SVM-Based Failure Detection of GHT Localizations

Thomas Blaffert, Cristian Lorenz, Hannes Nickisch, Jochen Peters, Jürgen Weese
Philips Research
March 3, 2016

innovation + you

PHILIPS

GHT Localization Classification

Aims

- Detect anatomical structures in an image (see below), e.g. in large data bases.
- Discriminate between correct and incorrect localizations, e.g. for Model Based Segmentation (MBS).
- Find better GHT solutions than just by voting, e.g. for improved MBS.

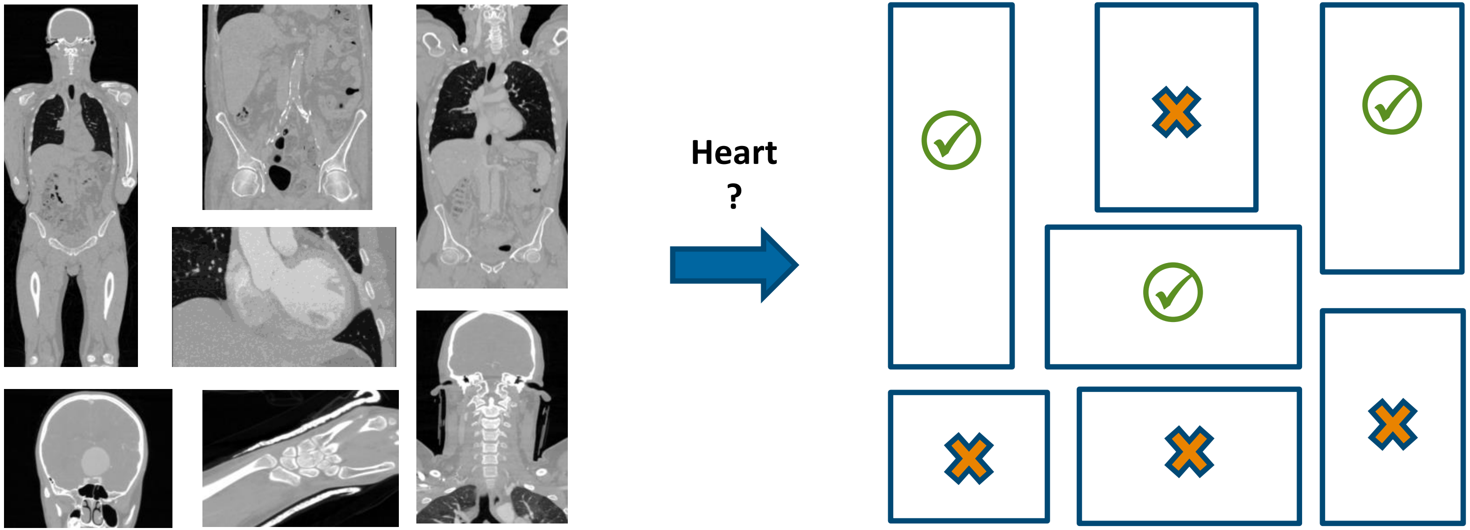

PHILIPS

GHT Localization Classification Method

- Input:
  GHT localization solution

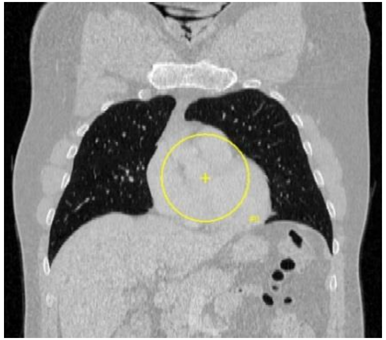

Heart localization

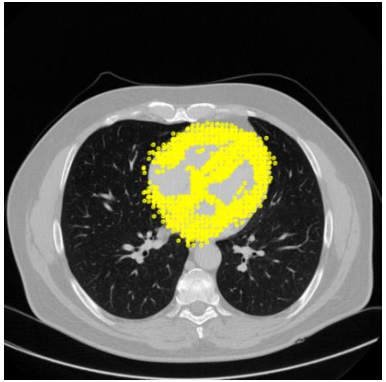

GHT shape model points

Voting shape model points

- New:
  Collective evaluation of voting GHT model point properties

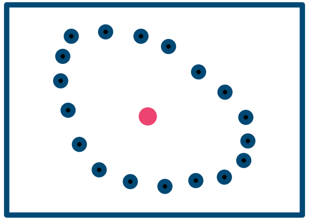

heart contained, high # votes

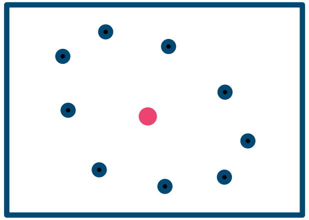

not contained, low # votes

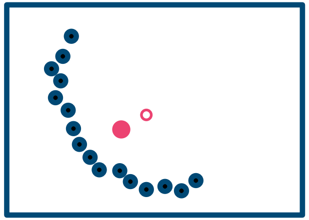

not contained, biased offset

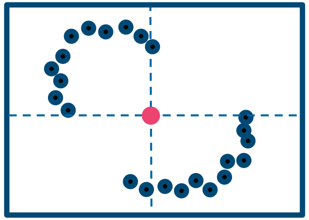

not contained, offset distribution

Generalized Hough Transform (GHT)

- Construct a shape model $\mathcal{M}$ with model points at
  - offsets $\boldsymbol{d}_i$ from a center
  - with strong edges in direction $\boldsymbol{n}_i$.
- Learn typical collections of $\boldsymbol{d}_i$ and $\boldsymbol{n}_i$ from a training set.
- Use a surface model (solid lines below) for restricting the selection of edges to relevant positions.

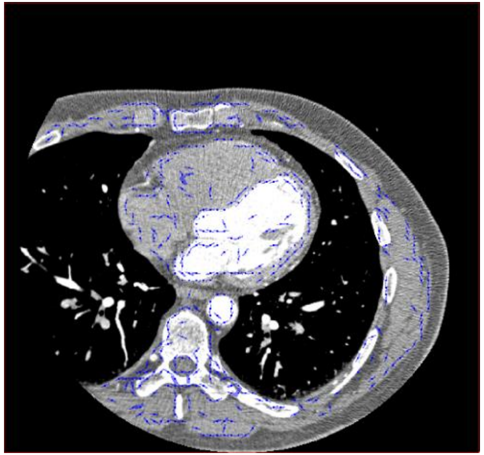

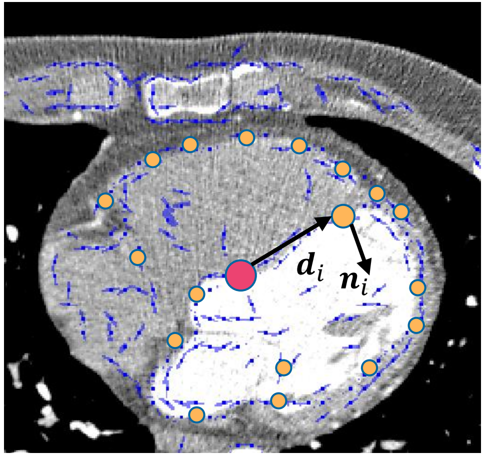

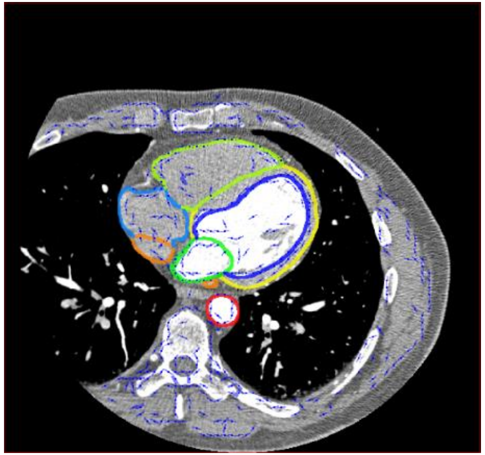

GHT Localization Algorithm

- Accumulate Hough votes at location $\boldsymbol{x}$:

  $$H(\boldsymbol{x}) = \sum_i h(\boldsymbol{x} + \boldsymbol{d}_i, \boldsymbol{n}_i)$$

- Choose $\boldsymbol{x}$ with highest vote count as localization solution (green area).

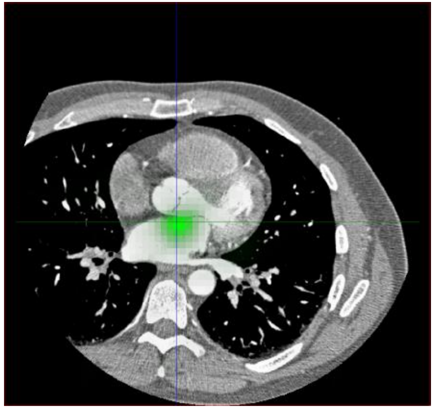

- Start with an image volume
- Calculate edge features
- Compare to reference model
  - Move model to test position
  - Count the number matching edges (votes)
- The votes are called the Hough space
- Choose the position with the highest number of votes.

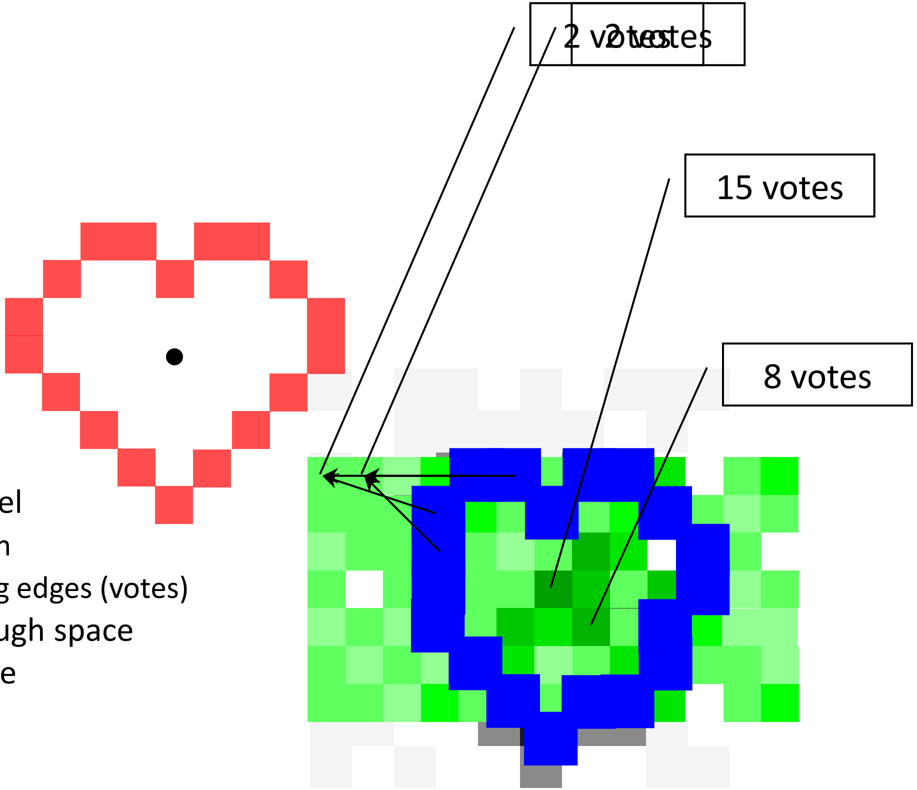

Confidence and Distance Features

- Confidence (vote count m relative to n shape points):

  $$f_c = m/n \, * 100$$

- Remark: Scores rather than counts possible, but not investigated.

- Offset distance:

  $f_d = \|\boldsymbol{o} - \boldsymbol{r}\|$, with

  $\boldsymbol{o} = 1/m \sum_{i=1}^{m} \boldsymbol{d}_i$ (average voting point offset),

  $\boldsymbol{r} = 1/n \sum_{j=1}^{n} \boldsymbol{d}_j$ (average model point offset)

- Gradient distance:

  $f_g = \|\boldsymbol{\omega} - \boldsymbol{\rho}\|$, with

  $\boldsymbol{\omega} = 1/m \sum_{i=1}^{m} \boldsymbol{n}_i$ (average voting gradient),

  $\boldsymbol{\rho} = 1/n \sum_{j=1}^{n} \boldsymbol{n}_j$ (average model gradient)

n shape points:

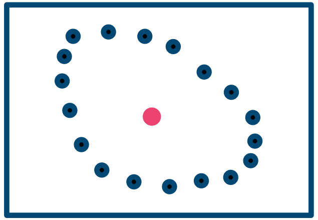

m votes:

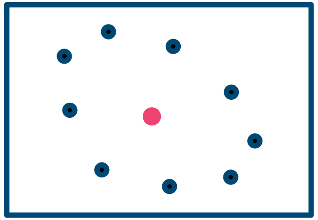

offset distance:

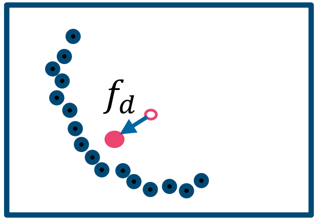

PHILIPS

Octant Distribution Features

- Model point offsets $\boldsymbol{d}_i$ are distributed over 8 spatial octants.
- Distribution of all m voting model points is stored in a histogram $\boldsymbol{h_o}$.
- A reference histogram $\boldsymbol{h_r}$ is calculated from all n shape model points.
- The new offset octant filling feature compares them by their difference.

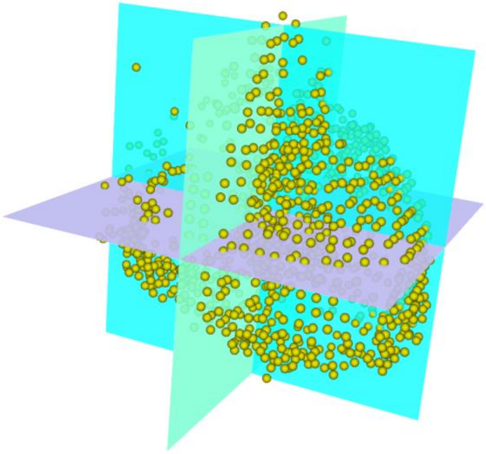

- Offset octants fill:

  $f_{od} = \sum_{l=0}^{7} |\boldsymbol{h_{ol}} - \boldsymbol{h_{rl}}|$, (l = histogram bin number)

- Similarly, histograms $\boldsymbol{h_\omega}$ and $\boldsymbol{h_\rho}$ are calculated and from the voting and shape model gradient vectors and compared.

- Gradient octants fill:

  $f_{og} = \sum_{l=0}^{7} |\boldsymbol{h_{\omega l}} - \boldsymbol{h_{\rho l}}|$, (l = histogram bin number)

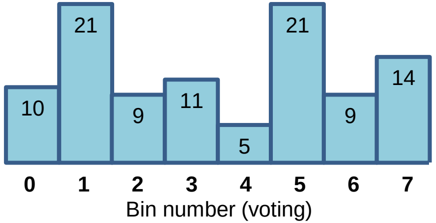

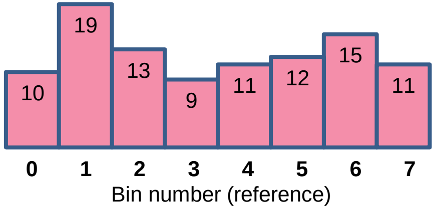

PHILIPS

Octant Distribution Features

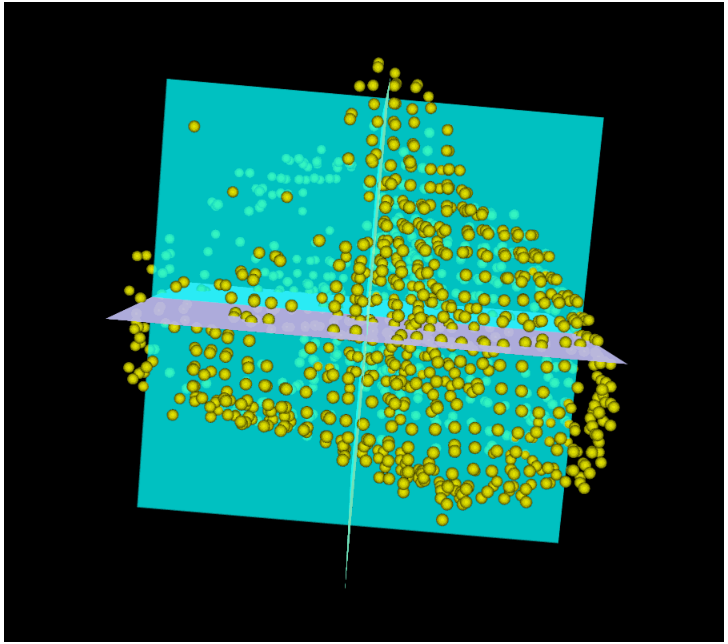

Voting GHT model points, offset distribution $\boldsymbol{h_o}$.

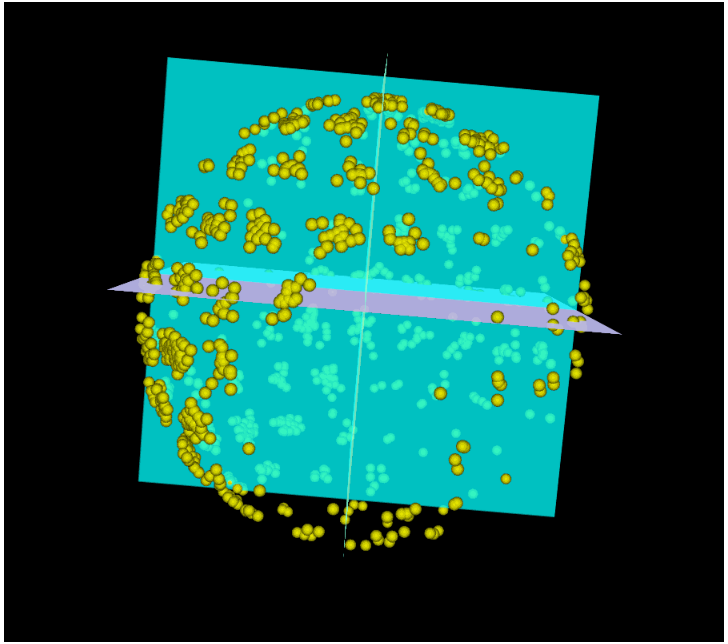

Voting GHT model points, gradient distribution $\boldsymbol{h_\omega}$.

PHILIPS

Support Vector Machine (SVM) Classifier Training

- Separate valid and invalid localizations by an optimal decision function

$$\mathrm{sgn}(\boldsymbol{w}^T \Phi(\boldsymbol{x}_i) + b)$$

$\boldsymbol{x}_i$	feature vector
$\Phi(\boldsymbol{x}_i)$	mapping function
$\boldsymbol{w}, b$	weights

- Solve the primal optimization problem

$$\min_{w,b,\xi} \quad \frac{1}{2}\boldsymbol{w}^T\boldsymbol{w} + C\sum_{i=1}^{l}\xi_i$$

$$\text{subject to} \quad y_i(\boldsymbol{w}^T\Phi(\boldsymbol{x}_i) + b) \geq 1 - \xi_i$$

$$\xi_i \geq 0, \quad i = 1, \dots, l$$

C Regularization parameter, penalty for wrong classifications.

- We use a Gaussian kernel function for the dual optimization problem

$$K(\boldsymbol{x}_i, \boldsymbol{x}_j) \equiv \Phi(\boldsymbol{x}_i)^T\Phi(\boldsymbol{x}_i) = \exp\left(-\gamma\left\|\boldsymbol{x}_i - \boldsymbol{x}_j\right\|^2\right), \gamma > 0$$

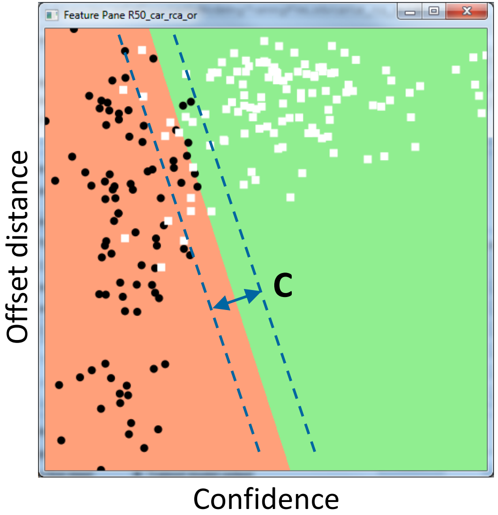

Works also for confidence only!

Grid search for optimal C, γ and feature combination

- For each SVM training run, the parameters C and γ are fixed.
- On a grid of C and γ optimal pairs are determined by the highest average accuracy in a 5-fold cross validation.
- Procedure iterates over all feature combinations (f_c , f_c+f_d, f_c+f_g , f_c+f_d+f_g , etc.).

PHILIPS

Experiments, Test Cases
Cardiac Substructures

- Test cases comprise GHT model of the full heart and 10 cardiac substructures.
- Each GHT model references a certain landmark (center, origin, ostium).
- Cardiac substructures were derived from the full heart model.

Anatomical structure / Landmark
Full heart center
Aortic valve
Pulmonary valve
Mitral valve
Tricuspid valve
Left coronary artery origin
Right coronary artery origin
Right inferior pulmonary vein (RIPV) ostium
Right superior pulmonary vein (RSPV) ostium
Superior vena cava (SVC) ostium

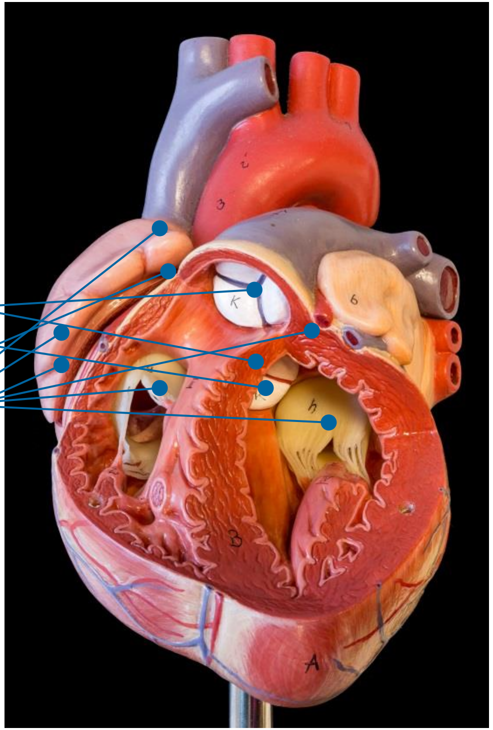

Experiments, Test Cases
Cardiac Substructures

- Test cases comprise GHT model of the full heart and 10 cardiac substructures.
- Each GHT model references a certain landmark (center, origin, ostium).
- Cardiac substructures were derived from the full heart model.

Anatomical structure / Landmark
Superior vena cava (SVC) ostium
Right coronary artery origin
Pulmonary valve
Right superior pulmonary vein (RSPV) ostium
Right inferior pulmonary vein (RIPV) ostium
Left coronary artery origin
Tricuspid valve
Full heart center
Aortic valve
Mitral valve

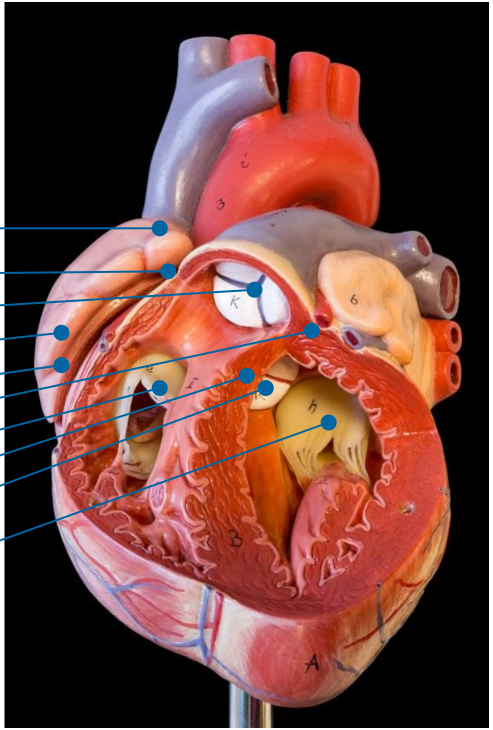

Cardiac Substructures
Landmarks of heart valves, heart vessels

Aortic Valve | Pulmon. Valve | Mitral Valve | Tricuspid Valve | Left Coronary | Right Coronary | RIPV ostium | RSPV ostium | SVC ostium

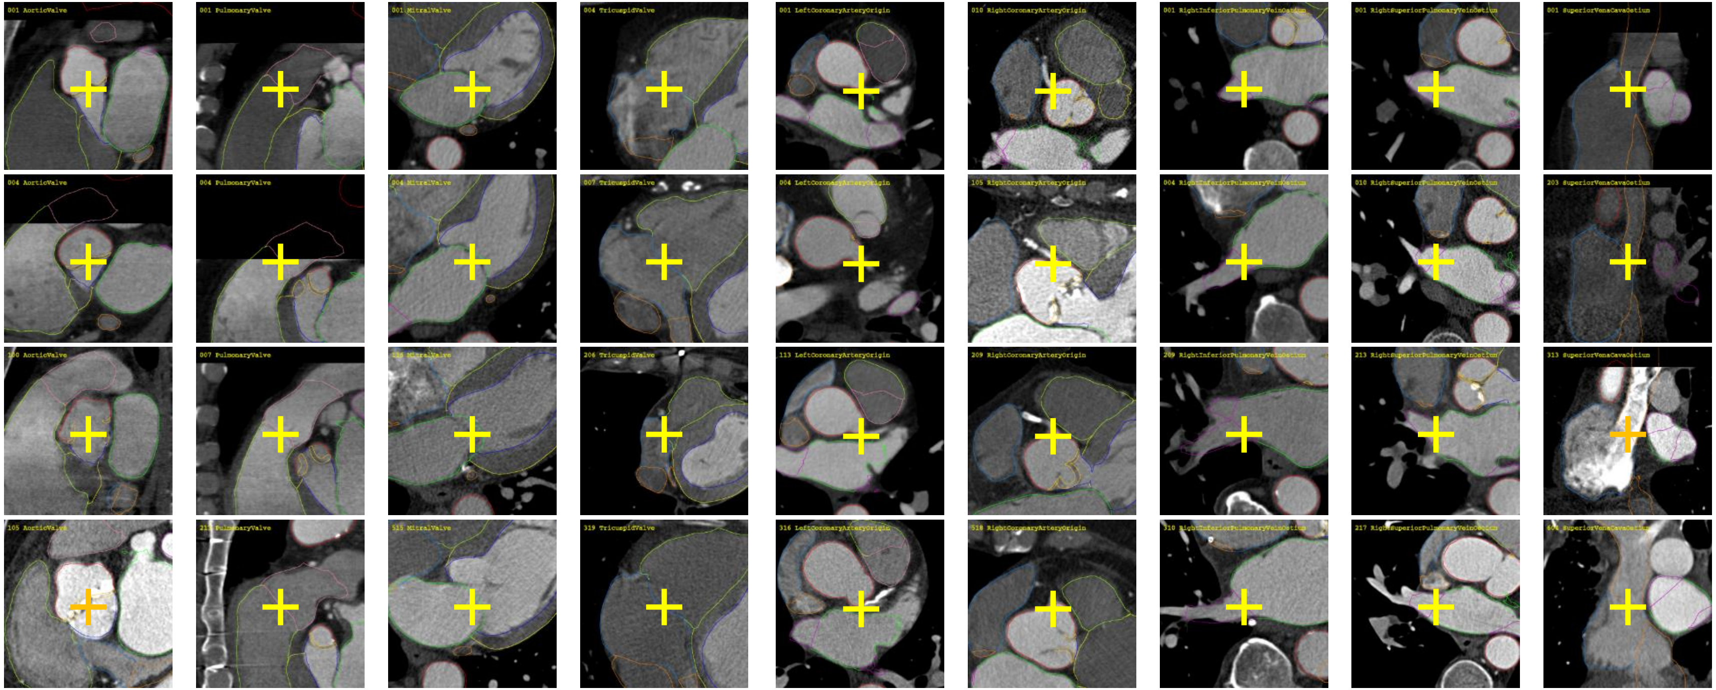

PHILIPS

Training Categories
Error Cases

- Invalid best GHT localization solutions can usually be clearly identified.
- Example: All 15 true error cases of mitral valve classification from the experiments.

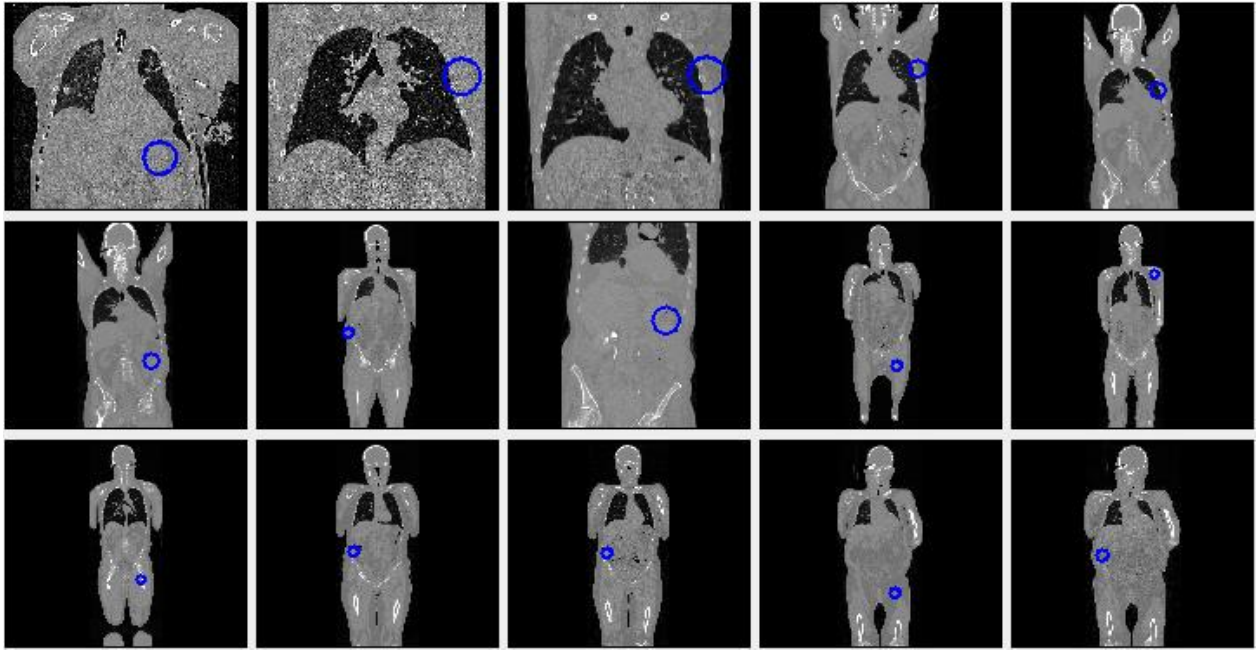

Training Categories
Valid and Negative Cases

- Valid best GHT localization solutions can also usually be clearly identified.
- Example: True valid cases of mitral valve classification from the experiments.

- For negative cases localizations are true negative or false positive by definition.
- Example: True negative cases of mitral valve classification from the experiments.

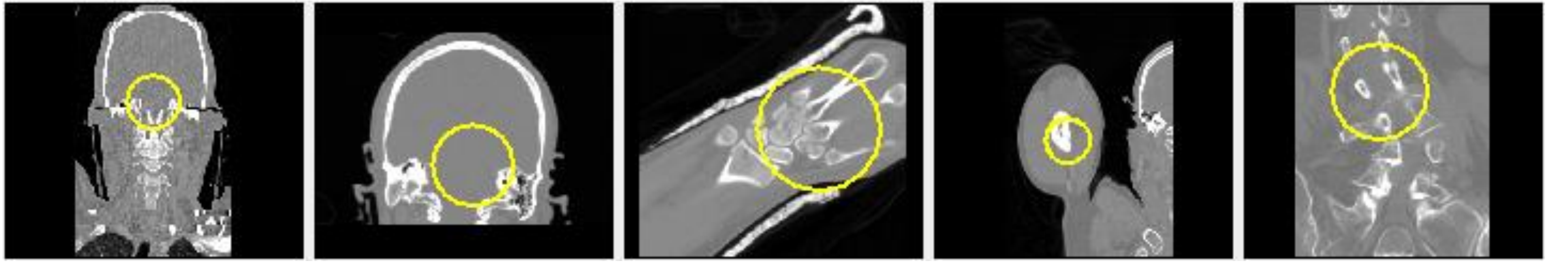

PHILIPS

Classification Categories

3 training categories (valid/error/negative) and
2 detection states (positive/negative) are assembled into
6 entries of an extended confusion matrix:

	Positive detection	Negative detection
Valid case	TV	FE
Error case	FV	TE
Negative case	FP	TN

Positive case: Landmark is contained in the image.

Valid case: Best GHT solution located at landmark, thus valid.

True valid (TV): Valid GHT solution is correctly classified as positive.

False error (FE): Valid GHT solution is incorrectly classified as negative.

Error case: Best GHT solution is not located at the landmark, thus invalid.

True error (TE): Invalid GHT solution is correctly classified as negative.

False valid (FV): Invalid GHT solution is incorrectly classified as positive.

Negative case: Landmark not contained in image, GHT solutions implicitly invalid.

True negative (TN): GHT solution is correctly classified as negative.

False positive (FP): GHT solution is incorrectly classified as positive.

PHILIPS

Experiments, 3 classifiers

- Classifier 1: Single confidence feature f_c, threshold search (ct)
- Classifier 2: SVM training on the confidence feature f_c (cs)
- Classifier 3: Optimized multi-feature SVM classification (ms)

- Results of the cross validation accuracy experiments with threshold classification
- Accuracy is calculated from the correctly classified (Tx) cases.
- 130 positive and 74 negative cases for each structure, 138 error cases in total.

Cardiac structure (landmark)	TV			FE			FV			TE			FP			TN		
	ct	cs	ms	ct	cs	ms	ct	cs	ms	ct	cs	ms	ct	cs	ms	ct	cs	ms
Full heart center	121	122	125	4	3	0	1	1	0	4	4	5	0	0	0	74	74	74
Aortic valve	113	113	115	5	5	3	1	1	1	11	11	11	1	0	0	73	74	74
Pulmonary valve	107	107	114	9	9	2	1	0	0	13	14	14	2	0	0	72	74	74
Mitral valve	105	105	112	10	10	3	4	3	0	11	12	15	1	0	0	73	74	74
Tricuspid valve	100	103	106	14	11	8	2	3	0	14	13	16	0	2	0	74	72	74
Left coronary artery	109	109	114	7	7	2	3	1	0	11	13	14	3	1	0	71	73	74
Right coronary artery	104	107	108	14	11	10	3	3	0	9	9	12	2	1	2	72	73	72
Right inf. pulmon. vein	95	94	103	14	15	6	4	2	3	17	19	18	2	0	0	72	74	74
Right sup. pulmon. vein	105	105	112	10	10	3	2	1	0	13	14	15	1	1	0	73	73	74
Superior vena cava	119	117	118	4	6	5	7	7	4	0	0	3	4	2	2	70	72	72
Sum	1078	1082	1127	91	87	42	28	22	8	103	109	123	16	7	4	724	733	736

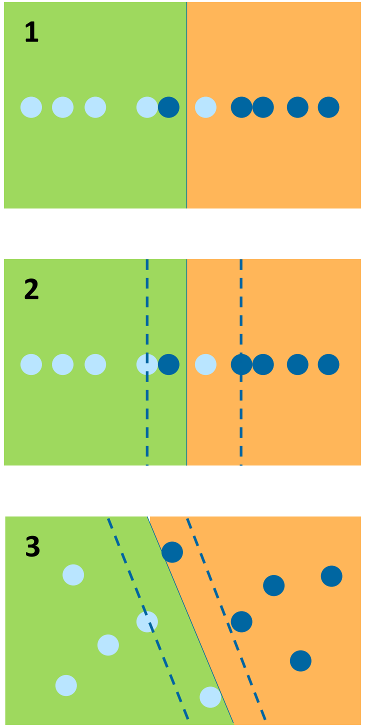

PHILIPS

Results

- Accuracy and error in percent for all three experiments.
- The multi-feature accuracy is obtained with the listed feature combination.
- Feature abbreviations: c = confidence, d = offset distance, g = gradient distance,
  od = offset octants fill, og = gradient octants fill.

Cardiac structure (landmark)	Confidence threshold (ct)		Confidence SVM (cs)		Multi feature SVM (ms)		Best feature combination
	Accuracy	Error	Accuracy	Error	Accuracy	Error	
Full heart center	95.59	4.41	98.04	1.96	100.00	0.00	c,g,og
Aortic valve	91.18	8.82	97.06	2.94	98.04	1.96	c,d,g,od
Pulmonary valve	87.75	12.25	95.59	4.41	99.02	0.98	c,og
Mitral valve	87.25	12.75	93.63	6.37	98.53	1.47	c,d,g,od,og
Tricuspid valve	85.29	14.71	92.16	7.84	96.08	3.92	c,d,og
Left coronary artery	88.24	11.76	95.59	4.41	99.02	0.98	c,og
Right coronary artery	86.27	13.73	92.65	7.35	94.12	5.88	c,d,g
Right inf. pulmon. vein	81.86	18.14	91.67	8.33	95.59	4.41	c,d,g
Right sup. pulmon. vein	87.25	12.75	94.12	5.88	98.53	1.47	c,d,g,od
Superior vena cava	92.65	7.35	92.65	7.35	94.61	5.39	c,g,og
Average	*88.33*	*11.67*	*94.31*	*5.69*	*97.35*	*2.65*	

Summary and Conclusions

- Aim: Distinction between valid and invalid GHT shape finder localizations.
- This distinction can be achieved by means of classification algorithms.
- Number of GHT voting counts are already a strong distinguishing feature.
- Introduction of additional collective voting model point features.
- Training of classifiers for these features with the SVM method.
- Confidence feature, threshold search vs. SVM training: Error reduction by ~50%.
- SVM training, confidence features vs. multi feature: Error reduction by ~50%.
- Best achievable error rate for valid/invalid localization classification: ~3%.
- Almost 3 times better than the intrinsic fraction of 11% invalid GHT localizations.

Acknowledgements

The authors would like to thank Katrina Read, Dirk Müller, Amnon Steinberg, and Mark Rabotnikov from Philips HealthTech for providing the CT Data.

PHILIPS

PHILIPS