# What!?! No Rubine Features?: Using Geometric-based Features to Produce Normalized Confidence Values for Sketch Recognition

Brandon Paulson<sup>†</sup>, Pankaj Rajan<sup>†</sup>, Pedro Davalos<sup>§</sup>, Ricardo Gutierrez-Osuna<sup>‡</sup>, Tracy Hammond<sup>†</sup>

Sketch Recognition  $Lab^{\dagger}$ 

Pattern Recognition and Intelligent Sensor Machines Lab<sup>‡</sup>

Spacecraft Technology Center<sup>§</sup>

Texas A&M University

3112 TAMU

College Station, TX 77843 USA {bpaulson, pankaj, p0d9861, rgutier, hammond}@cs.tamu.edu

# Abstract

As pen-based interfaces become more popular in today's applications, the need for algorithms to accurately recognize hand-drawn sketches and shapes has increased. In many cases, complex shapes can be constructed hierarchically as a combination of smaller primitive shapes meeting certain geometric constraints. However, in order to construct higher level shapes, it is imperative to accurately recognize the lower-level primitives. Two approaches have become widespread in the sketch recognition field for recognizing lower-level primitives: gesture-based recognition and geometric-based recognition. Our goal is to use a hybrid approach that combines features from both traditional gesturebased recognition systems and geometric-based recognition systems. In this paper, we show that we can produce a system with high recognition rates while providing the added benefit of being able to produce normalized confidence values for alternative interpretations; something most geometric-based recognizers lack. More significantly, results from feature subset selection indicate that geometric features aid the recognition process more than gesture-based features when given naturally sketched data.

# 1. Introduction

Hardware supporting pen-based input has become popular in recent years as Tablet PCs, SmartBoards, and touch screens are becoming common modes of input for many applications. Sketches are used in a variety of domains to help convey ideas and designs. Computer-aided design (CAD) tools have been created to allow the visualization of abstract ideas; however, traditional CAD applications use less than intuitive interfaces. Sketching is a very natural choice as an alternative to the multi-tool selection paradigm. Many tools have been created which allow sketching to be easily incorporated into user interfaces [12, 16]. However, in order for sketching to be an effective means of input, we must develop recognizers that can accurately determine the intention of a sketcher.

Two approaches have become standard to solving the sketch recognition problem. The first approach treats input sketches simply as two-dimensional gestures [1, 15, 18]. These gesture-based recognition techniques typically focus on how a sketch was drawn rather than on what the final sketch actually looks like. The typical goal of these systems is to take an input stroke (a sampling of points in the form of x, y, and time value) and classify each one into a set of pre-defined gestures. This approach has the benefit of using mathematically sound classifiers which produce fast classifications along with normalized confidence values, but has the disadvantage of using feature sets which are userdependent and require individual training by each user to give good recognition results. Furthermore, many of these gesture-based features produce systems which are very sensitive to changes in scale and rotation.

The second approach has been to describe shapes geometrically, focusing on what the sketch looks like and less on how it was actually drawn [14, 17, 19]. Essentially, these geometric-based techniques are meant to take a single stroke as input and classify it as one of the predefined geometric primitives. These techniques are geometric in nature because they compare a stroke



Figure 1. Example scenario in which higherlevel context can disambiguate a lower-level interpretation. In this example, the high-level recognizer may realize that the ambiguous stroke is more likely to be a circle (which represents a wheel in this domain) than an ellipse because of context.

to an ideal representation of each primitive using formulas based on geometry. Primitives can then be combined hierarchically to form more complex shapes using specialized grammars [7]. Since the geometric tests used by these systems focus more on what the sketch looks like, these recognizers are typically more userand style-independent. This also means that no individual (per user) training is necessary; the only training required is that which is necessary to determine numerous geometric thresholds. The disadvantage of such a system is that geometric-based recognizers typically use numerous thresholds and heuristic hierarchies which are difficult to analyze and optimize in a systematic fashion. Inferences about generalization are hard to determine because classification is not statistical. Furthermore, the use of multiple error measures on a per shape basis makes ranking alternative interpretations difficult [14].

Ranking alternative interpretations with a normalized confidence value can be important in aiding a higher-level recognition system, which often has access to context that can help resolve ambiguity in a lowerlevel interpretation [3, 6]. For example, imagine that a user draws a shape and the low-level recognizer returns a circle interpretation with an ellipse listed as an alternative interpretation. In addition, the low-level recognizer can also supply confidence values: 55% chance of an ellipse, 45% chance of a circle. The higher-level recognizer may now be more likely to choose the circle interpretation over the ellipse interpretation given that these confidence values are so close, if context indicates that an circle is a more likely interpretation. An example of this is shown in Figure 1.

Our initial goal for this work was to find a way to combine these gesture-based and geometric-based approaches in such a way as to take advantage of the positive aspects of each (accurate classification, user independent, mathematically sound, ability to produce normalized confidence values, etc.) while avoiding many of the disadvantages. During this process, we discovered that gesture-based features were less significant in aiding recognition on freely sketched data. This discovery, along with the production of normalized confidence values, is our main contribution.

As is the case in many previous works [14, 17, 19], we focused simply on classifying low-level primitive shapes. We have chosen to classify single strokes into one of nine different shape classes (arc, line, curve, circle, ellipse, helix, spiral, polyline, and complex). This is the shape set used in a previous geometric-based recognizer, PaleoSketch [14]. In this paper we show that combining the gesture-based features from Rubine [15] with geometric-based features and using a statistical classifier can yield accurate recognition results while producing normalized confidence values. We also perform feature subset selection to determine what features are the most significant for recognition. Our findings indicate that geometric-based features contribute more significantly to the recognition of naturally drawn sketch data than gesture-based features.

## 2 Previous Work

## 2.1 Gesture-based Recognition

In 1991, Dean Rubine introduced one of the first pen-based input gesture recognition systems, GRANDMA [15]. This toolkit allowed users to specify single stroke gestures that could be trained and learned through a simple linear classifier. Rubine proposed thirteen features which could be used to classify simple gestures with an accuracy of 98% on a fifteen-class gesture set when trained with at least fifteen examples per class.

Rubine's work was extended by [1]. In this paper, Long et al. added nine new features to Rubine's existing set. They performed multi-dimensional scaling to identify correlated features and ultimately found an optimal subset that contained eleven of Rubine's original thirteen features along with six of their own. Both of these works proved to be helpful in recognizing two-dimensional gestures, but when applied to natural sketch recognition problems, the accuracy of these approaches is not optimal.

The nature of the feature sets used by these recognizers requires that gestures be drawn the same way and to the same scale every time they are drawn. For example, a clockwise, circular gesture would not be the



Figure 2. Example of a higher-level shape (pendulum) being constructed from lower-level primitives (line, circle).

same as a counter-clockwise, circular gesture. When treating sketched shapes as gestures these approaches do not perform well because they put constraints on how users must draw. Our goal is to be able to create recognition systems which are user-independent and allow users to draw as they would naturally, without having to worry about issues such as where to start a stroke or which direction to draw certain shapes. This is typically referred to as "free" or "natural" sketch recognition [8].

#### 2.2 Geometric-based Recognition

Because of the drawing constraints imposed by gesture-based classifiers, the most recent shift in sketch recognition has been toward a geometric approach, which puts virtually no drawing constraints on the user. Essentially, shape grammars such as LADDER [7] can be used to define higher level shapes as a combination of lower level primitive shapes meeting certain geometric constraints. Figure 2 gives an example of how higher-level shapes can be constructed from lower-level primitives.

In order for these higher-level recognition schemes to be effective, it is important that the low-level primitive shapes are accurately recognized. Many geometricbased recognizers have been developed to recognize low-level primitive shapes [14, 17, 19]. Unlike gesturebased techniques, these recognizers do not use statistical classifiers. Instead, they focus on determining the error between a sketched shape and its ideal version using a series of geometric tests and formulas. Some recognizers focus on developing a universal error metric such as the feature area error metric [19]. However, universal error metrics can be hard to describe and compute for more complex primitive shapes such as spirals and helixes.

In order to support more primitives, some recognizers use different error metrics for each primitive [14]. These recognizers then rely upon heuristic hierarchies and numerous thresholds to order shape interpretations. Such an approach makes it hard to generalize and prove a recognizer's success for future sketches. It also makes producing normalized confidence values difficult (the biggest downfall of such an approach).

#### 3. Features

As mentioned before, our goal is to combine gesturebased and geometric-based approaches in such a way as to take advantage of the benefits of both techniques. We will use a statistical classifier - namely a quadratic classifier [5] - that uses a feature set containing both gestural and geometric features. Our hypothesis is that this approach will lend itself more naturally to systematic optimization than using heuristic hierarchies, will maintain user independence because of the addition of geometric features, and will still enable us to return multiple, ranked interpretations with normalized confidence values.

Our feature set initially contained 44 total features. The first 31 features came from geometric tests described in [14]. The remaining 13 features are the classical gesture-based features used by Rubine [15]. The descriptions of each of these features are beyond the scope of this paper, so the interested reader is referred to [14, 15]. Table 1 gives a list of the features we used.

# 4. Data

To perform our tests, we used a dataset consisting of 1800 total sketch examples - the same dataset from [14]. Each example consists of a single stroke and is labeled with the intention of the original sketcher. A stroke is defined as the set of points (x-coordinate, y-coordinate, and time stamp) sampled between pen-down and penup events. The data samples came from 20 different users. Each user provided 90 samples (10 of each shape class). Figure 3 shows some examples from the dataset.

For our experiments we split the data set into two halves. The first half consisted of 900 examples sketched by 10 distinct users. The second half consisted of 900 examples from 10 different users. Rather than use a random 50/50 split of the overall data, we chose to split based on user to more accurately reflect how a sketch system would be used (i.e. the classifier is trained with data offline and then new users interact with the system without providing their own training data). This testing procedure allows us to determine the extent to which our algorithm is user-independent. The first half of the data was used to perform feature subset selection and training while the second set was left untouched until the validation of our final model.

1. Endpoint to stroke	12. Curve least	23. Spiral fit: avg.	34. Length of bounding
length ratio (100%)	squares error $(90\%)$	radius/bounding box	box diagonal $(20\%)$
		radius ratio (60%)	
2. NDDE (90%)	13. Polyline fit: $\#$ of	24. Spiral fit: center	35. Angle of the
	sub-strokes (70%)	closeness error $(70\%)$	bounding box diagonal
			(40%)
<b>3. DCR (90%)</b>	14. Polyline fit:	25. Spiral fit: max	36. Distance between
	percent of sub-strokes	distance between	endpoints $(10\%)$
	pass line test $(50\%)$	consecutive centers $(20\%)$	
4. Slope of the direction	15. Polyline feature	26. Spiral fit: average	37. Cosine of angle
graph $(20\%)$	area error $(80\%)$	radius estimate $(10\%)$	between endpoints $(0\%)$
5. Maximum curvature	16. Polyline least squares	27. Spiral fit: radius test	38. Sine of angle between
(40%)	error $(30\%)$	passed $(1.0 \text{ or } 0.0) (40\%)$	endpoints $(10\%)$
6. Average curvature	17. Ellipse fit: major axis	28. Complex fit: $\#$ of	39. Total stroke length
(30%)	length estimate $(20\%)$	sub-fits $(60\%)$	(20%)
7. # of corners $(30\%)$	18. Ellipse fit: minor axis	<b>29.</b> Complex fit: $\#$ of	40. Total rotation
	length estimate $(30\%)$	non-polyline	(100%)
		primitives $(50\%)$	
8. Line least squares	19. Ellipse feature area	<b>30.</b> Complex fit:	41. Absolute rotation
error $(0\%)$	error $(10\%)$	percent of sub-fits	(10%)
		that are lines $(90\%)$	
9. Line feature area error	20. Circle fit: radius	31. Complex score /	42. Rotation squared
(40%)	estimate $(30\%)$	rank (50%)	(10%)
10. Arc fit: radius	21. Circle fit: major	32. Cosine of the starting	43. Maximum speed
estimate $(0\%)$	axis to minor axis	angle $(30\%)$	(20%)
	ratio (80%)		
11. Arc feature area error	22. Circle feature area	33. Sine of the starting	44. Total time $(30\%)$
(20%)	error $(0\%)$	angle $(10\%)$	

Table 1. Features used by our recognizer. Implementation details for features 1-31 can be found in [14]. Details for features 32-44 can be found in [15]. Bold features are those chosen as the optimal subset through feature subset selection. Percentage values indicate how often a feature was chosen as optimal through various folds of subset selection.

# 5. Results

Our goal was to determine whether or not we could use a statistical classifier to classify single-stroke sketched primitives using a combination of gesturebased and geometric-based features. In addition, we wanted to determine if our approach produced classification rates that are comparable to the current best low-level system, PaleoSketch. In order to compare our approach directly to the PaleoSketch system, we have presented our results using the same 50/50 user split from [14]. However, we also wanted to determine if our final classifier was robust to other splits. Therefore, we have also presented results based on 25 folds of crossvalidation where we combined the training and testing data into a single set and randomly selected 10 users for training and 10 users for testing. Table 2 shows the results of using the full feature set along with a quadratic classifier for each of the cases just mentioned. In this table, we also present the improved accuracy achieved through feature subset selection. Because of singularities in the data, the covariance matrix was regularized by adding a small value (0.001) to its diagonal [4].

As Table 2 shows, the full feature set alone did not provide the same accuracy reported in [14]. We can also see that the split used by the original PaleoSketch system is not favorable to the quadratic classifier using the full feature set. Fortunately, the quadratic classifier can be optimized by removing features which contribute negatively to recognition.

## 5.1. Feature Subset Selection

To determine relevant features, we employed a greedy, sequential forward selection (SFS) technique to perform feature subset selection [2]. The subset se-



Figure 3. Examples of each shape class from our dataset. From left to right, starting at the top: line, arc, ellipse, helix, complex (combination of other primitives), circle, polyline, curve, spiral.

lection algorithm used a quadratic classifier as a wrapper [10] to determine the order in which features were added. Optimal features are those which are in the subset that yielded the highest accuracy during the SFS process. In the event of a tie, the subset containing the fewest number of features was selected. We performed 10 folds of SFS using random 50/50 user splits (5 users to train, 5 users to test) with our 900 sample training set, yielding 10 different feature subsets from which the final subset was obtained. In order to determine the final optimal feature subset we performed 25 folds of classification using random 50/50 user splits on each of the 10 subsets. Furthermore, we looked at each of the 44 features and analyzed how often each feature appeared as an optimal feature in one of the 10 folds of subset selection. In addition to the 10 subsets found through feature subset selection, we also created 10 other subsets. The first of the 10 additional subsets contained only features that were present in all 10 folds (100%) of subset selection. The second contained the features from the first along with features that were present in at least 9 out of 10 folds (90%) of subset selection. The third contained features present in at least 8 out of 10 folds (80%) and so forth. Figure 4 shows the classification of each subset averaged across 25 folds.

The overall best performance was obtained with the subset that contained features which were present at

Table 2. Accuracy rates achieved using different feature sets along with a quadratic classifier.

least 50% of the time in our multi-fold subset selection process (bolded in Table 1). When using only this subset with a quadratic classifier we can produce higher classification rates that more closely resemble those given by the original PaleoSketch system, seen in Table 2. A t-test confirms that the classification difference between our approach and the PaleoSketch system is statistically insignificant (t = 0.0566, p = 0.9549).



Figure 4. Average accuracy of each subset averaged across 25 folds of cross-validation on our training set. Subsets 1-10 are the subsets determined through 10 different folds of feature subset selection. Subset 11 is the subset containing only features chosen 100% of the time in the subset selection process. Subset 12 contains features chosen at least 90%. Subset 13 contains features chosen at least 80% and so on.

# 6. Discussion

One interesting observation is that one can still achieve a high classification rate, around 93%, with only the top six features (those present in at least 90% of the subset selection folds). This result indicates that these features contain a wealth of information and are highly independent of one another. In addition to providing normalized confidence values for each interpretation, our approach also has the additional benefits of: (1) being easier to code and (2) performing faster classification than the original PaleoSketch system since fewer features need to be computed.

Of the 31 features from the geometric-based recognizer, fourteen were chosen to be in the optimal subset by the subset selection algorithm. Most of these features were based on the polyline and complex tests performed by the PaleoSketch system. This is not too surprising given that the classification decision between polylines and complex fits was one of the hardest ones to make for the heuristic hierarchy used in the original PaleoSketch system. We also noticed that both the NDDE (normalized distance between direction extremes) and DCR (direction change ratio) values were chosen to be significant. This supports the observations made in [14] that these are helpful for geometric classification. Other added features included the least squares error of a curve fit, the major axis to minor axis ratio (used to distinguish circles from other ellipses), and two tests used to distinguish spirals and helixes.

The most important observation, however, is that only one of the gesture-based features was chosen as an optimal feature during subset selection: total rotation. This feature is also used by the PaleoSketch system to compute the number of rotations that a stroke makes. All other gesture-based features were determined to be insignificant, which is not surprising given the manner in which we split our data. By performing our tests with splits based on user (a more accurate representation of how a true sketch recognition system would be used), we biased our subset selection algorithm toward user-independent features. The fact that geometric-based features proved to be more significant than gesture-based features is interesting given that so many sketch recognition systems have been built on gesture-based features alone: [9, 11] to name a few.

# 7. Future Work

The most important and immediate direction for future work is to see how our approach generalizes against other classes of complex shapes. For our tests, we used a complex shape consisting of one line and one arc: the

VL/HCC Workshop: Sketch Tools for Diagramming Herrsching am Ammersee, Germany 15 September 2008 Editors: Beryl Plimmer & Tracy Hammond complex shape used in [14]. Obviously, it would be impractical to create a class for every possible primitive combination, which is infinite. In theory, a complex shape is simply an outlier class that occurs when it is determined that a shape is not one of the basic primitive shapes. Therefore, it is possible that the addition of other complex shapes would still be recognized using our approach. However, there is also the possibility that our recognizer is simply recognizing one line, one arc interpretations and not a generalized complex interpretation. We hope to collect more complex shape data to determine whether or not this is the case. We also want to look into other features, such as the "ink features" proposed in [13], to see if they can further aid recognition.

## 8. Conclusion

We have taken two traditional approaches to lowlevel sketch recognition and combined them to create a hybrid recognition scheme. By combining gesturebased and geometric-based features into a statistical classifier, we have created a recognizer that can produce highly accurate classification while maintaining user independence and allowing users to sketch freely. Furthermore, our approach can produce normalized confidence values that can be used by a higherlevel recognition system when resolving ambiguity. Through feature subset selection, we have also shown that geometric-based features are more optimal than gesture-based features in recognizing freely sketched data.

## 9. Acknowledgments

This work supported in part by NSF IIS HCC Grant #0744150 Developing Perception-based Geometric Primitive-shape and Constraint Recognizers to Empower Instructors to Build Sketch Systems in the Classroom, as well as, NSF IIS Creative IT Grant #0757557 Pilot: Let Your Notes Come Alive: The SkRUI Classroom Sketchbook. The authors would also like to thank the members of the Sketch Recognition Lab, including Dr. Patricia Lombard, for their helpful comments and editing critiques.

## References

[1] J. A. Chris Long, J. A. Landay, L. A. Rowe, and J. Michiels. Visual similarity of pen gestures. In *CHI '00: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 360–367, New York, NY, USA, 2000. ACM Press.

- [2] D. W. Aha and R. L. Bankert. A comparative evaluation of sequential feature selection algorithms. In D. Fisher and H.-J. Lenz, editors, *Learning from Data: AI and Statistics V*, chapter 4, pages 199–206. Springer-Verlag, 1996.
- [3] C. Alvarado and R. Davis. Sketchread: a multidomain sketch recognition engine. In UIST '04: Proceedings of the 17th annual ACM symposium on User interface software and technology, pages 23–32, New York, NY, USA, 2004. ACM.
- [4] J. H. Friedman. Regularized discriminant analysis. Journal of the American Statistical Association, 84(405):165–175, mar 1989.
- [5] K. Fukunaga. Statistical pattern recognition. Handbook of pattern recognition & computer vision, pages 33-60, 1993.
- [6] L. Gennari, L. B. Kara, T. F. Stahovich, and K. Shimada. Combining geometry and domain knowledge to interpret hand-drawn diagrams. *Computers & Graphics*, 29(4):547–562, 2005.
- [7] T. Hammond and R. Davis. Ladder, a sketching language for user interface developers. Computers & Graphics, 29(4):518–532, 2005.
- [8] T. Hammond, B. Eoff, B. Paulson, A. Wolin, K. Dahmen, J. Johnston, and P. Rajan. Free-sketch recognition: putting the chi in sketching. In *CHI '08: CHI* '08 extended abstracts on Human factors in computing systems, pages 3027–3032, New York, NY, USA, 2008. ACM.
- [9] J. I. Hong and J. A. Landay. Satin: a toolkit for informal ink-based applications. In UIST '00: Proceedings of the 13th annual ACM symposium on User interface software and technology, pages 63–72, New York, NY, USA, 2000. ACM.
- [10] R. Kohavi and G. H. John. Wrappers for feature subset selection. Artif. Intell., 97(1-2):273–324, 1997.
- [11] R. L. Kullberg. Mark your calendar!: learning personalized annotation from integrated sketch and speech. In CHI '95: Conference companion on Human factors in computing systems, pages 302–303, New York, NY, USA, 1995. ACM.
- [12] A. C. Long, J. A. Landay, and L. A. Rowe. "those look similar!" issues in automating gesture design advice. In *PUI '01: Proceedings of the 2001 workshop on Perceptive user interfaces*, pages 1–5, New York, NY, USA, 2001. ACM Press.
- [13] R. Patel, B. Plimmer, J. Grundy, and R. Ihaka. Ink features for diagram recognition. In SBIM '07: Proceedings of the Fourth Eurographics Workshop on Sketch-Based Interfaces and Modeling. Eurographics Association, 2007.
- [14] B. Paulson and T. Hammond. Paleosketch: Accurate primitive sketch recognition and beautification. In *IUI* '08: Proceedings of the 13th international conference on Intelligent User Interfaces, pages 1–10, New York, NY, USA, 2008. ACM Press.
- [15] D. Rubine. Specifying gestures by example. In SIG-GRAPH '91: Proceedings of the 18th annual conference on Computer graphics and interactive techniques,

VL/HCC Workshop: Sketch Tools for Diagramming Herrsching am Ammersee, Germany 15 September 2008 Editors: Beryl Plimmer & Tracy Hammond pages 329–337, New York, NY, USA, 1991. ACM Press.

- [16] E. Saund, D. Fleet, D. Larner, and J. Mahoney. Perceptually-supported image editing of text and graphics. In UIST '03: Proceedings of the 16th annual ACM symposium on User interface software and technology, pages 183–192, New York, NY, USA, 2003. ACM Press.
- [17] T. M. Sezgin, T. Stahovich, and R. Davis. Sketch based interfaces: early processing for sketch understanding. In *PUI '01: Proceedings of the 2001 work*shop on *Perceptive user interfaces*, pages 1–8, New York, NY, USA, 2001. ACM Press.
- [18] J. O. Wobbrock, A. D. Wilson, and Y. Li. Gestures without libraries, toolkits or training: a \$1 recognizer for user interface prototypes. In UIST '07: Proceedings of the 20th annual ACM symposium on User interface software and technology, pages 159–168, New York, NY, USA, 2007. ACM.
- [19] B. Yu and S. Cai. A domain-independent system for sketch recognition. In GRAPHITE '03: Proceedings of the 1st international conference on Computer graphics and interactive techniques in Australasia and South East Asia, pages 141–146, New York, NY, USA, 2003. ACM Press.