

Object Set Matching With an Evolutionary Algorithm

Andrew R. Buck, *Student Member, IEEE*,
James M. Keller, *Fellow, IEEE*, and
Marjorie Skubic, *Member, IEEE*

Department of Electrical and Computer Engineering
University of Missouri
Columbia, Missouri, USA

Marcin Detyniecki and Thomas Baerecke
Laboratoire d'Informatique de Paris 6 (LIP6)
Université Pierre et Marie Curie (UPMC)
Paris, France

Abstract—In this paper, we present an improved evolutionary method for the task of locating a group of buildings based solely on their relative spatial relationships. This problem arises in the general text-to-sketch problem of conflating a hand or machine drafted sketch of building locations to a satellite image. We use the histograms of forces to capture the relative position information between buildings and develop a method to compare building sets. This represents an extension to our previous work, allowing for larger placement perturbations and changes in orientation.

Keywords—object set matching; histograms of forces; text-to-sketch; evolutionary algorithm

I. INTRODUCTION

Consider the problem of locating a group of buildings in a satellite image given an approximate description of their spatial relationships. This occurs in the general problem of text-to-sketch (T₂S), in which a linguistic description is used to generate an approximate sketch of building and object locations [1]. This sketch can be matched to a real-world location by searching within a segmented satellite image for a set of buildings whose relative positions match those of the sketch. Our problem can be extended to any situation in which we wish to find a specific spatial configuration of objects within a large search area.

We use the histograms of forces (HoF) [2], [3] to model the spatial relationships between objects for their descriptive power and predictable response to affine transformations. For the T₂S problem, the histograms of forces are particularly useful, as they can provide a model of uncertainty as a fuzzy descriptor. We rarely have a completely accurate representation of the set of objects we are trying to match, so our model must be tolerant of slight variations in sketches.

In [4], the matching problem is solved using an evolutionary method in which the search space is covered with a population of potential matches which evolve over time into better matches. A novel mutation operator is used to guide

candidate solutions toward the optimal match. The algorithm showed promising results for cases of direct resubstitution, but did not perform well when presented with object sets that had been rotated. We address many of the limitations of this approach in this paper and show that the histograms of forces allow us to match object sets that have undergone some basic transformation.

II. EVALUATING SPATIAL RELATIONS

The core of our problem lies in our representation of the spatial relationships between objects. The histograms of forces between two objects are a powerful measure of their spatial relationship as they incorporate direction, distance, and scale into a single framework. By utilizing the HoF relationships between all objects in a set, we can develop a method to compare two sets of objects.

A. Histograms of Forces

The spatial relationship between two objects, A and B , can be represented by the force histogram, F_r^{AB} . Here, we assume that A and B are crisp, two-dimensional objects in \mathbb{R}^2 with positive finite area. The function $F_r^{AB}(\theta)$ evaluates the amount of support for the proposition, “ A is in direction θ of B .” We refer to object A as the argument and object B as the referent. The force histogram provides a measure of the scalar resultant of elementary forces exerted by the points of A on those of B in a given direction, θ (Figure 1). We can evaluate the forces between objects as an inverse ratio of d^r where d represents the distance between corresponding points in A and B , and r is used to extract different features. Of particular interest to us are the F_0 and F_2 histograms which represent the histograms of constant and gravitational force respectively. The F_0 histogram is independent of the distance between objects, while the F_2 histogram is independent of scale. Details regarding the computation and properties of the histograms of forces can be found in [2] and [3].

B. Comparing Histograms

We now turn to the problem of evaluating the similarity between two pairs of spatial relationships. Given two pairs of objects, (A, B) and (A', B') , we can define the force histograms, $F_r^{AB}(\theta)$ and $F_r^{A'B'}(\theta)$, where $r = 0$ is the histogram of constant force and $r = 2$ is the histogram of gravitational force. If the

This work is funded by the U.S. National Geospatial-Intelligence Agency NURI grant HM 1582-08-1-0020.

A. R. Buck, J. M. Keller, and M. Skubic are with the Department of Electrical and Computer Engineering, University of Missouri, Columbia, MO 65211, USA (e-mail: arb9p4@mail.mizzou.edu; kellerj@missouri.edu; skubicm@missouri.edu). M. Detyniecki and T. Baerecke are with the Laboratoire d'Informatique de Paris 6, Université Pierre et Marie Curie, Paris France.




Fig. 1 The elementary forces exerted by the points of A on those of B , used to calculate $F_r^{AB}(\theta)$. Each one pulls B in direction θ .

two pairs of objects have a similar spatial configuration, their histograms should be similar whereas if the configurations are different, the histograms will reflect this difference. We can determine a similarity measure between two histograms by computing their cross-correlation,

$$\mu_c(F_1, F_2) = \frac{\sum_{\theta} F_1(\theta)F_2(\theta)}{\sqrt{\sum_{\theta} F_1^2(\theta)} \sqrt{\sum_{\theta} F_2^2(\theta)}} \quad (1)$$

We choose the cross-correlation as our evaluation tool for its invariance to the scale of the histograms [3].

When applied to the raw histograms, this similarity measure does not account for changes in orientation between the pairs (A, B) and (A', B') . Due to the rotational invariance of the histograms of forces, if pair (A, B) is rotated counter-clockwise by an angle, φ , its force histogram becomes

$$F_r^{rot(A)rot(B)}(\theta) = F_r^{AB}(\theta - \varphi) \quad (2)$$


Following the approach in [3], we can use this property to normalize the histograms prior to comparison such that

$$\overline{F_r^{AB}}(\theta) = m^{-1}F_r^{AB}(\theta + c), \quad (3)$$


where m is the mean value of F_r^{AB} and c is the centroid. These normalized histograms always have a mean value of 1 and a centroid of 0° . Computing the mean value is straightforward (and not actually necessary if we use the cross-correlation). Calculating the centroid, however, requires that we treat the histogram as a periodic function with a period of 360° . This can be evaluated by treating each histogram value as a vector in polar coordinates and computing the average of the sum of the vectors [5]. Since our histograms are derived from linguistic descriptions, it makes sense to calculate the centroid value, c , for a pair of histograms, F_0^{AB} and F_2^{AB} , as the centroid of the main direction histogram,

$$a^{AB}(\theta) = \max\{a_0^{AB}(\theta), \min\{a_2^{AB}(\theta), b_0^{AB}(\theta)\}\}. \quad (4)$$

Here, $a_r^{AB}(\theta)$ is the degree of truth for the proposition, “ A is in direction θ of B ,” according to F_r^{AB} and $b_r^{AB}(\theta)$ is the maximum degree of truth that can reasonably be attached to this proposition. The calculation of a_r^{AB} and b_r^{AB} requires the decomposition of F_r^{AB} into effective, contradictory and compensatory forces which Matsakis *et al.* describe in detail in



(a)



(b)

Fig. 2 The main direction histogram between the pair of objects in (a) utilizes both the histogram of constant force, F_0^{AB} , and the histogram of gravitational force, F_2^{AB} . In this configuration, the F_0^{AB} and F_2^{AB} histograms (b) have different centroids. The main direction histogram, a^{AB} , provides a linguistic interpretation of the support for the statement, “ A is in direction θ of B .”

[6]. While the original definition has the intent of measuring a^{AB} in only the four cardinal directions, we evaluate a^{AB} for each value of θ in the histograms. For many cases, this step would seem unnecessary. However, in situations where the centroids of the F_0 and F_2 histograms differ greatly, the main direction histogram provides an interpretation that gives a common centroid for both the constant and gravitational forces (Figure 2). This is important because the two forces represent the same pair of objects and should therefore be rotated equally when normalizing.

C. Comparing Sets of Histograms

We now extend the comparison of a single pair of objects to the comparison of two sets of N objects. While the following comparison method can be used in multiple contexts, we use the framework of comparing a chromosome of potential buildings to a sketch in order to guide our discussion. Let us define two sets of objects: the sketch, $\mathcal{S} = \{o_1, o_2, \dots, o_N\}$, and the chromosome, $\mathcal{C} = \{c_1, c_2, \dots, c_N\}$. Our goal is to assess the degree to which the objects of the chromosome match those of the sketch. We define the force histogram set between two objects as the pair of constant and gravitational force histo-

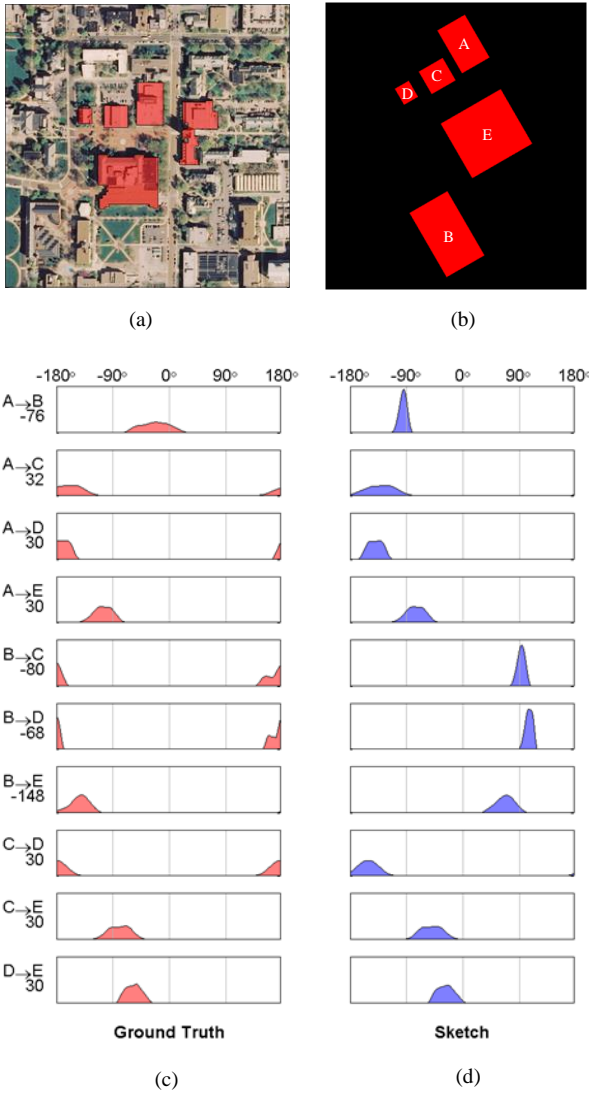


Fig. 3 Calculating the difference in main direction for a histogram set. The sketch in (b) is a simplification of the building set in (a), rotated counterclockwise 30° , and with object B moved to the lower-left of the sketch. (c) shows the histogram set for the ground truth, while (d) shows the histogram set for the sketch. The angular difference in main direction (measured in degrees) for each histogram pair is listed along the left-hand side of the ground truth. Together, these values constitute the set of angle differences, D .

grams, $\mathcal{F}^{AB} = \{F_0^{AB}, F_2^{AB}\}$. The set of all force histogram relationships for the sketch is then,

$$\mathcal{H}_S = \{\mathcal{F}^{o_i o_j} \mid (o_i, o_j) \in \mathcal{S} \times \mathcal{S} \mid i < j\} = \{h_1^S, h_2^S, \dots, h_M^S\}, \quad (5)$$

and similarly for the chromosome,

$$\mathcal{H}_C = \{\mathcal{F}^{c_i c_j} \mid (c_i, c_j) \in \mathcal{C} \times \mathcal{C} \mid i < j\} = \{h_1^C, h_2^C, \dots, h_M^C\}. \quad (6)$$

Here, we have used h_i^C and h_i^S to represent the force histogram set between a single pair of objects and made use of the semantic inverse between objects, which states that

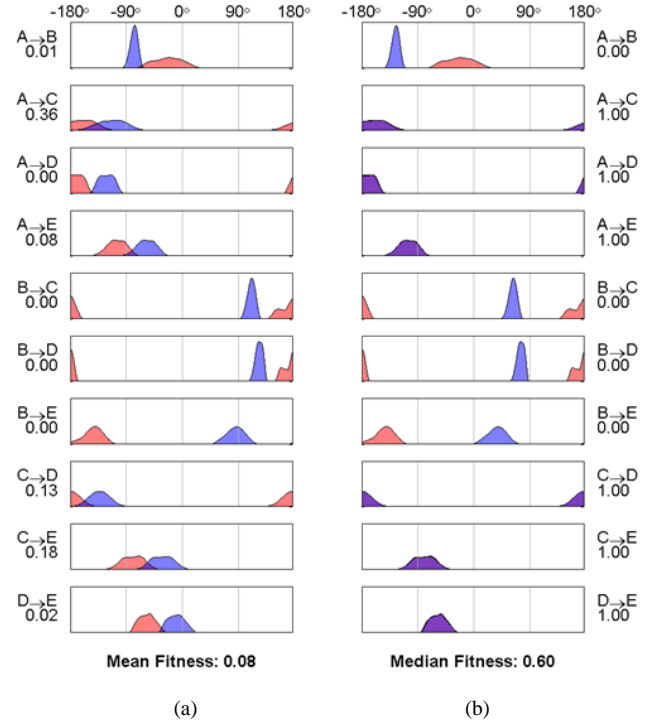


Fig. 4 Differences between calculating the optimal rotation angle for the histogram sets in Figure 3 using the mean (a) and the median (b) of the set of angle differences, D . The mean angle of D is -18° , while the median angle is 30° . Shifting the histograms of the sketch (blue) by -18° gives a very poor alignment, while shifting by 30° gives a perfect alignment for six of the histograms. The numbers along the sides of the histograms represent the cross-correlation between the ground truth histograms and those of the sketch. The total fitness is average of the cross-correlations, which is 0.08 for the mean angle rotation and 0.60 for the median angle rotation.

$$F_r^{BA}(\theta) = F_r^{AB}(\theta + 180^\circ), \quad (7)$$

allowing us to consider only one histogram set for each pair [2].

Our strategy will be to compute the similarity between each pair, h_i^C and h_i^S , for all i and take the average similarity between pairs as the overall similarity (or fitness value in the case of an evolutionary algorithm) of the two object sets. This assumes that we have already solved the assignment problem and are considering only a single mapping of \mathcal{S} onto \mathcal{C} . A more robust comparison method might iterate over several permutations to find the best ordering of \mathcal{S} and \mathcal{C} .

An additional feature of a set of more than two objects is the set of relative angles between objects. This information is lost if the histogram set for each pair is normalized independently. We wish to retain this information by rotating one set of histograms to the angle which best matches the other set. Although either set could be rotated, we choose to rotate the histograms of the sketch in our notation, mimicking how a map is rotated to match the topography of the real world. To find the optimal angle, we compute the difference in main direction for each pair of histograms,




Fig. 5 The weighting function, $S(\theta)$, used as a filter to assign high fitness values only to histograms with small differences in main direction.

$$D = \{\text{centroid}(h_i^S) - \text{centroid}(h_i^C) \mid 1 \leq i \leq M\} \quad (8)$$

$$= \{d_1, d_2, \dots, d_M\}$$

where $\text{centroid}(\mathcal{F}^{AB})$ is the centroid of the main direction histogram, a^{AB} . D represents the set of angle differences between the histograms of the sketch and the chromosome. Figure 3 shows an example where a sketch representing an actual location has been simplified, rotated, and drawn with one building in the wrong position. Such misrepresentations can easily occur when constructing a sketch from a linguistic description. The optimal rotation angle can be found from this set by selecting either the mean or median of the angles in D . Choosing the median angle gives more stable results (Figure 4), as it becomes less likely that a small number of poor histograms will influence an otherwise perfect set of matches.

Calculating the mean and median again requires that we treat D as a set of periodic values. We could use polar vector summation to find the mean, however this requires the use of trigonometric functions which significantly slow the runtime performance of the algorithm when many calculations must be made. We can instead treat these values as linear, so long as we shift the values to an appropriate range that does not cross the periodic boundary. We assume for the purposes of this paper that no object completely surrounds another object and that such a shift is possible. Calculating the median value requires only that we evaluate each element of D to find the value which minimizes the expression from [5],

$$m(\theta) = 180^\circ - \sum_{i=1}^M |180^\circ - |d_i - \theta||. \quad (9)$$

Having found the best rotation angle, θ_{Best} , we can rotate the histograms of the sketch to match those of the chromosome. The overall similarity then becomes

$$\mu(\mathcal{H}_S, \mathcal{H}_C) = \frac{1}{M} \sum_{i=1}^M f[h_i^S(\theta - \theta_{Best}), h_i^C(\theta), \beta] \quad (10)$$

$$f(\mathcal{F}^{AB}, \mathcal{F}^{A'B'}, \beta) = \beta \mu_C(F_0^{AB}, F_0^{A'B'}) + (1 - \beta) \mu_C(F_2^{AB}, F_2^{A'B'}) \quad (11)$$

where β is from the range $[0, 1]$ and represents the weighting between the histograms of constant and gravitational force. A value of 1 represents a complete bias toward the histogram of constant force, where a value of 0.5 is an equal weighting.

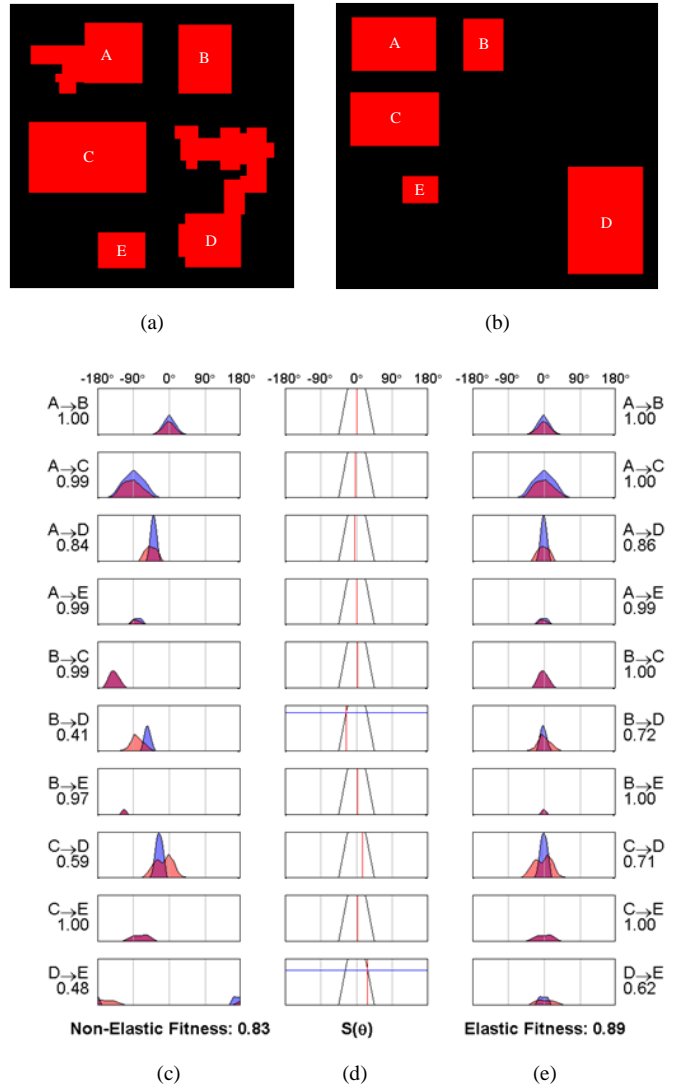


Fig. 6 Comparison of the elastic and non-elastic methods for evaluating fitness. The object set in (a) is taken from the ground truth of a real world location. The sketch in (b) is a simplification of (a) with object D significantly misplaced. (c) shows the computation of the non-elastic fitness, where the numbers to the left of each histogram represent the individual cross-correlation values. (d) and (e) show the computation of the elastic fitness. (d) is the weighting function $S(\theta)$, and (e) shows the normalized histograms, where the numbers to the right of each histogram represent the weighted cross-correlation values.

D. Elastic Angles

As an alternative to rotating each histogram of the sketch by the global best rotation angle, we can exercise a little more control over the match quality by rotating each histogram individually. We do this by first computing the best rotation angle as before, and then computing the overall similarity using the normalized histograms. We add a weighting term, $S(\theta)$, to ensure that only histograms which are close to the optimal rotation angle of the sketch are used in computing the average. $S(\theta)$ takes the form of a trapezoid centered at 0° and maps the range $[-180^\circ, 180^\circ]$ into $[0, 1]$ (Figure 5). The overall similarity is then

Algorithm I: Evolutionary Matching Algorithm

Input: \mathcal{H}_R and \mathcal{H}_S Create an initial population of chromosomes, P Set generation counter $t = 0$ **while** stopping criteria is not met **for each** chromosome, \mathcal{C}_P , in P Mutate \mathcal{C}_P and create \mathcal{C}_C Replace \mathcal{C}_P with \mathcal{C}_C according to (14) **end** **if** t is a multiple of $t_{\text{Repopulate}}$ Replace the least fit members of P with new random chromosomes **end** $t = t + 1$ **end****Output:** Top chromosomes in P

$$\mu_{\text{Elastic}}(\mathcal{H}_S, \mathcal{H}_C) = \frac{1}{M} \sum_{i=1}^M f(\bar{h}_i^S, \bar{h}_i^C, \beta) S(\theta_{\text{Best}} - d_i). \quad (12)$$

Figure 6 shows an example which compares both elastic and non-elastic fitness methods for a situation in which a building in the simplified sketch is misplaced. The elastic method is able to give somewhat higher fitness values, but as we shall see in the experiments section, this method tends to give universally higher fitness values, which may or may not be desirable.

III. AN EVOLUTIONARY ALGORITHM

Armed now with the tools to compare object sets using relative spatial locations, we develop our modified algorithm to match a sketch of building locations to a segmented satellite image. We build upon the evolutionary algorithm described in [4] with a few significant modifications. The general structure of the technique is described in Algorithm I.

A. Database Representation

The input to the algorithm is a sketch containing N objects, $\mathcal{S} = \{o_1, o_2, \dots, o_N\}$. We define the objects of the sketch to be the goal set, where each object, o_i , is defined as a polygon in \mathbb{R}^2 with positive finite area. We restrict our discussion here to sketches with no overlapping objects. Our search area is a map consisting of a set of L reference objects, $\mathcal{R} = \{x_1, x_2, \dots, x_L\}$, which must each satisfy the restrictions of the goal objects. Each reference object is linked to its k nearest neighbors by a set of force histograms, and the histograms between all linked reference objects are collectively gathered into a set of relationships,

$$\mathcal{H}_R = \{\mathcal{F}^{x_i x_j} \mid (x_i, x_j) \in \mathcal{R} \times \mathcal{R}\} \quad (13)$$




Fig. 7 The reference set, \mathcal{R} , used for our experiments. The set contains 2467 buildings from downtown Columbia, MO and the University of Missouri campus.


where x_j is one of the k nearest neighbors of x_i . Each chromosome represents a set of objects, $\mathcal{C} = \{c_1, c_2, \dots, c_N\} \subset \mathcal{R}$ which could potentially match the goal set, \mathcal{S} , such that $c_i = o_i$ for $1 \leq i \leq N$. To satisfy the requirements of our matching method, we require that each chromosome be fully connected, meaning that $\mathcal{F}^{c_i c_j} \in \mathcal{H}_R$ for all pairs (c_i, c_j) .

B. Search Process

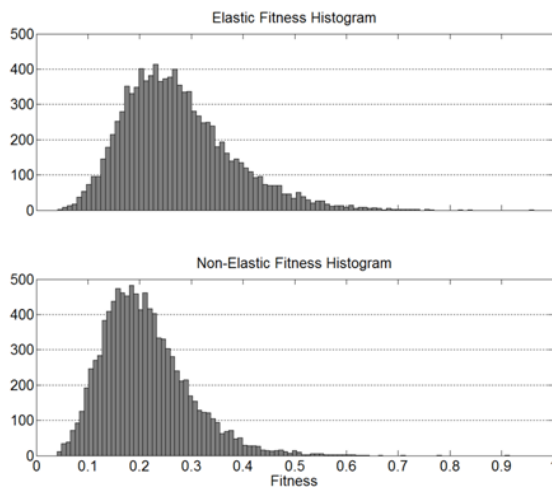
The search process is initialized by randomly choosing an initial population of chromosomes from the set of reference objects, and ensuring that each chromosome is fully connected. The fitness of each chromosome is evaluated by finding the similarity of its histograms with those of the sketch using the method described in Section II. At each generational cycle of the algorithm, we mutate each chromosome as in [4], by randomly choosing a building and replacing it with the building that best replaces it. The fitness of the child chromosome, \mathcal{C}_C , determines if it will replace its parent, \mathcal{C}_P , by the formula,

$$P(\mathcal{C}_C \text{ replaces } \mathcal{C}_P) = \frac{\mu(\mathcal{C}_C)}{\mu(\mathcal{C}_C) + \mu(\mathcal{C}_P)} \quad (14)$$

where $\mu(\mathcal{C})$ is the fitness of chromosome \mathcal{C} . This replacement method ensures that either the parent or child chromosome will survive to the next generation, and allows each potential solution the opportunity to evolve over time without being dropped from the population for having a fitness value which is too low. At regular intervals during the search process, new random chromosomes are added to the population, replacing the lowest scoring individuals. This keeps the algorithm from getting stuck in local minima, and forces it to search in new



(a)



(b)

Fig. 8 Comparison of the elastic and non-elastic angle methods. (a) shows the difference between the elastic and non-elastic fitness as building “D” from the sketch in Figure 6a is shifted downward. (b) shows the histograms generated by comparing the sketch in Figure 6a to 10,000 randomly generated object sets using both the elastic and non-elastic methods.

locations. Searching continues until a suitable match is found or some other stopping criterion is met, such as a maximum number of generations.

IV. EXPERIMENTS

Our experiments were conducted using a reference set of buildings from Columbia, MO, shown in Figure 7. We calculate a histogram relationship between each building and its 50 nearest neighbors using a 2° interval for each histogram. We first examine the usefulness of using elastic angles for the fitness calculation. We then evaluate several test sketches using two different versions of the algorithm.

A. Testing Elastic Angles

When discussing the calculation of the elastic angles in Section II, we made note of the fact that the elastic angle method tends to produce higher fitness values. The graph in

TABLE I

	Evolutionary Algorithm Parameters	
	Old Method	New Method
Max Generations	10,000	10,000
Population Size	493	493
Replacement Frequency	100 Generations	50 Generations
Percent of Replacement	10%	50%
Histogram Bias, β	1	0.5
Rotation Method	Mean Angle	Median Angle
Elastic Angles	Non-Elastic	Elastic

Figure 8a shows the elastic and non-elastic fitness values of the sketch in Figure 6a as object D is moved 100 pixels in a downward direction. With no displacement, both methods give a fitness of 1. As the building is moved, both fitness values decrease, with the elastic fitness remaining higher than the non-elastic fitness. The elastic fitness plot appears less smooth because in our implementation the best angle must be a multiple of 2° . This allows us to rotate histograms quickly by changing the indexing. A more accurate method could interpolate values between histogram bins, but at the cost of greater computation time.

If our goal is to evaluate the fitness of a single object set, the elastic angles provide a way to increase this value. However, if we are trying to locate the best match among many possible candidates we notice that the elastic angle method will increase the fitness value of all matches, making a very high fitness value less unique. To test this, we compare the sketch in Figure 6a to 10,000 randomly generated object sets and construct histograms of the computed fitness values using both the elastic and non-elastic methods (Figure 8b). The histograms show that the average fitness value is low in both cases, but distributed somewhat higher for the elastic angle method. This means that for a given fitness threshold, more matches will likely be returned.

B. Testing the Evolutionary Algorithm

We have made several improvements to the algorithm originally proposed in [4]. They are:

- Both histograms of constant and gravitational force can be used in calculating fitness values.
- The alignment of histograms is performed by calculating the centroid of the main direction histogram rather than taking the maximum value.
- The calculation of the best rotation angle is performed using the median of the angle differences rather than the mean.
- We have the option of computing the final fitness value using elastic angles.

To test these improvements, we created 100 random sketches from our reference set with five buildings each. The sketches were chosen from actual building locations and

TABLE II

	Matching Results	
	Old Method	New Method
Percent of tests which found correct match	89.4%	95.1%
Average Generations	3001	2104

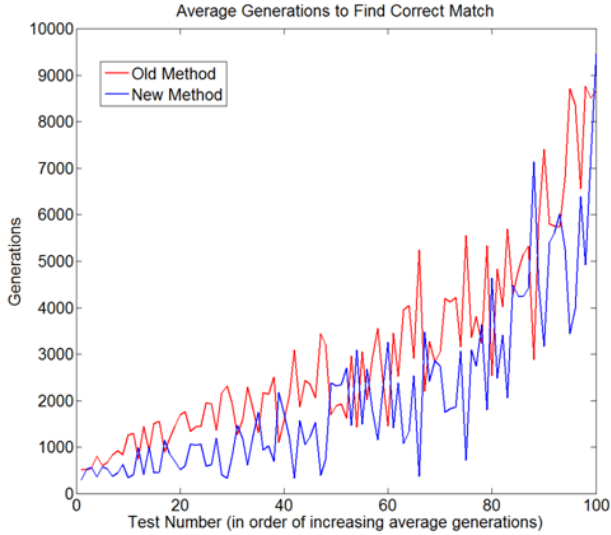


Fig. 9 The average number of generations required for the algorithm to find the correct match for both the old and new method. The new method required fewer average generations in most cases.

simplified such that each building was represented by its bounding box. Each sketch was then rotated to a random angle to give an arbitrary orientation. We ran our algorithm on each sketch using two different sets of parameters, defined in Table I. For each sketch, we take the average of 30 trials using each configuration.

The two methods are intended to model the old matching algorithm defined in [4] and the new method given in this paper. We chose to use 20% of the total buildings in the reference set as the population size and set the generation limit to 10,000. We then ran the algorithm until we found a chromosome that contains the same buildings used to generate the sketch. In a practical setting, we would not know these buildings beforehand. Our stopping criteria could use a fitness threshold, however we then risk never finding the match if the threshold is set too high. A qualitative analysis would be required to determine when a good solution is found. Since we are mainly interested in the algorithm's ability to find the real location where the sketch originated, we use the building comparison criteria to end our search. We note the average number of generations required to find the correct match, if the correct match is found at all.

Table II shows the matching results for the old and new methods. The new method using the median angle rotation and elastic angles had a higher matching rate, and required fewer average generations. Figure 9 shows a plot of the average number of generations required to find the correct match for each test. The new method had fewer average generations for most test cases. This suggests that while both methods can

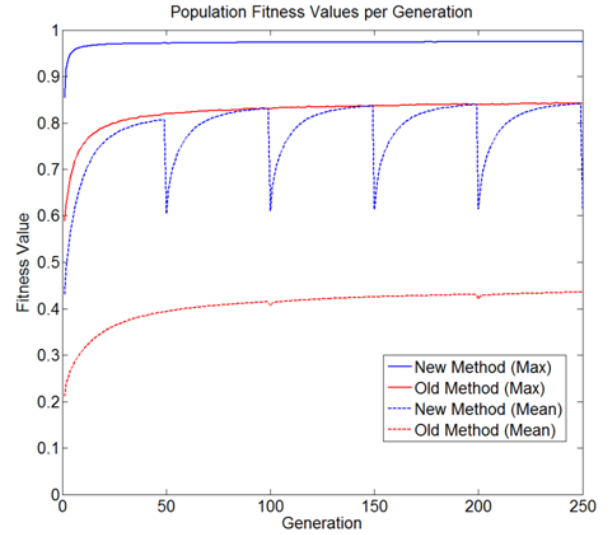


Fig. 10 The population fitness values for each generation, averaged over all of the trials. The new method (blue) has significantly higher fitness values due to the use of elastic angles. The mean values of the population drop every 50 generations for the new method and every 100 generations for the old method as new chromosomes are added.

give good results, the new method presented in this paper is more stable under varying conditions. Since the goal is to convert linguistic descriptions to sketches and then match to a geospatial database, there is a considerable likelihood that some of the linguistic descriptions are wrong (the building “to my right” was actually to the left), the parsing can give inaccurate spatial locations, or there are missing or incomplete descriptions; all making the matching process much harder. The modified approach ameliorates some of these problems.

As the search process proceeds, the average fitness values of the chromosomes in the population increases. Figure 10 shows the maximum and average fitness values of the population for both the old and new methods, averaged over all of the trials. The new method has significantly higher fitness values due to the use of elastic angles. The mean fitness values of the population drop every time new chromosomes are added. The new method repopulates 50% of the population every 50 generations, while the old method replaces 10% every 100 generations. This more aggressive approach to repopulation ensures that the algorithm always searches new locations without becoming stuck in local minima.

The values we found for the old method are much higher than those reported in [4]. Upon further analysis, we discovered that the calculation of the mean angle in [4] failed to account for the periodicity of the histograms. Additionally, histograms were aligned by their maximum values as opposed to their centroids or main direction. The convergence statistics also are based on a fitness threshold instead of determining if the algorithm actually found a correct match. These issues have been addressed by the method presented in this paper.

V. CONCLUSION

In this paper, we improved upon a method for comparing object sets using their relative spatial relationships. We found that the histograms of forces are a powerful analytical tool for representing the relationship between a single pair of objects. We used a set of histogram relationships to model a set of multiple objects and evaluated different approaches for comparing them. The median angle rotation gives stability to the search algorithm, while elastic angles can be used to increase the overall fitness values. We found that the histograms themselves allow for small variations in the object sets and can be used without the additional framework of the elastic angles if required.

The task of matching groups of buildings is a specific application of object set matching in general, and there are still many improvements to be made. Additional features such as object labels and road networks will improve matching by pruning the search space. We also only currently consider a single mapping between object sets. Future work in this area will explore multiple possible mappings between object sets, as well as other search methods such as subgraph isomorphism techniques.

REFERENCES

- [1] I. Sledge and J. Keller, "Mapping Natural Language to Imagery: Placing Objects Intelligently," *IEEE Proceedings, International Conference on Fuzzy Systems (FUZZ-IEEE)*, Jeju Island, Korea, August, 2009, pp. 518-524.
- [2] P. Matsakis and L. Wendling, "A New Way to Represent the Relative Position between Areal Objects," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, pp. 634-643, 1999.
- [3] P. Matsakis, J. Keller, O. Sjahputera, and J. Marjamaa, "The Use of Force Histograms for Affine-Invariant Relative Position Description," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, pp. 1-18, Jan. 2004.
- [4] A. Buck, J. Keller, and M. Skubic, "A Modified Genetic Algorithm for Matching Building Sets with the Histograms of Forces," *Proceedings IEEE Congress on Evolutionary Computation, World Congress on Computational Intelligence*, Barcelona, Spain, July 2010.
- [5] N. I. Fisher, *Statistical Analysis of Circular Data*, Cambridge University Press, 1993, pp. 30-37.
- [6] P. Matsakis, J. Keller, L. Wendling, J. Marjamaa, and O. Sjahputera, "Linguistic Description of Relative Positions in Images," *IEEE Transactions of Systems, Man, and Cybernetics—Part B: Cybernetics*, vol. 31, pp. 573-588, Aug. 2001.