

Relational reasoning network for anatomical landmarking

Neslisah Torosdagli,^a Syed Anwar,^{a,b,c} Payal Verma,^d Denise K. Liberton,^d
Janice S. Lee,^d Wade W. Han,^{e,f} and Ulas Bagci^{a,f,g}

^aUniversity of Central Florida, Orlando, Florida, United States

^bChildren's National Hospital, Sheikh Zayed Institute, Washington, District of Columbia, United States

^cGeorge Washington University, Washington, District of Columbia, United States

^dNational Institute of Dental and Craniofacial Research (NIDCR),
National Institutes of Health (NIH), Craniofacial Anomalies and Regeneration Section,
Bethesda, Maryland, United States

^eBoston Children's Hospital, Harvard Medical School,
Department of Otolaryngology - Head and Neck Surgery, Boston, Maryland, United States

^fTher-AI, LLC, Kissimmee, Florida, United States

^gNorthwestern University, Departments of Radiology, BME, and ECE,
Machine and Hybrid Intelligence Lab, Chicago, Illinois, United States

Abstract

Purpose: We perform anatomical landmarking for craniomaxillofacial (CMF) bones without explicitly segmenting them. Toward this, we propose a simple, yet efficient, deep network architecture, called relational reasoning network (RRN), to accurately learn the local and the global relations among the landmarks in CMF bones; specifically, mandible, maxilla, and nasal bones.

Approach: The proposed RRN works in an end-to-end manner, utilizing learned relations of the landmarks based on dense-block units. For a given few landmarks as input, RRN treats the landmarking process similar to a data imputation problem where predicted landmarks are considered missing.

Results: We applied RRN to cone-beam computed tomography scans obtained from 250 patients. With a fourfold cross-validation technique, we obtained an average root mean squared error of <2 mm per landmark. Our proposed RRN has revealed unique relationships among the landmarks that help us in inferring informativeness of the landmark points. The proposed system identifies the missing landmark locations accurately even when severe pathology or deformations are present in the bones.

Conclusions: Accurately identifying anatomical landmarks is a crucial step in deformation analysis and surgical planning for CMF surgeries. Achieving this goal without the need for explicit bone segmentation addresses a major limitation of segmentation-based approaches, where segmentation failure (as often is the case in bones with severe pathology or deformation) could easily lead to incorrect landmarking. To the best of our knowledge, this is the first-of-its-kind algorithm finding anatomical relations of the objects using deep learning.

© The Authors. Published by SPIE under a Creative Commons Attribution 4.0 International License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.JMI.10.2.024002](https://doi.org/10.1117/1.JMI.10.2.024002)]

Keywords: anatomical landmarking; craniomaxillofacial bones; deep relational learning; relational reasoning; surgical modeling.

Paper 22220GR received Aug. 24, 2022; accepted for publication Feb. 13, 2023; published online Mar. 6, 2023.

*Address all correspondence to Ulas Bagci, ulasbagci@gmail.com

1 Introduction

In the United States alone, >17 million patients suffer from developmental deformities of the jaw, face, and skull region due to trauma, deformities from tumor ablation, or congenital birth defects.¹ The number of patients who require orthodontic treatment is far beyond this number. An accurate anatomical landmarking on radiological scans (mostly it is volumetric computed tomography-CT-scans) is a crucial step in the deformation analysis and surgical planning of the craniomaxillofacial (CMF) bones. This, if done correctly and efficiently, would afford precise image-based surgical planning for patients. This is even more significant since such deformities are known to vary from patient to patient and hence need careful delineation.

As mentioned briefly, landmarking can be used for a variety of clinical applications including dental implant planning, orthodontic treatment planning, and assessment of temporomandibular joint disorders. In dental implant planning, for instance, accurate landmarking is more important than segmentation as it allows clinicians to determine the appropriate location for the implant based on the location of nearby anatomical structures, such as the maxillary and mandibular sinuses and the mental foramen. Inaccurate landmarking can lead to incorrect placement of the implant. In orthodontic treatment planning, accurate landmarking can help the clinician to assess the overall shape and size of the teeth and jaw, as well as the location and orientation of specific teeth. This can be useful for developing a treatment plan that takes into account the patient's specific anatomy. There is a significant need for developing automated landmarking procedure because manual landmarking in volumetric CT scans is a tedious process and prone to inter-operator variability. There are considerable efforts toward developing a fully-automated and accurate software for anatomical landmarking based on bone segmentation from CT scans.²⁻⁴ Despite this clinical need, very little progress has been made especially for bones with a high level of congenital and developmental deformations (~5% of the CMF deformities).

Deep learning-based approaches have become the standard choice for pixel-wise medical-image segmentation applications due to their high efficacy.^{2,5,6} However, it is difficult to generalize segmentation especially when there is a high degree of deformation or pathology,⁷ which is the case for treating CMF conditions. Figure 1 demonstrates two examples of challenging mandible cases where the patients have surgical intervention (left) and high variability in the bone (right), causing segmentation algorithms to fail (leakage or under-segmentation). Current state-of-the-art landmarking algorithms are mostly dependent on bone segmentation results,

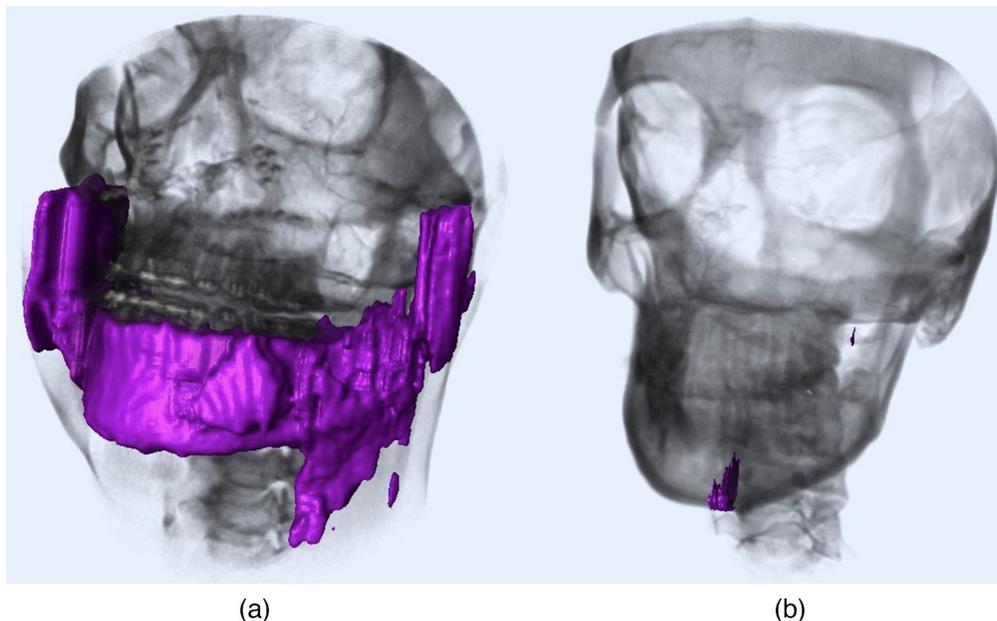


Fig. 1 CT segmentation results rendered in fuchsia which are scored as “unacceptable segmentation” at Ref. 5. (a) Patient with surgical intervention; (b) patient with high variability in the bone.

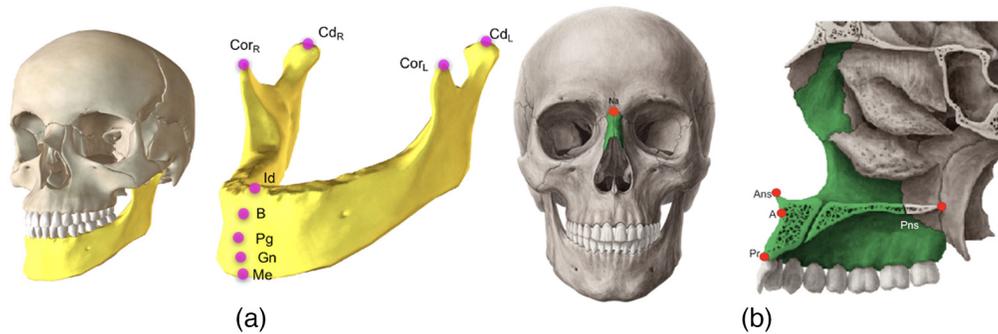


Fig. 2 Mandible and maxilla/nasal bone anatomies. (a) Mandibular landmarks: menton (Me), condylar left (Cd_L), condylar right (Cd_R), coronoid left (Cor_L), coronoid right (Cor_R), infradentale (Id), B point (B), pogonion (Pg), and gnathion (Gn); (b) maxillary landmarks: ANS, PNS, A-point (A), and prosthion (Pr), and nasal bones landmark: nasion (Na).

since locating landmarks could become easier once their parent anatomy (the bones they belong to) is precisely known.⁷ Figure 2 demonstrates mandible and maxilla/nasal bone anatomies along with the landmarks associated with those bones. If the underlying segmentation is poor, it is highly likely to have a large landmark localization error, directly affecting the quantification process (which could include severity measurement, surgical modeling, and treatment planning).

We hypothesize that if an explicit segmentation can be avoided for extremely challenging cases, landmark localization errors can be minimized. This will also lead to a widespread use of landmarking procedure in surgical planning and precision medicine. Since CMF bones are found in the same anatomical space even when there is deformity or pathology, the overall global relationship of the anatomical landmarks should still be preserved despite severe localized changes. Based on this rationale, we claim that utilizing local and global relations of the landmarks can help automatic landmarking without the extreme need for segmentation.

1.1 Background and Related Work

1.1.1 Landmarking

Anatomical landmark localization approaches can broadly be categorized into three main groups: registration-based (atlas-based),⁸ knowledge-based,^{9,10} and learning-based.^{7,11} Integration of shape and appearance increases the accuracy of the registration-based approaches. However, image registration is still an illposed problem, and when there are variations such as age (pediatrics versus adults), missing teeth (very common in certain age groups), missing bone or bone parts, severe pathology (congenital or trauma), and imaging artifacts, the performance can be quite poor.^{3,12,13} The same concerns apply to segmentation-based approaches too.

Gupta et al.¹⁰ developed a knowledge-based algorithm to identify 20 anatomical landmarks on cone-beam CT (CBCT) scans. Despite their promising results, a seed must be selected by using 3D template registration on the inferior–anterior region where fractures are most commonly found. An error in the seed localization may easily lead to a suboptimal outcome in such approaches. Zhang et al.¹⁴ developed a regression forest-based landmark detector to localize CMF landmarks on the CBCT scans. To address the spatial coherence of landmarks, image segmentation was used as a helper. The authors obtained a mean digitization error <2 mm for 15 CMF landmarks. The following year, to reduce the mean digitization error further, Zhang et al.² proposed a deep learning-based joint CMF bone segmentation and landmarking strategy. A context guided multitask fully convolutional neural (FCN) network was employed along with 3D displacement maps to perceive the spatial locations of the landmarks. A segmentation accuracy of $93.27 \pm 0.97\%$ and a mean digitization error of <1.5 mm for identifying 15 CMF landmarks was achieved. Further, a joint segmentation and landmark digitization framework was proposed, where two stages of FCN were cascaded to perform bone segmentation and landmark localization.⁷ The major disadvantage of this (one of the state-of-the-arts) method was the

memory constraint introduced by the redundant information in the 3D displacement maps such that only a limited number of the landmarks can be learned using this approach. Since the proposed strategy is based on joint segmentation and landmarking, it naturally shares other disadvantages of the segmentation-based methods: frequent failures for very challenging cases. The landmark localization problem was solved using an object detection method, where region proposals were used to identify landmark locations and a coarse-to-fine method was used to achieve landmark localization.¹⁵ It must be noted that the method does not use the relationships between the anatomical landmarks in the CMF bones.

Recently, we integrated the manifold information (geodesic) in a deep learning architecture to improve robustness of the segmentation-based strategies for landmarking,⁵ and obtained promising results, significantly better than the state-of-the-art methods. We also noticed that there is still room to improve landmarking process, especially when pathology or bone deformation is severe. To fill this research gap, in this study, we take a radically different approach by learning landmark relationships without segmenting bones. We hypothesize that the inherent relation of the landmarks in the CMF region can be learned by a relational reasoning algorithm based on deep learning. Although our proposed algorithm stems from this unique need of anatomical landmarking, the core idea of this work is inspired by the recent studies in artificial intelligence (AI), particularly in robotics and physical interactions of human/robots with their environments, as described in the following with further details.

1.1.2 Relational reasoning

The ability to learn relationship and infer reasons between entities and their properties is a central component of the AI field, however, it has been proven to be very difficult to learn through neural networks until recently.^{16,17} In 2009, Scarselli et al.¹⁸ introduced a graph neural network (GNN) by extending the neural network models to process graph data which encoded relationship information of the objects under investigation. Li et al.¹⁹ proposed a machine learning model based on gated recurrent units (GRUs) to learn the distributed vector representations from heap graphs. Despite this increase in use and promising nature of the GNN architectures,²⁰ there is a limited understanding of their representational properties, which is often a necessity in medical AI applications for their adoption in clinics.

Recently, DeepMind team(s) published four important studies on relational reasoning and explored how objects in complex systems can interact with each other.^{16,21–23} Battaglia et al.²¹ introduced interaction networks to reason about the objects and the relations in the complex environments. The authors proposed a simple, yet accurate system to reason about n -body problems, rigid-body collision, and nonrigid dynamics. The proposed system can predict the dynamics in the next step with an order of magnitude lower error and higher accuracy. Raposo et al.¹⁶ introduced a relational network (RN) to learn the object relations from a scene description, hypothesizing that a typical scene contains salient objects which are typically related to each other by their underlying causes and semantics. Following this study, Santoro et al.²² presented another relational reasoning architecture for tasks such as visual question-answering, text-based question-answering, and dynamic physical systems where the proposed model obtained most answers correctly. Lastly, Battaglia et al.²³ studied the relational inductive biases to learn the relations of the entities and presented the graph networks. These four studies show promising approaches to understanding the challenge of relational reasoning. To the best of our knowledge, such advanced reasoning algorithms have neither been developed for nor applied to the medical imaging applications yet. It must be noted that medical AI applications require fundamentally different reasoning paradigms than conventional computer vision and robotics fields have²⁴ (e.g., salient objects definitions). To address this gap, in this study we focus on the anatomy–anatomy and anatomy–pathology relationships in an implicit manner.

1.2 Summary of Our Contributions

- The proposed method is the first in the literature to successfully apply spatial reasoning of the anatomical landmarks for accurate and robust landmarking using deep learning.

- Many anatomical landmarking methods, including our previous works,^{5,10,25} use bone segmentation as a guidance for finding the location of the landmarks on the surface of a bone. The major limitation imposed by such an approach stems from the fact that it is not always possible to have an accurate segmentation. Our proposed RRN system addresses this problem by enabling accurate prediction of anatomical landmarks without employing explicit object segmentation.
- Since efficiency is a significant barrier for many medical AI applications, we explore new deep learning architecture designs for a better efficacy in the system performance. For this purpose, we utilize variational dropout²⁶ and targeted dropout²⁷ in our implementation for faster and more robust convergence of the landmarking procedure (~5 times faster than baselines).
- Our data set includes highly variable bone deformities along with other challenges of the CBCT scans with a larger number of scans (as compared to baselines). Hence, the proposed algorithm is considered robust and identifies anatomical landmarks accurately under varying conditions (Table 2). In our experiments, we find landmarks pertaining to mandible, maxilla, and nasal bones (Fig. 2).

The rest of this paper is organized as follows: we introduce our novel methodology and its details in Sec. 2. In Sec. 3, we present experiments and results and then we conclude the paper by discussing strengths and limitations of our study in Sec. 4.

2 Methods

2.1 Overview and Preliminaries

The most frequently deformed or injured CMF bone is the lower jawbone, or mandible, which is the only mobile CMF bone.²⁸ In our previous study,⁵ we developed a framework to segment mandible from CBCT scans and identify the mandibular landmarks in a fully-automated way. Herein, we focus on anatomical landmarking without the need for explicit segmentation, and extend the learned landmarks into other bones (maxilla and nasal). Overall, we seek answers to the following important questions:

- **Q1:** Can we automatically identify all anatomical landmarks of a bone if some of the landmarks are missing? If so, what is the least effort for performing this procedure? How many landmarks are necessary, and which landmarks are more informative to perform this whole procedure?
- **Q2:** Can we identify anatomical landmarks of nasal and maxilla bones if we only know locations of a few landmarks in the mandible and the rest is missing? Do relations of landmarks hold true even when they belong to different anatomical structures (manifold)?

In this study, we explore inherent relations among anatomical landmarks at the local and global levels in order to explore availability of structured data samples helping anatomical landmark localization. Inferred from the morphological integration of the CMF bones, we claim that landmarks of the same bone should carry common properties of the bone so that one landmark should give clues about the positions of the other landmarks with respect to a common reference. This reference is often chosen as segmentation of the bone to enhance information flow, but in our study, we leverage this reference point from the whole segmented bone into a reference landmark point. Throughout the text, we use the following definitions:

Definition 1: A landmark is an anatomically distinct point, helping clinicians to make reliable measurements related to a condition, diagnosis, modeling a surgical procedure, or creating a treatment plan.

Definition 2: A relation is defined as a geometric property between landmarks. Relations might include the following geometric features: size, distance, shape, and other implicit structural information. In this study, we focus on pairwise relations between landmarks as a starting point.

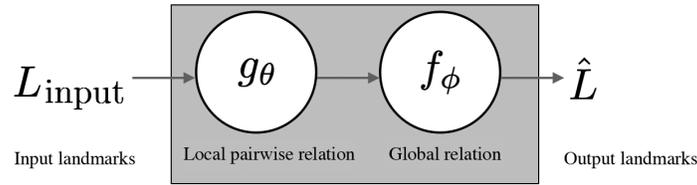


Fig. 3 Overview of the proposed RRN architecture: for a few given input landmarks, RRN utilizes both pairwise and combination of all pairwise relations to predict the remaining landmarks.

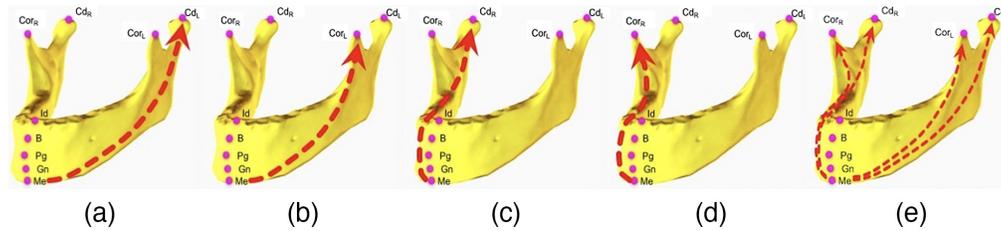


Fig. 4 For the input domain $L_{input} = \{Me, Cd_L, Cor_L, Cd_R, Cor_R\}$, (a)–(d) pairwise relations of landmark menton (Me): (a) menton-condylar left, (b) menton-coronoid left, (c) menton-condylar right, (d) menton-coronoid right, and (e) combined relations of menton.

Definition 3: A reason is defined as an inference about relationships of the landmarks. For instance, compared to closely localized landmarks (if given as input), a few of sparsely localized landmarks can help predicting landmarks better. The reason is that sparsely localized input landmark configuration captures the anatomy of a region of interest and infers better global relationships of the landmarks.

Once relationship among landmarks is learned effectively, we can use this relationship to identify the missing landmarks on the same or different CMF bones without the need for a precise segmentation. Toward this goal, we propose to learn a relationship between the anatomical landmarks in two stages (illustrated in Fig. 3) based on relational units (RUs). The first stage is shown as the function g , which learns the pairwise (local) relations. The second stage is shown as a function f , which combines pairwise relations (g) of the first stage into a global relation based on RUs.

Figure 4 shows an example of pairwise relations for different pairs of mandible landmarks. There are five sparsely localized landmarks. The basis/reference is chosen as menton (Me), in this example, hence, four pairwise relations are illustrated from Figs. 4(a)–4(d). Figure 4(e) illustrates combined relations Figs. 4(a)–4(d) of the landmark menton (reference) with respect to other four landmarks on the mandible.

2.2 Relational Reasoning Architecture

Anatomical landmarking has been an active research topic for several years in the medical imaging field. However, how to build a reliable/universal relationship between landmarks for a given clinical problem is an open problem. While anatomical similarities at the local and global levels could serve toward viable solutions, thus far, features that can represent anatomical landmarks from the medical images have not achieved the desired efficacy and interpretation.^{2,29–31}

We propose a new framework called relational reasoning network (RRN) to learn local and global relations of anatomical landmarks (o_i) through its units called RU (relationship unit). The proposed RRN architecture and its RU subarchitectures are summarized in Fig. 5. The relation between two landmarks is encoded via major spatial properties of the landmarks. We explore two architectures as RU: first one is a simple multilayer perceptron (MLP) (Fig. 5-bottom left) (similar to Ref. 16), the other one is more advanced architecture composed of dense-blocks (DBs) (Fig. 5 bottom middle). Both architectures are relatively simple compared to very dense complex deep-learning architectures. The rationale is simple when there is a less data (i.e., pairwise relation), it is natural to choose fully connected layers to keep the full spectrum of the data at hand.

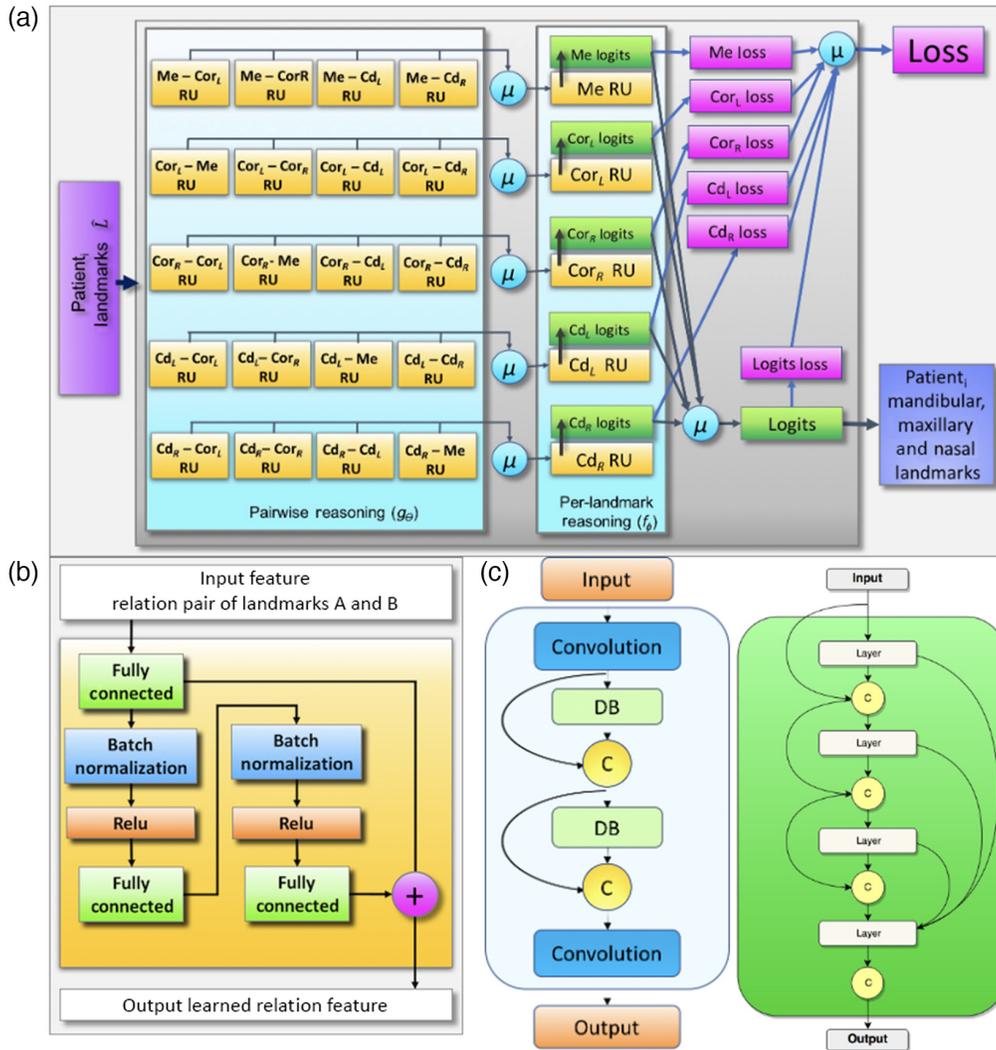


Fig. 5 (a) RRN architecture for five-input landmarks $RRN(L_{input})$: $L_{input} = \{Me, Cor_L, Cor_R, Cd_L, Cd_R\}$, $\hat{L} = \{Gn, Pg, B, Id, Ans, A, Pr, Pns, Na\}$ and μ is the average operator. (b) Content of the pairwise reasoning block, RU of the RRN. (c) RU composed of two DBs, convolution and concatenation (C) units. (d) DB architecture composed of four layers and concatenation layers.

Similarly, when more pairwise data are available for exploration of the more complex relation, it is natural to move into convolutional operation from fully connected layers to keep dominant information while reducing the redundant information and providing a computational feasibility. Our objective is to locate all anatomical landmarks by inputting a few landmarks to RRN, which provides reasoning inferred from the learned relationships of landmarks and locate all other landmarks automatically.

In the pairwise learning/reasoning stage (stage 1), five-landmarks-based system is assumed as an example network (other configurations are possible too, see experiments and results section). Sparsely spaced landmarks [Fig. 4(e)] and their pairwise relationships are learned in this stage (g_θ). These pairwise relationship(s) are later combined in a separate DB setting in (f_ϕ). It should be noted that this combination is employed through a joint loss function and an RU to infer an average relation information. In other words, for each individual landmark, the combined relationship vector is assigned a secondary learning function through a single RU.

The RU is the core component of the RRN architecture. Each RU is designed in an end-to-end fashion; hence, they are differentiable. For n landmarks in the input domain, the proposed RRN architecture learns $n \times (n - 1)$ pairwise and n combined relations (global) with a total of n^2 RUs. Therefore, depending on the number of input domain landmarks, RRN can be either

shallow or dense. Let L_{input} and \hat{L} indicate vectors of input and output anatomical landmarks, respectively. Then, two stages of the RRN of the input domain landmarks L_{input} can be defined as

$$G_{\theta_i} = \frac{1}{(n-1)} \sum_{j=1, j \neq i}^n (g_{\theta}(o_i, o_j)), \quad \text{RRN}(L_{\text{input}}; \theta, \phi) = \frac{1}{n} \sum_{i=1}^n f_{\phi_i}(G_{\theta_i}), \quad (1)$$

where G_{θ_i} is the mean pairwise relation vector of the landmark o_i to every other landmark $o_{j(j \neq i)} \in L_{\text{input}}$. The functions f_{ϕ} and g_{θ} are the functions with the free parameters ϕ and θ , and f_{ϕ} indicates a global relation (in other words, combined pairwise relations) of landmarks.

2.3 Pairwise Relation (g_{θ})

For a given a few input landmarks (L_{input}), our objective is to predict the 3D spatial locations of the target domain landmarks ($\in \hat{L}$) by using the 3D spatial locations of the input domain landmarks ($\in L_{\text{input}}$). With respect to relative locations of the input domain landmarks, we reason about the locations of the target domain landmarks. The RU function $g_{\theta}(o_i, o_j)$ represents the relation of two input domain landmarks o_i and o_j where $i \neq j$ [Figs. 4(a)–4(d)]. The output of $g_{\theta}(o_i, o_j)$ describes relative spatial context of two landmarks, defined for each pair of input domain landmarks (pairwise relation at Fig. 5). According to each input domain landmark o_i , the structure of the manifold is captured through mean of all pairwise relations [represented as G_{θ_i} at Eq. (1)].

2.4 Global Relation (f_{ϕ})

The mean pairwise relation G_{θ_i} is calculated with respect to each input domain landmark o_i , and it is given as input to the second stage where global (combined) relation f_{ϕ_i} is learned. f_{ϕ_i} is an RU function and the output of f_{ϕ_i} is the predicted 3D coordinates of the target domain landmarks ($\in \hat{L}$). In other words, each input domain landmark o_i learns and predicts the target domain landmarks by the RU function f_{ϕ_i} . The terminal prediction of the target domain landmarks is the average of individual predictions of each input domain landmark, represented by $\text{RRN}(L_{\text{input}}; \theta, \phi)$ at Eq. (1). There are totally n^2 RUs in the architecture. The number of trainable parameters used for each experimental configuration are directly proportional with n^2 (Fig. 6).

Experiment #	Configuration	L_{input}	\hat{L}	#RUs
1	5-landmarks	$Me, Cd_L, Cd_R, Cor_L, Cor_R$	$Gn, Pg, B, Id, Ans, A, Pr, Pns, Na$	25
2	3-Landmarks regular	Me, Cd_L, Cd_R	$Gn, Pg, B, Id, Cor_L, Cor_R, Ans, A, Pr, Pns, Na$	9
3	3-Landmarks cross	Me, Cd_R, Cor_L	$Gn, Pg, B, Id, Cd_L, Cor_R, Ans, A, Pr, Pns, Na$	9
4	6-landmarks	$Me, Cd_L, Cd_R, Cor_L, Cor_R, Na$	$Gn, Pg, B, Id, Ans, A, Pr, Pns$	36
5	9-landmarks	$Me, Cd_L, Cd_R, Cor_L, Cor_R, Gn, Pg, B, Id$	Ans, A, Pr, Pns, Na	81

Fig. 6 Five experimental landmark configurations for experimental explorations. L_{input} : input landmarks and \hat{L} : output landmarks, and #RUs indicate the number of relational units. Landmarks are visualized using reference standard bones for illustrative purposes; in our implementation there is no explicit segmentation exist.

Since all pairwise relations are leveraged under G_{θ_i} and f_{ϕ} with averaging operation, we can conclude that RRN is invariant to the order of input landmarks (i.e., permutation-invariant).

2.5 Loss Function

The natural choice for the loss function is the mean squared error (MSE) because it is a differentiable distance metric measuring how well landmarks are localized/matched, and it allows output of the proposed network to be real-valued functions of the input landmarks. For n input landmarks and m target landmarks, MSE simply penalizes large distances between the landmarks as follows:

$$\text{Loss}(W_{\Theta}, (\theta, \phi)) = \frac{1}{n * m} \sum_{i=1}^n \left(\sum_{k=1}^m \|(f_{\phi}(G_{\theta_i}))_k - o_k\|^2 \right), \quad (2)$$

where o_k is the target domain landmarks ($o_k \in \hat{L}$).

2.6 Variational Dropout

Dropout is an important regularizer employed to prevent overfitting at a cost of 2 to 3 times (on average) increase in training time.³² For efficiency reasons, speeding up dropout is critical and it can be achieved by a variational Bayesian inference on the model parameters.²⁶ Given a training input dataset $X = \{x_1, x_2, \dots, x_N\}$ and the corresponding output dataset $Y = \{y_1, y_2, \dots, y_N\}$, the goal in RRN is to learn the parameters ω such that $y = F_{\omega}(x)$. In the Bayesian approach, given the input and output datasets X, Y , we seek for the posterior distribution $p(\omega|X, Y)$, by which we can predict output y^* for a new input point x^* by solving the integral³³

$$p(y^*|x^*, X, Y) = \int p(y^*|X^*, \omega)p(\omega|X, Y)d\omega. \quad (3)$$

In practice, this computation involves intractable integrals.²⁶ To obtain the posterior distributions, a Gaussian prior distribution $N(0, I)$ is placed over the network weights³³ which leads to a much faster convergence.²⁶

2.7 Targeted Dropout

Alternatively, we also propose to use targeted dropout for better convergence.²⁷ Given a neural network parameterized by Θ , the goal is to find the optimal parameters $W_{\Theta}(\cdot)$ such that the loss $\text{Loss}(W_{\Theta})$ is minimized. For efficiency and generalization reasons, $|W_{\Theta}| \leq k$, only k weights of highest magnitude in the network are employed. In this regard, deterministic approach is to drop the lowest $|W_{\Theta}| - k$ weights. In targeted dropout, using a target rate γ and a drop out rate α , first a target set is generated with the lowest weights with the target rate γ . Next, weights are dropped out in a stochastic manner from the target set at a certain dropout rate α .

2.8 Landmark Features

Pairwise relations are learned through RU functions. Each RU accepts input features to be modeled as a pairwise relation. It is desirable to have such features characterizing landmarks and interactions with other landmarks. These input features can either be learned throughout a more complicated network design, or through feature engineering. In this study, for simplicity, we define a set of simple yet explainable geometric features. Since RUs model relations between two landmarks (o_A and o_B), we use 3D coordinates of these landmarks (both in pixel and spherical space), their relative positions with respect to a well-defined landmark point (reference), and approximate size of the mandible. The mandible size is estimated as the distance between the maximum and the minimum coordinates of the input domain mandibular landmarks (Table 1). Finally, a 19-dimensional feature vector is considered to be an input to local relationship function

Table 1 Input landmarks have the following feature(s) to be used only in stage I. 19D feature vector includes only structural information.

Pairwise feature (o_A, o_B)	
3D pixel-space position of the o_A	(A_x, A_y, A_z)
Spherical coordinate of the vector from landmark menton (o_1) to o_A	$(r_{me \rightarrow A}, \theta_{me \rightarrow A}, \phi_{me \rightarrow A})$
3D pixel-space position of the o_B	(B_x, B_y, B_z)
Spherical coordinate of the vector from landmark menton to o_B	$(r_{me \rightarrow B}, \theta_{me \rightarrow B}, \phi_{me \rightarrow B})$
3D pixel-space position of the landmark menton	(Me_x, Me_y, Me_z)
Spherical coordinate of the vector from o_A to o_B	$(r_{A \rightarrow B}, \theta_{A \rightarrow B}, \phi_{A \rightarrow B})$
Diagonal length of the bounding box capturing mandible roughly, computed as the distance between the minimum and the maximum spatial locations of the input domain mandibular landmarks (L_1) in the pixel space.	d_1

g. For a well-defined reference landmark, we used menton (Me) as the origin of the mandible region.

3 Experiments and Results

3.1 Data Description

Anonymized CBCT scans of 250 patients (142 female and 108 male, mean age = 23.6 years, standard deviation = 9.34 years) were included in our analysis through an IRB-approved protocol. The data set includes both pediatric and adult patients with craniofacial congenital birth defects, developmental growth anomalies, trauma to the CMF, and surgical interventions. CB MercurRay CBCT system (Hitachi Medical Corporation, Tokyo, Japan) was used to scan the data at 10 mA and 100 Kvp. The radiation dosage for each scan was around 300 mSv. To handle the computational cost, each patient's scan was resampled from $512 \times 512 \times 512$ to $256 \times 256 \times 512$. In-plane resolution of the scans was noted (in mm) either as $0.754 \times 0.754 \times 0.377$ or $0.584 \times 0.584 \times 0.292$. In addition, following image-based variations exist in the data set: aliasing artifacts due to braces, metal alloy surgical implants (screws and plates), dental fillings, and missing bones or teeth.⁵ Briefly, 3% of the whole data set was including CBCT scans with extreme deformation and artifacts, while 11% of the data set was including cases with large-scale tissue or bone deformations, artifacts, or missing bones. 16% of the data set was including minor tissue deformation and/or metal or other artifacts. The remaining 70% of the data was either no visible artifacts or minor problems in visual assessment. These statistics were obtained by participating two experts, blindly to each other, and qualitatively they were asked to evaluate the scans visually.

The data was annotated independently by three expert interpreters, one from the NIH team, and two from UCF team. Interobserver agreement values were computed as ~ 3 pixels. Experts used freely available 3D Slicer software for the annotations.⁵

3.2 Data Augmentation

Our data set includes fully annotated mandibular, maxillary, and nasal bones' landmarks. Due to insufficiency of 250 samples for a deep-learning algorithm to run, we applied data-augmentation approach. In our study, the common usage of random scaling or rotations for data-augmentation was not found to be useful for new landmark data generation because such transformations

would not generate new relations different from the original ones. Instead, we used random interpolation similar to active shape model's landmarks.³⁰ Briefly, we interpolated 2 (or 3) randomly selected scans with randomly computed weight. We merged the relation information at different scans to a new relation. We also added random noise to each landmark with a maximum in the range of ± 5 pixels, defined empirically based on the resolution of the images as well as the observed high deformity of the bones. We generated a dataset with ~ 100 K landmarks, which is empirically evaluated as a sufficiently large dataset.

3.3 Evaluation Methods

We used root-mean squared error (RMSE) in the anatomical space (in mm) to evaluate the goodness of the landmarking. Lower RMSE indicates successful landmarking process. For statistical significance comparisons of different methods and their variants, we used a P -value of 0.05 as a cut-off threshold to define significance and applied t -tests where applicable.

3.4 Input Landmark Configurations

In our experiments, there were three groups of landmarks (See Fig. 2) defined based on the bones they reside: Mandibular $L_1 = \{o_1, \dots, o_9\}$, Maxillary $L_2 = \{o_{10}, \dots, o_{13}\}$, and Nasal $L_3 = \{o_{14}\}$, where subscripts in o denote the specific landmark in that bone:

- $L_1 = \{\text{Me, Gn, Pg, B, Id, Cor}_L, \text{Cor}_R, \text{Cd}_L, \text{Cd}_R\}$,
- $L_2 = \{\text{Ans, A, Pr, Pns}\}$,
- $L_3 = \{\text{Na}\}$.

In each experiment, as detailed in Fig. 6, we designed a specific input set L_{input} where $L_{\text{input}} \subseteq L_1 \cup L_2$, $|L_{\text{input}}| = n$ and $1 < n \leq |L_1| + |L_2|$. The target domain landmarks for each experiment were $\hat{L} = (L_1 \cup L_2 \cup L_3) \setminus L_{\text{input}}$ and $|\hat{L}| = m$ such that $n + m = 14$. With carefully designed input domain configurations L_{input} , and pairwise relationships of the landmarks in the input set, we seek the answers to the following questions previously defined as Q1 and Q2 in Sec. 2:

- What configuration of the input landmarks can capture the manifold of bones so that other landmarks can be localized successfully?
- What is the minimum number and configuration of the input landmarks for successful identification of other landmarks?

Overall, we designed five different input landmark configurations called three-landmarks regular, three-landmarks cross, five-landmarks, six-landmarks, and nine-landmarks (Fig. 6). Each configuration is explained in Sec. 3.6.

3.5 Training

The MLP RU was composed of three fully connected layers, two batch normalizations, and two ReLUs (Fig. 5). The DB RU architecture contained 1 DB, which was composed of four layers with a growth rate of 4. We used a batch size of 64 for all experiments. For the five-landmarks configuration, there were 6,596,745 and 11,068,655 trainable parameters for the MLP and the DB architectures, respectively. We trained the network for 100 epochs on 1 NVIDIA Titan-XP GPU with 12 GB memory using the MLP architecture with the regular dropout compared to 20 epochs with the variational and targeted dropout implementations. For the DB architecture, it converged in around 20 epochs independent of the dropout implementation employed.

3.6 Experiments and Results

We ran a set of experiments to evaluate the performance of the proposed system using a fourfold cross-validation. We summarized the experimental configurations in Fig. 6, error rates in Table 2,

Table 2 Landmark localization errors (mm). The symbol “—” means not applicable (N/A). Bold values represent the best performance obtained from the ablation and benchmarking experiments.

Method	Mandibular landmarks													
	Cor _R	Cor _L	Cd _L	Gn	Pg	B	Id	Cor _R	Cor _L	Cd _L	Gn	Pg	B	Id
Three-landmarks regular (dense)	3.32 ± 0.30	3.03 ± 0.31	—	0.01 ± 0.03	0.09 ± 0.11	0.60 ± 0.15	0.56 ± 0.19	—	—	—	—	—	—	—
Three-landmarks cross (dense)	1.88 ± 0.24	—	1.70 ± 0.23	0.007 ± 0.03	0.10 ± 0.11	0.77 ± 0.18	0.58 ± 0.20	—	—	—	—	—	—	—
Five-landmarks var. dropout (MLP)	—	—	—	0.05 ± 0.05	0.22 ± 0.13	0.91 ± 0.16	0.95 ± 0.19	—	—	—	—	—	—	—
Five-landmarks (dense)	—	—	—	0.0002 ± 0.03	0.13 ± 0.11	0.87 ± 0.16	0.78 ± 0.19	—	—	—	—	—	—	—
Five-landmarks var. dropout (dense)	—	—	—	0.0008 ± 0.02	0.07 ± 0.02	0.76 ± 0.10	0.64 ± 0.18	—	—	—	—	—	—	—
Five-landmarks targeted dropout (dense)	—	—	—	0.004 ± 0.03	0.063 ± 0.11	0.71 ± 0.16	0.64 ± 0.20	—	—	—	—	—	—	—
Six-landmarks (dense)	—	—	—	1.52 ± 0.30	0.86 ± 0.29	1.07 ± 0.25	1.24 ± 0.24	—	—	—	—	—	—	—
Six-landmarks var. dropout (dense)	—	—	—	1.04 ± 0.30	1.18 ± 0.30	0.86 ± 0.28	1.06 ± 0.24	—	—	—	—	—	—	—
Six-landmarks targeted dropout (dense)	—	—	—	1.20 ± 0.29	0.92 ± 0.28	1.09 ± 0.24	1.21 ± 0.25	—	—	—	—	—	—	—
Nine-landmarks (dense)	—	—	—	—	—	—	—	—	—	—	—	—	—	—
Torosdagli et al. ⁵	0.03	0.27	1.01	0.41	1.36	0.68	0.35	—	—	—	—	—	—	—
Gupta et al. ¹⁰	—	—	3.20	1.62	1.53	2.08	—	—	—	—	—	—	—	—
Method	Maxillary-nasal bone landmarks													
	Ans	A	Pr	Pns	Na									
Three-landmarks regular (dense)	3.04 ± 0.39	3.04 ± 0.40	2.89 ± 0.40	2.04 ± 0.29	3.15 ± 0.34	—	—	—	—	—				
Three-landmarks cross (dense)	3.18 ± 0.39	3.14 ± 0.39	3.17 ± 0.38	2.61 ± 0.33	3.13 ± 0.37	—	—	—	—	—				
Five-landmarks var. dropout (MLP)	3.80 ± 0.44	3.95 ± 0.48	3.06 ± 0.01	3.85 ± 0.42	3.20 ± 0.34	—	—	—	—	—				

Table 2 (Continued).

Method	Mandibular landmarks								
	Cor _R	Cor _L	Cd _L	Gn	Pg	B	Id		
Five-landmarks (dense)	3.21 ± 0.27	3.16 ± 0.41	2.92 ± 0.42	2.37 ± 0.35	2.91 ± 0.40	—	—		
Five-landmarks var. dropout (dense)	3.15 ± 0.21	3.07 ± 0.38	3.09 ± 0.40	2.35 ± 0.32	3.14 ± 0.36	—	—		
Five-landmarks targeted dropout (dense)	3.17 ± 0.38	3.09 ± 0.39	2.85 ± 0.39	2.46 ± 0.32	3.14 ± 0.40	—	—		
Six-landmarks (dense)	0.79 ± 0.23	1.65 ± 0.29	1.51 ± 0.30	1.35 ± 0.34	—	—	—		
Six-landmarks var. dropout (dense)	1.16 ± 0.25	0.74 ± 0.22	1.60 ± 0.29	1.54 ± 0.31	—	—	—		
Six-landmarks targeted dropout (Dense)	0.76 ± 0.22	1.61 ± 0.28	1.51 ± 0.30	1.46 ± 0.36	—	—	—		
Nine-landmarks (Dense)	3.06 ± 0.37	3.05 ± 0.37	2.82 ± 0.35	2.42 ± 0.32	3.02 ± 0.33	—	—		
Torosdagli et al. ⁵	—	—	—	—	—	—	—		
Gupta et al. ¹⁰	1.42	1.73	—	2.08	—	—	1.17		

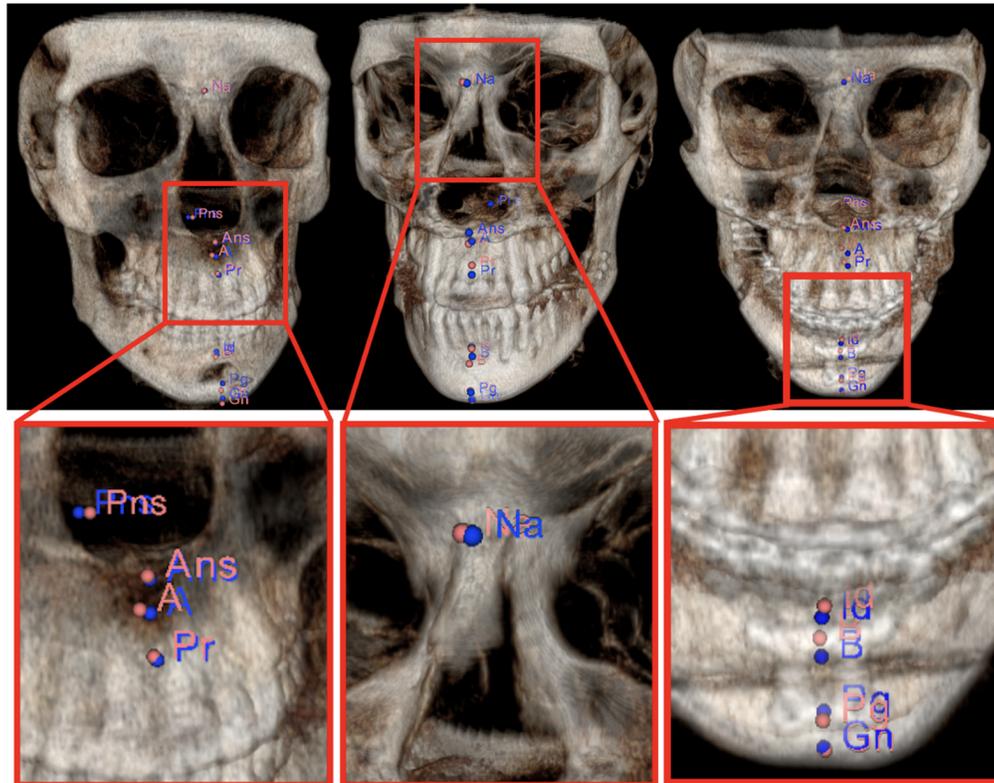


Fig. 7 Landmark annotations using the five-landmarks configuration: Ground truth in blue and computed landmarks in pink. (a) Genioplasty/chin advancement (male 43 year old), (b) malocclusion (mandibular hyperplasia, maxillary hypoplasia) surgery (male 19 year old), (c) malocclusion (mandibular hyperplasia, maxillary hypoplasia) surgery (female 14 year old). Note that landmarks are shown on the volume-rendered CBCT scans; there is no segmentation conducted.

and corresponding renderings in Fig. 7. The method achieving the minimum error for a corresponding landmark is colored the same as the corresponding landmark at Table 2. As shown by results, the minimum number of the input landmarks required for successful identification of other landmarks is determined as 3.

Among two different RU architectures, DB architecture was evaluated to be more robust and fast to converge as compared to the MLP architecture. To be self-complete, we provided the MLP experimental configuration performances only for the five-landmark experiment (See Table 2).

In the first experiment (Fig. 6, Experiment 1), to have an understanding of the performance of the RRN, we used the landmark grouping sparsely spaced and closely-spaced as proposed in Torosdagli et al.⁵ We named our first configuration as “five-landmarks” where closely spaced maxillary and nasal bones’ landmarks are predicted based on the relation of sparsely spaced landmarks (Fig. 6). In the five-landmarks RRN architecture, there were totally 25 RUs. In the second experiment (Fig. 6, Experiment 2), we explored the impact of a configuration with a smaller number of input mandibular landmarks on the learning performance. Compared to the five sparsely spaced input landmarks, we learned the relation of three landmarks, Me, Cd_L, and Cd_R, and predicted the closely-spaced landmark locations (as in the five-landmarks experiment) plus superior-anterior landmarks Cor_L and Cor_R and maxillary and nasal bones’ landmark locations. The network was composed of nine RUs. The training was relatively fast compared to the five-landmarks configuration due to small number of RUs. We named this method as “three-landmarks regular.”

After observing statistically similar accuracy compared to the five-landmarks method for the closely-spaced landmarks ($P > 0.005$), and high error rates at the superior–anterior landmarks Cor_L and Cor_R, we set up a new experiment, which we named “three-landmarks cross” (Fig. 6, Experiment 3). In this configuration, the third experiment, we used one superior–posterior and

one superior–anterior landmarks on the right and left sides, respectively. This network was similar to three-landmarks regular one in terms of number of RUs used.

In the fourth experiment (Fig. 6, Experiment 4), we evaluated the performance of the system in learning the closely-spaced mandibular landmarks (Gn, Pg, B, Id) and the maxillary landmarks (ANS, A, Pr, PNS) using the relation information of the sparsely spaced and the nasal-bones landmarks which is named as “six-landmarks.” There are a total of 36 RUs in this configuration.

In the last experiment (Fig. 6, Experiment 5), we aimed to learn the maxillary landmarks (ANS, A, Pr, PNS) and nasal bones landmark (Na) using the relation of the mandibular networks; hence, this network configuration is called “nine-landmarks.” The architecture was composed of 81 RUs. Owing to the high number of RUs in the architecture, the training of this network was the slowest among all the experiments performed.

For three challenging CBCT scans, Fig. 7 presents the ground-truth and the predicted landmarks with respect to the five-landmarks configuration DB architecture, annotated in blue and pink, respectively. We evaluated five-landmarks configuration for both MLP and the DB architectures using variational-dropout as regularizer (Table 2). For fourfolds, we observed that DB architecture was robust and fast-to-converge. Although, the performances were statistically similar for the mandibular landmarks, this was not the case for the maxillary and the nasal bone landmarks. The performance of the MLP architecture degrades notably compared to the decrease in the DB architecture for the maxilla and nasal bone landmarks.

Three-landmarks and five-landmarks configurations (Table 2) performed statistically similar for the mandibular landmarks. Interestingly, both three-landmarks configurations performed slightly better for the neighboring bone landmarks. This reveals the importance of optimum number of landmarks in the configuration.

In comparison of five-landmarks and six-landmarks configurations (Table 2), we observed that five-landmarks configuration is good at capturing the relations on the same bone. In contrast, six-landmarks configuration was good at capturing the relations on the neighboring bones. Although, the error rates were <2 mm, potentially redundant information induced by the Na landmark in the six-landmarks configuration caused the performance to decrease notably for the mandibular landmarks compared to the five-landmarks configuration.

Nine-landmarks configuration performed statistically similar to five-landmarks configuration, however, due to 81 RUs employed for the nine-landmarks, the training was slower.

Although direct comparison was not possible, we compared our results with Gupta et al.¹⁰ based on the landmark distances. We found that our results were significantly better for all landmarks except the Na landmark. The framework proposed at Ref. 10 uses an initial seed point using a 3D template registration at the inferior–anterior region where fractures are the most common. Eventually, any anatomical deformity that alters the anterior mandible may cause an error in the seed localization, which can lead to a suboptimal outcome.

We evaluated the performance of the proposed system when variational²⁶ and targeted²⁷ dropouts were employed. Although there was no statistically significant difference between dropouts in terms of accuracy, convergence of the systems was relatively fast (around 20 epochs compared to 100 when using regular dropout) for the MLP architecture. Hence, for the MLP architecture, in terms of computational resources, variational and targeted dropout implementations were far more efficient for our proposed system. This is particularly important because when there are a large number of RUs, one may focus more on the efficiency rather than accuracy. When the DB architecture was employed, we did not observe any performance improvement among different dropout implementations.

In landmarking, extreme performance would be very important. For example, the outliers would hamper the entire planning. Therefore, we have carefully checked the outliers for each landmarking problem, and found that there are <10 outliers in a total of 250 patients’ scans for each configuration, and the highest number of outliers was 7. In the best working experimental setup (six-landmark configuration), for both menton and condylar left, the highest errors we obtained with outliers were 1.5 mm. For coronoid left and right, errors of 2.75 and 0.3 mm were obtained, respectively. For infradentale, we obtained 1.5 and 1.7 mm errors by two outliers, all measured in volumetric spaces. Our results were consistent and robust to outliers,

as hypothesized before. In the three-landmark configuration, the highest error we obtained with an outlier was 5 mm to detect coronoid right.

It is also important to explore whether severity of the cases or metal artifacts can influence the final results, considered under the robustness measure. We specifically evaluated the performance of our method within the subsets of data containing significant artifacts and others (30% versus 70%, see data subsection). Statistically, there was no difference in the performance of our landmarking method when applied a *t*-test between these two groups, indicating a robustness of our method.

4 Discussion and Conclusion

We proposed the RRN framework which learns spatial dependencies between CMF landmarks in an end-to-end manner. Without the need for an explicit segmentation, we hypothesized that there is an inherent geometrical relation among CMF landmarks which can be learned using the relational reasoning architecture. Although, appearance-based deep-learning approaches are very strong alternative to what we proposed herein, generalization is still an unsolved and a very challenging problem, and reasoning is not directly applicable unlike geometric relations. For instance, authors in Ref. 34 used a two-step neural networks with head neck CT data, achieving an average of 2.64 mm localization error; however, their data set does not include any severe pathology, and still performance is inferior to what we have proposed here. The presented solution was shown to be effective in 2D images with normal anatomy. Further, appearance-based methods for landmark detection in CT scans,^{35,36} which can be considered related to our work, define landmarks as an anatomical region (ROI) comparatively larger than our landmark definition (25×25 versus 3×3), and again no deformation or pathology presence exist therein. In contrast to these methods, our method considers a very small area as landmark, and we use extremely challenging pathological cases, which also differentiates the current work from our previous work where we used a segmentation-based approach in the geodesic space.

Our relational reasoning framework, which is a model-based approach, can generalize well to the unseen data. Hence, once trained, RRN can be used at the same testing precision to detect the missing landmarks of the unseen data taken at completely different conditions. This would afford better outcomes for precision medicine and complex CMF deformities. In our experiments, we first evaluated this claim using a dataset with a high amount of bone deformities in addition to other CBCT challenges. We observed that (1) despite the large amount of deformities that may exist in the CMF anatomy, there is a functional relation between the CMF landmarks, and (2) RNN frameworks are strong enough to reveal this latent relation information. Next, we evaluated the detection performance of five different configurations of the input landmarks to find out the optimum configuration. We observed that not all landmarks are equally informative in the detection performance. Some landmark configurations are good in capturing the local information, while some have both good local and global prediction performance.

One may wonder how to choose which landmark configuration and number for a current task should be chosen by a user. Rationally, for 3D modeling and visualizations, a higher number of landmarks would benefit the final outcomes. However, our aim herein was to explore what the minimum number of landmarks and sufficient configuration of landmarks were to have successful landmarking. For example, we have found that cross landmark configuration keeps more information than regular configuration. Also, we found that five and/or six landmarks were often enough to capture the anatomical relationships compared with nine-landmark configurations. Our study reveals certain insights about how to create networks specific to anatomies and learn efficiently with minimal, but necessary data. In practical terms, we intend to predict even a higher number of landmarks, however, this was our current limitation in our study due to the availability of ground truth labels. Overall, per-landmark error for the six-landmarks configuration is <2 mm, which is considered as a clinically acceptable level of success. It should also be noted that landmarking error depends on the voxel size too.³⁷ Given the voxel size of 0.5 to 0.7 mm in our data set, the accuracy of our landmarking is around 1 or 2 voxels where 3 voxel sizes of error are often considered as a highly accurate landmarking procedure.³⁸ One should be aware that clinically accepted level of landmarking in CBCT scans may vary depending on the specific

application. For instance, in the context of dental implant planning, it is important to locate the maxillary and mandibular sinuses and the mental foramen, as these structures can affect the placement of the implant. High level of accuracy may be necessary to ensure the implant, 2 mm is considered as a safe, acceptable level in this manner.^{37,39} On the other hand, CBCT-based orthodontic treatment planning, a lower level of accuracy (i.e., 4 mm) may be acceptable, as the primary focus is on the overall shape, and size rather than precise location of specific landmarks.

In our implementation, we showed that other deep-learning networks can be integrated well into our platform as long as features are encoded via RUs. While one may argue whether changing specific parameters could make these predictions better. However, such incremental explorations are kept outside the main paper but are worth exploring in future studies from an optimization point of view. Moreover, for now RRN only employs spatial information (proof-of-concept stage), its extension could include using shape space learned landmark relationships as a conditional shape prior. Similarly, the use of those learned relationship as a look up table (atlas) is another venue that needs further exploration.

Our study has a number of limitations. For instance, we confined ourselves to manifold data (position of the landmarks and their geometric relations) without use of appearance information because one of our aims was to avoid explicit segmentation to be able to use simple geometric reasoning networks. As an extension of this study, we will incorporate appearance features from medical images to explore whether these features are superior to purely geometric features, or combined (hybrid) features can have additive value in this research domain. One alternative way to pursue the research that we initiated herein will be to explore deeper and more efficient networks. Hence explore how to scale up in to a much wider platform where large number of landmarks and various clinical problems are addressed. We believe that such advances will improve the current technology for 3D visualization and even afford embedding augmented reality in treatment and surgical planning.

Disclosures

No conflicts of interest.

Acknowledgments

We thank Mary McIntosh for helping data collection and landmarking. This work was partially supported by the NIH, grants R01-CA246704 and R01-CA240639.

Data, Materials, and Code Availability

Data was collected under IRB approved (PI: Janice S. Lee), and maybe accessible under certain agreement with the NIDCR. Code is available upon request.

References

1. J. J. Xia, J. Gateno, and J. F. Teichgraber, "A paradigm shift in orthognathic surgery: a special series part I," *J. Oral Maxillofacial Surg.* **67**(10), 2093–2106 (2009).
2. J. Zhang et al., *Joint Craniomaxillofacial Bone Segmentation and Landmark Digitization by Context-Guided Fully Convolutional Networks*, pp. 720–728, Springer International Publishing, Cham (2017).
3. N. Anuwongnukroh et al., "Accuracy of automatic cephalometric software on landmark identification," *IOP Conf. Ser.: Mater. Sci. Eng.* **265**(1), 012028 (2017).
4. Y. Lang et al., "Automatic localization of landmarks in craniomaxillofacial CBCT images using a local attention-based graph convolution network," *Lect. Notes Comput. Sci.* **12264**, 817–826 (2020).
5. N. Torosdagli et al., "Deep geodesic learning for segmentation and anatomical landmarking," *IEEE Trans. Med. Imaging* **38**(4), 919–931 (2018).

6. N. Torosdagli et al., “Robust and fully automated segmentation of mandible from CT scans,” in *IEEE 14th Int. Symp. Biomed. Imaging (ISBI 2017)*, pp. 1209–1212 (2017).
7. J. Zhang et al., “Context-guided fully convolutional networks for joint craniomaxillofacial bone segmentation and landmark digitization,” *Med. Image Anal.* **60**, 101621 (2020).
8. S. Shahidi et al., “The accuracy of a designed software for automated localization of craniofacial landmarks on cbct images,” *BMC Med. Imaging* **14**, 32 (2014).
9. J. Negrillo-Cárdenas et al., “Automatic detection of landmarks for the analysis of a reduction of supracondylar fractures of the humerus,” *Med. Image Anal.* **64**, 101729 (2020).
10. A. Gupta et al., “A knowledge-based algorithm for automatic detection of cephalometric landmarks on CBCT images,” *Int. J. Comput. Assist. Radiol. Surg.* **10**, 1737–1752 (2015).
11. D. Shen, G. Wu, and H.-I. Suk, “Deep learning in medical image analysis,” *Annu. Rev. Biomed. Eng.* **19**(1), 221–248 (2017).
12. S. E. Bahrapour et al., “The accuracy of a designed software for automated localization of craniofacial landmarks on CBCT images,” *BMC Med. Imaging* **14**, 32 (2014).
13. X. Li et al., “Comprehensive evaluation of ten deformable image registration algorithms for contour propagation between CT and cone-beam CT images in adaptive head and neck radiotherapy,” *PLoS One* **12**, 1–17 (2017).
14. J. Zhang et al., “Automatic craniomaxillofacial landmark digitization via segmentation-guided partially-joint regression forest model and multiscale statistical features,” *IEEE Trans. Biomed. Eng.* **63**, 1820–1829 (2016).
15. X. Chen et al., “Fast and accurate craniomaxillofacial landmark detection via 3D faster r-CNN,” *IEEE Trans. Med. Imaging* **40**(12), 3867–3878 (2021).
16. D. Raposo et al., “Discovering objects and their relations from entangled scene representations,” arxiv:1702.05068 (2017).
17. H. Yan and C. Song, “Multi-scale deep relational reasoning for facial kinship verification,” *Pattern Recognit.* **110**, 107541 (2021).
18. F. Scarselli et al., “The graph neural network model,” *IEEE Trans. Neural Netw.* **20**, 61–80 (2009).
19. Y. Li et al., “Gated graph sequence neural networks,” (2015).
20. K. Xu et al., “How powerful are graph neural networks?” arXiv:1810.00826 (2018).
21. P. Battaglia et al., “Interaction networks for learning about objects, relations and physics,” in *Proc. 30th Int. Conf. Neural Inf. Process. Syst., NIPS’16*, Curran Associates Inc., USA, pp. 4509–4517 (2016).
22. A. Santoro et al., “A simple neural network module for relational reasoning,” in *NIPS* (2017).
23. P. W. Battaglia et al., “Relational inductive biases, deep learning, and graph networks,” arXiv:1806.01261 (2018).
24. B. Zhou et al., “Temporal relational reasoning in videos,” *Lect. Notes Comput. Sci.* **11205**, 803–818 (2018).
25. F. Lalys et al., “Automatic aortic root segmentation and anatomical landmarks detection for tavi procedure planning,” *Minim. Invasive Ther. Allied Technol.* **28**(3), 157–164 (2018).
26. D. P. Kingma, T. Salimans, and M. Welling, “Variational dropout and the local reparameterization trick,” in *Proc. 28th Int. Conf. Neural Inf. Process. Syst. – Vol. 2, NIPS’15*, MIT Press, Cambridge, MA, USA, pp. 2575–2583 (2015).
27. A. N. Gomez et al., “Learning sparse networks using targeted dropout,” arXiv:1905.13678 (2019).
28. A. C. V. Armond et al., “Influence of third molars in mandibular fractures. Part 1: mandibular angle—a meta-analysis,” *Int. J. Oral Maxillofacial Surg.* **46**(6), 716–729 (2017).
29. U. Bagci, X. Chen, and J. K. Udupa, “Hierarchical scale-based multiobject recognition of 3-d anatomical structures,” *IEEE Trans. Med. Imaging* **31**, 777–789 (2011).
30. X. Chen and U. Bagci, “3D automatic anatomy segmentation based on iterative graph-CUT-ASM,” *Med. Phys.* **38**(8), 4610–4622 (2011).
31. S. Rueda, J. K. Udupa, and L. Bai, “Shape modeling via local curvature scale,” *Pattern Recognit. Lett.* **31**(4), 324–336 (2010).
32. N. Srivastava et al., “Dropout: a simple way to prevent neural networks from overfitting,” *J. Mach. Learn. Res.* **15**, 1929–1958 (2014).

33. Y. Gal and Z. Ghahramani, "A theoretically grounded application of dropout in recurrent neural networks," in *Proc. 30th Int. Conf. Neural Inf. Process. Syst., NIPS'16*, Curran Associates Inc., USA, pp. 1027–1035 (2016).
34. Y. Zheng et al., "3D deep learning for efficient and robust landmark detection in volumetric data," *Lect. Notes Comput. Sci.* **9349**, 565–572 (2015).
35. D. Yang et al., "Automatic vertebra labeling in large-scale 3D CT using deep image-to-image network with message passing and sparsity regularization," *Lect. Notes Comput. Sci.* **10265**, 633–644 (2017).
36. F. C. Ghesu et al., "Robust multi-scale anatomical landmark detection in incomplete 3D-CT data," *Lect. Notes Comput. Sci.* **10433**, 194–202 (2017).
37. K.-M. Lee et al., "Effect of voxel size on the accuracy of landmark identification in cone-beam computed tomography images," *J. Korean Dental Sci.* **12**(1), 20–28 (2019).
38. R. Ganguly, A. Ramesh, and S. Pagni, "The accuracy of linear measurements of maxillary and mandibular edentulous sites in cone-beam computed tomography images with different fields of view and voxel sizes under simulated clinical conditions," *Imaging Sci. Dent.* **46**(2), 93–101 (2016).
39. A. Ghowsi et al., "Automated landmark identification on cone-beam computed tomography: accuracy and reliability," *Angle Orthod.* **92**(5), 642–654 (2022).

Biographies of the other authors are not available.