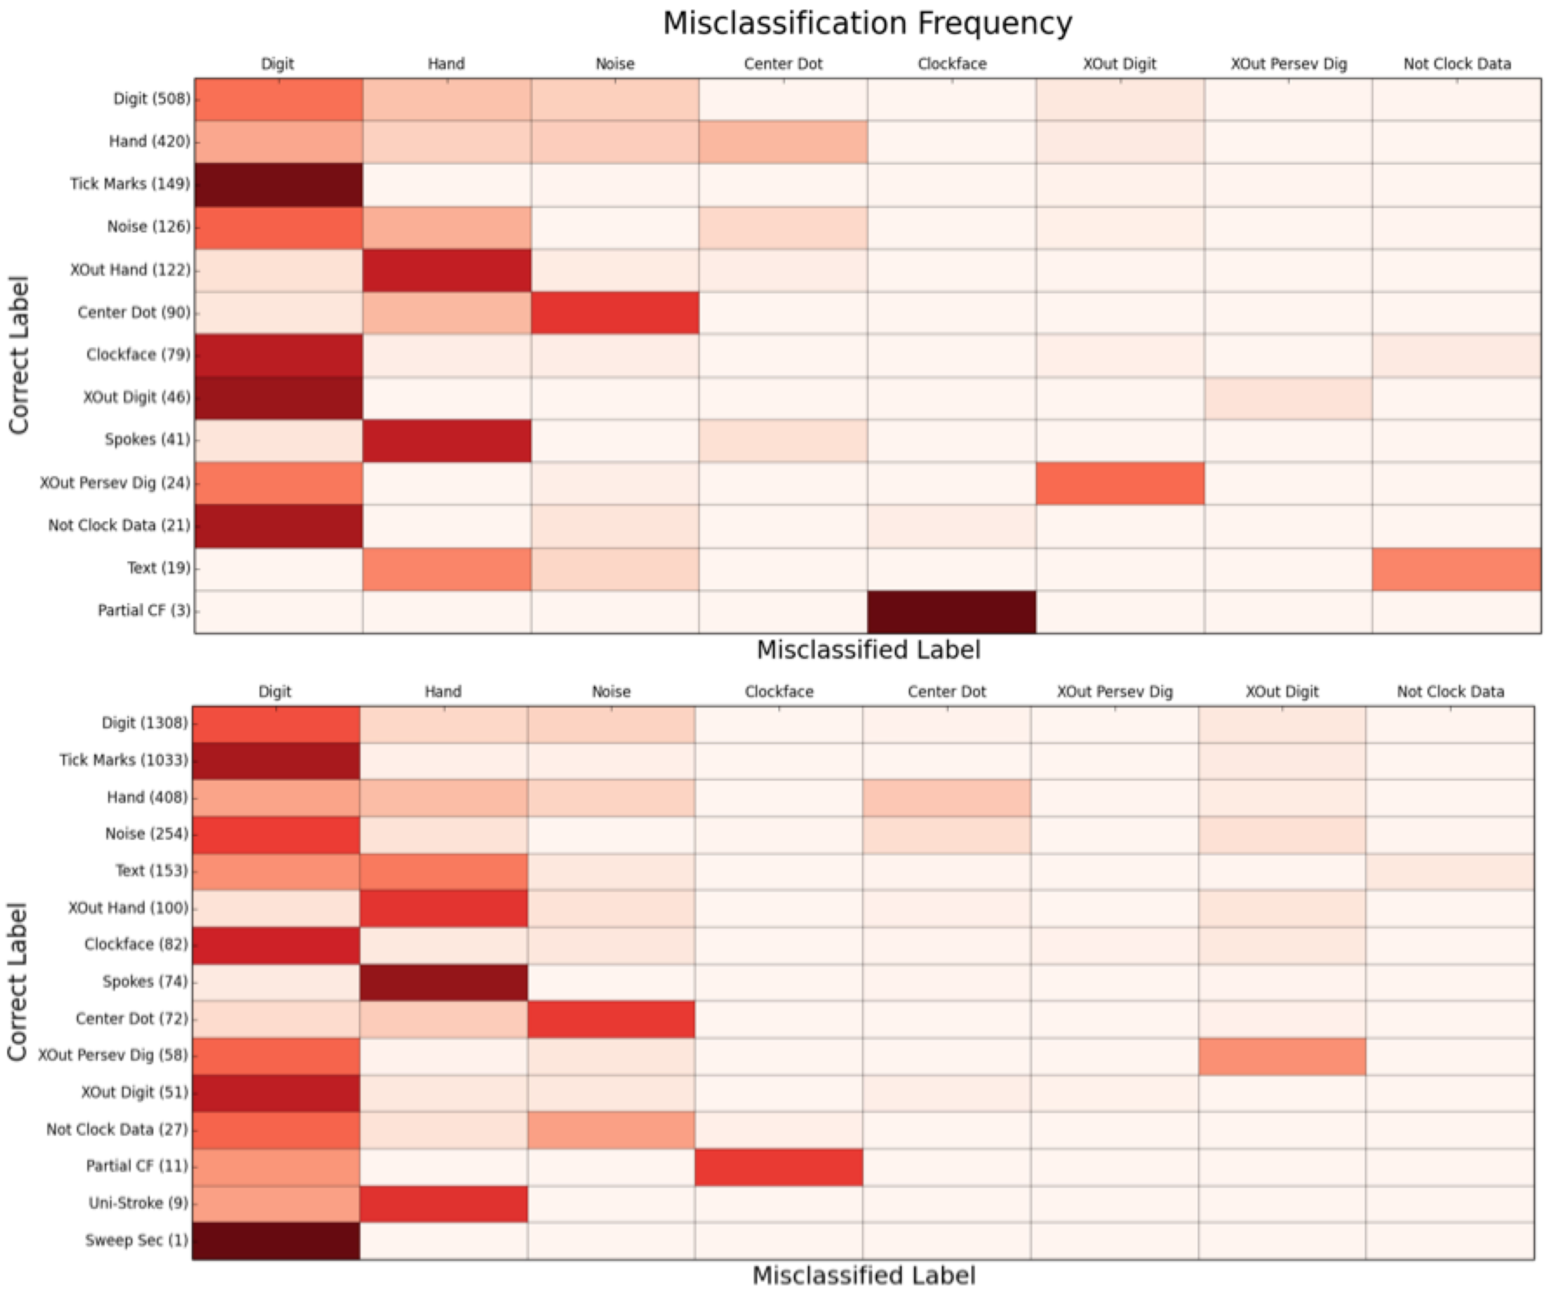

Figure 3-2: Heat maps depicting which label incorrectly classified strokes were assigned for drawings. The top and bottom heat maps show errors in the old and new data sets respectively. These maps only include misclassifications, not correctly classified strokes. Dark red indicates a high proportion of strokes with the label indicated by the row were mislabeled as a symbol indicated by the column while light red indicates a low proportion. The number in parentheses next to each row label is the number of misclassified strokes of that symbol type. For example, 508 digit strokes in the old data sets were misclassified mainly as other digits, but also as hands, noise, and crossed out digits to a lesser extent.

Chapter 4

Improvements to Previous ClockSketch Classifier

This chapter describes modifications designed to improve the previous performance of the ClockSketch classifier and evaluates the results of these modifications. One general focus of many of the improvements is to reduce the reliance on arbitrary parameter cutoffs like "0.6 times the average stroke length" by using more intrinsic properties such as shape. This approach makes fewer assumptions about how a symbol is drawn, which allows the classifier to robustly handle a wider variety of input.

4.1 Clockface Identification by Shape

4.1.1 Previous ClockSketch Behavior

Identification of the clockface is critical to the entire process of the previous ClockSketch classifier. For example, distance to the clockface center is used for clustering the inner and outer strokes, and angular position in the clockface is used for weighting the numeral's Ouyang scores and labeling the hands. As such, a mistake in identifying the clockface can cause many other strokes to be mislabeled.

Despite its importance, the current classifier handles only clockfaces that are drawn with one or two major strokes. In addition, it bases this determination strictly on the stroke's length and radius relative to other strokes, rather than looking for intrinsic properties of the symbol. In some drawings, the test taker writes down 11:10 at the top, relatively far away from where they end up drawing the clock. This greatly increases the drawing's bounding box size and average span, which can cause the clockface to not be recognized because of the text's presence even though nothing is different about the clockface stroke. Similar situations can occur if the test administrator writes the patient's information outside the designated area or if the test taker drops the pen.

Figure 4-1 shows three examples of failures to detect clockface strokes, ranging from a minor error to complete misidentification. Figure 4-2 shows an example of how failure to detect the clockface propagates into many other mislabels, affecting both numerals and hands.

Figure 4-1: Three cases where the classifier does not correctly identify all the clockface strokes.
Left: TCM7V3.5CIN2022489965. Despite the fact that the classifier is designed to recognize clockfaces composed of two strokes, it does not detect the smaller, highlighted stroke as the stroke's radius is less than 1/8 of the drawing's average span.
Middle: VIN0829757788. The clockface is composed of seven pieces, four of which are highlighted. The longest stroke is too short by the classifier's criteria, so the classifier concludes that there are no clockface strokes at all.
Right: CIN0300655383. The classifier labels the highlighted mantelpiece outline as the clockface as it is the longest stroke, mislabeling the actual circular clockface as a minute hand.

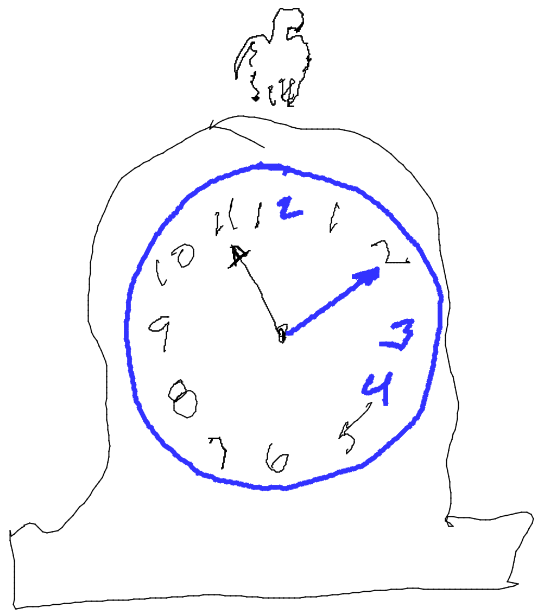

Figure 4-2: CIN0300655383. The outer mantelpiece outline was detected as the clockface causing many other mislabelings. All of the highlighted are are labeled as a single minute hand. The five, six, and seven digits are labeled as the center dot, due to the detected clockface's lower center compared to the actual clockface. Only four components of the clock are completely labeled correctly.

4.1.2 Improvement Using Inherent Stroke Properties

Our new clockface detection algorithm first tries to determine if there is a single stroke that makes up a complete clockface. This step covers the vast majority of our data set because the clockface is drawn as a single stroke in 89% of drawings. The algorithm first spatially resamples every stroke so that the number of points in each stroke remains constant. Resampling with too few points causes a loss in resolution, and resampling with too many points requires more data to be interpolated, so the current number of points is used as an approximation of the stroke's resolution. Next, we run an ellipse fitting algorithm as outlined in [Fitzgibbon et al., 1999] on each resampled stroke. If the goodness-of-fit or R^2 of the elliptical fit is at least 0.97, and the angular extent is at least 330 degrees, it classifies the resulting stroke as the Clockface.

If it does not detect a clockface in the previous step, it looks for a clockface composed of multiple arcs as follows:

1. For each stroke, compute the direction between successive points. For an arc

of a circle, the direction values should be well approximated by a line with non-zero slope. Figure 4-3 shows the direction graph and best fit line for two potential clockface strokes in a drawing.

2. If the direction values can be approximated by a line with Rˆ2 at least 0.9, we fit a circular arc through the stroke's first point, last point, and the point where the stroke intersects the perpendicular bisector of the line segment from the first to last point of the stroke. Figure 4-4 depicts an example of this process for two different arcs in a single drawing. For each fit arc, we note the extrapolated center, radius, and angular start and end.

3. The fitted arcs are sorted by their starting angular position. We now begin the process of merging arcs that overlap in angular position, starting with the overlapping pair of arcs that have the smallest distance between their extrapolated center points and radii. When a pair of strokes is merged, we refit an arc to the resulting merged stroke. If the resulting stroke has angular extent 330 degrees with elliptical Rˆ2 at least 0.97, we classify the set of strokes as the clockface.

4. If after merging is finished, we still do not meet our clockface criteria, then we check among all strokes to see if adding any single one of them would result in a clockface, with the intention of coveirng cases where a line or other non-curved stroke is drawn to repair the clockface.

In very rare cases where a patient crosses out the entirety of their previous work and starts anew, multiple clockfaces are detected. Since we do not handle crossed out clockfaces, we choose the later drawn clockface when multiple clockfaces are detected.

With the arc fitter, we can also recognize Partial Clockfaces, a label unhandled by the previous ClockSketch classifier. This occurs when the first stroke of a clock is a partial circle, and the test taker later draws a stroke that is a fully enclosed clockface (as opposed to drawing arcs to "stitch" the partial clockface). After detecting the clockface, we run our arc fitter on the remaining strokes to see if any arcs have a similar extrapolated center and radius to the actual clockface. Partial Clockfaces are

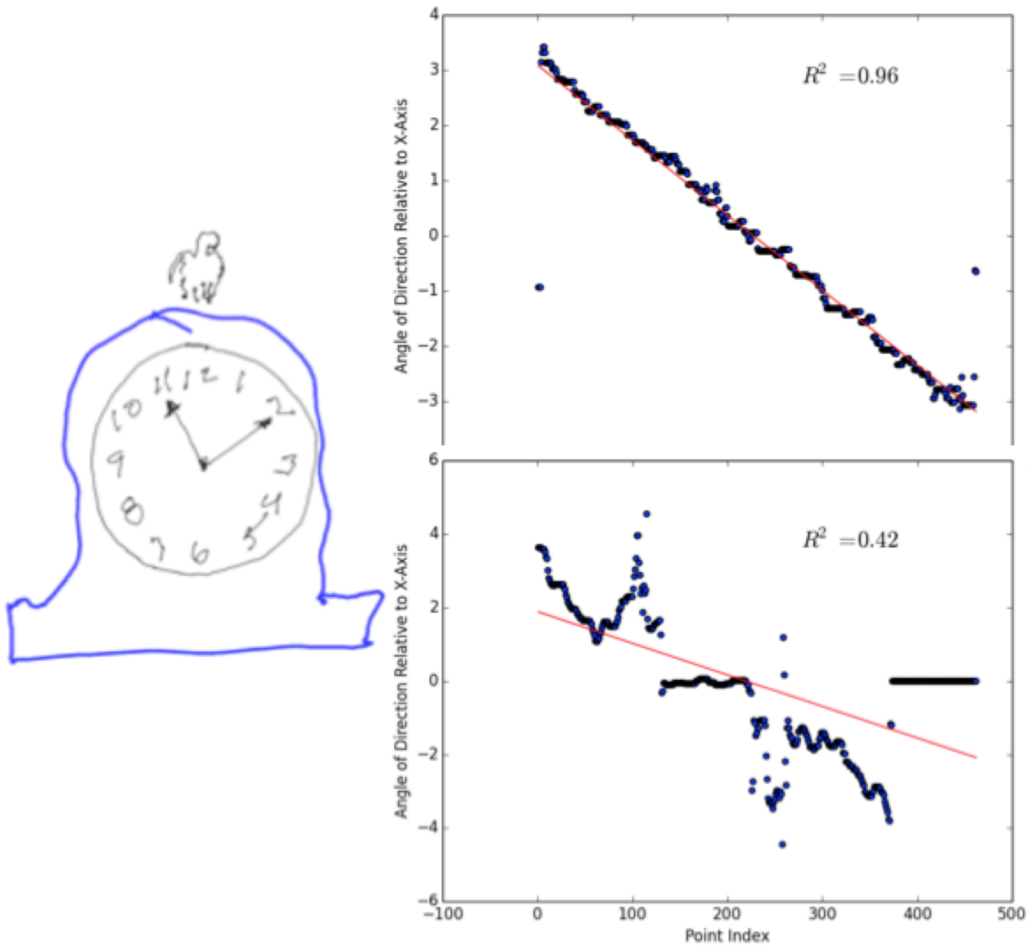

Figure 4-3: CIN0300655383. The angular direction between successive points is plotted for the clockface (above) and the highlighted stroke (below). The best fit lines are shown in red along with the R^2 value. While the previous classifier labeled the highlighted stroke as the clockface, our new classifier does not as it is a poor match for an ellipse. The circular clockface's direction values closely approximate a line.

a rare phenomenon, only appearing in 14 or 1% of our drawings, but we now have the capability to easily recognize them.

With this algorithm, we successfully identified a clockface in over 98% of drawings, and labeled over 92% of clockface strokes in the new clinical data sets correctly, an improvement from the previous classifier's 86%. Undetected clockfaces occur only in rare cases, where strokes were drawn with a lot of retracing (which confuses the direction graph). when the clockface did not meet our definition of a clockface, either because it was not an ellipse, or because its angular extent was less than 330 degrees. Figure 4-5 presents two examples of undetected clockfaces. The remaining undetected

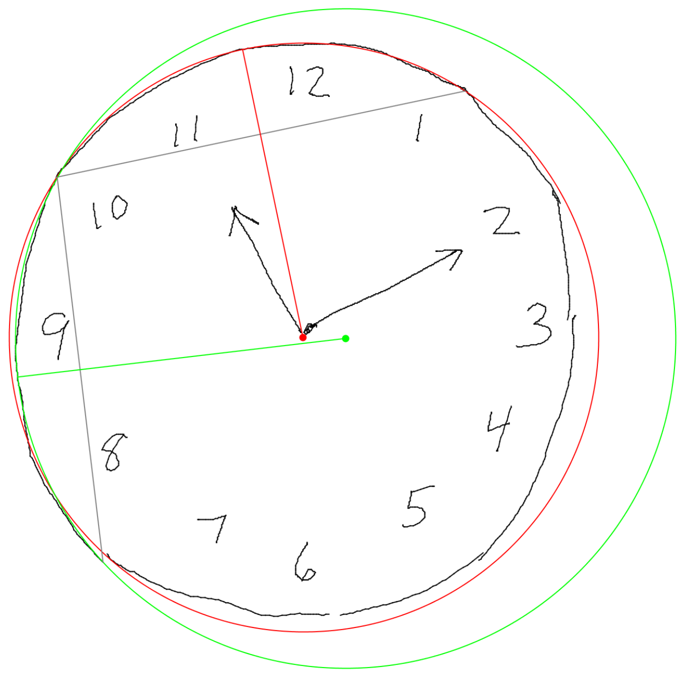

Figure 4-4: VIN0829757788. The extrapolated center and radius from our arc fitting algorithm are shown for the two longer arcs that make up the clockface. The top arc's radius and center are in red, and the left arc's are in green. The grey lines connect the stroke's first point to its last point. The radii perpendicularly bisect the grey lines as the third fitting point is the intersection between the grey line's perpendicular bisector and the stroke. We see that the top arc's center and radius are a very good approximation for the overall clockface's center and radius, and the left arc's center and radius is similar.

clockface strokes are relatively short or linear strokes that are not well approximated by an arc, and their presence does not make the difference between the clockface having an angular extent of 330 degrees or not.

4.1.3 Possible Further Improvements

Our new approach could be further refined to allow for retracing by looking at curvature, or the change in direction between successive pairs of points. If the curvature switches signs, it indicates the stroke has effectively switched directions (e.g. clockwise instead of counter-clockwise). We could deal with retraces by using the sign of the curvature to segment the stroke into different portions, and look only at the portion that spans the largest direction range. However, curvature can be sensitive to noise or small perturbations and would require window smoothing. Since a clockface

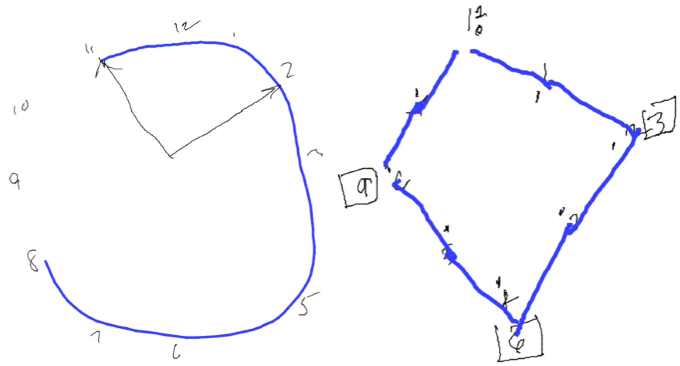

Figure 4-5: Left: CIN0246559773. The angular extent of the clockface is 270 degrees, lower than our requirement of 330 degrees. The stroke was long enough to be labeled as a clockface by the previous classifier.
Right: CIN1480529293. The clockface's shape is a diamond. Since it consists of nine short strokes, the old classifier also fails to label it properly.

is already found in over 98% of drawings, we did not pursue this refinement as it would likely require significant parameter tuning to pick up only on retraces without false positives.

The one case where the old approach is able to detect a clockface but the new approach is not is in the case of clockfaces that do not meet our new definition, which requires a round face with large angular extent. Since this case is rare, occurring in fewer than 0.5% of drawings, we do not account for it. However, we could handle this case by returning the longest stroke in the drawing provided it meets the previous classifier's clockface criterion if our new clockface detector fails. This would cover the left case in Figure 5-5, but not the right one. The right case is particularly challenging and is the only drawing in our data sets where both the previous and current algorithm fail to detect a single clockface stroke. It could thereotically be handled by an algorithm similar to the current classifier's arc merging when both detection algorithms fail. Instead of fitting arcs, we directly calculate each stroke's angular extent relative to the drawing's center of mass and merge based on that. Since the combined angular extent is greater than 330 degrees, this clockface should be detected, but it is not worth developing a new algorithm for a single case.

4.2 Elliptical Clockfaces

4.2.1 Previous ClockSketch Behavior

Distance to the clockface center is used for clustering the inner and outer strokes and determining whether an inner stroke is a center dot or hand. Similar to how incorrect clockface detection can propagate into many more errors, poor inner/outer clustering can add non-numeral strokes to numeral groups, messing up Ouyang scores and mislabeling digits.

Many clockfaces are drawn oblong rather than circular, as for example in Figure 4-6. The average ratio between a fitted ellipse's major and minor axes is 1.16 and ranges from 1 (circular) to just under 2 (squashed oval). In a highly eccentric oval, the distances between the clockface center and numerals drawn around the clockface perimeter will vary depending on a numeral's angular position. In some cases, a hand's arrowhead may actually be farther from the center than some numerals due to the clock's distortions. As the tips of the hands are generally close to the numerals, these are especially sensitive to small variations in distance. This can degrade the performance of the k-means clustering, leading to poor segmentations.

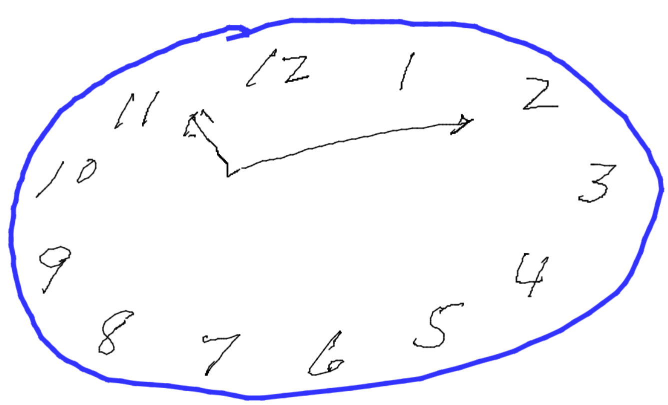

Figure 4-6: VIN1082466239. An oblong clockface where the 6 is closer to the center than the arrowhead of the minute hand.

4.2.2 Using a Distance Radius Ratio

To deal with problems caused by eccentric clockfaces, we fit an ellipse to each clockface and replace distance to the clockface center with the ratio between the distance to the center and the ellipse radius at that angular position. Strokes near the clock outline thus have a ratio near 1, regardless of the absolute distance. This accounts for oblique clockfaces as long as the clockface is recognized by our clockface detection algorithm.

For clocks without detected clockfaces, it is impossible to account for the clockface's obliqueness. Instead of the distance-radius ratio, we use the drawing's center of mass as the center and scale the mean distance to the center of all strokes to 1. This fallback provides values close to 1 when a clockface is not detected, which is important for consistency in tuning parameters such as the distance-to-time scaling parameter in k-means clustering. Since only relative distances matter for the purposes of clustering, this method still performs similarly to the old method of using absolute distances when we cannot detect a clockface.

As the repair steps and k-means weighting and parameters are specifically tailored to the current distance metric, it is difficult to quantify the exact improvement using the distance ratio alone brings, as the repair steps and k-means parameters would have to be reoptimized for this new metric. We note that the current k-means clustering system properly separates strokes into inner and outer categories for 77% of the 1242 drawings. Improperly clustered drawings have an average of 2.86 clustering errors per drawing. For clocks with a major to minor axis ratio of at least 1.2, the number of properly divided drawings drops down to 72% while the average number of clustering errors increases to 3.05. For even more elongated clocks with an axis ratio of at least 1.5, the clustering accuracy drops to 46% with an average of 3.86 clustering errors for incorrectly clustered drawings. This provides strong evidence that highly oblique clockfaces are not handled well in the current system. We use the distance ratio metric rather than absolute distance as a feature for later improvements and new classifiers.

4.3 Entropy

We incorporate entropy as a property to distinguish between different types of inner strokes and for training various classifiers as described later. Entropy is a measure of diversity; we apply it to the angles formed by consecutive points of a stroke. It has been used to successfully distinguish textual annotations from shapes in drawings with over 90% accuracy as most shapes have a much lower variation in angle between successive points than does text. [Bhat and Hammond, 2009] For our uses, we partition the possible angles into six bins (0 to 60 degrees, 60 to 120 degrees, etc.) and include stroke endpoints as the seventh bin. We calculate the angle formed by successive sets of three points after spatial resampling and record it in the corresponding bin. We measure information entropy of the distribution of the angles by computing:

$$H = -\sum_i x_i \log x_i$$

where x_i is the frequency of bin i.

A symbol like a scribble that has a uniform distribution of angles will have a high entropy while a symbol like a circle will have a low entropy because the angle formed by successive points is almost always close to 180.

We present a histogram of symbol entropies in Figure 4-7. As expected, clockfaces, which are generally circular, exhibit very low entropy. Hands, which are mostly a straight line but often include arrow heads are slightly higher. Center dots, which are often scribbles, have high entropy with numerals falling in the middle. The histogram illustrates that entropy can be an effective way of distinguishing between different types of strokes.

4.4 Tick Mark Labeling Support

The most common clock component that is completely unhandled by the previous classifier are tick marks. Tick marks are lines placed around the clockface outline in place of or in addition to numerals. While only a minority of clocks contain tick marks, clocks that do contain an average of 9.1 tick mark strokes. The presence of

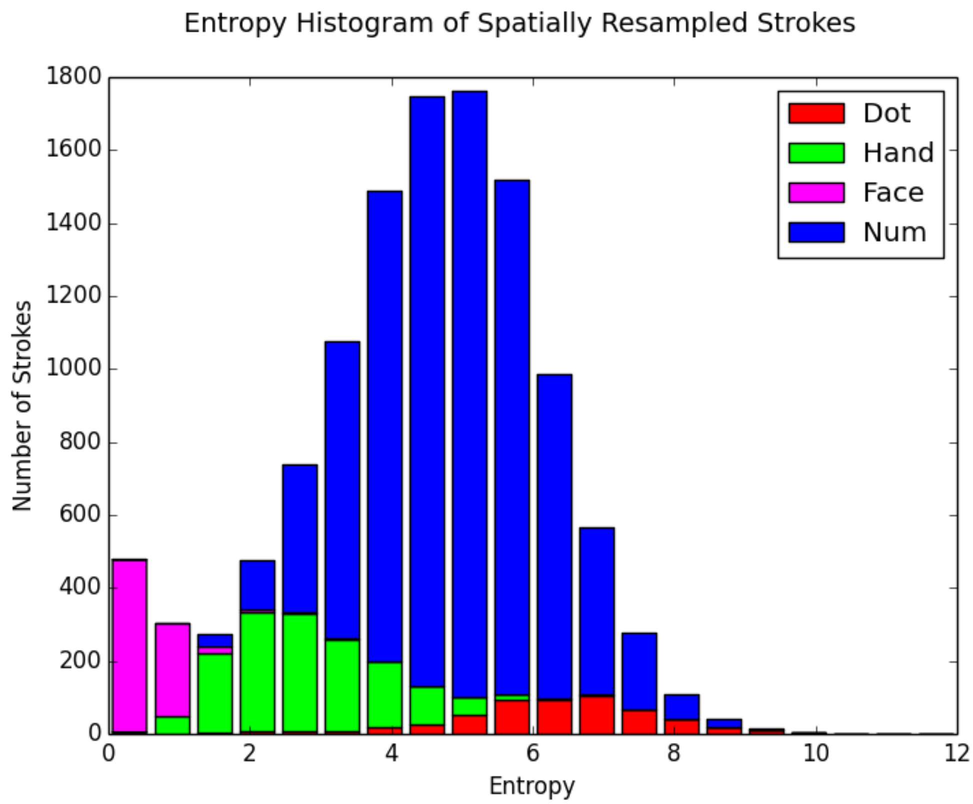

Figure 4-7: A histogram showing the number of strokes with a specific entropy. The entropy values are relative and were scaled to have a maximum value of 12. Dot refers to center dots. Hand refers to hour or minute hands. Face refers to the clockface. Num refers to numerals 1 to 12.

these additional strokes impedes many steps of the ClockSketch classifier, particularly numeral segmentation and labeling as the tick marks are drawn in the same position. Numeral, hand, and clockface classification accuracy for clocks with tick marks present ranges from 70% and 75%, with overall classification accuracy dropping to 49.8%.

4.4.1 Feature Selection

As tick marks, like numerals, are drawn around the clockface perimeter, we focus on features that allow us to distinguish between numerals and tick marks. We provide a list of the features, and a rationale for selecting each one below:

Ouyang scores to each of the 12 numerals : Numerals should have a low Ouyang score to their number. Since tick marks are often drawn perpendicular to the clockface, they are often rotated and would not necessarily have a low match distance for the numeral 1.

Mean and minimum distance to the clockface : Tick marks often intersect the clockface and are sometimes drawn so that they are bisected by the clockface, so they should have low mean and minimum distances to the clockface.

Closest and farthest distance ratios to clockface center : Numerals are almost always drawn entirely within the clock or outside of it, so both distance ratios should be less than or greater than 1 while tick marks often intersect the clockface.

Line goodness-of-fit : Tick marks are generally straight lines, which should be approximated well by a linear fit.

Entropy : Straight lines should also have low entropy, although entropy is significantly less sensitive to a single change in direction (as in a 7), which would ruin the linear fit, as it only looks at angles between successive points rather than relative positions.

Length : Numerals should have longer path lengths than tick marks of the same scale.

4.4.2 Performance

The feature vectors for all tick marks and numerals were extracted and used to train a probabilistic binary SVM with a RBF kernel with parameters C = 1 and gamma = 0.01. For a given set of features, the SVM outputs the probability that the stroke is a tick mark. If the probability is higher than 0.5, we classify the input a tick mark. The classifier was trained and tested using 10-fold cross validation. It labeled 91.2% of tick marks and 4.7% of numerals as tick marks, for an overall accuracy of 93.8%.

4.4.3 Further Improvements

We noted that tick marks often come in multiples, averaging over 9 in clocks with tick marks but did not use this fact in our SVM classifier. Likewise, we note that for a given drawing with tick marks, the tick marks are usually either consistently halfway between digits, or consistently at the same angular position as digits. By using angular spacing and considering the likely classifications of other numerals, we should be able to increase the accuracy of our tick mark classifier. Since both of these properties involve considering an individual stroke in the context of the other outer strokes, we incorporate them as contextual rules, which we discuss in the following chapter.

Chapter 5

Local Classifiers and Global Evaluation System

This chapter describes a global evaluation system that incorporates context to improve the current ClockSketch classification accuracy. It describes new symbol scorers and common sense rules, two major components of the new system.

5.1 Context and System Motivation

A powerful tool that has not been used so far is context. In many drawings by impaired patients, isolated recognition of certain symbols is impossible due to heavy distortions. In these cases, context is crucial in properly labeling the symbol. We characterize two types of context to improve the classification of strokes. The first is a stroke's local spatiotemporal context, which incorporates information about other strokes that were drawn directly before, afterwards, or spatially close by. The second is a global context, which uses information about every other stroke in the drawing. Three examples of situations where spatial context, temporal context, and global context respectively are important to correctly label a stroke are depicted in Figure 5-1.

Previous work on incorporating context has largely focused on conditional random fields (CRFs), which take into account neighboring inputs when determining a label

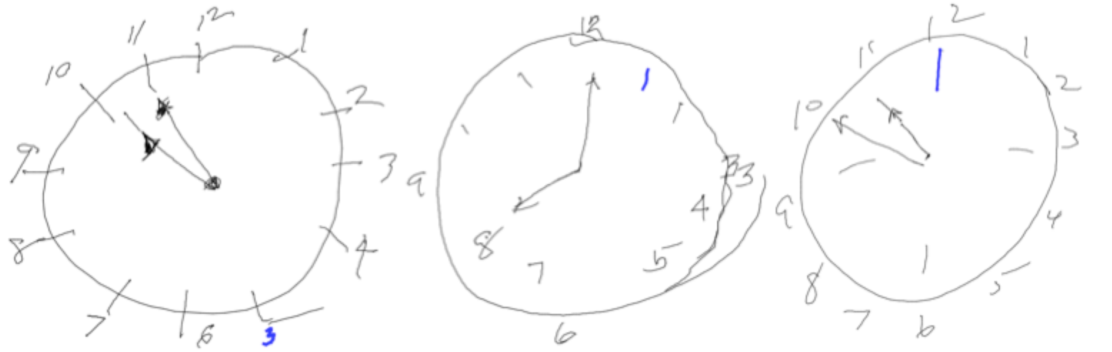

Figure 5-1: Three cases where context is important to identifying the highlighted stroke. Left: CIN1374758076. The highlighted stroke most resembles a 3. However, the adjacent numeral groups most resemble a 4 and 6. The spatial context correctly indicates that the highlighted stroke is probably part of a 5 symbol.
Middle: CIN0584487598. The highlighted stroke resembles a 1 and is also in the correct angular position to be a 1. The stroke was drawn immediately after the similar looking strokes in the 10 and 11 angular positions and immediately before the similar stroke in the 2 position. Since all four strokes are likely matches for tick marks, the temporal context indicates that the stroke is likely a tick mark.
Right: CIN1856010988. The highlighted stroke is the second stroke drawn and is drawn immediately after the clockface. Many patients draw the 1 stroke of the 12 symbol immediately after drawing the clockface, and this stroke resembles a 1. However, every numeral was drawn outside the clockface, and there are other tick marks present. The global context about the other numerals and tick marks indicates that the highlighted stroke is likely a tick mark.

for a given input. Work specifically on labeling properly segmented numerals from the clock drawing test achieved >99.9% performance on healthy data and 96.4% performance on impaired data.[Song and Davis, npub] We choose not to pursue the use of CRFs for the ClockSketch classifier for a few reasons. The first is the additional complication from non-numeral strokes. A hand's arrowhead or a tick mark are often closer to a numeral than the neighboring numerals are, introducing the additional problem of neighbor determination. Next, while the previous classifier's accuracy on numerals is 98.1% on the healthy YDU-100 data set and 91.3% overall, the errors are primarily due to segmentation. When applying the previous numeral labeling rules, which use angularly weighted Ouyang scores, on properly segmented numerals, we achieve 97.4% accuracy on the 14518 digits in our data sets. Since the large majority of our drawings comes from impaired patients, this indicates that the amount of improvement possible from a conditional random field approach is insignificant when compared to the errors caused by improper segmentation.

Instead, we propose incorporating a set of rules to capture a human's intuition or common sense about what a stroke is likely to be given the other symbols in the drawing. However, while a common sense rule like adjacent numerals should be consecutive in value is a good rule of thumb, there are exceptions to this rule, especially in cognitively impaired patients, so these rules must be appropriately weighted against what the stroke is likely to be based on appearance. To come up with a numeric weighting, we first need a score that represents an isolated symbol recognizer's confidence in a label. We can then run a parameter sweep on the magnitude of how much each rule should adjust the isolated scores and see which values results in optimal performance.

5.2 Hands, Center Dot, and Clockface Scorer

Since we already have a scorer for numerals and tick marks, only the hands, center dot, and clockface remain. Note that the scorer's job is to output a numeric score that represents the likelihood of a symbol's label being correct, not to actually label the symbol. While the two problems might seem analogous, as an algorithm that can perform the labeling could be modified to give probabilities for each candidate, the following example should illustrate why they are different problems.

Given a properly segmented group of strokes and the information that its label is either a center dot, hand, or clockface, it would probably not be difficult to come up with rules that achieve 100% labeling accuracy due to the large differences in shape and size between a small center dot, large clockface, and linear hand (relative angular positions and lengths of the hands could be used to decide between hour or minute hand). Now, augment the group with an additional small stroke that does not impact the shape fit, bounding box, or any other shape or size properties. A labeler would still label a center dot with the additional stroke as a center dot, a hand with the additional stroke as a hand, etc. However, a likelihood scorer may recognize that hands sometimes have a similar additional inconsequential stroke (e.g. a single part of the arrowhead) but center dots do not, and adjust the probabilities accordingly even if the magnitude of the adjustment is small. The desired system should output match scores as the Ouyang symbol recognizer does rather than relying on rules and cutoffs like the previous ClockSketch classifier.

5.2.1 Feature Selection

For a given symbol, we extract a feature vector as follows:

1. Determine the symbol's closest and farthest points to the clockface center (the drawing's center of mass is used if a clockface is not found).

2. Rotate the symbol so that the segment defined from the closest point to the farthest point points directly upright.

3. Scale and rasterize the symbol into a 100x24 raster. Apply Gaussian blurring and downsampling by a factor of 2 to the resulting raster as is done in the Ouyang symbol recognizer.

4. Reduce the dimensions of the image down to 5 by applying principal component analysis (PCA).

5. The angle of rotation, bounding box dimensions pre-rotation, and bounding box dimensions post-rotation are used along with the 5 image principal component values to construct a feature vector. Each feature is normalized to have zero mean and unit standard deviation.

5.2.2 Performance

The feature vectors for all center dots, hour hands, minute hands, and clockfaces in our data sets were extracted and used to train four probabilistic binary SVMs with a RBF kernel with parameters C = 10 and gamma = 0.1. For a given input and label, we use the probability of a positive labeling from the corresponding SVM as the score.

To quantify the performance of our scorer, we ran 5-fold cross validation and recorded the percentage of symbols where the corresponding SVM of the symbol's label outputted the highest score of the four SVMs. That percentage was 99.7% for clockfaces, 94.5% for hour hands, 95.1% for minute hands, and 97.8% for center dots. This performance is comparable to that of the Ouyang recognizer for properly segmented numerals.

5.3 Segmentation Repair with Perturbation

We now have symbol recognizers that can score numerals, hands, tick marks, the center dot, and the clockface. Using these symbol recognizers, we create a segmentation repairer that does a local search to maximize the average symbol score within a drawing.

While the previous ClockSketch classifier did attempt to repair joined, split, and incorrectly segmented numerals, it only covered numerals and also relied on parameters that did not necessarily apply to every patient. For example, any pause of greater than 0.3 seconds between strokes in a symbol was tested for splitting. However, the median time a patient pauses between strokes ranged from 112 milliseconds for the quickest patient to 1.8 seconds for the slowest. The repair strategy would have false negatives for quick patients, and false positives for slow patients.

We first normalize all our recognizers by evaluating them on the symbols they are trained to classify. We normalize it so that each recognizer's mean score for a symbol of its type is zero with unit standard deviation, resulting in z-scores. For the Ouyang numeral recognizer, we output the negative of its z-score, as a higher Ouyang score indicates a worse match. Our perturbation process has 7 steps:

1. Evaluate a score for each symbol using the scorer corresponding to the label assigned by the ClockSketch classifier.

2. Randomly choose a symbol among those that have negative scores (i.e. those whose label seems least definitive), weighting each symbol by the absolute value

of its negative score.

3. Determine the closest four symbols to the chosen symbol.

4. For each of the chosen and close by symbol pairs, iterate over every stroke in the pair. For each stroke, move the stroke to the other symbol and rerun each recognizer on the resulting symbols, taking the maximum score.

5. If the resulting two scores are higher than the initial ones, accept the change. If only one symbol remains as the only stroke of an original symbol was moved, compare the new score to the average of the initial ones. If the resulting scores are more than 1 point worse, reject the change. Otherwise, accept with probability 1-d, where d is the resulting drop in scores.

6. If the symbol can be split into two sub-symbols that are at least 3 millimeters apart, do so and rerun each recognizers on the resulting sub-symbols. If the resulting average is better than the initial score, accept the change. If the resulting scores are more than 1 point worse, reject the change. Otherwise, accept with probability 1-d, where d is the resulting drop in scores.

7. Keep track of the resulting classification with highest average score, and return after a set number of iterations.

The algorithm attempts splits, joins, and various segmentation possibilities. Since it reevaluates after each segmentation move, it handles errors that require multiple segmentation repairs without over perturbing errors that are solved by a single repair.

5.4 Global Context Evaluator

Now that we have a segmentation repairer that does a weighted-random search based on score, We now create a global context scorer that works in tandem with our individual symbol scorers. Instead of merely optimizing for maximizing the individual label scores, we can maximize for the sum of our global context and individual scores, allowing us to incorporate common sense contextual rules.

5.4.1 Common Sense Rules

We list the common sense rules below, along with a quantitative sense of how universal each rule is.

- The clockface should be a closed shape composed of arcs.
  1225/1237 (99%) clocks are more than 350 degrees closed.

- Adjacent digits should be consecutive in value.
  Of clocks with each digit exactly once, 981/1064 (92%) have all digits following this rule. 66/1064 (6%) have one violation, and 17/1064 (2%) have multiple.

- Each digit should be represented once before any digit is represented twice.
  Only 26/1242 (2%) clocks have a missing digit and another digit twice.

- All the digits generally lie entirely within or entirely outside the clockface.
  1138/1232 (92%) have all digits entirely within (n=1087) or outside (n=51) the clockface. Of the remaining 94, 30 have exactly one mismatch (e.g. 11 digits inside but 1 outside) and 60 have multiple.

- When digits are drawn outside the clockface, they are typically accompanied by tick marks in the same angular position.
  73/120 (61%) of clocks with more than one digit outside the clockface have tick marks.

- Both minute and hour hands should be represented once before any is represented twice.
  No clocks violate this rule. 63/1242 (5%) only have one hand and 52/1242 (4%) have more than two.

- The minute and hour hands should not differ significantly in number of strokes.
  Of the 1125 clocks with exactly one minute and one hour hand, 521 (46%) differ by 0 strokes, 413 (37%) by 1, 133 (12%) by 2, 41 (4%) by 3, and 17 (2%) by 4 or more.

- The hour and minute hand strokes should have endpoints that meet within the center dot if present.

  Of 744 clocks with center dots, 575 (77%) have both hands with endpoints within 2 millimeters of the center dot or within it. Of 381 clocks without center dots, 321 (84%) have both hands with endpoints within 2 millimeters of each other.

- Tick marks should come in groups, especially multiples of 4.

  Of 131 clocks with tick marks, 36 had exactly 12 tick mark strokes, 13 had 8, 12 had 1, 11 had 9, 10 had 4, 6 had 9, 9 had 11, 7 had 13, 6 had 2, with every other number appearing in 3 clocks or fewer. Note number of strokes is not exactly the same as number of tick marks as the sketcher could use two strokes to draw over or repair a single tick mark (so 13 strokes probably corresponds to 12 tick marks).

- Tick marks should consistently be between digits or in the same angular position as them.

- Tick marks should also consistently be roughly bisected by the clockface or not be.

  The above two rules are based on angular positions and lengths rather than a binary statement that holds true or not. As such, it is tricky to quantify how universal they are. Instead, what we do for the first rule is to measure the variance in angular spacing between successive tick marks. For the second rule, we take the location of intersection between the tick mark and the clockface normalized by the length of the tick mark (so a perfect bisection would intersect at 0.5). We treat a non-intersection as 1 and again quantify the variance in intersection locations.

5.4.2 Weighting Strategy

We determine appropriate weights for each rule by rerunning the segmentation repairer with one of the above common sense rules at a time. At each step, we apply the rule if it is relevant to the symbol type, penalizing as needed.

For example, the rule that adjacent digits should be consecutive in value is implemented as a potential penalty on the Ouyang numeral score. Call the value of the penalty x. Labeling a numeral between a 3 and 5 as a 5 would subtract 2*x from the 5's Ouyang score as it's 1 off from being consecutive for both sides. Similarly, a 6 would also subtract 2*x as it's 2 off from being consecutive on the side of a 3. A 7 would subtract 4*x while a 4 carries no penalty.

We write each rule as a penalty in proportion to the amount it is violated by. An example for each symbol type follows. To enforce the rule that the hands should not differ significantly in number of strokes, we penalize k*d to the average of the hand scores, where d is the number of strokes they differ by. To enforce the rule regarding tick mark's consistent angular positions, we penalize in proportion to the variance of the angular spacing. To enforce the rule that the clockface should be completely enclosed, we penalize in proportion to the difference between the clockface symbol's angular extent and 360.

Each rule has a parameter that must be tuned appropriately for optimal performance. To tune each parameter, we run a binary parameter search for values between 0 and 10 and until there is no longer an appreciable difference in performance for the rule's symbol group. There is the possibility that multiple rules encapsulate the same situation. In this case, optimally weighing each rule individually would result in a combined set of rules that is overly rigid toward satisfying the rules and does not adequately take into account the recognizer's scores.

To combat this, we classify each rule as applying to the clockface, numerals, hands and center dot, or tick marks. After using each rule individually, we now use each set one at a time. While we keep the relative weighting of parameters between rules in a set, we multiply the weights of every rule within the set by a parameter from 0 to 1, which is also found through a binary parameter search until the set's symbol group no longer has any appreciable change in performance.

With parameters for each rule in the global context evaluator, we report our final performance and results in the following chapter.

Chapter 6

Final System Performance and Conclusion

In this chapter, we present and discuss results for the new ClockSketch classifier. We also discuss future improvements and possibilities to improve the classification.

6.1 Overall Performance

Tables 6-1 and 6-2 summarize the performance of the new classifier on the old and the new data sets. Figure 6-1 depicts a heat map for aggregate errors made over all data sets.

Table 6.1: New Classifier Performance on Old Data Sets

Old Data Sets	**Accuracy**	**# Strokes**	**# Containing Clocks**
All strokes	90.2%	18803	754
Numerals	93.1%	12967	752
Hands & Center Dot	88.8%	3871	753
Crossed out digits	42.0%	145	57
Noise	78.1%	576	232
Tick marks	89.3%	149	14
Clockfaces	93.8%	862	749
Other strokes	0.0%	233	53

Table 6.2: New Classifier Performance on New Clinical Data Sets

New Data Sets	Accuracy	# Strokes	# Containing Clocks
All strokes	81.1%	13548	488
Numerals	84.9%	8039	473
Hands & Center Dot	81.3%	2214	450
Crossed out digits	35.4%	175	76
Noise	76.7%	1092	291
Tick marks	88.5%	1033	117
Clockfaces	92.9%	584	486
Other strokes	0.0%	411	91

6.2 Comparison of Results

We see that hands, center dot, clockface, and tick mark detection significantly improves across both data sets. Crossed out digits decrease across both data sets while numeral classification accuracy increases in the new data sets and decreases in the old ones. Of particular interest are the decreases in accuracy, especially as numerals are the vast majority of the data set.

The heat map makes it clear why there is a drop in numeral accuracy, We see that an almost equal number of digits are now misclassified as tick marks as they are other digits. While many of the digit mislabelings were corrected by the new stroke repairer and context evaluator, the fact that many digits are now being classified as tick marks cancels that effect out. The tick mark classifier was trained only on clocks containing tick marks, which are less than 25% of the data, even including the data set specifically targeting clocks with tick marks. As a result, the current tick mark classifier is probably significantly overaggressive in classifying strokes as tick marks. Numeral strokes outnumber tick mark strokes by a factor of over 17 to 1, so a 1% increase in numeral accuracy and 17% drop in tick mark accuracy will still improve our overall accuracy.

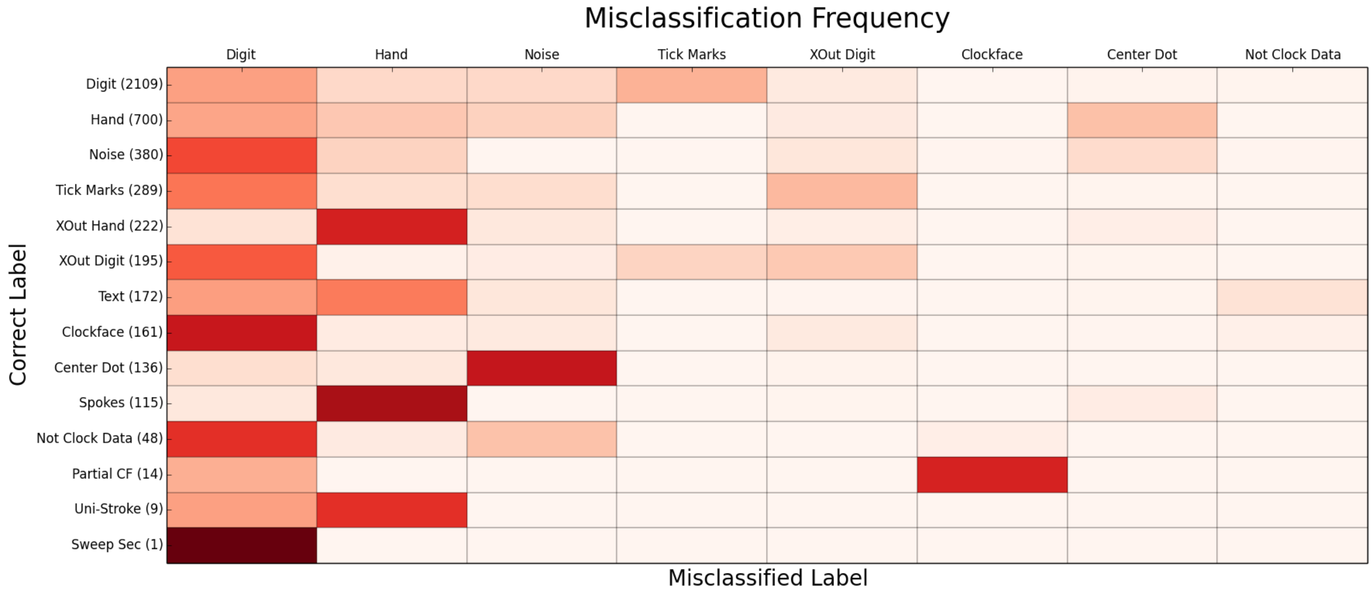

Figure 6-1: A heat map depicting which label incorrectly classified strokes were assigned by the new classifier for drawings in both data sets combined. This map only includes misclassifications, not correctly classified strokes. Dark red indicates a high proportion of strokes with the label indicated by the row were mislabeled as a symbol indicated by the column while light red indicates a low proportion. The number in parentheses next to each row label is the number of misclassified strokes of that symbol type.

6.3 Future Directions

6.3.1 Other Classes of Strokes

The only other stroke types that make up more than 0.1% of the data set are crossed out hands (0.7%), text (0.54%), and spokes (0.36%). While it is impossible to account for every type of symbol that a patient could draw, a general cross-out detector would be of interest given that it could encompass these two as well as boost cross out digit detection, which is still poor at the moment.

Cross out detection is a tricky problem since the content of both the underlying stroke and overwriting stroke affect the resulting classification. Three examples of overwriting a 0 numeral are scribbling over it, drawing a different numeral over it, or augmenting it with another 0 forming an 8. In the first case, the scribble must be classified along with the 0 as a crossed out numeral. In the second case, the 0 alone is a crossed out numeral but the new digit is labeled as normal. In the final case, neither stroke is a crossed out numeral. A generic cross out detector would need to focus on temporal separation of strokes that overlap spatially, and separately score the underlying, overwriting, and combined symbol.

Text is present in a very small number of drawings, and is probably not worth pursuing. Spokes are even rarer, but are somewhat analogous to tick marks except they appear in the hand region. Similar to tick marks, they often appear in multiples, are angularly equally spaced, and run from the center of the clock to the clockface. While a detector for spokes would probably not be difficult to build due to these unique properties, one would need to be careful to not make the spoke detector overly aggressive and misclassify hands as spokes, as hands outnumber spokes by over 40 to 1.

6.3.2 State-based Segmentation Approach

The previous ClockSketch classifier approach focused on spatiotemporal k-means clustering to segment the strokes. The new approach added spatial perturbation, joining, and splitting for every single stroke after the clustering takes place. While it is infeasible to build a large Markov model to capture every symbol and transition probability of a sketch, we do note that in over 90% of drawings, the first stroke is a clockface. After the clockface, most drawers start drawing the numerals starting from the 12 and finish drawing all the numerals before moving onto other components. While some drawers do just draw 12, 3, 6, and 9 at first and intersperse non-numeral symbols in the middle, it may still be promising to make a simplified Hidden Markov model, where some or all numerals as well as hands are grouped together as a single state.

This would mostly exist in parallel with the current clustering method as the two are very different approaches. While it would eventually breaks down for complex drawings with overwrites, it still may be useful for segmentation and labeling to weight a symbol's classification by the Markov model's state probabilities based on initial stroke properties and drawing order.

6.3.3 Parameter Analysis and Other Probabilistic Algorithms

Much of the new classifier's behavior is still reliant on empirically tuned parameters. In addition, these parameters were never optimized in conjunction with every other parameter for two reasons. The first is that the search space for changing every combination of parameters is prohibitively large. The second is that empirically tuning every parameter at once to optimize the performance on this data set would effectively overfit the current classifier to our data sets. Overtuning would bias the parameters toward our specific set of data sets, and the resulting performance would not generalize well to other clinical data.

While we have a much larger sample of clinical data than the previous ClockSketch classifier, work could still be down to quantify how sensitive our performance is to each of the parameters, and how sensitive the parameters are to the input training data. While performance and parameter variance analysis via nested cross validation would be prohibitively slow with the algorithms in their present state, the algorithms could be optimized, perhaps by picking a different perturbation approach such as simulated annealing. Simulated annealing would only make optimal changes after a certain number of iterations and would thus converge and terminate faster. It would also be interesting to characterize the resulting segmentation and classification performance for a variety of probabilistic algorithms.

Appendix A

Clock Drawing Data Sets

A.1 YDU-51 Healthy Training Set

YDU0015872230-2012-04-13-09-22-13ScoredV3.7.csk

YDU0024943741-2012-04-12-11-58-18ScoredV3.7.csk

YDU0029629324-2012-04-26-14-55-31ScoredV3.8.csk

YDU0030990860-2012-07-02-15-24-30ScoredV4.0.csk

YDU0053889897-2012-06-08-10-14-47ScoredV3.8.csk

YDU0057362444-2012-12-20-07-56-35ScoredV4.1.csk

YDU0064762880-2012-12-20-12-05-22ScoredV4.1.csk

YDU0065892741-2011-10-13-09-25-34ScoredV4.0.csk

YDU0089151951-2012-11-28-10-39-20ScoredV4.1.csk

YDU0092317699-2013-03-08-09-42-13ScoredV4.4.csk

YDU0093793509-2012-11-27-15-05-48ScoredV4.4.csk

YDU0112592176-2012-07-31-16-43-38ScoredV4.0.csk

YDU0123041202-2011-11-15-09-21-19ScoredV3.8.csk

YDU0129091497-2011-12-19-18-51-20ScoredV3.8.csk

YDU0144341018-2012-04-24-15-12-09ScoredV3.8.csk

YDU0160769184-2012-05-09-12-21-12ScoredV3.8.csk

YDU0169902394-2011-11-08-17-54-10ScoredV3.8.csk

YDU0181639157-2012-11-16-09-34-50ScoredV4.1.csk

YDU0187125954-2013-05-23-10-36-18ScoredV4.4.csk

YDU0189863547-2012-06-13-08-04-00ScoredV3.8.csk

YDU0191480696-2013-04-11-10-09-41ScoredV4.4.csk

YDU0196742056-2012-04-27-12-43-28ScoredV4.0.csk

YDU0211097573-2011-11-28-18-48-15ScoredV4.0.csk

YDU0213031281-2013-02-28-08-05-50ScoredV4.4.csk

YDU0236027528-2011-10-17-18-39-30ScoredV3.7.csk

YDU0240350443-2012-07-12-08-52-55ScoredV4.0.csk

YDU0247161988-2012-07-23-10-55-23ScoredV4.0.csk

YDU0252721858-2011-11-01-15-16-12ScoredV3.7.csk

YDU0257649605-2012-09-24-13-21-16ScoredV4.1.csk

YDU0257793879-2012-02-17-10-00-09ScoredV3.7.csk

YDU0262424844-2012-12-04-09-43-24ScoredV4.1.csk

YDU0283089966-2012-11-13-08-23-52ScoredV4.1.csk

YDU0283670083-2012-11-29-10-12-58ScoredV4.1.csk

YDU0283970673-2012-06-11-15-21-01ScoredV4.0.csk

YDU0286312782-2013-01-18-10-17-23ScoredV4.4.csk

YDU0295441515-2012-07-17-14-54-08ScoredV4.0.csk

YDU0305577176-2012-08-17-10-39-04ScoredV4.1.csk

YDU0308807254-2012-11-08-13-06-00ScoredV4.1.csk

YDU0310771230-2011-11-14-18-10-16ScoredV3.8.csk

YDU0312745603-2012-06-11-19-10-54ScoredV3.8.csk

YDU0320124394-2013-04-29-11-04-18ScoredV4.4.csk

YDU0323716285-2012-10-15-08-14-30ScoredV4.1.csk

YDU0331939947-2012-09-12-08-15-15ScoredV4.1.csk

YDU0342530934-2011-10-13-11-48-27ScoredV3.7.csk

YDU0354521219-2011-10-07-09-38-42ScoredV3.7.csk

YDU0371045035-2012-02-13-15-17-44ScoredV3.7.csk

YDU0372786121-2013-02-11-09-35-09ScoredV4.4.csk

YDU0375115358-2012-07-17-14-37-32ScoredV4.0.csk

YDU0395245091-2012-10-03-13-33-44ScoredV4.1.csk

YDU0408173575-2012-10-23-07-47-51ScoredV4.1.csk

YDU0410561867-2012-01-19-09-24-47ScoredV3.7.csk

A.2 YDU-100 Healthy Testing Set

YDU1458697332-2013-04-26-11-09-13ScoredV4.4.csk

YDU1460920966-2011-11-18-07-47-33ScoredV3.8.csk

YDU1469606748-2012-12-11-10-08-26ScoredV4.1.csk

YDU1473040393-2012-04-12-10-33-21ScoredV3.7.csk

YDU1473494971-2011-12-06-15-30-15ScoredV3.8.csk

YDU1476838616-2012-08-21-08-26-06ScoredV4.1.csk

YDU1481304857-2012-01-25-11-24-17ScoredV3.7.csk

YDU1482261301-2012-07-10-09-36-47ScoredV4.0.csk

YDU1483910122-2012-03-08-10-38-01ScoredV3.7.csk

YDU1487544303-2012-09-17-15-44-06ScoredV4.1.csk

YDU1493344152-2012-12-17-11-51-55ScoredV4.1.csk

YDU1494051224-2011-12-06-19-18-57ScoredV3.8.csk

YDU1494484873-2012-05-07-13-12-19ScoredV3.8.csk

YDU1503656639-2013-01-29-10-14-25ScoredV4.4.csk

YDU1515921483-2013-01-10-10-41-40ScoredV4.4.csk

YDU1522045959-2013-05-23-09-15-56ScoredV4.4.csk

YDU1522657394-2012-03-23-11-48-11ScoredV3.7.csk

YDU1523118076-2011-11-22-17-18-56ScoredV3.7.csk

YDU1524067713-2013-02-15-10-01-03ScoredV4.4.csk

YDU1536706745-2012-01-23-15-24-32ScoredV3.7.csk

YDU1546193696-2012-02-27-16-37-29ScoredV3.7.csk

YDU1551141438-2012-08-16-10-27-41ScoredV4.1.csk

YDU1552370799-2013-03-01-10-04-20ScoredV4.4.csk

YDU1558755772-2012-04-02-17-21-13ScoredV3.7.csk

YDU1589894957-2012-07-23-11-24-59ScoredV4.0.csk

YDU1596983648-2012-07-31-09-26-01ScoredV4.0.csk

YDU1608374391-2012-08-22-09-19-22ScoredV4.1.csk

YDU1622808830-2012-01-31-08-21-08ScoredV3.7.csk

YDU1625107784-2012-06-07-09-46-51ScoredV3.8.csk

YDU1637367575-2012-07-13-08-46-50ScoredV4.0.csk

YDU1640468945-2012-02-24-09-36-36ScoredV3.7.csk

YDU1663829053-2012-02-21-17-58-58ScoredV3.7.csk

YDU1664510618-2012-02-23-08-40-34ScoredV4.0.csk

YDU1698831024-2012-02-21-09-41-37ScoredV3.7.csk

YDU1700123170-2012-07-26-08-10-07ScoredV4.0.csk

YDU1726209825-2011-11-08-16-03-35ScoredV3.8.csk

YDU1728485575-2011-11-22-08-32-13ScoredV3.7.csk

YDU1731592614-2013-05-24-12-46-50ScoredV4.4.csk

YDU1737670214-2013-04-08-09-22-25ScoredV4.4.csk

YDU1737982982-2012-06-22-12-55-42ScoredV4.4.csk

YDU1758879357-2011-10-03-18-06-51ScoredV3.7.csk

YDU1758978507-2012-07-05-10-16-53ScoredV4.0.csk

YDU1766128431-2012-07-12-09-10-32ScoredV4.0.csk

YDU1766227533-2012-12-18-09-58-19ScoredV4.1.csk

YDU1771628421-2012-04-24-19-17-14ScoredV3.8.csk

YDU1780078212-2013-01-10-10-13-15ScoredV4.4.csk

YDU1793443167-2012-10-02-10-08-58ScoredV4.1.csk

YDU1796066114-2012-02-10-10-14-59ScoredV4.0.csk

YDU1798309932-2012-04-20-10-39-25ScoredV3.8.csk

YDU1803059071-2011-11-08-09-16-21ScoredV3.8.csk

YDU1824770682-2012-04-10-07-54-02ScoredV3.8.csk

YDU1826778250-2012-12-18-15-37-16ScoredV4.1.csk

YDU1830830601-2012-06-26-17-22-37ScoredV3.8.csk

YDU1839401766-2011-11-22-10-14-37ScoredV3.7.csk

YDU1839553265-2012-07-02-13-10-27ScoredV4.0.csk

YDU1839604771-2011-12-09-10-30-08ScoredV3.8.csk

YDU1857163369-2012-03-13-09-39-37ScoredV4.0.csk

YDU1870990911-2012-01-17-07-57-45ScoredV3.7.csk

YDU1878818512-2011-11-15-16-02-23ScoredV4.0.csk

YDU1886670830-2012-11-16-08-08-23ScoredV4.1.csk

YDU1891448893-2012-04-03-17-46-56ScoredV4.0.csk

YDU1903803704-2012-09-14-11-16-58ScoredV4.1.csk

YDU1905167133-2013-04-18-10-19-48ScoredV4.4.csk

YDU1915587439-2012-11-09-10-35-42ScoredV4.1.csk

YDU1921640868-2012-03-15-09-21-53ScoredV3.7.csk

YDU1922832223-2013-04-23-08-10-45ScoredV4.4.csk

YDU1950620409-2011-10-27-09-30-57ScoredV3.7.csk

YDU1953159052-2011-12-13-10-02-21ScoredV3.8.csk

YDU1956444688-2011-12-13-08-06-06ScoredV3.8.csk

YDU1965420037-2011-10-18-12-57-16ScoredV3.7.csk

YDU1966488010-2011-10-17-17-06-12ScoredV4.4.csk

YDU1966654365-2012-02-02-09-44-24ScoredV3.7.csk

YDU1980877338-2012-06-21-10-42-42ScoredV3.8.csk

YDU1980935602-2012-02-09-10-18-13ScoredV3.7.csk

YDU1980985828-2012-08-31-07-34-34ScoredV4.1.csk

YDU2005771610-2012-09-18-09-31-10ScoredV4.1.csk

YDU2010621421-2012-04-18-10-02-57ScoredV3.8.csk

YDU2012172477-2012-02-06-18-45-26ScoredV3.7.csk

YDU2015396656-2012-06-11-10-56-48ScoredV3.8.csk

YDU2018030890-2012-10-19-08-25-27ScoredV4.1.csk

YDU2032502788-2012-02-13-17-10-52ScoredV3.7.csk

YDU2045447907-2012-11-05-17-59-55ScoredV4.1.csk

YDU2060021469-2012-05-23-10-53-28ScoredV3.8.csk

YDU2061476618-2011-10-14-08-58-36ScoredV3.7.csk

YDU2061880442-2012-01-27-08-27-04ScoredV3.7.csk

YDU2071144394-2012-10-31-09-19-46ScoredV4.1.csk

YDU2075451915-2012-12-18-09-17-12ScoredV4.1.csk

YDU2076742150-2011-10-11-14-57-26ScoredV3.7.csk

YDU2084403185-2012-02-06-17-30-39ScoredV3.7.csk

YDU2090469603-2013-03-22-08-45-26ScoredV4.4.csk

YDU2092354574-2013-05-31-11-14-08ScoredV4.4.csk

YDU2092607378-2012-03-14-10-24-53ScoredV3.7.csk

YDU2100188187-2012-06-26-09-47-45ScoredV4.4.csk

YDU2100492527-2013-04-12-09-27-06ScoredV4.4.csk

YDU2134735600-2012-02-09-09-54-31ScoredV3.7.csk

YDU2139094310-2012-12-13-07-51-40ScoredV4.1.csk

YDU2140910798-2013-03-28-11-23-55ScoredV4.4.csk

YDU2141877161-2012-02-28-15-45-54ScoredV4.0.csk

YDU2144586219-2012-11-19-15-05-43ScoredV4.1.csk

A.3 VIN-96 Clinical Test Set

VIN0063271152-2011-05-13-10-31-30ScoredV3.6.csk

VIN0078085048-2009-05-05-13-30-03ScoredV3.6.csk

VIN0091872821-2007-07-30-16-10-59ScoredV3.0.csk

VIN0098771936-2011-11-22-12-04-09ScoredV3.6.csk

VIN0102603075-2011-08-03-10-08-28ScoredV3.6.csk

VIN0150986070-2011-11-10-15-03-08ScoredV3.6.csk

VIN0176770694-2006-02-20-13-19-09ScoredV3.0.csk

VIN0212724106-2006-04-05-11-46-30ScoredV3.0.csk

VIN0216801333-2011-02-25-16-32-24ScoredV3.6.csk

VIN0250487960-2009-04-22-10-05-59ScoredV3.0.csk

VIN0292653489-2009-08-21-12-37-33ScoredV3.0.csk

VIN0320646684-2006-03-13-13-40-53ScoredV3.0.csk

VIN0360962039-2009-07-08-14-40-12ScoredV3.0.csk

VIN0398883252-2007-03-19-09-45-19ScoredV3.0.csk

VIN0408418879-2011-03-30-11-48-21ScoredV3.6.csk

VIN0441352371-2007-04-25-09-55-48ScoredV3.0.csk

VIN0487443598-2009-12-09-13-13-59ScoredV3.0.csk

VIN0500200077-2011-03-04-16-28-30ScoredV3.6.csk

VIN0516652623-2011-11-15-08-44-46ScoredV3.6.csk

VIN0525310967-2009-02-18-09-47-22ScoredV3.6.csk

VIN0536748850-2009-08-28-11-18-54ScoredV3.0.csk

VIN0542469514-2011-02-18-16-37-10ScoredV3.6.csk

VIN0588654376-2009-07-22-13-06-22ScoredV3.0.csk

VIN0624107039-2011-07-25-13-17-30ScoredV3.6.csk

VIN0676446680-2006-03-22-09-56-04ScoredV3.0.csk

VIN0711833317-2011-11-14-10-52-39ScoredV3.6.csk

VIN0762360910-2007-04-11-14-40-56ScoredV3.0.csk

VIN0775888172-2006-07-19-10-06-33ScoredV3.0.csk

VIN0828812527-2007-06-25-10-52-45ScoredV3.0.csk

VIN0829757788-2009-10-29-10-28-11ScoredV3.0.csk

VIN0833140908-2011-07-25-09-05-47ScoredV3.6.csk

VIN0842680555-2007-03-27-09-11-57ScoredV3.0.csk

VIN0847029653-2007-01-23-14-22-06ScoredV3.0.csk

VIN0847870313-2007-06-14-15-49-26ScoredV3.0.csk

VIN0873509087-2007-04-12-10-37-10ScoredV3.6.csk

VIN0880699204-2011-12-21-11-26-50ScoredV3.6.csk

VIN0902133667-2006-12-14-10-15-57ScoredV3.0.csk

VIN0903667079-2009-10-21-09-16-28ScoredV3.0.csk

VIN0929937545-2011-11-14-13-24-56ScoredV3.6.csk

VIN0934868080-2011-07-19-09-54-30ScoredV3.6.csk

VIN0969671570-2011-07-27-15-32-26ScoredV3.6.csk

VIN0983535487-2009-12-02-11-21-16ScoredV3.0.csk

VIN1019479827-2007-03-12-09-32-31ScoredV3.0.csk

VIN1021133203-2006-06-21-09-56-44ScoredV3.0.csk

VIN1030668968-2009-02-23-09-39-31ScoredV3.6.csk

VIN1052780286-2007-05-22-13-18-35ScoredV3.0.csk

VIN1082466239-2006-04-18-10-02-09ScoredV3.0.csk

VIN1106145152-2006-05-10-13-15-06ScoredV3.0.csk

VIN1152904937-2009-03-16-10-18-38ScoredV3.0.csk

VIN1158639961-2009-02-09-10-22-03ScoredV3.0.csk

VIN1181803228-2011-10-11-15-41-25ScoredV3.6.csk

VIN1185197344-2006-03-20-14-01-00ScoredV3.6.csk

VIN1202464077-2011-12-16-12-51-14ScoredV3.6.csk

VIN1245859164-2011-11-26-13-10-18ScoredV3.6.csk

VIN1250546032-2009-12-15-12-09-06ScoredV3.0.csk

VIN1263579883-2007-01-04-09-26-53ScoredV3.0.csk

VIN1270987934-2010-07-26-13-29-18ScoredV3.0.csk

VIN1290845432-2011-04-01-13-50-59ScoredV3.6.csk

VIN1303699403-2006-11-15-09-12-48ScoredV3.0.csk

VIN1313750347-2006-05-04-16-19-12ScoredV3.0.csk

VIN1335376767-2011-11-16-15-15-23ScoredV3.6.csk

VIN1341672470-2007-01-11-13-42-38ScoredV3.0.csk

VIN1355430577-2006-03-06-14-25-58ScoredV3.0.csk

VIN1381455645-2011-06-10-08-41-19ScoredV3.6.csk

VIN1475529716-2009-07-07-09-55-26ScoredV3.0.csk

VIN1487156944-2007-07-17-13-35-36ScoredV3.0.csk

VIN1520402976-2007-03-05-13-32-22ScoredV3.0.csk

VIN1542794379-2007-04-10-13-31-20ScoredV3.0.csk

VIN1582668235-2007-01-02-09-33-30ScoredV3.0.csk

VIN1638973340-2006-05-24-10-45-45ScoredV3.0.csk

VIN1660017801-2009-03-17-15-35-05ScoredV3.0.csk

VIN1687212776-2006-04-20-11-13-56ScoredV3.0.csk

VIN1699673802-2009-11-06-14-37-36ScoredV3.0.csk

VIN1715687822-2007-07-19-11-09-04ScoredV3.0.csk

VIN1724539676-2006-02-21-14-45-09ScoredV3.0.csk

VIN1724927675-2009-09-29-14-56-48ScoredV3.0.csk

VIN1728476201-2006-05-30-10-23-30ScoredV3.0.csk

VIN1737987862-2011-11-16-12-59-56ScoredV3.6.csk

VIN1745613379-2007-05-31-15-43-49ScoredV3.0.csk

VIN1786441936-2009-12-08-12-12-26ScoredV3.0.csk

VIN1792676811-2007-01-18-11-31-56ScoredV3.0.csk

VIN1842240487-2006-03-27-16-17-04ScoredV3.0.csk

VIN1851737753-2007-06-27-09-39-27ScoredV3.6.csk

VIN1926831987-2011-10-24-08-52-21ScoredV3.6.csk

VIN1931807909-2009-11-06-11-29-55ScoredV3.0.csk

VIN1973134116-2009-04-29-16-12-05ScoredV3.0.csk

VIN1989294807-2009-12-04-15-30-48ScoredV3.0.csk

VIN1996974147-2011-03-17-18-43-37ScoredV3.6.csk

VIN1998384886-2006-03-06-15-43-42ScoredV3.0.csk

VIN2015542200-2011-02-25-12-03-29ScoredV3.6.csk

VIN2016358114-2009-09-16-14-38-42ScoredV3.6.csk

VIN2099133980-2006-11-16-13-46-07ScoredV3.0.csk

VIN2100335296-2009-04-21-10-11-50ScoredV3.0.csk

VIN2107627400-2006-03-22-16-34-03ScoredV3.0.csk

VIN2109969259-2009-10-19-13-35-21ScoredV3.0.csk

VIN2123733155-2011-04-07-11-30-15ScoredV3.6.csk

A.4 EGE/ORU-112 Clinical Test Set

EGE0048723202-2010-06-21-12-57-03ScoredV3.0.csk

EGE0056676390-2010-07-14-09-25-41ScoredV3.0.csk

EGE0091443850-2010-05-21-09-22-22ScoredV3.0.csk

EGE0092099221-2010-05-26-12-56-40ScoredV3.0.csk

EGE0143851472-2010-07-20-10-14-13ScoredV3.0.csk

EGE0167599527-2010-07-12-12-57-41ScoredV3.0.csk

EGE0270005059-2010-06-18-13-15-38ScoredV3.5.csk

EGE0271665761-2010-07-06-14-23-16ScoredV3.0.csk

EGE0348579752-2010-05-13-08-59-58ScoredV3.0.csk

EGE0375558151-2010-07-13-09-29-48ScoredV3.0.csk

EGE0430106867-2010-01-22-13-22-42ScoredV3.5.csk

EGE0458758770-2010-06-01-09-30-50ScoredV3.0.csk

EGE0468185867-2009-05-13-10-43-23ScoredV3.0.csk

EGE0477008340-2010-06-10-09-35-01ScoredV3.0.csk

EGE0477340547-2009-04-23-11-04-36ScoredV3.0.csk

EGE0488148287-2010-08-09-09-19-33ScoredV3.0.csk

EGE0489394608-2010-07-14-14-09-05ScoredV3.9.csk

EGE0495314770-2010-03-10-13-26-24ScoredV3.0.csk

EGE0528412617-2010-04-13-14-22-10ScoredV3.9.csk

EGE0556887562-2010-08-31-09-35-00ScoredV3.0.csk

EGE0603076589-2010-05-27-13-13-45ScoredV3.0.csk

EGE0628057102-2010-03-31-13-33-01ScoredV3.0.csk

EGE0706902923-2010-07-09-13-09-46ScoredV3.0.csk

EGE0719336199-2010-04-05-09-22-44ScoredV3.0.csk

EGE0720764651-2010-05-13-14-25-42ScoredV3.0.csk

EGE0744818560-2010-06-30-14-07-11ScoredV3.0.csk

EGE0766567854-2010-07-30-13-44-54ScoredV3.0.csk

EGE0813752608-2010-04-21-14-17-02ScoredV3.0.csk

EGE0838358279-2010-06-11-09-28-04ScoredV3.9.csk

EGE0882081772-2010-05-27-09-14-35ScoredV3.0.csk

EGE0887135619-2010-03-31-09-31-58ScoredV3.0.csk

EGE0922540745-2010-07-15-13-20-16ScoredV3.0.csk

EGE0960117541-2010-04-19-10-16-16ScoredV3.9.csk

EGE0969601126-2010-06-16-09-20-44ScoredV3.9.csk

EGE1059111135-2010-06-23-09-59-34ScoredV3.0.csk

EGE1077000824-2010-07-20-09-41-05ScoredV3.0.csk

EGE1152023406-2010-04-14-09-24-23ScoredV3.0.csk

EGE1165700469-2010-08-02-13-24-32ScoredV3.5.csk

EGE1168812095-2010-03-30-14-08-48ScoredV3.0.csk

EGE1174719586-2010-07-12-10-09-13ScoredV3.0.csk

EGE1194738423-2010-06-09-13-35-23ScoredV3.0.csk

EGE1199219204-2010-04-28-13-21-55ScoredV3.5.csk

EGE1210175927-2010-04-12-09-47-36ScoredV3.0.csk

EGE1227383573-2009-04-22-04-57-14ScoredV3.0.csk

EGE1229216930-2010-04-26-09-17-58ScoredV3.0.csk

EGE1230811213-2010-07-01-13-08-52ScoredV3.5.csk

EGE1238003390-2010-06-02-12-44-08ScoredV3.0.csk

EGE1240346762-2009-04-20-05-25-08ScoredV3.0.csk

EGE1252519043-2010-06-23-13-15-35ScoredV3.0.csk

EGE1318066790-2009-04-15-05-45-20ScoredV3.0.csk

EGE1331665504-2010-07-08-10-15-42ScoredV3.0.csk

EGE1370827434-2009-04-16-05-31-15ScoredV3.0.csk

EGE1377705409-2010-04-26-13-27-01ScoredV3.0.csk

EGE1418164226-2010-07-28-09-27-05ScoredV3.0.csk

EGE1466972375-2010-04-16-11-25-25ScoredV3.0.csk

EGE1477402218-2010-06-17-10-05-57ScoredV3.0.csk

EGE1516019668-2010-06-14-13-17-14ScoredV3.0.csk

EGE1528314806-2010-07-08-13-19-36ScoredV3.0.csk

EGE1551806340-2010-04-23-13-25-45ScoredV3.0.csk

EGE1555519969-2010-06-16-13-00-56ScoredV3.0.csk

EGE1564463898-2010-06-25-10-23-04ScoredV3.0.csk

EGE1579069590-2010-08-24-14-29-28ScoredV3.0.csk

EGE1595650202-2010-05-19-13-19-21ScoredV3.0.csk

EGE1635278100-2010-07-22-09-39-32ScoredV3.5.csk

EGE1663533090-2010-03-01-09-54-36ScoredV3.0.csk

EGE1703572234-2010-04-20-14-31-03ScoredV3.0.csk

EGE1724348272-2010-08-27-13-27-54ScoredV3.9.csk

EGE1729607156-2010-02-02-14-06-26ScoredV3.0.csk

EGE1731929614-2010-08-10-09-21-32ScoredV3.9.csk

EGE1765900491-2010-06-07-13-13-50ScoredV3.0.csk

EGE1813478853-2010-04-30-10-52-35ScoredV3.5.csk

EGE1865516148-2010-07-13-14-42-27ScoredV3.0.csk

EGE1923413357-2010-07-27-14-19-36ScoredV3.0.csk

EGE1940454417-2010-06-28-13-21-41ScoredV3.0.csk

EGE1962770219-2010-03-17-09-19-15ScoredV3.0.csk

EGE1969144806-2009-04-29-09-27-31ScoredV3.0.csk

EGE2038384853-2010-05-12-12-40-20ScoredV3.0.csk

EGE2089890253-2010-07-16-15-11-01ScoredV3.0.csk

EGE2100566174-2010-06-28-14-19-49ScoredV3.0.csk

ORU0033873510-2010-06-04-09-56-39ScoredV3.0.csk

ORU0141409120-2011-06-27-09-45-21ScoredV3.7.csk

ORU0214002724-2009-11-06-11-01-07ScoredV3.0.csk

ORU0234651310-2010-07-29-10-50-19ScoredV3.0.csk

ORU0260103282-2010-03-01-11-18-51ScoredV3.0.csk

ORU0263782536-2010-08-25-11-10-07ScoredV3.7.csk

ORU0271454768-2009-12-17-14-27-21ScoredV3.0.csk

ORU0278498759-2009-09-15-10-52-37ScoredV3.0.csk

ORU0355087110-2009-06-23-15-22-47ScoredV3.0.csk

ORU0377348240-2011-09-09-11-05-39ScoredV3.7.csk

ORU0435765892-2009-05-11-11-45-40ScoredV3.7.csk

ORU0526107987-2010-02-15-16-27-57ScoredV3.0.csk

ORU0566807646-2011-06-14-15-30-35ScoredV3.0.csk

ORU0577384976-2010-05-20-14-27-09ScoredV3.5.csk

ORU0702838292-2011-06-22-14-58-57ScoredV3.0.csk

ORU1184517177-2009-06-29-10-43-39ScoredV3.6.csk

ORU1298744135-2009-05-27-11-17-58ScoredV3.7.csk

ORU1504064473-2010-07-08-14-59-54ScoredV3.0.csk

ORU1513723722-2010-03-02-10-56-57ScoredV3.0.csk

ORU1523209376-2010-09-23-10-39-28ScoredV3.0.csk

ORU1545747048-2009-05-19-16-26-10ScoredV3.7.csk

ORU1572369380-2009-05-13-11-09-35ScoredV3.7.csk

ORU1579143314-2010-05-27-12-06-45ScoredV3.0.csk

ORU1594504278-2010-04-20-15-29-54ScoredV3.0.csk

ORU1614808862-2009-07-15-09-44-15ScoredV3.0.csk

ORU1823106156-2009-11-17-11-01-55ScoredV3.0.csk

ORU1852069863-2009-09-29-14-27-11ScoredV3.0.csk

ORU1952963752-2011-10-04-10-49-41ScoredV3.0.csk

ORU1967217631-2009-04-07-11-48-39ScoredV3.7.csk

ORU1979847017-2010-11-02-13-57-42ScoredV3.0.csk

ORU2120290419-2009-10-13-15-02-54ScoredV3.0.csk

ORU2123368665-2009-10-26-10-52-34ScoredV3.0.csk

ORU2124971720-2009-04-24-11-07-33ScoredV3.6.csk

A.5 EMD-20 Clinical Test Set

TCD1V3.5CIN2109964091-2010-10-15-15-31-17ScoredV3.0.csk

TCD2V3.5CIN0939054722-2011-01-07-12-17-59ScoredV3.0.csk

TCD3V3.5CIN2076705627-2011-03-16-14-48-47ScoredV3.0.csk

TCD4V3.5CIN0781482116-2011-05-10-15-10-49ScoredV3.0.csk

TCD5V3.5CIN1370460189-2011-05-11-15-02-05ScoredV3.0.csk

TCE1V3.5CIN1080838032-2010-05-07-12-10-52ScoredV3.0.csk

TCE2V3.5CIN0665001820-2010-06-03-13-55-53ScoredV3.0.csk

TCE3V3.5CIN1792930152-2010-06-08-08-59-22ScoredV3.0.csk

TCE4V3.5CIN1841775235-2010-06-23-11-32-41ScoredV3.0.csk

TCE5V3.5CIN0752776689-2010-12-09-15-56-30ScoredV3.5.csk

TCE6V3.5CIN1970844532-2011-01-25-14-31-04ScoredV3.0.csk

TCE7V3.5CIN0988448711-2011-04-05-15-25-31ScoredV3.0.csk

TCE8V3.5CIN0668952640-2011-05-13-10-48-08ScoredV3.0.csk

TCM2V3.5CIN0788977646-2010-07-08-13-35-31ScoredV3.0.csk

TCM3V3.5CIN1788079352-2011-03-11-13-07-53ScoredV3.0.csk

TCM4V3.5CIN0196784358-2011-04-08-16-19-15ScoredV3.0.csk

TCM5V3.5CIN1878643352-2011-05-11-10-12-04ScoredV3.0.csk

TCM6V3.5CIN1029212004-2011-05-11-12-10-16ScoredV3.0.csk

TCM7V3.5CIN2022489965-2011-05-13-14-41-23ScoredV3.0.csk

A.6 CIN-170 Clinical Test Set

CIN0000034213-2013-10-09-14-13-35ScoredV5.0.csk

CIN0003155252-2014-05-01-15-15-47ScoredV4.6.csk

CIN0009532850-2013-02-01-11-04-41ScoredV4.6.csk

CIN0010839927-2010-01-12-15-04-08ScoredV4.6.csk

CIN0028584735-2013-05-02-11-53-49ScoredV4.6.csk

CIN0039693431-2010-02-11-15-42-48ScoredV4.6.csk

CIN0055911696-2013-01-04-11-22-05ScoredV4.6.csk

CIN0059530813-2009-01-27-09-51-44ScoredV4.6.csk

CIN0079657991-2011-04-29-14-48-00ScoredV3.0.csk

CIN0082624001-2014-01-17-09-07-56ScoredV4.6.csk

CIN0083950870-2011-04-28-15-08-00ScoredV3.0.csk

CIN0093206168-2011-07-19-14-40-25ScoredV4.6.csk

CIN0105093392-2012-05-22-10-45-46ScoredV4.6.csk

CIN0105520977-2007-06-14-10-32-37ScoredV4.4.csk

CIN0105982484-2013-05-07-15-42-15ScoredV4.6.csk

CIN0123862361-2013-11-15-14-22-36ScoredV4.6.csk

CIN0124675803-2013-06-19-09-53-45ScoredV4.6.csk

CIN0128420664-2012-02-24-10-38-29ScoredV4.6.csk

CIN0132590717-2008-04-03-11-05-51ScoredV4.4.csk

CIN0145662587-2012-11-13-12-17-29ScoredV3.8.csk

CIN0164742787-2009-12-10-15-35-03ScoredV4.6.csk

CIN0167812185-2010-03-16-10-24-39ScoredV3.0.csk

CIN0168920098-2013-05-24-15-56-12ScoredV4.6.csk

CIN0175203147-2011-11-02-14-28-50ScoredV4.6.csk

CIN0176652109-2011-06-03-11-37-44ScoredV3.6.csk

CIN0213849691-2012-04-24-09-10-19ScoredV3.6.csk

CIN0217094361-2010-03-05-13-52-53ScoredV4.6.csk

CIN0228374989-2013-11-19-14-39-10ScoredV4.6.csk

CIN0235899179-2012-01-11-14-55-22ScoredV3.6.csk

CIN0245955862-2014-01-29-09-12-13ScoredV4.6.csk

CIN0246559773-2013-10-22-11-49-49ScoredV4.6.csk

CIN0251433661-2012-05-08-10-55-50ScoredV4.6.csk

CIN0253153890-2013-06-10-14-47-34ScoredV4.6.csk

CIN0253370351-2014-03-11-14-46-06ScoredV4.6.csk

CIN0257085842-2010-12-08-10-53-59ScoredV3.5.csk

CIN0263503813-2005-07-22-12-51-42ScoredV3.0.csk

CIN0268059704-2010-02-19-10-22-14ScoredV3.5.csk

CIN0268598879-2012-05-25-11-11-36ScoredV4.4.csk

CIN0274302801-2014-02-28-14-40-52ScoredV4.6.csk

CIN0276246219-2010-04-16-10-57-20ScoredV4.6.csk

CIN0281970435-2012-01-13-12-50-22ScoredV4.6.csk

CIN0300655383-2005-08-05-13-44-55ScoredV3.0.csk

CIN0302941166-2013-07-18-14-09-07ScoredV4.6.csk

CIN0319786108-2013-09-06-11-04-17ScoredV4.4.csk

CIN0324015245-2012-01-27-11-35-30ScoredV4.6.csk

CIN0327850012-2006-05-05-12-46-51ScoredV3.0.csk

CIN0335734233-2011-10-11-10-42-48ScoredV3.6.csk

CIN0338325221-2008-09-24-09-00-32ScoredV4.4.csk

CIN0341880680-2012-05-22-09-27-55ScoredV4.6.csk

CIN0346719245-2014-07-22-14-26-49ScoredV4.7.csk

CIN0353394385-2012-07-31-15-33-49ScoredV4.6.csk

CIN0370890411-2013-03-26-10-49-27ScoredV4.4.csk

CIN0379627862-2013-09-11-11-22-14ScoredV4.4.csk

CIN0381839798-2005-07-22-10-02-39ScoredV3.0.csk

CIN0388064539-2011-05-06-11-13-37ScoredV4.6.csk

CIN0389960265-2014-02-17-14-21-48ScoredV4.6.csk

CIN0403856442-2013-10-16-13-28-19ScoredV4.6.csk

CIN0413002535-2010-01-13-14-47-20ScoredV3.0.csk

CIN0416841585-2009-10-01-16-20-40ScoredV3.0.csk

CIN0431929809-2010-01-07-10-44-47ScoredV4.6.csk

CIN0437392888-2012-01-11-11-11-20ScoredV3.6.csk

CIN0451985477-2011-09-20-10-11-44ScoredV3.6.csk

CIN0455388352-2012-03-14-14-48-28ScoredV4.6.csk

CIN0462467944-2012-11-27-11-03-21ScoredV4.4.csk

CIN0479823172-2012-05-15-14-16-18ScoredV4.6.csk

CIN0482526350-2013-11-26-13-38-34ScoredV4.6.csk

CIN0492479409-2011-06-28-11-26-55ScoredV3.5.csk

CIN0508401083-2011-03-01-14-49-55ScoredV4.6.csk

CIN0509887985-2012-06-08-09-28-19ScoredV3.8.csk

CIN0517439397-2012-02-08-14-46-04ScoredV3.7.csk

CIN0550945481-2012-04-27-10-43-35ScoredV4.6.csk

CIN0560040151-2007-11-05-08-27-20ScoredV4.4.csk

CIN0584487598-2010-09-17-15-49-01ScoredV4.6.csk

CIN0601270254-2014-03-13-10-50-49ScoredV4.6.csk

CIN0604497858-2013-02-19-12-11-03ScoredV4.5.csk

CIN0605679142-2008-01-22-15-15-55ScoredV4.6.csk

CIN0639609890-2014-01-07-12-23-10ScoredV4.6.csk

CIN0642296171-2014-02-19-12-12-05ScoredV4.7.csk

CIN0649042983-2012-12-12-12-14-01ScoredV4.4.csk

CIN0653161572-2010-12-16-14-47-41ScoredV3.0.csk

CIN0670219764-2014-04-22-08-20-39ScoredV4.7.csk

CIN0694525168-2012-05-23-10-58-54ScoredV4.6.csk

CIN0699897986-2008-05-02-09-19-47ScoredV4.6.csk

CIN0700659550-2010-01-28-14-45-26ScoredV4.6.csk

CIN0723152914-2012-12-18-14-53-31ScoredV4.6.csk

CIN0742533336-2010-02-17-11-25-06ScoredV3.0.csk

CIN0743332986-2012-10-30-13-34-14ScoredV4.6.csk

CIN0760212440-2012-12-11-13-47-35ScoredV4.6.csk

CIN0771982657-2008-04-03-14-01-42ScoredV3.7.csk

CIN0781482116-2011-05-10-15-10-49ScoredV4.6.csk

CIN0791852929-2014-01-29-15-30-15ScoredV4.6.csk

CIN0797570370-2009-09-01-10-52-29ScoredV4.6.csk

CIN0816762614-2009-06-23-14-30-04ScoredV4.6.csk

CIN0816782378-2011-11-03-11-50-34ScoredV3.6.csk

CIN0818723200-2010-01-13-11-54-28ScoredV4.6.csk

CIN0822866441-2009-11-24-14-52-30ScoredV4.6.csk

CIN0824158583-2014-01-15-09-20-32ScoredV4.6.csk

CIN0829359923-2010-02-09-15-23-40ScoredV4.6.csk

CIN0854992145-2011-03-31-14-42-06ScoredV4.6.csk

CIN0931051869-2013-08-29-15-15-03ScoredV4.6.csk

CIN0939054722-2011-01-07-12-17-59ScoredV4.6.csk

CIN0942116121-2011-02-15-13-30-06ScoredV4.6.csk

CIN0965876830-2011-10-20-13-28-18ScoredV3.6.csk

CIN0985338557-2011-06-23-15-23-18ScoredV4.6.csk

CIN0993769749-2007-10-23-10-32-51ScoredV4.4.csk

CIN1015587754-2010-02-25-10-52-57ScoredV4.6.csk

CIN1040066317-2011-01-14-11-14-59ScoredV4.6.csk

CIN1047511993-2007-04-03-11-06-19ScoredV3.7.csk

CIN1061551666-2011-07-06-11-56-17ScoredV4.4.csk

CIN1082366723-2011-09-07-10-09-19ScoredV3.6.csk

CIN1088914040-2009-05-22-11-32-22ScoredV4.6.csk

CIN1089443927-2013-01-24-15-24-25ScoredV4.6.csk

CIN1105631229-2013-10-01-12-14-05ScoredV4.4.csk

CIN1111077033-2013-10-15-16-14-00ScoredV4.6.csk

CIN1112350489-2013-09-05-15-49-54ScoredV4.6.csk

CIN1118890191-2012-11-13-15-46-25ScoredV4.6.csk

CIN1136301330-2005-07-21-11-15-28ScoredV3.0.csk

CIN1206532380-2010-08-31-10-17-51ScoredV4.6.csk

CIN1220067567-2011-09-28-14-41-56ScoredV4.6.csk

CIN1229225404-2011-04-12-12-03-21ScoredV3.0.csk

CIN1239163249-2012-05-04-15-50-11ScoredV4.6.csk

CIN1300692833-2013-02-27-09-58-34ScoredV4.4.csk

CIN1361356030-2012-06-29-11-14-49ScoredV4.6.csk

CIN1371523966-2012-02-22-09-37-57ScoredV3.6.csk

CIN1393917190-2012-02-29-14-34-45ScoredV4.6.csk

CIN1436950986-2013-03-19-08-59-27ScoredV4.6.csk

CIN1440212394-2011-05-25-14-54-16ScoredV3.6.csk

CIN1469526899-2012-11-02-15-47-20ScoredV4.6.csk

CIN1507887534-2008-11-14-12-06-06ScoredV4.6.csk

CIN1525915365-2011-04-15-09-31-49ScoredV4.6.csk

CIN1530958939-2010-12-21-09-54-00ScoredV3.0.csk

CIN1573515168-2013-09-24-10-08-53ScoredV4.4.csk

CIN1579029831-2010-02-03-10-03-23ScoredV4.5.csk

CIN1582152442-2012-10-17-09-53-55ScoredV4.3.csk

CIN1671055220-2012-06-01-10-02-39ScoredV4.6.csk

CIN1687732000-2013-08-06-09-40-21ScoredV4.6.csk

CIN1735040052-2010-05-25-10-19-47ScoredV4.6.csk

CIN1737373792-2011-12-08-16-18-31ScoredV4.6.csk

CIN1739530392-2010-04-07-12-17-24ScoredV3.0.csk

CIN1760264257-2010-06-09-14-26-29ScoredV4.6.csk

CIN1792641094-2009-05-29-15-34-23ScoredV4.6.csk

CIN1793327664-2012-12-04-10-24-37ScoredV4.6.csk

CIN1800007620-2010-03-12-10-29-43ScoredV4.6.csk

CIN1830288560-2010-10-05-14-35-05ScoredV3.0.csk

CIN1841368545-2013-05-13-15-43-48ScoredV4.6.csk

CIN1853578308-2010-04-27-12-12-57ScoredV4.6.csk

CIN1868194358-2010-02-24-14-19-24ScoredV3.0.csk

CIN1887040604-2012-06-06-10-01-11ScoredV4.6.csk

CIN1890265273-2010-01-13-16-01-27ScoredV3.6.csk

CIN1897577607-2013-08-27-14-29-18ScoredV4.4.csk

CIN1897605712-2011-11-29-11-22-57ScoredV4.6.csk

CIN1905170544-2009-12-09-11-17-29ScoredV4.6.csk

CIN1909600964-2010-10-26-10-49-10ScoredV4.6.csk

CIN1914662806-2010-11-16-10-56-20ScoredV4.6.csk

CIN1936333279-2013-08-27-11-35-55ScoredV4.4.csk

CIN1942433978-2012-01-17-09-50-10ScoredV4.6.csk

CIN1943645286-2010-11-09-16-01-25ScoredV4.6.csk

CIN1986393573-2009-09-09-09-29-05ScoredV3.7.csk

CIN2006928793-2011-02-01-09-25-29ScoredV4.6.csk

CIN2012822029-2009-09-10-15-03-26ScoredV4.6.csk

CIN2050302313-2012-06-14-11-30-39ScoredV4.6.csk

CIN2060293230-2010-01-26-09-51-39ScoredV4.6.csk

CIN2065300116-2009-11-05-15-31-23ScoredV4.6.csk

CIN2065475662-2010-07-22-15-46-04ScoredV4.6.csk

CIN2072315286-2012-08-14-13-36-41ScoredV4.6.csk

CIN2087014604-2008-04-03-10-50-34ScoredV3.0.csk

CIN2109964091-2010-10-15-15-31-17ScoredV4.6.csk

CIN2126605636-2011-09-30-10-03-51ScoredV4.6.csk

CIN2146016349-2011-06-24-11-13-27ScoredV3.5.csk

CIN2147142001-2011-10-12-12-18-07ScoredV3.6.csk

A.7 TICK-74 Clinical Test Set

CIN0087127573-2013-07-03-13-03-33ScoredV4.6.csk

CIN0100603242-2013-09-11-09-43-06ScoredV4.4.csk

CIN0103347419-2012-05-29-10-24-36ScoredV4.6.csk

CIN0136113639-2013-02-19-09-15-07ScoredV4.6.csk

CIN0152572861-2011-09-20-13-43-27ScoredV4.6.csk

CIN0184968821-2007-06-07-14-05-53ScoredV4.6.csk

CIN0248655962-2012-10-17-12-03-48ScoredV4.5.csk

CIN0249837095-2013-05-10-10-23-12ScoredV4.6.csk

CIN0252802715-2012-05-30-09-45-10ScoredV4.6.csk

CIN0263633495-2010-02-04-14-39-45ScoredV3.0.csk

CIN0273928773-2009-06-11-15-11-41ScoredV4.6.csk

CIN0278986977-2014-03-05-11-43-09ScoredV4.6.csk

CIN0281783343-2014-03-18-10-01-21ScoredV4.6.csk

CIN0371975662-2014-05-01-11-10-45ScoredV4.6.csk

CIN0381315269-2010-11-19-14-41-33ScoredV4.6.csk

CIN0381774299-2013-06-04-11-30-06ScoredV4.3.csk

CIN0400188137-2012-06-20-10-39-38ScoredV4.6.csk

CIN0414452951-2008-11-07-09-05-32ScoredV4.6.csk

CIN0428397456-2014-04-09-12-03-58ScoredV4.6.csk

CIN0481557520-2012-07-18-12-19-27ScoredV3.8.csk

CIN0482970969-2012-08-01-11-50-44ScoredV4.4.csk

CIN0510508378-2013-09-24-15-06-54ScoredV4.6.csk

CIN0530974982-2010-01-15-16-30-45ScoredV3.5.csk

CIN0568585115-2008-07-10-08-55-08ScoredV3.0.csk

CIN0613043161-2011-10-21-14-09-58ScoredV4.6.csk

CIN0626337772-2009-08-25-13-39-45ScoredV4.6.csk

CIN0637371952-2010-10-21-14-01-27ScoredV4.6.csk

CIN0655752254-2012-07-10-10-15-27ScoredV4.6.csk

CIN0678593951-2012-03-09-14-56-58ScoredV4.6.csk

CIN0684048327-2012-10-04-13-22-18ScoredV4.6.csk

CIN0802634318-2014-03-25-10-56-28ScoredV4.6.csk

CIN0873531566-2014-05-20-09-50-41ScoredV4.7.csk

CIN0881543816-2013-11-22-10-31-09ScoredV4.6.csk

CIN0928346331-2007-04-04-15-27-39ScoredV3.7.csk

CIN0980486179-2014-03-06-12-43-14ScoredV4.6.csk

CIN1134532484-2013-05-10-14-26-56ScoredV4.4.csk

CIN1136641253-2010-11-12-11-06-57ScoredV4.6.csk

CIN1162209701-2008-02-28-08-18-23ScoredV4.4.csk

CIN1176605293-2012-06-08-15-02-31ScoredV4.6.csk

CIN1220912665-2009-10-09-14-32-37ScoredV4.6.csk

CIN1229671350-2014-03-31-10-20-47ScoredV4.6.csk

CIN1237389173-2013-07-25-12-29-48ScoredV4.6.csk

CIN1243239826-2013-05-01-12-03-29ScoredV4.5.csk

CIN1268835105-2014-07-30-09-49-26ScoredV4.7.csk

CIN1276579891-2014-04-09-12-25-59ScoredV4.6.csk

CIN1374758076-2011-05-27-10-54-10ScoredV4.6.csk

CIN1436063098-2007-10-11-15-19-15ScoredV4.6.csk

CIN1437981433-2013-10-23-14-41-16ScoredV4.6.csk

CIN1441223030-2013-06-28-09-55-08ScoredV4.6.csk

CIN1469615084-2012-05-29-10-16-25ScoredV4.6.csk

CIN1470707530-2010-10-29-11-11-48ScoredV4.6.csk

CIN1480529293-2013-05-31-10-12-53ScoredV4.6.csk

CIN1503842164-2012-11-20-12-05-41ScoredV4.4.csk

CIN1518111233-2013-03-13-15-04-26ScoredV4.4.csk

CIN1523269884-2008-10-30-14-27-49ScoredV4.6.csk

CIN1542627735-2013-10-16-10-31-56ScoredV4.6.csk

CIN1550123002-2014-02-27-16-25-05ScoredV4.6.csk

CIN1578339852-2013-01-16-14-42-56ScoredV4.6.csk

CIN1649736986-2013-09-11-13-04-24ScoredV4.6.csk

CIN1744559777-2012-05-23-12-22-54ScoredV3.6.csk

CIN1788188325-2013-02-07-11-13-00ScoredV4.6.csk

CIN1791895795-2013-02-06-14-16-50ScoredV4.6.csk

CIN1807983379-2009-11-24-10-53-31ScoredV4.6.csk

CIN1814016875-2013-09-20-11-57-44ScoredV4.6.csk

CIN1817602571-2013-09-18-12-59-11ScoredV4.6.csk

CIN1824350152-2014-03-24-15-05-09ScoredV4.6.csk

CIN1856010988-2014-02-26-10-15-11ScoredV4.6.csk

CIN1861646671-2013-10-09-15-06-37ScoredV4.6.csk

CIN1971691545-2011-05-25-15-49-22ScoredV4.6.csk

CIN1980597825-2010-07-13-15-06-38ScoredV4.6.csk

CIN1993089748-2011-10-18-16-41-14ScoredV4.6.csk

CIN1999330564-2005-12-16-14-37-42ScoredV3.5.csk

CIN2041745182-2013-01-08-12-16-21ScoredV4.5.csk

CIN2146196895-2010-03-04-13-35-04ScoredV4.6.csk

Bibliography

[Bhat and Hammond, 2009] Bhat, A. and Hammond, T. (2009). Using entropy to distinguish shape versus text in hand-drawn diagrams. In *Proceedings of the 21st International Jont Conference on Artifical Intelligence*, IJCAI'09, pages 1395–1400, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.

[Ciresan et al., 2012] Ciresan, D. C., Meier, U., and Schmidhuber, J. (2012). Multi-column deep neural networks for image classification. *CoRR*, abs/1202.2745.

[Dickmann et al., 2010] Dickmann, L., Lensing, T., Porzel, R., Malaka, R., and Lischka, C. (2010). Sketch-based interfaces: Exploiting spatio-temporal context for automatic stroke grouping. In Taylor, R., Boulanger, P., KrÃijger, A., and Olivier, P., editors, *Smart Graphics*, volume 6133 of *Lecture Notes in Computer Science*, pages 139–151. Springer Berlin Heidelberg.

[Fitzgibbon et al., 1999] Fitzgibbon, A., Pilu, M., and Fisher, R. B. (1999). Direct least square fitting of ellipses. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21(5):476–480.

[Ispas et al., 2011] Ispas, A., Signer, B., and Norrie, M. C. (2011). An extensible digital ink segmentation and classification framework for natural notetaking. In *Proceedings of the 3rd ACM SIGCHI Symposium on Engineering Interactive Computing Systems*, EICS '11, pages 231–240, New York, NY, USA. ACM.

[Lee et al., 2008] Lee, W., de Silva, R., Peterson, E. J., Calfee, R. C., and Stahovich, T. F. (2008). Newton's pen: A pen-based tutoring system for statics. *Computers & Graphics*, 32(5):511–524.

[Ma, 2014] Ma, K. (2014). Robust dynamic symbol recognition : the clocksketch classifier. Master's thesis, Massachusetts Institute of Technology.

[Ma and Liu, 2008] Ma, L.-L. and Liu, C.-L. (2008). A new radical-based approach to online handwritten chinese character recognition. In *Pattern Recognition, 2008. ICPR 2008. 19th International Conference on*, pages 1–4. IEEE.

[Sezgin and Davis, 2008] Sezgin, T. M. and Davis, R. (2008). Sketch recognition in interspersed drawings using time-based graphical models. *Computers & Graphics*, 32(5):500–510.

[Shulman, 2000] Shulman, K. I. (2000). Clock-drawing: is it the ideal cognitive screening test? *International journal of geriatric psychiatry*, 15(6):548–561.

[Song and Davis, npub] Song, Y. and Davis, R. (unpub). Combining spatial and temporal representations for sketch interpretation in the real world.

[Yin and Sun, 2005] Yin, J. and Sun, Z. (2005). An online multi-stroke sketch recognition method integrated with stroke segmentation. In Tao, J., Tan, T., and Picard, R., editors, *Affective Computing and Intelligent Interaction*, volume 3784 of *Lecture Notes in Computer Science*, pages 803–810. Springer Berlin Heidelberg.

[Yu and Cai, 2003] Yu, B. and Cai, S. (2003). A domain-independent system for sketch recognition. In *Proceedings of the 1st International Conference on Computer Graphics and Interactive Techniques in Australasia and South East Asia*, GRAPHITE '03, pages 141–146, New York, NY, USA. ACM.

[Zhou et al., 2009] Zhou, X.-D., Wang, D.-H., and Liu, C.-L. (2009). A robust approach to text line grouping in online handwritten japanese documents. *Pattern Recogn.*, 42(9):2077–2088.