

Efficient Conformal Volumetry for Template-Based Segmentation

Matt Y. Cheung¹[0000–0002–7846–3297], Ashok Veeraraghavan¹[0000–0001–5043–7460], and Guha Balakrishnan¹[0000–0001–8703–1368]

Rice University, Houston TX 77005, USA
mattyc@rice.edu

Abstract. Template-based segmentation, a widely used paradigm in medical imaging, propagates anatomical labels via deformable registration from a labeled atlas to a target image, and is often used to compute volumetric biomarkers for downstream decision-making. While conformal prediction (CP) provides finite-sample valid intervals for scalar metrics, existing segmentation-based uncertainty quantification (UQ) approaches either rely on learned model features, often unavailable in classic template-based pipelines, or treat the registration process as a black box, resulting in overly conservative intervals when applied directly in output space. We introduce ConVOLT, a CP framework that achieves efficient volumetric UQ by conditioning calibration on properties of the estimated deformation field from template-based segmentation. ConVOLT calibrates a learned volumetric scaling factor from deformation space features. We evaluate ConVOLT on template-based segmentation tasks involving global, regional, and label volumetry across multiple datasets and registration methods. ConVOLT achieves target coverage while producing substantially tighter intervals than output-space conformal baselines. Our work paves way to exploit the registration process for efficient UQ in medical imaging pipelines. Our code is available.

1 Introduction

Obtaining statistical guarantees on clinically meaningful quantities derived from medical image segmentation maps, such as anatomical volumes, is crucial to support reliable downstream decision-making. Recent work has demonstrated that Conformal Prediction (CP) [17,6,13] can provide finite-sample valid coverage guarantees for scalar metrics derived from black-box imaging pipelines [3,4]. But while this black-box approach guarantees marginal coverage, it generally yields inefficient (wide) prediction intervals. Recent work has shown that we can produce far more efficient intervals for metrics derived from medical image segmentation deep networks by conditioning on their internal learned feature representations (“feature CP”) [5,15]. However, many clinical pipelines rely on a different segmentation paradigm called *template-based segmentation* [8,2], where a new image is segmented by densely registering it to a labeled template and




Fig. 1. Overview of ConVOLT. We show a standard template-based segmentation pipeline for case i (black) where a moving image is registered to a fixed image, resulting in a deformation field. The fixed labels are warped to compute predicted labels and volumes $\hat{Y}_{i,l}^0$ on the moved image. ConVOLT (blue) trains a label-specific regressor on the training dataset D_{tr} with statistical features extracted from the deformation field features to predict a multiplicative scaling factor $\hat{\beta}_{i,l}$. The scaling factor is conformalized using the calibration dataset D_{cal} to attain valid prediction intervals.

propagating the template’s labels through the deformation field [14,9,1]. Because this paradigm does not use an explicit segmentation model, it is unclear how to readily extend feature CP methods to obtain efficient prediction intervals.

In this work, we bridge this gap by showing that conditioning CP on properties of the *deformation field* can yield substantially tighter intervals while preserving coverage. We hypothesize that deformation-aware conditioning leads to improved efficiency over output-space CP baselines because many downstream metrics are deterministic functionals of the deformation. In fact, volumetric change can be written as an integral of the Jacobian determinant over spatial regions, so uncertainty in the biomarker is induced by uncertainty in the deformation itself. The deformation field encodes informative geometric cues – such as local expansion, contraction, and spatial heterogeneity – that govern variability in volume estimates, with highly heterogeneous Jacobians typically inducing greater volumetric uncertainty than smooth, near-identity transformations.

Based on these observations, we introduce **ConVOLT** (Conformal Volumetry for Template-Based Segmentation) (Fig. 1), a CP framework for template-based segmentation pipelines that calibrates a learned volumetric scaling factor from deformation space features. The scaling preserves non-negativity and scales

uncertainty with volume. This yields finite-sample valid prediction intervals for volumetric biomarkers while exploiting geometric structure to improve efficiency. We evaluate ConVOLT on template-based segmentation tasks involving global, regional, and label volumetry across multiple datasets and registration methods, and find that ConVOLT achieves target coverage while producing tighter intervals than output-space CP baselines. These results demonstrate that exploiting deformation structure can significantly improve the efficiency of uncertainty quantification for template-based segmentation pipelines.

2 Method

The key idea in ConVOLT is to calibrate *multiplicative* deviations induced by the deformation field rather than additive residuals in output space. Because many volumetric biomarkers are deterministic functionals of the deformation (e.g., integrals of the Jacobian determinant), multiplicative errors in the deformation propagate multiplicatively to the scalar target. ConVOLT exploits this structure by learning deformation-conditioned scaling factors and conformalizing them via split CP, yielding closed-form volumetric intervals with finite-sample coverage.

Problem setup. For each case $i \in \{1, \dots, N\}$, a registration algorithm produces a displacement field $u_i : \Omega \rightarrow \mathbb{R}^3$ on a voxel lattice Ω . The induced local volume change is given by the Jacobian determinant $J_i(x) = \det(\mathbf{I} + \nabla u_i(x))$. Let $m_{i,l}^F : \Omega \rightarrow \{0, 1\}$ denote a fixed-space mask for label $l \in \{1, \dots, L\}$, where labels may represent the global ROI, spatial regions (e.g., shells), or anatomical labels. All quantities are evaluated on the fixed grid. Let dV denote voxel volume. Integrating the Jacobian over a label mask yields the deformation-derived baseline prediction $\hat{Y}_{i,l}^0 = \sum_{x \in \Omega} m_{i,l}^F(x) J_i(x) dV$ with ground-truth scalar $Y_{i,l}$.

Split CP. We seek valid prediction intervals for $Y_{i,l}$ that adapt to deformation variability. We utilize the following lemma to prove validity: Let $(X_{i,l}, Y_{i,l})$ be exchangeable pairs and split the data into a training set D_{tr} and calibration set $D_{\text{cal}} = \{(X_{i,l}, Y_{i,l})\}_{i=1}^n$. Given non-conformity scores $R_{i,l}$ on D_{cal} , define $\hat{q}$ as the $\lceil (1 - \alpha)(n + 1) \rceil$ -th order statistic. For a test pair $(X_{\text{test},l}, Y_{\text{test},l})$, the conformal set $\{y : R(X_{\text{test},l}, y) \leq \hat{q}\}$ satisfies $P(Y_{\text{test},l} \in S_{\hat{q}}(X_{\text{test},l})) \geq 1 - \alpha$ under exchangeability. This is a standard result of split CP [17,13,6]. The remainder of this section focuses on how to construct a novel volumetric non-conformity score and its valid prediction interval that leverage deformation field features from template-based segmentation.

ConVOLT. We start with the observation that multiplicative perturbations in the deformation induce multiplicative deviations in $\hat{Y}_{i,l}^0$. This is because volumetric targets are integrals of $J_i(x)$. We model this via $\hat{Y}_{i,l}^\beta = \beta_{i,l} \hat{Y}_{i,l}^0$ where $\beta_{i,l} \geq 0$, and define the oracle ratio $\beta_{i,l}^* = Y_{i,l} / \hat{Y}_{i,l}^0$. This parameterization preserves non-negativity and scales uncertainty with volume. Next, using the training set D_{tr} , we learn a predictor for deformation-conditioned ratios $\hat{\beta}(X_{i,l})$ from

deformation-derived features $X_{i,l}$ (e.g., Jacobian statistics, spatial heterogeneity) to $\beta_{i,l}^*$ on the training dataset D_{tr} . On D_{cal} , define nonconformity scores $R_{i,l} = |\beta_{i,l}^* - \hat{\beta}(X_{i,l})|$ and compute $\hat{q}$ by taking the $\lceil (1 - \alpha)(n + 1) \rceil$ -th empirical quantile of the scores. For a test case with features x and baseline $\hat{Y}_l^0(x)$, the conformal set in ratio space $\{y : |y/\hat{Y}_l^0(x) - \hat{\beta}(x)| \leq \hat{q}\}$ yields the interval:

$$\mathcal{C}_l(x) = [(\hat{\beta}(x) - \hat{q})\hat{Y}_l^0(x), (\hat{\beta}(x) + \hat{q})\hat{Y}_l^0(x)]. \quad (1)$$

The validity of ConVOLT relies on the following proposition: Assume $(X_{i,l}, Y_{i,l})$ are exchangeable across calibration and test cases. Then Eq. 1 satisfies $P(Y_{n+1,l} \in \mathcal{C}_l(X_{n+1,l})) \geq 1 - \alpha$. The result follows from the split CP lemma, $P(|\beta_{n+1,l}^* - \hat{\beta}(X_{n+1,l})| \leq \hat{q}) \geq 1 - \alpha$. Since $Y_{n+1,l} = \beta_{n+1,l}^* \hat{Y}_{n+1,l}^0$ and the mapping $\beta \mapsto Y$ is monotone for $\beta \geq 0$, the event transfers directly to $\mathcal{C}_l$.

When multiple labels share a deformation field, we aggregate label-wise scores into a case-level score $S_i = A(\{R_{i,l}\}_{l=1}^L)$ where $A : \mathbb{R}^L \rightarrow \mathbb{R}$ is an aggregation function (e.g., $\max_l$ for simultaneous guarantees or quantile for weaker guarantees) and apply split CP to $\{S_i\}_{i \in D_{cal}}$. The resulting $\hat{q}$ is shared across labels, yielding case-level coverage under case-level exchangeability.

3 Experimental Setup

Datasets and Volume Metrics. We evaluated ConVOLT on intra-patient and inter-patient template-based segmentation tasks using datasets from the Learn2Reg challenge [7]: 1) NLST: 209 lung CT inhale-exhale pairs with lung masks, 2) ThoraxCBCT: 30 chest CT pairs across breathing phases with lung masks, and 3) OASIS: 409 brain MRI scans with 35 anatomical labels. In the intra-patient setting (ThoraxCBCT and NLST), each example is a longitudinal image pair and the target is volume change. In the inter-patient setting (OASIS), atlas images are registered to a target patient and warped labels yield absolute-volume targets. To prevent data leakage, we excluded atlases used for training from calibration and testing sets. We evaluate global volume/volume change (all datasets), regional volume change (ThoraxCBCT and NLST), and label volume (OASIS). For regional volume change, we partitioned each lung mask into $L = 5$ concentric radial shells from boundary to interior. We defined each shell as label l with target $Y_{i,l}$ equal to regional volume change. We evaluated patient-level guarantees by aggregating label-wise scores using $\max$ and $Q_{0.9}(\cdot)$ aggregation. We performed calibration at the patient level. For label volume, we define each of the 35 anatomical structures as label l with target absolute volume $Y_{i,l}$ and atlas-derived prediction $\hat{Y}_{i,l}^0$. We extracted deformation-conditioned features near predicted label boundaries and fit a separate ridge regressor per label to predict multiplicative ratios. We applied calibration independently per label.

Registration Algorithms. We used two popular, representative registration methods: Demons [16] and VoxelMorph [1]. We used a standard multi-resolution configuration for Demons. We trained VoxelMorph with normalized cross-correlation similarity and smoothness regularization. We kept all preprocessing and training settings consistent across methods and datasets.

Experimental Protocol and Baselines. For each dataset, we randomly split the patient IDs into training/calibration/test sets with 0.4/0.4/0.2 proportions for NLST and OASIS, and 10/15/5 for ThoraxCBCT. We repeated experiments 100 times with random calibration-test splits at miscoverage level $\alpha = 0.1$. We report mean and standard deviation of empirical coverage and interval size. For global volume/volume change, we compared the following methods:

- **SCP** [10]: The absolute residual score $R_{i,l} = |Y_{i,l} - \hat{Y}_{i,l}^0|$.
- **CQR** [12]: conformalized quantile regression conditioned on baseline volume/change with pinball loss and ridge regularization ($\lambda = 0.01$).
- **LCP** [11]: A locally calibrated residual CP conditioned on baseline volume/change using $k = 50$ nearest calibration neighbors based on $|\hat{Y}^0|$.
- **ConVOLT**: Our proposed deformation-space method uses ridge regression to predict multiplicative ratios. For each target mask or region, we choose features that reflect registration difficulty. All feature summaries are computed on the fixed grid; μ , σ , Q_α , $\text{Frac}[\cdot]$, $\max$, and $\|\cdot\|$ denote the empirical mean, standard deviation, α -th percentile, voxel fraction satisfying a condition, maximum, and Euclidean norm, respectively. We use deformation features from four groups: (i) log-Jacobian statistics, $\mu(\log J)$, $\sigma(\log J)$, and $\mu(|\log J|)$, which capture average expansion/contraction and its variability; (ii) Jacobian percentiles and near-folding/extreme-compression fractions, including $Q_\alpha(J)$ for $\alpha \in \{10, 50, 90\}$ and $\text{Frac}[J < 0.1]$, $\text{Frac}[J < 0.01]$, which capture tail behavior and localized deformation failures that may be missed by averages; (iii) displacement and spatial-irregularity summaries, including statistics of $\|u\|$, $\|\nabla \log J\|$, $\nabla \cdot u$, and $\|\nabla \times u\|$, which capture motion magnitude, sharp spatial changes in volume change, local expansion/compression, and rotational deformation components; and (iv) fixed-warped image mismatch features, including $\text{MAE}(I^F, I_\phi^M)$, $\text{MSE}(I^F, I_\phi^M)$, and $\text{corr}(I^F, I_\phi^M)$, which capture residual image disagreement after registration, where I^F is the fixed image and I_ϕ^M is the warped moving image.

4 Results and Discussion

ConVOLT produces efficient CP intervals. For total volume change (ThoraxCBCT and NLST) and volume (OASIS), we compared ConVOLT to baseline methods in Tab. 1. Across most datasets and registration algorithms, ConVOLT achieved the desired coverage while producing smaller or comparable intervals relative to baselines. These results suggest that conditioning on deformation-field features can yield efficiency gains beyond those obtained from adaptive

Table 1. ConVOLT achieves high efficiency compared to baselines for global volume guarantees. For $\alpha = 0.1$, we show interval size and coverage (Mean $\pm$ STD) for SCP, CQR, LCP, and ConVOLT. While ConVOLT performs well on the majority of benchmarks, it achieves slightly worse/comparable performance on NLST with Demons and ThoraxCBCT with VoxelMorph.

Dataset	Method	Demons		VoxelMorph	
		Coverage	Interval Size (mL)	Coverage	Interval Size (mL)
ThoraxCBCT	CQR	0.90 $\pm$ 0.14	2707 $\pm$ 951	0.91 $\pm$ 0.15	2355 $\pm$ 830
	ConVOLT	0.93 $\pm$ 0.12	1222$\pm$357	0.93 $\pm$ 0.14	2354$\pm$1018
	LCP	0.91 $\pm$ 0.13	4631 $\pm$ 986	0.91 $\pm$ 0.15	2400 $\pm$ 542
	SCP	0.91 $\pm$ 0.13	4631 $\pm$ 986	0.91 $\pm$ 0.15	2400 $\pm$ 542
NLST	CQR	0.91 $\pm$ 0.06	496$\pm$84	0.91 $\pm$ 0.06	446 $\pm$ 117
	ConVOLT	0.91 $\pm$ 0.06	511 $\pm$ 89	0.92 $\pm$ 0.05	405$\pm$86
	LCP	0.90 $\pm$ 0.06	685 $\pm$ 77	0.94 $\pm$ 0.05	736 $\pm$ 65
	SCP	0.91 $\pm$ 0.06	895 $\pm$ 121	0.91 $\pm$ 0.06	721 $\pm$ 56
OASIS	CQR	0.90 $\pm$ 0.04	130 $\pm$ 6	0.90 $\pm$ 0.05	129 $\pm$ 6
	ConVOLT	0.91 $\pm$ 0.04	105$\pm$7	0.91 $\pm$ 0.06	106$\pm$11
	LCP	0.90 $\pm$ 0.04	152 $\pm$ 8	0.89 $\pm$ 0.05	134 $\pm$ 5
	SCP	0.91 $\pm$ 0.03	159 $\pm$ 6	0.91 $\pm$ 0.05	138 $\pm$ 5

output-space scores alone. However, in a few settings (e.g., NLST with Demons and ThoraxCBCT with VoxelMorph), ConVOLT and CQR achieved similar interval sizes. In these cases, residual volumetric variability is largely explained by volume magnitude itself, and the additional deformation features provided limited predictive signal for correcting errors. We also compared interval size for regional volume change in ThoraxCBCT and NLST against CQR in Tab. 2 for aggregation functions max and $Q_{0.9}$. On ThoraxCBCT, ConVOLT consistently produced tighter intervals while maintaining coverage, indicating that deformation features captured informative structure for regional errors. On NLST, results were mixed: ConVOLT improved efficiency for VoxelMorph but was less efficient than CQR for Demons under the max aggregation. This suggests that for NLST with Demons, the deformation-derived features were weakly correlated with regional registration errors, so conditioning on them introduced additional variance without improving bias correction. Finally, we compared interval size for label volume in OASIS against CQR in Fig. 2. Our results indicated that ConVOLT produced tighter intervals while maintaining coverage for the majority of labels. Overall, these findings indicate that deformation-aware calibration is most beneficial when deformation features meaningfully explain volumetric error variability, and remains competitive with strong adaptive baselines otherwise.

Interpretability. ConVOLT learns multiplicative correction factors $\hat{\beta}_{i,l}$ using ridge regression on deformation-derived features, allowing direct interpretation through the learned regression coefficients. We plotted the top-5 coefficient magnitudes over 100 randomized patient-level train/calibration/test splits and

Table 2. ConVOLT achieves high efficiency compared to CQR for regional guarantees. For $\alpha = 0.1$, we show interval size and coverage (Mean $\pm$ STD) for CQR (best performing baseline) and ConVOLT. While ConVOLT performs well on the majority of benchmarks, it has noticeably worse performance on NLST.

Dataset	Agg	Method	Demons		VoxelMorph	
			Coverage	Interval Size (mL)	Coverage	Interval Size (mL)
ThoraxCBCT	max	CQR	0.91 $\pm$ 0.14	1381 $\pm$ 362	0.92 $\pm$ 0.13	1148 $\pm$ 273
		ConVOLT	0.92 $\pm$ 0.14	1117$\pm$146	0.95 $\pm$ 0.10	1128$\pm$295
	Q _{0.9}	CQR	0.91 $\pm$ 0.14	1139 $\pm$ 187	0.92 $\pm$ 0.14	898 $\pm$ 154
		ConVOLT	0.92 $\pm$ 0.14	743$\pm$111	0.95 $\pm$ 0.11	822$\pm$176
NLST	max	CQR	0.91 $\pm$ 0.05	670$\pm$39	0.91 $\pm$ 0.05	376 $\pm$ 34
		ConVOLT	0.91 $\pm$ 0.05	1083 $\pm$ 91	0.91 $\pm$ 0.05	357$\pm$40
	Q _{0.9}	CQR	0.91 $\pm$ 0.05	619$\pm$28	0.92 $\pm$ 0.05	338 $\pm$ 30
		ConVOLT	0.91 $\pm$ 0.05	744 $\pm$ 58	0.91 $\pm$ 0.05	273$\pm$29




Fig. 2. ConVOLT achieves high efficiency compared to CQR for label volume guarantees. For $\alpha = 0.1$, we show interval size inflation of CQR (best performing baseline) in % for each label in OASIS with ConVOLT as the baseline. We find that ConVOLT achieves higher efficiency on the majority of labels.

reported the mean and standard deviation of the coefficient magnitude for each feature in Fig. 3. Larger magnitudes indicate higher contribution. On Thorax-CBCT, multiple deformation features exhibited moderate magnitudes for both Demons and VoxelMorph, suggesting that volumetric errors are partially explained by deformation heterogeneity. This aligns with the observed efficiency gains of ConVOLT on this dataset. On OASIS, Jacobian-based quantile features dominated with consistently large coefficients, indicating that spatial expansion statistics are strong predictors of volumetric bias. Correspondingly, ConVOLT achieved substantial interval reductions. In contrast, for NLST with Demons, coefficient magnitudes were relatively small, suggesting that deformation-derived features provide limited explanatory power for volumetric error. This explains why ConVOLT offers little or no efficiency improvement in this setting: when features are weakly correlated with residual error, deformation-aware condition-




Fig. 3. Coefficient magnitudes explain ConVOLT’s efficiency improvements. We show Mean $\pm$ STD of absolute ridge coefficients for features across 100 random calibration-test splits. Large and stable magnitudes indicate features are consistently used to predict multiplicative correction, while small and unstable magnitudes indicate the limited predictive value of features and explains weaker performance.

Table 3. Multiplicative modeling, learned ratios, and localized features improve ConVOLT efficiency. We report interval inflation (%) relative to ConVOLT when replacing multiplicative modeling with additive corrections, removing learning or features, or using global instead of local deformation features.

Dataset	Ablation	Demons		VoxelMorph	
		Coverage	Inflation (%)	Coverage	Inflation (%)
ThoraxCBCT	Additive	0.93 $\pm$ 0.13	41.4 $\pm$ 56.1	0.94 $\pm$ 0.14	2.8 $\pm$ 60.4
	No Features	0.91 $\pm$ 0.14	93.4 $\pm$ 77.3	0.91 $\pm$ 0.15	-22.2 $\pm$ 38.5
	No Learning	0.91 $\pm$ 0.13	218.1 $\pm$ 108.4	0.90 $\pm$ 0.14	-7.1 $\pm$ 45.5
	Global Features	0.92 $\pm$ 0.14	73.8 $\pm$ 32.9	0.95 $\pm$ 0.11	65.7 $\pm$ 57.8
NLST	Additive	0.91 $\pm$ 0.06	3.3 $\pm$ 26.3	0.91 $\pm$ 0.06	0.2 $\pm$ 28.5
	No Features	0.91 $\pm$ 0.06	11.9 $\pm$ 27.9	0.92 $\pm$ 0.05	12.8 $\pm$ 26.0
	No Learning	0.92 $\pm$ 0.06	75.1 $\pm$ 36.2	0.91 $\pm$ 0.06	64.0 $\pm$ 36.6
	Global Features	0.91 $\pm$ 0.05	31.7 $\pm$ 11.5	0.91 $\pm$ 0.05	72.2 $\pm$ 21.3
OASIS	Additive	0.90 $\pm$ 0.04	0.0 $\pm$ 9.4	0.90 $\pm$ 0.06	-0.9 $\pm$ 14.0
	No Features	0.90 $\pm$ 0.03	27.6 $\pm$ 9.7	0.90 $\pm$ 0.05	26.4 $\pm$ 16.2
	No Learning	0.91 $\pm$ 0.03	53.3 $\pm$ 11.7	0.90 $\pm$ 0.05	30.2 $\pm$ 14.3

ing cannot meaningfully outperform adaptive output-space baselines.

Ablations. ConVOLT’s efficiency can be attributed to two factors: 1) learned scores from deformation field features and 2) adaptive scores using the multiplicative formulation. We made several intentional design choices for ConVOLT. To justify them, we performed ablations on 1) multiplicative ($\beta_i \hat{Y}_{0,i}$) versus additive scalars ($\hat{Y}_{0,i} + \beta_i$), 2) no learning ($\hat{\beta}(x) = 1$), 3) no features ($\hat{\beta}(x) = \mu(\hat{\beta}_i^*)$), and 4) localized features (region-restricted deformation features) versus global features (whole-mask deformation features) for regional guarantees. We show set size inflation relative to ConVOLT in Tab. 3. We found that multiplicative factors produced more efficient intervals than additive scalars, learning and sample-specific features improved efficiency, and localized features improved ef-

iciency for region based guarantees.

5 Conclusion

We demonstrated with ConVOLT that conditioning CP on deformation-field features can produce substantially more efficient CP intervals for deformation-derived volumetric biomarkers. Our results further showed that its effectiveness depends on the predictive strength of geometric features. More broadly, these findings suggest that exploiting structure within medical imaging pipelines, beyond black-box outputs, can enable more efficient and interpretable UQ for clinically-relevant downstream metrics.

Acknowledgements. MC would like to acknowledge support from a fellowship from the Gulf Coast Consortia on the NLM Training Program in Biomedical Informatics and Data Science T15LM007093.

Disclosure of Interests. The authors have no competing interests to disclose.

References

1. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Gutttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging* **38**(8), 1788–1800 (2019)
2. Cabezas, M., Oliver, A., Lladó, X., Freixenet, J., Cuadra, M.B.: A review of atlas-based segmentation for magnetic resonance brain images. *Computer methods and programs in biomedicine* **104**(3), e158–e177 (2011)
3. Cheung, M.Y., Netherton, T.J., Court, L.E., Veeraraghavan, A., Balakrishnan, G.: Metric-guided conformal bounds for probabilistic image reconstruction. In: *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging: 7th International Workshop, UNSURE 2025, Held in Conjunction with MICCAI 2025, Daejeon, South Korea, September 27, 2025, Proceedings.* p. 169. Springer Nature (2025)
4. Cheung, M.Y., Netherton, T.J., Court, L.E., Veeraraghavan, A., Balakrishnan, G.: Bias-aware conformal prediction for metric-based imaging pipelines. In: *2026 IEEE 23rd International Symposium on Biomedical Imaging (ISBI).* pp. 01–05. IEEE (2026)
5. Cheung, M.Y., Veeraraghavan, A., Balakrishnan, G.: Compass: Robust feature conformal prediction for medical segmentation metrics. *arXiv preprint arXiv:2509.22240* (2025)
6. Fontana, M., Zeni, G., Vantini, S.: Conformal prediction: a unified review of theory and new challenges. *Bernoulli* **29**(1), 1–23 (2023)
7. Hering, A., Hansen, L., Mok, T.C., Chung, A.C., Siebert, H., Häger, S., Lange, A., Kuckertz, S., Heldmann, S., Shao, W., et al.: Learn2reg: comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning. *IEEE Transactions on Medical Imaging* **42**(3), 697–712 (2022)

8. Iglesias, J.E., Sabuncu, M.R.: Multi-atlas segmentation of biomedical images: a survey. *Medical image analysis* **24**(1), 205–219 (2015)
9. Klein, A., Andersson, J., Ardekani, B.A., Ashburner, J., Avants, B., Chiang, M.C., Christensen, G.E., Collins, D.L., Gee, J., Hellier, P., et al.: Evaluation of 14 nonlinear deformation algorithms applied to human brain mri registration. *Neuroimage* **46**(3), 786–802 (2009)
10. Lei, J., G’Sell, M., Rinaldo, A., Tibshirani, R.J., Wasserman, L.: Distribution-free predictive inference for regression. *Journal of the American Statistical Association* **113**(523), 1094–1111 (2018)
11. Papadopoulos, H., Proedrou, K., Vovk, V., Gammerman, A.: Inductive confidence machines for regression. In: *European conference on machine learning*. pp. 345–356. Springer (2002)
12. Romano, Y., Patterson, E., Candes, E.: Conformalized quantile regression. *Advances in neural information processing systems* **32** (2019)
13. Shafer, G., Vovk, V.: A tutorial on conformal prediction. *Journal of Machine Learning Research* **9**(3) (2008)
14. Sotiras, A., Davatzikos, C., Paragios, N.: Deformable medical image registration: A survey. *IEEE transactions on medical imaging* **32**(7), 1153–1190 (2013)
15. Teng, J., Wen, C., Zhang, D., Bengio, Y., Gao, Y., Yuan, Y.: Predictive inference with feature conformal prediction. *arXiv preprint arXiv:2210.00173* (2022)
16. Thirion, J.P.: Non-rigid matching using demons. In: *Proceedings CVPR IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. pp. 245–251. IEEE (1996)
17. Vovk, V., Gammerman, A., Shafer, G.: *Algorithmic learning in a random world*. Springer (2005)