UNIMORE
UNIVERSITÀ DEGLI STUDI DI
MODENA E REGGIO EMILIA

XXXVIII cycle of the International Doctorate School in
"Information and Communication Technologies (ICT)"

Scaling AI for Oral and Dental Image Analysis

Candidate
Luca Lumetti

Supervisor
Prof. Costantino Grana

Co-Supervisor
Prof. Federico Bolelli

Coordinator
Prof. Luigi Rovati

Oral diseases are common, costly, and under-served by AI

3.5 billion

People affected by oral diseases in 2019

> $380 B.

Spent on the main oral diseases in 2019

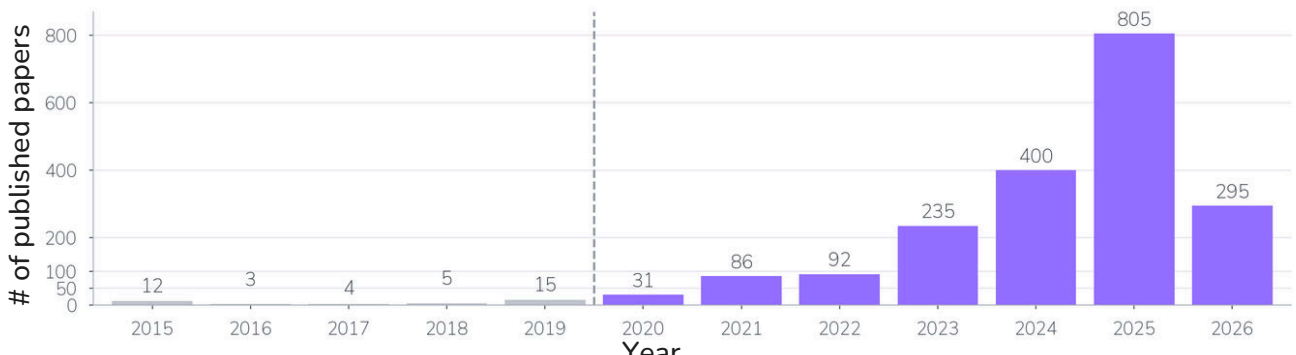

Sources: https://www.who.int/data/gho/data/themes/oral-health-data-portal, https://scopus.com

AI research starts from (open) data

Author	Year	#data	Public	Country
Jaskari et al.	2020	637	N	Finland
Jang et al.	2021	38	N	China
Lahoud et al.	2022	235	N	Belgium
Usman et al.	2022	1010	N	S. Korea
Cipriano et al.	2022	347	Y	Italy
Cui et al.	2022	4938	N	China
Cui et al.	2022	22	Y	China
Dou et al.	2022	40	N	China
Chun et al.	2023	50	N	S. Korea

- Medical datasets are difficult to release because of privacy and ethical constraints.

- Without public data, reproducibility and fair comparison are limited.

- Private studies are still valuable, but open benchmarks are essential for progress.

Data requires annotations

Providing accurate labels for medical data requires high expertise in the field and performing 3D annotations is hard and very time-consuming.

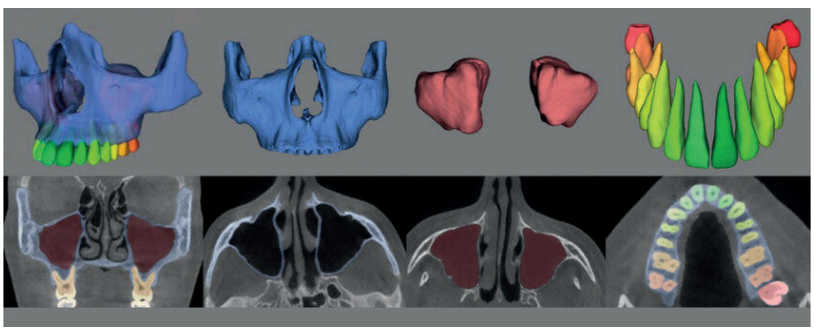

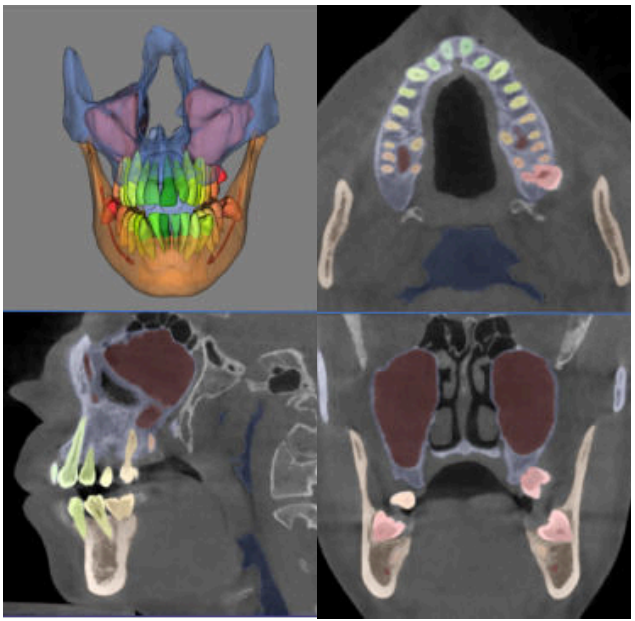

Images source: "The Use of Artificial Intelligence in Third Molar Surgery Risk Assessment", F. Van der Cruyssen, Dental Update

Annotations require specific tools

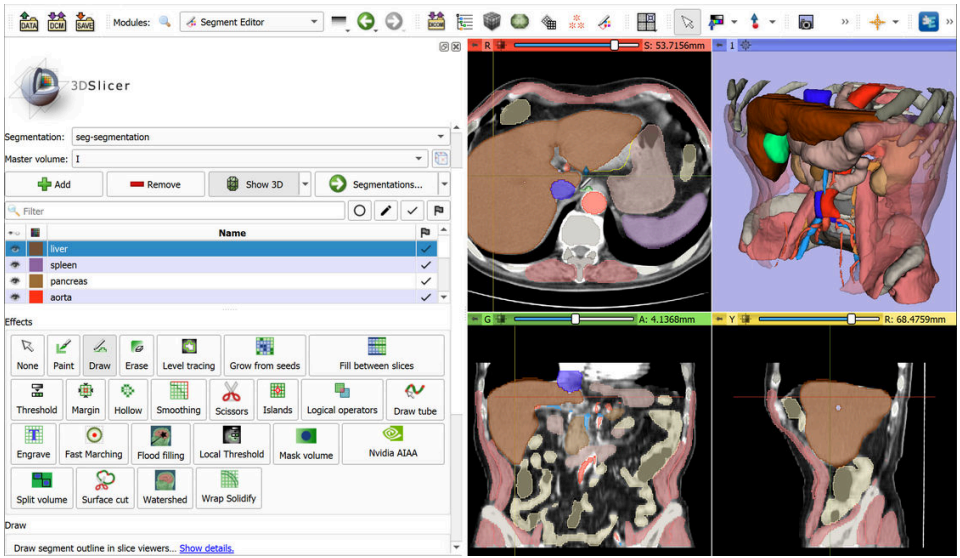

Open tools for annotating medical data do exist, but they are rarely specific for a given field, and rather general for 3D medical data.

Specific structures require high precision and specialized tools can simplify the annotation process, reducing the costs.

Scaling dental AI by jointly improving annotation tools, datasets, and models

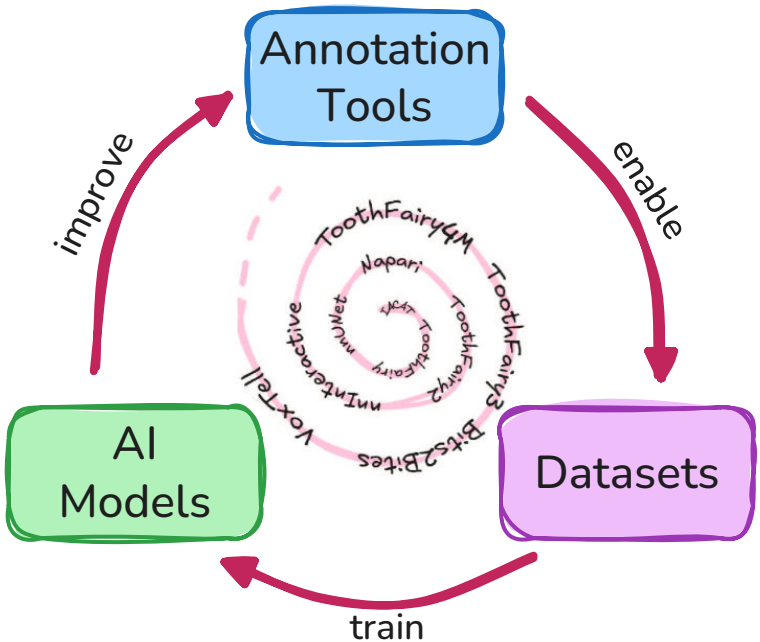

Where do we start?

Author	Year	#data	Public	Country
Jaskari et al.	2020	637	N	Finland
Jang et al.	2021	38	N	China
Lahoud et al.	2022	235	N	Belgium
Usman et al.	2022	1010	N	S. Korea
Cipriano et al.	2022	347	Y	Italy
Cui et al.	2022	4938	N	China
Cui et al.	2022	22	Y	China
Dou et al.	2022	40	N	China
Chun et al.	2023	50	N	S. Korea

We started from a limited public dataset: only 91/347 volumes had 3D IAC labels, yet this was enough to bootstrap the first tool/model loop.

IAC Annotation Tool v2

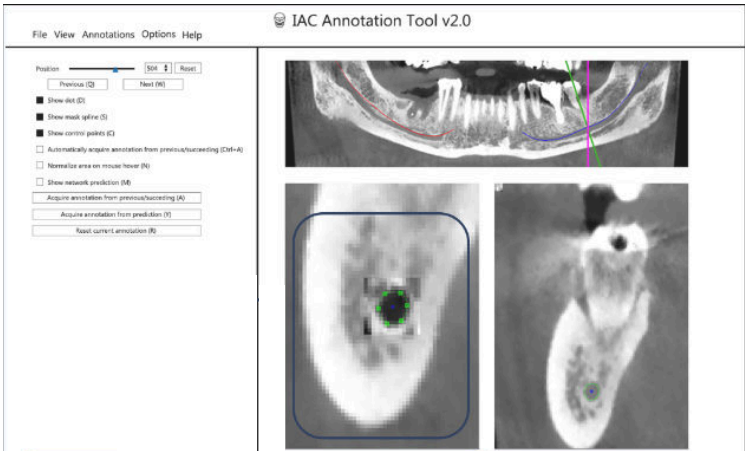

We corrected previously created annotations in regions that were difficult to identify by improving the IACAT software and including a deep learning model as support for the segmentation.

- Corrected 40 under-segmented volumes.
- Added 62 new 3D annotations.

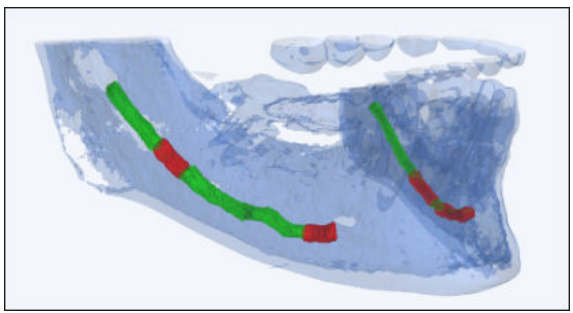

[1] Lumetti, L., Pipoli, V., Bolelli, F., & Grana, C. (2023, September). Annotating the Inferior Alveolar Canal: The Ultimate Tool. In International Conference on Image Analysis and Processing (pp. 525-536). Cham: Springer Nature Switzerland.

[2] Lumetti, L., Pipoli, V., Bolelli, F., Ficarra, E., & Grana, C. (2024, December). Location Matters: Harnessing Spatial Information to Enhance the Segmentation of the Inferior Alveolar Canal in CBCTs. In International Conference on Pattern Recognition (pp. 108-123). Cham: Springer Nature Switzerland.

ToothFairy: A Cone-Beam Computed Tomography Segmentation Challenge

MICCAI 2023 Vancouver CANADA
TOOTH FAIRY

Final Rank	ID	Participant	Country	Dice		HD95	
				Score	Rank	Score	Rank
1	A	Liu et al.	China	0.796 ± 0.093	1	4.49 ± 6.08	1
2	B	Wang et al.	China	0.769 ± 0.106	2	7.45 ± 10.90	2
3	C	Wodzinski et al.	Switzerland	0.760 ± 0.088	3	9.22 ± 14.40	3
4	D	Kirchoff et al.	Germany	0.739 ± 0.146	4	13.90 ± 28.70	5
5	E	Huang	China	0.715 ± 0.110	8	13.10 ± 13.50	4
6	F	Su	China	0.734 ± 0.099	5	18.00 ± 25.30	8
7	G	Ye	China	0.728 ± 0.094	7	18.90 ± 25.00	9
8	H	Yang	China	0.731 ± 0.151	6	28.60 ± 47.40	12
8	I	Han et al.	South Korea	0.700 ± 0.148	11	15.70 ± 24.80	7
10	J	Szczepański et al.	Poland	0.675 ± 0.208	13	14.90 ± 29.30	6
11	K	Wu	China	0.712 ± 0.142	10	19.40 ± 35.60	10
11	L	Zheng	China	0.713 ± 0.129	9	19.50 ± 28.00	11
*13	M	Lumetti et al.	Italy	0.699 ± 0.151	12	38.60 ± 48.00	13
14	N	Han et al.	South Korea	0.643 ± 0.179	14	40.20 ± 49.70	14
15	O	Huang	China	0.642 ± 0.187	15	inf	17
16	P	Szczepański et al.	Poland	0.613 ± 0.201	16	inf	17
17	Q	Pang	China	0.612 ± 0.178	17	inf	17
18	R	Gamal	Egypt	0.326 ± 0.172	21	42.70 ± 43.90	15
18	S	Li	China	0.327 ± 0.218	20	65.90 ± 48.10	16
18	T	Caselles Ballester et al.	Spain	0.507 ± 0.209	19	inf	17
21	U	Kirchoff et al.	Germany	0.565 ± 0.259	18	inf	21

ToothFairy set the first public benchmark for segmentation of the Inferior Alveolar Canal.

Challenge test-set is still private and only accessible via Grand-Challenge.

[3] Bolelli, F., Lumetti, L., Vinayahalingam, S., Di Bartolomeo, M., Pellacani, A., Marchesini, K., ... & Grana, C. (2024). Segmenting the Inferior Alveolar Canal in CBCT Volumes: the ToothFairy Challenge. IEEE Transactions on Medical Imaging, 44(4), 1890-1906.

[4] Lumetti, L., Pipoli, V., Bolelli, F., Ficarra, E., & Grana, C. (2024). Enhancing patch-based learning for the segmentation of the mandibular canal. IEEE Access, 12, 79014-79024.

ToothFairy: A Cone-Beam Computed Tomography Segmentation Challenge

MICCAI 2023 Vancouver CANADA
TOOTH FAIRY

2778
Downloads

88
Citations

40
Countries

36
Submissions

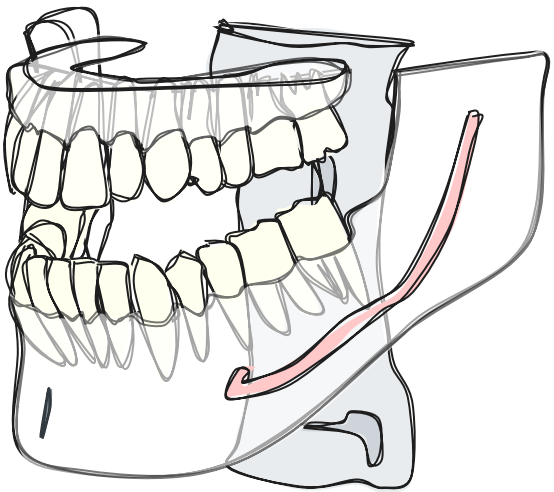

Source: https://ditto.ing.unimore.it/dataset_view/?dataset=toothfairy

ToothFairy2: Multi-Structure Segmentation in CBCT Volumes

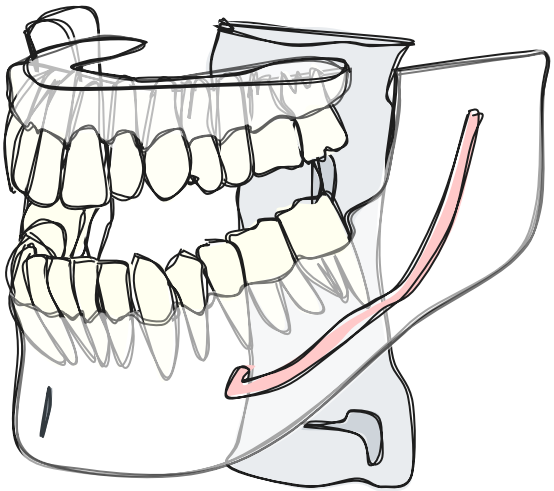

2 labels per volume

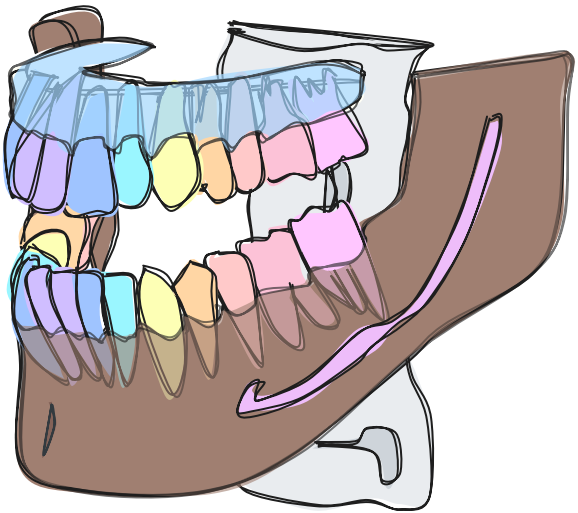

42 labels per volume

ToothFairy2: Multi-Structure Segmentation in CBCT Volumes

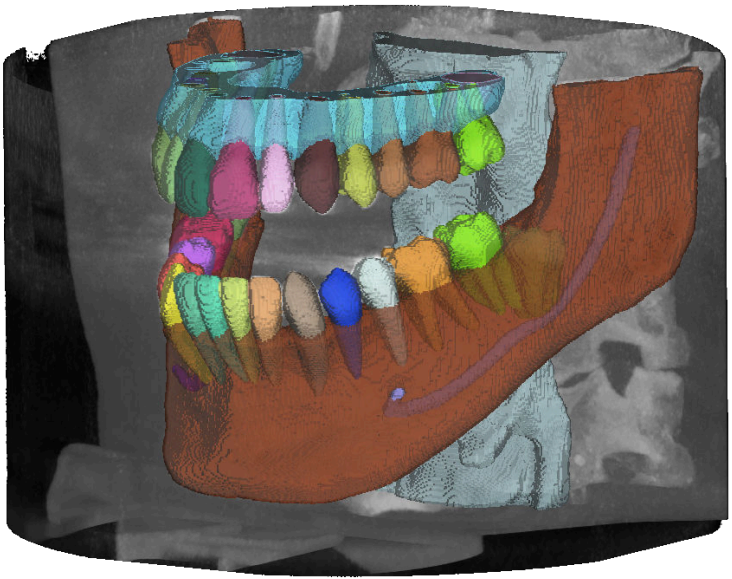

- 480 volumes, all with 3D annotations of all 42 structures

- From 306 segmentation masks in ToothFairy1 to more than 15,000 in ToothFairy2.

ToothFairy2: Multi-Structure Segmentation in CBCT Volumes

Group	Metric	(A) F. Isensee, Y. Kirchhoff et al.	(B) Y. Jiang et al.	(C) H. Wang et al.	(D) A. Gao et al.	(E) P. Shi et al.
Overall	DSC	**0.925 ± 0.171**	0.917 ± 0.175	0.911 ± 0.188	0.911 ± 0.171	0.902 ± 0.207
	HD95	18.869 ± 106.469	17.647 ± 102.822	**17.564 ± 96.074**	18.576 ± 106.897	26.794 ± 130.444
L/R IAC	DSC	0.897 ± 0.047	0.891 ± 0.040	**0.898 ± 0.037**	0.872 ± 0.054	0.892 ± 0.040
	HD95	1.606 ± 2.034	1.548 ± 1.118	**1.474 ± 0.911**	2.022 ± 2.146	1.557 ± 1.028
L/R Sinus	DSC	**0.891 ± 0.214**	0.813 ± 0.265	0.837 ± 0.265	0.836 ± 0.238	0.886 ± 0.226
	HD95	33.724 ± 142.266	**30.996 ± 109.426**	38.994 ± 138.839	37.028 ± 141.887	49.945 ± 156.102
Teeth	DSC	**0.930 ± 0.171**	0.925 ± 0.170	0.915 ± 0.191	0.918 ± 0.166	0.904 ± 0.211
	HD95	17.878 ± 106.926	16.930 ± 106.305	18.603 ± 101.543	**16.892 ± 106.509**	25.851 ± 132.694
Jawbones	DSC	0.931 ± 0.086	0.946 ± 0.065	0.956 ± 0.050	0.938 ± 0.050	**0.957 ± 0.044**
	HD95	14.492 ± 27.625	6.901 ± 17.914	**4.202 ± 12.159**	12.258 ± 21.708	4.693 ± 14.132
Pharynx	DSC	**0.955 ± 0.027**	0.944 ± 0.039	0.951 ± 0.032	0.947 ± 0.033	0.936 ± 0.051
	HD95	**3.911 ± 3.182**	5.308 ± 5.571	4.061 ± 4.466	4.168 ± 4.075	9.982 ± 23.277
Others	DSC	**0.896 ± 0.241**	0.887 ± 0.240	0.894 ± 0.229	0.883 ± 0.259	0.851 ± 0.288
	HD95	38.479 ± 148.126	38.172 ± 136.603	**16.670 ± 77.282**	43.731 ± 155.026	58.273 ± 178.434

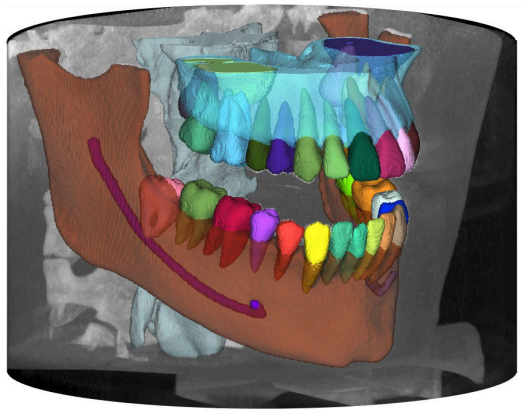

[5] Lumetti, L., Bolelli, F., Van Nistelrooij, N., ... & Grana, C. (2026). Multi-Structure Segmentation in CBCT Volumes: the ToothFairy2 Challenge. Medical Image Analysis, 2026 (Under Review)

[6] Bolelli, F., Marchesini, K., Van Nistelrooij, N., Lumetti, L., Pipoli, V., Ficarra, E., ... & Grana, C. (2025). Segmenting maxillofacial structures in cbct volumes. In Proceedings of the Computer Vision and Pattern Recognition Conference (pp. 5238-5248).

ToothFairy2: Multi-Structure Segmentation in CBCT Volumes

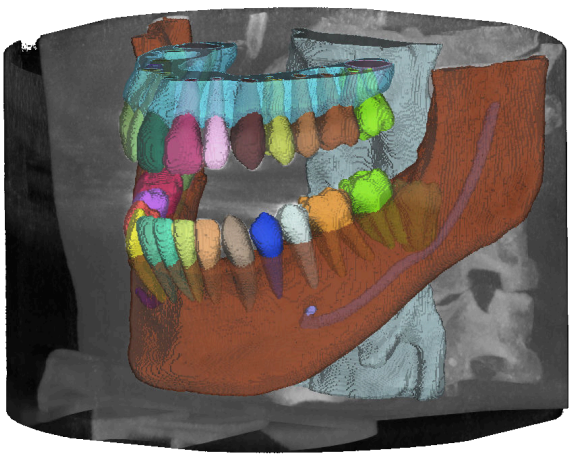

4202
Downloads

37
Citations

49
Countries

30
Submissions

Source: https://ditto.ing.unimore.it/dataset_view/?dataset=toothfairy2

ToothFairy3: Multi-class and Interactive segmentations

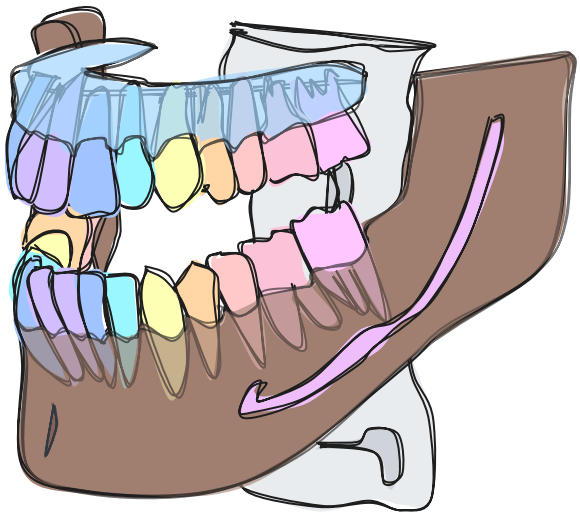

42 labels per volume

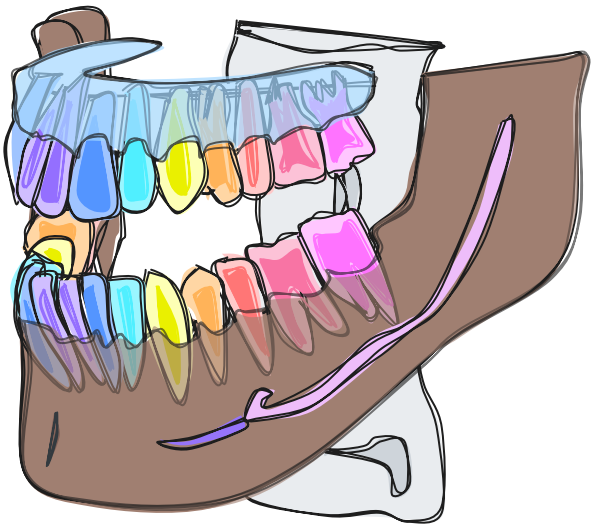

77 labels per volume

ToothFairy3: Fast Multiclass and Interactive Segmentations

MICCAI2025 Daejeon REPUBLIC OF KOREA
TOOTH FAIRY3

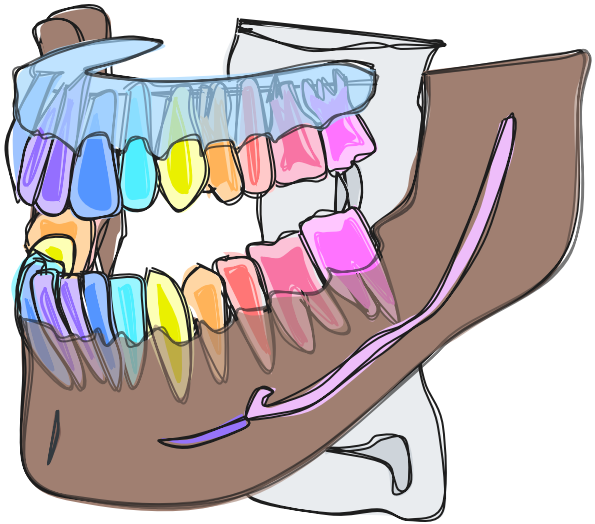

77 labels per volume

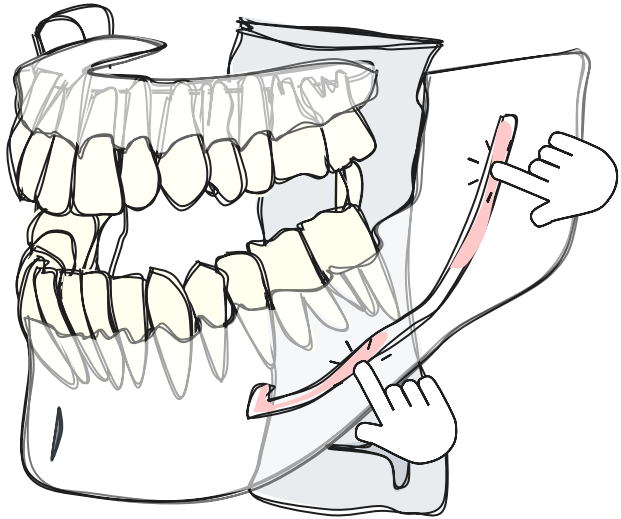

2 labels + simulated clicks

ToothFairy3: Multi-class and Interactive segmentations

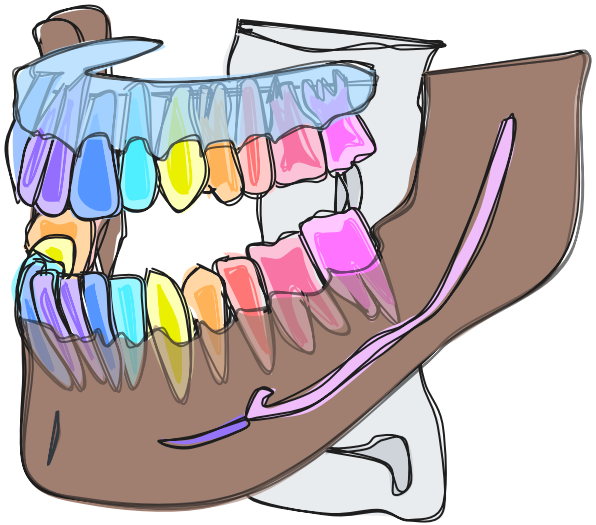

4886
Downloads

NA
Citations

50
Countries

22
Submissions

Source: https://ditto.ing.unimore.it/dataset_view/?dataset=toothfairy3

Efficient Segmentation Model for 77 Classes

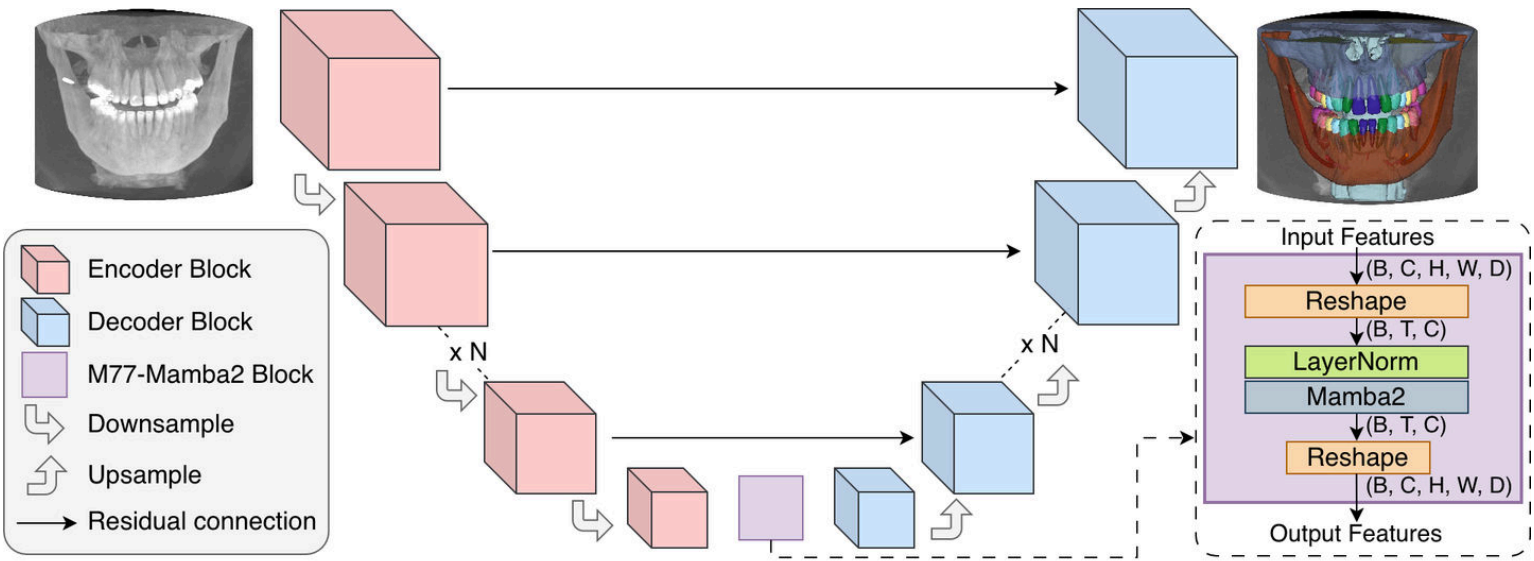

[7] Lumetti, L., Qin Tan, Z., Borghi, L., Addison O., Li, Y., Rosati, G., ... & Bolelli, F. Maxillo-77: Scaling CBCT Maxillofacial Segmentation to 77 Classes with M77-Mamba2. MICCAI 2026. (Under Review)

ODIN 2025: Oral and Dental Image aNalysis

- We organized the ODIN 2025 workshop to help build a research community around AI for dentistry.

- It also provided an opportunity to organize talks by inviting high-profile speakers from the medical AI field.

- International collaboration across academia and industry.

- Workshop proceedings and sponsored support.

Website: https://odin-workshops.org/2025/

ToothFairy4M – Web-based data collection and annotations

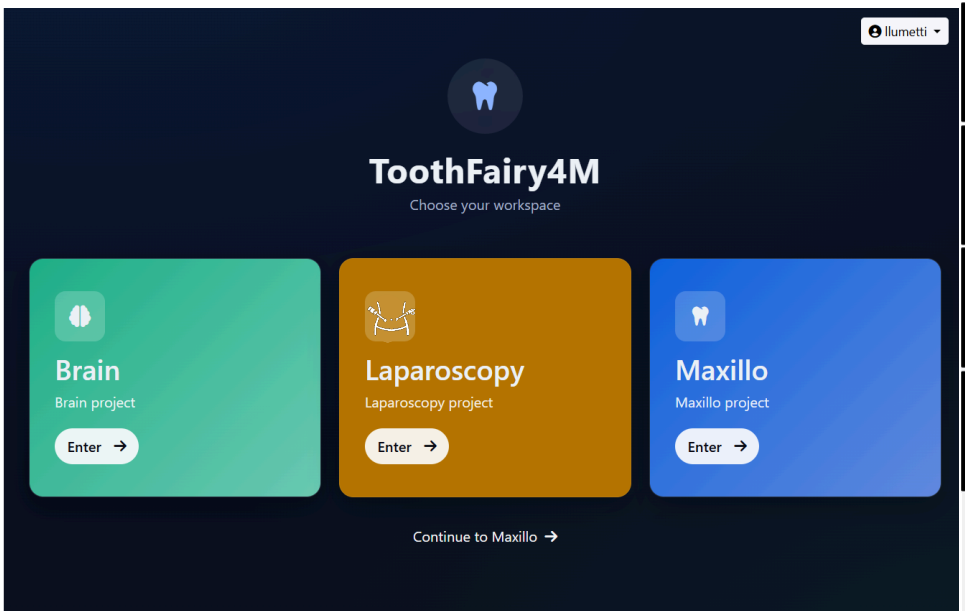

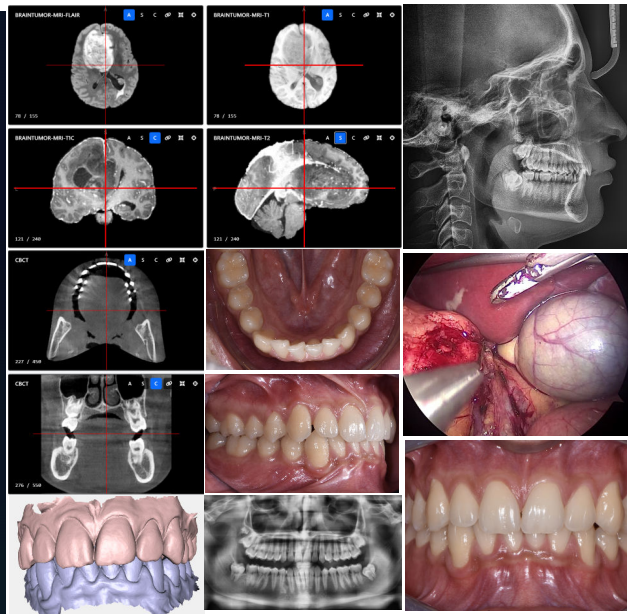

Website: https://toothfairy4m.ing.unimore.it

Intraoral scans open a second major research line

Ferrara School of Orthodontics

In orthodontics, intraoral scans (IOS) provide a complementary modality to CBCT.

They are fast, low-cost, and radiation-free, which makes them routine in clinical practice.

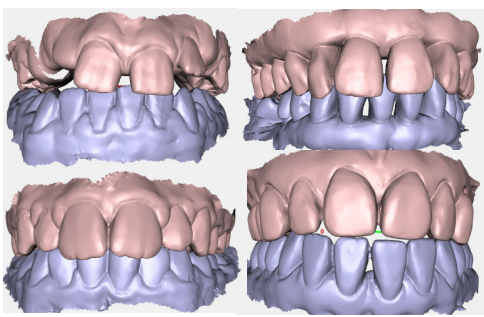

Open data for intraoral-scans.

Name	Year	#data	Public	Country
Ben-Hamadou	2020	1800	Y	Tunisia
Vinayahalingam	2021	1750	N	Netherlands

They also organized MICCAI Challenges!

Bits2Bites: Intraoral Scans Occlusal Classification

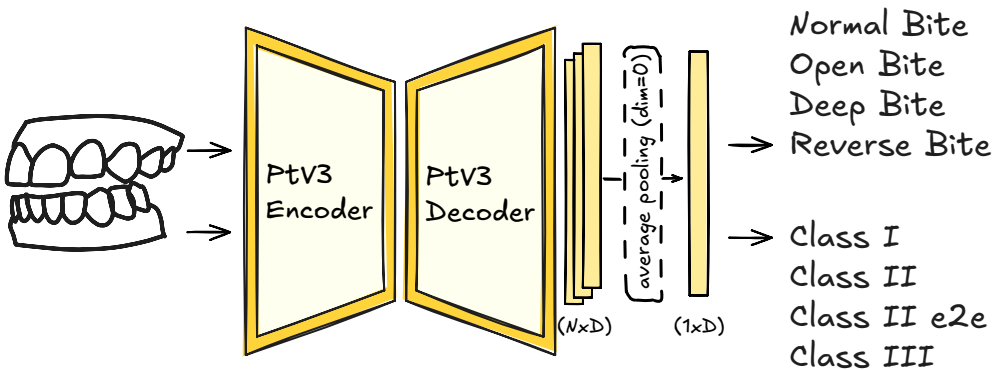

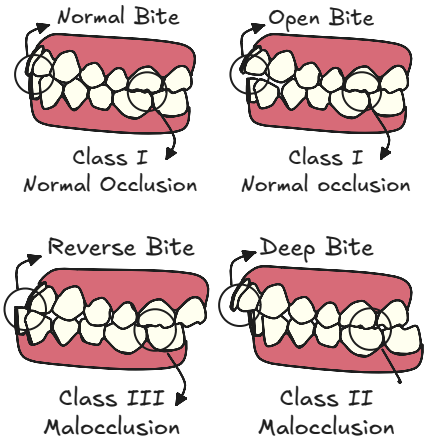

[8] Lumetti, L., Borghi, L., Cremonini, F., Rizzo, F., Grana, C., Lombardo, L., & Bolelli, F. (2025). Bits2Bites: Intra-oral Scans Occlusal Classification. In Oral and Dental Image aNalysis Workshop.

Brackets – Orthodontic Bracket Placement

- Bracket placement is a routine but time-consuming orthodontic task.

- Even for one arch, it can take more than 15 minutes for a single expert.

- OrtoFE have a digital tool for manual bracket placement.

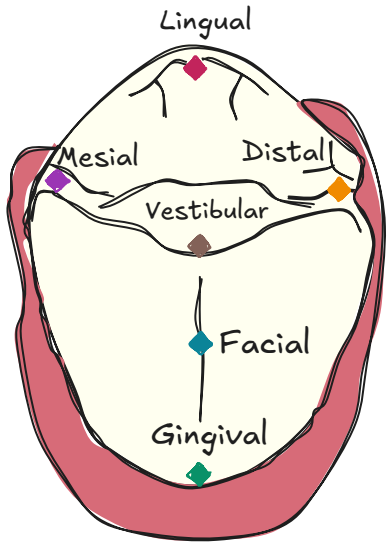

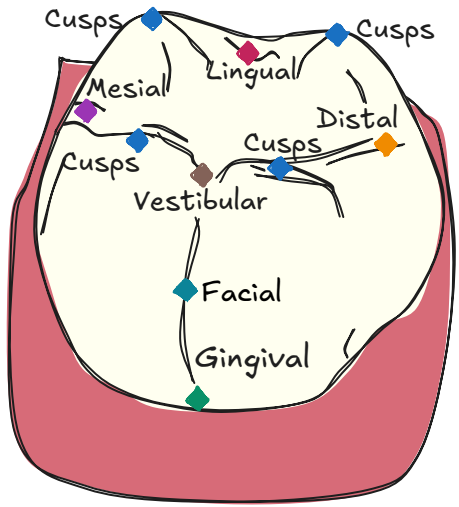

Automated Intraoral Scan Analysis for Accurate Orthodontic Bracket Placement

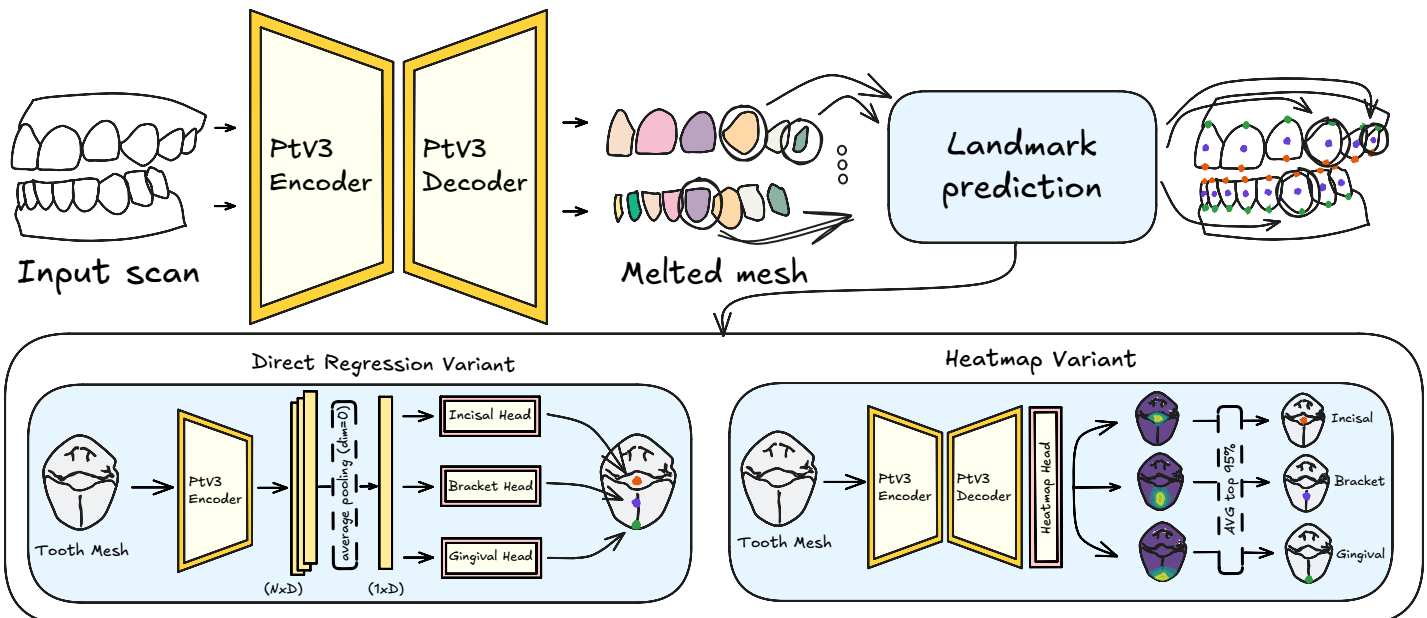

[9] Lumetti, L., Lugli, M., Cremonini, F., Brucculeri, L., Grana, C., Lombardo, L., Bolelli, F. Automated Intra-Oral Scan Analysis for Accurate Orthodontic Bracket Placement. MICCAI 2026. (Under Review)

AutoBonding – Closing the loop between research and clinic

- Deployed in a real digital workflow.
- Collects corrections from orthodontists.
- Supports continuous retraining.
- Exposes edge cases early.

The time was reduced from more than 15 minutes to 4 minutes.

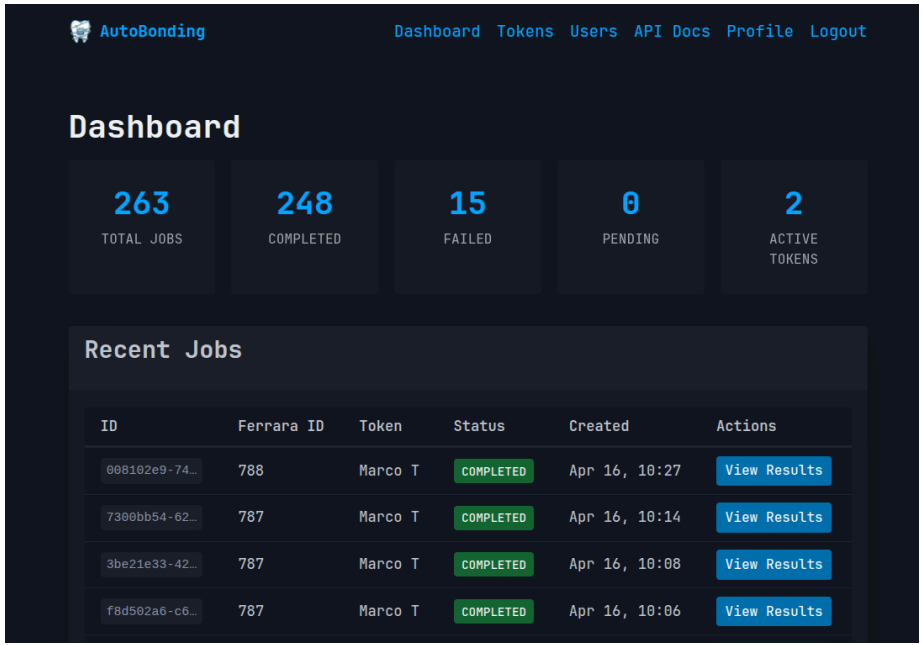

Website: https://autobonding.ing.unimore.it

U-Net Transplant: The Role of Pre-training for Model Merging in 3D Medical Segmentation

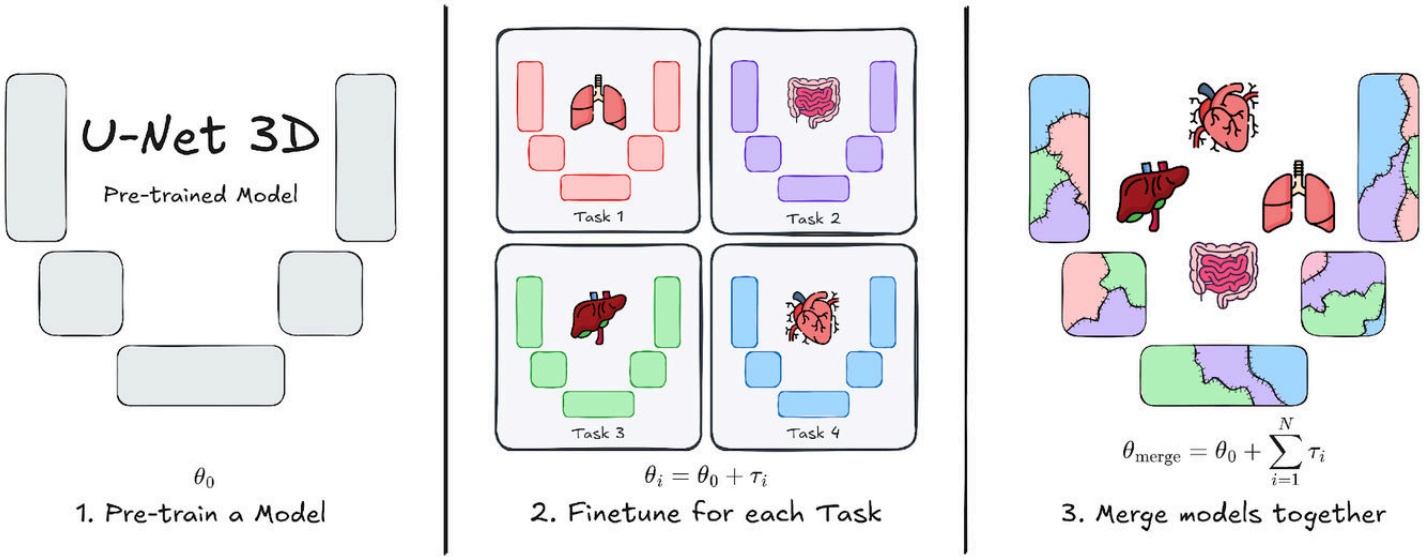

[10] Lumetti, L., Capitani, G., Ficarra, E., Calderara, S., Grana, C., Porrello, A., & Bolelli, F. (2025, September). U-Net transplant: the role of pre-training for model merging in 3D medical segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention (pp. 618-628). Cham: Springer Nature Switzerland.

Model Merging and Task Arithmetic

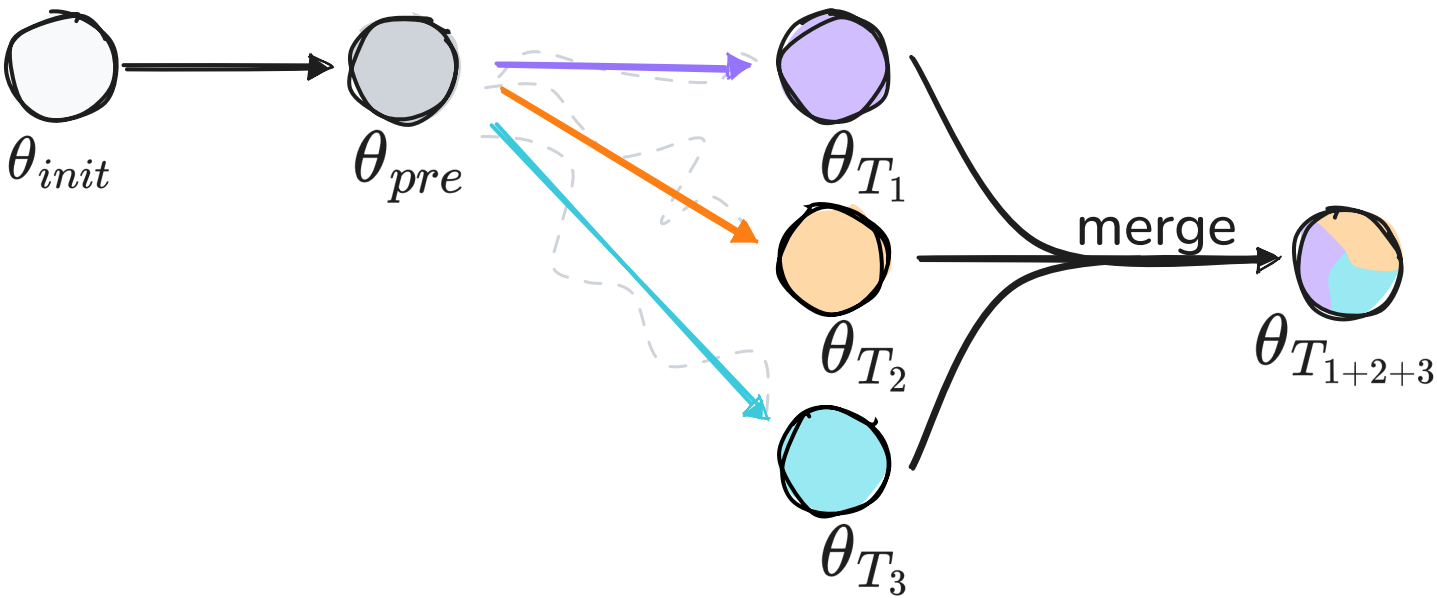

Stable Pre-training Allows Better Merging

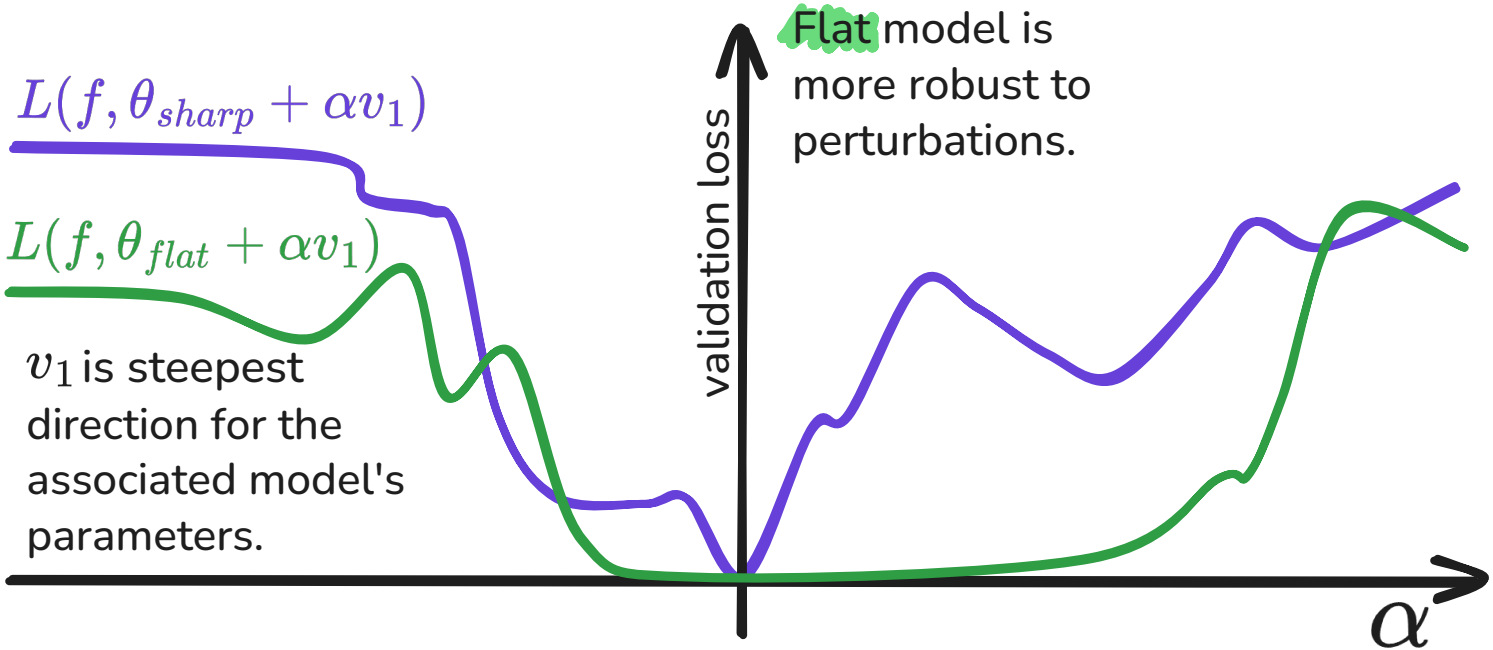

Stable Pre-training Allows Better Merging

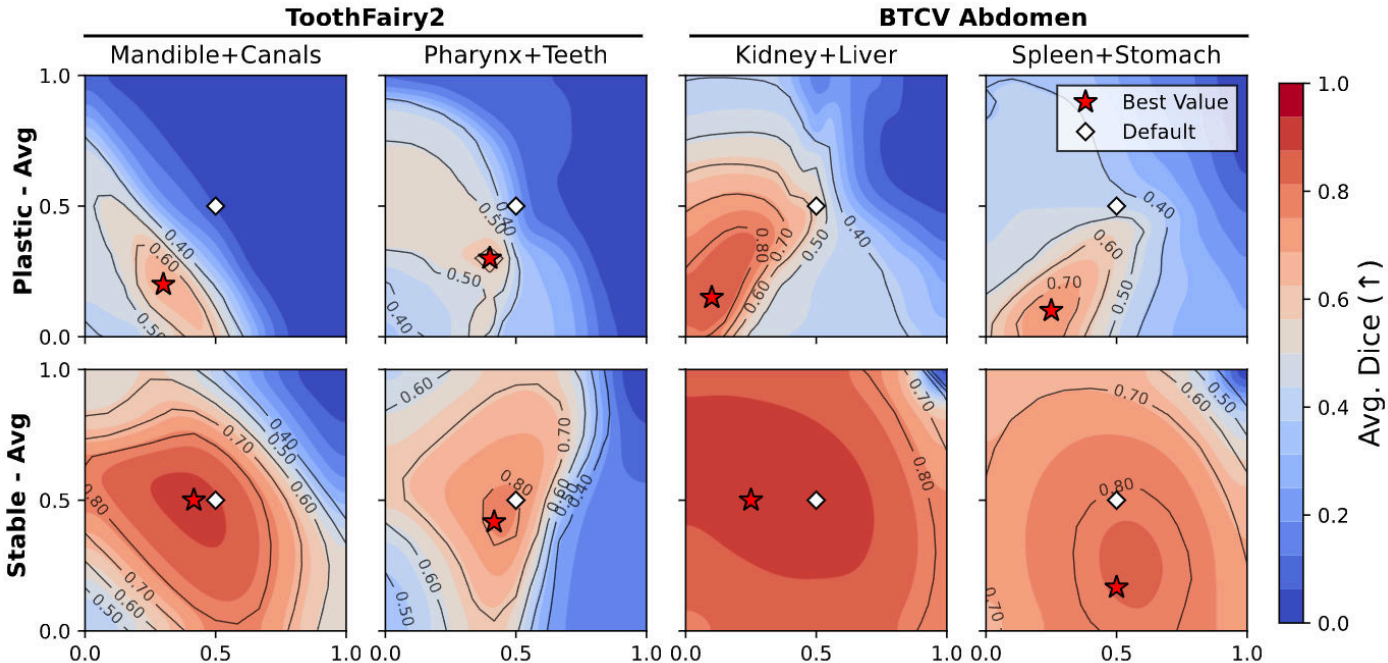

Stable Pre-training Allows Better Merging

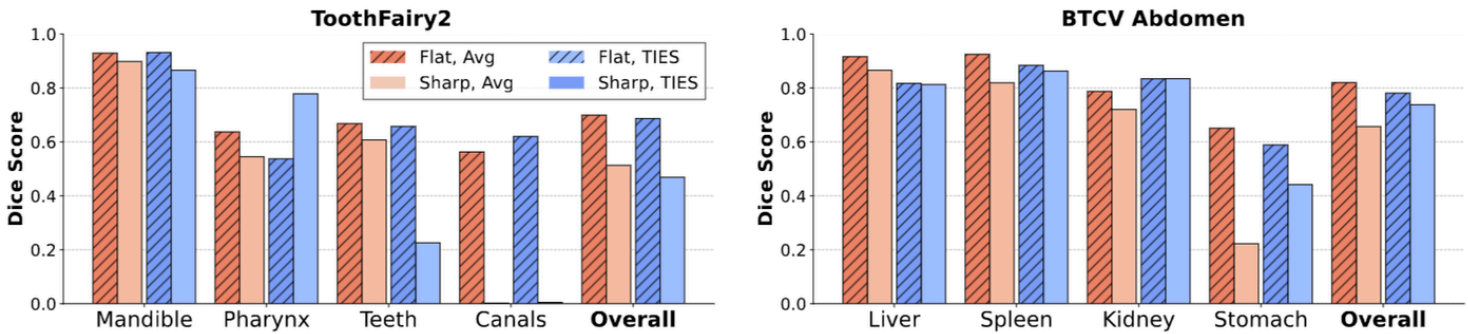

CBCT Report Generation

- Natural images vary wildly in what they represent.
- Medical reports often follow a relatively fixed template.
- The output space is relatively constrained.
- Few annotators and limited data.
- Hard to measure output correctness.

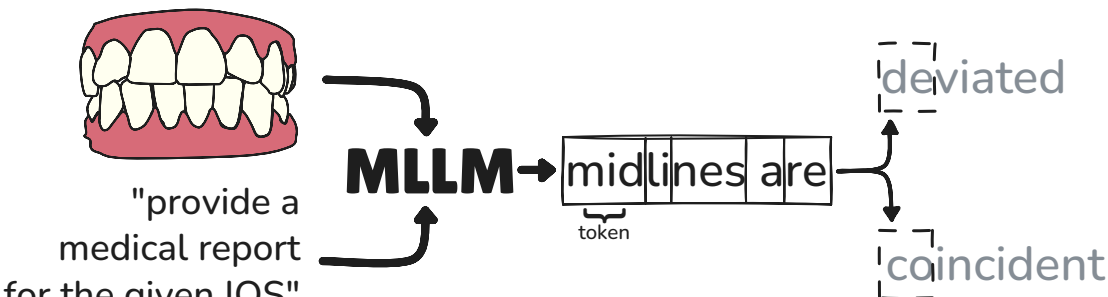

Ontology–driven Report Generation

```
:Finding a owl:Class .
:FindingType a owl:Class .
:Side a owl:Class .
:CanalTrajectory a owl:Class .
:CanalToothContact a owl:Class .

:findingType a owl:ObjectProperty ;
  rdfs:domain :Finding ;
  rdfs:range :FindingType .

:onSide a owl:ObjectProperty ;
  rdfs:domain :Finding ;
  rdfs:range :Side .

:hasCanalTrajectory a owl:ObjectProperty ;
  rdfs:domain :Finding ;
  rdfs:range :CanalTrajectory .

:hasCanalToothContact a owl:ObjectProperty ;
  rdfs:domain :Finding ;
  rdfs:range :CanalToothContact .

:hasToothId a owl:DatatypeProperty ;
  rdfs:domain :Finding ;
  rdfs:range xsd:string .

:MandibularCanalCourse a owl:NamedIndividual, :FindingType .
:ToothAbsence a owl:NamedIndividual, :FindingType .
:Left a owl:NamedIndividual, :Side .
:Right a owl:NamedIndividual, :Side .
:CentroLingual a owl:NamedIndividual, :CanalTrajectory .
:IntimateContact a owl:NamedIndividual, :CanalToothContact .
```

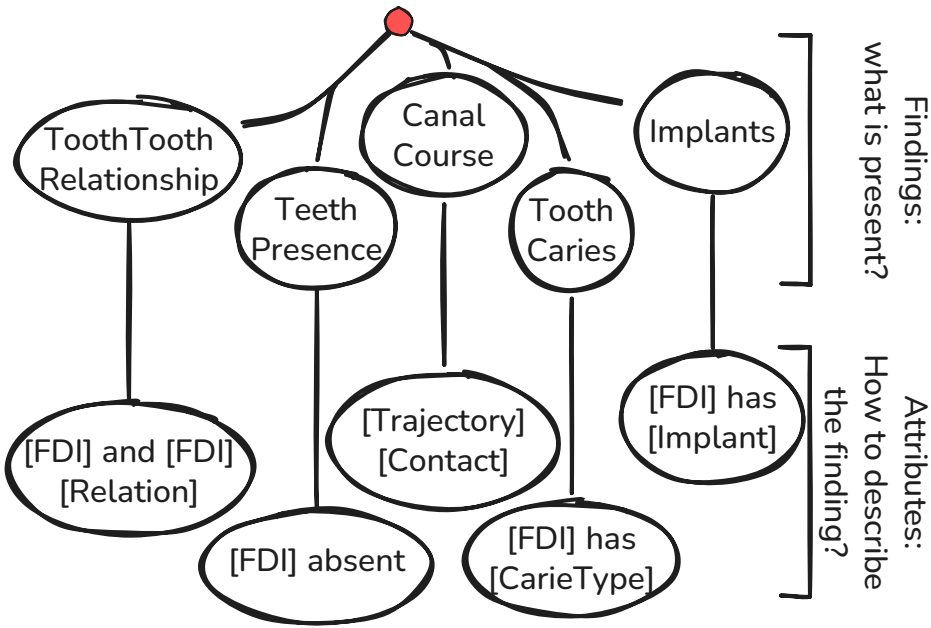

[11] Lumetti, L., Di Bartolomeo, M., Pellacani, A., Anesi, A., Grana, C., Bolelli, F. Ontology-Grounded Structured Prediction for Dental CBCT Reporting: Dataset and Baseline. MICCAI, 2026. (Under Review)

Ontology–driven Report Generation

Ontology definition (.ttl)

```
:Finding a owl:Class .
:FindingType a owl:Class .
:Side a owl:Class .
:CanalTrajectory a owl:Class .
:CanalToothContact a owl:Class .

:findingType a owl:ObjectProperty ;
  rdfs:domain :Finding ;
  rdfs:range :FindingType .

:onSide a owl:ObjectProperty ;
  rdfs:domain :Finding ;
  rdfs:range :Side .

:hasCanalTrajectory a owl:ObjectProperty ;
  rdfs:domain :Finding ;
  rdfs:range :CanalTrajectory .

:hasCanalToothContact a owl:ObjectProperty ;
  rdfs:domain :Finding ;
  rdfs:range :CanalToothContact .

:hasToothId a owl:DatatypeProperty ;
  rdfs:domain :Finding ;
  rdfs:range xsd:string .

:MandibularCanalCourse a owl:NamedIndividual, :FindingType .
:ToothAbsence a owl:NamedIndividual, :FindingType .
:Left a owl:NamedIndividual, :Side .
:Right a owl:NamedIndividual, :Side .
:CentroLingual a owl:NamedIndividual, :CanalTrajectory .
:IntimateContact a owl:NamedIndividual, :CanalToothContact .
```

Instance (.ttl)

```
:Finding_001 a :Finding ;
  rdfs:comment "Canale mandibolare destro a
decorso regolare" ;
  :findingType :MandibularCanalCourse ;
  :onSide :Right ;
  :hasCanalTrajectory :Regular .

:Finding_002 a :Finding ;
  rdfs:comment "Canale mandibolare sinistro a
decorso prevalentemente linguale" ;
  :findingType :MandibularCanalCourse ;
  :onSide :Left ;
  :hasCanalTrajectory :PrevalentlyLingual .

:Finding_004 a :Finding ;
  rdfs:comment "Assenza di elemento 46" ;
  :findingType :ToothAbsence ;
  :hasToothId "46" ;
  :onSide :Right .

:Finding_005 a :Finding ;
  rdfs:comment "Assenza di elemento 18" ;
  :findingType :ToothAbsence ;
  :hasToothId "18" ;
  :onSide :Right .
```

Ontology-driven Report Generation

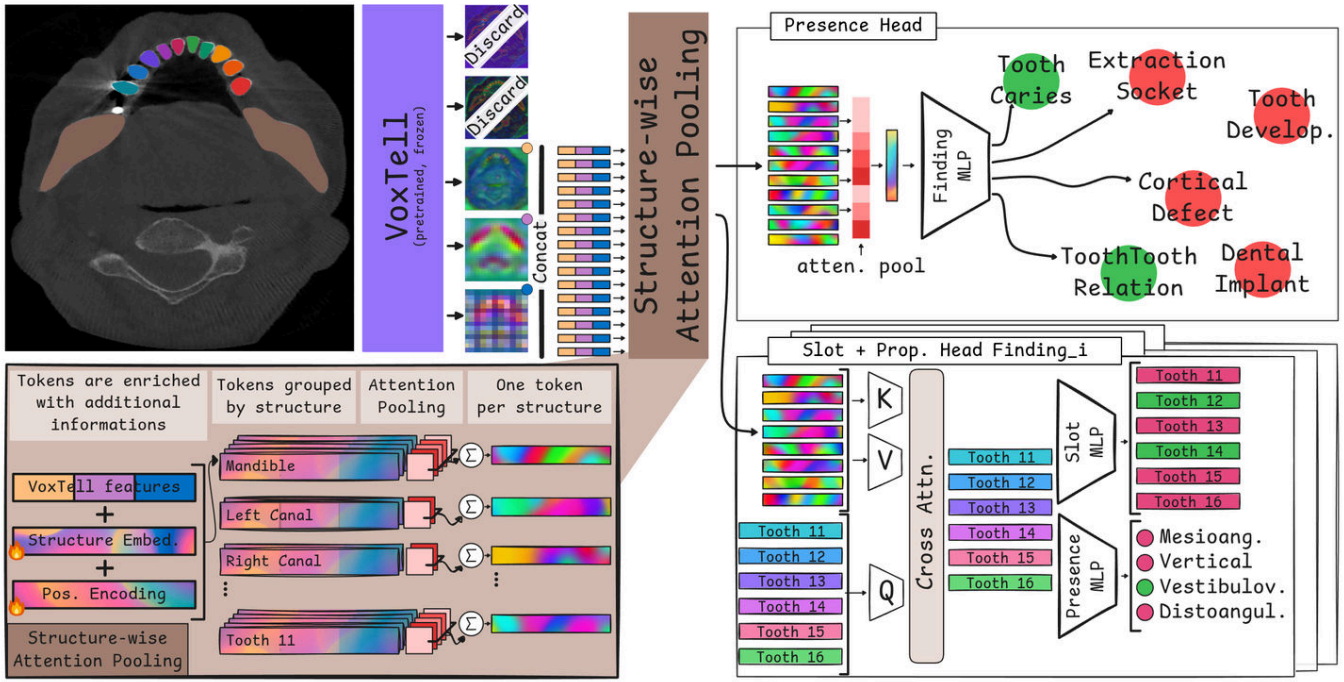

What's Next?

MICCAI 2026 Abu Dhabi UNITED ARAB EMIRATES

ODIN 2026: Oral and Dental Image aNalysis

Bite2Text Challenge: Orthodontic Report Generation from Intraoral Scans and Photographs

Bite2Text

ToothFairy4 Challenge: Maxillofacial and Surgical Report Generation from CBCT

https://odin-workshops.org/2026/

Internship Experiences

6 months in Leuven, Belgium

MILIARIS

6 months in Modena, Italy

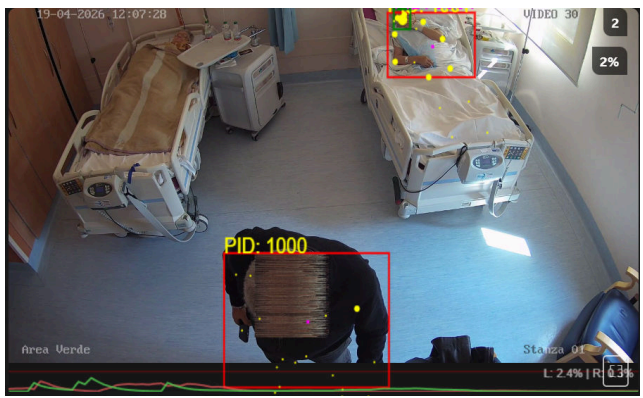

Additional Research

MedShapeNet – a large-scale dataset of 3D medical shapes for computer vision

Li, J., Zhou, Z., ..., Lumetti, L., ... & Reyes, M. (2025). Biomedical Engineering, 70(1), 71-90.

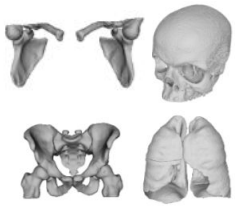

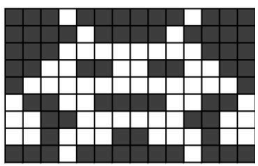

A State-of-the-Art Review with Code about Connected Components Labeling on GPUs

Bolelli, F., Allegretti, S., Lumetti, L., & Grana, C. (2024). IEEE Transactions on Parallel and Distributed Systems.

Buffer-MIL: Robust multi-instance learning with a buffer-based approach

Bontempo, G., Lumetti, L., Porrello, A., Bolelli, F., Calderara, S., & Ficarra, E. (2023, September). In International Conference on Image Analysis and Processing (pp. 1-12). Cham: Springer Nature Switzerland.

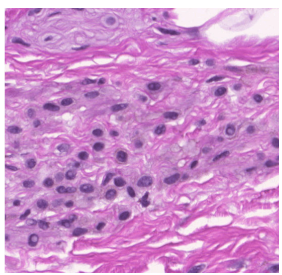

Additional Research

Enhancing Testicular Ultrasound Image Classification Through Synthetic Data and Pretraining Strategies

Morelli, N., Marchesini, K., Lumetti, L., Santi, D., Grana, C., & Bolelli, F. (2025, September). In International Conference on Image Analysis and Processing (pp. 26-39). Cham: Springer Nature Switzerland.

TesticulUS

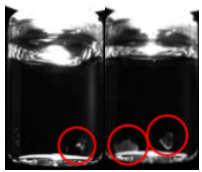

Identifying Impurities in Liquids of Pharmaceutical Vials

Rosati, G., Marchesini, K., Lumetti, L., Sartori, F., Balboni, B., Begarani, F., ... & Grana, C. (2024, December). In International Conference on Pattern Recognition (pp. 269-283). Cham: Springer Nature Switzerland.

Investigating the ABCDE Rule in Convolutional Neural Networks

Bolelli, F., Lumetti, L., Marchesini, K., Candeloro, E., & Grana, C. (2024, December). In International Conference on Pattern Recognition (pp. 94-111). Cham: Springer Nature Switzerland.

People that contributed to my PhD

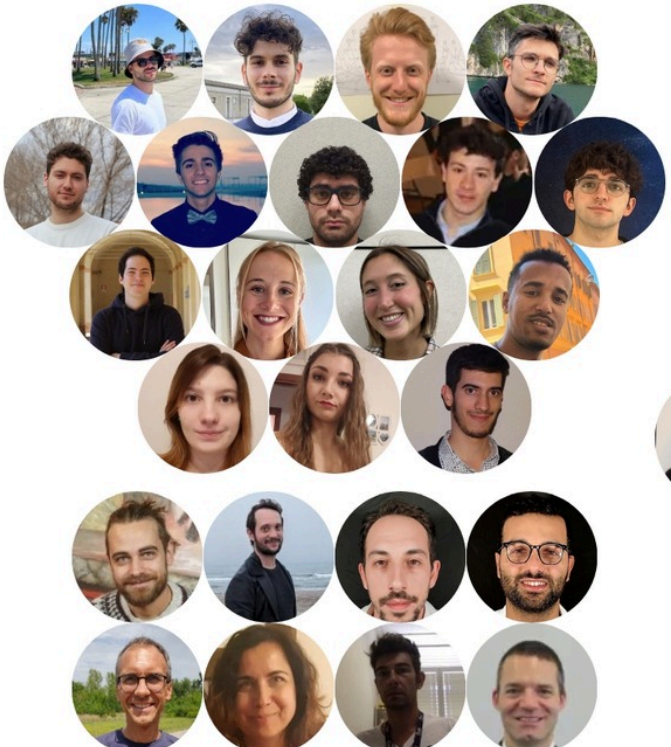

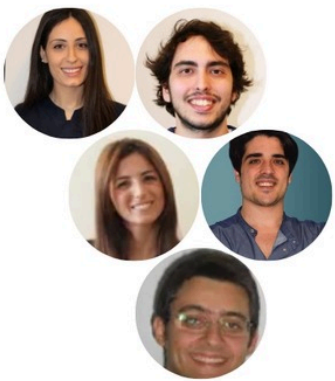

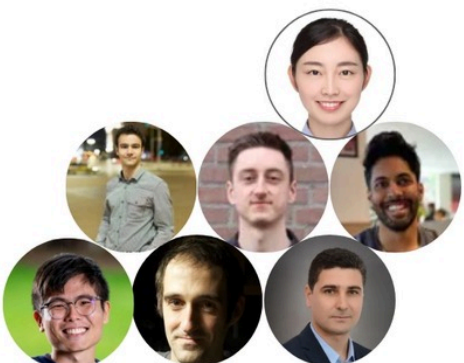

People that contributed to my PhD

Questions?

Publications – 1/3

[1] Lumetti, Luca, et al. "U-Net transplant: the role of pre-training for model merging in 3D medical segmentation." International Conference on Medical Image Computing and Computer-Assisted Intervention. Cham: Springer Nature Switzerland, 2025.

[2] Morelli, Nicola, et al. "Enhancing Testicular Ultrasound Image Classification Through Synthetic Data and Pretraining Strategies." International Conference on Image Analysis and Processing. Cham: Springer Nature Switzerland, 2025.

[3] Wang, Yaqi, et al., eds. Supervised and Semi-supervised Multi-structure Segmentation and Landmark Detection in Dental Data: MICCAI 2024 Challenges: ToothFairy 2024, 3DTeethLand 2024, and STS 2024, Held in Conjunction with MICCAI 2024, Marrakesh, Morocco, October 6, 2024, Proceedings. Springer Nature, 2025.

[4] Lumetti, Luca, et al. "Taming mambas for 3D medical image segmentation." IEEE Access (2025).

[5] Lumetti, Luca, et al. "Accurate 3D medical image segmentation with Mambas." 2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI). IEEE, 2025.

[6] Li, Jianning, et al. "Medshapenet–a large-scale dataset of 3d medical shapes for computer vision." Biomedical Engineering/Biomedizinische Technik 70.1 (2025): 71-90.

[7] Lumetti, Luca, et al. "Bits2Bites: Intra-oral Scans Occlusal Classification." Oral and Dental Image aNalysis Workshop. 2025.

Publications – 2/3

[8] Bolelli, Federico, et al. "Segmenting maxillofacial structures in cbct volumes." Proceedings of the Computer Vision and Pattern Recognition Conference. 2025.

[9] Segmenting the inferior alveolar canal in cbcts volumes: the toothfairy challenge

[10] Ma, Jun, et al. "Efficient medsams: Segment anything in medical images on laptop." arXiv preprint arXiv:2412.16085 (2024).

[11] Lumetti, Luca, et al. "Location matters: Harnessing spatial information to enhance the segmentation of the inferior alveolar canal in CBCTs." International Conference on Pattern Recognition. Cham: Springer Nature Switzerland, 2024.

[12] Bolelli, Federico, et al. "Investigating the ABCDE Rule in Convolutional Neural Networks." International Conference on Pattern Recognition. Cham: Springer Nature Switzerland, 2024.

[13] Rosati, Gabriele, et al. "Identifying Impurities in Liquids of Pharmaceutical Vials." International Conference on Pattern Recognition. Cham: Springer Nature Switzerland, 2024.

[14] Bolelli, Federico, et al. "A State-of-the-Art Review with Code about Connected Components Labeling on GPUs." IEEE Transactions on Parallel and Distributed Systems (2024).

[15] Lumetti, Luca, et al. "Enhancing patch-based learning for the segmentation of the mandibular canal." Ieee Access 12 (2024): 79014-79024.

Publications – 3/3

[16] Bontempo, Gianpaolo, et al. "Buffer-MIL: Robust multi-instance learning with a buffer-based approach." International Conference on Image Analysis and Processing. Cham: Springer Nature Switzerland, 2023.

[17] Lumetti, Luca, et al. "Annotating the inferior alveolar canal: The ultimate tool." International Conference on Image Analysis and Processing. Cham: Springer Nature Switzerland, 2023.

[18] Di Bartolomeo, Mattia, et al. "Inferior alveolar canal automatic detection with deep learning CNNs on CBCTs: development of a novel model and release of open-source dataset and algorithm." Applied Sciences 13.5 (2023): 3271.

Papers currently under review

[1] Lumetti, Luca, et al. "Do Multimodal LLMs Understand Intraoral Dental Data? Dataset, Platform, and Baselines." ECCV 2026. (Under Review)

[2] Talks, Joshua, et al. "Unsupervised and Source-Free Ranking of Biomedical Segmentation Models." ECCV 2026. (Under Review)

[3] Lumetti, Luca, et al. "Ontology-Grounded Structured Prediction for Dental CBCT Reporting: Dataset and Baseline" International Conference on Medical Image Computing and Computer-Assisted Intervention. Cham: Springer Nature Switzerland, 2026.

[4] Bolelli, Federico, et al. "The Paper has a GitHub, the GitHub has a README, the README has Nothing: Reproducibility Signals for Review Support" International Conference on Medical Image Computing and Computer-Assisted Intervention. Cham: Springer Nature Switzerland, 2026.

[5] Lumetti, Luca, et al. "Maxillo-77: Scaling CBCT Maxillofacial Segmentation to 77 Classes with M77-Mamba2" International Conference on Medical Image Computing and Computer-Assisted Intervention. Cham: Springer Nature Switzerland, 2026.

[6] Lumetti, Luca, et al. "Automated Intra-Oral Scan Analysis for Accurate Orthodontic Bracket Placement" International Conference on Medical Image Computing and Computer-Assisted Intervention. Cham: Springer Nature Switzerland, 2026.

[7] Lumetti, Luca et al. "Multi-Structure Segmentation in CBCT Volumes: the ToothFairy2 Challenge.", Medical Image Analysis (Under Review).