



(19) **United States**

(12) **Patent Application Publication**  
**Birkbeck et al.**

(10) **Pub. No.: US 2012/0069003 A1**

(43) **Pub. Date: Mar. 22, 2012**

(54) **METHOD AND SYSTEM FOR EVALUATION USING PROBABILISTIC BOOSTING TREES**

**Publication Classification**

(75) Inventors: **Neil Birkbeck**, Plainsboro, NJ (US); **Michal Sofka**, Franklin Park, NJ (US); **Shaohua Kevin Zhou**, Plainsboro, NJ (US)

(51) **Int. Cl.**  
**G06K 9/62** (2006.01)  
**G06T 15/00** (2011.01)

(73) Assignee: **Siemens Corporation**, Iselin, NJ (US)

(52) **U.S. Cl. .... 345/419; 382/160**

(21) Appl. No.: **13/228,505**

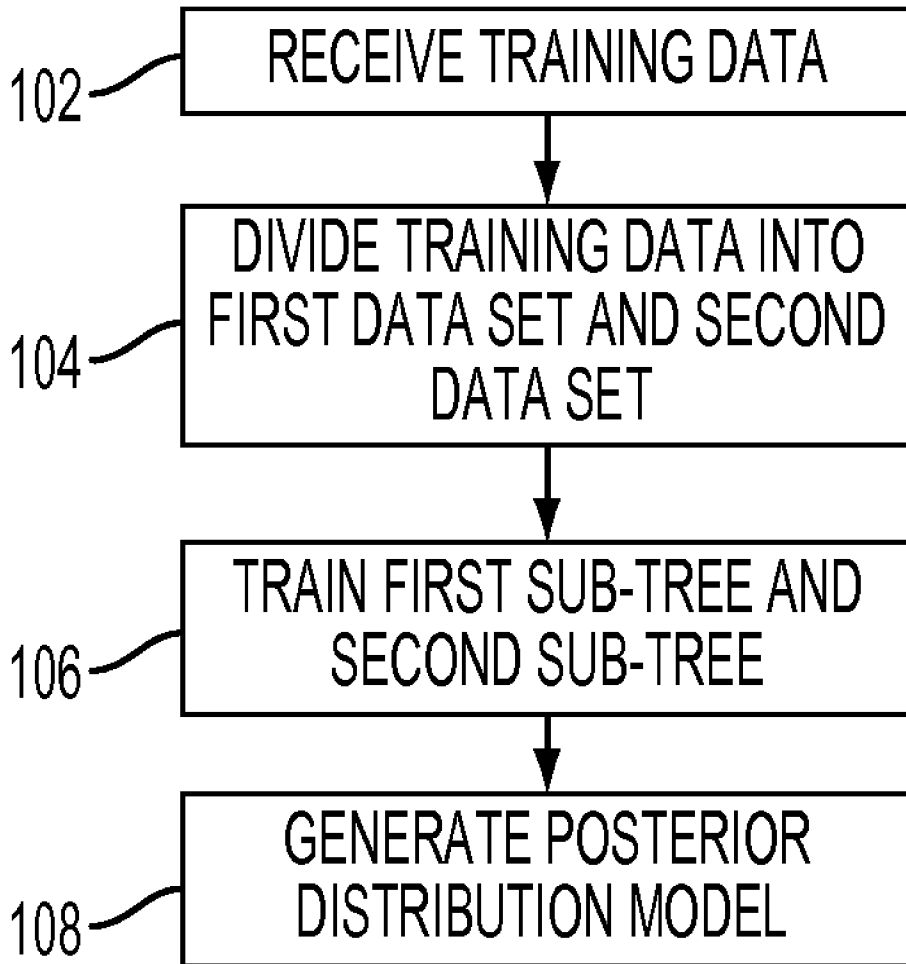
(57) **ABSTRACT**

(22) Filed: **Sep. 9, 2011**

A method and system for evaluating probabilistic boosting trees is disclosed. In an embodiment, input data is received at a graphics processing unit. A weighted empirical distribution associated with each node of the probabilistic boosting tree is determined using a stack implementation. The weighted empirical distribution associated with each node is added to a total posterior distribution value.

**Related U.S. Application Data**

(60) Provisional application No. 61/385,240, filed on Sep. 22, 2010, provisional application No. 61/424,715, filed on Dec. 20, 2010.



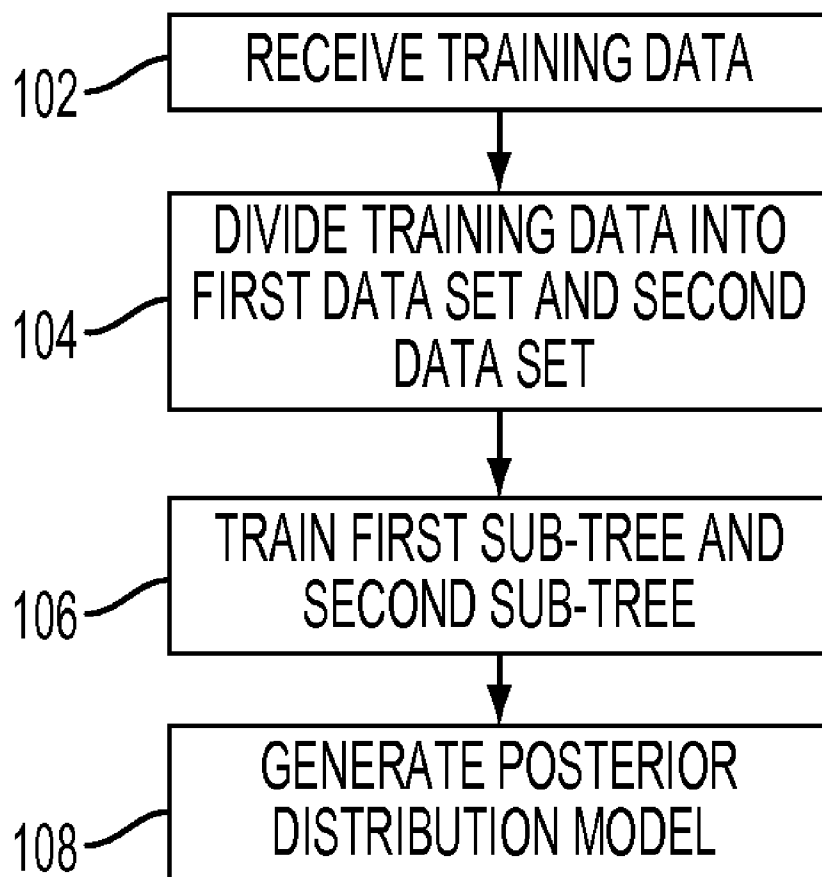


FIG. 1

---

**Algorithm 1:** PbtPosterior

---

**Data:**  $N$  node of the tree

**Data:**  $L = \text{left}(N)$ ,  $R = \text{right}(N)$

**Result:**  $p_N(+1|x)$  the posterior probability for tree ← 206  
 rooted at  $N$

**if** Leaf( $N$ ) **then** return  $q_N(+1)$ ;

$p = \hat{q}_N(+1|x)$ ; ← 202

**if**  $p > (1 - e_1)$  **then** return PbtPosterior( $R$ );

**else if**  $p < e_1$  **then** return PbtPosterior( $L$ );

**else if**  $p > 0.5 + e_2$  **then**

  └ return  $(1-p)q_L(+1) + p\text{PbtPosterior}(R)$

**else if**  $p < 0.5 - e_2$  **then**

  └ return  $(1-p)\text{PbtPosterior}(L) + p q_R(+1)$

**else**

  └ return  $(1-p)\text{PbtPosterior}(L) + p \text{PbtPosterior}(R)$

← 204

FIG. 2

---

**Algorithm 2:** PbtPosteriorStackBased

---

```

Data:  $N$  node of the tree
Result:  $p(+1|X)$  the posterior of the tree ← 310
 $S = \{\langle \text{root}, 1.0 \rangle\}$ ;
total ← 0
while  $\| S \| \neq 0$  do ← 306
     $\langle N, w \rangle = \text{pop}(S)$ ; ← 308
    if isLeaf( $N$ ) then
        total ← total +  $w * q_N(+1)$ ; ← 302
        continue;
     $p \leftarrow \hat{q}_N(+1, \chi)$ ;
    if  $p > (1 - e_1)$  then
         $S.\text{push}(\langle \text{right}(N), w \rangle)$ 
    else if  $p < e_1$  then
         $S.\text{push}(\langle \text{left}(N), w \rangle)$ 
    else if  $p > 0.5 + e_2$  then ← 312
         $L = \text{left}(N)$ ;
        total = total +  $w * (1 - p) q_L(+1)$ ;
         $S.\text{push}(\langle \text{right}(N), w * p \rangle)$ ;
    else if  $p < 0.5 - e_2$  then
         $R = \text{right}(N)$ ;
        total = total +  $w * p * q_R(+1)$ ;
         $S.\text{push}(\langle \text{left}(N), w * (1 - p) \rangle)$ 
    else
        // Descend down both nodes ;
         $S.\text{push}(\langle \text{right}(N), w * p \rangle)$ ;
         $S.\text{push}(\langle \text{left}(N), w * (1 - p) \rangle)$ 
return total;

```

---

FIG. 3




FIG. 4




FIG. 5

```

template<float feature_func(float3& pixel,
                            float nx,
                            float ny)>
__device__ float
pbt_evaluate_tree(const float3& pixel) {
    float4 stack [kMaxQueueSize];
    int numInStack = 1;

    // Insert root (@ 0,0) with weight 1.
    ....
    while(numInStack>0) {
        __numInQueue;
        float4 node = stack[numInQueue];
        float prob =
            pbt_eval_classifier
                (<feature_func>(pixel, node.x, node.y));
        if (prob > 1.0 - e1)
            ...
    }
}

```

FIG. 6




FIG. 7






FIG. 8

```
float4 numCubes = tex2D(texPBT,x,y);  
for (int i=0; i<numCubes.x; i++) {  
    float4 cubemin = tex2D(PBT,x,y+2*i+1);  
    float4 cubemax = tex2D(PBT,x,y+2*i+2);  
    val += cubemin.w *  
           haar_eval_cube(cubemin+pos,  
                           cubemax+pos);  
}  
return val;
```

FIG. 9

$$\begin{aligned}
 & I, \sqrt{I}, I^{0.333}, I^2, I^3, \max(0.0001, \log(I)), I_x, I_y, I_z \\
 & \qquad \qquad \qquad |\nabla I|^2 \\
 & f = \langle |\nabla I|, \mathbf{d} \rangle, |f|, \sqrt{|f|}, f^2, |f|^3, \log(\max(10^{-3}, |f|)) \\
 & \qquad \qquad \qquad \sqrt{|\nabla I|^2 - f^2} \\
 & \theta = \cos^{-1} \left( \left\langle \frac{\nabla I}{|\nabla I|}, \mathbf{d} \right\rangle \right), \theta^2, \theta^3, \log(\max(10^{-3}, \theta))
 \end{aligned}$$

FIG. 10

```
float sf eval(float4 p, float4 d, int type) {  
    float val;  
    int flags = flagTableShared[type];  
  
    // Extract data  
    if (flags & Norm) {  
        p.x = 0, val = dot4(p, p);  
    }  
    else if (flags & Projection) {  
        p.x = 0, val = dot4(d, p);  
    }  
    else if (flags & Angle) {  
        p.x = 0;  
        float l = max(1e-10, sqrt(dot4(p, p)));  
        val = acos(min(1.0,  
            max(-1.0, dot4(d, p)/l)));  
    }  
    else if (flags & NormalDistance) {  
        p.x = 0; float proj= dot4(d, p);  
        val = dot4(p, p)- proj * proj;  
    }  
    else val = dot4(component[type], p);  
  
    // Transform the pixel.  
    if (flags & Absolute) value = fabsf(val);  
    if (flags & Log) val = log(max(1e-3, val));  
    else val = pow(val, powTable[type]);  
    return val;  
}
```

The diagram shows three callout boxes on the right side of the code block, each with an arrow pointing to a specific line of code. Callout 1102 points to the line 'int flags = flagTableShared[type];'. Callout 1104 points to the line 'float l = max(1e-10, sqrt(dot4(p, p)));'. Callout 1106 points to the line 'else val = pow(val, powTable[type]);'.

FIG. 11




FIG. 12

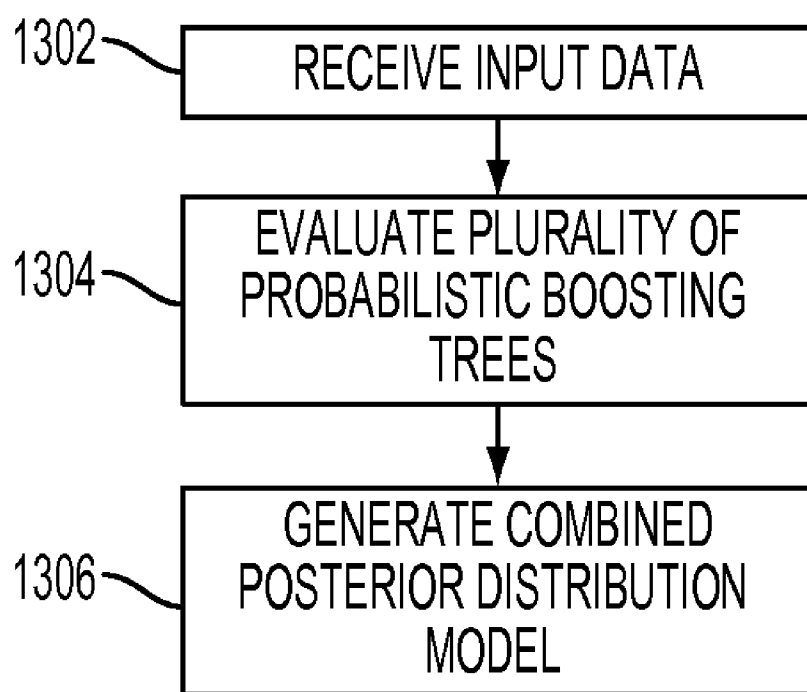


FIG. 13




FIG.14




FIG. 15



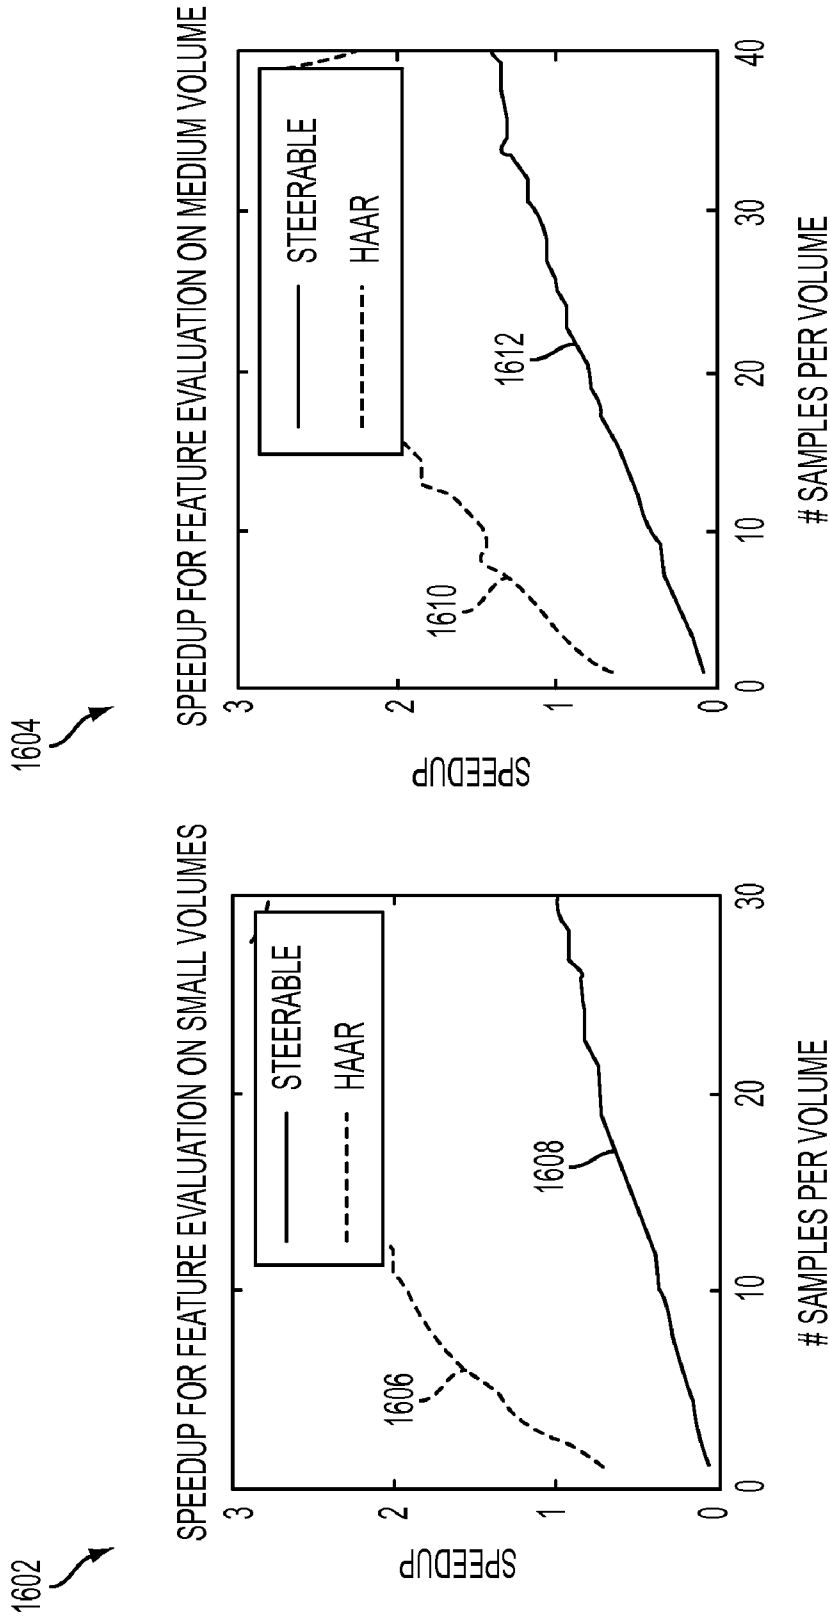


FIG. 16




FIG. 17

## METHOD AND SYSTEM FOR EVALUATION USING PROBABILISTIC BOOSTING TREES

[0001] This application claims the benefit of U.S. Provisional Application No. 61/385,240, filed Sep. 22, 2010, and U.S. Provisional Application No. 61/424,715, filed Dec. 20, 2010, the disclosures of which are herein incorporated by reference. This application is related to U.S. application Ser. No. 12/180,696, filed Jul. 28, 2008, and issued Apr. 20, 2010 as U.S. Pat. No. 7,702,596, and U.S. application Ser. No. 12/248,536, filed Oct. 9, 2008.

### BACKGROUND OF THE INVENTION

[0002] The present invention relates to evaluation of data using probabilistic boosting trees.

[0003] Discriminative classifiers are often a bottleneck in 3D pose detection routines. Tree-based classifiers, such as Probabilistic Boosting Trees (PBT) and Random Forests, are discriminative models used for vision-based classification and object detection. The classifier is typically evaluated at every pixel in an image, which can be inefficient. The PBT is a general type of decision tree that uses strong classifiers to make fuzzy decisions at internal nodes. Generally, using PBT requires multiple recursive calls, which slows down object detection.

[0004] Efficiency can be improved using hierarchical methods or cascades, but 3D medical applications and real-time applications require further efficiency improvements.

### BRIEF SUMMARY OF THE INVENTION

[0005] The present invention provides a method and system for evaluation of probabilistic boosting trees. In an embodiment, input data is received at a graphics processing unit. A weighted empirical distribution associated with each node of the probabilistic boosting tree is determined using a stack implementation. The weighted empirical distribution associated with each node is added to a total posterior distribution value.

[0006] In an embodiment, posterior distribution of a probabilistic boosting tree is determined by determining a weighted empirical distribution associated with each node of the probabilistic tree using a stack implementation, and adding the weighted empirical distribution associated with each node to a total posterior distribution value. A root node of the probabilistic boosting tree is pushed onto a stack and then is determined whether to descend thru a node in a left sub-tree or a node in a right sub-tree.

[0007] In an embodiment, a discriminative classifier of the root node is determined. Based on the discriminative classifier satisfying certain conditions, a left node, a right node, or both the left node and the right node of the root node are descended. Descending a node comprises determining a discriminative classifier of that node. If the discriminative classifier is a leaf node, the weighted empirical distribution of the node is added to the total posterior value. Otherwise, the node is pushed onto a stack and the left node, right node, or both the left and the right node of the node are descended.

[0008] In an embodiment, the probabilistic tree is evaluated using a parallel computing architecture. The parallel computing architecture may be Compute Unified Device Architecture (CUDA). The stack may be associated with one of a plurality of threads running in parallel.

[0009] In an embodiment, a method and system for evaluating a forest of probabilistic boosting trees is disclosed. Input data is received at a graphics processing unit. The plurality of probabilistic boosting trees is evaluated using a stack implementation. A combined posterior distribution based on a posterior distribution of each of the plurality of probabilistic boosting trees is generated.

[0010] These and other advantages of the invention will be apparent to those of ordinary skill in the art by reference to the following detailed description and the accompanying drawings.

### BRIEF DESCRIPTION OF THE DRAWINGS

[0011] FIG. 1 illustrates a method for training a probabilistic boosting tree (PBT), in accordance with an embodiment of the present invention;

[0012] FIG. 2 shows pseudocode for determining the posterior value of a node of a PBT using recursion;

[0013] FIG. 3 shows pseudocode for determining the posterior value of a node of a PBT using a stack implementation according to an embodiment of the present invention;

[0014] FIG. 4 illustrates a method for determining the posterior distribution of a node of a probabilistic boosting tree, in accordance with an embodiment of the present invention;

[0015] FIG. 5 illustrates an exemplary PBT data structure that is arranged into a texture image;

[0016] FIG. 6 illustrates an exemplary Compute Unified Data Architecture (CUDA) stack-based implementation of evaluating a PBT, in accordance with an embodiment of the present invention;

[0017] FIG. 7 illustrates exemplary cubes showing Haar feature types, in accordance with an embodiment of the present invention;

[0018] FIG. 8 illustrates a histogram classifier and feature data packed into a single column of a texture image, in accordance with an embodiment of the present invention;

[0019] FIG. 9 shows an exemplary CUDA algorithm for implementing feature computation, in accordance with an embodiment of the present invention;

[0020] FIG. 10 shows a table of specific steerable features used, in accordance with an embodiment of the present invention;

[0021] FIG. 11 shows an algorithm for implementation in CUDA for steerable feature evaluation, in accordance with an embodiment of the present invention;

[0022] FIG. 12 shows a CUDA implementation integrated into an existing software library, in accordance with an embodiment of the present invention;

[0023] FIG. 13 illustrates a method for evaluating a forest of probabilistic boosting trees, in accordance with an embodiment of the present invention;

[0024] FIG. 14 illustrates a graph showing execution times when threads evaluate different classifiers and when thread-blocks evaluate different classifiers, in accordance with an embodiment of the present invention;

[0025] FIG. 15 illustrates total detection time, central processing unit (CPU) execution time, and wait time for a graphics processing unit (GPU), in accordance with an embodiment of the present invention;

[0026] FIG. 16 illustrates speedup comparisons for small volumes and medium volumes, in accordance with an embodiment of the present invention; and

[0027] FIG. 17 is a high level block diagram of a computer capable of implementing the present invention.

#### DETAILED DESCRIPTION

[0028] The present invention is directed to a method and system for evaluation of probabilistic boosting trees. Embodiments of the present invention are described herein to give a visual understanding of the method for evaluation of probabilistic boosting trees. A digital image is often composed of digital representations of one or more objects (or shapes). The digital representation of an object is often described herein in terms of identifying and manipulating the objects. Such manipulations are virtual manipulations accomplished in the memory or other circuitry/hardware of a computer system. Accordingly, it is to be understood that embodiments of the present invention may be performed within a computer system using data stored within the computer system.

[0029] FIG. 1 illustrates a method for training a probabilistic boosting tree (PBT), according to an embodiment of the present invention. FIG. 1 more specifically describes accelerated training of a probabilistic boosting tree using a graphics processing unit (GPU).

[0030] At step 102, training data is received. The training data may be annotated training data including images having annotations representing locations of landmarks of objects in the image. The training data may consist of images having annotations representing locations of landmarks of objects in the image. Training data can be medical image data such as computerized tomography (CT), magnetic resonance imaging (MRI), X-ray, or Ultrasound image data.

[0031] A PBT models the posterior distribution of a data set. In training a PBT, the aim is to enable the use of the PBT as a discriminative model for classification and detection of an object in image data. A PBT is a binary decision tree with a fuzzy decision taken at each internal node of the tree depending on the output of the node's strong classifiers. The posterior value of any given individual node is determined by combining the posterior values of its child nodes using a weighted distribution. Weights used in the determination are determined by evaluating a node's strong (learned) classifier,  $\hat{q}_N(y|x)$ .

[0032] Each node N contains a strong classifier,  $\hat{q}(y|x)$ , and the empirical distribution of its leaf nodes  $q_N(y)$ , where  $y \in \{-1, +1\}$ , and  $x$  is an input point. The strong classifier can be any classifier that uses a problem specific feature. For example, an AdaBoost classifier can be used, which combines several binary weak classifiers to produce a strong estimate.

[0033] At step 104, the training data is divided into a first data set and a second data set by a classifier. The classifier that divides the training data is a strong (learned) classifier associated with the top parent node of the PBT.

[0034] At step 106, a first sub-tree and a second sub-tree are trained. The first data set is used to train a first sub-tree of the parent node and the second data set is used to train a second sub-tree of the parent node. A classifier at each node of the PBT is trained using a feature matrix.

[0035] During training of a PBT, the PBT is constructed using a set of positive and negative examples, which may originate from different 3D images. Each node of the PBT includes a strong classifier, such as the AdaBoost classifier, which may be trained for that particular node. A feature matrix is computed (e.g., every possible feature is evaluated

for every input sample), and then each of the weak classifiers that constitute the strong classifier is trained using the feature matrix.

[0036] Using the feature matrix, the weak classifiers of the strong classifier may be trained sequentially by choosing each weak classifier using a greedy algorithm by considering how well each feature acts as a classifier. The algorithm for choosing the weak classifier during training includes three passes. First, feature bounds must be computed. Afterwards, a feature value is mapped to each sample within a histogram bin. Then, the histogram bins are incremented. Each sample may then be sampled with the trained classifier to compute an error. All of the aforementioned steps may be performed on a GPU. The operations for a classifier are independent, so each weak classifier can be trained by a different thread. The feature matrix column indexes the feature type, while the row indexes the sample. Each thread processes a single column.

[0037] The feature matrix is stored as a single component floating-point texture. The input is then split up into several textures of a maximum width and processed in chunks. Through CUDA texture limits, the restriction is a maximum width of  $32768/\text{sizeof(float)}=8192$ . The feature data stays on the GPU for each of the weak training classifier passes. For each pass, the weights of the samples are updated.

[0038] Returning to FIG. 1, at step 108, a trained posterior distribution model of the PBT is generated based on the feature matrices of each node. The posterior distribution model represents a set of classifications that may be used for object classification and object detection. The set of classifications from the posterior distribution model of the PBT may be used thereafter by a detector to perform detection on image data. For example, the trained tree can be used to estimate the posterior probability of unseen data using a stack implementation on a GPU, as discussed in further detail in the following paragraphs.

[0039] The data parallel nature of evaluation using the PBT means that it is advantageous to utilize the computation power of a GPU instead of a central processing unit (CPU). Thus, in order to implement evaluation of a PBT on a GPU, multiple recursive calls, which are used in a traditional evaluation of a PBT may be replaced using a stack-based implementation. Cached texture memory is used to represent features and tree data structures. Since GPUs do not support recursive calls, the stack-based implementation is necessary to remove recursion.

[0040] For comparison, FIG. 2 shows pseudocode for determining the posterior value of a node of a PBT using recursion. FIG. 3 shows pseudocode for determining the posterior value of a node of a PBT using a stack implementation according to an embodiment of the present invention.

[0041] The posterior probability of a node, is recursively computed as a weighted combination of the posterior probabilities of the child nodes, referenced by reference numeral 204. The weight factor of this combination is determined by evaluating the node's strong classifier, referenced by reference numeral 202. The recursion terminates at leaf nodes of the tree (i.e., nodes with no child nodes), which simply return their empirical distribution  $q_N(+1)$ . The final result, or posterior distribution for each node is a sum of weighted empirical distributions of the child nodes, represented by a result referenced by reference numeral 206. The total weight given to any node's empirical distribution is the product of the weights associated with the path from the root to the node.

**[0042]** The algorithm for using a stack implementation is shown in FIG. 3. As illustrated in FIG. 3, a stack holds a list of to-be traversed nodes and their respective weights. When a leaf node is visited, the empirical distribution of that leaf node with its weight is added to a total empirical distribution of the root node, represented by reference numeral 302.

**[0043]** The root node and corresponding weight is first pushed into a stack. Each node in the stack is then processed sequentially, represented by reference numeral 306. If the node is a leaf node, the empirical distribution of the node is weighted by the corresponding weight and accumulated into the total posterior of the tree, represented by reference numeral 302. If a particular node has child nodes, then either the left or right child nodes are placed into the stack with weights that are the product of the parent node's weight and a factor dependent on the value of the parent node's strong classifier, represented by reference numeral 312. The empirical distribution of child nodes that are not pushed on the stack is weighted and accumulated to the total posterior distribution of the tree. The weight is proportional to the weight of the parent node and the value of the strong classifier evaluated at the parent. The sum of all of these products represents the total posterior distribution model for the PBT, represented by reference numeral 310.

**[0044]** FIG. 4 illustrates a method for determining the posterior distribution of a node of a probabilistic boosting tree using a GPU, in accordance with an embodiment of the present invention.

**[0045]** At step 402, input data is received at a GPU. Input data may be a 2D image or a 3D volume received from an image acquisition device. The 3D volume can be a C-arm CT volume, computed tomography (CT) volume, magnetic resonance imaging (MRI) volume, etc. The 3D volume can be received from an image acquisition device, such as a C-arm image acquisition system, or can be a previously stored volume loaded from memory or storage of a computer system, or some other computer readable medium.

**[0046]** At step 404, a weighted empirical distribution associated with each node of the probabilistic boosting tree is determined using a stack implementation. The posterior probability  $\hat{p}(y|x)$  is determined using the principles described above with a stack implementation. Specifically, a root node of the probabilistic boosting tree is first pushed onto a stack. Then a determination is made whether to descend down the PBT thru a node in a left sub-tree or thru a node in a right sub-tree.

**[0047]** In determining whether to descend down a node, a discriminative classifier of the root node is first determined. If the discriminative classifier of the root node satisfies a first condition signifying that a left node must be descended, then the method proceeds to descend down a left node in the left sub-tree. If the discriminative classifier of the root node satisfies a second condition signifying that a right node must be descended, then the method proceeds to descend down a right node in the left sub-tree. If the discriminative classifier of the root node satisfies a third condition, then the method descends down both the aforementioned left node and the aforementioned right node.

**[0048]** After descending into a node, the discriminative classifier of that node is determined. If the node is determined to be a leaf node, then the weighted empirical distribution of that node is added to a total posterior distribution value. If the node satisfies the first condition, then the node will be pushed onto the stack, and the left child node of the node will be

descended. If the node satisfies the second condition, then the node will be pushed onto the stack, and the right child node of the node will be descended. If the node satisfies the third condition, then the node is pushed onto the stack, and both the left and right child nodes are descended.

**[0049]** At step 406, the weighted empirical distribution of each node of the probabilistic boosting tree is added to the total posterior distribution value. The total posterior distribution represents a posterior distribution model that can be used in the detection of objects or landmarks in 2D images or 3D volumes.

**[0050]** In an advantageous embodiment, the above described method for evaluating a PBT can be implemented using a Compute Unified Device Architecture (CUDA) programming model. In the CUDA programming model, a parallel problem is decomposed into a grid of treaded blocks, with each block containing many threads. Blocks are assigned to the GPU's streaming multi-processors, which breaks down the block and schedules it in groups of 32 threads. In detection or classification, the grid of thread blocks will overly the input space (e.g., all pixels in the volume), and each thread will evaluate the classifier for a different pixel in the volume or a different orientation/scale hypothesis.

**[0051]** The programming model exposes the memory architecture, which includes a 16 kb low-latency shared memory (accessible from threads within the same block), high-latency local memory (per-thread), global memories (accessible by all threads), and cached global accesses through texture memory. Designing the algorithm to use these memory types appropriately is key to efficiency.

**[0052]** On a GPU implementation of evaluating a probabilistic boosting tree, several threads will descend the tree in parallel. Thus, neighboring threads may access the tree data structure in different regions, so placing the tree structure in global memory may cause slow uncoalesced memory access. However, since the tree data structure is too large to fit entirely in a 16 kb shared data memory, the tree data structure may be packed into a texture image. Instead of node pointers, 2D location indices that reference the node's 2D region are placed inside the texture image.

**[0053]** For example, the node data may contain AdaBoost classifiers, which include several weak classifiers. Each node contains either a simple threshold classifier or a histogram classifier whose decision is made on the basis of a single feature. Thus, the PBT node data that is placed inside the texture image must store the sum of alpha (or weighting) values, the number of weak classifiers, and the data for each of the classifiers. The weak classifiers are stored in adjacent columns of the texture image, and each column containing an associated weight  $\alpha_i$ . FIG. 5 illustrates an exemplary PBT data structure that is packed into a texture image. All weak classifiers are arranged along the columns 502a-502g, and each PBT node stores 2D texture coordinate indices to their child nodes, represented by 504a for node 0, 504b for node 1, and 504c for node 3.

**[0054]** FIG. 6 illustrates an exemplary CUDA stack-based implementation of evaluating a PBT using the algorithm shown in FIG. 3.

**[0055]** In order to introduce new feature types into a PBT evaluation, the CUDA PBT evaluation function shown in FIG. 6 is represented by feature\_func, which allows new

feature types to be added. Both 3D sample box information and 2D texture location of the feature data are passed to this feature evaluation function.

**[0056]** Marginal space learning (MSL) is an efficient method used to perform object localization. In MSL, a large 9-dimensional search space for a similarity transformation is decomposed into a series of smaller search spaces from translation to full similarity. During detection, a discriminative classifier is evaluated on pixels in the volume to determine a small set (e.g., 100-1000) of candidate positions. Orientation is detected by evaluation each of these candidate boxes with a series of hypothesis orientations and similarly for scale. The final list of candidates are aggregated to obtain a single pose estimate.

**[0057]** A Hierarchical Detection Network (HDN) to extend MSL to multiple structure detection may be used. The HDN decomposes a joint multi-object detection problem into sequential detections with spatial priors used for object prediction. The hierarchical detection of multiple structures is broken down into a network of nodes for detection position, orientation, and scale, for each structure. Spatial dependencies are also represented as arcs in this detection network.

**[0058]** Although MSL is an efficient framework for detection, the main computational burden comes from evaluating a discriminative classifier in each phase of the detection. In position detection, the classifier is either evaluated by looping over the 3D locations in the image (if the node has no ancestor) or by looping over a set of candidate locations provided by a preceding node. For orientation and scale detection, the classifier is evaluated over each of the input candidate locations using each of the possible orientation (or scale) hypothesis. The same classifier is evaluated on different data, which makes it an ideal detection algorithm for the data parallel power of a GPU. GPU accelerated implementations are possible, using a PBT as the classifier for MSL.

**[0059]** During detection, a PBT tree is evaluated with the context surrounding a specific location (and pose). Features may be computed on the fly from an input image. In MSL, position detection utilizes Haar features, and the subsequent orientation and scale detection utilize efficient steerable features that sample the image intensity and gradient in a volume using the candidate pose of the box.

**[0060]** Haar features are weighted combinations of the sums of cubic regions of an image. These sums are efficiently calculated using the integral image. The features use no more than four cubes. FIG. 7 illustrates exemplary cubes showing Haar feature types. Cubes 702 indicate negative weights.

**[0061]** The features possible for a given location consists of various translated and scaled versions of these boxes. Each possible feature can be described by the size, weights, and locations of these cubes (relative to a testing point). FIG. 8 illustrates a histogram classifier and feature data packed into a single column of a texture image. Column 802 shows a layout for Haar features with a histogram classifier, where each cube must be stored in the texture. The histogram contains 64-bins and are packed into two float components that are interpreted as integers. Column 804 shows a steerable feature with a threshold classifier requires only the threshold for the classifier and a single texel to store the feature information.

**[0062]** Evaluation on the GPU then proceeds by looking up cubes for the feature, evaluating the integral image, and combining the results. FIG. 9 shows an exemplary CUDA algorithm for implementing the feature computation described

above. To avoid conditionals for boundary testing, the boundary of the integral image is padded with extra planes filled with zeros.

**[0063]** Parallel work efficient algorithms for computing the integral image in a 3D case on a GPU requires mimicking a typical CPU implementation of the 3D integral image computations that proceeds in 3 passes. While performing the padding of the integral images with zeros, as stated above, the accumulation of the image in the x-direction is also performed on the CPU. Once this padded image is transferred to the GPU, a grid of thread blocks on the x-z plane are used to accumulate in the y-direction. The same step is taken for the z direction, where the threads are on the x-y plane.

**[0064]** FIG. 10 shows a table of specific steerable features used.  $I$  represents the image intensity,  $\nabla I = \{I_x, I_y, I_z\}$ , and  $d$  is an input sample direction. For a specific sample location, the features are transformations on either a gray value, the gradient, or the projection of the gradient onto an input direction. In the case of oriented box detection, the input direction comes from the x-axis of the box, and a discrete sampling of 3D locations is considered for these features within the detected box. Each of these features can be calculated on one of several image resolutions of an image pyramid. Therefore, each feature is completely described by its position in the 3D sampling pattern (3 integers), the type of feature (1 integer), and the discrete image scale it is computed at (1 integer). By packing the type and scale into a single value (e.g.,  $\text{scale} * 32 + \text{type}$ ), this data can be packed into as few as 1 RGBA pixels of a texture image.

**[0065]** To avoid a large conditional or switch statement which treats all of the basic feature types as unique elements, a series of common extractions is extracted. Extraction takes place in two phases, data extraction, and modification. During data extraction, the initial datum to be extracted for modification is extracted. This requires a group of conditionals corresponding to those shown in FIG. 10. After data extraction, three modifiers can be applied: a power, an absolute value, or a logarithm. Both the type of data to be extracted and the subsequent modifier are stored as flags in a table of length 25.

**[0066]** FIG. 11 shows an algorithm for implementation in CUDA for steerable feature evaluation. In the algorithm shown in FIG. 11,  $p = \{I_x, I_y, I_z\}$ , and input direction,  $d = \{0, d_x, d_y, d_z\}$ , defined at reference number 1102. If threads in a block diverge, the path each thread takes from the others is short (and not the entire evaluation of the feature). The data extraction and storage of flags is represented by reference numeral 1104, and the pixel transformation is represented by reference numeral 1106.

**[0067]** As texture arrays cannot be dynamically indexed in CUDA, the multiple scales of the image are packed into a single texture. An offset table stores the lower left corner of each of the resolutions in this texture.

**[0068]** FIG. 12 shows a CUDA implementation integrated into an existing software library. The GPU implementation is mostly isolated from the detection package, with the interface being exposed by deriving from the SrcObjectPositionDetector 1202 and SrcBoxDetector 1204. The detector may utilize a classifier, such as a PBT, or a random Forest™, which is discussed below.

**[0069]** In order to interface with CUDA, an object position detector, SrcCudaPositionDetector 1206 replaces the main classifier loop within the CPU position detector. This loop iterates either over entire 3D volumes or a set of previously

generated position candidates. The evaluation of the posterior distribution from the underlying classifier is scheduled on the GPU over these locations. The results are then sorted on the GPU and only the required number of candidate locations are transferred back to main memory.

**[0070]** The case is similar for orientation and scale detection which is handled by SrcBoxDetector **1204**. A common interface is derived to replace routines responsible for detecting orientation and scale. In these instances, the input is a set of  $m$  candidate box locations for which there are also a set of  $n$  hypothesis orientations (or scales). During orientation detection, the 3D axes corresponding to the hypothesis orientations are computed and transferred to the GPU as a texture. The CUDA kernel for orientation detection, represented by reference number **1208**, uses  $m*n$  threads to evaluate the PBT for all possible candidate boxes using each of the potential orientations. The situation is similar for scale detection although the hypothesis scales are transferred to the GPU instead of the orientations.

**[0071]** FIG. **13** illustrates a method for evaluating a forest of probabilistic boosting trees, in accordance with an embodiment of the present invention. Decision trees, such as the probabilistic boosting tree may be used in a forest algorithm, such as the random forest algorithm. Evaluation of a forest algorithm is similar to that of the PBT. For example, the forest may comprise a plurality of PBTs, and the posterior distribution output of the forest is a combination of the output of each PBT. At step **1302**, input data is received at a GPU.

**[0072]** At step **1304**, a plurality of PBTs are evaluated using a stack implementation. The stack implementation may be the one described above with respect to the method of FIG. **4**.

**[0073]** At step **1306**, a combined posterior distribution model of each of the PBTs is generated. The posterior distribution model may be represented by

$$\frac{1}{T} \sum_{t=1}^T p^{(t)}(+1 | x),$$

where  $T$  represents the number of trees. The posterior distribution model represents a set of classifications that may be used for object classification and object detection. The set of classifications from the posterior distribution model of the PBT may be used thereafter by a detector to perform detection on image data.

**[0074]** Experiments using the methods described herein were performed. Specifically, the experiments were performed on two hardware setups: Config1) a 1 GB Nvidia 9800 GT with an Intel Core(2) Duo with 3.5 GB of RAM; and Config2) A 1.5 GB Nvidia 480 GTX with an Intel Core Quad with 16 GB of RAM. Results for evaluating PBT are presented, namely feature evaluation, weak classifier evaluation, strong classifier evaluation, and full classifier evaluation. The GPU-PBT implementation was compared to similar implementations of random forests in terms of accuracy and improvements in efficiency.

**[0075]** Unless otherwise noted, all timing results compare the GPU version to an OpenMP implementation, where there are 2 threads for Config1 and 8 threads for Config2. In all examples, the following datasets were used:

**[0076]** 1) A Semantic Indexing (SI) data set of ultrasound images of fetal heads used in a semantic indexing project. 990 volumes with 1 mm resolution and dimensions 143×90×110

to 231×161×208 were used for training. 215 were used for testing. The HDN network is built on 6 structures, including, e.g., the Corpus Callosum (CC) and Cerebellum (CER). The HDN network encodes spatial relationships between structures, and performs detection on volume resolutions of 4 mm, 2 mm, and 1 mm.

**[0077]** 2) A data set including CT scans of 247 hips with image dimensions ranging from 70×70×89 to 125×125×334. The HDN network consists of position detection for the hip and is performed on resolutions of 16 mm, 8 mm, and 4 mm.

**[0078]** As discussed above, Haar features use the integral image, which can be calculated on the GPU. Table 1, shown below, summarizes the speedups that are possible with GPU accelerated integral image calculation. Table 1 shows speedups in the table comparing the GPU accelerated version to a multi-threaded OpenMP version. For each configuration, the right column shows a time when reading back an image is not required. The timings include the time to pad the image and transfer the results on the GPU. For a volume of 256<sup>3</sup>, a five-fold speedup is attained on both systems when the integral image does not need to be read back (as is the case during detection).

TABLE 1

| size             | Config1 |          | Config2 |          |
|------------------|---------|----------|---------|----------|
|                  | w/ read | w/o read | w/ read | w/o read |
| 64 <sup>3</sup>  | 0.79    | 1.58     | 1.49    | 2.49     |
| 128 <sup>3</sup> | 3.47    | 5.07     | 2.15    | 3.44     |
| 256 <sup>3</sup> | 4.04    | 5.57     | 3.11    | 5.20     |
| 400 <sup>3</sup> | 3.89    | 5.35     | 2.95    | 4.86     |

**[0079]** Table 2, shown below, shows speedups gained throughout the evaluation process for both steerable and Haar features. The Feature row considers only evaluating the features. In these tests, the same feature was evaluated on all pixels in a volume. It is apparent that the steerable features give a better speedup, possibly because their implementation involves more computation, whereas the Haar features mostly involve texture lookups and only a few additions (e.g. bandwidth limited). As expected, these speedups propagate into the weak classifier evaluation, which only evaluates a single features and looks up a bit in a histogram bin. However, the strong classifier evaluates and combines the results of roughly 40 weak classifiers. In this case, the two feature types start to behave more similarly, where a 26× speedup is achieved with Config1. On Config2, a strong classifier with steerable features is still 1.5 times faster than the Haar features.

TABLE 2

|         | Config1 |           | Config2 |           |
|---------|---------|-----------|---------|-----------|
|         | Haar    | Steerable | Haar    | Steerable |
| Feature | 3.2x    | 15x       | 4.9x    | 21.5x     |
| Weak    | 10.5x   | 25.7x     | 38.6x   | 112.4x    |
| Strong  | 26x     | 26x       | 22.3x   | 36.5x     |
| PBT     | 4.7x    | 7x        | 11.5x   | 12.18x    |

**[0080]** When it comes to evaluating an entire tree (the PBT row of Table 2), there is only a 4.7× speedup on Config 1 and 12× on Config 2. The speedup of the strong classifier limits the attainable speedup on the PBT tree. An explanation for the different speedup times may be due to cache consistency.

When evaluating a single strong classifier on every pixel, each pixel will access the same relative location as its neighboring pixel. As the tree is descended, neighboring pixels may take different paths down the tree, meaning that they will be evaluating different strong classifiers, which in turn use different features. For both Haar features and steerable features, this means that neighboring pixels will sample from different relative locations in the texture, and different control points may be taken.

[0081] FIG. 14 illustrates a graph showing execution times when threads evaluate different classifiers and when thread-blocks evaluate different classifiers. As nodes in a PBT are descended, threads in the same thread block will evaluate different nodes. The thread curve 1402 represents when threads evaluate different classifiers and the block curve 1404 represents when different thread-blocks evaluate the different classifiers. Thread curve 1402 shows that there is a longer execution time when evaluating different classifiers with threads. FIG. 14 also shows that evaluating the classifier by blocks do not affect the run-times. In contrast, when evaluating by threads, when the tree is deeper, the evaluation is 5x slower.

[0082] The results of experiments shown above illustrate the potential benefit of using GPU accelerated PBT evaluation. However, these experiments were carried out in ideal execution environments where there is enough work for the GPU and the same program is being run on all voxels in an image. In the following paragraphs, the PBT is evaluated in the context of the already optimized hierarchical detection system.

[0083] Table 3 shows timings and speedups for the detection of 6 structures in the SI data set. These times were averaged over 201 volumes. Most of the execution is in orientation and scale detection (which uses steerable features). On Config1, the overall speedup of 4.8x is slightly lower than the results from the steerable feature PBT evaluation shown in Table 2. Similarly, with Config2, the 9.73x speedup for hierarchical detection was slightly lower than the results from the steerable feature PBT evaluation in Table 2. This is due to hierarchical detection having regions that are inherently serial, which must be performed on the host CPU (e.g., results are read back, candidates are pruned, and detection results are written to files). Furthermore, some of the phases of detection only need to evaluate the PBT on as few as 1000s of elements, meaning that the GPU cannot be fully utilized.

TABLE 3

|     | Config1 |      |       | Config2 |      |       |
|-----|---------|------|-------|---------|------|-------|
|     | CPU     | GPU  | S-up  | CPU     | GPU  | S-up  |
| LV  | 13.9    | 2.63 | 5.28x | 7.22    | 0.57 | 12.7x |
| CC  | 1.42    | 0.33 | 4.21x | 0.88    | 0.12 | 7.33x |
| CP  | 2.90    | 0.63 | 4.60x | 1.74    | 0.15 | 11.5x |
| HC  | 1.11    | 0.32 | 3.46x | 0.77    | 0.15 | 5.20x |
| CER | 6.48    | 1.83 | 3.52x | 4.40    | 0.67 | 6.54x |
| CSP | 7.57    | 1.20 | 6.30x | 3.44    | 0.24 | 14.6x |
| Sum | 33.4    | 6.96 | 4.80x | 18.5    | 1.90 | 9.73x |

[0084] While code is executing on the CPU, it is possible to concurrently execute some of the code on the CPU. This trade-off was also investigated via experiments on the orientation detector, which is the most time consuming component of hierarchical detection.

[0085] In this experiment, a portion of the work was selected and scheduled on the GPU, and then the remaining work was given to the CPU. Once the CPU finished work, execution waited for the preceding GPU call to complete. FIG. 15 illustrates the total detection time, the CPU execution time, and the wait time for the GPU against varying proportions of work assigned to the GPU. The total amount of time for detection is represented by line 1502, the total CPU time is represented by line 1504, and the GPU sync time (wait time for the GPU) is represented by line 1506. The GPU work is started asynchronously after the CPU work is initialized. A non-zero GPU sync time indicates that the GPU does not have enough work. The optimal ratio, from FIG. 15, for the detection line appears to be roughly 0.85.

[0086] For Config1, the optimal proportion is expected to be in accordance with speedups obtained in the PBT evaluation. For steerable features, the GPU was 7x faster than the dual core for PBT tree evaluation, so the optimal proportion of work on the GPU is expected to be around 0.875, as confirmed by FIG. 15. In this case, a 10-15% improvement can be expected in the runtime by utilizing the CPU at the same time as the GPU.

[0087] Using Config 2, the speedup for orientation evaluation was above 12x, meaning that utilizing the CPU at the same time gives about 5% improvement.

[0088] In a further set of experiments, the accuracy and speed-ups attained on the GPU PBT implementation vs. the random Forest™ implementation was compared. A first experiment compares the speed-up of the GPU implementation over a single thread CPU implementation in a classification context. For this example, a random Forest™ of 10 trees of a depth 10 on every pixel of a 96x97x175 3D volume was evaluated. The speedup on Config1 is 26x and Config2 is 114x. Table 4 illustrates the comparison of timings and speed-ups for the GPU implementation of a random Forest™.

TABLE 4

|         | CPU time [s] | GPU time [s] | Speed-up |
|---------|--------------|--------------|----------|
| Config1 | 29.3         | 1.1          | 26.6x    |
| Config2 | 33.7         | 0.294        | 114x     |

[0089] A second experiment compares the results on hierarchical position, orientation, and scale detection using the CER detection from the SI dataset. In this case, the classifier evaluation makes up a large portion of the total run-time. The PBT was trained to a maximum depth of 6, with 40 weak classifiers, and the forest implementation had a maximum depth of 10 with 50 trees. In terms of total number of weak classifiers, a descent through the forest implementation evaluates more weak classifiers (maximum of 500), whereas the PBT has at most 270 classifiers. However, the PBT uses histogram classifiers, while the forest implementation uses a threshold-based classifier. When testing on the training data, both models achieved similar accuracy. Table 5 illustrates the detection accuracy when testing on training data for the CER data structure on the SI hierarchical data set.

TABLE 5

| Classifier | Center (mm) | Line (mm) |
|------------|-------------|-----------|
| RF         | 1.82        | 2.87      |
| RF(cuda)   | 2.00        | 3.09      |



TABLE 5-continued

| Classifier | Center (mm) | Line (mm) |
|------------|-------------|-----------|
| PBT        | 1.81        | 2.83      |
| PBT(cuda)  | 1.91        | 3.05      |

[0090] Two measurements were used as shown in Table 5: the distance between box centers (Center (mm)), and the maximum distance between end-points of an axial line on the recovered box and the ground truth (Line (mm)). From Table 5, it is apparent that both models behave similarly, and that in each case, the CUDA implementation gives slightly worse results due to a different feature evaluation code.

[0091] On both systems, the CUDA accelerated implementation of the forest implementation is more than 10x faster than the multi-core CPU implementation. However, the CPU version of the RF is 2x slower than the PBT with similar accuracy on Config 1. The forest implementation of CUDA achieves a better speed-up compared to the PBT, making the GPU version of the forest implementation 2% slower than the GPU-PBT on Config 1.

[0092] Table 6 shows the timing results for the CER detection on the SI data for the two hardware configurations and the PBT and forest models that achieve similar accuracy. The forest implementation is slower on the CPU, but the GPU implementation achieves a greater speedup than the PBT. However, the GPU implementation on the forest is still slower than the GPU implementation of the PBT.

TABLE 6

|       | RF model (depth 10, num trees 50) |        |       |         |        |       |
|-------|-----------------------------------|--------|-------|---------|--------|-------|
|       | Config1                           |        |       | Config2 |        |       |
|       | Cpu(s)                            | Gpu(s) | S-up  | Cpu(s)  | Gpu(s) | S-up  |
| 1 mm  | 7.57                              | 0.68   | 11.16 | 4.65    | 0.31   | 14.79 |
| 2 mm  | 2.05                              | 0.21   | 10.02 | 1.49    | 0.13   | 11.17 |
| 4 mm  | 12.62                             | 1.02   | 12.41 | 7.39    | 0.25   | 30.11 |
| Total | 22.25                             | 1.90   | 11.70 | 13.53   | 0.69   | 19.52 |

|        | PBT model (depth 6, num. weak 40) |        |      |         |        |      |
|--------|-----------------------------------|--------|------|---------|--------|------|
|        | Config1                           |        |      | Config2 |        |      |
|        | Cpu(s)                            | Gpu(s) | S-up | Cpu(s)  | Gpu(s) | S-up |
| 1 mm   | 4.02                              | 0.78   | 5.12 | 1.63    | 0.29   | 5.55 |
| 2 mm   | 1.62                              | 0.36   | 4.49 | 0.81    | 0.14   | 5.9  |
| 4 mm   | 2.98                              | 0.71   | 4.20 | 1.39    | 0.17   | 8.15 |
| Totals | 8.62                              | 1.85   | 4.65 | 3.83    | 0.60   | 6.37 |

[0093] On the PDT data, a model was trained using 161 training samples and tested on 86 unseen cases. The PBT contained 4 tree levels with 40, 42, 44, and 48 weak classifiers on each level. For the forest implementation, the maximum depth was set to 8 and trained two sets of models with either 10 or 30 trees. The forest implementation with 10 trees performs better than the PBT, which has a similar number of weak classifiers.

[0094] Table 7 shows the comparison of timings, speed-ups, and accuracy on the 16-8-4 mm translation hierarchy. The CPU version of the forest implementation is slower; and because better speedups occur on the GPU, the two GPU

implementations take the same amount of time. The forest implementation is superior in this example due to superior detection accuracy. Using 30 trees improves accuracy because the GPU accelerated portions occupy more time and better speed-ups can be seen.

TABLE 7

|                      | Config1 |        |      | Config2 |        |      | Testing |
|----------------------|---------|--------|------|---------|--------|------|---------|
|                      | Gpu(s)  | Cpu(s) | S-up | Gpu(s)  | Cpu(s) | S-up | error   |
| PBT                  | 0.31    | 2.25   | 7.36 | 0.20    | 0.98   | 4.92 | 10.07   |
| Ranfor<br>(10 trees) | 0.31    | 2.56   | 8.34 | 0.21    | 1.20   | 5.61 | 7.47    |
| Ranfor<br>(30 trees) | 0.60    | 6.44   | 10.7 | 0.27    | 2.88   | 10.6 | 4.93    |

[0095] The following paragraphs discuss the evaluation of improvements to the training procedure of the PBT. As discussed, the training of AdaBoost classifiers for a PBT could be performed by computing a feature matrix for feature evaluation, and then training the weak classifiers in parallel on the GPU. In the hierarchical network, after the classifier for a PBT is trained, a batch detection phase is performed. During batch detection, the newly trained classifier is used to perform detection on each input volume, and the results are then used by the later nodes in the network. The detection code can be immediately reused to speed-up the phases.

[0096] When training a strong classifier, there are often too many positive and negative training samples to build a full feature matrix, so a number of input samples are chosen randomly to build the feature matrix (around 2500 positive and 2500 negative samples). The input samples come from different volumes, and there are often too many volumes to keep on the GPU at all times. Therefore, if the features are to be evaluated on the GPU, the total number of input volumes used by the samples would need to be transferred to the GPU.

[0097] In order to evaluate the performance of the feature matrix computation during training, 3 volumes are used and the speed-up is measured in evaluating and reading back all features. FIG. 16 shows speedups for Config2 using small volumes in graph 1602 and speedups for Config2 using medium volumes on graph 1604. The Haar features line for graph 1602 is represented by line 1606 and the Steerable features line for graph 1602 is represented by line 1608. The Haar features line for graph 1604 is represented by line 1610 and the Steerable features line for graph 1604 is represented by line 1612. From FIG. 16 it is apparent that if only a few samples share the same volume, the speed-up will be low. For Haar features, the speed-up also depends on the size of the image (as the integral image is computed on the GPU).

[0098] In a benchmark, a hierarchical model was trained. The hierarchical model had 3 resolutions of data, each with translation, orientation, and scale detectors using 300 images from the SI data set. For each PBT trained, there were roughly 2400 positive samples and 140K negative samples. Training of a node used a maximum of 2500 negative samples. Timing results were all computed on Config2. The results after including the CUDA enhancements (not including feature matrix computation) are shown by Table 8. Table 8 shows the total times for phases of the training on Config2. The weak classifier training was sped-up, and the overall computation time gives a 1.5x speedup overall. The 3.5x speed-up on weak classifier training gives a roughly 2.1x speed-up for training a strong classifier. The speed-up on the entire PBT training was

1.7x (which includes loading data, determining which samples to use, propagating samples down the tree, and creating the tree data structure).

TABLE 8

| Component        | GPU (s) | CPU (S) | Speed-up |
|------------------|---------|---------|----------|
| •• Feature       | 1675.3  | 1623.4  | 1.0      |
| •• Weak classifs | 1346.6  | 4695.7  | 3.5      |
| • Strong         | 3021.8  | 6319.1  | 2.1      |
| All Train        | 4269.4  | 7243.3  | 1.7      |
| Train Detect     | 277.9   | 398.7   | 1.4      |
| Other            | 1552.7  | 1649.0  | 1.1      |
| Total            | 6100.0  | 9291.0  | 1.5      |

[0099] The above-described methods for evaluation of a probabilistic boosting tree using a stack implementation and for anatomical landmark, position estimation, and object segmentation in a 3D volume, may be implemented on a computer using well-known computer processors, memory units, storage devices, computer software, and other components. A high level block diagram of such a computer is illustrated in FIG. 17. Computer 1702 contains a processor 1704 which controls the overall operation of the computer 1702 by executing computer program instructions which define such operation. The computer program instructions may be stored in a storage device 1712, or other computer readable medium (e.g., magnetic disk, CD ROM, etc.) and loaded into memory 1710 when execution of the computer program instructions is desired. Thus, the steps of the methods of FIGS. 1, 6, and 13 may be defined by the computer program instructions stored in the memory 1710 and/or storage 1712 and controlled by the processor 1704 executing the computer program instructions. An image acquisition device 1720 can be connected to the computer 1702 to input images to the computer 1702. For example the image acquisition device 1720 may be a C-arm image acquisition system capable of inputting 3D C-arm CT images and 2D fluoroscopic images to the computer 1702. It is possible to implement the image acquisition device 1720 and the computer 1702 as one device. It is also possible that the image acquisition device 1720 and the computer 1702 communicate wirelessly through a network. The computer 1702 also includes one or more network interfaces 1706 for communicating with other devices via a network. The computer 1702 also includes other input/output devices 1708 that enable user interaction with the computer 1702 (e.g., display, keyboard, mouse, speakers, buttons, etc.). One skilled in the art will recognize that an implementation of an actual computer could contain other components as well, and that FIG. 17 is a high level representation of some of the components of such a computer for illustrative purposes.

[0100] The foregoing Detailed Description is to be understood as being in every respect illustrative and exemplary, but not restrictive, and the scope of the invention disclosed herein is not to be determined from the Detailed Description, but rather from the claims as interpreted according to the full breadth permitted by the patent laws. It is to be understood that the embodiments shown and described herein are only illustrative of the principles of the present invention and that various modifications may be implemented by those skilled in the art without departing from the scope and spirit of the invention. Those skilled in the art could implement various other feature combinations without departing from the scope and spirit of the invention.

1. A method for training a probabilistic boosting tree, comprising:
  - receiving training data at a graphics processing unit (GPU);
  - dividing the training data into a first dataset and a second dataset using a classifier;
  - training a first sub-tree and a second sub-tree at the GPU, the first sub-tree using the first dataset and the second sub-tree using the second dataset;
  - generating a posterior distribution model based on the trained first sub-tree and the trained second sub-tree.
2. The method of claim 1, wherein training a first sub-tree and a second sub-tree at the GPU comprises:
  - training a classifier at each node of the first sub-tree and the second sub-tree.
3. The method of claim 2, wherein training a classifier at each node of the first sub-tree and the second sub-tree comprises:
  - computing a feature matrix for each node.
4. The method of claim 3, computing a feature matrix for each node comprises:
  - computing feature bounds;
  - mapping feature values to samples within a histogram bin based on the feature bounds; and
  - computing an error associated with each classifier.
5. The method of claim 1, further comprising:
  - determining a set of classifications based on the posterior distribution model.
6. The method of claim 1, wherein the method for evaluating a probabilistic boosting tree is implemented using a parallel computing architecture.
7. The method of claim 7, wherein the parallel computing architecture is Compute Unified Device Architecture (CUDA).
8. The method of claim 1, further comprising:
  - detecting structures in a 3D volume using the trained probabilistic boosting tree.
9. A method for determining the posterior distribution of a probabilistic boosting tree, comprising:
  - receiving input data at a graphics processing unit (GPU)
  - determining a weighted empirical distribution associated with each node of the probabilistic boosting tree using a stack implementation;
  - adding the weighted empirical distribution associated with each node to a total posterior distribution value.
10. The method of claim 9, wherein determining a weighted empirical distribution associated with each node of the probabilistic boosting tree using a stack implementation comprises:
  - pushing a root node of the probabilistic boosting tree onto a stack;
  - determining whether to descend down the probabilistic boosting tree thru a node in a left sub-tree or thru a node in a right sub-tree.
11. The method of claim 10, wherein determining whether to descend down the probabilistic boosting tree thru a node in a left sub-tree or thru a node in a right sub-tree comprises:
  - determining a discriminative classifier of the root node;
  - if the discriminative classifier of the root node satisfies a first condition, descending a left node in the left sub-tree;
  - if the discriminative classifier of the root node satisfies a second condition, descending a right node in the right sub-tree; and

- if the discriminative classifier of the root node satisfies a third condition, descending both the left node and the right node;
- wherein the first condition signifies that a left node must be descended, the second condition signifies that a right node must be descended, and the third condition signifies that both a left node and a right node must be descended.
- 12.** The method of claim **11**, wherein descending to a left node in the left sub-tree comprises:
- determining a discriminative classifier of the left node;
  - if the left node is a leaf node, adding a weighted empirical distribution of the left node to the total posterior distribution value;
  - if the left node satisfies the first condition, pushing the left node onto the stack and descend a left child node of the left node;
  - if the left node satisfies the second condition, pushing the left node onto the stack and descend a right child node of the left node;
  - if the left node satisfies the third condition, pushing the left node onto the stack and descend both the left child node and the right child node.
- 13.** The method of claim **11**, wherein descending to a right node in the right sub-tree comprises:
- determining a discriminative classifier of the right node;
  - if the right node is a leaf node, adding a weighted empirical weighted distribution of the left node to the total posterior distribution value;
  - if the right node satisfies the first condition, pushing the right node onto the stack and descend a left child node of the right node;
  - if the right node satisfies the second condition, pushing the right node onto the stack and descend a right child node of the right node;
  - if the right node satisfies the third condition, pushing the right node onto the stack and descend both the left child node and the right child node.
- 14.** The method of claim **9**, wherein the method for evaluating a probabilistic boosting tree is implemented using a parallel computing architecture.
- 15.** The method of claim **9**, wherein the parallel computing architecture is Compute Unified Device Architecture (CUDA).
- 16.** The method of claim **11**, wherein the stack is associated with one of a plurality of threads running in parallel.
- 17.** A method for evaluating a forest of probabilistic boosting trees, comprising:
- receiving input data at a graphics processing unit (GPU);
  - evaluating the plurality of probabilistic boosting trees using a stack implementation;
  - generating a combined posterior distribution based on a posterior distribution of each of the plurality of probabilistic boosting trees.
- 18.** A system for training a probabilistic boosting tree, comprising:
- means for receiving training data at a graphics processing unit (GPU);
  - means for dividing the training data into a first dataset and a second dataset using a classifier;
  - means for training a first sub-tree and a second sub-tree at the GPU, the first sub-tree using the first dataset and the second sub-tree using the second dataset;
  - means for generating a posterior distribution model based on the trained first sub-tree and the trained second sub-tree.
- 19.** The system of claim **18**, wherein means for training a first sub-tree and a second sub-tree at the GPU comprises:
- means for training a classifier at each node of the first sub-tree and the second sub-tree.
- 20.** The system of claim **19**, wherein means for training a classifier at each node of the first sub-tree and the second sub-tree comprises:
- means for computing a feature matrix for each node.
- 21.** The system of claim **20**, wherein means for computing a feature matrix for each node comprises:
- means for computing feature bounds;
  - means for mapping feature values to samples within a histogram bin based on the feature bounds; and
  - means for computing an error associated with each classifier.
- 22.** The system of claim **18**, further comprising:
- means for determining a set of classifications based on the posterior distribution model.
- 23.** The system of claim **18**, wherein the system for evaluating a probabilistic boosting tree is a parallel computing architecture.
- 24.** The system of claim **23**, wherein the parallel computing architecture is Compute Unified Device Architecture (CUDA).
- 25.** The system of claim **18**, further comprising:
- means for detecting structures in a 3D volume using the trained probabilistic boosting tree.
- 26.** A system for determining the posterior distribution of a probabilistic boosting tree, comprising:
- means for receiving input data at a graphics processing unit (GPU)
  - means for determining a weighted empirical distribution associated with each node of the probabilistic boosting tree using a stack implementation;
  - means for adding the weighted empirical distribution associated with each node to a total posterior distribution value.
- 27.** The system of claim **26**, wherein means for determining a weighted empirical distribution associated with each node of the probabilistic boosting tree using a stack implementation comprises:
- means for pushing a root node of the probabilistic boosting tree onto a stack;
  - means for determining whether to descend down the probabilistic boosting tree thru a node in a left sub-tree or thru a node in a right sub-tree.
- 28.** The system of claim **27**, wherein means for determining whether to descend down the probabilistic boosting tree thru a node in a left sub-tree or thru a node in a right sub-tree comprises:
- means for determining a discriminative classifier of the root node;
  - means for descending a left node in the left sub-tree if the discriminative classifier of the root node satisfies a first condition;
  - means for descending a right node in the right sub-tree if the discriminative classifier of the root node satisfies a second condition; and
  - means for descending both the left node and the right node if the discriminative classifier of the root node satisfies a third condition;

- wherein the first condition signifies that a left node must be descended, the second condition signifies that a right node must be descended, and the third condition signifies that both a left node and a right node must be descended.
- 29.** The system of claim **28**, wherein means for descending to a left node in the left sub-tree comprises:
- means for determining the discriminative classifier of the left node;
  - means for adding a weighted empirical weighted distribution of the left node to the total posterior distribution value if the left node is a leaf node;
  - means for pushing the left node onto the stack and descend a left child node of the left node if the left node satisfies the first condition;
  - means for pushing the left node onto the stack and descend a right child node of the left node if the left node satisfies the second condition;
  - means for pushing the left node onto the stack and descend both the left child node and the right child node if the left node satisfies the third condition.
- 30.** The system of claim **28**, wherein means for descending to a right node in the right sub-tree comprises:
- means for determining the discriminative classifier of the right node;
  - means for adding a weighted empirical distribution of the left node to the total posterior distribution value if the right node is a leaf node;
  - means for pushing the right node onto the stack and descend a left child node of the right node if the right node satisfies the first condition;
  - means for pushing the right node onto the stack and descend a right child node of the right node if the right node satisfies the second condition;
  - means for pushing the right node onto the stack and descend both the left child node and the right child node if the right node satisfies the third condition.
- 31.** The system of claim **26**, wherein the system for evaluating a probabilistic boosting tree is implemented using a parallel computing architecture.
- 32.** The system of claim **31**, wherein the parallel computing architecture is Compute Unified Device Architecture (CUDA).
- 33.** The system of claim **28**, wherein the stack is associated with one of a plurality of threads running in parallel.
- 34.** A system for evaluating a forest of probabilistic boosting trees, comprising:
- means for receiving input data at a graphics processing unit (GPU);
  - means for evaluating the plurality of probabilistic boosting trees using a stack implementation;
  - means for generating a combined posterior distribution based on a posterior distribution of each of the plurality of probabilistic boosting trees.
- 35.** A non-transitory computer readable medium encoded with computer executable instructions for training a probabilistic boosting tree, the computer executable instructions defining steps comprising:
- receiving training data at a graphics processing unit (GPU);
  - dividing the training data into a first dataset and a second dataset using a classifier;
  - training a first sub-tree and a second sub-tree at the GPU, the first sub-tree using the first dataset and the second sub-tree using the second dataset;
  - generating a posterior distribution model based on the trained first sub-tree and the trained second sub-tree.
- 36.** The computer readable medium of claim **35**, wherein the computer executable instructions defining the step of training a first sub-tree and a second sub-tree at the GPU comprises computer executable instructions defining the step of:
- training a classifier at each node of the first sub-tree and the second sub-tree.
  - computing an empirical distribution to determine the first dataset to send to the first sub-tree and to determine the second dataset to send to the second sub-tree.
- 37.** The computer readable medium of claim **36**, wherein the computer executable instructions defining the step of training a classifier at each node of the first sub-tree and the second sub-tree comprises computer executable instructions defining the step of:
- computing a feature matrix for each node.
- 38.** The computer readable medium of claim **37**, wherein the computer executable instructions defining the step of computing a feature matrix for each node comprises wherein the computer executable instructions defining the steps of:
- computing feature bounds;
  - mapping feature values to samples within a histogram bin based on the feature bounds; and
  - computing an error associated with each classifier.
- 39.** The computer readable medium of claim **35**, further comprising computer executable instruction defining the step of:
- determining a set of classifications based on the posterior distribution model.
- 40.** The computer readable medium of claim **35**, wherein the computer executable instructions for evaluating a probabilistic boosting tree are implemented using a parallel computing architecture.
- 41.** The computer readable medium of claim **40**, wherein the parallel computing architecture is Compute Unified Device Architecture (CUDA).
- 42.** The computer readable medium of claim **35**, further comprising computer executable instructions defining the step of:
- detecting structures in a 3D volume using the trained probabilistic boosting tree.
- 43.** A non-transitory computer readable medium encoded with computer executable instructions for determining the posterior distribution of a probabilistic boosting tree, the computer executable instructions defining steps comprising:
- receiving input data at a graphics processing unit (GPU);
  - determining a weighted empirical distribution associated with each node of the probabilistic boosting tree using a stack implementation;
  - adding the weighted empirical distribution associated with each node to a total posterior distribution value.
- 44.** The computer readable medium of claim **43**, wherein the computer executable instructions defining the step of determining a weighted empirical distribution associated with each node of the probabilistic boosting tree using a stack implementation comprises computer executable instructions defining the steps of:
- pushing a root node of the probabilistic boosting tree onto a stack;
  - determining whether to descend down the probabilistic boosting tree thru a node in a left sub-tree or thru a node in a right sub-tree.

45. The computer readable medium of claim 44, wherein the computer executable instructions defining the step of determining whether to descend down the probabilistic boosting tree thru a node in a left sub-tree or thru a node in a right sub-tree comprises computer executable instructions defining the steps of:

- determining a discriminative classifier of the root node;
  - if the discriminative classifier of the root node satisfies a first condition, descending a left node in the left sub-tree;
  - if the discriminative classifier of the root node satisfies a second condition, descending a right node in the right sub-tree; and
  - if the discriminative classifier of the root node satisfies a third condition, descending both the left node and the right node;
- wherein the first condition signifies that a left node must be descended, the second condition signifies that a right node must be descended, and the third condition signifies that both a left node and a right node must be descended.

46. The computer readable medium of claim 45, wherein the computer executable instructions defining the step of descending to a left node in the left sub-tree comprises computer executable instructions defining the steps of:

- determining the discriminative classifier of the left node;
- if the left node is a leaf node, adding a weighted empirical weighted distribution of the left node to the total posterior distribution value;
- if the left node satisfies the first condition, pushing the left node onto the stack and descend a left child node of the left node;
- if the left node satisfies the second condition, pushing the left node onto the stack and descend a right child node of the left node;
- if the left node satisfies the third condition, pushing the left node onto the stack and descend both the left child node and the right child node.

47. The computer readable medium of claim 45, wherein the computer executable instructions defining the step of descending to a right node in the right sub-tree comprises computer executable instructions defining the steps of:

- determining the discriminative classifier of the right node;
- if the right node is a leaf node, adding a weighted empirical weighted distribution of the left node to the total posterior distribution value;
- if the right node satisfies the first condition, pushing the right node onto the stack and descend a left child node of the right node;
- if the right node satisfies the second condition, pushing the right node onto the stack and descend a right child node of the right node;
- if the right node satisfies the third condition, pushing the right node onto the stack and descend both the left child node and the right child node.

48. The computer readable medium of claim 43, wherein the computer executable instructions for evaluating a probabilistic boosting tree are implemented using a parallel computing architecture.

49. The computer readable medium of claim 48, wherein the parallel computing architecture is Compute Unified Device Architecture (CUDA).

50. The computer readable medium of claim 45, wherein the stack is associated with one of a plurality of threads running in parallel.

51. A non-transitory computer readable medium encoded with computer executable instructions for evaluating a forest of probabilistic boosting trees, the computer executable instructions defining steps comprising:

- receiving input data at a graphics processing unit (GPU);
- evaluating the plurality of probabilistic boosting trees using a stack implementation;
- generating a combined posterior distribution based on a posterior distribution of each of the plurality of probabilistic boosting trees.

\* \* \* \* \*