A comparison of some predictors of stereoscopic match correctness

Val Petran*[ab], 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

Previously we introduced the concept of continuous quantification of uniqueness, as a general purpose technique designed to be applicable to any situation in which there is a need to decide which of several equally effective objects to choose for a task, that requires recognition of the chosen object, in a variety of contexts, by comparing attributes which contain a non trivial amount of context dependent variability. We defined that uniqueness assessment as an algorithm that computes a fuzzy set membership function that measures some but not all aspects of the probability that the sought after object will not be confused with other objects in the space being searched. We evaluated the usefulness of that concept by experimentally assessing the extent to which the uniqueness of the SAD global minimum of locally computed image subset dissimilarity was both a predictor of bidirectional match compliance with the Epipolar Constraint, and a predictor of bidirectional match disparity correctness, for the classical stereoscopic correspondence problem of computer vision, and in that context found the uniqueness of the aforementioned global minimum to be a useful but imperfect predictor of success. In this paper we compare the usefulness of the uniqueness of the aforementioned global minimum to that of, the magnitude of that same global minimum, the magnitude of variability across contributors to that global minimum, uniqueness of that variability, and co-occurrence of the global minimum of local image subset dissimilarity and global minimum of variability across contributors to local image subset dissimilarity.

Keywords: Uniqueness, Distinctiveness, Stereoscopic Vision, Computer Vision, Stereoscopic Correspondence, Subjective Probability, Epistemological Probability, Objective Probability, Bayesian Viewpoint, Frequentist Viewpoint

1. BACKGROUND INFORMATION

The classical stereoscopic correspondence problem of computer vision has been described extensively [1,2,3], as have numerous techniques [1] for solving it, along with explanations of why the problem is difficult, why there is good reason to believe it can almost always be solved in spite of that difficulty, and why algorithms which attempt to solve it may not be able to avoid at least implicitly resorting to probabilistic techniques that are equivalent to judicious guessing. The work [3] of David Marr, Tomaso Poggio and Bela Julesz, with random dot stereograms comprised of images that are completely devoid of conceptually meaningful objects, opportunities to infer shape from shading, depth from focus, or depth from perspective cues, suggested that solving this problem suffices for the type of depth perception that allows the human binocular visual system to perceive high quality photographs of three dimensional objects as incontrovertibly flat surfaces in spite the wealth of depth cues that such photographs contain.

The research we describe in this paper is focused on empirically exploring some predictors that can help us identify as many accurate stereoscopic matches as possible, without having to resort to computationally demanding global optimization techniques. In other words, the goal is to identify and explore that subset of the stereoscopic correspondence problem, which can be solved by algorithms that lie closer to the greedy, in other words locally optimal decision making end of the spectrum.

One can envision a non trivial number of applications, for techniques that can reliably and quickly identify even relatively small numbers of highly accurate stereoscopic matches. These include ensuring that fovea centers of independently rotating left and right verging cameras that comprise a stereo rig form a correspondence pair at all times, recomputing the epipolar constraint [4,5,6] pertaining to such a stereo rig after each camera rotation, seeding greatest confidence first algorithms that search for matches in disparity space [1] neighborhoods of known correct matches, and the application of stereoscopic vision to image reconstruction, for example to stitch together subsets of several images of the same scene that have each been damaged in a different way.

Visual Information Processing XXI, edited by Mark Allen Neifeld, Amit Ashok, Proc. of SPIE
Vol. 8399, 83990S · © 2012 SPIE · CCC code: 0277-786X/12/$18 · doi: 10.1117/12.918893

Previously [7] we hypothesized that magnitude of the global minimum of image subset dissimilarity, measures of the unusualness of that global minimum, and variability across contributors to image subset dissimilarities would turn out to be complementary predictors of whether or not stereoscopic matches comprised of such global minima comply with ground truth disparity. Shortly thereafter Yoon and Kweon demonstrated that a combination [8] of magnitude of image subset dissimilarity and a measure of image subset unusualness called Distinctiveness [9] , that was introduced by Manduchi and Tomasi, could be used to obtain high rates of ground truth disparity compliance. More recently Mordohai developed a technique [10] based on computing correlations of goodness of match curves that allowed him to work around some drawbacks of Distinctiveness, including the fact that real world stereograms contain many points that are not particularly Distinctive and also, that Distinctiveness assessment requires deciding what precisely constitutes a local minimum that is sufficiently well defined, to be eligible for designation as the local minimum most similar to the global minimum of image subset dissimilarity. The latter is an integral part of Distinctiveness assessment. Last year, we introduced an alternative object unusualness measure that we called Continuous Quantification of Uniqueness [11] which differs qualitatively from Distinctiveness in several respects, including that like Mordohai's SAMM measure, it avoids the prerequisite of quantifying the extent to which local minima are well defined, and also in that inspired by the work of Quiroz [12] it takes into account the variability of image subset appearance as a function of changing point of view. We suspect that Distinctiveness and Continuous Quantification of Uniqueness are complementary, and in future work are interested in empirically comparing their performances on different types of tasks. For the sake of brevity, we often use the term Uniqueness rather than Continuous Quantification of Uniqueness, throughout much of this and other work.

2. HYPOTHESES

In this paper we experimentally compare the effectiveness of the uniqueness [11] of the global minimum of SAD (Sum of Absolute Value of Differences) image subset dissimilarity (UD) to that of the magnitude of that global minimum (MD) , the magnitude of variability across contributors to that global minimum (MV), the uniqueness of that variability (UV) and co-occurrence of the global minimum of local image subset dissimilarity and global minimum of variability across contributors to local image subset dissimilarity, as predictors of agreement of bidirectional stereoscopic matches comprised of global minima of SAD image subset dissimilarity, with ground truth disparity data.

The data we subsequently present and examine was obtained by performing the same types of experiments we described in detail in our most recent publication [11] . Specifically, we carried out 100 random experiments using each of the same Middlebury Laundry, Art and Moebius rectified stereograms [13,14,15,16] that we used in our previous publication [11] for which Scharstein, Szeliski, Hirschmuller, Pal et all have produced and disseminated highly accurate ground truth disparity data. Each experiment was comprised of choosing 900 image locations at random, attempting to find SAD bidirectional matches for each of them, and computing the aforementioned predictors for each of those bidirectional matches . On average it was possible to find bidirectional matches for roughly one third of the randomly chosen locations.

We explored the resulting data in order to determine the extent to which extreme values of the aforementioned hypothesized predictors and/or combinations of them, identified bidirectional matches that agreed with published [13,14,15,16] ground truth disparity information. All data presented in this paper was obtained by using SAD to compare rectangular 17 by 17 pixel subsets of color images. This choice was motivated by our previous work [11] and other results that we have not yet published, which suggest that the aforementioned predictors are most effective when the image subsets being compared are relatively large. During the course of previous work, we concluded that the Laundry, Art, and Moebius stereograms are of high, medium and low difficulty respectively, in the sense that percentages of the set of all bidirectional matches which comply with ground truth disparity pertaining to those stereograms are low, medium and high respectively. This is not surprising, since the Laundry stereogram depicts a scene with repetition, many depth discontinuities, occlusions, and a significant absence of surface texture, the Moebius stereogram, with its many cloth draped surfaces is the exact opposite, and the Art stereogram lies between those two extremes.

While conducting the aforementioned experiments we noticed that the uniqueness of the global minimum of image subset dissimilarity seemed to be a much better predictor of stereoscopic match compliance with ground truth disparity, if search spaces were not strictly confined to corresponding epipolar lines. Slightly thicker search spaces comprised of groups of three adjacent epipolar lines centered on the corresponding epipolar lines that contained the locations for which matches were sought, seemed to give the uniqueness quantification measure an opportunity to distinguish between edges and corners, and also to give the matching processes more opportunities to fail to find a bidirectional match.

This paper explores the hypotheses that such thicker search spaces have a positive impact on the effectiveness of the aforementioned predictors, and that it is both possible and worthwhile to combine some of those predictors with each other.

3. DATA

Each of the plots that are presented in figures 1 through 21 are comprised of dots, that each represent one bidirectional match, that was found for a randomly chosen image location. The vertical coordinate of each dot describes the extent to which the bidirectional match represented by that dot agrees with ground truth disparity data, and it's horizontal coordinate describes the value of a hypothesized predictor of compliance with ground truth disparity (one of UD, MD, UV, MV, whose identity is documented by the horizontal axis label) that was computed for that bidirectional match. Each plot presents an aggregate of results obtained across 100 repetitions of the experiment described in the previous section. Data pertaining to the Laundry stereogram, that is presented in Figure 1 supports the hypothesis that high uniqueness of global minimum of image subset dissimilarity is a good predictor of the extent to which SAD bidirectional matches agree with ground truth disparity. Furthermore it supports the hypothesis that this predictor benefits substantially from use of a thicker search space and very little from restriction of attention to only those bidirectional matches for which the global minimum of SAD image subset dissimilarity, and the global minimum of variability across contributors to each SAD image subset dissimilarity, occur at the same location.

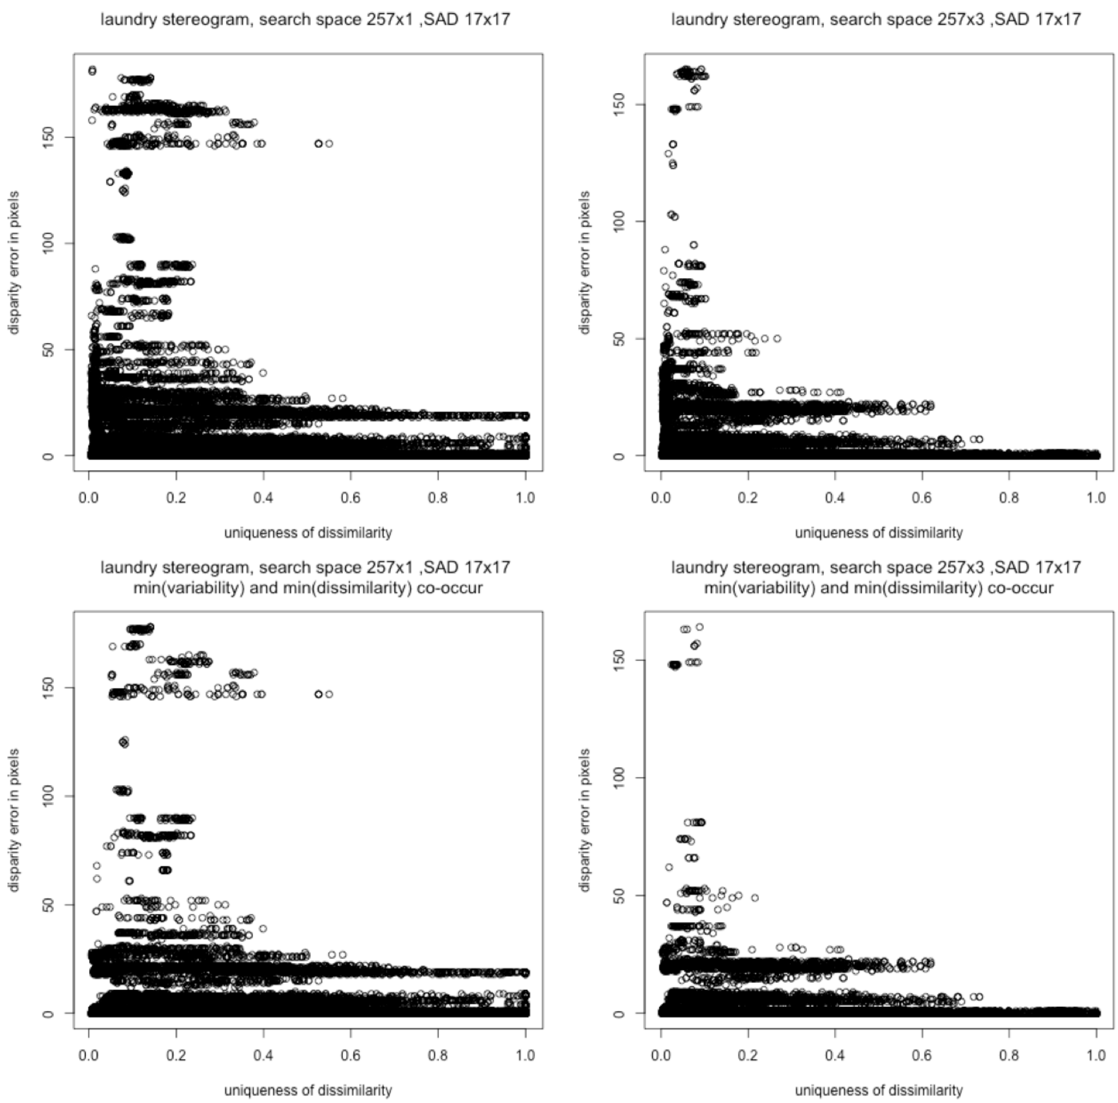

Figure 1. Impact of co-occurrence and search space thickness on UD (uniqueness of dissimilarity) predictor effectiveness

Magnified data pertaining to the Laundry stereogram, that is presented in Figure 2 supports the hypothesis that small magnitude of global minimum of image subset dissimilarity is a predictor of the extent to which SAD bidirectional matches agree with ground truth disparity. Furthermore it supports the hypothesis that this predictor benefits slightly from the use of a thicker search space and substantially from restriction of attention to only those bidirectional matches for which the global minimum of SAD image subset dissimilarity, and the global minimum of variability across contributors to each SAD image subset dissimilarity, occur at the same location.

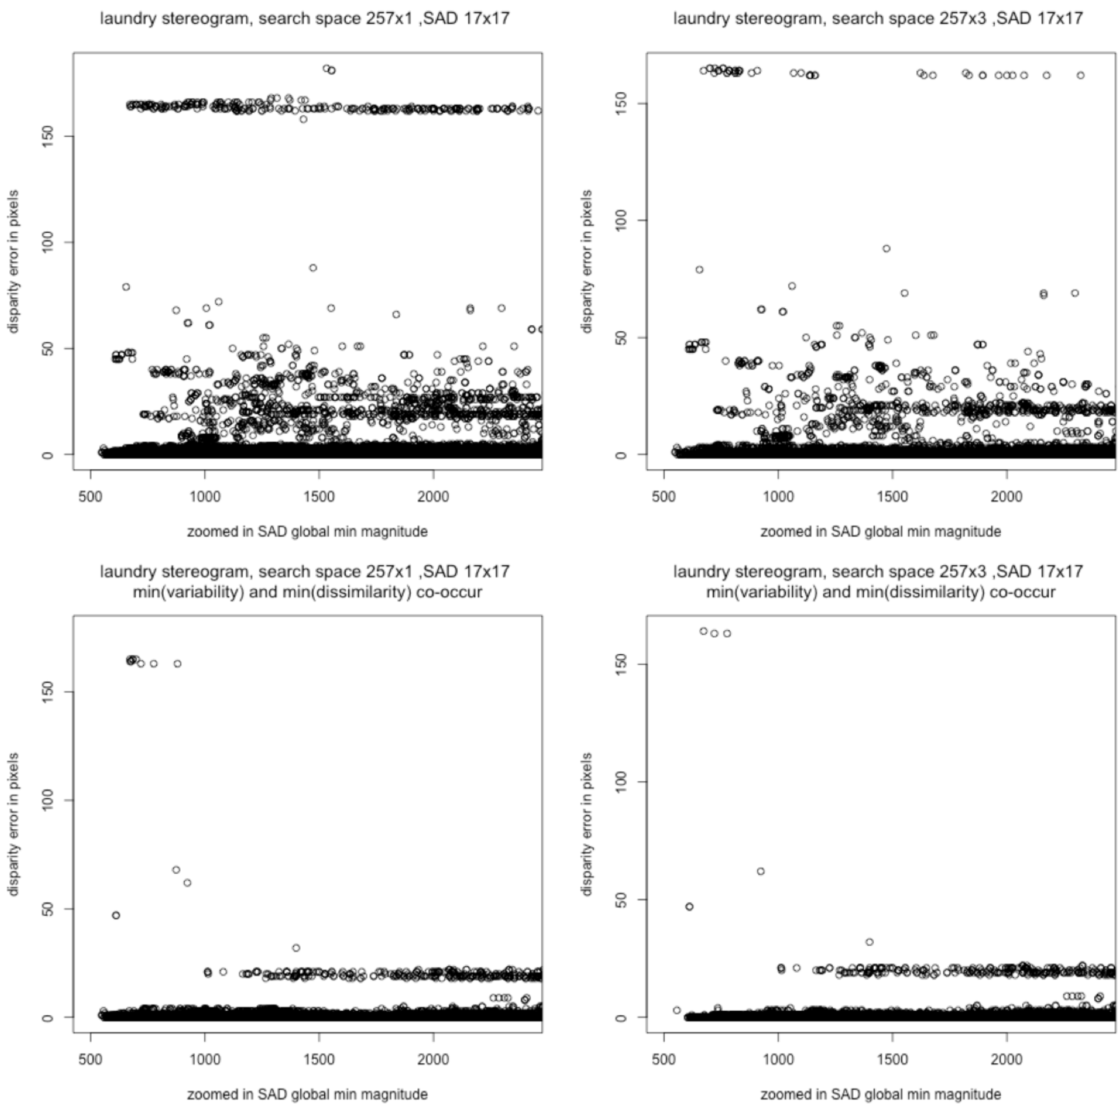

Figure 2. Impact of co-occurrence and search space thickness on MD (magnitude of dissimilarity) predictor effectiveness

The above data is presented in magnified form, because the range (from around 500 to 1000) in which low magnitude of image subset dissimilarity is a predictor of bidirectional match agreement with ground truth disparity for the case of the Laundry Stereogram is only a small fraction of the range (from around 500 to 20,000) of all observed image subset dissimilarities, and is difficult to discern on a plot of the entire data set.

Figure 3 presents an unmagnified version of the MD predictor data that was presented in Figure 2.

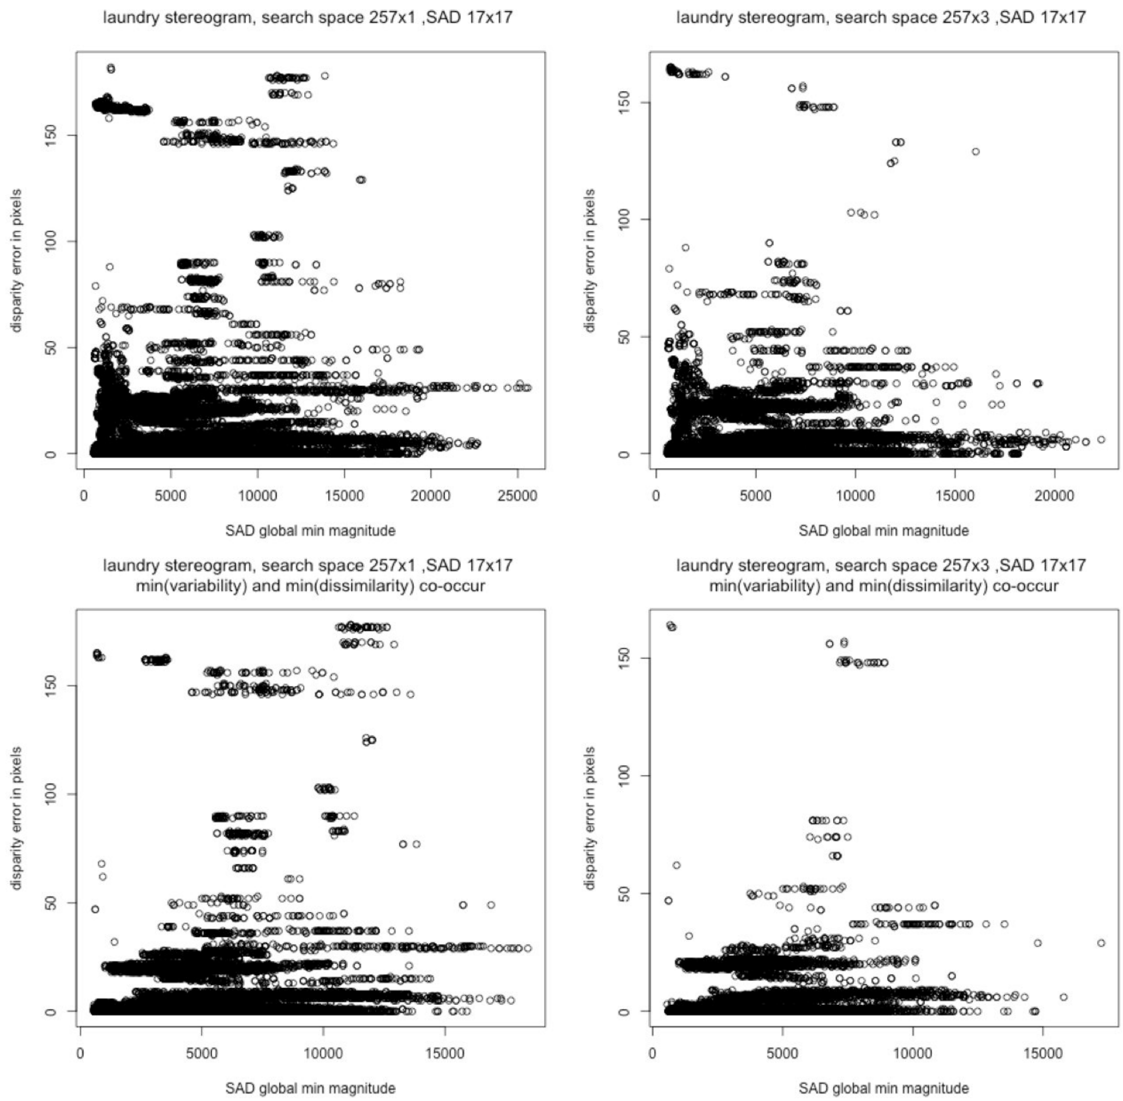

Figure 3. Impact of co-occurrence and search space thickness on MD (magnitude of dissimilarity) predictor effectiveness

Magnified data pertaining to the Laundry stereogram, that is presented in Figure 4 supports the hypothesis that small variability across contributors to the global minimum of image subset dissimilarity is a predictor of the extent to which SAD bidirectional matches agree with ground truth disparity. Furthermore it supports the hypothesis that this predictor benefits slightly from the use of a thicker search space and substantially from restriction of attention to only those bidirectional matches for which the global minimum of SAD image subset dissimilarity, and the global minimum of variability across contributors to each SAD image subset dissimilarity, occur at the same location.

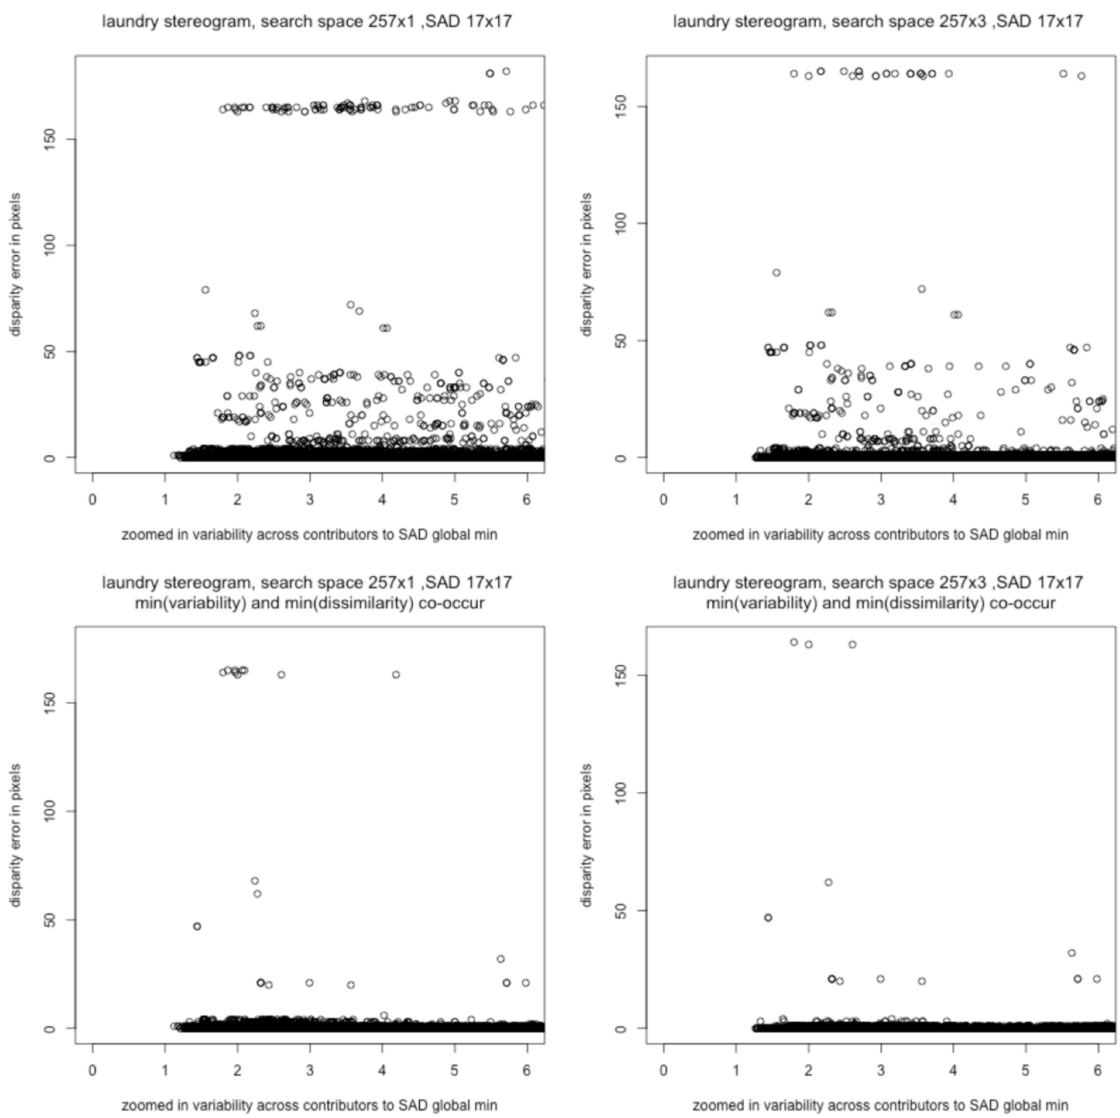

Figure 4. Impact of co-occurrence and search space thickness on MV (magnitude of variability) predictor effectiveness

The above data is presented in magnified form, because the range in which low variability across contributors to image subset dissimilarity is a predictor of bidirectional match agreement with ground truth disparity for the case of the Laundry Stereogram, is only a small fraction of the range of all observed image variabilities, and is difficult to discern on a plot of the entire data set.

Figure 5 presents an unmagnified version of the MV predictor data that was presented in Figure 4.

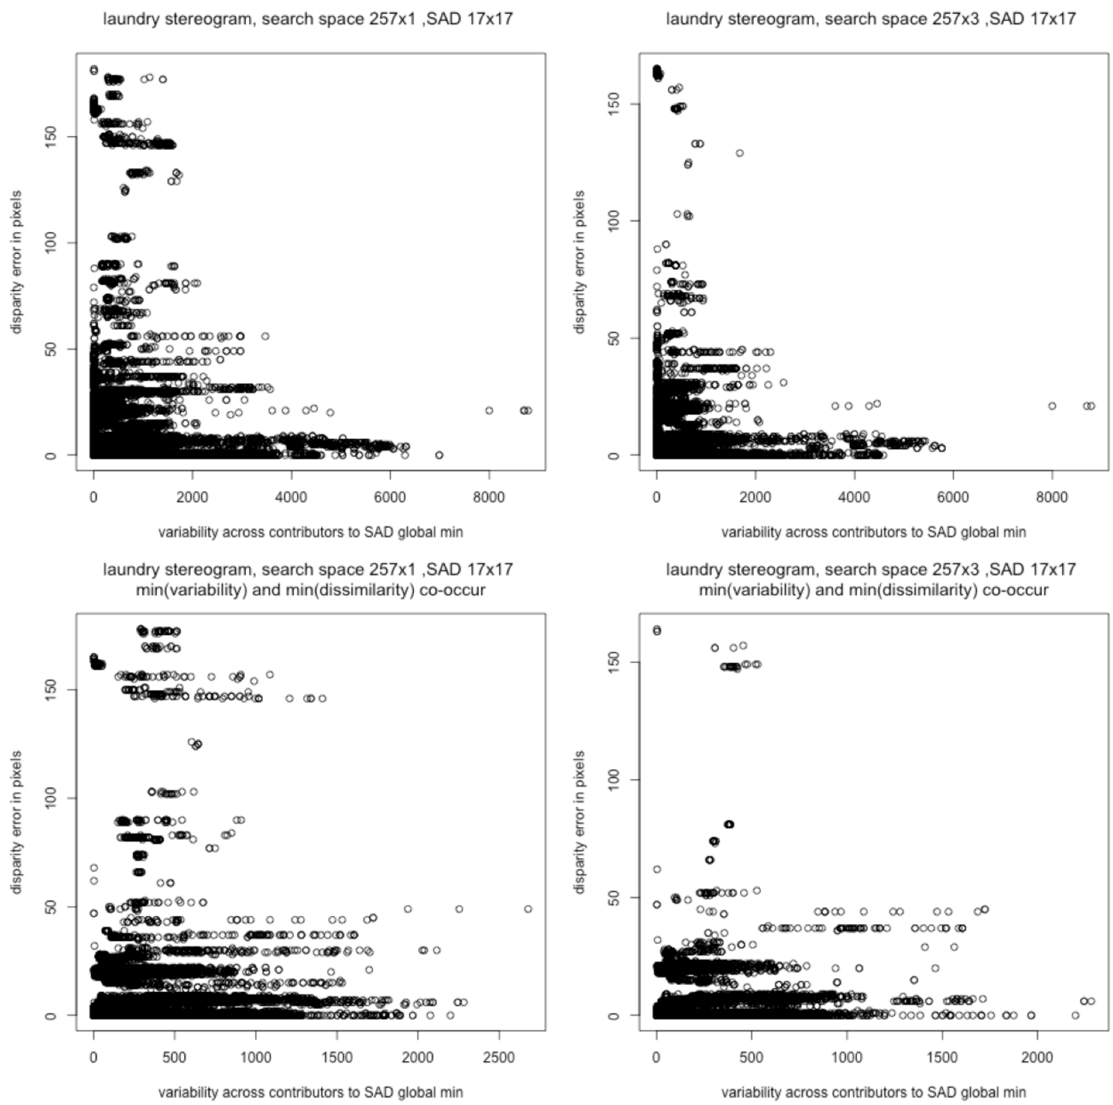

Figure 5. Impact of co-occurrence and search space thickness on MV (magnitude of variability) predictor effectiveness

Magnified data pertaining to the Laundry stereogram, that is presented in Figure 6 supports the hypothesis that in contrast to the Uniqueness of Dissimilarity predictor , low uniqueness of the global minimum of variability across contributions to image subset dissimilarity is a slight predictor of the extent to which SAD bidirectional matches agree with ground truth disparity. Furthermore it supports the hypothesis that this predictor benefits slightly from both use of a thicker search space and from restriction of attention to only those bidirectional matches for which the global minimum of SAD image subset dissimilarity, and the global minimum of variability across contributors to each SAD image subset dissimilarity, occur at the same location.

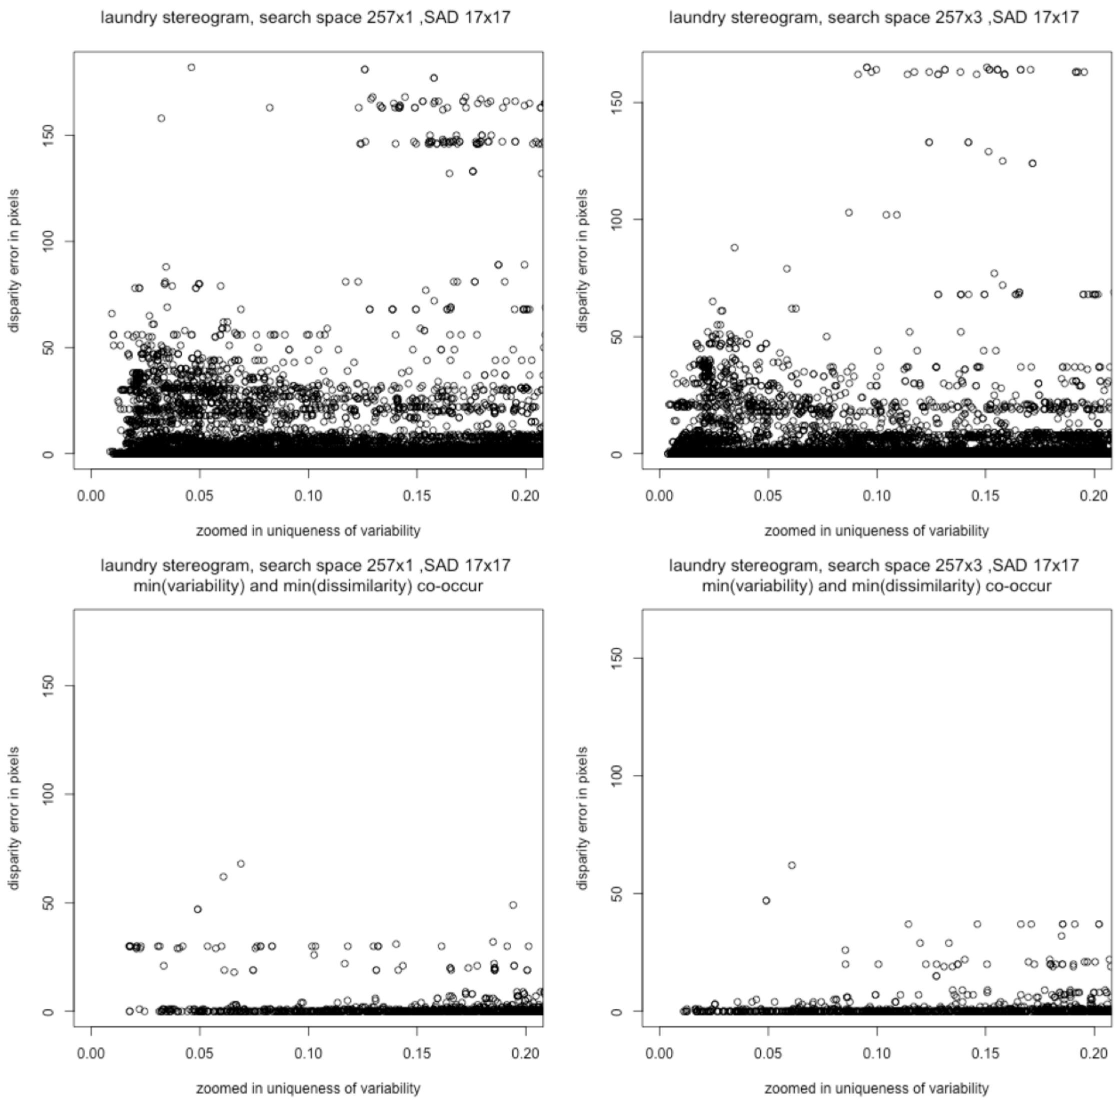

Figure 6. Impact of co-occurrence and search space thickness on UV (uniqueness of variability) predictor effectiveness

The above data is presented in magnified form, because the range in which low uniqueness of variability across contributors to image subset dissimilarity is a predictor of bidirectional match agreement with ground truth disparity for the case of the Laundry Stereogram, is only a small fraction of the range of all observed values and is difficult to discern on a plot of the entire data set, albeit much less so for the set of bidirectional matches for which the global minimum of image subset dissimilarity and the global minimum of variability across contributors to image subset dissimilarity occur at the same location.

Figure 7 presents an unmagnified version of the UV predictor data that was presented in Figure 6.

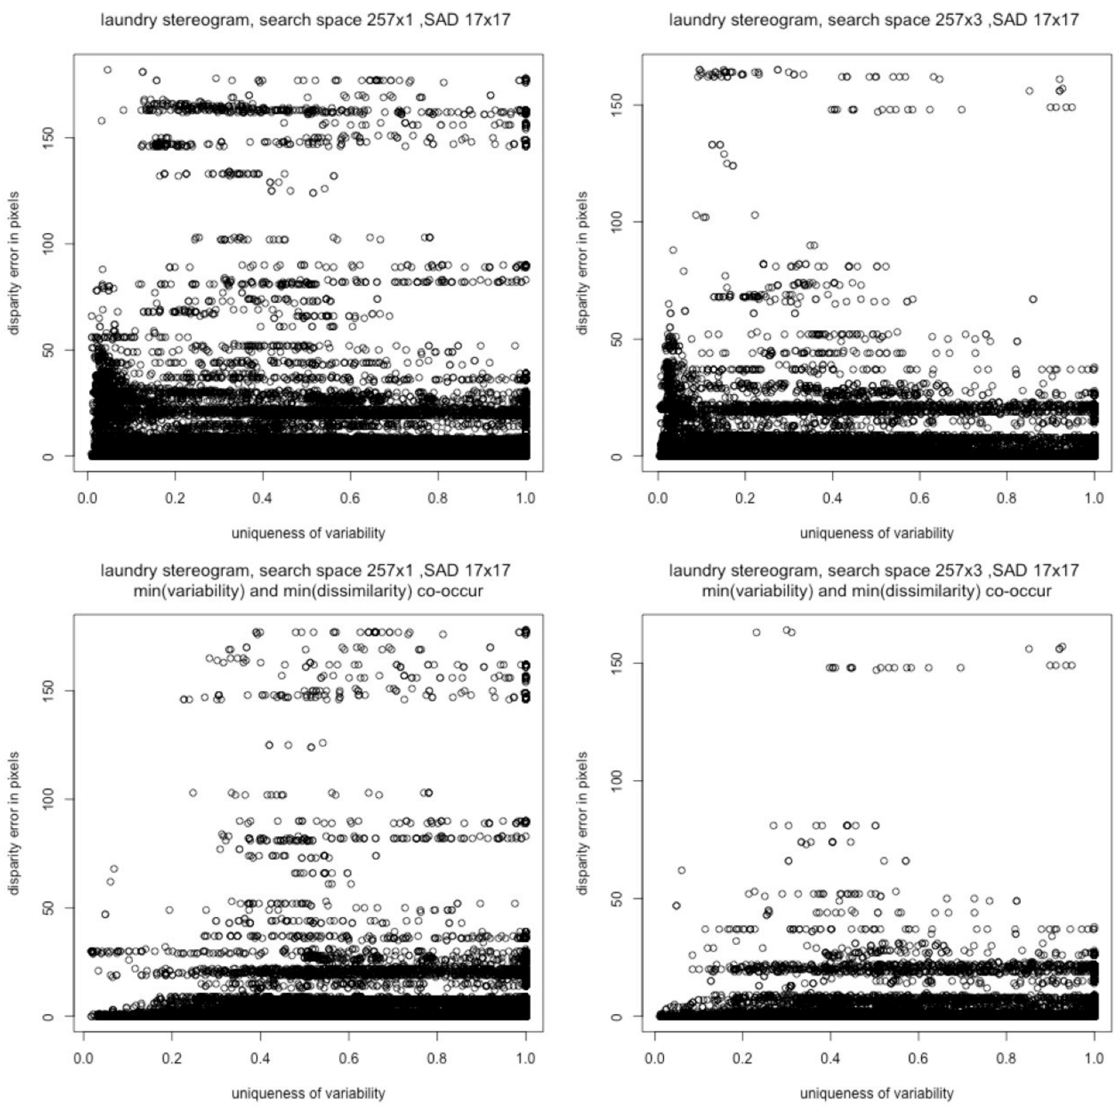

Figure 7. Impact of co-occurrence and search space thickness on UV (uniqueness of variability) predictor effectiveness

The combination of low variability across contributors to image subset dissimilarity and low uniqueness of that variability could be construed as an indication of a region of object surface continuity, in other words, of a portion of a stereogram that depicts a portion of a scene which does not contain many depth discontinuities.

Data pertaining to the Art stereogram, that is presented in Figure 8 supports the hypothesis that high uniqueness of global minimum of image subset dissimilarity is a good predictor of the extent to which SAD bidirectional matches agree with ground truth disparity. Furthermore it supports the hypothesis that this predictor benefits substantially from use of a thicker search space and very little from restriction of attention to only those bidirectional matches for which the global minimum of SAD image subset dissimilarity, and the global minimum of variability across contributors to each SAD image subset dissimilarity, occur at the same location.

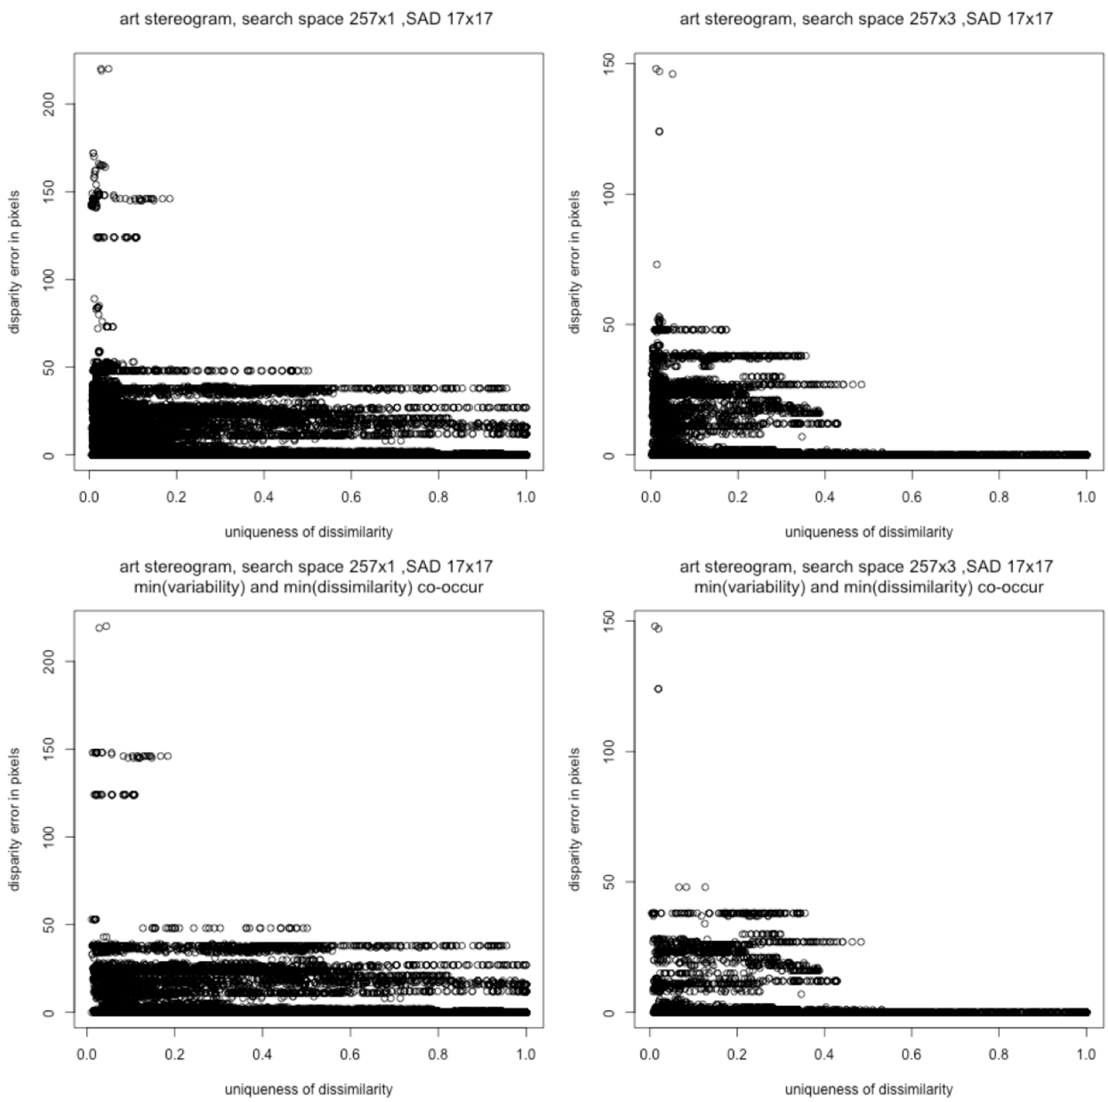

Figure 8. Impact of co-occurrence and search space thickness on UD (uniqueness of dissimilarity) predictor effectiveness

The above data also illustrates a sense in which the Art stereogram is less difficult than the Laundry stereogram. For the case of the Art stereogram and thick search spaces, Uniqueness of global minimum of Dissimilarity values that are merely greater than 0.6 suffice to ensure 100% agreement of bidirectional matches with ground truth disparity. This does not suffice for the case of the Laundry stereogram.

Magnified data pertaining to the Art stereogram, that is presented in Figure 9 supports the hypothesis that small magnitude of global minimum of image subset dissimilarity is a predictor of the extent to which SAD bidirectional matches agree with ground truth disparity. Furthermore it supports the hypothesis that this predictor benefits slightly from the use of a thicker search space and substantially from restriction of attention to only those bidirectional matches for which the global minimum of SAD image subset dissimilarity, and the global minimum of variability across contributors to each SAD image subset dissimilarity, occur at the same location.

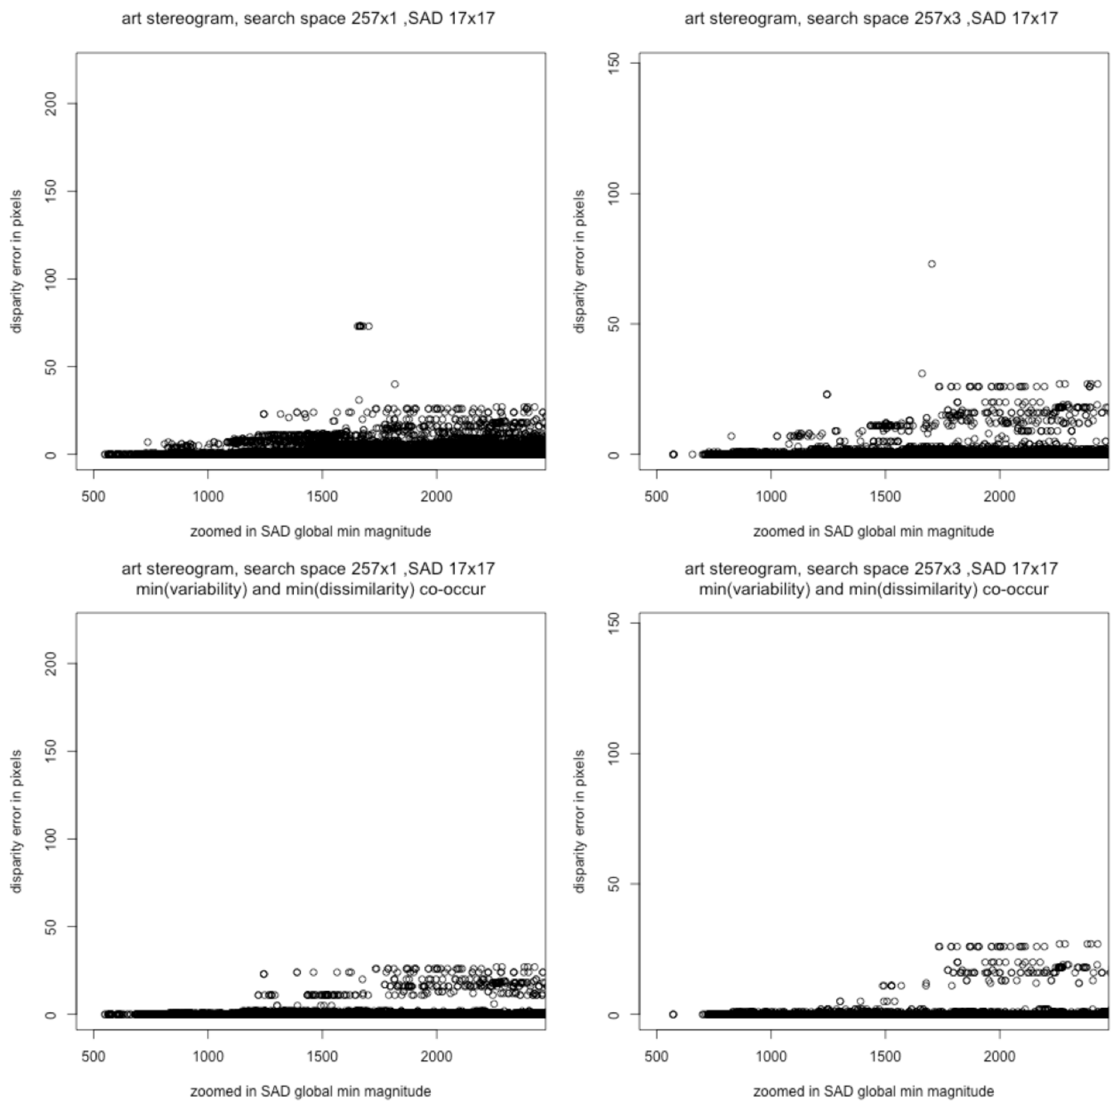

Figure 9. Impact of co-occurrence and search space thickness on MD (magnitude of dissimilarity) predictor effectiveness

Figure 10 presents an unmagnified version of the MD predictor data that was presented in Figure 9.

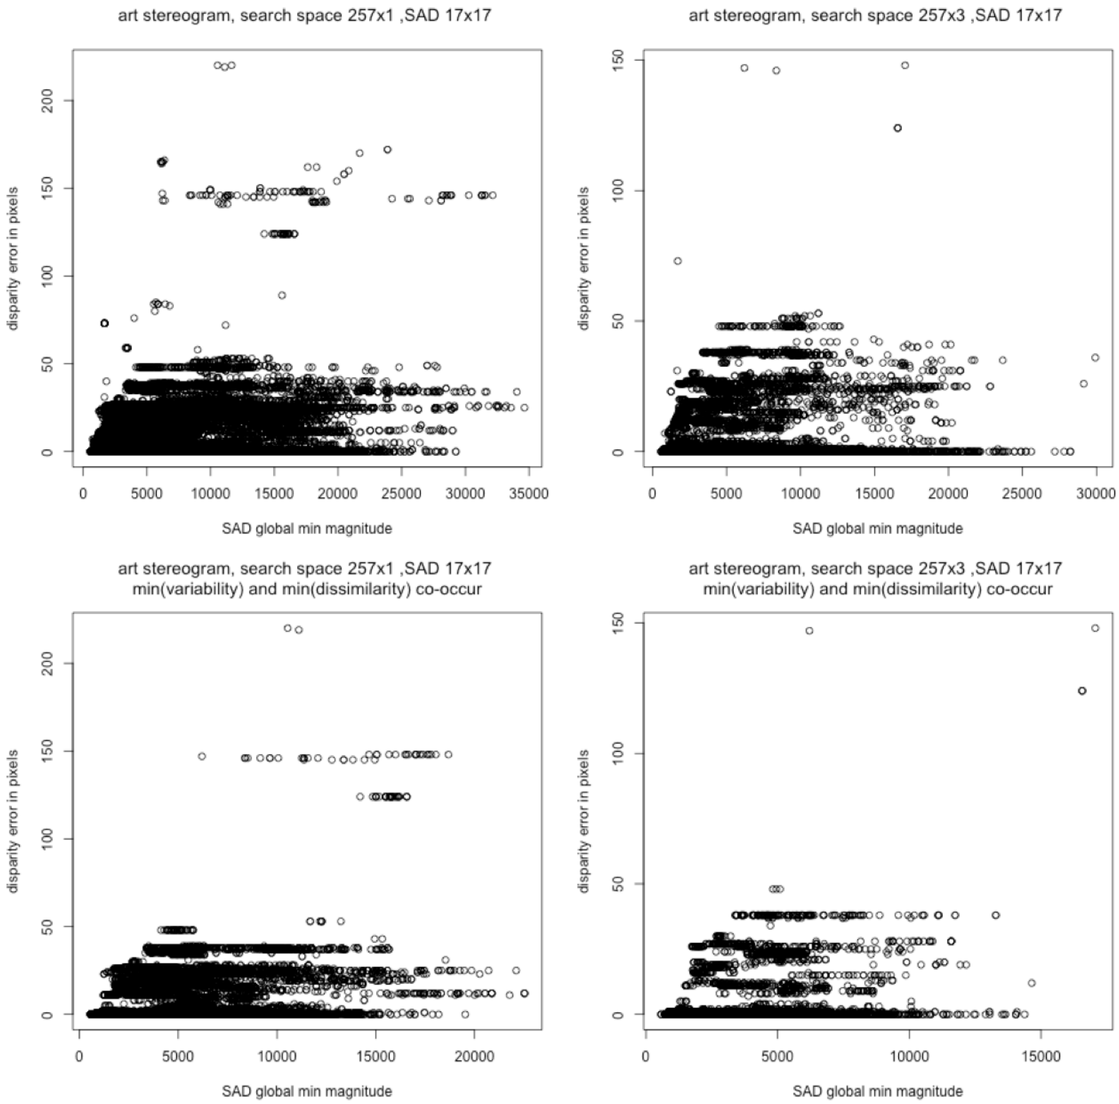

Figure 10. Impact of co-occurrence and search space thickness on MD (magnitude of dissimilarity) predictor effectiveness

The above data illustrates another sense in which the Art stereogram is less difficult than the Laundry stereogram. For the case of the Art stereogram it is not necessary to resort to a magnified view of the data in order to be able to clearly discern that low Magnitude of global minimum of image subset Dissimilarity is a predictor of bidirectional match compliance with ground truth disparity data.

Magnified data pertaining to the Art stereogram, that is presented in Figure 11 supports the hypothesis that small variability across contributors to the global minimum of image subset dissimilarity is a predictor of the extent to which SAD bidirectional matches agree with ground truth disparity. Furthermore it supports the hypothesis that this predictor benefits slightly from the use of a thicker search space and substantially from restriction of attention to only those bidirectional matches for which the global minimum of SAD image subset dissimilarity, and the global minimum of variability across contributors to each SAD image subset dissimilarity, occur at the same location.

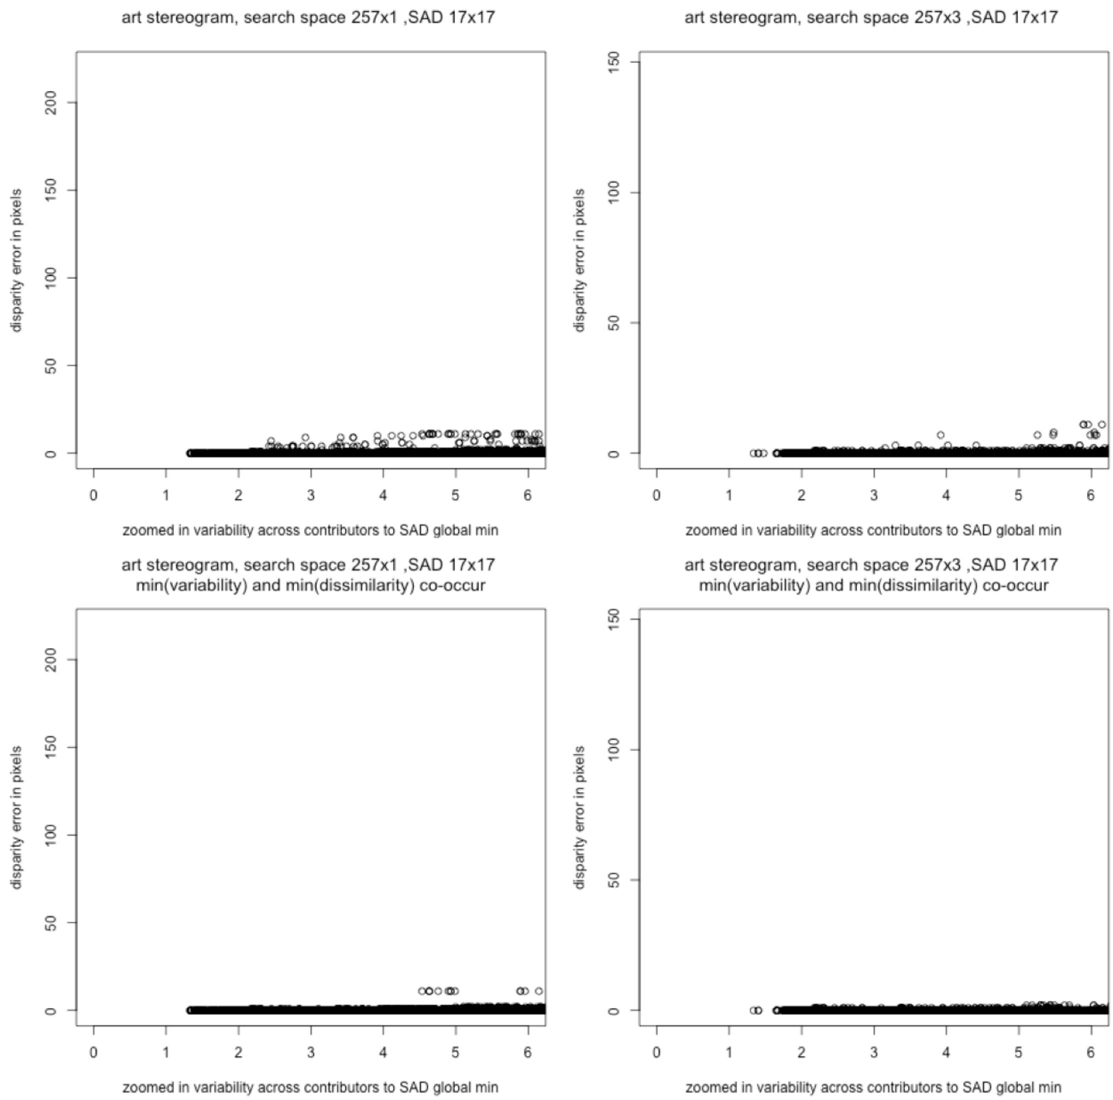

Figure 11. Impact of co-occurrence and search space thickness on MV (magnitude of variability) predictor effectiveness

Figure 12 presents an unmagnified version of the MV predictor data that was presented in Figure 11.

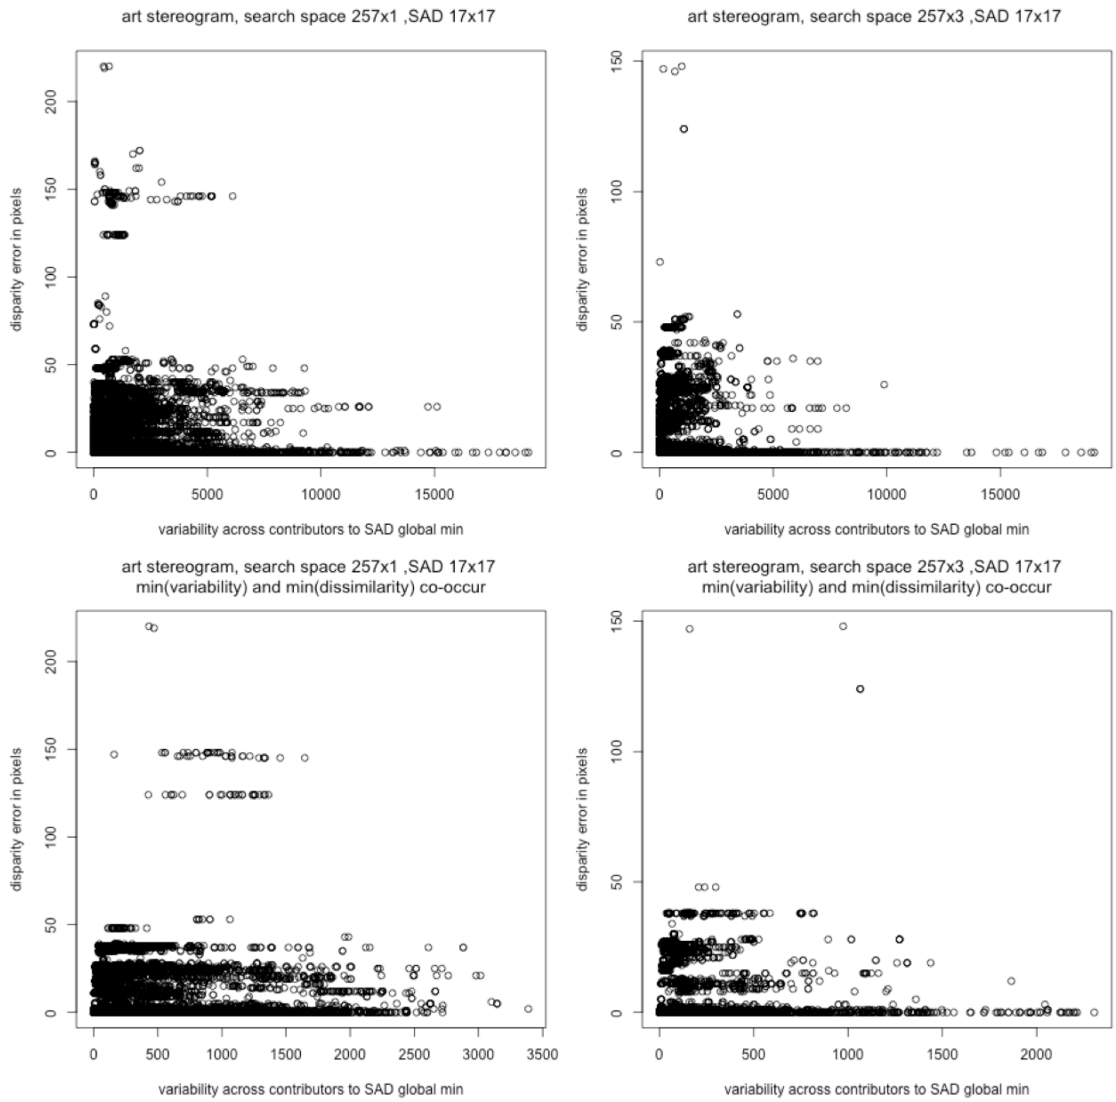

Figure 12. Impact of co-occurrence and search space thickness on MV (magnitude of variability) predictor effectiveness

The above data illustrates a phenomenon that we were somewhat surprised by, namely, that extremely high variability across contributors to the global minimum of image subset dissimilarity also seems to be a predictor of bidirectional match compliance with ground truth disparity. We examined some stereogram locations where such bidirectional matches occurred, and found them to lie in regions of high image complexity that were near a depth discontinuity. In future research we are interested in systematically exploring the hypothesis that fractional image subset match that is sufficiently high to ensure a bidirectional match in spite of simultaneous fractional image subset match that is sufficiently poor to yield the very highest MV predictor values, is a predictor of bidirectional match compliance with ground truth disparity.

Magnified data pertaining to the Art stereogram, that is presented in Figure 13 supports the hypothesis that in contrast to the Uniqueness of Dissimilarity predictor , low uniqueness of the global minimum of variability across contributions to image subset dissimilarity is a slight predictor of the extent to which SAD bidirectional matches agree with ground truth disparity. Furthermore it supports the hypothesis that this predictor benefits slightly from both use of a thicker search space and from restriction of attention to only those bidirectional matches for which the global minimum of SAD image subset dissimilarity, and the global minimum of variability across contributors to each SAD image subset dissimilarity, occur at the same location.

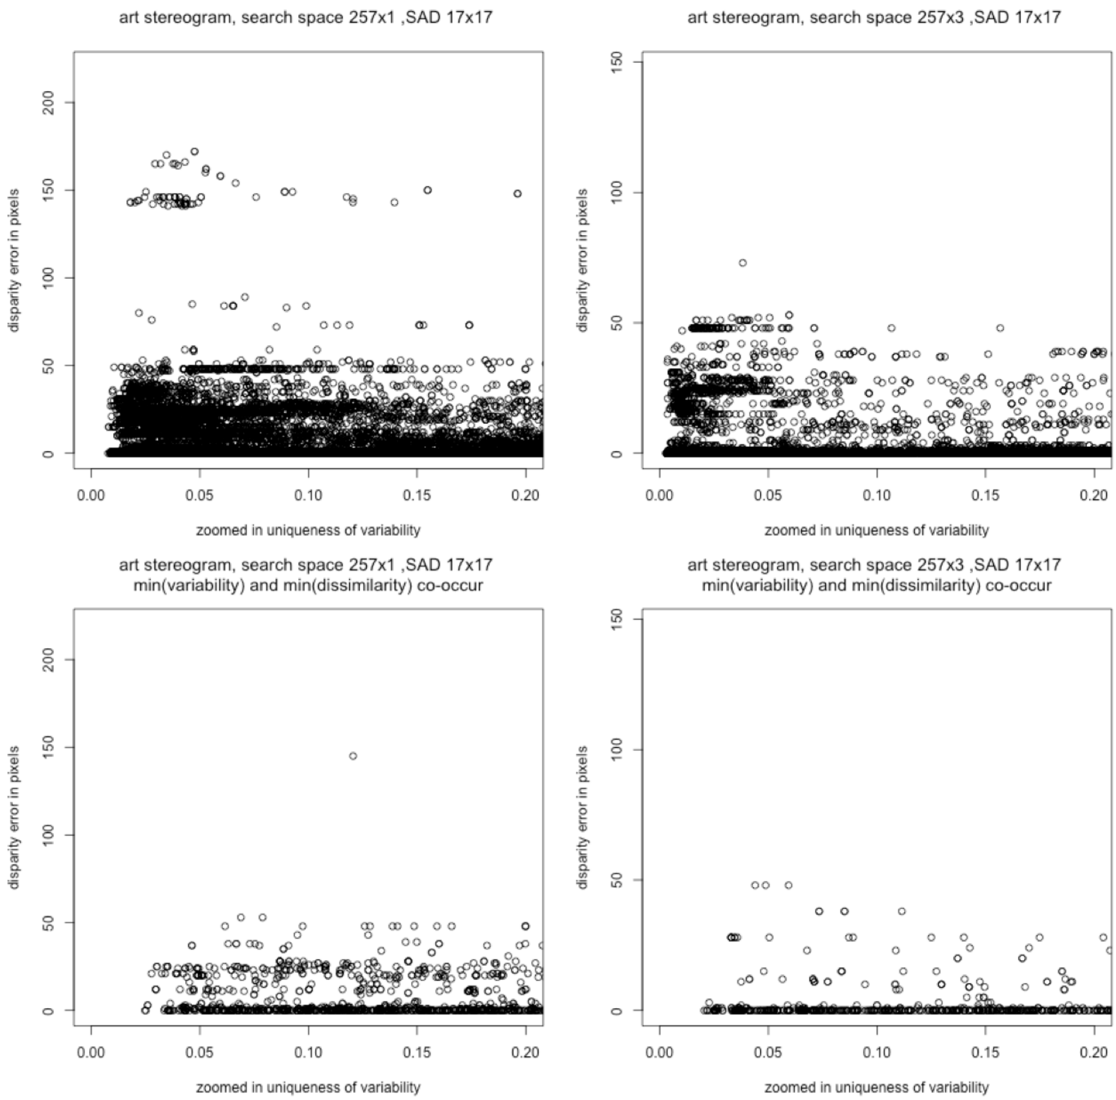

Figure 13. Impact of co-occurrence and search space thickness on UV (uniqueness of variability) predictor effectiveness

Figure 14 presents an unmagnified version of the UV predictor data that was presented in Figure 13.

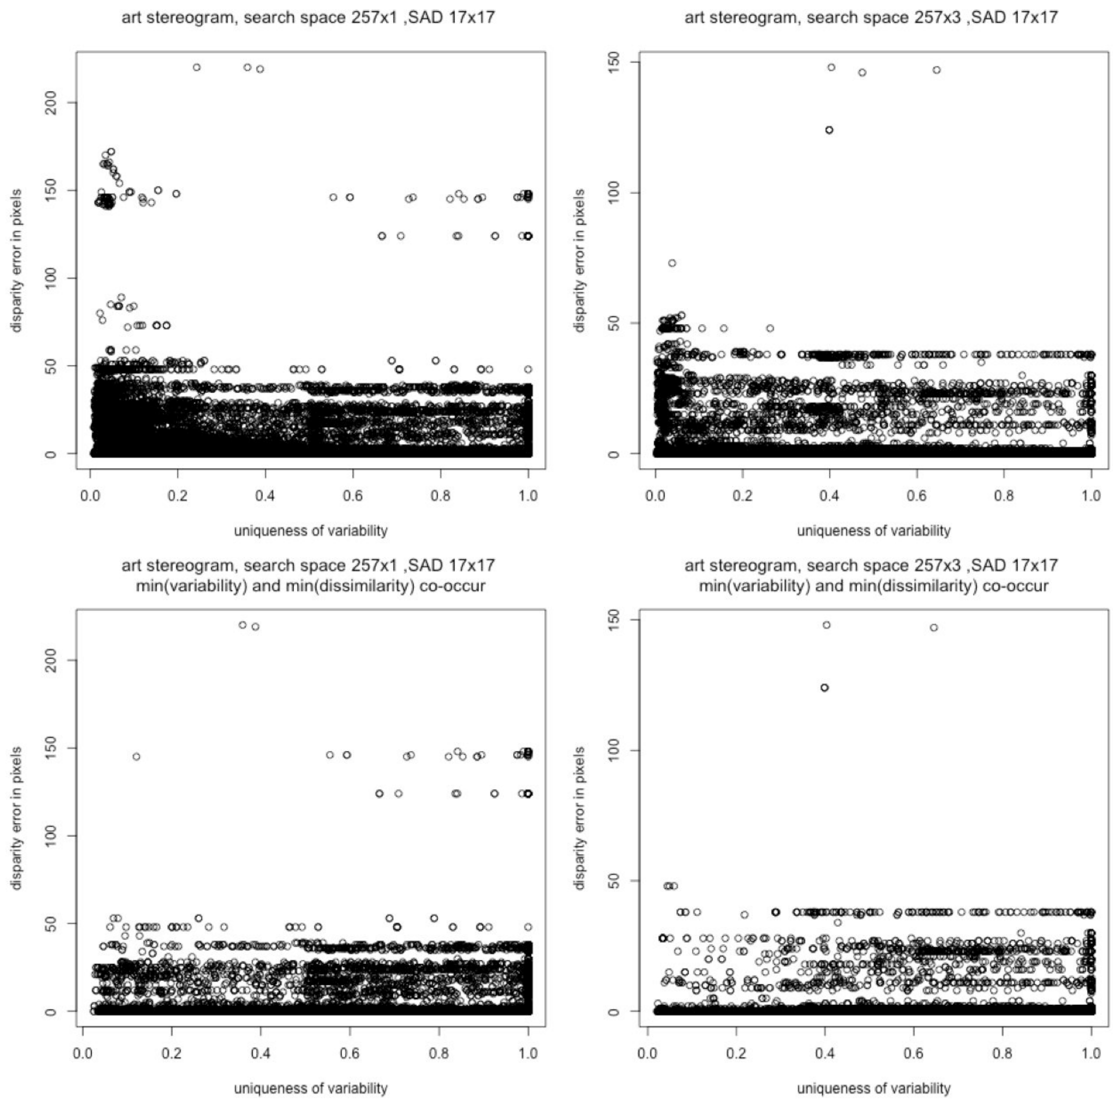

Figure 14. Impact of co-occurrence and search space thickness on UV (uniqueness of variability) predictor effectiveness

Data pertaining to the Moebius stereogram, that is presented in Figure 15 supports the hypothesis that high uniqueness of global minimum of image subset dissimilarity is a good predictor of the extent to which SAD bidirectional matches agree with ground truth disparity. Furthermore it supports the hypothesis that this predictor benefits substantially from use of a thicker search space and very little from restriction of attention to only those bidirectional matches for which the global minimum of SAD image subset dissimilarity, and the global minimum of variability across contributors to each SAD image subset dissimilarity, occur at the same location.

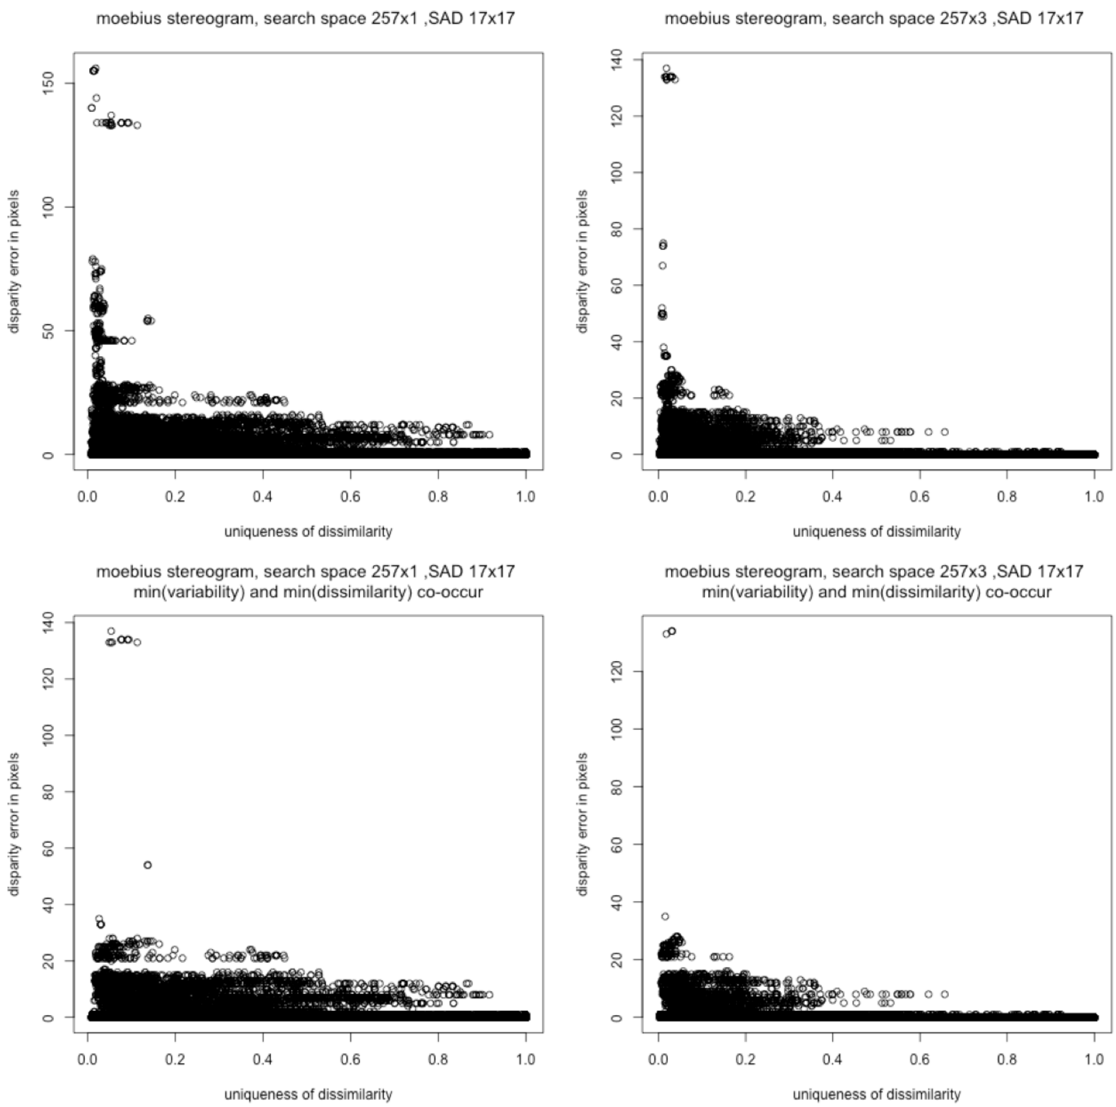

Figure 15. Impact of co-occurrence and search space thickness on UD (uniqueness of dissimilarity) predictor effectiveness

Magnified data pertaining to the Moebius stereogram, that is presented in Figure 16 supports the hypothesis that small magnitude of global minimum of image subset dissimilarity is a predictor of the extent to which SAD bidirectional matches agree with ground truth disparity. Furthermore it supports the hypothesis that this predictor benefits slightly from the use of a thicker search space and substantially from restriction of attention to only those bidirectional matches for which the global minimum of SAD image subset dissimilarity, and the global minimum of variability across contributors to each SAD image subset dissimilarity, occur at the same location.

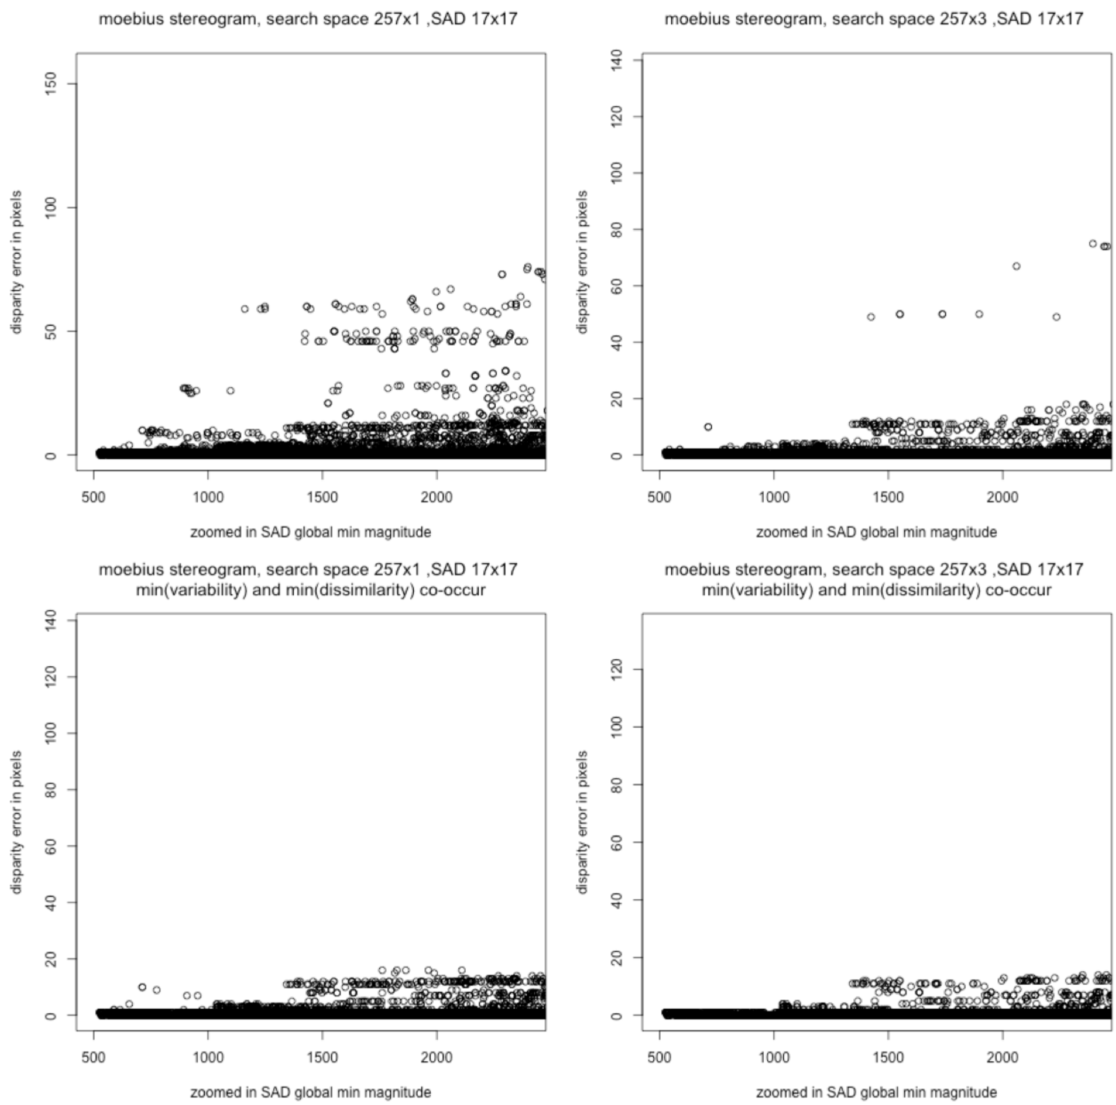

Figure 16. Impact of co-occurrence and search space thickness on MD (magnitude of dissimilarity) predictor effectiveness

Figure 17 presents an unmagnified version of the MD predictor data that was presented in Figure 16.

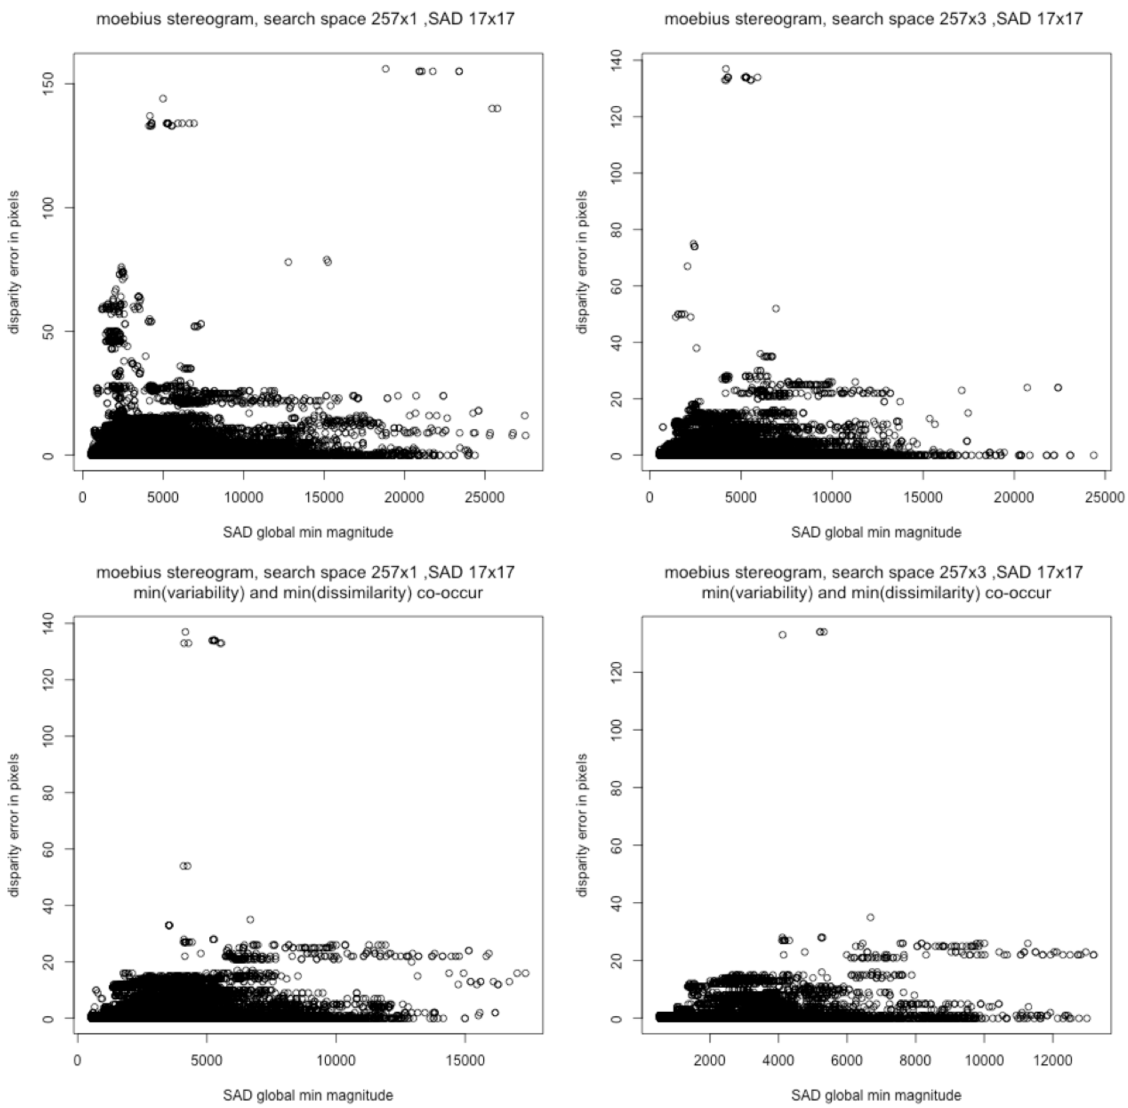

Figure 17. Impact of co-occurrence and search space thickness on MD (magnitude of dissimilarity) predictor effectiveness

Magnified data pertaining to the Moebius stereogram, that is presented in Figure 18 supports the hypothesis that small variability across contributors to the global minimum of image subset dissimilarity is a predictor of the extent to which SAD bidirectional matches agree with ground truth disparity. Furthermore it supports the hypothesis that this predictor benefits slightly from the use of a thicker search space and substantially from restriction of attention to only those bidirectional matches for which the global minimum of SAD image subset dissimilarity, and the global minimum of variability across contributors to each SAD image subset dissimilarity, occur at the same location.

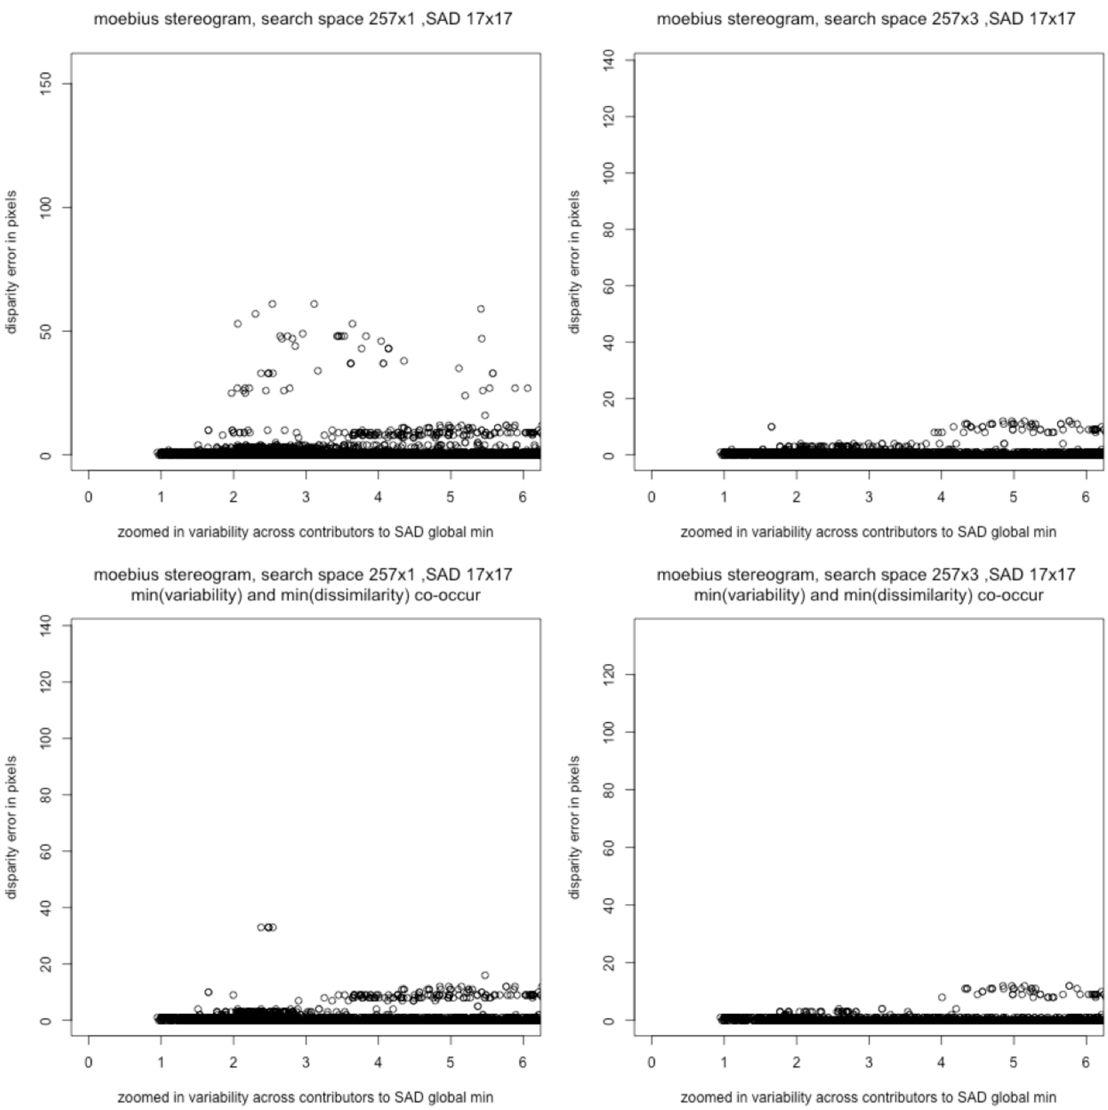

Figure 18. Impact of co-occurrence and search space thickness on MV (magnitude of variability) predictor effectiveness

Figure 19 presents an unmagnified version of the MV predictor data that was presented in Figure 18.

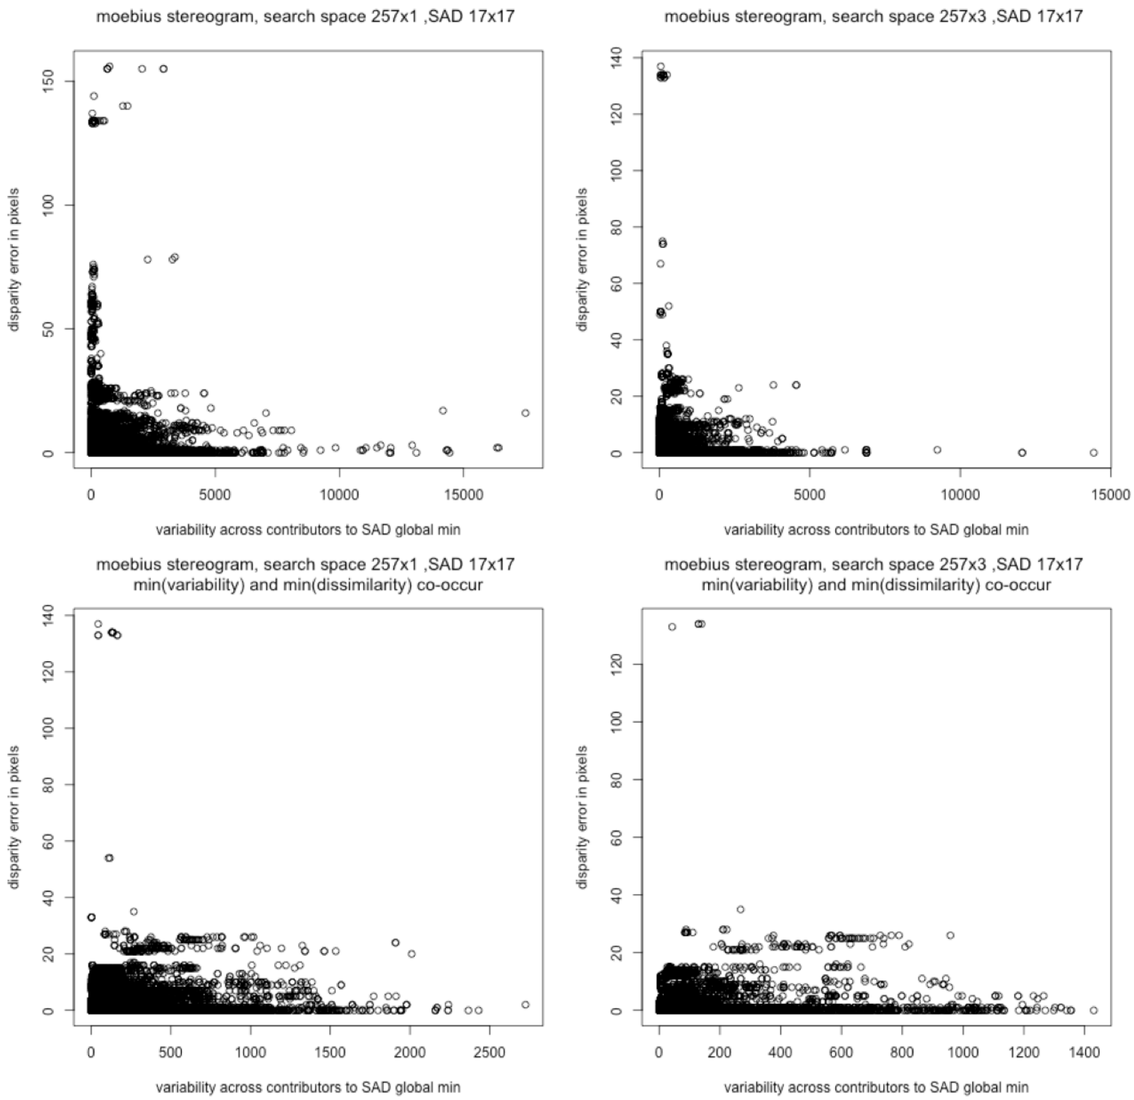

Figure 19. Impact of co-occurrence and search space thickness on MV (magnitude of variability) predictor effectiveness

Magnified data pertaining to the Moebius stereogram, that is presented in Figure 20 supports the hypothesis that in contrast to the Uniqueness of Dissimilarity predictor , low uniqueness of the global minimum of variability across contributions to image subset dissimilarity is a slight predictor of the extent to which SAD bidirectional matches agree with ground truth disparity. Furthermore it supports the hypothesis that this predictor benefits slightly from both use of a thicker search space and from restriction of attention to only those bidirectional matches for which the global minimum of SAD image subset dissimilarity, and the global minimum of variability across contributors to each SAD image subset dissimilarity, occur at the same location.

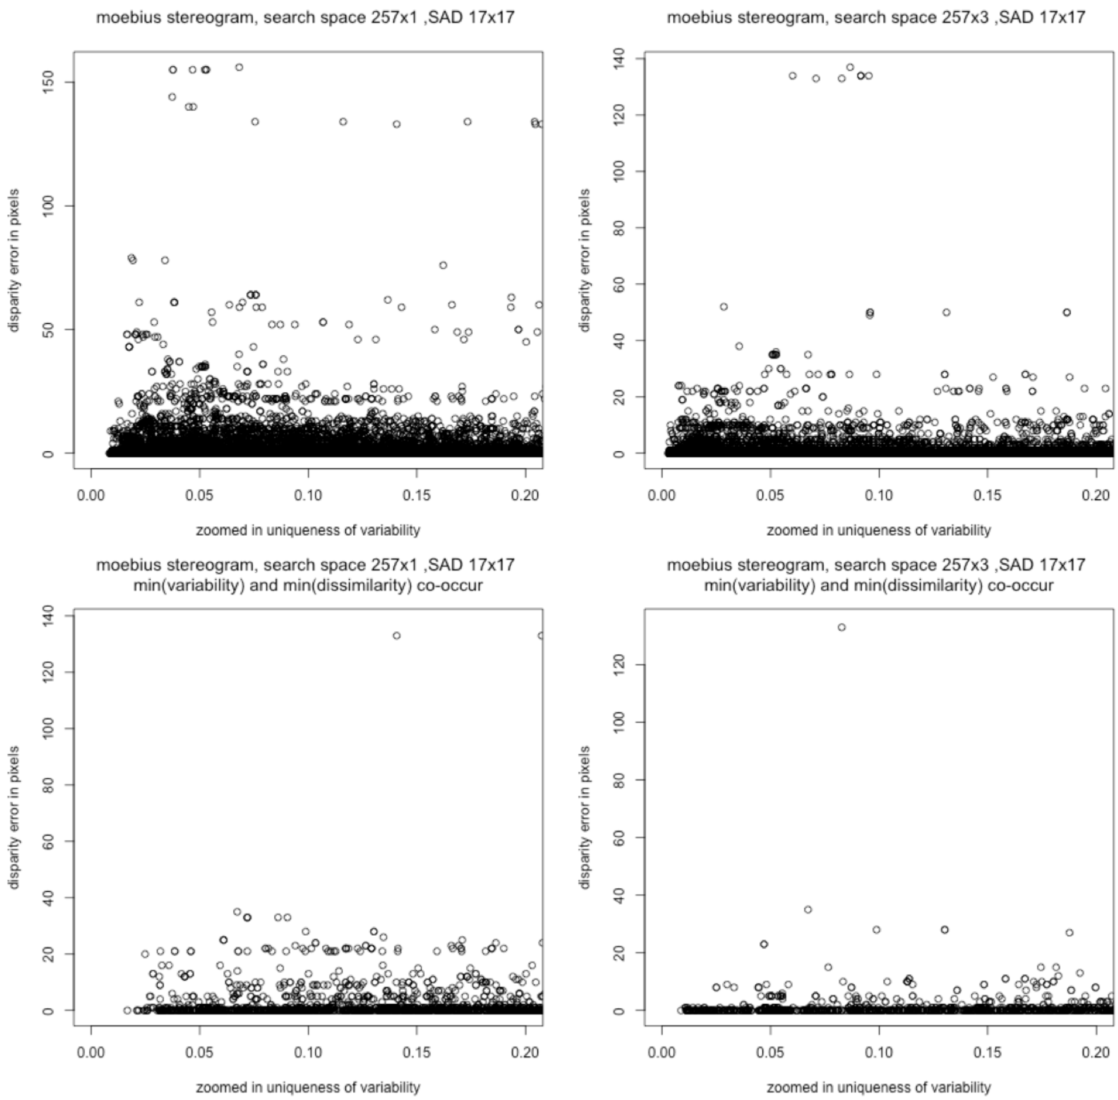

Figure 20. Impact of co-occurrence and search space thickness on UV (uniqueness of variability) predictor effectiveness

Figure 21 presents an unmagnified version of the UV predictor data that was presented in Figure 20.

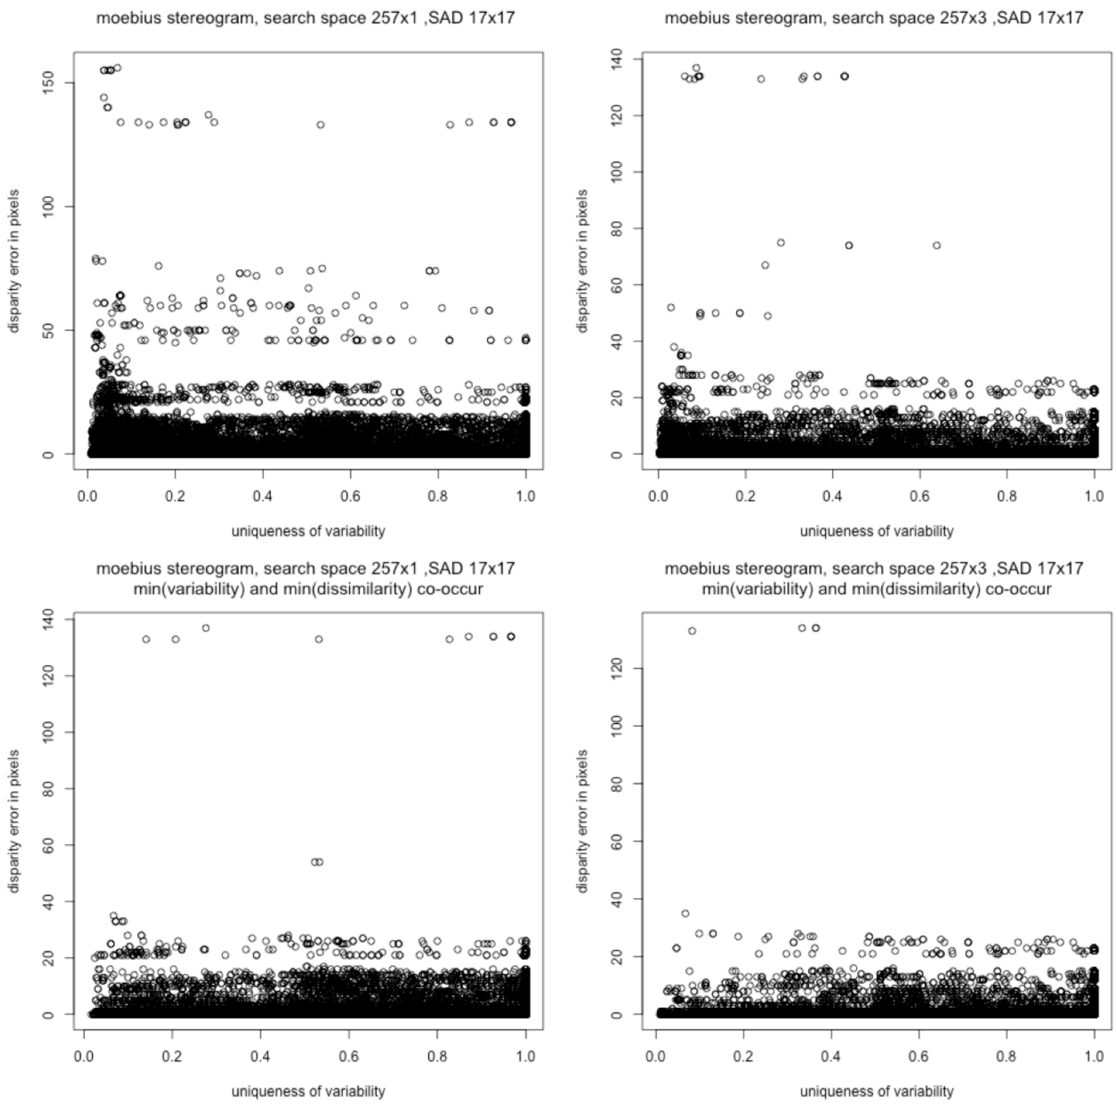

Figure 21. Impact of co-occurrence and search space thickness on UV (uniqueness of variability) predictor effectiveness

The data presented in Figures 1 through 21 suggests that the UD, MD, and MV predictors are each fairly effective when combined with either a thick search space or the co-occurrence of global min of variability and global min of dissimilarity constraint. In the figures which follow, we present data describing the effectiveness of at least one way in which UD, MD and MV can be combined with each other. Specifically, for each of the Laundry, Art, and Moebius stereograms, we present Disparity Error versus Uniqueness of Dissimilarity for the entire set of computed bidirectional matches , for the set of bidirectional matches from which high dissimilarity matches have been removed, for the set of bidirectional matches from which high variability matches have been removed, and for the set of bidirectional matches from which both high variability and high dissimilarity matches have been removed. What constitutes "high" in this context is somewhat arbitrary, however, we have tried a variety of thresholds and have repeatedly observed that regardless of which stereogram is used, removing high variability data (which we call MV filtering) improves the effectiveness of the UD predictor, while removing high dissimilarity data, (which we call MD filtering) does not have an impact on match correctness that can not also be obtained through MV filtering. We have however observed instances in which the sets of matches obtained through these alternative techniques are not comprised of the same elements, hence there remains a reason to explore use of both techniques and to consider merging the results. The UV predictor did not turn out to be particularly effective, and although we plan to continue studying it in the future, we will not report further data pertaining to it in this paper. All data presented in the figures 22 through 25 pertains to thick search spaces for which the global min of dissimilarity and the global min of variability across contributors to dissimilarity occur at the same search space location. Combined predictor data pertaining to the Laundry stereogram is presented below.

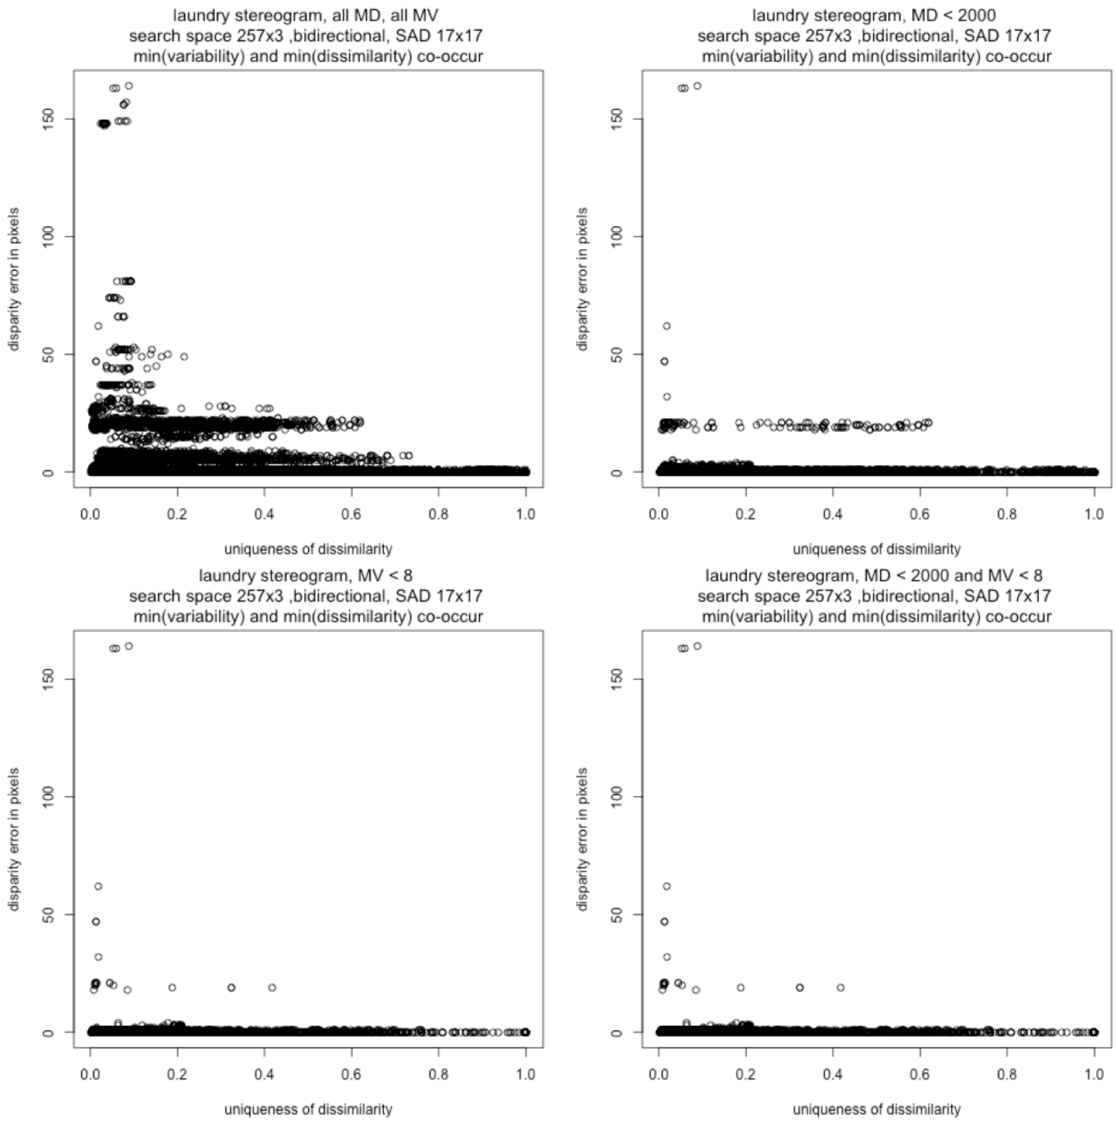

Figure 22. Impact of MD filtering and MV filtering on effectiveness of UD predictor (Laundry Stereogram).

The difficulty of the Laundry stereogram is further illustrated by the fact that the UD predictor was not able to achieve 100% compliance of bidirectional matches with ground truth disparity without assistance from the MV predictor.

Combined predictor data pertaining to the Art stereogram is presented below.

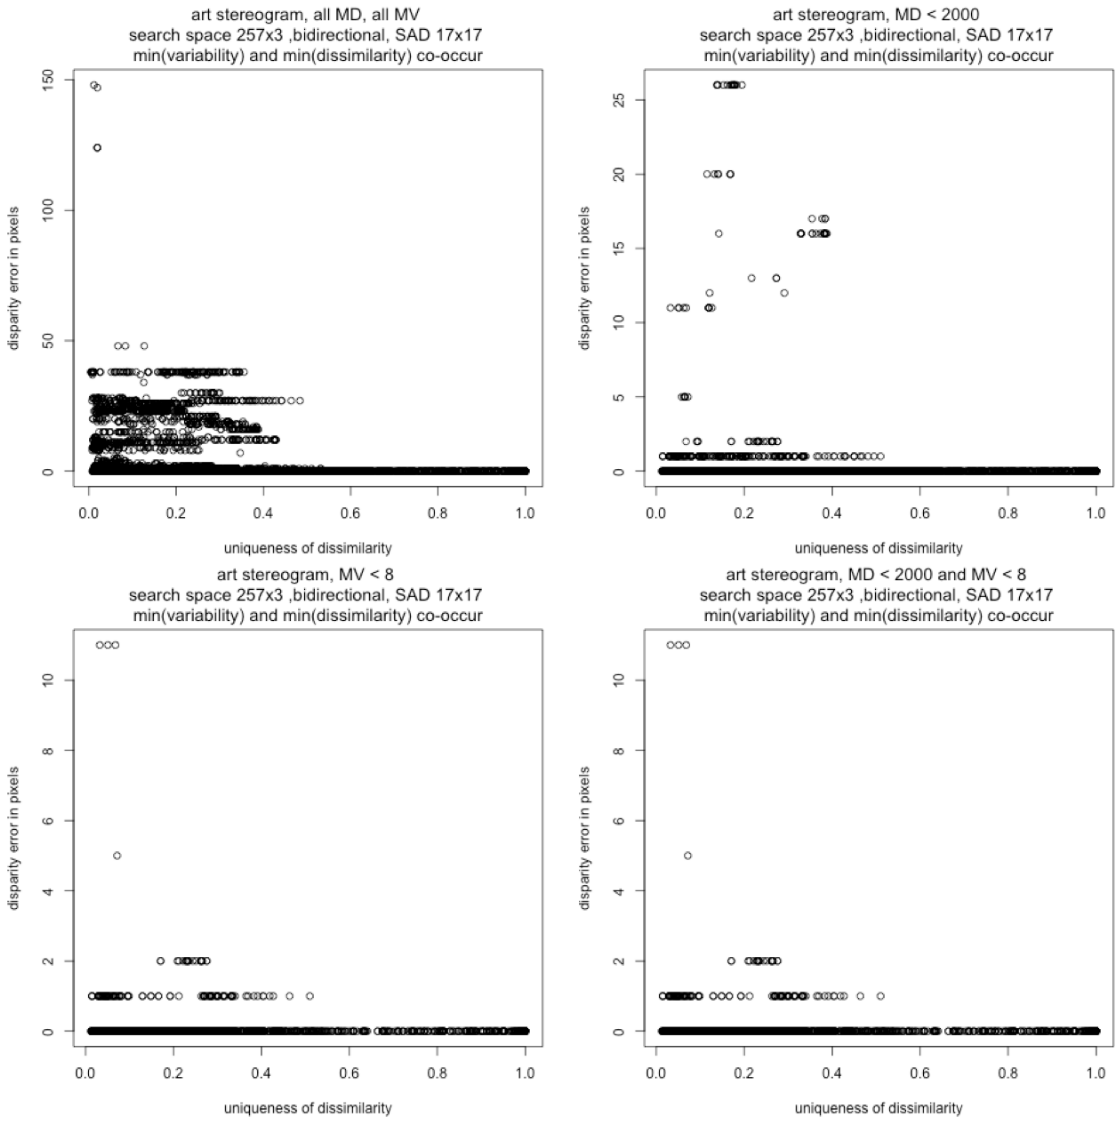

Figure 23. Impact of MD filtering and MV filtering on effectiveness of UD predictor (Art Stereogram).

The above data suggests that depending on what one is attempting to accomplish, and depending on the difficulty of the stereograms one is processing, combining the UD predictor with the MV predictor is not necessarily a good thing. It is true that for the case of the Art stereogram, combining them removes many highly incorrect matches from the UD < 0.6 range, but it does not remove all of them, and furthermore, it also removes many perfect matches from the UD > 0.6 range. In future research we are interested in exploring techniques for attempting to forecast various aspects of stereogram difficulty.

Combined predictor data pertaining to the Moebius stereogram is presented below.

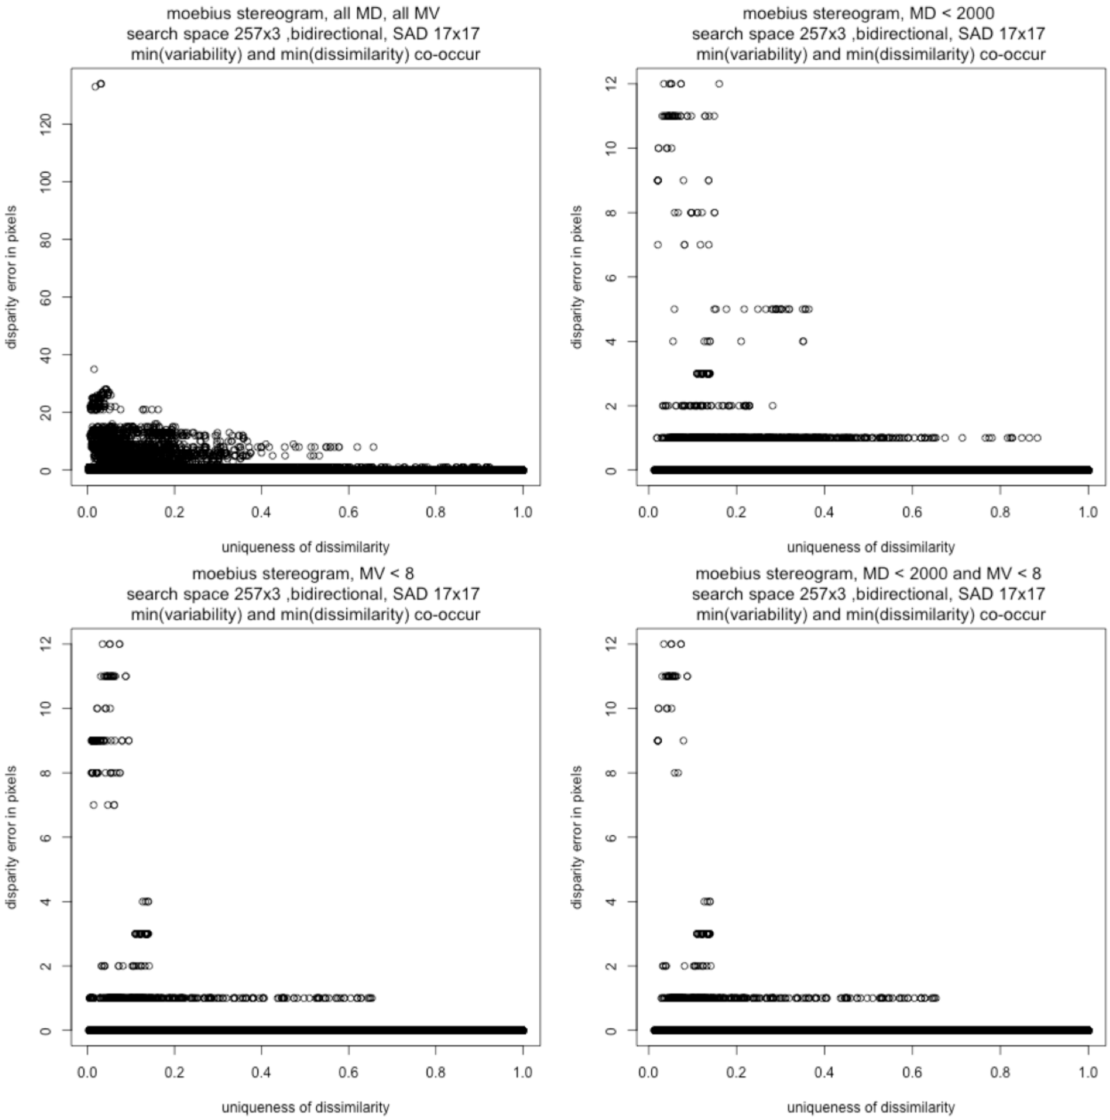

Figure 24. Impact of MD filtering and MV filtering on effectiveness of UD predictor (Moebius Stereogram).

Unlike for the case of the Art stereogram, the above Moebius stereogram data, like the earlier Laundry stereogram data, presents an example of a situation in which combining UD and MV helps us to lower the UD threshold beyond which bidirectional matches comply perfectly with ground truth disparity. Together the three sets of data suggest that, barring availability of stereogram difficulty prediction techniques, whenever we need small data sets that comply perfectly with ground truth disparity, it might be safest to combine UD and MV.

4. CONCLUSION

The data we have presented suggests that using search spaces that are slightly thicker than pairs of corresponding epipolar lines greatly improves the ability of the Uniqueness of the global minimum of image subset Dissimilarity (UD) to predict the extent to which bidirectional matches comprised of such global minima agree with ground truth disparity. The data suggests that some of the predictors we studied also perform better when limited to the set of bidirectional matches for which the global minimum of image subset dissimilarity and global minimum of variability across contributors to image subset dissimilarity occur at the same location. Finally the data suggests that the UD, MD, and MV predictors can be used individually and also that they are somewhat complementary, i.e. not completely correlated, and can be combined to yield benefits that exceed those that can be obtained using any one predictor alone. We are interested in exploring alternative ways to combine them in the future. It is somewhat surprising that the MD and MV predictors worked as well as they did when used alone. For stereograms depicting scenes containing large perfectly homogeneous surfaces, we would expect that it would be difficult to make MD or MV work well without filtering (for example using UD) to remove the least unusual image subsets from consideration.

Our application of continuous quantification of uniqueness to the classical stereoscopic correspondence problem of computer vision, made use of the concept of probability, in an environment that is completely deterministic in the sense that although the process which computes stereoscopic match candidates does not know whether or not those candidates comply with ground truth disparity, there is no uncertainty about whether or not they do so. Is use of probability in such a context defensible? It is our contention that at the very least, such a use is not unprecedented, and is consistent with the Bayesian and Classical interpretations of the mathematics of probability, as quantification of forecasters' ignorance [17] regarding a predetermined outcome, rather than quantification of the likelihood of an outcome that is not predetermined. By definition, opportunities to do the latter, do not arise in deterministic environments. Our exploration of this area introduced us to the nuances of the disagreement between the Bayesian and Frequentist positions regarding interpretation of the axiomatized mathematics of probability, and led us to suspect that the two viewpoints are in fact not incompatible, and that a key to reconciling them might be to:

a) Consider leaving to the experimental physical sciences the task of deciding whether or not any specific experiment (or phenomenon) is deterministic in the sense that repeating the experiment in exactly the same way, always produces the same outcome, or truly random a.k.a. nondeterministic, in the sense that it does not. (Note that having a human gambler roll the exact same unbiased dice twice in a casino is unlikely to constitute exact repetition of an experiment, because in all likelihood the gambler does not impart the same initial velocities, angular momentums and so on to them each time, nor is it likely that the molecules which constitute the portion of the universe with which those dice interact during their trajectories, are in the same state during each roll)

b) Consider leaving to Mathematics and Philosophy only the conditional task of proposing various ways to quantify different hypothetically plausible types of uncertainty, for example, the task of exploring uncertainty quantification proposals which do not require making assumptions about whether or not all phenomena in the universe are deterministic. (The physical sciences may very well produce evidence supporting the hypothesis that some phenomena are deterministic and others are nondeterministic, or that there exist nondeterministic phenomena which conform to constant probability distributions and nondeterministic phenomena which do not. Steven Pinker presents a nice introduction [18] to the disagreement between those who are satisfied that Quantum Mechanics has already done some of this, those who expect future research to reveal subatomic determinism, and those who predict a lack of subatomic determinism that almost always averages out to yield deterministic macroscopic phenomena.)

c) Consider drawing a distinction between uncertainty, it's classification, and its quantification, and reserve use of the word probability only to refer to it's quantification.

d) Consider that many things which can be quantified objectively, can also be quantified subjectively and vice versa, that this possibility is in no way paradoxical, and that whenever subjective quantification suffices because it is possible to complete tasks without resorting to precise consistent techniques, the benefit of subjective quantification is that it requires less effort expenditure both on computation and on the gathering of data and/or making of measurements.

We are not the only ones to suspect that a pluralist approach to the quantification of uncertainty is defensible. Donald Gillies [19] discusses this topic in great detail. He proposes a satisfying resolution to Humphreys' Paradox, as well as an alternative to the Subjective and Propensity interpretations of the axiomatized mathematics of probability, which he calls "Intersubjective". He reviews the Ramsey de Finetti theorem (according to which a gambler can avoid becoming the victim of a Dutch book if and only if the odds he assigns to various gambles conform to the Kolmogorov axioms of probability) and discusses the usefulness of that theorem, both as justification of the subjective interpretation of the axiomatized mathematics of probability, and to motivate formulation of the Kolmogorov axioms in the first place.

During the course of future research we are interested in exploring the hypotheses that Gillies' proposals and those of Laplace, Keynes, von Mises, de Finetti, Popper, Miller, and Fetzer that he critiques, can be further reconciled, that the Laplace classical formulation of probability is more widely applicable than von Mises suggested, that the Bayesian interpretation of probability as quantification of ignorance need not be subjective, that the reference class problem does not apply to nondeterministic phenomena since repeatedly performing the same experiment in exactly the same way can yield a probability distribution rather than producing the same outcome each time, that even for deterministic phenomena the reference class problem can be avoided by restricting attention to quantification of ignorance of forecasters who are attempting to predict outcomes of well defined experiments rather than attempting to predict occurrences of events with no constraints on how those events may come about, and that probability, in other words the quantification of uncertainty, is no more "Janus Faced" in other words "Two Faced" than the quantification of anything else. The latter hypothesis is based the observation that, although uncertainty can be quantified both subjectively and objectively no one quantification procedure is in some paradoxical sense simultaneously objective and subjective. We do not say that "length is Janus faced" because on the one hand we can measure a piece of furniture objectively using a tape measure, and on the other hand we can measure it subjectively by asking a neighbor whether he believes that it will easily fit through most doors.

Several people have asked what the image subsets to which our uniqueness quantification measure assigns high values look like. Our answer is that they do not have a specific kind of appearance. Neither large size, nor a high degree of complexity are guarantees of image subset uniqueness, nor are low complexity or small size a guarantee of it's absence. Having said that, it is true that in most images of natural scenes, in other words scenes which do not contain many identical precision manufactured objects, we would expect image subsets that are larger and that contain more complexity, to have a higher likelihood of being unique than those that are not. An object's uniqueness is a measure of it's relationship to other objects. Any object that is not unique, can be made unique by destroying all other objects that resemble it, and any object that is highly unique can be made not unique by creating replicas of it. It is the assessment of this conceptual property, that Distinctiveness and Continuous Quantification of Uniqueness have in common with each other, and that distinguishes them from interest operators which preceded them. We suspect more measures that assess this conceptual property will emerge in the future, and suggest the term Unusualness Assessments to refer to all of them as a group.

REFERENCES

[1] Scharstein, D., Szeliski, R.,"A taxonomy and evaluation of dense two-frame stereo correspondence algorithms", IJCV 47, 7-42 (2002).
[2] Marr, D., Poggio, T., "Cooperative computation of stereo disparity.",Science 194, 283-287 (1976).
[3] Marr, D. ,[Vision], W. H. Freeman and and Company , San Francisco, California, (1982)
[4] Faugeras, O., Luong Q. T. , Maybank S., "Camera self-calibration: theory and experiments", Proceedings of European Conference on Computer Vision, (1992).
[5] 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).
[6] Luong, Q. T., Faugeras O. D., "The fundamental matrix: theory, algorithms, and stability analysis", IJCV 17(1), 43-75 (1996).
[7] Petran, V. , Merat, F. , "Stereoscopic vision without continuity assumptions" , Proc of IEEE SSIAI , (2006)
[8] Yoon, K., Kweon, I. , "Stereo matching with distinctive similarity measure", ICCV, (2007)
[9] Manduchi, R. , Tomasi C., "Distinctiveness maps for image matching", ICIAP Venice Italy, (1999)
[10] Mordohai, P., "The Self-aware matching measure for stereo", ICCV Kyoto Japan (2009)
[11] Petran, V., Merat, F., "Continuous quantification of uniqueness and stereoscopic vision", Proc of SPIE Vol. 8056, (2011)
[12] 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).
[13] Hirschmuller, H. ,Scharstein, D. "Evaluation of cost functions for stereo matching", IEEE CVPR , 195-202 (2003).
[14] Hirschmuller, H. ,Scharstein, D. "Evaluation of stereo matching costs on images with radiometric differences", IEEE TPAMI 31(9), 1582-1599 (2009).
[15] Scharstein, D., Pal, C.,"Learning conditional random fields for stereo", IEEE CVPR Minneapolis MN, (2007).
[16] Scharstein, D., Szeliski, R.,"High-accuracy stereo depth maps using structured light", IEEE CVPR Madison WI 1,195-202 (2003).
[17] Sivia, D. S.,[Data Analysis: A Bayesian Tutorial], Oxford University Press, New York, (1996)
[18] Pinker, S., [How The Mind Works], W. W. Norton & Company, New York, (2009)
[19] Gillies, D., [Philosophical Theories of Probability] , Routledge, London, (2000)