A Rigid-Body Target Design Methodology for Optical Pose-Tracking Systems

Thomas Pintaric* Hannes Kaufmann†

Interactive Media Systems Group
Institute of Software Technology and Interactive Systems
Vienna University of Technology

Abstract

The standard method for estimating the rigid-body motion of arbitrary interaction devices with an infrared-optical tracking system involves attaching pre-defined geometric constellations of retro-reflective or light-emitting markers, commonly referred to as "targets", to all tracked objects. Optical markers of the same type are typically indistinguishable from each other, requiring the tracking system to establish their identities through known spatial relationships. Consequently, the specific geometric arrangement of markers across multiple targets has a considerable impact on the system's overall performance and robustness.

In this paper, we propose a simple new methodology for constructing optically tracked rigid-body targets. Our practically-oriented approach employs an optimization heuristic to compute near-optimal marker arrangements. Using prefabricated mounting fixtures, the assembly step requires only basic hobbyist tools and skills.

CR Categories: H.5.2 [Multimedia Information Systems]: User Interfaces—Input devices and strategies G.1.6 [Numerical Analysis]: Optimization—Constrained optimization I.5.2 [Pattern Recognition]: Design Methodology—Classifier design and evaluation

Keywords: model-based object tracking, rigid body, marker constellation, target design, 6-DOF pose-tracking, optical tracking

1 Introduction

Optical motion-tracking technologies are generally appreciated for their ability to cover room-sized working volumes while delivering measurements with a relatively high spatial accuracy, typically in the sub-millimeter range. A prominent subset of existing commercial and non-commercial optical tracking systems require objects to be instrumented with geometric arrangements of (three or more) optical markers before they can be tracked. Such rigid constellations of optical markers are customarily referred to as "targets", an example of which is shown in Figure 1. Commonly used optical markers include light-emitting diodes ("active markers") and small plastic spheres with a retro-reflective coating ("passive markers"). Markers of the same type are typically indistinguishable to the tracking system, hence requiring a model-fitting algorithm to establish their identities through known spatial relationships within a target.

*e-mail: pintaric@ims.tuwien.ac.at
†e-mail: kaufmann@ims.tuwien.ac.at

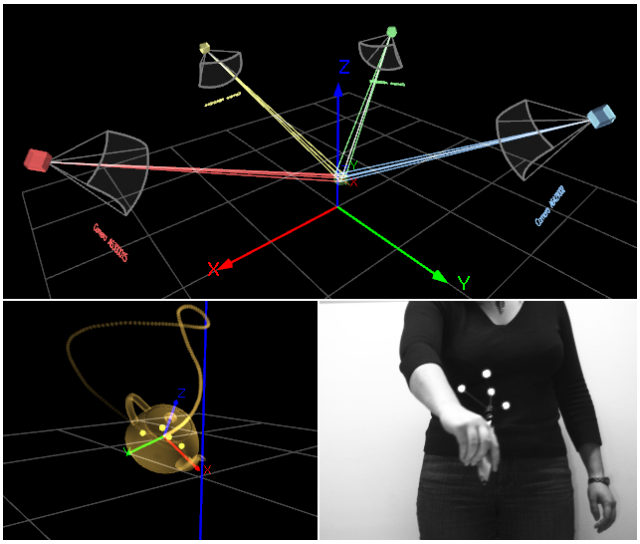

Figure 1: *Optically tracked rigid-body targets are typically made from geometric arrangements of retro-reflective spherical markers.*

Consequently, the chosen geometric arrangement of optical markers will directly affect the system's ability to recognize and distinguish between targets. Whenever multiple rigid-body targets (or parts thereof) exhibit a high degree of "similarity" (see Section 2.1 for a mathematical definition), the system may be unable to correctly label those targets, with effects ranging from incorrect pose estimates to a complete loss of tracking. In fact, the correct labeling of optical markers is generally considered a prerequisite for obtaining usable pose estimates from an optical motion-tracking system.

Yet, there is little published work on the problem of rigid-body target construction. Exceptions include [Jansen et al. 2007], who emphasize theoretical over practical aspects, and [Davis et al. 2004], whose "Viewpoint Algorithm" requires the use of a three-dimensional CAD model of the instrumented object. Both techniques focus on a single rigid-body target. To the best of our knowledge, none of the published approaches consider the joint use of multiple targets.

Unfortunately, commercial manufacturers tend to refrain from publicly disclosing their design procedures. As a result, end users who are unwilling or unable to purchase pre-assembled rigid-body targets from commercial vendors are forced to rely on a mixture of intuition and experience when building their customized targets.

In the remainder of this paper, we propose a new practically-oriented approach for constructing optically tracked rigid-body targets. We provide the blueprint of a plastic base fixture that can be used to arrange multiple markers in a wide range of geometric patterns using only hobbyist tools and skills. A configurable heuristic optimization algorithm is used to calculate near-optimal marker-arrangements for different use-cases and numbers of targets, taking into account problem-specific constraints such as occlusions and the shape of the instrumented interaction device.

2 Target Design

2.1 Defining Optimality

At the model-fitting stage of the data processing pipeline, optical pose-tracking systems attempt to uniquely identify the markers of each pre-calibrated target ("model") within a larger, unstructured set of observed markers ("observation"). This is usually achieved by comparing the Euclidean distances between all pairs of observed markers with the known distances between markers of each target. To account for the presence of noise and different types of process errors, two distances d_i and d_j are considered equal if $|d_i - d_j| < \varepsilon$.

Representative examples of model-fitting algorithms can be found in [Dorfmüller-Ulhaas 2002], [Ribo et al. 2001] or [Pintaric and Kaufmann 2007]. It can be stated that, regardless of which specific model-fitting algorithm is used, the common goal of a good target design is to minimize the number of distance-matching ambiguities during the model-fitting stage.

Fortunately, it is relatively straight-forward to devise a mathematical description of "target optimality". In essence, one needs to identify the specific marker-arrangements for which the performance of the model-fitting algorithm is optimal, which we formulate as follows:

Let D_i be the set of pair-wise Euclidean distances between all markers of the i^{th} tracked rigid-body target. Now let the vector u contain all elements of every D_i for $i = 1..n$, where n denotes the total number of targets. We then construct a new set $S = \bigcup_{j \neq k} |u_j - u_k|$, which contains the pair-wise differences of all elements in u. (Note that u_j and u_k denote the j^{th} and k^{th} component of the vector u.) The goal of our optimization algorithm is to maximize $min(S)$. In other words, the smallest difference between every pair of marker-to-marker distances across multiple targets should be maximal.

Is is not hard to notice that this is exactly the case when all marker-distances are equi-distantly spaced. The harder task is to compute the location of markers in such a way, which can be modeled as a combinatorial optimization problem.

2.2 Optimization Procedure

In theory, we could model the problem of computing an optimal arrangement of m optical markers, belonging to n different rigid-body targets, as a global optimization in $(3m)$ dimensions, where every parameter corresponds to a specific marker's translation along one coordinate axis (X,Y or Z) in $\mathbb{R}$. We could then attempt to construct a target according to the optimal solution. However, such an approach is not feasible for the following reasons:

- It is hard to construct a customized support structure to hold an arbitrary spatial arrangement of markers in place.
- The accuracy with which a person can affix an optical marker to some carrier objects is limited, and in most cases relatively low (i.e. at best 1 mm).
- Global optimization is hard, and it takes significant experience to tune an algorithm to perform well on a specific problem.

Instead, we opted to devise a simplified construction procedure first, and then formulate an optimization algorithm for it.

2.2.1 Practical Aspects

Using only hobbyist tools, we found it much harder to measure an object's angle of inclination than its length. Thus, we designed a CNC-machined polycarbonate mounting fixture (see Figure 2) with 25 pre-drilled receptor holes at equi-distant angles between 0° (vertical) and 90° (horizontal), which have an angular accuracy of $\pm 0.1°$. The holes are intended to hold thin carbon-fiber rods ($\varnothing$1.5 mm), to which retro-reflective optical markers can be attached (see Figure 3). This enables designers to freely arrange markers within a hemispherical region centered around the mounting fixture by trimming each carrier-rod to the desired length and inserting it into one of the pre-drilled holes.

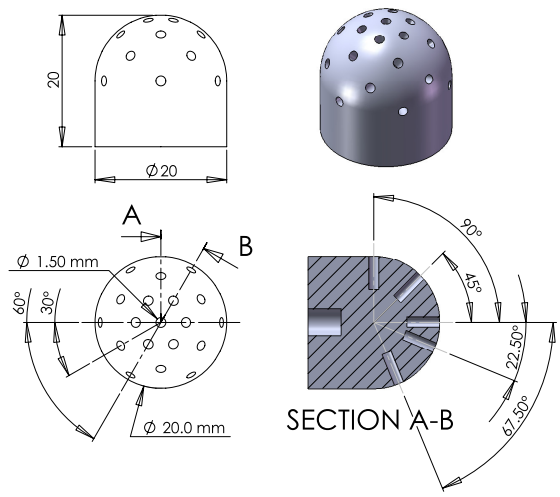

(a) *Technical drawing with dimensions and cross-section view.*

(b) *CNC-machined sample from the initial production run.*

Figure 2: *Mounting fixture for optically tracked rigid-body targets.*

2.2.2 Combinatorial Optimization by Random Local Search

We can now reduce the optimization problem to a discrete combinatorial optimization in $(2m)$ dimensions (for m markers across all n rigid-body targets), where every marker is assigned two parameters:

- The length of its carrier rod in integer multiples of 1.0 mm.
- The hole number into which its carrier rod is inserted, which is an integer $\in \{1, 2, \ldots, 25\}$.

We chose a randomized local search (RLS) algorithm to compute optimal marker-arrangements for a given number of target mounting fixtures. RLS is a well-known heuristic for solving optimization problems by starting at a randomly chosen point in the solution space, then iteratively moving towards better solutions by examining its immediate neighborhood region. Three concepts must be defined: How to generate the initial random starting point, the neighborhood operator and an objective function that defines a mapping from solution space to $\mathbb{R}$.

(a) *We are using a Dremel rotory tool with an attached diamond cut-off disc to cleanly trim our solid carbon-fiber rods. All rod lengths are verified with a sliding caliper to within ±0.5 mm.*

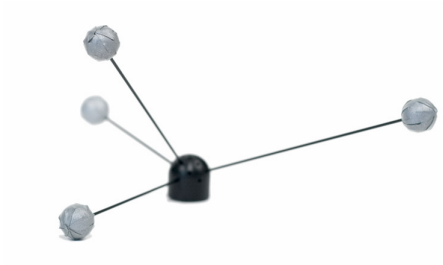

(b) *Assembled four-marker rigid-body target.*

Figure 3: *Rigid-body target assembly procedure.*

Random starting point: For each rigid-body target, three or more holes are selected at random (depending on the desired number of markers per target) and populated with randomly-sized rods from a user-specified length interval $[l_{min}, l_{max}]$. We commonly choose $l_{min} = 50\,\text{mm}$ and $l_{max} = 120\,\text{mm}$. No further constraints are imposed at this time.

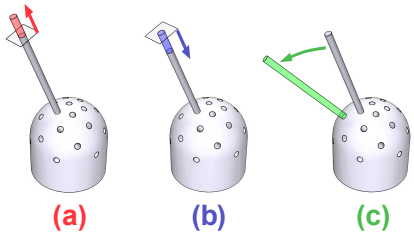

Figure 4: *Local neighborhood operator.*

Neighborhood operator: We define three basic operations (see Figure 4) to generate the local neighborhood region of a given solution:

(a) Lengthen a rod by a specific amount (l_{step}).
(b) Shorten a rod by a specific amount (l_{step}).
(c) Move a rod to an adjacent unoccupied hole.

Initially, we set l_{step} to a value between 10 mm and 15 mm. Once we reach a point during the optimization at which no solution in the local neightborhood improves upon the current solution's objective function score, we keep regenerating the neighborhood region with a reduced step width ($l_{step} = l_{step} - 1\,\text{mm}$). The optimization is stopped when l_{step} reaches zero.

Objective function: Our formulation of the objective function f_{obj} (see below) was motivated by an attempt to strike a balance between optimizing the worst-case (represented by d_{min}) and average-case (represented by d_{avg}) target quality. It is not hard to see that maximizing only d_{min} would also implicitly maximize d_{avg}. Unfortunately, the resulting objective function would turn out "too flat" to generate any meaningful results with a gradient descent method. We thus chose to optimize a linear combination of d_{min} and d_{avg}. The goal of the optimization is to maximize f_{obj}.

$$f_{obj} = \begin{cases} \alpha * d_{min} + (1-\alpha) * d_{avg} & \text{if} \quad f_{con} = 0, \\ -f_{con} & \text{if} \quad f_{con} > 0. \end{cases}$$

f_{obj} ... objective function to be maximized
f_{con} ... constraint penalty function (see below)
d_{min} ... minimal difference of pair-wise marker distances
d_{avg} ... average difference of pair-wise marker distances

$\alpha = \frac{2}{3}$ (obtained experimentally).

Constraints: In order to be able to enforce certain types of user-specified constraints, our definition of the objective function includes a separate "constraint penalty" function f_{con}. If no constraints are violated, f_{con} evaluates to zero. Otherwise, f_{con} returns a positive value that is proportional to the "severity" of the violation. (The exact formulation of f_{con} is of minor importance for the outcome of the optimization and will not be discussed in further detail.) So far, we have implemented the following constraint types:

- Minimal and maximal rod length (l_{min}, l_{max}). Can be defined separately for each hole.
- Minimal marker-to-marker distance (within a target).
- Anti-collinearity constraint (see Figure 5). Intended to prevent the optimization algorithm from placing three or more markers on the same (imaginary) line. This will mitigate occlusion problems and improve the stability of the tracking system's orientation estimator.

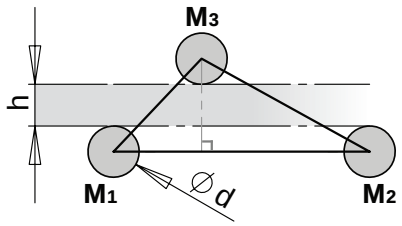

Figure 5: *Anti-collinearity constraint: $h > d$.*

3 Results

For testing purposes, we ran our RLS algorithm on seven representative problem sets, each containing between 2 and 8 rigid-body targets. Apart from the number of targets, identical constraints were applied to each optimization problem (see Table 1).

We computed between 500 and 1500 solutions for each of the seven optimization problems and examined the 90^{th} percentile range of results. Figure 6 shows a plot of d_{avg} (a good indicator of the "overall quality" of rigid-body targets) against the increasing problem size.

Markers per target:	4
Minimal rod length (l_{min}):	50 mm
Maximal rod length (l_{max}):	120 mm
Min. marker-to-marker distance:	40 mm
Anti-collinearity constraint:	$h > 14$ mm

Table 1: *Target optimization constraints.*

To assess the effectiveness of our approach, we analytically approximated the upper bound for d_{avg} by dividing a target's theoretical span width (i.e. twice the maximal rod length) by the total number of pair-wise marker combinations. Our mechanical target design (we can only place markers along fixed lines, instead of freely positioning them in space) precludes our algorithm from ever reaching this upper bound. Yet, the 90^{th} percentile of d_{avg} falls within 75% to 85% of this upper bound. We consider these results very encouraging.

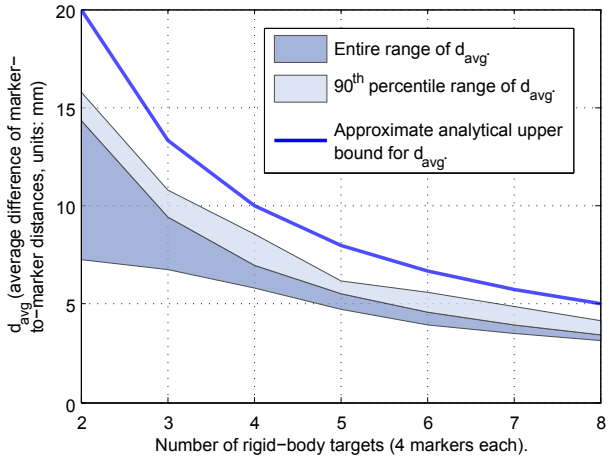

Figure 6: *Optimization results.*

To assess the practical applicability of our results, we assembled several optimized rigid-body target sets (one of which is shown in Figure 8) and ran repeated 30-minute test sessions with a collaborative Virtual Reality CAD application [Kaufmann and Schmalstieg 2006]. In our test scenario, multiple users were equipped with up to three tracked input/output devices (as shown in Figure 7), and were free to move around in a room-sized (3x3x2m) interaction volume covered by an ***iotracker*** infrared-optical tracking system [Pintaric and Kaufmann 2007].

Preliminary findings from these tests indicate significant improvements in tracking stability (fewer mislabeled targets, increased robustness against partial occlusions) compared to "homebrew" rigid-body targets. Future work will include a more thorough quantitative analysis of these improvements.

Acknowledgements

This research was funded in part by the EU (IST project FP7-215839) and the Austrian Science Fund (FWF contract P19265).

References

DAVIS, L., HAMZA-LUP, F. G., AND ROLLAND, J. P. 2004. A method for designing marker-based tracking probes. In *Proceedings of the 3th IEEE/ACM International Symposium on Mixed and Augmented Reality (ISMAR'04)*, IEEE Computer Society, Washington, DC, USA, 120–129.

DORFMÜLLER-ULHAAS, K. 2002. *Optical Tracking - From User Motion To 3D Interaction*. PhD thesis, Institute of Computer Graphics and Algorithms, Vienna University of Technology, Favoritenstrasse 9-11/186, A-1040 Vienna, Austria.

JANSEN, C., STEINICKE, F., HINRICHS, K. H., VAHRENHOLD, J., AND SCHWALD, B. 2007. Performance improvement for optical tracking by adapting marker arrangements. In *IEEE VR 2007 Workshop on Trends and Issues in Tracking for Virtual Environments*, Shaker-Verlag, G. Zachmann, Ed., 28–33.

KAUFMANN, H., AND SCHMALSTIEG, D. 2006. Designing Immersive Virtual Reality for Geometry Education. In *Proceedings of the IEEE conference on Virtual Reality (VR '06)*, IEEE Computer Society, Washington, DC, USA, 51–58.

PINTARIC, T., AND KAUFMANN, H. 2007. Affordable infrared-optical pose-tracking for virtual and augmented reality. In *Proceedings of Trends and Issues in Tracking for Virtual Environments Workshop, IEEE VR 2007*.

RIBO, M., PINZ, A., AND FUHRMANN, A. L. 2001. A new Optical Tracking System for Virtual and Augmented Reality Applications. In *Proceedings of IEEE Instrumentation and Measurement Technology Conference, IMTC 2001*, vol. 3, 1932–1936.

Figure 7: *Typical use-cases of our method include VR applications that require the use of multiple tracked input/ouput devices per user.*

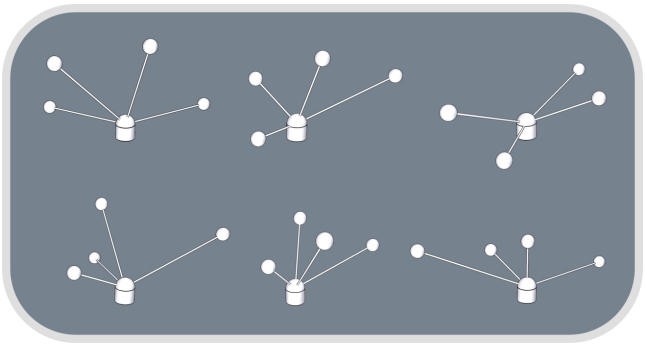

Figure 8: *Set of six rigid-body targets (after optimization).*