Certified Dental Biometric Verification Under Partial Overlap

Krishi Attri
Seoul National University · Independent research
krishiattriwork@gmail.com github.com/Archerkattri/toothprint

Abstract

Dental identification is usually evaluated as closed-set retrieval: given a dental scan, rank the enrolled gallery and report whether the correct subject appears first. This framing hides two operational requirements. First, forensic and clinical deployments need an *accept/abstain* decision with a controlled false-match rate (FMR), not only a Rank-1 score. Second, practical queries are often partial: field-of-view limits, trauma, extractions, and missing teeth break the rigid alignment assumption used by most 3D dental matching pipelines. We present ToothPrint, a dental biometric verification pipeline for 3D intraoral scans and 2D radiographs that combines rigid surface verification, learned partial-overlap correspondence, and split-conformal calibration. On 200 public 3D dental arches, rigid surface verification achieves Rank-1 0.995 and EER 0.005 under full coverage. Under realistic whole-tooth dropout, however, rigid GICP falls to Rank-1 0.23 at 50% tooth retention. A learned point-correspondence verifier lifts this regime to Rank-1 0.867 (AUC 0.984), while a crop-hardened embedding baseline reaches 0.635. The same score family supports a conformal open-set decision: full-coverage FNIR at FMR=1% is 0.00, while heavy tooth loss is correctly exposed as an abstention regime. We report DET curves, hard-negative calibration, an untrained-correspondence control, a cross-dataset transfer test on Teeth3DS+, and auxiliary 2D radiograph and restoration-pattern results. The central limitation is explicit: because available public 3D dental scans are mostly single-session, identity queries are synthetic reacquisitions. The method is therefore evidence for a research direction, not a clinical claim.

1 Introduction

Dental structures are among the most durable biometric signals. Forensic odontology has long used dental charting, restorations, and radiographs for post-mortem identification; modern intraoral scanners now provide dense 3D crown and gingival surfaces. Recent work has shown that 3D intraoral scans can identify individuals with high closed-set accuracy using registration pipelines based on handcrafted descriptors and ICP refinement [13]. Yet a high Rank-1 number is not the same as a deployable biometric decision.

Two gaps motivate this paper. The first is *certification*. A forensic examiner or clinical reviewer needs to know when a system should accept a match, reject it, or decline to decide. Standard dental-identification papers typically report Rank-N retrieval, not a finite-sample bound on the false-match rate of an individual accept decision. The second is *partial overlap*. A full dental arch is the easy case: a rigid surface alignment can exploit the entire crown geometry. A real query can be incomplete because the scan did not cover the full arch, because teeth are missing, or because a fragment is available. Rigid methods then fail before scoring: the geometry used to initialize and refine the alignment is no longer the same geometry as the enrolled scan.

We study dental biometric verification under these two requirements. ToothPrint is a research pipeline with three layers. A rigid point-to-surface verifier provides a strong full-coverage baseline. A learned correspondence verifier emits a descriptor per point and verifies partial queries by mutual nearest-neighbor correspondences followed by weighted Procrustes alignment. A split-conformal calibration layer converts similarity scores into accept/abstain decisions with a target FMR.

The contribution is not a new clinical device. All public 3D identity experiments here use synthetic reacquisition of single-session scans. The contribution is a reproducible benchmark and method study showing that (i) certified dental verification is feasible under full coverage, (ii) partial overlap is the dominant failure mode of rigid dental biometrics, and (iii) learned point correspondence substantially raises performance in that regime while still exposing when open-set rejection must abstain.

Contributions.

- We formulate dental identity as open-set verification with a conformal accept/abstain rule, rather than only closed-set retrieval.
- We introduce a point-correspondence verifier for partial dental arches and compare it against rigid GICP and crop-hardened embedding baselines under realistic whole-tooth dropout.
- We report full DET curves, hard-negative calibration, untrained-correspondence controls, and cross-dataset transfer to Teeth3DS+.

- We provide a reproducible open-source implementation, result artifacts, and a clear statement of the data limitation: real cross-session 3D dental identity validation remains the missing gate.

2 Related Work

Dental biometrics. Dental identification is traditionally performed through expert comparison of dental charts, restorations, radiographs, and anatomical features. Dedicated 3D intraoral-scan biometrics is newer. The closest prior system is the digital dental biometrics framework of Zhou et al. [13], which uses FPFH, SAC-IA, ICP, and RMSE on private real re-scan data and reports saturated closed-set performance. We share the registration family but study a different operating point: public data, open-set calibration, partial overlap, and explicit abstention.

3D dental scan datasets. Public intraoral scan datasets have primarily been released for segmentation, labeling, and landmarking rather than biometric reacquisition. Poseidon3D contains challenging orthodontic surface scans with crowding and missing-teeth cases [5]; Teeth3DS+ is a large benchmark for intraoral 3D scan analysis and MICCAI dental challenges [2]. These datasets make dental-shape learning possible, but they do not solve longitudinal identity validation: they are not designed as repeated scans of the same subject across visits.

Point-cloud registration and descriptors. ICP and its variants remain a standard surface-registration tool. Generalized ICP [10] combines point-to-plane and ICP ideas in a probabilistic framework and is a natural baseline for dental surfaces. Learned point-cloud features such as DGCNN/EdgeConv [12] and metric-learning heads such as ArcFace and sub-center ArcFace [3, 4] motivate our embedding baselines. Correspondence networks and transformer-based point matching [7] show that local descriptors can support partial matching; we adapt this idea to dental arches where crop-induced correspondences are available without manual labels.

Open-set recognition and conformal calibration. Open-set recognition formalizes the requirement that a classifier reject queries from classes unseen at enrollment rather than force every input into a known label [9]; certified dental verification is the biometric instance of this reject option. We realize it with conformal prediction, which provides finite-sample, distribution-free calibration under exchangeability [1, 11]. We use split conformal calibration as a decision wrapper over dental similarity scores. The guarantee is intentionally modest: it controls the false-match rate under the calibration distribution, and it should be recalibrated for a new site, scanner, or population.

3 Problem Formulation

Let an enrolled gallery contain subjects $G = \{(y_i, P_i)\}_{i=1}^{N}$, where $P_i \subset \mathbb{R}^3$ is a point cloud sampled from an upper dental arch and y_i is the subject identity. A query $Q \subset \mathbb{R}^3$ may be full coverage or partial. A verifier produces a dissimilarity score

$$s(Q, P_i) \in \mathbb{R}_{\geq 0}, \tag{1}$$

where smaller is more similar. Closed-set retrieval ranks gallery subjects by $s(Q, P_i)$. Open-set verification additionally chooses whether the best candidate is acceptable.

We evaluate three regimes. In *full coverage*, Q is a synthetically reacquired version of the full arch. In *partial overlap*, Q retains only a fraction ρ of teeth or points. We report $\rho = 0.5$ and $\rho = 0.3$ as the main stress tests. In *cross-dataset transfer*, a model trained on Poseidon3D is evaluated on Teeth3DS+ without retraining.

For a threshold τ, accepting the nearest candidate when $\min_i s(Q, P_i) \leq \tau$ induces a false-match rate (FMR) on impostor queries and a false-negative identification rate (FNIR) on genuine queries. We report Rank-1 accuracy, area under the ROC curve (AUC), equal error rate (EER), detection-error-tradeoff (DET) curves, and FNIR at FMR=1%.

4 Methods

4.1 Rigid Surface Verification

The rigid verifier aligns a query Q to each gallery surface P_i. We initialize with PCA principal axes and enumerate proper-rotation hypotheses to reduce sensitivity to the arch's bilateral symmetry. Each candidate is refined with multi-scale GICP. The score is the mean point-to-surface distance from aligned query points to the gallery surface:

$$s_{\text{rigid}}(Q, P_i) = \frac{1}{|Q|} \sum_{q \in Q} d(R^\star q + t^\star, \mathcal{S}(P_i)), \tag{2}$$

where $(R^\star, t^\star)$ is the best rigid transform and $\mathcal{S}(P_i)$ is the triangle or nearest-neighbor surface approximation. We do not allow scale changes; the biometric signal is shape, not size-normalized pose.

4.2 Crop-Hardened Embedding

The embedding baseline uses a DGCNN encoder $f_\theta(P) \in \mathbb{R}^d$ with L2 normalization and a sub-center ArcFace classification head. During training, each arch is randomly cropped with retention $\rho \geq 0.35$ so the descriptor sees partial coverage. At test time, identities are ranked by cosine distance. This baseline tests whether a global descriptor can survive missing teeth without explicit correspondence.

Figure 1: The unified ToothPrint pipeline. A query is detected into an arch point cloud or per-tooth landmarks, retrieved against the gallery by a global embedding, verified on the shortlist by learned point correspondence with a rigid Procrustes residual, and certified by a split-conformal threshold $\hat{\tau}_\alpha$: the nearest candidate is accepted as subject $\hat{i}$ only if its score clears the threshold (false-match rate $\leq \alpha$), otherwise the system abstains.

4.3 Point-Correspondence Verification

The partial-overlap verifier learns local descriptors rather than a single pooled descriptor. Given an arch point cloud $P = \{p_j\}_{j=1}^{m}$, a DGCNN backbone emits unit descriptors $z_j \in \mathbb{R}^d$. Training pairs are generated by cropping the same canonical point set twice; point indices provide positive correspondences for free. For anchor point a and positive point p^+, we use an InfoNCE loss over candidate descriptors:

$$\mathcal{L}_{\mathrm{NCE}} = -\log \frac{\exp(z_a^\top z_{p^+}/T)}{\sum_{b\in\mathcal{B}} \exp(z_a^\top z_b/T)}. \tag{3}$$

At test time, we compute descriptors for Q and P_i, keep mutual nearest-neighbor descriptor matches, estimate a weighted Procrustes transform, and score the residual over all query points:

$$s_{\mathrm{corr}}(Q, P_i) = \frac{1}{|Q|} \sum_{q\in Q} \|R_i q + t_i - \Pi_{P_i}(R_i q + t_i)\|_2 . \tag{4}$$

Scoring all query points, not only matched points, penalizes an impostor that finds a few accidental local correspondences but fails to align the full partial arch.

4.4 Conformal Accept/Abstain Rule

Let $\{u_k\}_{k=1}^{n}$ be calibration impostor scores (smaller is more match-like), assumed exchangeable with the score $u_\star$ of a new impostor query. To control the false-match rate at level α, we set the accept threshold to the order statistic

$$\hat{\tau}_\alpha = u_{(\lfloor (n+1)\alpha \rfloor)}, \tag{5}$$

the $\lfloor (n+1)\alpha \rfloor$-th smallest calibration impostor score. Under exchangeability the rank of $u_\star$ among the $n+1$ scores is uniform, which yields the finite-sample, distribution-free guarantee

$$\Pr\big(u_\star \leq \hat{\tau}_\alpha\big) \leq \alpha. \tag{6}$$

A query is then accepted as subject

$$\hat{i} = \arg\min_i s(Q, P_i) \quad \text{iff} \quad s(Q, P_{\hat{i}}) \leq \hat{\tau}_\alpha, \tag{7}$$

and the system abstains otherwise. The bound is the precise sense in which an accept decision is *certified*: it caps the population false-match rate at the target level from a finite calibration set, with no distributional assumption beyond exchangeability. It is, however, a *marginal* guarantee tied to the calibration distribution—it does not condition on a specific gallery subject, and it must be recalibrated for a new site, scanner, or population. We also evaluate a hard-negative variant in which calibration impostors are each subject's nearest impostors rather than random impostors, which holds the FMR target against the worst-case neighbors at the cost of genuine accept rate.

4.5 Unified Retrieve–Verify–Certify Pipeline

The full pipeline uses the embedding for recall, keeping the top-k gallery candidates, then verifies that shortlist with correspondence scores and applies the conformal threshold. This separates the roles of representation learning: global descriptors retrieve candidates quickly; correspondence supplies the final geometric evidence. Figure 1 summarizes the stages and the accept/abstain decision.

5 Experimental Protocol

Datasets. Poseidon3D provides 3D intraoral surface scans with challenging orthodontic variation [5]. We use 200 arches, with 150 subjects for training learned models and 50 held out for learned partial-overlap evaluation; the rigid full-coverage analysis uses all 200 as a public benchmark. Teeth3DS+ is used only for cross-dataset transfer of the learned correspondence model [2]. DenPAR periapical radiographs support auxiliary 2D landmark-constellation identity experiments [8]. A paired CBCT+IOS dataset supports auxiliary multimodal and restoration-pattern analyses [6].

Synthetic reacquisition. Because public datasets do not provide repeated 3D scans of the same subject

Dataset	Role	Subjects	Modality
Poseidon3D	3D train/test identity	200	IOS mesh
Teeth3DS+	transfer	80 eval	IOS mesh
DenPAR	2D identity	400	radiograph
CBCT+IOS	fusion	55	CBCT + IOS

Table 1: Datasets used in the experiments. The primary paper claim is on 3D dental identity; 2D and fusion results are secondary evidence that the signal appears across dental modalities.

across visits, genuine queries are synthetic reacquisitions. We apply rigid repositioning, acquisition noise, voxel/subsampling changes, and partial crops. We distinguish planar cuts from realistic whole-tooth dropout. The latter removes discrete teeth or tooth regions and is the primary reported partial-overlap setting.

Metrics and leakage control. Learned models are trained on the training subjects only and evaluated on held-out subjects. Cross-dataset transfer uses a Poseidon3D-trained model without Teeth3DS+ retraining. We report mean Rank-1 over repeated crops, AUC, EER, DET curves, and open-set FNIR at FMR=1%. All result JSONs and scripts are committed in the repository.

6 Results

6.1 Full-Coverage Dental Identity

Under full coverage, rigid surface verification is already strong. On 200 Poseidon3D arches, Rank-1 is 0.995, AUC is 0.997, and EER is 0.005. Point-to-surface fidelity separates genuine re-scans from impostors: mean genuine surface distance is 0.060 mm, while the nearest impostor tail starts around 0.82 mm in the point-to-surface analysis. This confirms that the full dental arch is a strong biometric signal.

The 2D radiograph constellation experiment is auxiliary but consistent: on 400 DenPAR subjects, scale-normalized per-tooth landmarks reach Rank-1 1.000 in the main protocol and remain robust to moderate jitter and magnification. We do not present this as a deployment-ready radiograph system; it is evidence that dental geometry and arrangement carry identity beyond 3D meshes.

6.2 Rigid Matching Fails Under Whole-Tooth Dropout

Partial overlap is the dominant failure mode. At 50% retention, rigid GICP falls to Rank-1 0.23; at 30% retention it falls to 0.10. The failure is not only the distance statistic. If the initialization and correspondence are wrong, robust statistics over the wrong alignment cannot recover the identity. Table 2 and Figure 3 show the main comparison.

Method	keep-1.0	keep-0.5	keep-0.3
Rigid GICP	1.00	0.23	0.10
Embedding baseline	0.93	0.46	0.10
Crop-hardened embedding	0.92	0.635	0.26
CorrNet, planar crop	–	0.913	0.800
CorrNet, tooth dropout	–	**0.867**	**0.567**

Table 2: Rank-1 under partial coverage. Whole-tooth dropout is the primary realistic setting. CorrNet improves the 50%-retention regime by 3.8× over rigid GICP.

6.3 Correspondence Helps Beyond Architecture Alone

The untrained correspondence control reaches Rank-1 0.72 at keep-0.5. This is important: some of the gain comes from the architecture of local matching plus rigid verification, not only from learned descriptor semantics. Training still contributes substantially, raising the realistic keep-0.5 result to 0.867 and AUC to 0.984.

Planar crops are easier than tooth dropout. CorrNet reaches 0.913 at keep-0.5 and 0.800 at keep-0.3 under planar cuts, but those numbers overstate performance on missing-teeth queries. We therefore report whole-tooth dropout as the main partial-overlap protocol.

6.4 Open-Set Certification and Abstention

Full-coverage conformal verification is strong. With the unified retrieve–verify–certify pipeline, FNIR at FMR=1% is 0.00 under full coverage. Hard-negative calibration remains close to the target FMR but costs genuine accept rate, as expected. Under 50% tooth retention, FNIR at FMR=1% rises to approximately 0.735. This is not a failure of calibration; it is the correct signal that a half-arch often cannot support safe open-set acceptance. The desired behavior in this regime is abstention rather than a forced identification.

6.5 Cross-Dataset Transfer

CorrNet trained on Poseidon3D transfers imperfectly to Teeth3DS+. At realistic keep-0.5, Rank-1 drops from 0.867 to 0.425; at keep-0.3 it drops to 0.242. These results are still well above chance and above the rigid partial-overlap baseline, but they show that learned local descriptors are dataset-specific. We report this as an experimental finding and a stated limitation, not a hidden weakness.

6.6 Robustness to Sensor Perturbation

The rigid full-coverage pipeline is not fragile to ordinary scanner perturbations in the synthetic protocol. Rank-1 remains 1.0 through 0.4 mm added sensor noise and 4× voxel coarsening after randomized rotation and translation. This distinction matters: the hard case is not small pose noise or modest resolution loss, but missing anatomy.

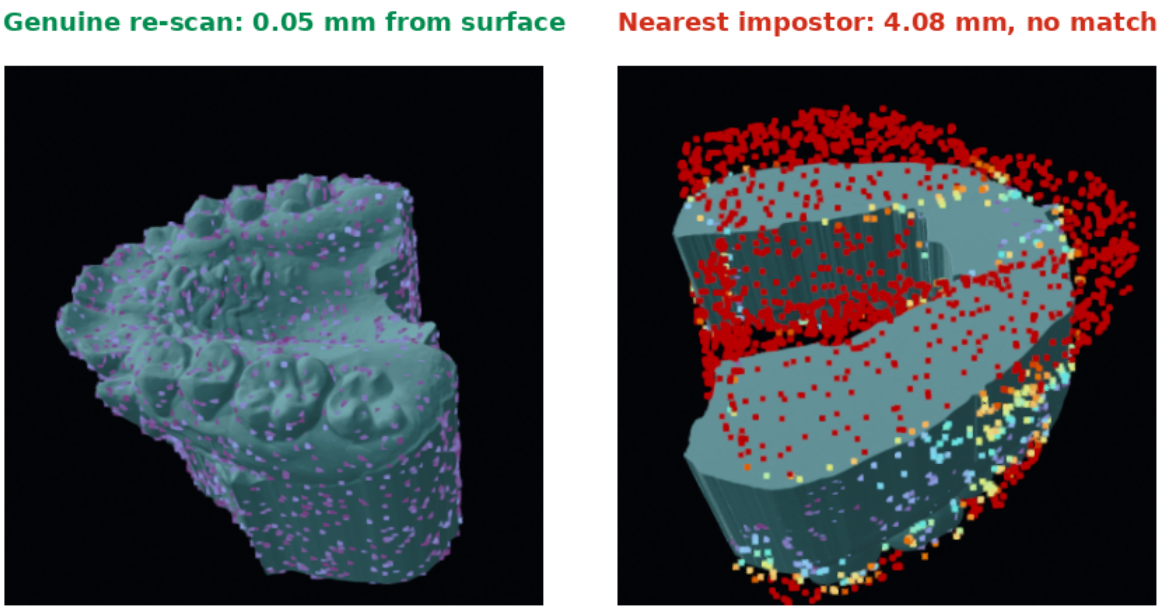

Figure 2: Qualitative identity evidence. A genuine synthetic re-scan aligns tightly to the enrolled arch, while the nearest impostor remains visibly inconsistent. The figure is included to make the biometric evidence inspectable rather than only tabular.

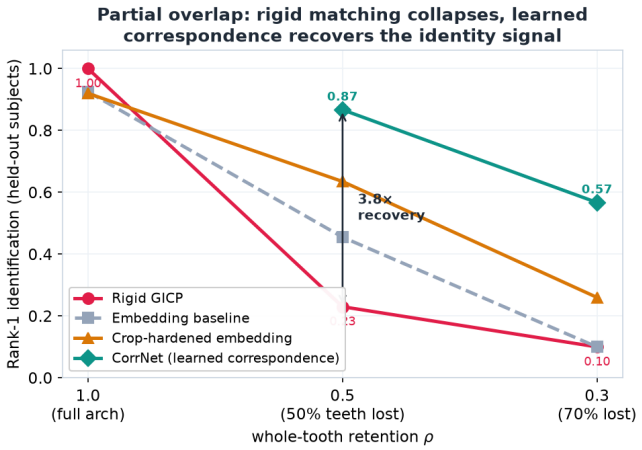

Figure 3: Identification under realistic whole-tooth dropout on held-out subjects. Rigid GICP collapses from Rank-1 1.00 at full coverage to 0.23 at 50% tooth loss, while learned point correspondence (CorrNet) recovers it to 0.867 (a 3.8× improvement) and degrades gracefully to 0.567 at 70% loss. Values match Table 2.

The method therefore focuses learning capacity on partial overlap rather than replacing a rigid matcher that is already adequate when the full arch is present.

6.7 Auxiliary Dental-Work and Multimodal Evidence

Restoration patterns provide an additional forensic cue. On CBCT+IOS paired data, CBCT dental-work geometry reaches Rank-1 0.927. On DenPAR radiographs, a per-tooth local-contrast restoration extractor identifies 165 restoration-bearing subjects with Rank-1 0.91–0.99 under jitter/dropout perturbations. These are auxiliary results because they rely on restoration-bearing subsets and synthetic reacquisition, but they support the broader claim that dental identity is multi-cue.

Multimodal fusion is complementary but not uniformly beneficial. In high-quality paired data, IOS crowns already saturate. In a degraded hard regime, simple fusion can dilute evidence. Quality weighting improves Rank-1 in some settings but can regress AUC. We therefore keep fusion out of the primary claim.

7 Ablations and Negative Results

Partial-overlap protocol. Planar crops are too easy. Whole-tooth dropout better reflects missing teeth and field-of-view artifacts; it lowers keep-0.3 from 0.800 to 0.567.

Global descriptors. Crop-hardening improves global embeddings under tooth loss, but it does not close the gap. At keep-0.5 it reaches 0.635 compared with 0.867 for correspondence verification.

Open-set partial identity. Even the best partial-overlap matcher cannot make a half-arch uniquely identifiable in every open-set case. Reporting FNIR at controlled FMR is therefore more informative than reporting only closed-set Rank-1.

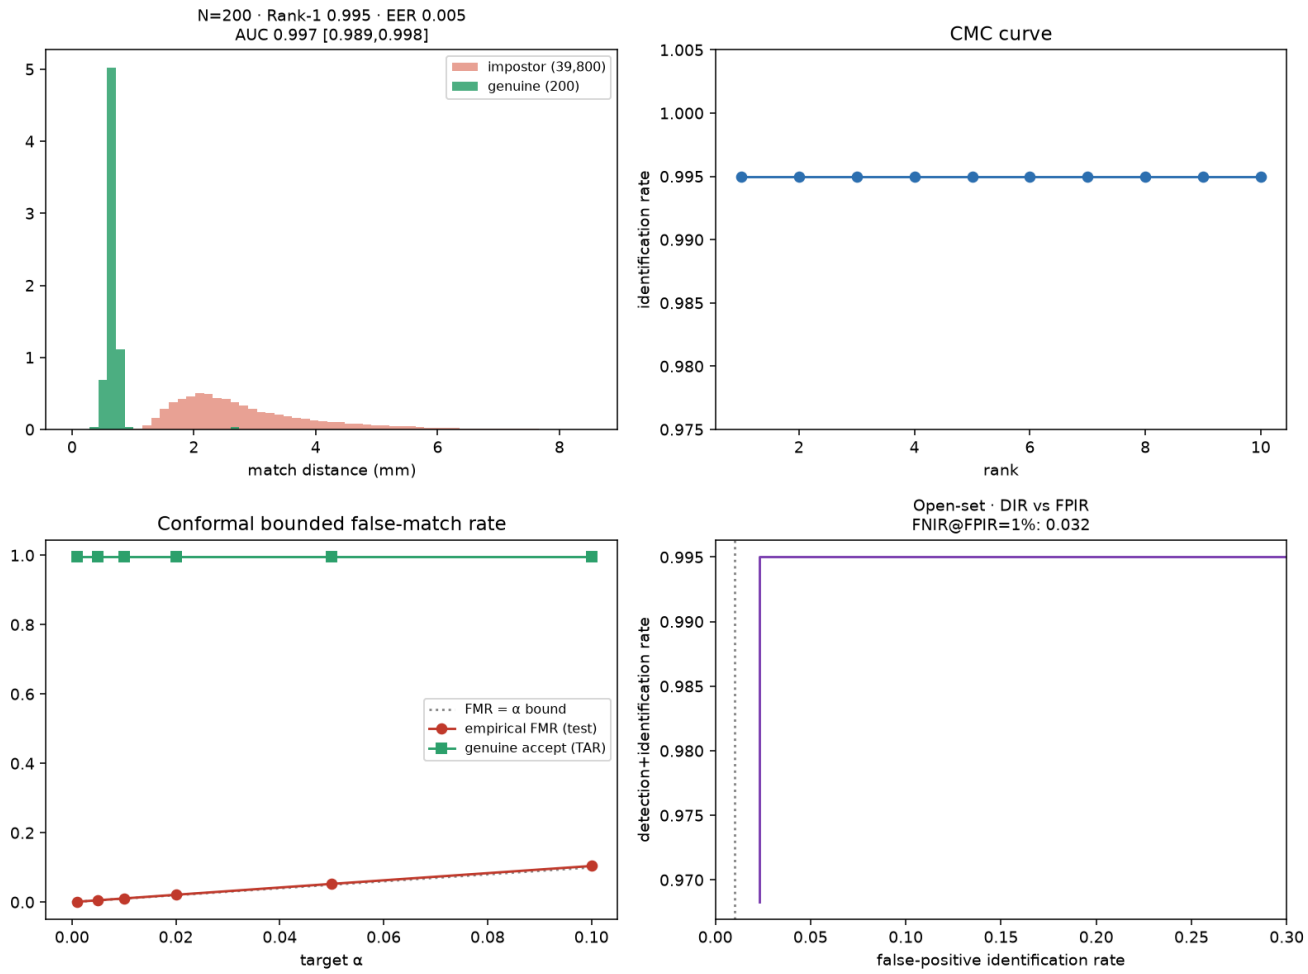

Figure 4: Full-coverage identity metrics: score separation, CMC, conformal FMR tracking, and open-set decision behavior on Poseidon3D.

8 Limitations

No public real cross-session 3D identity benchmark. The largest limitation is data. The experiments synthesize reacquisition from single-session public scans. This permits controlled method comparison but cannot establish real longitudinal performance across scanners, operators, patient changes, restorations, or time.

Domain shift. The Teeth3DS+ transfer result shows that learned dental descriptors should be trained and calibrated across multiple acquisition sources before any serious deployment study.

Clinical and forensic use. This work does not support autonomous clinical or forensic decisions. A real deployment would require consented galleries, site-specific calibration, access control, expert review, prospective validation, and regulatory governance.

9 Ethics and Security

Dental biometrics are sensitive personal data. A false identification can have severe forensic or clinical consequences. The system is designed to abstain, to show score evidence, and to avoid silent fallbacks, but those design choices do not remove the need for human review. The released code stores no patient database and includes hardened medical-file loaders; deployment security, retention policy, and legal basis remain the responsibility of any integrator.

10 Reproducibility

The repository includes source code, evaluation scripts, result JSONs, figures, and a fixture smoke test that runs without off-machine data. External datasets are license-

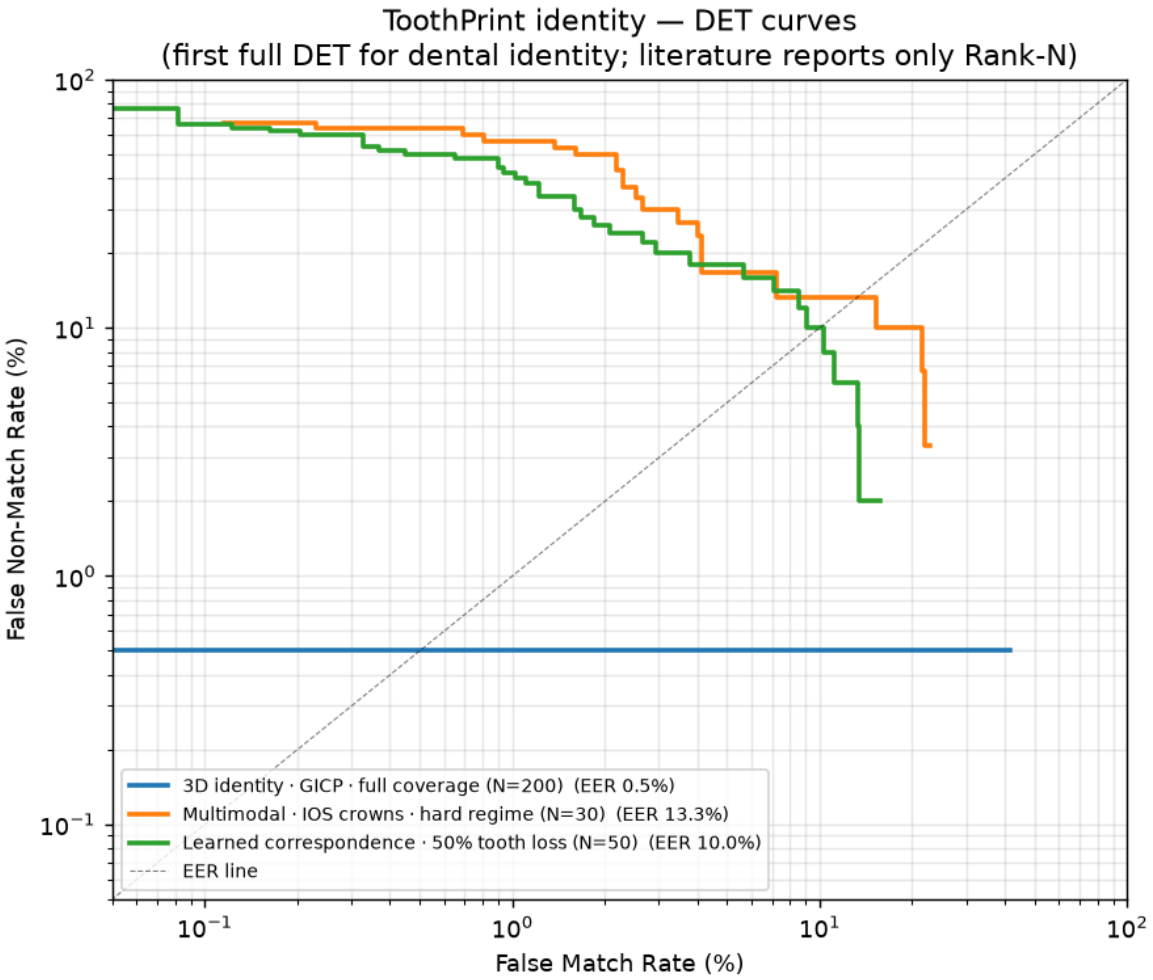

Figure 5: Detection-error tradeoff curves for identity pillars. The full-coverage 3D curve is near-saturated; learned correspondence under 50% tooth loss remains usable but exposes the open-set difficulty.

gated or large and are intentionally not committed. Paths are configured through environment variables, and the result tables in this manuscript correspond to committed artifacts under `evaluation/results`. The paper should be read together with `REPRODUCE.md`.

Code and data availability. The implementation, evaluation scripts, and result artifacts are available at https://github.com/Archerkattri/toothprint; every table in this manuscript corresponds to a committed file under `evaluation/results`. The external datasets (Poseidon3D, Teeth3DS+, DenPAR, and the paired CBCT+IOS set) are obtained from their original providers under their respective licenses and are not redistributed here.

11 Conclusion

Dental arches are strong biometric signals under full coverage, but the research problem becomes more interesting and more realistic under partial overlap and open-set operation. ToothPrint shows that learned point correspondence can recover much of the identity signal lost by rigid registration under missing teeth, while conformal calibration turns retrieval scores into accept/abstain decisions with an explicit FMR target. The contribution is concrete: it identifies an open gap in dental biometrics—certified partial-overlap verification—and supports it with reproducible evidence, ablations, and negative results. Its central open requirement is equally clear: real cross-session dental data is needed before any clinical or forensic claim can be made.

References

[1] Anastasios N. Angelopoulos and Stephen Bates. Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning*, 16(4):494–591, 2023. doi: 10.1561/2200000101.

[2] Achraf Ben-Hamadou, Nour Neifar, Ahmed Rekik, Oussama Smaoui, Firas Bouzguenda, Sergi Pujades, Edmond Boyer, and Edouard Ladroit. Teeth3DS+: An extended benchmark for intraoral 3D scans analysis. *arXiv preprint arXiv:2210.06094*, 2022.

[3] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.

[4] Jiankang Deng, Jia Guo, Tongliang Liu, Mingming Gong, and Stefanos Zafeiriou. Sub-center ArcFace: Boosting face recognition by large-scale noisy web faces. In *European Conference on Computer Vision*, 2020.

[5] Tibor Kubik and Michal Spanel. LMVSegRNN and Poseidon3D: Addressing challenging teeth segmentation cases in 3D dental surface orthodontic scans. *Bioengineering*, 11(10):1014, 2024. doi: 10.3390/bioengineering11101014.

[6] Xiang Li. 3D Multimodal Dental Dataset Based on CBCT and Oral Scan, 2024.

[7] Zheng Qin, Hao Yu, Changjian Wang, Yulan Guo, Yuxing Peng, and Kai Xu. Geometric Transformer for fast and robust point cloud registration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11143–11152, 2022.

[8] Sumudu Rasnayaka, Dhanushka Leuke Bandara, Amali Jayasundara, Ruwan Jayasinghe, Chathura Wimalasiri, Piumal Rathnayake, Shamod Wijerathne, Roshan Ragel, Vajira Thambawita, and Isuru Nawinne. DenPAR: Annotated intra-oral periapical radiographs dataset for machine learning. *Scientific Data*, 12:1615, 2025. doi: 10.1038/s41597-025-05906-9.

[9] Walter J. Scheirer, Anderson de Rezende Rocha, Archana Sapkota, and Terrance E. Boult. Toward open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(7):1757–1772, 2013.

[10] Aleksandr Segal, Dirk Haehnel, and Sebastian Thrun. Generalized-ICP. In *Robotics: Science and Systems*, 2009. doi: 10.15607/RSS.2009.V.021.

[11] Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer, 2005.

[12] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph CNN for learning on point clouds. *ACM Transactions on Graphics*, 38(5):146, 2019. doi: 10.1145/3326362.

[13] Yu Zhou, Li Yuan, Yanfeng Li, and Jiannan Yu. Digital dental biometrics for human identification based on automated 3D point cloud feature extraction and registration. *Bioengineering*, 11(9):873, 2024. doi: 10.3390/bioengineering11090873.