# An Improved Solver for Square-Piece, Mixed-Bag Jigsaw Puzzles

## CS297 Final Report

Zayd Hammoudeh
(zayd.hammoudeh@sjsu.edu)

May 15, 2016

# Table of Contents

ii

# List of Figures

# List of Tables

# 1 Introduction

Jigsaw puzzles were first introduced in the 1760s when they were made from wood; their name derives from the jigsaws that were used to carve the wooden pieces. The 1930s saw the introduction of the modern jigsaw puzzle where an image was printed on a cardboard sheet that was cut into a set of interlocking pieces [1, 2]. Although jigsaw puzzles had been solved by children for two centuries, it was not until 1964 that the first automated jigsaw puzzle solver was proposed by Freeman & Gardner [3]. While an automated jigsaw puzzle solver may seem trivial, the problem has been shown by Altman [4] and Demaine & Demaine [5] to be strongly NP-complete when pairwise compatibility between pieces is not a reliable metric for determining adjacency.

Jig swap puzzles are a specific type of jigsaw puzzle where all pieces are equally sized, non-overlapping squares. Jig swap puzzles are substantially more challenging to solve since piece shape cannot be considered when determining affinity between pieces. Rather, only the image information on each individual piece is used when solving the puzzle.

Solving a jigsaw puzzle simplifies to reconstructing an object from a set of component pieces. As such, techniques developed for jigsaw puzzles can be generalized to many practical problems. Examples where jigsaw puzzle solving strategies have been used include: reassembly of archaeological artifacts [6, 7], forensic analysis of deleted files [8], image editing [9], reconstruction of shredded documents [10], DNA fragment reassembly [11], and speech descrambling [12]. In most of these practical applications, the original, also known as "ground-truth," input is unknown. This significantly increases the difficulty of the problem as the structure of the complete solution must be determined solely from the bag of component pieces.

This thesis proposes an improved jig swap puzzle solver that is able to solve multiple puzzles simultaneously. What is more, it defines a set of new metrics for measuring the quality of outputs of such solvers. Lastly, this thesis proposes enhancements to existing techniques to improve solver performance on computer generated images.

# 2  Previous Work

Computational jigsaw puzzle solvers have been studied since the 1960s when Freeman & Gardner proposed a solver that relied only on piece shape and could puzzles with up to nine pieces [3]. Since then, the focus of research has gradually shifted from traditional jigsaw puzzles to jig swap puzzles.

Cho *et al.* [13] proposed in 2010 one of the first modern computational jig swap puzzle solvers; their approach relied on a graphical model built around a set of one or more "anchor piece(s)," which are pieces whose position is fixed in the correct location before the solver began. Cho *et al.*'s solver required that the user specify the puzzle's actual dimensions. Future solvers would improve on Cho *et al.*'s results while simultaneously reducing the amount of information (beyond the set of pieces) passed to the solver.

A significant contribution of Cho *et al.* is that they were first to use the LAB (Lightness and the A/B opponent color dimensions) colorspace to encode image pixels. LAB was selected due to its property of normalizing the lightness and color variation across all three pixel dimensions. Cho *et al.* also proposed a measure for quantifying the pairwise distance between two puzzle pieces that became the basis of most of the future work (see Section 3).

Pomeranz *et al.* [14] proposed an iterative, greedy jig swap puzzle solver in 2011. Their solver did not rely on anchor pieces, and the only information passed to the solver were the pieces, their orientation, and the size of the puzzle. Pomeranz *et al.* also generalized and improved on Cho *et al.*'s piece pairwise distance measure by proposing a "predictive distance measure." Finally, Pomeranz *et al.* introduced the concept of "best buddies," which are any two pieces that are more similar to each other than they are to any other piece. Best buddies have served as both an estimation metric for the quality of solver result as well as the foundation of some solvers' placers [15].

An additional key contribution of Pomeranz *et al.* is the creation of three image benchmarks. The first benchmark is comprised of twenty 805 piece images; the sizes of the images in the second and third benchmarks are 2,360 and 3,300 pieces respectively.

In 2012, Gallagher [16] formally categorized jig swap puzzles into four primary types. The following is Gallagher's proposed terminology; his nomenclature is used throughout this thesis.

- **Type 1 Puzzle**: The dimensions of the puzzle (i.e., the width and height of the ground-truth image in number of pixels) is known. The orientation of each piece is also known, which means that there are exactly four pairwise

relationships between any two pieces. A single anchor piece, with a known, correct, location is required with additional anchor pieces being optional. This type of puzzle is used by [13, 14].

- **Type 2 Puzzle**: This is an extension of a Type 1 puzzle, where pieces may be rotated in 90° increments (e.g., 0°, 90°, 180°, or 270°); in comparison to a Type 1 puzzle, this change alone increases the number of possible solutions by a factor of $4^n$, where $n$ is the number of puzzle pieces. What is more, no piece locations are known in advance; this change eliminates the use of anchor piece(s). Lastly, the dimensions of the ground-truth image may be unknown.

- **Type 3 Puzzle**: All puzzle piece locations are known and only the rotation of the pieces is unknown. This is the least computationally complex of the puzzle variants and is generally considered the least interesting. Type 3 puzzles are not explored as part of this thesis.

- **Mixed-Bag Puzzles**: The input set of pieces are from multiple puzzles, or there are extra pieces in the input set that belong to no puzzle. The solver may output either a single, merged puzzle, or it may separate the input pieces into disjoint sets that ideally align the set of ground-truth puzzles. This type of puzzle is the primary focus of this thesis.

Sholomon *et al.* [17] in 2013 proposed a genetic algorithm based solver for Type 1 puzzles. By moving away from the greedy approach used by Pomeranz *et al.*, Sholomon *et al.*'s approach is more immune to suboptimal decisions early in the placement process. Sholomon *et al.*'s algorithm is able to solve puzzles of significantly larger size than previous techniques (e.g., greater than 23,000 pieces). What is more, Sholomon *et al.* defined three new large image (e.g., 5,015, 10,375, and 22,834 piece) benchmarks [18].

Paikin & Tal [15] published in 2015 a greedy solver that handles both Type 1 and Type 2 puzzles, even if those puzzles are missing pieces. What is more, their algorithm is one of the first to support solving Mixed-Bag Puzzles. Paikin & Tal's algorithm is used as the basis for much of this thesis and is discussed in significant depth in Section 5.

# 3 Puzzle Piece Pairwise Affinity

All jigsaw puzzle solvers require a means to measure the congruity of individual puzzle pieces. Pairwise affinity quantifies the similarity between the sides of two puzzle pieces. $D(x_i, s_a, x_j, s_b)$ and $C(x_i, s_a, x_j, s_b)$ represent the distance and compatibility (i.e., similarity) respectively between side $s_a$ of puzzle piece $x_i$ and side $s_b$ of puzzle piece $x_j$.

## 3.1 Cho *et al.* Pairwise Affinity

As mentioned in Section 2, Cho *et al.* [13] proposed one of earliest edge-based pairwise affinity measures for two puzzle pieces. Equation (1) defines the distance between the right side, $R$, of piece $x_i$ and the left side, $L$, of piece $x_j$ using Cho *et al.*'s approach. Note that $K$ is both the width and height of a puzzle piece in number of pixels.[1]

$$D(x_i, R, x_j, L) = \sum_{k=1}^{K} \sum_{d=1}^{3} (x_i(k, K, d) - x_j(k, 1, d))^2 \tag{1}$$

Since the LAB colorspace has three dimensions (e.g., lightness and A/B opponent color dimensions), $d$ ranges between 1 and 3. $x_i(k, K, d)$ represents the LAB value in dimension "$d$" for the pixel in row "$k$" and column "$K$" in puzzle piece $x_i$. Hence, Equation (1) is the sum of the squares of the pixel value differences between the right side of piece $x_i$ and the left side of piece $x_j$.

## 3.2 Pomeranz *et al.* Pairwise Affinity

One of the disadvantages of Cho *et al.*'s metric is that it simply squares the difference between the pieces' pixel values. In some cases, solver performance may improve if a different exponent is used. Pomeranz *et al.* in [14] generalize Equation (1) using the $(L_p)^q$ norm as shown in Equation (2).[2]

$$D(x_i, R, x_j, L) = \left( \sum_{k=1}^{K} \sum_{d=1}^{3} (|x_i(k, K, d) - x_j(k, 1, d)|)^p \right)^{\frac{q}{p}} \tag{2}$$

Equation (2) is identical to that of Cho *et al.* when $p$ and $q$ are equal to 2.

---

[1] Cho *et al.* used 7 for $K$.
[2] Pomeranz *et al.* used 28 for $K$. This has since become the standard.

Another disadvantage of the metric proposed by Cho *et al.* is that it only considers border pixels. Hence, if there is a gradient in the ground-truth image, two pieces may appear artificially dissimilar. To address this issue, Pomeranz *et al.* proposed predictive compatibility; Equation (3) is the predictive compatibility between the right side, $R$, of piece $x_i$ and the left side, $L$, of piece $x_j$.

$$C(x_i, R, x_j, L) = \sum_{k=1}^{K} \sum_{d=1}^{3} \left[ ([2x_i(k, K, d) - x_i(k, K-1, d)] - x_j(k, 1, d))^p \right.$$
$$\left. - ([2x_j(k, 1, d) - x_j(k, 2, d)] - x_i(k, K, d))^p \right]^{\frac{q}{p}} \tag{3}$$

By including in Equation (3) the difference between the column of pixels adjacent to each edge (e.g., $x_i(k, K-1, d)$ and $x_j(k, 2, d)$), predictive compatibility adjusts for local gradients across puzzle pieces.

## 3.3  Paikin & Tal Pairwise Affinity

Paikin & Tal in [15] used Pomeranz *et al.*'s predictive compatibility as the foundation of their asymmetric distance measure, which is shown in Equation (4); it is referred to "asymmetric" since unlike Equation (3), Paikin & Tal's approach only considers the gradient on one side of the image (i.e., $x_i(k, K-1, d)$).

$$D(x_i, R, x_j, L) = \sum_{k=1}^{K} \sum_{d=1}^{3} \| [2x_i(k, K, d) - x_i(k, K-1, d)] - x_j(k, 1, d) \| \tag{4}$$

Paikin & Tal set $p$ and $q$ from Equation (2) equal to 1 as it both only increased the accuracy of their solver and significantly reduced the computational time. It is important to note that since this distance is asymmetric, in most cases $D(x_i, R, x_j, L)$ does not equal $D(x_j, L, x_i, R)$.

Pomeranz *et al.*'s metric in Equation (3) only considers the relationship between two individual pieces when determining pairwise compatibility. In images or areas of images with little variation (e.g., a cloudless sky), non-adjacent puzzle pieces may have artificially low pairwise distances. Paikin & Tal's asymmetric compatibility, shown in Equation (5), corrects for this. Note that $secondD(x_i, R)$ is the second best asymmetric distance between the right side of piece $x_i$ and all other pieces' sides.

Figure 1: A Computer Manipulated Image with Misleadingly High Asymmetric Compatibility due to a White Background

$$C(x_i, R, x_j, L) = 1 - \frac{D(x_i, R, x_j, L)}{secondD(x_i, R)} \tag{5}$$

By normalizing compatibility with respect to the second best distance, pairings that have uniquely high compatibility are boosted, and speciously high compatibility pairings from areas of the image with low variation are penalized.

## 3.4 An Improved Asymmetric Compatibility

It is expected that images generated from an analog input (e.g., photographs), will have some degree of natural variation due to differences in brightness, the object itself, and the image capture sensor. In contrast, such variation can be trivially removed in images or parts of images that are computer generated or manipulated.

Figure 1 shows a berry in front of a white background; note that the puzzles pieces along the edge of the image and in the background have all white borders. Hence, regardless of which of the three previously described metrics is used, the distance between all of these pieces' sides is zero. However, Equation (5) is undefined when $secondD(x_i, s_a)$ is equal to zero. Therefore, to better handle computer generated images, this thesis proposes a revised definition of asymmetric compatibility shown in Equation (6).

$$C(x_i, s_a, x_j, s_b) = \begin{cases} 1 - \frac{D(x_i, s_a, x_j, s_b)}{secondD(x_i, s_a)} & secondD(x_i, s_a) \neq 0 \\ -\alpha & secondD(x_i, s_a) = 0 \end{cases} \tag{6}$$

Note that $\alpha$ is a tunable, penalty factor representing low compatibility.

## 3.5   Best Buddies

Inter-piece distance and compatibility can be used to define specific, useful relationships between pieces. For example, Pomeranz *et al.* define that two pieces, $x_i$ and $x_j$ are "best buddies" on their respective sides $s_a$ and $s_b$ if and only if they are more compatible (i.e., similar) to each other than they are to any other piece. This is shown formally in Equation (7). The definition of best buddies is identical irrespective of the compatibility function used.

$$\forall x_k \forall s_c, C(x_i, s_a, x_j, s_b) \geq C(x_i, s_a, x_k, s_c)$$

$$\text{and} \tag{7}$$

$$\forall x_k \forall s_c, C(x_j, s_b, x_i, s_a) \geq C(x_j, s_b, x_k, s_c)$$

$x_k$ is any puzzle piece from the set of input pieces, and $s_c$ is one of the four sides of $x_k$.

It is relatively rare that two pieces are best buddies and are not actually neighbors [15]. When considering all sides of a piece, it is rarer still that a piece has more non-neighbor best buddies than neighbor best buddies. These attributes make best buddies a critical tool for many solvers.

### 3.5.1   Unique Best Buddies

It is relatively unlikely that in an photograph there will be a meaningful number pieces that have multiple best buddies on the same side. However, similar to the phenomenon described in Section 3.4, multiple best buddies may occur in computer generated or computer manipulated images. In such cases, best buddies become a far less discerning tool for determining piece adjacency. This thesis addresses this by modifying the definition of best buddies as shown in Equation (8).

$$\forall x_k \neq x_j \forall s_c, C(x_i, s_a, x_j, s_b) > C(x_i, s_a, x_k, s_c)$$

$$\text{and} \tag{8}$$

$$\forall x_k \neq x_i \forall s_c, C(x_j, s_b, x_i, s_a) > C(x_j, s_b, x_k, s_c)$$

Rather than relying on best buddies being "greater than or equal" to all other

(a) Original Image       (b) Best Buddy Visualization

Figure 2: Visualization of Best Buddies in an Image

| No Best Buddy | Non-Adjacent Best Buddy | Adjacent Best Buddy | No Piece Present |
|---|---|---|---|
| | 🟥 | 🟩 | ⬛ |

Table 1: Color Scheme for Puzzles Piece Sides in Best Buddy Visualizations

pieces as in Equation (7), the modified requirement is that pairings must be strictly "greater than." Hence, best buddy pairings are exclusive.

### 3.5.2   Visualizing Best Buddies

There is currently no standard for visualizing an image's best buddies. This thesis proposes such a standard for the first time. As a nomenclature, any best buddies that are neighbors in the ground-truth image are referred to as "adjacent best buddies" while any best buddies that are not neighbors are referred to as "non-adjacent best buddies."

In a jig swap puzzle, a piece may have best buddies on up to four sides (since the pieces are square). As such, each piece in the best buddy visualization is divided into four isosceles triangles; the base of each triangle is along the side of the puzzle piece whose best buddy is being represented. A puzzle piece's four isosceles triangles all share a common, non-base vertex at the piece's center.

Figure 2 shows an image and its best buddy visualization. The color scheme used to denote the different best buddy relationships is shown in Table 1.

### 3.5.3   Interior and Exterior Best Buddies

In previous research, all best buddies (in particular best buddy errors) were treated the same. However, non-adjacent best buddies are expected to naturally occur more along pieces that are on the puzzle's edge or that are next to a missing piece since those pieces lack a true neighbor; this causes those pieces to couple with unrelated pieces at a higher rate. For example, in Figure 2, the image has 4 interior and 14 exterior non-adjacent best buddies despite there being 16-times more interior sides.

When using best buddy accuracy as an estimation metric, this thesis differentiates between interior and exterior best buddy errors by giving interior best buddy errors a higher weight.

# 4 Quantifying the Quality of a Solver Output

Modern jig swap puzzle solvers are not able to perfectly reconstruct the ground-truth input in most cases. As such, quantifiable metrics are required to objectively compare the quality of outputs from different solvers. Cho *et al.* [13] defined two such metrics namely: direct accuracy and neighbor accuracy. These metrics have been used by others including [14, 15, 16, 17, 19]. This section describes the existing quality metrics, their weaknesses, and proposes enhancements to these metrics to make them more meaningful for Type 2 and Mixed-Bag puzzles.

In the final subsection, tools developed as part of this thesis to visualize the solver output quality are discussed.

## 4.1 Direct Accuracy

Direct accuracy is a relatively naïve quality; it is defined as the fraction of pieces placed in the same location in the ground-truth (i.e., original) and solved image with respect to the total number of pieces. Equation (9) shows the formal definition of direct accuracy ($DA$), where $n$ is the total number of pieces and $c$ is the number of pieces placed in their original (i.e., correct) location.

$$DA = \frac{c}{n} \tag{9}$$

Direct accuracy is vulnerable to shifts in the solved image where even a few misplaced pieces can cause a significant decrease in accuracy. As shown in Section 4.1.2, this can be particularly true when the ground-truth image's dimensions are not known by the solver.

This thesis proposes two new direct accuracy metrics namely: Enhanced Direct Accuracy Score (EDAS) and Shiftable Enhanced Direct Accuracy Score (SEDAS). They are described in the following two subsections; the complementary relationship between EDAS and SEDAS is described in the third subsection.

### 4.1.1 Enhanced Direct Accuracy Score

The standard direct accuracy metric does not account for the possibility that there may be pieces from multiple puzzles in the same solver output. For a given a puzzle $P_i$ in the set of input puzzles $P$ (where $P_i \in P$) and a set of solved

puzzles $S$ where $S_j$ is in $S$, Enhanced Direct Accuracy Score (EDAS) is defined as shown in Equation (10).

$$EDAS_{P_i} = \arg\max_{S_j \in S} \frac{c_{i,j}}{n_i + \sum_{k \neq i}(m_{k,j})} \qquad (10)$$

$c_{ij}$ is the number of pieces from input puzzle $P_i$ correctly placed (with no rotation for Type 2 puzzles) in solved puzzle $S_j$ while $n_i$ is the number of pieces in puzzle $P_i$. $m_{k,j}$ is the number of pieces from an input puzzle $P_k$ (where $k \neq i$) that are also in $S_j$.

When solving only a single puzzle, EDAS and standard direct accuracy as defined Equation (9) are equivalent. When solving multiple puzzles simultaneously, EDAS necessarily marks as incorrect any pieces from $P_i$ that are not in $S_j$ by dividing by $n_i$. What is more, the summation of $m_{k,j}$ in EDAS is used to penalize for any puzzle pieces not from $P_i$. Combined, these two factors enable EDAS to penalize for both extra and misplaced pieces.

It is important to note that EDAS is a score and not a measure of accuracy. While its value is bounded between 0 and 1 (inclusive), it is not specifically defined as the number of correct placements divided by the total number of placements since the denominator of Equation (10) is greater than or equal to the number of pieces in both $P_i$ and $S_j$.

### 4.1.2 Shiftable Enhanced Direct Accuracy Score

As mentioned previously, the direct accuracy decreases if there are shifts in the solved image. In many cases such, direct accuracy is overly punitive.

Figure 3 shows a ground-truth image and an actual solver output when the puzzle boundaries were not fixed. Note that only a single piece is misplaced; this shifted all other pieces to the right one location causing the direct accuracy to drop to zero. Had this same piece been misplaced along either the right or bottom side of the image, the direct accuracy would have been largely unaffected. The fact that direct accuracy can give such vastly differing results for essentially the same error shows that direct accuracy has a serious flaw. This thesis proposes Shiftable Enhanced Direct Accuracy Score (SEDAS) to address the often misleadingly punitive nature of direct accuracy.

Let $d_{min}$ be the Manhattan distance between the upper left corner of the solved image and the nearest placed puzzle piece. Also let $L$ be the set of all puzzle piece locations within distance $d_{min}$ of the upper left corner of the image. Given that $l$ is a location in $L$ that is used as the reference point for

(a) Ground-Truth Image

(b) Solver Output

Figure 3: Solver Output where a Single Misplaced Piece Catastrophically Affects the Direct Accuracy

determining the absolute location of all pieces, then SEDAS is defined as shown in Equation (11).

$$SEDAS_{P_i} = \arg\max_{l \in L} \left( \arg\max_{S_j \in S} \frac{c_{i,j,l}}{n_i + \sum_{k \neq i}(m_{k,j})} \right) \tag{11}$$

In the standard definition of direct accuracy proposed by Cho *et al.*, $l$ is fixed at the upper left corner of the image. In contrast, SEDAS shifts this reference point within a radius of the upper left corner of the image in order to find a more meaningful value for direct accuracy.

Rather than defining SEDAS based off the distance $d_{min}$, an alternative approach is to use the point anywhere in the image that maximizes Equation (11). However, that approach can be significantly more computationally complex in particular in large puzzles with several thousand pieces. Hence, this thesis' approach balances finding a meaningful direct accuracy score with computational efficiency.

### 4.1.3 The Necessity to Use Both EDAS and SEDAS

While EDAS can be misleadingly punitive, it cannot be wholly replaced by SEDAS. Rather, EDAS and SEDAS serve complementary roles. First, EDAS must necessarily be calculated as part of SEDAS since the upper left corner location is inherently a member of the set $L$. Hence, there is no additional time required to calculate EDAS. What is more, by continuing to use EDAS along with SEDAS, some shifts in the solved image may be quantified; this would not be possible if SEDAS was used alone.

## 4.2 Neighbor Accuracy

Cho *et al.* [13] defined neighbor accuracy as the ratio of the number of puzzle piece sides adjacent to the same piece's side in both the ground-truth and solved image versus the total number of puzzle piece sides. Formally, let $q$ be the number of sides each piece has (i.e., four in a jig swap puzzle) and $n$ be the number of pieces. If $a$ is the number of puzzle piece sides adjacent in both the ground-truth and solved images, then the neighbor accuracy, $NA$, is defined as shown in Equation (12).

$$NA = \frac{a}{n \; q} \tag{12}$$

Unlike direct accuracy, neighbor accuracy is largely unaffected by shifts in the solved image since it considers only a piece's neighbors and not its absolute location. However, the standard definition of neighbor accuracy cannot encompass the case where pieces from multiple input puzzles may be present in the same solver output.

### 4.2.1 Enhanced Neighbor Accuracy Score

Enhanced Neighbor Accuracy Score (ENAS) improves the neighbor accuracy metric by providing a framework to quantify the quality of Mixed-Bag solver outputs.

Let $n_i$ be the number of puzzles pieces in the input puzzle $P_i$ and $a_{i,j}$ be the number of puzzle piece sides adjacent in $P_i$ and $S_j$. If $m_{k,j}$ is the number of puzzle pieces from an input puzzle $P_k$ (where $k \neq i$) in $S_j$, then the ENAS for $P_i$ is defined as shown in Equation (13).

$$ENAS_{P_i} = \underset{S_j \in S}{\arg\max} \; \frac{a_{i,j}}{q \; (n_i + \sum_{k \neq i}(m_{k,j}))} \tag{13}$$

In the same fashion as the technique described for EDAS in Section 4.1.1, ENAS divides by the number of pieces $n_i$ in input puzzle $P_i$. By doing so, it effectively marks as incorrect any pieces from $P_i$ that are not in $S_j$. What is more, by including a summation of all $m_{k,j}$ in the denominator of (13), ENAS marks as incorrect any pieces not from $P_i$ that are in $S_j$. The combination of these two factors allows ENAS to account for extra and misplaced pieces.

## 4.3 Visualizing Solver Output Quality

In images with thousands of pieces, it is often difficult to visually determine the location of individual pieces that are incorrectly placed. What is more, visual tools help developers quickly detect and fix latent bugs.

The following two subsections describe the tools developed as part of this thesis for visualizing direct accuracy and neighbor accuracy.

### 4.3.1 Visualizing EDAS and SEDAS

In standard direct accuracy, EDAS, and SEDAS, each puzzle piece is assigned a single value (i.e., correct or incorrect). Due to that, the direct accuracy visualization represents each puzzle by a square filled with a solid color. One additional refinement used in this thesis is to subdivide the "incorrect" placements into a set of subcategories; they are (in order of precedence): wrong puzzle, wrong location, and wrong rotation. Table 2 shows the colors assigned to puzzle pieces depending on their direct accuracy classification.

| Wrong Puzzle | Wrong Location | Wrong Rotation | Correct Location | No Piece Present |
|---|---|---|---|---|
|  |  |  |  |  |

Table 2: Color Scheme for Puzzles Pieces in Direct Accuracy Visualizations

Figure 4 shows a Type 2 solver output along with the associated EDAS and SEDAS visualizations. Since four puzzle pieces were erroneously placed on the left of the image, almost all pieces had the wrong location according to EDAS; the only exception is a single piece that had the right location but wrong rotation. In contrast, almost all pieces have the correct location in the SEDAS representation; note that the piece in the correct location but with wrong rotation in EDAS has the wrong location in SEDAS.

### 4.3.2 Visualizing ENAS

The visualization for neighbor accuracy is very similar to the techniques described in Section 3.5.2 for visualizing best buddies where each puzzle piece is divided into four equal-sized isosceles triangles (one for each side). The triangles are assigned colors depending on whether their neighbors in the solver output and ground-truth image match. The visualization includes a subcategory known as

(a) Ground-Truth Image

(b) Type 2 Solver Output

(c) EDAS Visualization

(d) SEDAS Visualization

Figure 4: Example Solver Output Visualizations for EDAS and SEDAS

(a) Input Image # 1 - Rainforest House [20]



(b) Input Image # 2 - Building Exterior [21]



(c) Solver Output



(d) ENAS Visualization

Figure 5: Example Solver Output Visualization for ENAS

"wrong puzzle" which is a special case that occurs when solving Mixed-Big puzzles and some of the pieces in the solved puzzle are not from the puzzle of interest, $P_i$. Table 3 defines the colors used to represent the different classifications of puzzle piece sides in neighbor accuracy visualizations.

| Wrong Puzzle | Wrong Neighbor | Correct Neighbor | No Piece Present |
|---|---|---|---|
| <span style="color:blue">██</span> | <span style="color:red">██</span> | <span style="color:green">██</span> | <span style="color:black">██</span> |

Table 3: Color Scheme for Puzzles Piece Sides in Neighbor Accuracy Visualizations

Figure 5 shows an actual output when solving a Mixed-Bag puzzle with two images. Note that that the puzzle of interest $P_i$ is the glass and stone building while the other puzzle $P_k$ is a rainforest house.

All pieces that came from the rainforest house image are shown as blue despite being assembled correctly; this is because they are not from the puzzle of interest. All neighbors from the puzzle of interest (i.e., the glass and stone

16

building) that are placed next to their original neighbor are represented by green triangles while all incorrect neighbors, such as those bordering the rainforest house image, are shown with red triangles.

# 5    Paikin & Tal Solver

This section reviews the solver proposed by Paikin & Tal [15]; their Java implementation is not open-source. As such, this section also describes a complete implementation of their approach that was developed as part of this thesis.

## 5.1    Overview of Paikin & Tal's Algorithm

Paikin & Tal's solver was inspired by the work of Pomeranz *et al.* in [14]. Paikin & Tal's uses a deterministic, greedy algorithm that places the puzzle piece that has the maximum confidence score at each iteration. Paikin & Tal's approach is able to handle puzzles of unknown size with missing pieces and where piece orientation is not known. The only required input to the algorithm is the expected number puzzles.

Paikin & Tal's algorithm has three distinct phases namely: inter-puzzle piece distance calculation, selection of the seed piece, and placement. These stages are described in the following subsections. The modification required to solve Mixed-Bag puzzles is described in an additional subsection.

### 5.1.1    Inter-Puzzle Piece Distance Calculation

The first stage of Paikin & Tal's algorithm is to calculate the distance between all pairs of pieces. This is done using the asymmetric distance measure described in Section 3.3. The distance information is stored in an $n$ by $n$ matrix (where $n$ is the number of puzzle pieces); after being initially calculated, the asymmetric distances never need to be recalculated throughout the remainder of the algorithm's execution.

As explained in Section 3.3, Paikin & Tal normalize the asymmetric distances using the asymmetric compatibility function in Equation (5). This has the effect of amplifying truly unique pairings while penalizing speciously similar pairings that arise from low variation areas of the ground-truth image. The asymmetric compatibility is then used to find the best buddies for each piece (if any).

The last step is to calculate the mutual compatibility ($\tilde{C}$) between pairs of pieces using Equation (14). $C(x_i, s_a, x_j, s_b)$ is the asymmetric compatibility between side $s_a$ of piece $x_i$ and side $s_b$ of piece $x_j$; $C(x_j, s_b, x_i, s_a)$ is defined similarly. Note that mutual compatibility is symmetric.

$$\tilde{C}(x_i, s_a, x_j, s_b) = \tilde{C}(x_j, s_b, x_i, s_a) = \frac{C(x_i, s_a, x_j, s_b) + C(x_j, s_b, x_i, s_a)}{2} \quad (14)$$

### 5.1.2  Selecting the Seed Piece

Similar to Pomeranz *et al.*, Paikin & Tal use a kernel growing algorithm. Hence, a single seed piece is selected, and all pieces are placed around that seed. Since the algorithm is greedy, the selection of a poor seed can have a significant, deleterious impact on the final solution. Due to this, Paikin & Tal select a seed piece that is itself "distinctive" and comes from a "distinctive region."

Paikin & Tal define a seed piece as "distinctive" if it has best buddies on each of its sides. To ensure that a piece comes from a distinctive region, all of the seed piece's best buddies must also have four best buddies. In a puzzle, there may be multiple pieces that satisfy the "distinctive" piece in a "distinctive region" criteria; in such cases, ties are broken by selecting the piece that has the maximum sum of mutual compatibilities with its direct neighbors.

### 5.1.3  Placement

Paikin & Tal utilize an iterative, greedy placer that places pieces around an expanding seed. Pseudocode for their placer is shown in Algorithm 1. Placement continues all pieces have been placed.

---
**Algorithm 1** Paikin & Tal Placer

---
1: **while** $|UnplacedPieces| > 0$ **do**
2:    **if** $|BestBuddyPool| > 0$ **then**
3:       Get best candidate from the $BestBuddyPool$
4:    **else**
5:       Recalculate the asymmetric and mutual compatibility
6:       Select piece with the highest asymmetric compatibility
7:    Place the best piece
8:    Add the best piece's unplaced best buddies to the $BestBuddyPool$

---

The selection of the next (i.e., best) piece to place emphasizes pieces about which there is the highest confidence of correctness. As mentioned previously, two pieces are much more likely to be adjacent if they are best buddies. Due to this, the algorithm keeps a pool of the already placed pieces best buddies; whenever a new piece is placed, all of that piece's unplaced best buddies are added to the $BestBuddyPool$.

As long as the *BestBuddyPool* is not empty, candidates for placement only come from that pool; the best candidate from the *BestBuddyPool* is the one with the maximum mutual compatibility with an open slot in the puzzle. This rule prioritizes placing the best buddy pairings upon which there is the greatest confidence.

If the *BestBuddyPool* is ever empty, the algorithm recalculates the asymmetric and mutual compatibilities between the unplaced pieces and the open puzzle slots. The piece that has the maximum asymmetric compatibility with an open slot is then selected as the next piece for placement.

### 5.1.4 Solving Mixed-Bag Puzzles

As mentioned in Section 5, the only input to the Paikin & Tal algorithm is the expected number of puzzles. When solving more than one puzzle at a time, only a minor change to the placer described in Algorithm 1 is required.

Their algorithm spawns a new puzzle any time the mutual compatibility between the next piece to place its associated open slot drops below a predefined threshold[3]; this rule applies as long the current number of puzzles is less than the specified number passed to the algorithm. When a new puzzle is spawned, the *BestBuddyPool* is cleared, and a seed piece is selected using the approach previously described in Section 5.1.2. Placement then continues simultaneously across all puzzles.

## 5.2 A Python Implementation of Paikin & Tal's Algorithm

Since no open-source implementation of the Paikin & Tal solver exists, one was developed as part of this thesis. It is written entirely using Python 2.7 [22]. The following subsections describe the implementation's two Python packages.

### 5.2.1 The `hammoudeh_puzzle` Package

This package is a generic infrastructure that is independent of the solver used; it consists of two primary classes namely: `Puzzle` and `PuzzlePiece`.

The `Puzzle` class encapsulates all attributes of a puzzle including: an identification number, dimensions, type (e.g., Type 1, Type 2 - this is an enumerated

---

[3]Paikin & Tal use 0.5 as the threshold to spawn new boards

type: "PuzzleType"), and all associated puzzle pieces. A puzzle can be created either from an image file or from a set of puzzle pieces. When parsing image files and exporting a solved puzzle, the Puzzle class uses the OpenCV Python package [23].

Individual puzzle pieces are represented using the PuzzlePiece class. Each object of type PuzzlePiece has attributes: identification number, width (in number of pixels), rotation (via the class PuzzlePieceRotation), the LAB image information (stored in a NumPy array [24]), and the location in the puzzle.

Additional features included in the hammoudeh_puzzle package include calculating and visualizing EDAS, SEDAS, and ENAS as well as performing best buddy analysis on an image.

### 5.2.2   The paikin_tal_solver Package

The paikin_tal_solver package implements the Paikin & Tal algorithm described in Section 5.1. The primary interface for the user is through the PaikinTalSolver class; objects of this type are created using a constructor that takes as parameters: the expected number of puzzles, a set of PuzzlePiece objects, a distance function, the puzzle type (via class "PuzzleType"), and optionally a single set of fixed puzzle dimensions. A PaikinTalSolver object is composed of a set of component objects, which are described in the following paragraphs.

An object of type InterPieceDistance calculates the asymmetric distance, asymmetric compatibility, and mutual compatibility between all pieces. What is more, since the PaikinTalSolver takes a distance function in its constructor, the user is able to tune the solver's performance using different metrics. While calculating the asymmetric compatibilities, the InterPieceDistance class also finds each piece's best buddies; this made the InterPieceDistance class the natural choice to identify the seed piece(s).

Paikin & Tal do not identify the data structure used to implement the *BestBuddyPool* (see Algorithm 1). However, the choice of data structure is critical as the algorithm must be able to quickly remove the best candidates from the pool and quickly insert new candidates into the pool. This thesis' implementation of the *BestBuddyPool* relies on a combination three data structures. They are:

- **Dictionary of the Best Buddies in this Pool** – This stores the identification numbers of best buddies currently in the pool. A dictionary was

used here as it enables new best buddies to be quickly added to the pool and for placed pieces to be quickly deleted from the pool.

- **Dictionary of the Open Slots in the Puzzle** – It contains the set of valid locations where pieces can be placed. As with the previous data structure, this dictionary enables valid locations to be quickly added and removed.

- **Best Candidate Max Heap** – The best candidate is a pairing of a best buddy from the pool and an open puzzle location. Every time a new best buddy is added to the pool, pairings between that new pool member and all open slots are added to the best candidate heap. Similarly, whenever a new slot in the puzzle is opened, pairings between that new slot and all best buddies currently in the pool are also added to the heap. A max heap allows the best candidate to be selected in $O(lg(n))$ time.

As best buddies are placed and puzzle slots are filled, items in the best buddy pool will become invalid. However, to reduce overhead, the best candidate heap is not cleaned after each placement. Rather, as items are popped off the heap, they are checked for validity. If the heap grows too large (currently set to one million elements), this thesis' implementation periodically cleans the entire heap.

# 6  Conclusions

Significant progress was made over the course of this semester. Major accomplishments include: a thorough review of existing solvers, implementing a jigsaw puzzle solver that improves on state-of-the art techniques, defining new solution quality metrics, and building visualization tools. Paikin & Tal's algorithm has limitations; the following subsection describes some of them in depth. The final subsection outlines tasks that will be performed in the next semester (i.e., CS298) to address these limitations.

## 6.1  Limitations of Paikin & Tal's Algorithm

Despite being state-of-the-art, Paikin & Tal's algorithm still has many short-comings. For example, the user must pass to the solver the expected number of puzzles. In many real world applications, this may not be known.

Another limitation is that when performing placement, Paikin & Tal's algorithm only considers a single side of the puzzle piece. While this may only lead to a handful of poor placements, it is known that a single suboptimal decision in a greedy algorithm can have serious, negative effects on the final result.

Paikin & Tal's algorithm also has two significant limitations when spawning new puzzles. First, whenever a new puzzle is spawned, Paikin & Tal's algorithm always use the same approach to choose the seed piece (see Section 5.1.2). This can lead to significant issues if multiple puzzles are spawned using seed pieces from the same ground-truth input. What is more, Paikin & Tal spawn a new puzzle when the mutual compatibility between the best candidate and associated the open slot falls below a preset threshold. This can cause the algorithm to prematurely spawn boards which further complicates the process of selecting the next seed piece since there are more pieces from which to choose.

## 6.2  Next Semester Tasks to Address the Limitations of Paikin & Tal's Algorithm

The focus of the next semester will be to use the existing infrastructure to improve on Paikin & Tal's solver. As explained in Section 5.1.3, only a single side is considered when placing each piece. Rather than using that single side, the improved placer will prioritize placing pieces that have multiple best buddies first.

Currently, Paikin & Tal's algorithm is a single pass algorithm. While this

has runtime advantages, it significantly restricts the ability of the algorithm to recover from poor decisions. By making multiple passes, it is expected that the algorithm selection of seeds for mixed bag puzzles would improve. The approach that will be implemented in the next semester will be to solve the puzzle first using Paikin & Tal's existing technique. The algorithm will then use the region growing segmentation algorithm previously used by Pomeranz *et al.* and described in [25] to select new $n$ seeds, where $n$ is the expected number of input puzzles; the solver would then be repeated with placement proceeding across multiple puzzles simultaneously. The algorithm would terminate after either a specific number of iterations or once the maximum best buddy accuracy has stabilized.

Another improvement to the solver that will be implemented next semester is to eliminate the requirement that the user specify the number of input puzzles. The algorithm will achieve this by running the solver multiple times with different expected number of input puzzles. The solver will then select the solution with the maximum best buddy accuracy.

Lastly, the image datasets used by existing solvers are well established and include [13, 18, 20]; however, these datasets are relatively small and consist of only about 100 total images. This thesis will widen that dataset to several hundred images. The images will be from image repositories including [21, 26] that use the Creative Commons License.

# Bibliography

[1] A. D. Williams, *Jigsaw Puzzles: An Illustrated History and Price Guide.* Wallace-Homestead Book Co., 1990.

[2] A. D. Williams, *The Jigsaw Puzzle: Piecing Together a History.* Berkley Books, 2004.

[3] H. Freeman and L. Gardner, "Apictorial jigsaw puzzles: The computer solution of a problem in pattern recognition," *IEEE Transactions on Electronic Computers*, vol. 13, pp. 118–127, 1964.

[4] T. Altman, "Solving the jigsaw puzzle problem in linear time," *Applied Artificial Intelligence*, vol. 3, pp. 453–462, Jan. 1990.

[5] E. D. Demaine and M. L. Demaine, "Jigsaw puzzles, edge matching, and polyomino packing: Connections and complexity," *Graphs and Combinatorics*, vol. 23 (Supplement), pp. 195–208, June 2007.

[6] B. J. Brown, C. Toler-Franklin, D. Nehab, M. Burns, D. Dobkin, A. Vlachopoulos, C. Doumas, S. Rusinkiewicz, and T. Weyrich, "A system for high-volume acquisition and matching of fresco fragments: Reassembling Theran wall paintings," *ACM Transactions on Graphics*, vol. 27, Aug. 2008.

[7] M. L. David Koller, "Computer-aided reconstruction and new matches in the Forma Urbis Romae," *Bullettino Della Commissione Archeologica Comunale di Roma*, vol. 2, pp. 103–125, 2006.

[8] S. L. Garfinkel, "Digital forensics research: The next 10 years," *Digital Investigation*, vol. 7, Aug. 2010.

[9] T. S. Cho, M. Butman, S. Avidan, and W. T. Freeman, "The patch transform and its applications to image editing," *Proceedings of the 2008 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.

[10] L. Zhu, Z. Zhou, and D. Hu, "Globally consistent reconstruction of ripped-up documents," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, pp. 1–13, 2008.

[11] W. Marande and G. Burger, "Mitochondrial DNA as a genomic jigsaw puzzle," *Science*, vol. 318, no. 5849, pp. 415–415, 2007.

[12] Y.-X. Zhao, M.-C. Su, Z.-L. Chou, and J. Lee, "A puzzle solver and its application in speech descrambling," in *Proceedings of the 2007 International Conference on Computer Engineering and Applications*, pp. 171–176, World Scientific and Engineering Academy and Society, 2007.

[13] T. S. Cho, S. Avidan, and W. T. Freeman, "A probabilistic image jigsaw puzzle solver," in *Proceedings of the 2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, CVPR '10, pp. 183–190, IEEE Computer Society, 2010.

[14] D. Pomeranz, M. Shemesh, and O. Ben-Shahar, "A fully automated greedy square jigsaw puzzle solver," in *Proceedings of the 2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, CVPR '11, pp. 9–16, IEEE Computer Society, 2011.

[15] G. Paikin and A. Tal, "Solving multiple square jigsaw puzzles with missing pieces," in *Proceedings of the 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, CVPR '15, IEEE Computer Society, 2015.

[16] A. C. Gallagher, "Jigsaw puzzles with pieces of unknown orientation," in *Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, CVPR '12, pp. 382–389, IEEE Computer Society, 2012.

[17] D. Sholomon, O. David, and N. S. Netanyahu, "A genetic algorithm-based solver for very large jigsaw puzzles," in *Proceedings of the 2013 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, CVPR '13, pp. 1767–1774, IEEE Computer Society, 2013.

[18] D. Sholomon, O. David, and N. S. Netanyahu, "Datasets of larger images and GA-based solver's results on these and other sets." `http://u.cs.biu.ac.il/~nathan/Jigsaw/`, 2013. (Accessed on 05/01/2016).

[19] K. Son, J. Hays, and D. B. Cooper, "Solving square jigsaw puzzles with loop constraints," in *Proceedings of the 2014 European Conference on Computer Vision (ECCV)*, pp. 32–46, Springer, 2014.

[20] D. Pomeranz, M. Shemesh, and O. Ben-Shahar, "Computational jigsaw puzzle solving." `https://www.cs.bgu.ac.il/~icvl/icvl_projects/automatic-jigsaw-puzzle-solving/`, 2011. (Accessed on 05/01/2016).

[21] A. Olmos and F. A. A. Kingdom, "McGill calibrated colour image database." `http://tabby.vision.mcgill.ca/`, 2005. (Accessed on 05/01/2016).

[22] P. S. Foundation, "Python language reference, version 2.7." `https://www.python.org/`.

[23] G. Bradski, "The OpenCV library," *Dr. Dobb's Journal of Software Tools*, Nov. 2000.

[24] S. v. d. Walt, S. C. Colbert, and G. Varoquaux, "The NumPy array: A structure for efficient numerical computation," *Computing in Science and Engineering*, vol. 13, pp. 22–30, Mar. 2011.

[25] I. Pitas, *Digital Image Processing Algorithms and Applications*. New York, NY, USA: John Wiley & Sons, Inc., 1st ed., 2000.

[26] H. Braxmeier and S. Steinberger, "Pixabay." `https://pixabay.com/`. (Accessed on 05/15/2016).