Continuous quantification of uniqueness and stereoscopic vision

Val Petran*[a], Frank Merat[b]
[a]Artificial Perception Technologies Inc., Croton On Hudson, NY, USA
[b]Dept. of Electrical Engineering and Computer Science , Case Western Reserve University, USA

ABSTRACT

In this paper we introduce the concept of continuous quantification of uniqueness. Our approach is to construct an algorithm that computes a fuzzy set membership function, which given any inter-object dissimilarity metric and it's variability, measures the probability that an entity of interest will not be confused with other similar entities in a search space. We demonstrate use of this algorithm by applying it to stereoscopic computer vision, in order to identify which of several sub-problems pertaining to solution of the classic stereoscopic correspondence problem are least likely to be solved incorrectly, and hence are most well suited to greatest confidence first approaches.

Keywords: Uniqueness, Uniqueness Quantification, Stereoscopic Vision, Correspondence Problem, Fuzzy Logic, Fuzzy Set, Computer Vision, Correspondence

1. INTRODUCTION

Suppose a manufacturer mails a variety of acoustic musical instruments to a store that is located several thousand miles away, and along with each one, includes a compact disc that contains an audio recording of that specific instrument being played by a famous and talented musician. During transit the recordings somehow get separated from the instruments and all paperwork, packaging, or markings that could be used to re-unite them is damaged or lost. A clerk at the store is faced with the task of deciding which instruments to attempt to hastily re-unite with their audio recordings in preparation for an imminent trade show, by comparing how they sound when he plays them to how each of the audio recordings sound. Which instruments should he focus on if it is his goal to minimize the likelihood that he will make a mistake? One strategy could be, to focus on those instruments that have the most unique timbres. For example if the shipment contains 20 violins but only one cow bell, it is unlikely that he will fail to pair up the cow bell with the correct audio recording. How unusual of a timbre is sufficiently unusual in order for him to succeed? That depends not only on how similar the timbres of the instruments in the shipment are to each other, but also, on the variability of each instrument's timbre as a function of who is playing it, temperature, humidity, the impact of a long journey in a shipping container, and how adept the clerk or the diagnostic equipment available to him are at producing dissimilarity estimates with low variability when repeatedly comparing two identical signals. If there were no such sources of variability, even very minor differences between timbres of instruments in the shipment would suffice to ensure success. On the other hand, even if all such sources of variability were completely absent, if two of the instruments had identical timbres, the probability of correctly re-uniting them with their factory audio recordings would be merely 1/2.

The above discussion presents a specific instance of a general problem, comprised of needing to decide which of several equally effective objects to use, for a purpose that requires recognition of the chosen object in different environments that contain objects with which it could be confused, by performing comparisons using a metric that relies on properties whose measurements do unfortunately contain a lot of environment dependent variability.

In this paper we propose a strategy for solving such problems. It is comprised of deciding to use those objects that are deemed most unique by a measure we introduce in the next section. Here is a one sentence summary of what that measure does: The measure deems an object to be highly unique if and only if the distribution of differences between repeated measurements of the object's properties does not substantially overlap distributions obtained by subjecting the objects with which it could be confused to the same procedure.

We empirically assessed the effectiveness of this uniqueness quantification measure by examining it's ability to predict the circumstances in which some very basic and widely used stereoscopic vision techniques compute correct matches, and the circumstances in which they compute matches that regardless of whether or not they are correct, are useful, in the sense that they are compatible with the Epipolar Constraint[1,2,8].

Visual Information Processing XX, edited by Zia-ur Rahman, Stephen E. Reichenbach, Mark A. Neifeld,
Proc. of SPIE Vol. 8056, 80560E · © 2011 SPIE · CCC code: 0277-786X/11/$18 · doi: 10.1117/12.883620

2. PROPOSED UNIQUENESS QUANTIFICATION MEASURE

2.1 Definition

Let $\mathbf{x}$ denote an object of interest, whose uniqueness we want to assess by comparing it to other objects in a search space S using a metric d. Assume that the search space S contains $\mathbf{x}$. We define the uniqueness $U(\mathbf{x})$ of $\mathbf{x}$ in the search space S with respect to the metric d as:

$$U(\mathbf{x}) = 1/\left(1 + \sum_{\substack{\mathbf{a}_j \in S \\ \mathbf{a}_j \neq \mathbf{a}_{gm}}} Area(H(\mathbf{a}_{gm}) \bigcap H(\mathbf{a}_j)) / Area(H(\mathbf{a}_{gm}))\right) \tag{1}$$

where $\mathbf{a}_1$ through $\mathbf{a}_k$ denote k objects that comprise the search space S and $\mathbf{a}_{gm}$ denotes any fixed but arbitrary element of the search space S at which $d(\mathbf{a}_{gm},\mathbf{x})$ is equal to the global minimum of the values $\{d(\mathbf{a}_j,\mathbf{x})\}$. (Usually the global minimum occurs at only one location in the search space, but that is not guaranteed to always be the case, and would most certainly not occur if the output of the metric d were constant. The latter could happen if all of the objects involved were identical and the act of measuring them were completely free of variability)

In the musical instrument example, $\mathbf{a}_1$ through $\mathbf{a}_k$ could denote the musical instruments that were in the shipment that was received by the music store clerk and $\mathbf{x}$ could denote one of those instruments whose identity the clerk does not know, that was used to produce a specific factory audio recording. The metric d could be an algorithm that compares audio recordings that were shipped to the store, to audio recordings of the shipped instruments that were produced by the clerk in the store. Given these conventions, whenever the clerk evaluates $d(\mathbf{a}_j,\mathbf{x})$ he is comparing the sound of an instrument $\mathbf{a}_j$ that he has chosen from the shipment, to the sound of an instrument $\mathbf{x}$ whose identity he is attempting to discover, and which he selected implicitly when he decided to compare a specific factory audio recording to a recording of $\mathbf{a}_j$ that he made in the store. For all he knows, $\mathbf{x}$ and $\mathbf{a}_j$ may very well be the same instrument.

For each object $\mathbf{a}_j$, $H(\mathbf{a}_j)$ denotes a histogram of the values that can be obtained by repeatedly using the metric d to compare each of many measurements of $\mathbf{a}_j$ that are obtained in various environments of interest, to each other. The histograms pertaining to objects that comprise the search space S need to be comparable to each other, so, either they need to be normalized, or they each need to have been created using contributions from precisely the same number of measurements.

In the musical instrument example, each $H(\mathbf{a}_j)$ needs to include numerous dissimilarity measurements obtained by comparing recordings of the instrument $\mathbf{a}_j$ that were made in the store by the clerk, to recordings of the musical instrument $\mathbf{a}_j$ that were made at the factory when it was played by a talented and famous musician. $H(\mathbf{a}_j)$ denotes an entity that is only theoretically obtainable. It could be estimated by repeatedly sending identical sets of instruments and recordings of those instruments, to the clerk without allowing the recordings to become disassociated from the instruments during shipping, however it is not a quantity that the clerk can possibly compute after having received only one shipment in which the correct matching between factory audio recordings and shipped instruments has been lost. Fortunately, as suggested by empirical evidence presented later in this paper, it is expected the clerk can productively harness the uniqueness quantification measure U by very crudely estimating the histograms $H(\mathbf{a}_j)$ using information that is available to him.

$Area(H(\mathbf{a}_{gm}))$ represents the area of the histogram $H(\mathbf{a}_{gm})$

$Area(H(\mathbf{a}_{gm}) \bigcap H(\mathbf{a}_j))$ represents the area of the intersection of the histograms $H(\mathbf{a}_{gm})$ and $H(\mathbf{a}_j)$

To facilitate theoretical analysis we also present an alternative definition of $U(\mathbf{x})$ that is comprised of the formula which is obtained by replacing all of the histograms that appear in equation (1) above, with the probability distributions that best describe them, the intersection of any two histograms with a function that assigns to each point x, the minimum of D1(x) and D2(x), where D1 and D2 denote the probability distributions that have replaced those histograms, and the area computing operation with integration. When taking this latter approach $Area(H(\mathbf{a}_{gm}) \bigcap H(\mathbf{a}_j))$ is replaced by the integral of the minimum of the two distributions that best describe the histograms $H(\mathbf{a}_{gm})$ and $H(\mathbf{a}_j)$, and $Area(H(\mathbf{a}_{gm}))$ is replaced by the value 1.0.

2.2 Properties

The uniqueness quantification measure U has the following properties:

The outputs it produces are real numbers that are greater than 0 and less than or equal to 1.

If d looks like an inverted impulse function, for example if $d(\mathbf{a}_{gm},\mathbf{x})$ is a small value, and all of the other $d(\mathbf{a}_j,\mathbf{x})$ are large, in the sense that the histograms $H(\mathbf{a}_{gm})$ and $H(\mathbf{a}_j)$ do not intersect as long as $\mathbf{a}_{gm}$ is not $\mathbf{a}_j$, then $U(\mathbf{x})=$ 1.0 which we interpret as meaning that $\mathbf{x}$ is highly unique in the search space S with respect to the metric *d*.

If d looks like a sum of a finite number of inverted impulse functions, for example if n of the $d(\mathbf{a}_j,\mathbf{x})$ are all equal to the same value $d(\mathbf{a}_{gm},\mathbf{x})$ and and the remaining $d(\mathbf{a}_j,\mathbf{x})$ are large in the sense that their associated histograms do not intersect $H(\mathbf{a}_{gm})$ at all, then it is likely that $U(\mathbf{x})$ will be equal to a value that is close to $1/n$. Furthermore, if the histograms $H(\mathbf{a}_{gm})$ and each of the $H(\mathbf{a}_j)$ all have the exact same shape and size and if the values $d(\mathbf{a}_{gm},\mathbf{x})$ and each of the $d(\mathbf{a}_j,\mathbf{x})$ are in fact the means of those histograms, then $U(\mathbf{x})$ will be equal to $1/n$.

A direct corollary of the above property is that if *d* looks like a constant function then $U(\mathbf{x})$ approaches 0 as the number of elements in *S* approaches infinity. We interpret values of $U(\mathbf{x})$ that are close to zero as indications that $\mathbf{x}$ is not at all unique in the search space *S* with respect to the metric *d*.

If an object $\mathbf{a}_{gm}$ is very different from each element $\mathbf{a}_j$ of a collection of objects { $\mathbf{a}_j$ }, in the sense the histogram $H(\mathbf{a}_j)$ does not intersect $H(\mathbf{a}_{gm})$ and the objects { $\mathbf{a}_j$ } change, but not enough to eliminate this absence of intersections, then the value of $U(\mathbf{x})$ remains unchanged. In summary the uniqueness of an object is not impacted by changes in it's similarity to other objects that are, and after the change continue to be, very different from it.

3. APPLICATION TO STEREOSCOPIC VISION

3.1 Overview

Next, we present an example that illustrates how to apply the above described uniqueness quantification measure, to the classic stereoscopic vision correspondence problem, and begin by providing a brief overview of what that problem is.

Given two images of the same scene taken from slightly different points of view, for each pixel in one image, the goal is to find the pixel in the other image, which corresponds to it in the sense that photons which contributed to creation of those pixels originated from the same small volume of space. Here the term "small volume of space" describes that small volume of space containing opaque or at least translucent matter with which photons collided immediately prior to colliding with the sensor arrays (for example retinas or CCDs) that captured the images, rather than a small volume of space containing original sources of photons (for example portions of light-bulbs or the sun) . Each of the above described pixel specific problems is called a single point correspondence problem, and the union of these sub-problems is called the correspondence problem. Since as a function of changing point of view, images of distant objects move more slowly across a camera's sensor array than images of objects that are close to the camera, solution to the correspondence problem is all that is needed in order to determine which of any two pixels represents matter that is further from the camera, in other words, to perceive depth. Due to occlusions (which due to parallax are not guaranteed to be the same in both images) it is not always possible to obtain such knowledge. Any pair of images, of the above described variety, is called a stereogram, and since it can be shown that it is possible to construct stereograms which equally well describe more than one spatial matter distribution, all algorithms focused on solving the correspondence problem must at least implicitly engage in some sort of judicious guessing, that is not guaranteed to succeed. The empirical success of the human stereoscopic vision system does however guarantee that there exist algorithms which will almost always guess correctly when processing real world images, rather than pathological contrived counter-examples. To date, a large number of algorithms that do a good job solving this problem have been developed, however none are perfect. Scharstein and Szeliski[12] provide an extensive taxonomy that summarizes and compares many of them.

We applied uniqueness quantification to the stereoscopic vision correspondence problem by using it to decide which of many single point correspondence problems were least likely to be solved incorrectly. For each small volume of space $\mathbf{x}$ that pertained to a single point correspondence problem, our estimate of $U(\mathbf{x})$ made use of a metric *d* that compared the image of $\mathbf{x}$ in one photograph, to images in a different photograph taken from a slightly different point of view, of elements of a set (of small volumes of space $\mathbf{a}_1$ through $\mathbf{a}_k$) that was likely to contain $\mathbf{x}$.

We present continuous quantification of uniqueness neither as a stand alone algorithm for solving the stereoscopic vision correspondence problem, nor as an alternative intended to replace other techniques, but rather, as yet another addition to a collection of basic techniques (which include bidirectional matching and occlusion detection[15], automated computation of the Epipolar Constraint[1,2,8] from image data and its application to image rectification, a variety of image subset comparison metrics[5], multi[3] resolution processing, extraction of edges through multi scale LOG[3] filtering, cooperative[9,15] processing, and dynamic[6] image data driven adjustment of the size of the area that is scrutinized by image subset comparison metrics) which can each fail to produce correct matches when used alone, but which together can, when judiciously combined, be used to greatly reduce the incidence of errors in solutions to correspondence problems pertaining to stereograms that are devoid of unusual pathologies.

3.2 Experimental assessment of uniqueness quantification efficacy pertaining to prediction of compatibility of computed matches with the Epipolar Constraint

Although all correct matches comply with the Epipolar Constraint, matches do not need to be correct in order to do so. Matches which are Epipolar Constraint compliant are legitimate descriptions of actual locations in space. Whether or not those matches are correct depends on whether or not matter that is visible to the cameras that captured a stereogram actually exists at those locations. Regardless of whether or not Epipolar Constraint compliant matches are correct, they are useful in the sense that they can provide builders of stereo rigs that are equipped with independently rotating cameras with the raw data that is needed for computation of the Epipolar Constraint. Faugeras, Luong and Maybank [1,2,8] described robust techniques for accomplishing this even if a small fraction of the matches used as inputs are incorrect. We regard uniqueness quantification as a way to provide their technique with a way to enhance screening of matches that are used as input data ,either to improve accuracy or reduce computational burden.

Our empirical assessment of the effectiveness of the uniqueness quantification measure U included execution of one hundred repetitions, of eight experiments, for each of three un-rectified stereograms (depicting a chess set, a fern, and a park using 320x240 pixel images), pertaining to which ground truth disparity was not known.

Each experiment began by randomly choosing a collection of single point correspondence problems. This was done by choosing 900 locations in one of the two images that comprise a stereogram. For some of the experiments the choices of locations were completely random, and for others they were randomly chosen locations on edges. Edges were computed merely by subtracting the image from a version of itself that was shifted both vertically and horizontally by one pixel (a crude but computationally efficient approximation to LOG filtering) and selection of edge resident locations was carried out randomly, by repeatedly raster scanning the entire edge map and if edge map brightness exceeded a very high threshold, using a random number generator to decide whether or not that location would be included. After each such pass over the entire edge map, the threshold was decreased, thereby providing less intense edge resident locations with an opportunity to be included. Our intent was to ensure that edge resident locations everywhere in the image had an opportunity to participate and that the most well defined edges would be favored but not exclusively.

In each experiment an attempt was made to compute a bidirectional match for each location in the above described data set, using the SAD (Sum of Absolute Value of Differences)[5] metric to compare a rectangular image subset centered on each pixel in one image for which a match was sought, to a subset having the same size and geometry and centered on any pixel in another image that was being evaluated as a candidate match. Like our choice of edge extraction technique, SAD is not the most sophisticated available tool, but it is not computationally demanding , has been widely studied, and performs well especially when used in conjunction with bidirectional matching. The sizes of pixel neighborhood image subsets (a.k.a "image patches"), that were compared by SAD in any one experiment were all identical. Experiments using 3x3, 5x5, 7x7, and 17x17 pixel neighborhoods were carried out.

In each experiment, the uniqueness of each bidirectional match was assessed as follows: The metric d was assumed to be SAD applied to image patches and configured in one of the specific above described ways. We very crudely approximated $H(\mathbf{a}_{gm})$ by an isosceles triangle whose base and height were both equal to $d(\mathbf{a}_{gm},\mathbf{x})$ and whose maximum value occurred at the location $d(\mathbf{a}_{gm},\mathbf{x})$. Even more crudely we approximated each $H(\mathbf{a}_j)$ by an isosceles triangle whose width and height were also equal to $d(\mathbf{a}_{gm},\mathbf{x})$, but whose maximum value was located at $d(\mathbf{a}_j,\mathbf{x})$. Note that this latter simplification implies that it is highly likely that a different estimate of $H(\mathbf{a}_j)$ will be used by each single point correspondence problem specific uniqueness assessment, and that the resulting uniqueness assessments are therefore not as comparable to each other as they would be if a more computationally intense approach were taken in which results obtained from an attempt to find a match for each specific $\mathbf{a}_j$ were used to guide estimation of $H(\mathbf{a}_j)$.

The uniqueness of each bidirectional match, was set equal to the minimum of the uniqueness measures computed for each direction of search. These two values were often comparable, but not always. Such sporadic asymmetry, was also displayed by the magnitudes of the SAD global minimums associated with the two directions of search. As long as a match turned out to be bidirectional, it was included in subsequent processing, regardless of differences between the magnitudes of computed uniquenesses or SAD global minimums associated with the two directions of search.

As demonstrated by data presented in the results section, it turned out that uniqueness quantification provided useful information in spite of the egregiously crude approximations described above, and this suggests that an opportunity to improve it's performance by using more refined estimates of the quantities it requires, still exists.

For each experiment, compliance of bidirectional matches (obtained by performing two dimensional search of un-rectified images,) with a pre-computed Epipolar Constraint was assessed for both a test group and a control group. The control group was comprised of all bidirectional matches whose uniquenesses did not exceed mean uniqueness by more than a multiple of one standard deviation, and the test group was comprised of those bidirectional matches that did. Various multiples were tried, and resulting outcomes are presented in this paper's results section.

3.3 Experimental assessment of uniqueness quantification efficacy pertaining to prediction of computed disparity correctness

Our empirical assessment of the effectiveness of the uniqueness quantification measure *U* also included execution of one hundred repetitions, of eight experiments, for each of three of Scharstein, Pal and Hirschmuller's half size rectified stereograms[4,11] (depicting art, laundry, and a Moebius strip, whose images we cropped to a size of 640x480 pixels), and pertaining to which Scharstein et all published ground truth disparity information[4,11] computed using a structured lighting technique[13] developed by Scharstein and Szeliski. These latter experiments differed from those described in the above section in only two respects: First, the success criterion (rather than being defined as compliance with the Epipolar Constraint,) was defined to be correctness of match, assessed by comparing computed matches to Scharstein et all's published ground truth data[4,11]. Second, search spaces were defined to be one dimensional subsets of corresponding Epipolar Lines[1,2,8] (which were all horizontal and parallel to each other since the images being processed were rectified images).

3.4 Results

Each screenshot depicted in Figures 1 through 6 below, graphically presents results we obtained from one iteration of one of the above described experiments. Except for the Chess and Fern stereograms, all of the images we experimented with were in color. The screenshots present grayscale renditions of those images in order to make perception of computed bidirectional matches (represented by colored plus signs superimposed onto images) more easy. Although evaluation of compliance with success criteria was objective and carried out computationally using software, some readers of this paper might enjoy subjectively assessing correctness of bidirectional matches presented in our screenshots and/or their compliance with the Epipolar Constraint by crossing their eyes, and fusing stereograms to perceive three dimensional surfaces. Correct matches will be perceived as plus signs painted onto surfaces, incorrect matches that are compatible with the Epipolar Constraint will be perceived as plus signs either floating in mid air or embedded in the interiors of solid objects, and incorrect matches that are incompatible with the Epipolar Constraint will be perceived as sources of eye strain. Readers who are viewing an electronic color version of this paper may be interested to know that the color of each pair of corresponding plus signs communicates some additional potentially interesting but nonessential information, namely the uniqueness of each presented match using the following conventions:

$0.99 \leq$ green ≤ 1, $0.40 \leq$ white < 0.99, $0.20 \leq$ yellow < 0.40, $0.05 \leq$ purple < 0.20, $0 <$ red ≤ 0.05

Each row of Tables 1, 2 and 3 that are presented after the screenshots, summarizes results obtained from 100 repetitions of one experiment. A key documenting the precise meaning of each column in those tables, appears at the end of this section.

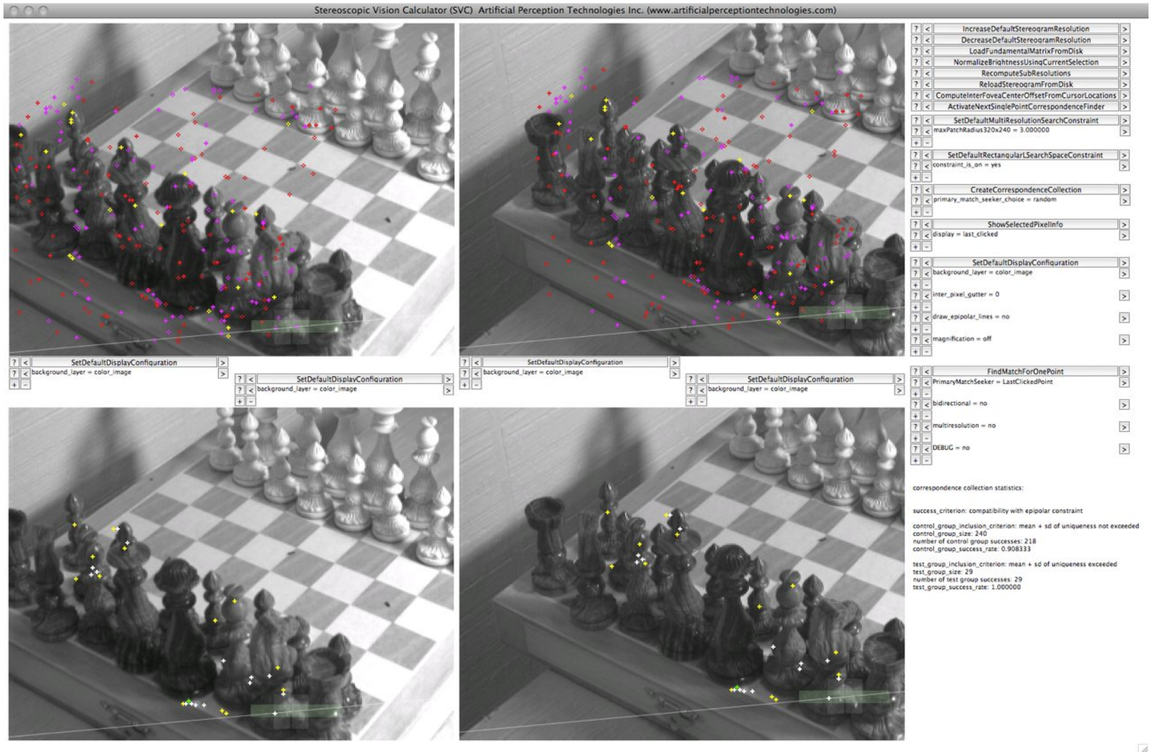

Figure 1. Outcome of 1 iteration of experiment with Chess stereogram, randomly chosen locations, and 7x7 pixel SAD

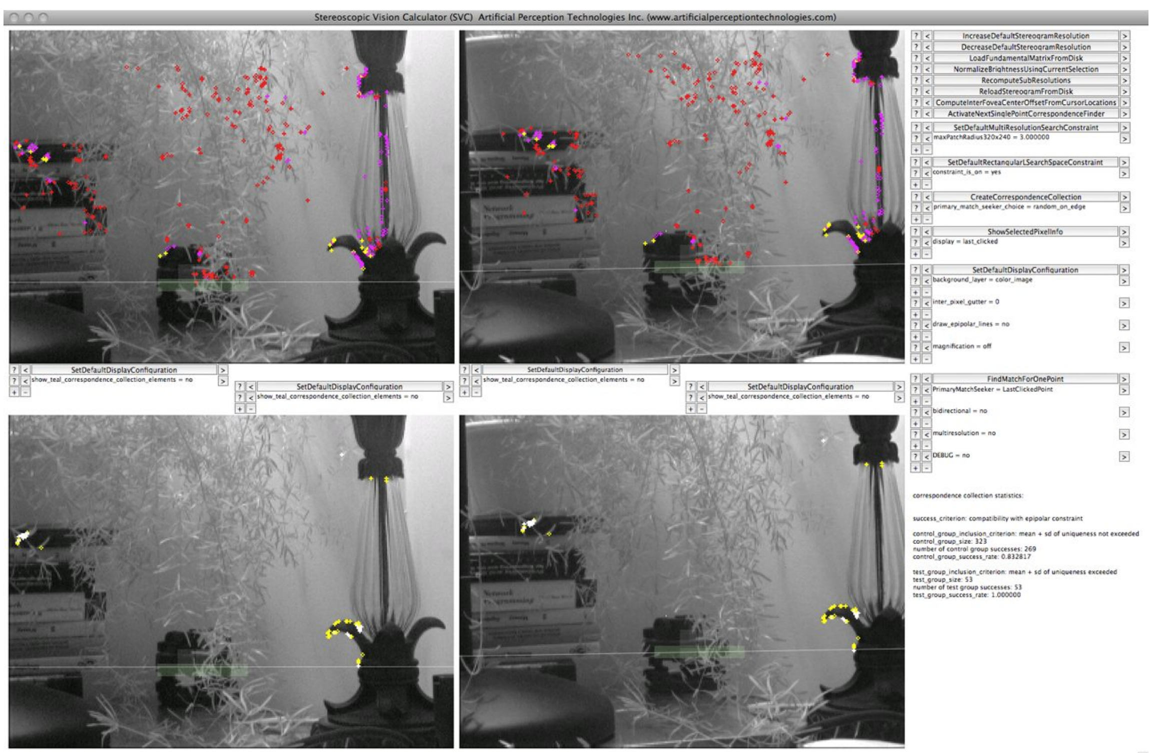

Figure 2. Outcome of 1 iteration of experiment with Fern stereogram, randomly chosen edge locations, and 7x7 pixel SAD

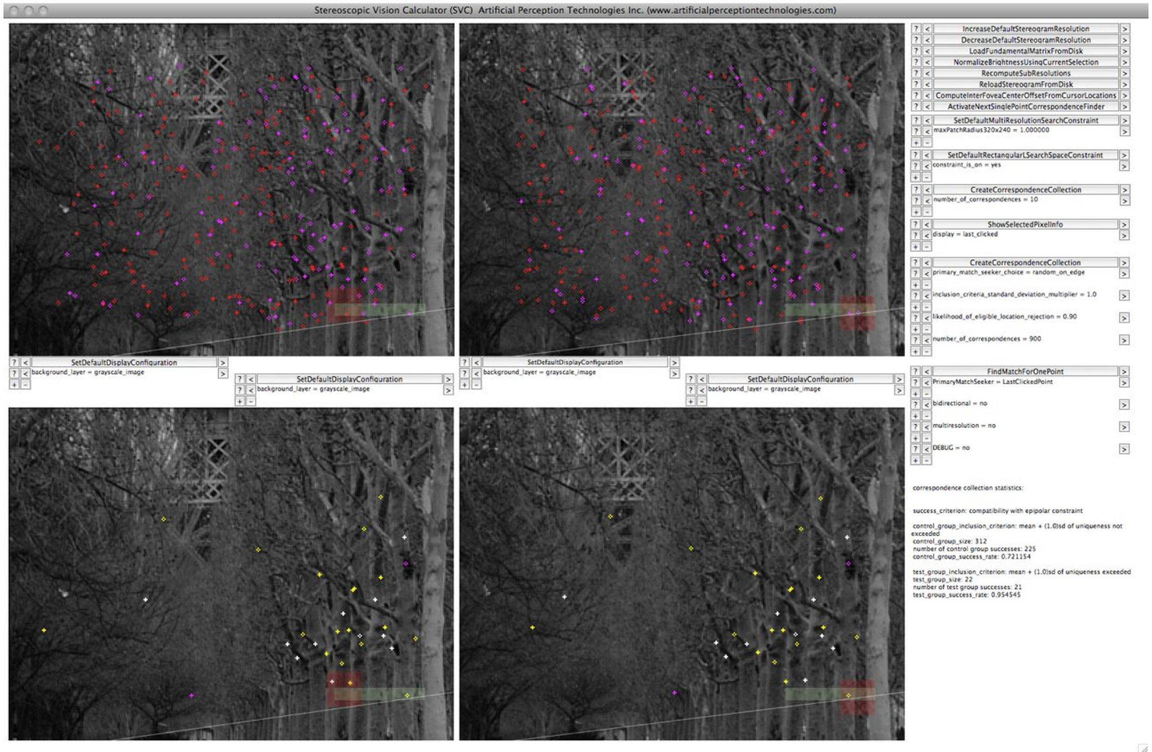

Figure 3. Outcome of 1 iteration of experiment with Park stereogram, randomly chosen edge locations, and 3x3 pixel SAD

Figure 4. Outcome of 1 iteration of experiment with Art stereogram, randomly chosen locations, and 7x7 pixel SAD

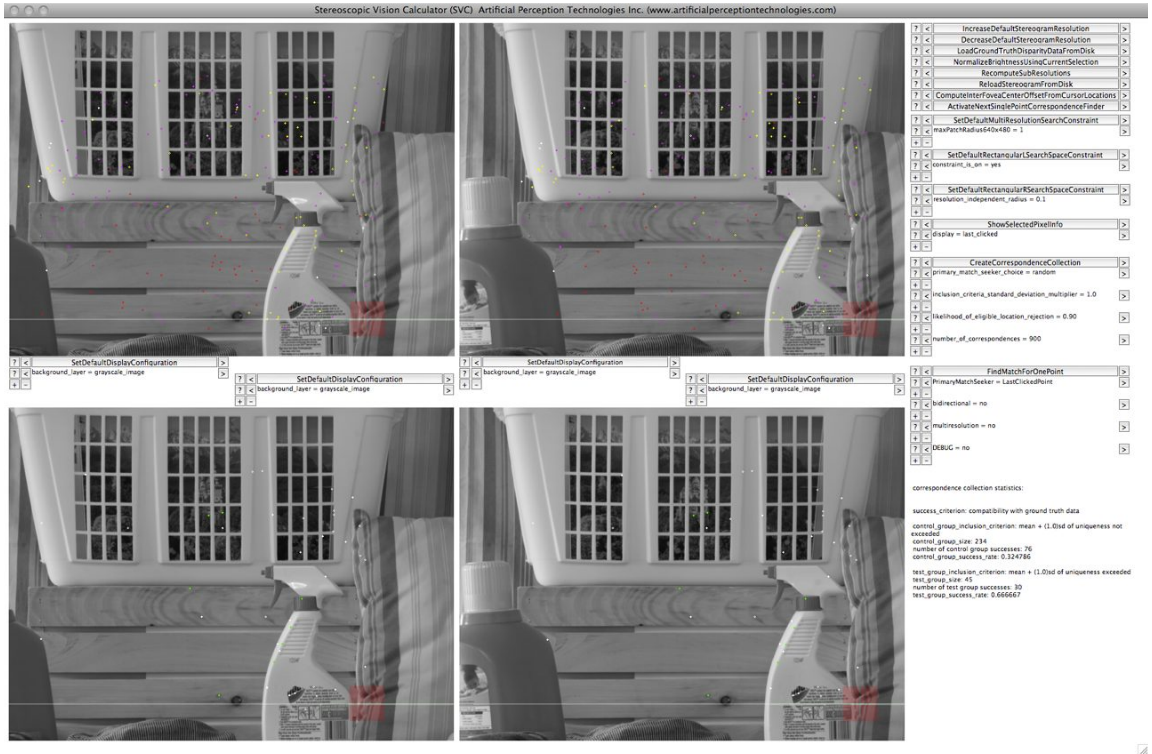

Figure 5. Outcome of 1 iteration of experiment with Laundry stereogram, randomly chosen locations, and 3x3 pixel SAD

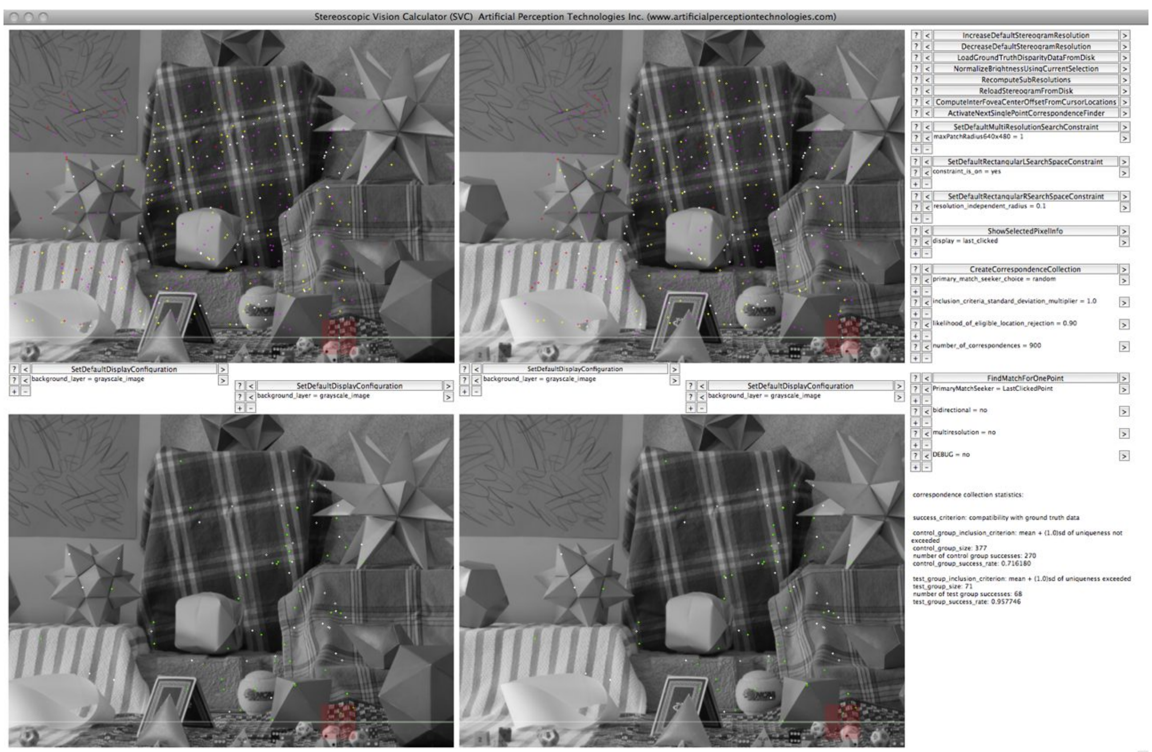

Figure 6. Outcome of 1 iteration of experiment with Moebius stereogram, randomly chosen locations, and 3x3 pixel SAD

Table 1. Empirical assessment of the usefulness of uniqueness quantification for predicting whether or not computed matches comply with the Epipolar Constraint.

Stereogram	Success Criterion = compatibility with	Search Space	Location Selection	Patch Size	Mean Control Group Size	Mean Control Group Success Rate	Control Group Success Rate Stdev	Uniqueness Boundary Between Test Group And Control Group	Mean Test Group Size	Mean Test Group Success Rate	Test Group Success Rate Stdev
Fern	Epipolar Constraint	2D	rand	3x3	226.57	78.49%	2.52%	1.0 $\sigma_U + \mu_U$	25.59	85.46%	7.17%
Fern	Epipolar Constraint	2D	rand	5x5	200.53	83.00%	2.5%	1.0 $\sigma_U + \mu_U$	15.81	97.91%	3.43%
Fern	Epipolar Constraint	2D	rand	7x7	194.54	86.90%	2.45%	1.0 $\sigma_U + \mu_U$	16.34	99.78%	1.1%
Fern	Epipolar Constraint	2D	rand	17x17	255.41	95.69%	2.11%	1.0 $\sigma_U + \mu_U$	19.77	99.82%	0.91%
Fern	Epipolar Constraint	2D	rand-edg	3x3	272.55	87.12%	1.61%	1.0 $\sigma_U + \mu_U$	44.82	96.88%	2.44%
Fern	Epipolar Constraint	2D	rand-edg	5x5	277.12	92.83%	1.84%	1.0 $\sigma_U + \mu_U$	40.64	99.34%	1.27%
Fern	Epipolar Constraint	2D	rand-edg	7x7	281.05	93.63%	1.33%	1.0 $\sigma_U + \mu_U$	46.99	100%	0%
Fern	Epipolar Constraint	2D	rand-edg	17x17	372.49	97.16%	0.69%	1.0 $\sigma_U + \mu_U$	47.56	100%	0%
Chess	Epipolar Constraint	2D	rand	3x3	203.79	80.39%	2.77%	1.0 $\sigma_U + \mu_U$	27.09	92.92%	4.59%
Chess	Epipolar Constraint	2D	rand	5x5	218.69	88.26%	2.14%	1.0 $\sigma_U + \mu_U$	27.89	99.77%	0.77%
Chess	Epipolar Constraint	2D	rand	7x7	241.60	91.83%	1.66%	1.0 $\sigma_U + \mu_U$	30.11	100%	0%
Chess	Epipolar Constraint	2D	rand	17x17	322.90	96.99%	2.19%	1.0 $\sigma_U + \mu_U$	38.86	99.92%	0.83%
Chess	Epipolar Constraint	2D	rand-edg	3x3	262.82	88.43%	1.51%	1.0 $\sigma_U + \mu_U$	46.92	99.37%	1.05%
Chess	Epipolar Constraint	2D	rand-edg	5x5	322.91	92.80%	1.02%	1.0 $\sigma_U + \mu_U$	47.89	99.86%	0.51%
Chess	Epipolar Constraint	2D	rand-edg	7x7	362.73	94.53%	1.23%	1.0 $\sigma_U + \mu_U$	51.25	100%	0%
Chess	Epipolar Constraint	2D	rand-edg	17x17	370.13	98.50%	0.5%	1.0 $\sigma_U + \mu_U$	65.82	100%	0%
Park	Epipolar Constraint	2D	rand	3x3	305.93	70.14%	3.04%	1.0 $\sigma_U + \mu_U$	32.41	89.95%	5.34%
Park	Epipolar Constraint	2D	rand	5x5	344.71	87.30%	1.78%	1.0 $\sigma_U + \mu_U$	33.7	99.88%	0.61%
Park	Epipolar Constraint	2D	rand	7x7	406.11	93.80%	1.79%	1.0 $\sigma_U + \mu_U$	40.28	99.84%	1.56%
Park	Epipolar Constraint	2D	rand	17x17	579.07	99.28%	0.33%	1.0 $\sigma_U + \mu_U$	46.44	100%	0%
Park	Epipolar Constraint	2D	rand-edg	3x3	254.72	81.53%	2.28%	1.0 $\sigma_U + \mu_U$	30.45	91.58%	4.35%
Park	Epipolar Constraint	2D	rand-edg	5x5	269.72	91.82%	1.33%	1.0 $\sigma_U + \mu_U$	30.29	100%	0%
Park	Epipolar Constraint	2D	rand-edg	7x7	303.82	93.86%	1.03%	1.0 $\sigma_U + \mu_U$	29.77	100%	0%
Park	Epipolar Constraint	2D	rand-edg	17x17	463.78	99.72%	0.22%	1.0 $\sigma_U + \mu_U$	33.34	100%	0%

The time required to perform one iteration of each of the above described experiments (using an Apple Macintosh laptop) increased substantially as a function of patch size. Approximate durations per experiment were less than 1 second when using 3x3 pixel neighborhoods, 2 seconds when using 5x5 pixel neighborhoods, 3 seconds when using 7x7 pixel neighborhoods, and 20 seconds when using 17x17 pixel neighborhoods.

The combination of this latter fact with the above presented data suggests that uniqueness quantification can be used either to reduce error rates, or if a certain rate of error is acceptable, to reduce the computational burden that must be incurred to achieve that rate.

Table 2. Empirical assessment of the usefulness of uniqueness quantification for predicting whether or not computed matches are correct

Stereogram	Success Criterion = Compatibility With:	Search Space	Location Selection	Patch Size	Mean Control Group Size	Mean Control Group Success Rate	Control Group Success Rate Stdev	Uniqueness Boundary Between Test Group And Control Group	Mean Test Group Size	Mean Test Group Success Rate	Test Group Success Rate Stdev
Art	ground truth disparity	1D	rand	3x3	267.14	44.47%	2.82%	1.0 $\sigma_U+\mu_U$	50.53	82.26%	4.98%
Art	ground truth disparity	1D	rand	5x5	284.86	53.55%	2.61%	1.0 $\sigma_U+\mu_U$	55.24	85.32%	4.49%
Art	ground truth disparity	1D	rand	7x7	298.07	57.2%	2.96%	1.0 $\sigma_U+\mu_U$	59.79	84.8%	4.93%
Art	ground truth disparity	1D	rand	17x17	316.48	52.27%	2.69%	1.0 $\sigma_U+\mu_U$	50.76	83.44%	4.8%
Art	ground truth disparity	1D	rand-edg	3x3	363.03	38.63%	2.36%	1.0 $\sigma_U+\mu_U$	69.33	48.49%	5.48%
Art	ground truth disparity	1D	rand-edg	5x5	385.91	46.85%	2.08%	1.0 $\sigma_U+\mu_U$	68.78	55.47%	5.43%
Art	ground truth disparity	1D	rand-edg	7x7	392.75	49.17%	2.20%	1.0 $\sigma_U+\mu_U$	61.24	58.14%	6.07%
Art	ground truth disparity	1D	rand-edg	17x17	350.04	43.25%	2.59%	1.0 $\sigma_U+\mu_U$	65.84	58.36%	5.4%
Laundry	ground truth disparity	1D	rand	3x3	233.52	30.73%	3.3%	1.0 $\sigma_U+\mu_U$	44.16	70.67%	6.49%
Laundry	ground truth disparity	1D	rand	5x5	256.86	39.68%	2.99%	1.0 $\sigma_U+\mu_U$	49.46	72.12%	6.24%
Laundry	ground truth disparity	1D	rand	7x7	279.88	44.33%	2.95%	1.0 $\sigma_U+\mu_U$	52.33	72.32%	6.18%
Laundry	ground truth disparity	1D	rand	17x17	400.01	50.84%	2.28%	1.0 $\sigma_U+\mu_U$	63	70.89%	5.54%
Laundry	ground truth disparity	1D	rand-edg	3x3	271.71	17.27%	2.03%	1.0 $\sigma_U+\mu_U$	42.19	31.96%	6.39%
Laundry	ground truth disparity	1D	rand-edg	5x5	282.46	21.31%	2.05%	1.0 $\sigma_U+\mu_U$	44.28	37.83%	6.11%
Laundry	ground truth disparity	1D	rand-edg	7x7	300.32	24.33%	2.25%	1.0 $\sigma_U+\mu_U$	50.21	37.23%	6.99%
Laundry	ground truth disparity	1D	rand-edg	17x17	434.61	36.9%	2.09%	1.0 $\sigma_U+\mu_U$	73.4	44.38%	5.33%
Moebius	ground truth disparity	1D	rand	3x3	383.35	69.52%	2.11%	1.0 $\sigma_U+\mu_U$	75.55	91.19%	3.73%
Moebius	ground truth disparity	1D	rand	5x5	444.45	75.01%	1.96%	1.0 $\sigma_U+\mu_U$	84.88	93.85%	2.47%
Moebius	ground truth disparity	1D	rand	7x7	476.85	76.3%	1.66%	1.0 $\sigma_U+\mu_U$	88.09	95.15%	2.38%
Moebius	ground truth disparity	1D	rand	17x17	530.78	74.4%	1.56%	1.0 $\sigma_U+\mu_U$	92.24	96.93%	1.65%
Moebius	ground truth disparity	1D	rand-edg	3x3	393.65	61.58%	2.15%	1.0 $\sigma_U+\mu_U$	156.9	79.53%	2.51%
Moebius	ground truth disparity	1D	rand-edg	5x5	466.31	66.09%	2%	1.0 $\sigma_U+\mu_U$	110.8	84.6%	2.64%
Moebius	ground truth disparity	1D	rand-edg	7x7	501	66.45%	1.6%	1.0 $\sigma_U+\mu_U$	89.12	89.71%	3.18%
Moebius	ground truth disparity	1D	rand-edg	17x17	537.22	63.85%	1.77%	1.0 $\sigma_U+\mu_U$	73.99	75.17%	4.38%

The data in Table 2 above suggests that unlike for the case of compliance with the Epipolar Constraint described in Table 1, correctness of match comparable to that obtainable by quantifying uniqueness can not easily be obtained merely by increasing patch size.

This is not surprising since for many popular metrics, like SAD[5], SSD[5], and Normalized Cross Correlation[5], a very large pixel neighborhood can in certain situations, have a high likelihood of producing matches that conform to the Epipolar Constraint, but are incorrect. Consider for example the case of attempting to find a match for a pixel that represents a very small object in the foreground, (for example a gnat which is close to cameras that are capturing the stereograms being processed) , against a more distant flat background (for example a billboard depicting a vegetable garden) for which, due to its highly unique pixel neighborhoods, and lack of depth discontinuities and associated occlusions, it is very likely that correct disparities will be computed. When using a large pixel neighborhood (like for example 17x17) the contribution to image subset dissimilarities computed by metrics like SAD, that is made by the comparatively small set of pixels which represent the small object (like the gnat), can easily be overwhelmed by the substantially larger set of contributions from pixels representing the more distant background (like the billboard). The net effect is that an effort focused on finding a match for a small foreground object (like the gnat), ends up finding a match for a location on the more distant large object (like the billboard). As we suggested before, this is not a problem if the goal is merely to produce input data for an Epipolar Constraint computation. If however the goal is to combine uniqueness quantification with knowledge of an Epipolar Constraint that has already been computed, in order to attempt to correctly solve all single point correspondence problems pertaining to a stereogram, then there can exist a reason other than mere avoidance of large computational expenditure, to use pixel neighborhoods that are as small as possible.

Table 3. Impact of choice of Uniqueness boundary between test group and control group

Stereogram	Success Criterion = Compatibility With:	Search Space	Location Selection	Patch Size	Mean Control Group Size	Mean Control Group Success Rate	Control Group Success Rate Stdev	Uniqueness Boundary Between Test Group And Control Group	Mean Test Group Size	Mean Test Group Success Rate	Test Group Success Rate Stdev
Moebius	ground truth disparity	1D	rand	3x3	75.6	43.45%	5.75%	-1.0 $\sigma_U+\mu_U$	383	78.91%	2%
Moebius	ground truth disparity	1D	rand	3x3	122.25	50.29%	4.6%	-0.8 $\sigma_U+\mu_U$	335.5	81.42%	2.13%
Moebius	ground truth disparity	1D	rand	3x3	164.18	53.95%	3.78%	-0.6 $\sigma_U+\mu_U$	296.4	83.53%	2.19%
Moebius	ground truth disparity	1D	rand	3x3	200.4	57.64%	3.58%	-0.4 $\sigma_U+\mu_U$	256.6	85.17%	2%
Moebius	ground truth disparity	1D	rand	3x3	234.48	60.4%	3.18%	-0.2 $\sigma_U+\mu_U$	227.2	86.15%	2.44%
Moebius	ground truth disparity	1D	rand	3x3	264.15	62.77%	2.93%	0.0 $\sigma_U+\mu_U$	198.1	87.22%	2.65%
Moebius	ground truth disparity	1D	rand	3x3	289.57	64.22%	2.48%	0.2 $\sigma_U+\mu_U$	171.1	87.63%	2.56%
Moebius	ground truth disparity	1D	rand	3x3	312.16	65.98%	2.82%	0.4 $\sigma_U+\mu_U$	147.5	88.81%	2.62%
Moebius	ground truth disparity	1D	rand	3x3	343.36	67.29%	2.4%	0.6 $\sigma_U+\mu_U$	117.5	89.93%	2.81%
Moebius	ground truth disparity	1D	rand	3x3	366.64	68.4%	2.39%	0.8 $\sigma_U+\mu_U$	93.7	90.85%	3%
Moebius	ground truth disparity	1D	rand	3x3	384.67	69.21%	2.34%	1.0 $\sigma_U+\mu_U$	75.11	91.43%	2.7%
Moebius	ground truth disparity	1D	rand	3x3	397.42	70.03%	2.21%	1.2 $\sigma_U+\mu_U$	60.24	92.61%	3.42%
Moebius	ground truth disparity	1D	rand	3x3	414.74	70.84%	2.39%	1.4 $\sigma_U+\mu_U$	46.43	92.82%	3.47%
Moebius	ground truth disparity	1D	rand	3x3	417.77	71.21%	2.26%	1.6 $\sigma_U+\mu_U$	42.46	92.05%	4.14%
Moebius	ground truth disparity	1D	rand	3x3	418.23	71.34%	2.27%	1.8 $\sigma_U+\mu_U$	43.63	92.37%	3.59%
Moebius	ground truth disparity	1D	rand	3x3	416.38	71.44%	2.18%	2.0 $\sigma_U+\mu_U$	42.75	93.31%	4.25%

The data presented in the three preceding tables (specifically the differences between Mean Test Group Success Rate, and Mean Control Group Success Rate) suggests that uniqueness quantification can provide tangible benefits.

Examination of how the difference between Mean Control Group Size, and Mean Test Group Size varies across the rows of Table 3 suggests that when uniqueness quantification is used to filter out all but the most unique matches, the resulting data set has a relatively high probability of containing correct matches, but contains few elements. When however uniqueness quantification is used to filter out all but the least unique matches, it still eliminates many incorrect matches yet does not decrease the size of the resulting data set substantially more than a widely used technique like bidirectional matching.

In each of the above tables:

1.The Stereogram column describes to which stereogram each row of data pertains.
2.The Success Criterion = Compatibility With: column describes whether success was defined as compliance of computed matches with the Epipolar Constraint, or agreement of computed matches with Scharstein et all's ground truth disparity data[4,11].
3.The Search Space column describes whether the Epipolar Constraint was used to constrain bidirectional search to a one dimensional image subset (indicated by 1D) or whether the search process did not make use of the Epipolar Constraint (indicated by 2D).
4.The Location Selection column describes whether matches were sought for 900 locations selected completely at random (indicated by *rand*), or for 900 randomly selected locations on image edges (indicated by *rand-edg*).
5.The Patch Size column describes the sizes in pixels of image subsets that were compared by SAD when bidirectional matches were computed.
6.The Mean Control Group Size column describes the average (across 100 experiment repetitions) number of the 900 randomly selected locations that met the inclusion criteria required for admission to the control group.
7.The Mean Control Group Success Rate column, describes the average (across 100 experiment repetitions) percentage of the bidirectional matches in the control group that fulfilled the success criterion.
8.The Control Group Success Rate Stdev column describes the standard deviation of the control group success rates pertaining to 100 repetitions of an experiment.
9.The Uniqueness Boundary Between Test Group And Control Group column, describes what threshold was used to split the set of all bidirectional matches into a test group and a control group. For example, a value of $2.0\sigma_U+\mu_U$ would mean that the control group was comprised of all bidirectional matches for which uniqueness exceeded mean uniqueness across all bidirectional matches, by more than two standard deviations, and the control group is comprised of all bidirectional matches for which it did not.
10.The Mean Test Group Size column describes the average (across 100 experiment repetitions) number of the 900 randomly selected locations that met the inclusion criteria required for admission into the test group.
11.The Mean Test Group Success Rate column describes the average (across 100 experiment repetitions) percentage of bidirectional matches in the test group that fulfilled the success criterion.
12.The Test Group Success Rate Stdev column describes the standard deviation of the test group success rates pertaining to 100 repetitions of an experiment.

4. CONCLUSION

At some point during the course of our past research, we needed to choose a technique for transforming each function (that described dissimilarity of an object for which a match was being sought, to each object in the space being searched), into a single floating point number that summarized the extent to which the function contained a single well defined global minimum that was very different from all of it's other local minima. We decided to call this number the uniqueness of the global minimum, and planned to use it to compare several dissimilarity of appearance functions whenever we needed to decide for which one it was least likely that the global minimum would fail to be a correct match. The Equivalence focused work of Quiroz[10] inspired us to attempt to formulate a definition of this uniqueness assessment as an estimate of the probability of making an error, which in turn led to the uniqueness quantification measure described by this paper.

The relationship between Equivalence and Uniqueness is obvious: An object is unique if it is equivalent to no other objects, or if we decide to allow the transition from "unique" to "not unique" to vary continuously, it is "somewhat unique", if it is "not very equivalent" to "most" of the competing objects of interest to which it is being compared using a metric of interest.

The empirical evidence we presented in this paper, suggests uniqueness quantification can be useful, however the measure we constructed contains so many implicit assumptions and simplifications that it no longer seems reasonable to regard the values it computes as even estimates of probabilities. There seem however to exist countless precedents for the kind of computation it carries out in the Fuzzy Logic community, and so we decided to consider interpreting the values it computes, not as estimates of probabilities but rather as degrees of membership in a fuzzy set of entities which for a particular purpose are deemed unique.

This led us to consider differences and similarities between the concept of probability and the Fuzzy Logic concept of degree of membership. Is Fuzzy Logic a branch of Logic, a branch of Mathematics, or merely alternative nomenclature for the concepts of Probability? Compelling arguments disputing the assertion that Fuzzy Logic is mere "Probability In Disguise", were presented by Zadeh[14] and Kosko[7] . What follows is yet another proposal to reconcile these viewpoints, which can be regarded either as our understanding of Zadeh and Kosko's proposals, or as a small contribution that we are able to make to the Fuzzy Logic community because of Zadeh and Kosko's work. In either case, we present it merely as a point of view that at the present time is satisfactory for our needs and which seems to have a lot in common with many other points of view pertaining to this topic.

Whenever we need to quickly explain what Fuzzy Logic is and why anyone should care about it, to a person whose time and interest is limited, we say that the answer is still subject to debate, but that we are satisfied that a defensible point of view is to regard it as the set of all consequences of modifying the axioms of the Set Theory that underlies all Mathematics, to replace the concept of binary set membership with the concept of continuous membership. In other words, to transform a theory which only applies to objects that either completely belong or completely do not belong to collections of objects called sets, into a theory that allows degree of belonging to vary continuously between 0% and 100%. One consequence of doing this is that the concept of binary truth, (which requires that if the truth of a statement can be assessed, that statement must be deemed to be either true or false rather than a little of both), is also replaced by a continuous alternative. Proposing such a drastic change to the axioms of Set Theory, is motivated by the observation that for any two real world objects, like for example a chair and a table, it is possible to envision arbitrarily many intermediate objects and it is unreasonable to insist on claiming that one object is a chair, but another object that is almost identical to it is not, merely because they lie on opposite sides of an arbitrarily chosen threshold.

Extension of binary logic to continuous logic requires non trivial changes to some very well entrenched ideas and it is not surprising, that such proposals generated controversy among people who care about these things. Still, although creating a new branch of Mathematics by changing an axiom, is less trivial than creating a new game by modifying the rules of an old one, it is not unprecedented. Non-Euclidean Geometry is the outcome of one such effort. Throughout history, Mathematics has grown, as new problems in science and technology have emerged.

In non mathematical contexts, the term "fuzzy" has among other things been used to describe confused thinking. Were the pioneers of Fuzzy Logic, deliberately attempting to antagonize their colleagues by incorporating that word into the name they gave to their field of study? Why not call it Continuous Logic, or choose some similarly informative name which fewer people might be prone to associate with a lack of rigor ? There are many precedents in the history of Mathematics, for modifying a useful binary or finite concept to create a more widely applicable continuous one. Would different terminology have made Fuzzy Logic less controversial ?

A strong case in favor of the assertion that more than mere choice of terminology is responsible for the controversy, can be made on the grounds that assessing degree of membership of an object in one or more sets, is equivalent to quantifying its usefulness for one or more purposes, and that doing the latter, regardless of what it is called, seems feasible without requiring modifications to Set Theory. For example, percentages can be used to transform partially true statements that would be allowed by Continuous Logic, into more convoluted statements that are either 100% true or 100% false. The assertion that a statement like "the tomato I am now eating is a baseball" is 1% true and which also communicates that the degree of membership of a specific tomato in the set of all baseballs is 1%, can be replaced by use of less jarring but more convoluted statements that are 100% true, and hence for which there is no need to explicitly communicate degrees of truth, like for example, "the tomato I am now eating will suffice in more than 0% and less than 1% of situations that require a baseball" or, "the probability that the tomato I am now eating will suffice for completion of a randomly chosen task that requires a baseball is greater than 0% and less than 1%" or, " more than 0% but less than 1% of the properties of the tomato I am now eating are indistinguishable from those of the average baseball".

Arguments in favor of the assertion that Fuzzy Logic, is merely probability in disguise, could point out that percentages describing degrees of membership, and/or usefulness for a purpose can be described in terms of probabilities, as for example in the second of the three convoluted tomato statements above. A counter argument presented by Kosko, is that not all percentages need to be interpreted as quantifications of uncertainty, as for example in the first and third of those same convoluted tomato related statements. Either way, there is more to Fuzzy Logic than mere continuity of set membership. Zadeh wrote "...for humans it is generally much easier to estimate grades of membership or degrees of possibility rather than probabilities..." We interpreted that as a concise articulation of the hypothesis that humans are more adept at estimating object properties that can be perceived without accumulation of evidence over time (like for example weight) than at estimating frequencies of occurrence of events that stem from those properties (like for example probability of a heart attack), more adept at classifying objects into broad categories defined by proximity to class exemplars which differ qualitatively from one another and pertaining to whose precise definition there is a lack of agreement across people (for example categories denoted by natural language terms like thin and obese) than at producing quantitative estimates of object properties (like weight in micrograms), and more adept at deciding which of two events with very different frequencies of occurrence is more likely than at estimating what those frequencies of occurrence are. Kosko's references to the success of Fuzzy Systems built using such assumptions can be construed as evidence that the hypothesis is correct.

From this perspective, Fuzzy Logic could be regarded not as Probability in disguise, but rather, as equivalent to Probability without precise universally standardized quantification of likelihood.

The term “Fuzzy” is not a bad choice to describe this concept since unlike for the case of computations involving experimentally measurable frequencies of occurrence, it is not necessary for different groups of people that are assessing degrees of membership pertaining to the same types of objects in the same types of sets, to agree on what those degrees of membership are in order to build useful systems. They merely need to ensure that the internal representations used by any one system are consistent, in other words to ensure that orderings of relative likelihoods, and relative similarities are approximately correct. Success of the “imprecise, not universally standardized” approach to quantification of uncertainty seems prevalent not only among artificial Fuzzy Logic Controllers but also in biological systems, like for example human brains, which are effective decision makers in spite of the fact that the concepts and natural languages on which they rely are to a great extent neither precisely defined nor universally standardized. It is neither an abuse of language nor an impediment to good decision making when two people in different environments, have different opinions regarding what it means to be a “tall person”. Since behavioral decisions computed by devices which use these computational techniques are useful, and can be obtained using less computation than approaches which attempt to quantify uncertainty using experimentally measurable frequencies of occurrence, it is arguably more accurate to regard decision making based on Fuzzy Logic , not as imprecise, but rather, as devoid of unnecessary complexity. While it may very well be feasible to also express these ideas using the language of probability and convergence of approximations, there is a point beyond which the amount of work needed to transform one representation into another is great enough to justify the existence of both, each as Zadeh put it, for use in the context to which it is best suited. Chinese is not usually regarded as Spanish in disguise merely because one can be translated into the other, nor do Mathematicians insinuate that Fourier Transforms are useless because they can after all, be regarded as mere time varying signals in disguise.

Are all “imprecise not universally standardized” quantifications of uncertainty equally useful ? We suspect the answer is no, and that there exist opportunities to classify them and study their properties. It would not be the first time that the systematic elimination of properties of concrete ideas yielded opportunities to theoretically analyze concepts that are more abstract. Consider for example, the emergence of Abstract Algebra in general, and Group Theory in particular.

The combination of the above described point of view and our experimental data, suggests that it would be neither unreasonable nor unprecedented to interpret the types of values that are computed by our uniqueness quantification algorithm as degrees of membership in a fuzzy set. In our particular case, that fuzzy set is comprised of objects which for a particular purpose can be regarded as unique. The associated degrees of membership are values which were obtained by estimating some but not all aspects of probabilities and arguably they are not estimates of experimentally measurable frequencies of occurrence. The relationship between values computed by our uniqueness quantification measure and actual probabilities seems similar to the relationship between estimates of stereoscopic disparities and actual distances to objects. In order for our uniqueness quantification algorithm to be useful, it does not need to compute actual probabilities. It merely needs to compute values that can be used to sort problems of interest into roughly the same order as knowledge of actual probabilities that they will be solved incorrectly, would allow us to do. A less rough approximation of orderings would of course be preferable if it can be obtained, in other words, this point of view does not preclude us from expecting that the utility of our uniqueness quantification algorithm as a predictor of incorrect matches would be further enhanced by introduction of any techniques (for example pertaining to more accurate estimation of the variability of the chosen dissimilarity metric as a function of changing environment in which measurements are made) which cause the values it computes to be more comparable to experimentally measurable frequencies of occurrence, provided (and here is the hard part) computational tractability is not sacrificed.

Many stereoscopic vision publications contain images of disparity maps. This one contains none. That is because, this paper describes a stereoscopic vision algorithm component, rather than a complete stereopsis algorithm. We presented a preliminary estimate of it’s potential usefulness, that we obtained by integrating it into a very simple environment. In future work we plan to explore it’s integration into more complex mechanisms. For example, although many small image subsets fail to be unique, we would expect that a sufficiently large union of such subsets could turn out to be highly unique, and we are interested in exploring their properties. This approach suggests a more abstract feature extraction technique in which the features of interest are not predefined geometric shapes like points lines or circles, or for that matter edges and corners, but rather, arbitrarily shaped, not necessarily connected, image subsets that happen to be unique in the stereogram being processed. One could argue that edges and corners have historically been a focus of attention precisely because they tended to be more unique than other features available at the time, but that it was their uniqueness that made them interesting and that this uniqueness is a more fundamental concept, or rather a concept more essential to stereoscopic matching, than the concept of being an edge or the concept of being a corner. Consider for example a stereogram which depicts a scene composed entirely of randomly distributed dots, for example a stereogram depicting an arctic ice storm or a random dot stereogram. Using edges and corners as features of interest in that context might not turn out to be as useful as using collections of image patches that happened to be unique.

The computational approach used by the experiments presented in this paper, might be useful to builders of systems that desire a scalable approach to stereoscopic vision, in which systems benefit from having the means to adjust their computational expenditures, for example, from computing only a few matches at low resolution to computing dense high resolution disparity maps in regions of interest, or anywhere in between. Like cooperative[9,15] algorithms, this computational approach does make use of a group effort, however that effort is competitive rather than cooperative, in the sense that mechanisms which attempt to solve individual single point correspondence problems do not help each other. This computation differs from a collection of independent individual efforts only in that when it is unlikely that all members of the group will perform equally well, the existence of the group provides the final stage of the computation with a pool from which to choose best performers. Most parts of this type of competitive computation are highly parallelizable, with little need for competing processes to communicate with each other. A disadvantage of this approach is that there is little opportunity to coordinate efforts of participants, derive value from a conflict resolution driven process of elimination or ensure that competing problem solvers do not contradict each other. For this reason we present this approach not as an alternative to cooperative computation but rather as a technique system designers can integrate with cooperative approaches to an extent determined by the magnitude of the computational expenditures they are able to make and the region in the continuum from sparse to dense disparity map computation that their systems must operate in at any given time.

In this paper we described one type of useful incorrect match, namely an incorrect match that is compatible with the Epipolar Constraint. In future work, we are interested in further exploring the concept of "useful incorrect matches" along with algorithms that are able to find them without making large computational expenditures. We suspect for example, that many incorrect matches that lie on edges that happen to represent depth discontinuities are benign, in the sense that their disparity is equal to the correct disparity of an adjacent pixel that is part of a large connected region of pixels in which disparity varies smoothly. The impact of this type of incorrectness is merely a slightly inaccurate computation of where depth discontinuities lie. We suspect that the performance of many systems which make use of stereoscopic vision would not be adversely impacted by such inaccuracy, especially since, if they require great accuracy in a particular region of interest, they can always zoom in their cameras. In summary we expect that many system designers could benefit from algorithms which require substantially less computational expenditure than alternatives that make few errors, provided the vast majority of the errors made by these undemanding algorithms are benign when the algorithms are used for their intended purposes.

REFERENCES

[1] Faugeras, O., Luong Q. T. , Maybank S., "Camera self-calibration: theory and experiments", Proceedings of European Conference on Computer Vision, (1992).

[2] Faugeras, O., Luong, Q.T., [The Geometry of Multiple Images, The Laws That Govern the Formation of Multiple Images of a Scene and Some of Their Applications], MIT Press, Cambridge, Massachusetts, (2001).

[3] Grimson, W.E.L., "Computational experiments with a feature based stereo algorithm", IEEE TPAMI 7(1), 17-34 (1985).

[4] Hirschmuller, H. ,Scharstein, D. "Evaluation of cost functions for stereo matching", IEEE CVPR , 195-202 (2003).

[5] Hirschmuller, H. ,Scharstein, D. "Evaluation of stereo matching costs on images with radiometric differences", IEEE TPAMI 31(9), 1582-1599 (2009).

[6] Kanade, T., Okutomi, M., "A stereo matching algorithm with an adaptive window: theory and experiment", IEEE TPAMI 16(9), 920- 932 (1994).

[7] Kosko, B., "Fuzziness vs. probability" ,Int. J. General Systems 17, 211-240 (1990).

[8] Luong, Q. T., Faugeras O. D., "The fundamental matrix: theory, algorithms, and stability analysis", IJCV 17(1), 43-75 (1996).

[9] Marr, D., Poggio, T., "Cooperative computation of stereo disparity.",Science 194, 283-287 (1976).

[10] Quiroz, J., "A method for determining equivalence in industrial settings: defining and testing the equivalence of two methods or two laboratories", Quality Engineering 18, 5-14 (2006).

[11] Scharstein, D., Pal, C.,"Learning conditional random fields for stereo", IEEE CVPR Minneapolis MN, (2007).

[12] Scharstein, D., Szeliski, R.,"A taxonomy and evaluation of dense two-frame stereo correspondence algorithms", IJCV 47, 7-42 (2002).

[13] Scharstein, D., Szeliski, R.,"High-accuracy stereo depth maps using structured light", IEEE CVPR Madison WI 1,195-202 (2003).

[14] Zadeh, A. L.,"Probability theory and fuzzy logic are complementary rather than competitive", Technometrics 37(3), 271-276 (1995).

[15] Zitnick, L., Kanade, T., "A cooperative algorithm for stereo matching and occlusion detection", IEEE TPAMI 22(7), 675 - 684 (2000).