Three-Dimensional Shape Modeling for Automated Assessment of Orthopedic Implant Positioning and Anatomical Compatibility

Tymon Rej
Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
hellotymon@buffalo.edu

Abstract

Automated assessment of orthopedic implant positioning requires more than surface alignment: the selected pose must respect patient-specific anatomy, preserve safe clearance from cortical bone, avoid unintended load transfer, and remain stable under image and segmentation uncertainty. This manuscript develops a coherent three-dimensional shape-modeling framework that couples statistical anatomical modeling, surface registration, signed-distance clearance analysis, pose optimization, and uncertainty-aware decision logic for implant compatibility assessment. The central contribution is a formulation in which anatomical shape variation and implant pose are represented in a shared geometric space, and compatibility is evaluated through continuous clearance and fit fields rather than isolated landmark errors. Statistical shape models provide a low-dimensional representation of plausible anatomy and support correspondence across cases. Registration estimates the transform between preoperative or intraoperative geometry and a canonical anatomical space, while signed distance fields convert pose hypotheses into interpretable penetration, clearance, and contact measures. A constrained objective combines data fit, minimum clearance, and penalization of cortical breach. The framework also separates segmentation, registration, shape-truncation, and model-discrepancy uncertainty, and it converts the resulting distribution of compatibility scores into an acceptance rule with explicit safety margins. Analytical discussion focuses on objective-function behavior, failure modes such as correspondence drift and over-constrained clearance penalties, and validation endpoints suitable for prospective clinical evaluation. The formulation supports reproducible, patient-specific implant assessment while preserving clinical interpretability and avoiding dependence on a single imaging modality or proprietary planning system.

Keywords

statistical shape modeling; implant positioning; three-dimensional registration; signed distance fields; orthopedic surgery; uncertainty quantification

1. Introduction

Orthopedic implant placement is a geometric and mechanical decision made under anatomical variability. A femoral stem, acetabular cup, or patient-specific plate must occupy a pose that restores function without violating cortical boundaries, producing adverse contact, or creating an unstable fixation scenario. Conventional preoperative planning often reduces this decision to a small set of angles, offsets, and implant sizes measured on two-dimensional projections or selected three-dimensional views. Such reductions are clinically useful but lose the spatial relationships that determine whether an implant is actually compatible with the host bone. Three-dimensional shape modeling offers a more complete representation: it

describes the patient's anatomy as a surface or signed-distance field, represents the implant as a parameterized rigid object, and evaluates pose hypotheses through continuous geometric criteria.

Surface registration provides the technical foundation for aligning preoperative models, intraoperative scans, and implant templates [1]. Statistical shape models extend this foundation by encoding plausible anatomical variation across a population and by supplying a low-dimensional space in which incomplete or noisy observations can be interpreted [2]. Reviews of statistical shape modeling in medical imaging show that the quality of such models depends on correspondence, shape parameterization, and the relationship between model variability and the imaging evidence [3]. These dependencies matter for implant assessment because an incorrect correspondence can shift the apparent cortical boundary and therefore change the recommended implant size or pose.

Correspondence construction remains a central methodological issue. Landmark-based, mesh-based, and optimization-based approaches trade anatomical specificity against reproducibility [4]. Minimum description length formulations seek compact models by penalizing correspondence complexity [5], while coherent point drift treats registration as a probabilistic density estimation problem and can accommodate non-rigid deformation [6]. Segmentation systems that combine statistical shape knowledge with convolutional networks demonstrate how anatomical priors can stabilize bone and cartilage boundaries in noisy images [7]. Evaluation of three-dimensional segmentation also requires metrics that account for surface distance and volume overlap rather than a single scalar [8]. These developments are relevant to implant assessment, but they do not by themselves answer the pose question: a well-segmented bone may still be paired with an implant pose that creates local breach, excessive clearance, or mechanically unfavorable contact.

The clinical need is substantial because primary and revision arthroplasty volumes continue to rise [9]. Automated planning therefore cannot be limited to segmentation accuracy; it must produce a defensible compatibility judgment. Landmark detection on three-dimensional scans is one enabling component, and segmentation-guided semi-supervised strategies have been proposed to reduce annotation burden when landmarks are sparse [10]. Robotic and navigated arthroplasty systems have increased interest in quantitative placement, but their accuracy depends on the planning target and on the geometric criteria used to define that target [11]. Mechanical context is equally important: hip contact forces during routine activities establish load paths that influence implant design and fixation [12], while three-dimensional preoperative planning studies show that reconstruction accuracy depends on how position and orientation are specified [13]. A shape-modeling framework that unifies these geometric, mechanical, and uncertainty considerations is therefore needed.

The central contribution of this manuscript is a methodological framework for automated implant positioning assessment that represents anatomy and implant geometry in a common coordinate system, evaluates pose through signed-distance clearance and fit fields, and incorporates uncertainty into an explicit acceptance rule. The framework is developed analytically rather than as a report of a completed clinical trial. Its purpose is to specify the variables, constraints, objective functions, and failure modes that a rigorous implementation must address, and to identify validation endpoints that can distinguish geometric accuracy from clinical compatibility.

2. Clinical and Technical Requirements

The assessment problem begins with a patient-specific anatomical model derived from computed tomography, magnetic resonance imaging, or another three-dimensional source. The model may be a triangulated surface, a volumetric label map, or a signed-distance field. Each representation has consequences. Surface meshes are efficient for visualization and contact queries but can be sensitive to meshing artifacts. Volumetric labels preserve interior information but require an explicit boundary extraction step. Signed-distance fields provide continuous queries and stable gradients near the surface, which is advantageous for optimization, but they require careful sign conventions and consistent resolution.

The implant is represented as a rigid surface with known geometry and a pose parameterized by rotation and translation. Size selection may be discrete, while pose is continuous. Compatibility has at least three geometric components. First, the implant should align with the planned cavity or host surface within a tolerance that reflects surgical preparation. Second, the implant should maintain a minimum clearance from critical cortical regions unless controlled contact is intended. Third, contact should be distributed in mechanically plausible zones rather than concentrated at a single edge. These requirements are not interchangeable. A pose that minimizes average surface error may still create a focal breach because the average suppresses local violations. Conversely, a pose that eliminates all clearance violations may be over-constrained and may sacrifice alignment with the intended mechanical axis.

Anatomical variability further complicates the decision. Femoral anteversion, canal curvature, cortical thickness, and acetabular coverage vary across patients [13]. A population-level shape model can regularize an incomplete observation, but it can also bias the estimate toward the mean anatomy. The framework therefore treats shape-model regularization as a prior rather than as ground truth. The model should be sufficiently expressive to represent the observed anatomy, and its residual should be reported as a measure of model discrepancy.

3. Methodological Framework

3.1 Shape representation and anatomical compatibility field

The anatomical surface is represented by a signed-distance function. For a surface S, the field is defined so that its zero level set is the surface, negative values lie inside the host bone, and positive values lie in accessible space. This convention provides a continuous query for clearance and penetration without repeatedly computing mesh intersections.

$$\varphi_S(x) \;=\; s_S(x) \cdot min_\{y \in S\}\, ||x \;-\; y||_2$$

Here x is a query point in image or planning space, y is a point on the anatomical surface S, $||x - y||_2$ is the Euclidean distance, and s_S(x) is a sign function that is negative inside the host bone and positive outside it. The field φ_S supports local clearance evaluation, but it does not encode material properties or load transfer. Mechanical compatibility is therefore assessed separately through contact distribution, implant design, and bone-quality information.

3.2 Statistical shape modeling and correspondence

A statistical shape model represents each anatomical instance as a mean shape plus a weighted combination of principal modes. Let x be a vector of corresponding surface points. The model is written as a constrained linear expansion.

$$x \;=\; \bar{x} \;+\; \sum_\{k=1\}\^\{m\}\, b_k\, u_k,\; |b_k| \;\leq\; 3\sqrt{\lambda_k}$$

In this expression, x̄ is the mean shape, u_k is the kth mode of variation, b_k is the shape coefficient, λ_k is the corresponding eigenvalue, and m is the number of retained modes. The bound limits each coefficient to three standard deviations along its mode. This constraint prevents implausible extrapolation when the image evidence is weak. The choice of m controls a bias-variance trade-off: too few modes cannot represent the patient, while too many modes allow noise to be absorbed into shape variation.

Correspondence quality determines whether the modes are anatomically meaningful. Landmark-based, closest-point, and optimization-based correspondence methods can produce different models from the same training set [4]. Minimum description length criteria provide one principled way to select among correspondences by favoring compact models [5]. In an implant workflow, correspondence should be evaluated not only by model compactness but also by whether landmarks near implant-host interfaces remain stable under perturbation. A correspondence that is compact globally but unstable at the femoral neck or acetabular rim is unsuitable for pose assessment.

3.3 Registration and pose estimation

Registration aligns the patient-specific anatomy, the statistical shape model, and the implant template in a common space. A general registration energy combines a data term with a regularization term.

$$E_reg(T) = \sum_\{i = 1\}\^\{n\} w_i \, ||T(x_i) - y_i||_2^2 + \lambda R(T)$$

Here T is the transform, x_i are model or source points, y_i are target points or surface correspondences, w_i are robust weights that reduce the influence of outliers, R(T) is a regularization term that penalizes implausible deformation or large rigid deviation, and λ balances data fidelity against regularity. For rigid implant placement, T is restricted to rotation and translation. For anatomical registration, a non-rigid extension may be required, but it should be separated from the rigid implant pose so that anatomical deformation is not mistakenly used to justify an incompatible implant position.

Medical image registration surveys emphasize that the choice of transformation model, similarity metric, and optimization strategy must match the clinical task [14]. Point-set registration methods such as coherent point drift or non-rigid point matching can accommodate missing or noisy correspondences [6, 15], but they require safeguards against correspondence drift. In the proposed framework, anatomical registration supplies a canonical coordinate system, while implant pose is optimized in that system using explicit clearance constraints.

3.4 Implant compatibility and clearance analysis

The implant surface is sampled at points z_j in implant coordinates. For a pose parameter p, the rigid transform T_p maps these points into planning space. The signed clearance at sample j is the value of the host-bone distance field at the transformed sample.

$$d_j(p) = \varphi_B\left(T_p z_j\right)$$

Here φ_B is the signed-distance field of the prepared host bone, positive in the accessible cavity and negative inside cortical bone, and T_p is the transform determined by rotation and translation parameters p. A positive d_j indicates clearance, while a negative value indicates penetration. This definition makes local violations explicit and avoids the averaging effect of a single surface error.

Pose selection is formulated as a constrained objective. The data term rewards alignment with the planned cavity or anatomical reference. The clearance term penalizes samples that fall below a minimum safe margin. An additional penalty suppresses cortical breach and edge loading.

$$E_pose(p) = D(I_p, S_B) + \lambda_c \sum_\{j=1\}\^\{q\} [\delta_c - d_j(p)]_+ \^ 2$$

In this expression, I_p is the implant surface at pose p, S_B is the prepared bone surface or planned cavity, D(I_p, S_B) is a surface-to-surface fit measure, q is the number of implant samples, δ_c is the minimum clearance margin, [·]_+ denotes the positive part, and λ_c controls the penalty for margin violations. The positive-part operator ensures that samples with adequate clearance do not dominate the objective. The fit measure should be robust to localized surface irregularities; otherwise a single segmentation artifact can force the implant into an anatomically implausible pose.

A continuous compatibility score summarizes the spatial distribution of clearance and contact. It is useful for comparing pose candidates and for communicating risk.

$$C(p) = (1/|\partial I|) \int_\{\partial I\} w(x)\, g(d(x;p))\, dA$$

Here ∂I is the implant surface, |∂I| is its area, w(x) is a local weight reflecting anatomical importance, g is a monotone function that maps clearance to a compatibility value, and d(x;p) is the signed clearance at surface point x under pose p. The weight w(x) can emphasize regions near cortical bone, fixation zones, or articular surfaces. The integral form makes the score sensitive to spatial distribution rather than only to the minimum. It also permits direct interpretation: a high score requires that mechanically and anatomically important regions maintain acceptable clearance and contact.

Mechanical compatibility requires additional considerations. Stress shielding and bone resorption around total hip stems depend on the relationship between implant stiffness, load transfer, and bone remodeling [18]. Femoral component design must respect the anatomic basis of the proximal femur, including canal geometry and metaphyseal fit [19]. These factors cannot be reduced to a single clearance threshold. The framework therefore separates geometric compatibility, which is computed from shape and pose, from mechanical compatibility, which requires load cases and material assumptions. A pose that is geometrically admissible may still be mechanically unfavorable if it creates abnormal load transfer.

3.5 Uncertainty and decision logic

Automated assessment must represent uncertainty rather than presenting a single pose as exact. Segmentation error, registration error, shape-model truncation, and model discrepancy all contribute to uncertainty in clearance and compatibility. A first-order variance decomposition can be used when components are approximately independent.

$$\sigma^2_total = \sigma^2_seg + \sigma^2_reg + \sigma^2_shape + \sigma^2_mod$$

Here σ²_seg represents segmentation-induced variance, σ²_reg represents registration variance, σ²_shape represents variance from truncated or misspecified shape modes, and σ²_mod represents model discrepancy from geometric simplifications or material assumptions. The independence assumption is a working approximation; when dependencies are strong, covariance terms should be estimated or the components should be propagated jointly.

Modern uncertainty quantification for learned models distinguishes aleatoric uncertainty, which arises from observation noise, from epistemic uncertainty, which arises from limited model knowledge [16, 17]. In implant assessment, segmentation variability is partly aleatoric and partly epistemic, while shape-model truncation is primarily epistemic. The decision rule should therefore use a lower confidence bound on the compatibility score and a hard constraint on minimum clearance.

$$accept\ if\ C_LB(p)\ \geq\ C_min\ and\ d_min(p)\ \geq\ \delta_s$$

In this rule, C_LB(p) is the lower confidence bound of the compatibility score, C_min is the minimum acceptable score, d_min(p) is the minimum signed clearance over implant samples, and δ_s is the safety threshold. The two conditions serve different purposes. The compatibility score reflects overall fit and contact distribution, while the minimum clearance condition prevents a locally unsafe pose from being accepted on the basis of a favorable average.

4. Analytical Assessment and Design Trade-offs

4.1 Objective-function behavior

The pose objective combines a data term and a clearance penalty. If the data term dominates, the optimization may align the implant closely with a planned cavity while allowing small cortical breaches. If the clearance penalty dominates, the optimizer may push the implant away from the planned surface and create excessive clearance or unsatisfactory contact. The balance parameter λ_c therefore controls a clinically meaningful trade-off between alignment and safety. It should not be selected solely by optimization convenience. A defensible approach is to calibrate λ_c against expert-defined acceptability regions on retrospective cases and to report sensitivity to its value.

The positive-part penalty also affects gradients. Samples far above the clearance threshold contribute nothing, while samples below the threshold contribute a quadratic penalty. This sparsity is desirable because most implant surface points may be adequately clear, but it can produce discontinuous active sets during optimization. A smooth approximation or a continuation strategy can reduce oscillation. The minimum clearance constraint should still be checked exactly after optimization, because a smoothed penalty does not guarantee feasibility.

The compatibility integral introduces spatial weighting. Uniform weights treat all implant surface regions equally, which is rarely appropriate. Cortical regions and fixation zones require different tolerances from non-load-bearing regions. The choice of w(x) is therefore a modeling decision that should be documented and tested. If the weights are too sharply localized, the score becomes sensitive to segmentation noise; if they are too diffuse, the score loses anatomical specificity.

4.2 Failure modes

Several failure modes are intrinsic to shape-based implant assessment. Correspondence drift can misalign anatomical landmarks and shift the apparent cortical boundary. Registration ambiguity can produce a visually plausible transform that is incorrect along the axis of a cylindrical bone, where surfaces provide weak constraints. Shape-model bias can pull an atypical patient toward the mean, especially when the imaging field of view is incomplete. Signed-distance discretization can create staircase artifacts that affect clearance queries near thin cortical regions. Finally, implant sampling density can miss local penetrations if samples are too sparse.

These failure modes suggest specific safeguards. Correspondence should be evaluated with landmark error, surface distance, and stability under perturbation. Registration should report multiple plausible alignments when the geometry is ambiguous. Shape-model residuals should be inspected rather than suppressed. Clearance queries should use a continuous interpolant or a sufficiently fine grid. Implant sampling should be adaptive, with denser sampling near high-curvature regions and expected contact zones. None of these safeguards eliminates uncertainty, but they make its sources visible.

4.3 Validation strategy and endpoints

Validation should distinguish geometric accuracy from clinical compatibility. Geometric endpoints include surface distance between the predicted and reference implant pose, minimum clearance, penetration volume, contact area, and contact distribution. These endpoints can be computed from expert-approved planning data or phantom studies without requiring clinical outcome claims. Clinical endpoints include intraoperative correction rate, implant size change, and adverse events, but these require prospective study designs and appropriate governance.

A rigorous evaluation should also test sensitivity to segmentation and registration perturbations. If a small perturbation changes the recommended implant size or violates the clearance constraint, the assessment is not robust. Calibration of the compatibility score is equally important: a score threshold that accepts unsafe poses is unacceptable even if average agreement with expert planning is high. The decision rule should therefore be evaluated by its false-accept and false-reject behavior under clinically relevant tolerances, not only by correlation with a reference pose.

5. Implementation Considerations and Clinical Translation

An implementation of the framework requires a consistent geometric pipeline. Imaging data must be resampled to isotropic spacing where possible, and segmentation must produce both a label and a confidence estimate. Surface extraction should preserve cortical boundaries and avoid smoothing that removes thin but clinically important structures. The implant library must be represented in a common coordinate system with verified dimensions. Pose optimization should be initialized from anatomical landmarks or a coarse registration and then refined under the clearance objective.

Computational cost is dominated by repeated distance queries and surface sampling. A spatial acceleration structure such as a bounding volume hierarchy or an octree can reduce query time. However, the choice of data structure should not alter the clearance field near the surface. Validation should include numerical consistency checks between the accelerated and direct distance queries.

Clinical translation also requires interpretability. The output should not be a single unqualified pose. It should include the selected pose, clearance maps, contact maps, uncertainty summaries, and the binding constraints that determined the decision. When a pose is rejected, the system should indicate whether the limiting factor is penetration, insufficient clearance, poor fit, or high uncertainty. Patient-specific implants and three-dimensional printing add another layer: the digital model used for manufacturing must be traceable to the approved planning geometry and must preserve tolerances across the manufacturing workflow [20]. The framework therefore treats the planning model, the assessment model, and the manufacturing model as linked but separately verified artifacts.

6. Conclusion

Three-dimensional shape modeling can make implant positioning assessment more complete than landmark-based planning alone, but only if the geometry is connected to explicit compatibility criteria and uncertainty. This manuscript developed a framework that represents patient anatomy through statistical shape models and signed-distance fields, estimates registration and implant pose in a common space, evaluates local clearance and contact through continuous fields, and combines fit and safety through a constrained objective. The compatibility score summarizes spatial distribution, while the minimum clearance condition prevents locally unsafe poses from being accepted on average. Variance decomposition and a lower-confidence decision rule make uncertainty part of the assessment rather than an afterthought. The principal design trade-offs concern correspondence stability, shape-model expressiveness, clearance-penalty weighting, and the separation of geometric from mechanical compatibility. These trade-offs are clinically consequential and should be evaluated through perturbation studies, expert-defined tolerances, and prospective validation. The framework provides a technically coherent basis for automated, interpretable, and patient-specific implant assessment without reducing compatibility to a single alignment error.

References

1. Besl, P. J., & McKay, N. D. (1992). A method for registration of 3-D shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence, 14(2), 239–256. https://doi.org/10.1109/34.121791

2. Cootes, T. F., Taylor, C. J., Cooper, D. H., & Graham, J. (1995). Active shape models—their training and application. Computer Vision and Image Understanding, 61(1), 38–59. https://doi.org/10.1006/cviu.1995.1004

3. Heimann, T., & Meinzer, H.-P. (2009). Statistical shape models for 3D medical image segmentation: A review. Medical Image Analysis, 13(4), 543–563. https://doi.org/10.1016/j.media.2009.05.004

4. Styner, M. A., Rajamani, K. T., Nolte, L. P., Zsemlye, G., Székely, G., Taylor, C. J., & Davies, R. H. (2003). Evaluation of 3D correspondence methods for model building. Information Processing in Medical Imaging, 2732, 63–75.

5. Davies, R. H., Twining, C. J., Cootes, T. F., Waterton, J. C., & Taylor, C. J. (2002). A minimum description length approach to statistical shape modeling. IEEE Transactions on Medical Imaging, 21(5), 525–537. https://doi.org/10.1109/TMI.2002.1009388

6. Myronenko, A., & Song, X. (2010). Point set registration: Coherent point drift. IEEE Transactions on Pattern Analysis and Machine Intelligence, 32(12), 2262–2275. https://doi.org/10.1109/TPAMI.2010.46

7. Ambellan, F., Tack, A., Ehlke, M., & Zachow, S. (2019). Automated segmentation of knee bone and cartilage combining statistical shape knowledge and convolutional neural networks: Data from the Osteoarthritis Initiative. Medical Image Analysis, 52, 109–118. https://doi.org/10.1016/j.media.2018.11.009

8. Taha, A. A., & Hanbury, A. (2015). Metrics for evaluating 3D medical image segmentation: Analysis, selection, and tool. BMC Medical Imaging, 15, 29. https://doi.org/10.1186/s12880-015-0068-x

9. Kurtz, S., Ong, K., Lau, E., Mowat, F., & Halpern, M. (2007). Projections of primary and revision hip and knee arthroplasty in the United States from 2005 to 2030. Journal of Bone and Joint Surgery, 89(4), 780–785. https://doi.org/10.2106/JBJS.F.00222

10. Liu, S., Jin, J., Li, W., & Yan, X. (2026). Segmentation-Guided Semi-Supervised Landmark Detection on Intraoral 3D Scans. Preprints. https://doi.org/10.20944/preprints202609.1722.v1

11. Jacofsky, D. J., & Allen, M. (2016). Robotics in arthroplasty: A comprehensive review. The Journal of Arthroplasty, 31(10), 2353–2363. https://doi.org/10.1016/j.arth.2016.01.074

12. Bergmann, G., Deuretzbacher, G., Heller, M., Graichen, F., Rohlmann, A., Strauss, J., & Duda, G. N. (2001). Hip contact forces and gait patterns from routine activities. Journal of Biomechanics, 34(7), 859–871. https://doi.org/10.1016/S0021-9290(01)00040-9

13. Sariali, E., Mouttet, A., Pasquier, G., & Durante, E. (2009). Accuracy of reconstruction of the hip using computerised three-dimensional pre-operative planning and a cementless modular neck. Journal of Bone and Joint Surgery. British Volume, 91(3), 333–340. https://doi.org/10.1302/0301-620X.91B3.20937

14. Maintz, J. B. A., & Viergever, M. A. (1998). A survey of medical image registration. Medical Image Analysis, 2(1), 1–36. https://doi.org/10.1016/S1361-8415(01)80026-8

15. Chui, H., & Rangarajan, A. (2003). A new point matching algorithm for non-rigid registration. Computer Vision and Image Understanding, 89(2–3), 114–141. https://doi.org/10.1016/S1077-3142(03)00009-2

16. Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian approximation: Representing model uncertainty in deep learning. Proceedings of the 33rd International Conference on Machine Learning, 48, 1050–1059.

17. Kendall, A., & Gal, Y. (2017). What uncertainties do we need in Bayesian deep learning for computer vision? Advances in Neural Information Processing Systems, 30, 5574–5584.

18. Huiskes, R., Weinans, H., & van Rietbergen, B. (1992). The relationship between stress shielding and bone resorption around total hip stems and the effects of flexible materials. Clinical Orthopaedics and Related Research, 274, 124–134.

19. Noble, P. C., Alexander, J. W., Lindahl, L. J., Yew, D. T., Granberry, W. M., & Tullos, H. S. (1988). The anatomic basis of femoral component design. Clinical Orthopaedics and Related Research, 235, 148–165.

20. Rengier, F., Mehndiratta, A., von Tengg-Kobligk, H., Zechmann, C. M., Unterhinninghofen, R., Kauczor, H.-U., & Giesel, F. L. (2010). 3D printing based on imaging data: Review of medical applications. International Journal of Computer Assisted Radiology and Surgery, 5(4), 335–341. https://doi.org/10.1007/s11548-010-0476-x