

Enhancing motion trajectory segmentation of rigid bodies using a novel screw-based trajectory-shape representation

Arno Verduyn^{1,2}, Maxim Vochten¹, Joris De Schutter¹

Abstract—Trajectory segmentation refers to dividing a trajectory into meaningful consecutive sub-trajectories. This paper focuses on trajectory segmentation for 3D rigid-body motions. Most segmentation approaches in the literature represent the body’s trajectory as a point trajectory, considering only its translation and neglecting its rotation. We propose a novel trajectory representation for rigid-body motions that incorporates both translation and rotation, and additionally exhibits several invariant properties. This representation consists of a geometric progress rate and a third-order trajectory-shape descriptor. Concepts from screw theory were used to make this representation time-invariant and also invariant to the choice of body reference point. This new representation is validated for a self-supervised segmentation approach, both in simulation and using real recordings of human-demonstrated pouring motions. The results show a more robust detection of consecutive sub-motions with distinct features and a more consistent segmentation compared to conventional representations. We believe that other existing segmentation methods may benefit from using this trajectory representation to improve their invariance.

I. INTRODUCTION

Trajectory segmentation aims to divide a trajectory into consecutive sub-trajectories that have a specific meaning for the considered application. It is a valuable tool to reduce the dimensionality of trajectories, thereby improving the efficiency and reliability of trajectory recognition and learning algorithms. Trajectory segmentation has found applications in various research fields, including monitoring urban traffic [1], tracking changes in animal behavior [2], analyzing robot-assisted minimally invasive surgical procedures [3], and learning motion primitives for robots from human-demonstrated object manipulation tasks [4]–[7].

In the literature, three types of segmentation approaches can be distinguished: supervised segmentation, unsupervised segmentation, and self-supervised segmentation.

Supervised segmentation algorithms rely on prior expert knowledge, such as a predefined library of template segments, such as in [8], or a set of predefined segmentation criteria. The approaches in [1], [9] segment trajectories based on criteria that assess the homogeneity of movement profiles, including the location, speed and heading of the object.

Unsupervised segmentation algorithms do not rely on prior expert knowledge. These approaches are typically event-based, searching for discrete segmentation points along the trajectory. The approach in [5] identifies segmentation points

by fitting a Gaussian Mixture Model to the trajectory data and identifying regions of overlap between the confidence ellipsoids of consecutive Gaussians. Another approach involves extracting segmentation points from local extrema in curvature and torsion along the trajectory, such as in [6].

Self-supervised segmentation can be viewed as a supervised approach in which the segmentation criteria are learned from the data rather than being predetermined by expert knowledge. The segmentation approach in [10] detects transitions in movement profiles and labels them as ‘features’. Similar features are then clustered together, and subsequently, the trajectory is segmented into a sequence of cluster labels. The approach in [7] initially segments trajectories in an unsupervised way. It then learns a library of motion primitives from the segmented data, which is subsequently used to improve the initial segmentation accuracy.

The literature on trajectory segmentation has two main shortcomings. The first shortcoming is that the object is typically approximated as a moving point, considering only the translational motion while neglecting the rotational motion. Hence, not all trajectory information is taken into account.

The second shortcoming is that supervised and self-supervised segmentation approaches based on template matching typically are dependent on the motion profile of the trajectory and the choice of references. These references include both the coordinate frame in which the trajectory coordinates are expressed and the reference point on the body being tracked. This dependency limits the approach’s capability to generalize across different setups.

The main objective of this work is to enhance template-based approaches in supervised and self-supervised trajectory segmentation by incorporating invariance with respect to time, coordinate frame, and body reference point.

Our approach consists of first reparameterizing the trajectory to achieve time-invariance, i.e. invariance to changes in motion profile. This is achieved by defining a novel geometric progress rate using *screw theory*, inspired by [11]. The progress rate combines both rotation and translation, while also being invariant to the chosen body reference point. The defined progress rate is regularized to better cope with pure translations compared to [11]. After reparameterization, the trajectory is represented using a novel trajectory-shape descriptor based on first-order kinematics of the trajectory. Finally, invariance to changes in the coordinate frame is achieved by performing a spatial alignment before matching the trajectory descriptor with template segments in the library. The latter template segments are also represented using the same trajectory-shape descriptor.

This result is part of a project that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 788298).

¹Department of Mechanical Engineering and Flanders Make at KU Leuven, 3001 Leuven, Belgium.

²Corresponding author (arno.verduyn@kuleuven.be)

The contribution of this work is (1) the introduction of a novel screw-based geometric progress rate for rigid-body trajectories, and (2) a novel trajectory-shape descriptor. These concepts were applied in a self-supervised segmentation approach using simulations and real recordings of human-demonstrated pouring motions. The results show that more consistent segmentation results can be achieved by using an invariant trajectory descriptor compared to conventional descriptors that are not invariant to changes in execution speed, coordinate frame, and reference point.

The outline of the paper is as follows. Section II reviews essential background on rigid-body trajectories and screw theory. Section III introduces the novel screw-based geometric progress rate and trajectory-shape descriptor for rigid-body trajectories. Section IV explains the implementation of these novel concepts in a self-supervised segmentation approach. Section V explains the validation of the segmentation approach using simulations and real recordings of human-demonstrated pouring motions. Section VI concludes with a discussion and suggests future work.

II. PRELIMINARIES

The displacement of a rigid body in 3D space is commonly represented by attaching a body frame to the rigid body and expressing the position \mathbf{p} and orientation R of this frame with respect to a static world frame. The position coordinates \mathbf{p} represent the relative position of the origin of the body frame with respect to the origin of the world frame. The rotation matrix R , which is part of the Special Orthogonal group $SO(3)$, represents the relative orientation of the body frame with respect to the world frame¹. The position \mathbf{p} and orientation R of the rigid body can be combined into the homogeneous transformation matrix $T = \begin{bmatrix} R & \mathbf{p} \\ \mathbf{0} & 1 \end{bmatrix}$, which is part of the Special Euclidean group $SE(3)$. The homogeneous transformation matrix as a function of time $T(t)$ represents a temporal rigid-body trajectory.

The first-order kinematics of the rigid-body trajectory $T(t)$ are commonly represented by a 6D twist \mathfrak{t} consisting of two 3D vectors: a rotational velocity vector $\boldsymbol{\omega}$ and a linear velocity vector \mathbf{v} . Based on the coordinate frame in which these velocities are expressed and based on the reference point for the linear velocity \mathbf{v} , three twists are commonly defined in the literature [12], [13], i.e. the *pose twist*, *spatial twist*, and *body twist*. A twist is considered left- or right-invariant when it is invariant to changes of the world or body frame, respectively. The *body twist* is left-invariant, the *spatial twist* is right-invariant, and the *pose twist* is neither left- nor right-invariant.

In rigid-body kinematics, the Mozzi-Chasles' theorem [14], [15] states that a rigid-body velocity can always be represented as a rotation about an axis in space and a translation parallel to this axis, referred to as the *Instantaneous Screw Axis* (ISA). The direction of the ISA is uniquely defined by the direction of the rotational velocity $\boldsymbol{\omega}$,

¹Note that other possibilities exist for representing the orientation such as the quaternion or the axis-angle representation.

while the location of the ISA is defined by any point on the ISA. A unique choice is to take the point on the ISA that lies closest to the reference point for the linear velocity \mathbf{v} .

The translational velocity $\boldsymbol{\nu}$ parallel to the ISA, and the position vector \mathbf{p}_\perp of the point on the ISA that lies closest to the body frame's origin are calculated from the pose twist components $\boldsymbol{\omega}$ and \mathbf{v} as follows:

$$\mathbf{p}_\perp = \frac{\boldsymbol{\omega} \times \mathbf{v}}{\|\boldsymbol{\omega}\|^2} \quad \text{and} \quad \boldsymbol{\nu} = \mathbf{v} + \boldsymbol{\omega} \times \mathbf{p}_\perp. \quad (1)$$

The magnitudes of the rotational and translational velocities $\|\boldsymbol{\omega}\|$ and $\|\boldsymbol{\nu}\|$ are closely related to the 'SE(3) invariants' in [16] and it has been shown that they are both left- and right-invariant, also referred to as *bi-invariant*.

III. SCREW-BASED REPRESENTATION

This section introduces a novel screw-based geometric progress rate for defining the progress over a trajectory, independent of time. Next, a new invariant trajectory-shape descriptor is introduced for rigid-body trajectories, which is both time-invariant and invariant to the choice of the reference point on the body. Invariance to changes in coordinate frame is obtained by proposing a spatial alignment algorithm.

A. Screw-based geometric progress rate

To obtain a time-invariant representation of the trajectory, a geometric progress rate has to be defined first. We propose to combine the magnitude of the rotational velocity $\|\boldsymbol{\omega}\|$ and the translational velocity $\|\boldsymbol{\nu}\|$ parallel to the ISA:

$$\dot{s} = \sqrt{L^2 \|\boldsymbol{\omega}\|^2 + \|\boldsymbol{\nu}\|^2}, \quad (2)$$

where L is a weighting factor with units [m]. The progress rate \dot{s} [m/s] can be interpreted as the magnitude of the linear velocity of any point on the body, at a distance L from the ISA (see Fig. 1). The progress parameter s [m] is then found as the integral of the progress rate over time, and can be considered as a scalar value signifying the geometric progress in translation and rotation over the trajectory.

Since $\|\boldsymbol{\omega}\|$ and $\|\boldsymbol{\nu}\|$ are bi-invariant properties of the trajectory, as mentioned in Section II, the resulting geometric progress rate \dot{s} is also bi-invariant. However, it is important to clarify that the proposed progress rate does not comply with the conditions for being a bi-invariant metric on $se(3)$ [17]. For example, consider a serial kinematic chain of two twist motions \mathfrak{t}_1 and \mathfrak{t}_2 generating a resulting motion \mathfrak{t}_{1+2} . Then, the triangle inequality $\dot{s}_{1+2} \leq \dot{s}_1 + \dot{s}_2$ (a necessary condition for being a metric) does not hold for all cases on $se(3)$. Consider the example of a pure translation $\dot{s}_{1+2} = \|\boldsymbol{\nu}\|$. This translation can be generated by a couple of rotations with magnitude $\|\boldsymbol{\omega}\| = \frac{\|\boldsymbol{\nu}\|}{2a}$, where a is half the distance between the two rotation axes, as shown in Figure 2. Since the sum of the progress rates of these two rotations equals:

$$\dot{s}_1 + \dot{s}_2 = 2L\|\boldsymbol{\omega}\| = \frac{L}{a}\|\boldsymbol{\nu}\|, \quad (3)$$

it can be concluded that $\dot{s}_1 + \dot{s}_2 < \dot{s}_{1+2}$, if $L < a$. Since L is a predefined weighting factor and a can take any value in \mathbb{R}_0^+ (because $\boldsymbol{\omega}$ can always be decreased by increasing a

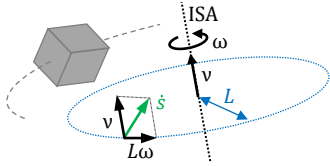


Fig. 1. Interpretation of the proposed progress rate \dot{s} as the linear velocity of a point on the moving body at a distance L from the ISA.

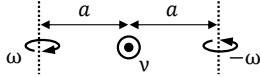


Fig. 2. Illustration of a pure couple of rotations ω generating a pure translation ν out of the plane. The two rotation axes, distanced by two times a , are depicted with dotted lines.

while still resulting in the same ν), there is no choice for L for which the triangle inequality holds for each value of a .

B. Regularized geometric progress rate

The calculation of \mathbf{p}_\perp using (1) is degenerate for pure translations, i.e. when $\|\omega\| = 0$. Because of this, approaching a pure translation can result in $\|\mathbf{p}_\perp\|$ approaching infinity:

$$\text{if } \lim_{\|\omega\| \rightarrow 0} \frac{\omega \times \mathbf{v}}{\|\omega\|} \neq \mathbf{0}, \text{ then } \lim_{\|\omega\| \rightarrow 0} \|\mathbf{p}_\perp\| = \infty. \quad (4)$$

To ensure that the calculated translational velocity of the rigid body and corresponding progress rate remain well-defined in these degenerate cases, a regularized translational velocity $\tilde{\mathbf{v}}$ is introduced, such that:

$$\dot{s} = \sqrt{L^2 \|\omega\|^2 + \|\tilde{\mathbf{v}}\|^2} \quad \text{with} \quad \tilde{\mathbf{v}} = \mathbf{v} + \omega \times \tilde{\mathbf{p}}, \quad (5)$$

where $\tilde{\mathbf{p}}$ is a regularized version of \mathbf{p}_\perp , defined as follows:

$$\tilde{\mathbf{p}} = \begin{cases} \mathbf{p}_\perp & \text{if } \|\mathbf{p}_\perp\| \leq b \\ b \frac{\mathbf{p}_\perp}{\|\mathbf{p}_\perp\|} & \text{if } \|\mathbf{p}_\perp\| > b \\ \mathbf{0} & \text{if } \|\omega\| = 0. \end{cases} \quad (6)$$

In other words, a sphere with radius b is defined with the origin of the body frame as its center, so that when the position \mathbf{p}_\perp is outside of the sphere, it will be projected onto the sphere's surface. This is also visualized in Fig. 3.

The regularization in (6) has two important properties. First, the profile of $\tilde{\mathbf{v}}$ remains continuous over the threshold where $\|\mathbf{p}_\perp\| = b$, since:

$$\lim_{\|\mathbf{p}_\perp\| \rightarrow b^+} \tilde{\mathbf{p}} = \lim_{\|\mathbf{p}_\perp\| \rightarrow b^-} \tilde{\mathbf{p}}. \quad (7)$$

Second, the regularization is independent of the execution speed of the motion, since \mathbf{p}_\perp is a geometric property.

In the Appendix, it is shown that choosing $b \leq L$ ensures that the triangle inequality $\dot{s}_{1+2} \leq \dot{s}_1 + \dot{s}_2$ holds within the bi-invariant region $\|\mathbf{p}_\perp\| \leq b$, for the example discussed in Section III-A. As a result, we propose to choose $b = L$, since this results in the largest region of bi-invariance of the progress rate \dot{s} without violating the triangle inequality.

Remark that when $\|\mathbf{p}_\perp\| > b$, the regularized translational velocity $\tilde{\mathbf{v}}$ becomes dependent on the location of the body

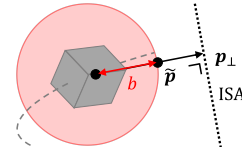


Fig. 3. Point $\tilde{\mathbf{p}}$ is defined as a characteristic reference point on the body. When \mathbf{p}_\perp is within the spherical region with radius b , then $\tilde{\mathbf{p}} = \mathbf{p}_\perp$. Outside the sphere, $\tilde{\mathbf{p}}$ is the projection of \mathbf{p}_\perp on the sphere's surface.

reference point and hence loses its *bi-invariant* property. Fortunately, since the object is dominantly translating in this case, the sensitivity of the velocity $\tilde{\mathbf{v}}$ to the location of the body reference point will remain limited.

Algorithm 1 contains pseudocode to calculate the regularized progress rate \dot{s} assuming $b = L$.

Algorithm 1: Calculation regularized progress rate \dot{s}

Data: $\omega \in \mathbb{R}^3, \mathbf{v} \in \mathbb{R}^3, L \in \mathbb{R}$
Result: $\dot{s} \in \mathbb{R}, \tilde{\mathbf{p}} \in \mathbb{R}^3, \tilde{\mathbf{v}} \in \mathbb{R}^3$
if $\|\omega\| = 0$ **then**
 $\tilde{\mathbf{p}} \leftarrow \mathbf{0}; \tilde{\mathbf{v}} \leftarrow \mathbf{v};$
else
 $\tilde{\mathbf{p}} \leftarrow (\omega \times \mathbf{v}) / (\omega \cdot \omega);$
 if $\|\tilde{\mathbf{p}}\| > L$ **then** $\tilde{\mathbf{p}} \leftarrow L\tilde{\mathbf{p}} / \|\tilde{\mathbf{p}}\|;$
 $\tilde{\mathbf{v}} \leftarrow \mathbf{v} + \omega \times \tilde{\mathbf{p}};$
end
 $\dot{s} \leftarrow \sqrt{L^2 \omega \cdot \omega + \tilde{\mathbf{v}} \cdot \tilde{\mathbf{v}}};$

C. Reparameterization to geometric domain

The temporal rigid-body trajectory $T(t)$ can now be reparameterized to a geometric trajectory $T(s) = T(t(s))$ using the geometric progress rate \dot{s} defined in (5). As a result, the execution speed will be separated from the geometric path of the trajectory. Such a geometric trajectory $T(s)$ is also referred to as a *unit-speed* trajectory, since the geometric derivative s' ($= \frac{ds}{ds}$) of the progress s along this trajectory is equal to one:

$$s' = 1 = \sqrt{L^2 \|\omega(s)\|^2 + \|\tilde{\mathbf{v}}(s)\|^2}. \quad (8)$$

In practice, for discrete measurement data, this reparameterization is performed in three steps. Firstly, the temporal pose twist \mathcal{E} is calculated from the temporal trajectory $T(t)$ by numerical differentiation with the matrix logarithm operator [12]. Secondly, the progress rate \dot{s} is determined using Algorithm 1. The progress s along the discrete trajectory is then found by a cumulative sum on \dot{s} . Thirdly, the trajectory T is reparameterized from time t to progress s with fixed progress step Δs using *Screw Linear Interpolation* [18].

D. Screw-based trajectory-shape descriptor

This subsection introduces a screw-based trajectory-shape descriptor based on the rotational and translational velocities $\omega(s)$ and $\tilde{\mathbf{v}}(s)$ along the reparameterized trajectory. These velocities can be calculated similarly to the calculation of $\omega(t)$ and $\tilde{\mathbf{v}}(t)$, but now starting from the reparameterized trajectory $T(s)$ instead of $T(t)$. Afterwards, a normalization step is included to ensure that property (8) holds.

The proposed trajectory-shape descriptor at a given trajectory sample i consists of the rotational and translational velocities at three successive sample points, centered at the sample point i , and stacked into a 3×6 matrix S_i :

$$S_i = [L\omega_{i-1} \quad L\omega_i \quad L\omega_{i+1} \quad \tilde{\nu}_{i-1} \quad \tilde{\nu}_i \quad \tilde{\nu}_{i+1}]. \quad (9)$$

By using three successive sample points (instead of only the center one), the descriptor S_i contains information of the magnitudes and local change in direction of ω_i and $\tilde{\nu}_i$.

The velocities in (9) have their coordinates expressed with respect to the static world frame. Hence, to be able to compare two local trajectory-shape descriptors in an invariant way, a local rotation alignment has to be performed. Similarly to the rotation alignment proposed in [19], shape descriptor S_2 can be aligned with S_1 in three steps:

- 1) Obtain the singular value decomposition of the relative 3×3 matrix $S_1 S_2^T = U \Sigma V$.
- 2) Calculate the rotation matrix $R = V U^T$, while ensuring that $R \in SO(3)$ by changing the sign of the third column of U when $\det(V U^T) = -1$.
- 3) Align S_1 with S_2 by left-multiplying S_1 with R .

The difference $\Delta_1^2 S$ between two descriptors S_1 and S_2 is then defined as the Frobenius norm of the difference between the descriptors after alignment:

$$\Delta_1^2 S = \|R S_1 - S_2\|_F. \quad (10)$$

This difference $\Delta_1^2 S$ is a value with units [m], measuring the difference in the local trajectory shape of two trajectories.

IV. APPLICATION TO SEGMENTATION

This section applies the proposed screw-based progress rate \dot{s} and trajectory-shape descriptor S to a trajectory segmentation approach. Envisioned towards incremental learning applications, a self-supervised segmentation approach based on incremental clustering was devised. This segmentation approach consists of two phases: an offline template learning phase where a library of trajectory-shape primitives is learned, and a trajectory segmentation phase.

Offline template learning: Trajectory-shape primitives are learned by an incremental clustering approach, where the mean of each cluster represents a learned primitive. The clustering approach consists of four steps:

1) *Initialization:* Use the first sample point S_1 to generate the first cluster with mean S_1 and a chosen initial guess for its standard deviation σ_0 .

2) *Cluster growing:* Given a sample point i with corresponding trajectory-shape descriptor S_i , calculate the difference ΔS between the descriptor S_i and the mean \bar{S} of every learned cluster. Afterwards, add the sample point to the cluster with the smallest difference in trajectory-shape, if this difference is smaller than three times its standard deviation σ , and update the mean \bar{S} and standard deviation σ of the cluster accordingly. If no such cluster exists, create a new cluster with mean S_i and initial standard deviation σ_0 .

3) *Cluster parameter update:* After all the available data is clustered, update the value for σ_0 based on the mean

of the standard deviations of all the learned clusters $\bar{\sigma}$: $\sigma_{0,next} = \bar{\sigma} + \hat{\sigma}$, with $\hat{\sigma}$ a tuning parameter related to the process noise of the clusters. Iterate steps 1 to 3 until σ_0 converges to a steady value.

4) *Outlier removal:* After convergence, remove sparse clusters that do not represent at least $\beta\%$ of the data, with β being a chosen value.

Trajectory segmentation phase: Each sample point gets associated with a learned primitive (using a 1-Nearest-Neighbor classifier) as long as the difference in trajectory-shape ΔS is smaller than 3σ of the respective cluster. Otherwise, the sample point is labeled as ‘non-classified’. Afterwards, segments are formed by grouping consecutive sample points along the trajectory that were associated with the same trajectory-shape primitive.

V. EXPERIMENTS

A key property of the proposed trajectory representation is its invariance to time and the choice of body reference point. The main aim of the experiments is to demonstrate the advantages of these invariant properties for trajectory segmentation. This is done both in simulation and using real recordings of human-demonstrated pouring motions.

The simulated trajectories represent temporal rigid-body trajectories $T(t)$ of pouring motions performed with two types of objects: a teakettle and a bottle. The strategy of the pouring motion consists of a sequence of six intuitive sub-motions, also visualized in Fig. 4:

- *slide+* : the object is slid across the table from its initial position to a fixed location on the table
- *lift+* : the object is lifted and reoriented so that the spout is directed towards the glass
- *tilt+* : the object is tilted to pour the liquid
- *tilt-* : the object’s tilt is undone
- *lift-* : the object is placed back on the table
- *slide-* : the object is slid back

We aim to robustly segment the pouring motion into these sub-motions. To simulate sensor noise, white noise with a standard deviation of 2° and 1 mm were added to the object’s orientation and reference point’s location, respectively.

To study robustness to changes in the body reference point, each of the three trajectories was simulated with a different reference point: the first was near the spout (**P1**), the second near the handle (**P2**), and the third near the center of mass (**P3**). Fig. 4 on the left illustrates that, when rotation of the object is involved, the different reference points result in significantly different point trajectories.

The generalization capability is tested by applying the primitives that were learned for the kettle motion to another object. The object was chosen to be a bottle of which the reference point was chosen in the center of its opening.

The proposed approach is compared to other methods based on the literature, shown in Table I. *Method A* does not transform the object’s trajectory to a geometric domain. *Methods A and B* neglect information on the object’s orientation by choosing $L = 0$. *Methods B and C* use the arclength of the reference point’s trajectory as the geometric

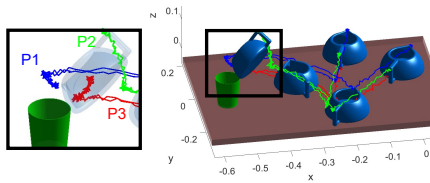


Fig. 4. Visualization of three rigid-body trajectories (red, green, and blue) representing simulated pouring motions performed with a kettle. Different body reference points (P1, P2 and P3) were considered.

TABLE I

SUMMARY OF METHODS AND CORRESPONDING TUNING PARAMETERS

method	progress-type	\dot{s}	L [m]	Δs	$\hat{\sigma}$	β
A	time	1	0	0.1s	0.1m/s	2%
B	arclength [19]	$\ \mathbf{v}\ $	0	0.02m	0.02m	5%
C	arclength [19]	$\ \mathbf{v}\ $	0.3	0.02m	0.1m	5%
D	angle [20]	$\ \boldsymbol{\omega}\ $	0.3	3°	0.1m	5%
E	combined [17]	$\sqrt{L^2\ \boldsymbol{\omega}\ ^2 + \ \mathbf{v}\ ^2}$	0.3	0.02m	0.1m	5%
F	screw-based [11]	$L\ \boldsymbol{\omega}\ + \ \mathbf{v}\ $	0.3	0.02m	0.1m	5%
G	screw-based	$\sqrt{L^2\ \boldsymbol{\omega}\ ^2 + \ \mathbf{v}\ ^2}$	0.3	0.02m	0.1m	5%

progress parameter, such as in [19]. *Method D* uses the traveled angle of the moving object as the geometric progress parameter, such as in [20]. *Method E* uses a combination of the rotational velocity and linear velocity of the reference point on the body as the progress rate, such as in [17]. *Method F* uses a screw-based progress rate without the proposed regularization, such as in [11]. *Method G* uses the proposed screw-based geometric progress rate.

To validate that the proposed method works in practice, it was tested on recordings of real pouring motions. These motions were recorded using an HTC VIVE motion capture system, consisting of a tracker attached to the kettle (see Fig. 5a), and two base stations. The VIVE system recorded the pose trajectories with a frequency of 60 Hz and an accuracy in the order of a few mm and a few degrees. To introduce contextual variations in the measurements, the tracker was physically attached to the kettle at two different locations, one at the side of the kettle (P1) and one near the top of the handle (P2). For each tracker location, three trials were recorded. Fig. 5b depicts the kettle’s pose trajectory for the first trial with the tracker attached to the side.

A. Data processing

The simulated data and the recorded data were processed in the same way. The rigid-body trajectories were first preprocessed using a Kalman smoother with a constant acceleration model to deal with the effects of the measurement noise. For the methods *B* to *G*, the trajectories were first reparameterized as in Section III-A according to the chosen definition of the progress rate \dot{s} . Then, the local shape descriptor was calculated as in Section III-D. Finally, the same segmentation algorithm (Section IV) is applied for all approaches. For reproducibility reasons, all software used in the experiments is made publicly available [21].

The values of the tuning parameters are reported in Table I. A good choice for L depends on numerous factors, including the scale of the motion and the scale of the moving object. Given the scale of the objects and motions of interest, a value



Fig. 5. (a) Human demonstration of a pouring motion using a teakettle to which an HTC VIVE tracker is attached. (b) Visualization of the first trial within a batch of six trials in the same simulation environment as Fig. 4.

of $L = 30$ cm seemed reasonable. The other parameters in Table I were manually tuned. For method *A*, the parameters Δs and $\hat{\sigma}$ have different units compared to the ones of the other methods since method *A* does not transform the trajectories to a geometric domain before segmentation.

B. Results

Figure 6 visualizes the segmentation results of all methods for the simulated data. The ground-truth segmentation points are indicated by the vertical lines. To compare between methods, the segmented trajectories were transformed back to the time domain by re-applying the motion profile $s(t)$ that was extracted from the trajectories. This was done by inverting the reparameterization procedure of Section III-C. The segmentation results are also evaluated quantitatively by reporting the number of *detected sub-motions* and *consistent segments*. A sub-motion was considered ‘detected’ when a corresponding segment was formed. A sub-motion was considered ‘consistently segmented’ when the corresponding segments were associated to the same trajectory-shape primitive across the three trials.

The results of the simulations are interpreted as follows. Method *A* generated a relatively high number of segments. The segmentation is mainly based on differences in magnitude of the reference point’s velocity. The gray segments represent regions of standstill. The light and dark blue segments represent segments of low and high magnitude in velocity, respectively. Methods *B* to *G* generated segments based on differences in shape of the rigid-body trajectory.

For methods *A-E*, the learning of the primitives and segmentation of the trajectories was dependent on the location of the reference point on the object. Furthermore, methods *A-C* could not deal well with pure rotations of the object. Method *A* classified these segments either as stationary or as segments with low magnitude in velocity. Methods *B* and *C* treated these segments as outliers. The reason for this is that during these pure rotations, the traveled arclength of the reference point remained relatively small, resulting in a small number of geometric sample points representing these pure rotations. Hence, for pure rotations, sparse clusters were created, which were seen as outliers. Following a similar reasoning, method *D* could not deal with pure translations.

Method *F* performed the segmentation in a reference-point invariant way, but could not deal well with pure translations, since \mathbf{p}_\perp is degenerate in this case.

The proposed method *G* dealt well with pure translations and performed the segmentation in a reference-point invariant way. The trajectories of the kettle were consistently seg-

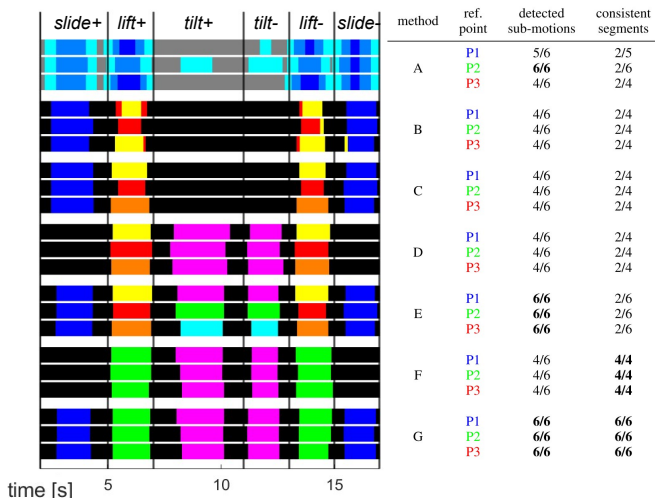


Fig. 6. Comparison of segmentation results between methods A-G for the simulated pouring motion data. Segments associated with different primitives are indicated with different colors. The black regions contain non-classified samples.

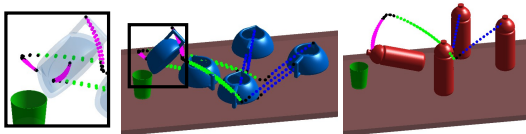


Fig. 7. Visualization of the segmented trajectories using the proposed method (method G).

mented into six segments, corresponding to the six intuitive sub-motions. More in detail, all sub-motions were detected (6/6) and all sub-motions were consistently segmented (6/6). Figure 7 visualizes the segmented trajectories.

The proposed approach also succeeded to segment the simulated pouring motion performed with a bottle (with significantly different location of the reference point) using the trajectory-shape primitives learned from the trajectories of the kettle. This illustrates the capability of the approach to generalize to different objects with different geometries.

Figure 8 visualizes the segmentation results of the proposed method G for the real recorded pouring motions. To illustrate the generated segments in the geometric domain, the extra transformation back to the time domain was not performed. The same values for the tuning parameters as reported in Table I were used for this experiment.

Three primitives were learned from the real data. The corresponding values of the mean trajectory-shape descriptor \bar{S} of the three clusters are reported in Table II. The learned primitives represent 1D motions, since for each primitive, the three columns are almost identical. The first primitive represents a 1D translation (slide). The second primitive represents a rotation with a non-zero pitch (lift). The third primitive represents a pure rotation (tilt). The segmentation approach created segments conform the six intuitive sub-motions, apart from some short segments near transition regions. These short segments can be avoided by implementing a postprocessing step or a more advanced segmentation algorithm from the literature, which is part of future work.

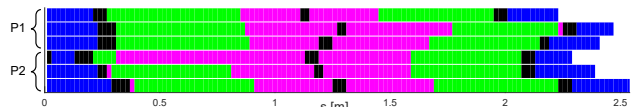


Fig. 8. Segmentation results of proposed method G for the real recorded pouring motion data. P1 and P2 correspond to trials with a different location of the motion tracker on the kettle.

TABLE II
CLUSTER MEANS LEARNED FROM THE REAL POURING MOTION DATA.

	\bar{S}_1			\bar{S}_2			\bar{S}_3		
	$s - \Delta s$	s	$s + \Delta s$	$s - \Delta s$	s	$s + \Delta s$	$s - \Delta s$	s	$s + \Delta s$
$L\omega_x$	-0.02	-0.02	-0.02	0.97	0.97	0.97	1.00	1.00	1.00
$L\omega_y$	-0.01	-0.01	-0.01	-0.03	-0.01	-0.02	0.01	-0.00	-0.01
$L\omega_z$	0.12	0.11	0.11	0.02	0.01	0.01	0.03	-0.00	-0.03
ν_x	0.99	0.99	0.99	0.26	0.26	0.26	0.06	0.06	0.06
ν_y	0.01	0.01	0.01	-0.00	0.01	0.01	0.00	0.00	-0.00
ν_z	-0.11	-0.11	-0.11	-0.01	-0.02	-0.01	0.00	0.00	-0.00

VI. DISCUSSION AND CONCLUSION

The objective of this work was to enhance template-based trajectory segmentation approaches by incorporating invariance. Time-invariance was achieved by reparameterizing the trajectory using a novel geometric progress parameter. By considering the translation along the screw axis, the progress parameter was made invariant to the choice of body reference point. Based on the reparameterized trajectory, a screw-based trajectory-shape descriptor was proposed to characterize the local geometry of the trajectory.

For the devised self-supervised segmentation scheme, the results showed a more robust detection of consecutive sub-motions with distinct features and a more consistent segmentation thanks to the invariant properties of the screw-based progress parameter and the trajectory-shape descriptor.

The proposed approach also has a more practical advantage. Due to the invariance, the formation of the segments becomes invariant to changes in sensor setup (i.e. changes in the location and angle of the camera, changes in the location of markers or trackers on the object, etc.). Therefore, sensor calibration efforts can be reduced.

Future work is to examine the benefits of the invariant segmentation approach for other types of object manipulation tasks and to verify the extent to which other segmentation methods may benefit from the invariant approach.

APPENDIX

This appendix shows that the triangle inequality property $\dot{s}_{1+2} \leq \dot{s}_1 + \dot{s}_2$ in Section III-B is always satisfied for the regularized progress rate \dot{s} if $b \leq L$ in (6).

Consider again the case of a couple of rotations generating a translation as explained in Section III-B. Additionally consider that the rotation axes of the couple are symmetrically positioned w.r.t. the reference point, such that $a = \|\mathbf{p}_\perp\|$. Equation (3) then remains valid under the proposed regularization action when $\|\mathbf{p}_\perp\| \leq b$. From (3), it was derived that the triangle inequality holds when $a \leq L$. Hence, given that $a = \|\mathbf{p}_\perp\|$ and $\|\mathbf{p}_\perp\| \leq b$, then $a \leq L$ is always satisfied when $b \leq L$.

REFERENCES

- [1] S. Dodge, R. Weibel, and E. Foroootan, "Revealing the physics of movement: Comparing the similarity of movement characteristics of different types of moving objects," *Computers, Environment and Urban Systems*, vol. 33, no. 6, pp. 419–434, 2009.
- [2] M. Teimouri, U. G. Indahl, H. Sickel, and H. Tveite, "Deriving animal movement behaviors using movement parameters extracted from location data," *ISPRS International Journal of Geo-Information*, vol. 7, no. 2, 2018.
- [3] S. Krishnan, A. Garg, S. Patil, C. Lea, G. Hager, P. Abbeel, and K. Goldberg, "Transition state clustering: Unsupervised surgical trajectory segmentation for robot learning," *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1595–1618, 2017.
- [4] M. Wächter and T. Asfour, "Hierarchical segmentation of manipulation actions based on object relations and motion characteristics," in *International Conference on Advanced Robotics*, 2015, pp. 549–556.
- [5] S. H. Lee, I. H. Suh, S. Calinon, and R. Johansson, "Autonomous framework for segmenting robot trajectories of manipulation task," *Autonomous Robots*, vol. 38, pp. 107–141, 2014.
- [6] C. Song, G. Liu, X. Zhang, X. Zang, C. Xu, and J. Zhao, "Robot complex motion learning based on unsupervised trajectory segmentation and movement primitives," *ISA Transactions*, vol. 97, pp. 325–335, 2020.
- [7] D. Kulic, W. Takano, and Y. Nakamura, "Online segmentation and clustering from continuous observation of whole body motions," *IEEE Transactions on Robotics*, vol. 25, no. 5, pp. 1158–1166, 2009.
- [8] F. Meier, E. Theodorou, F. Stulp, and S. Schaal, "Movement segmentation using a primitive library," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011, pp. 3407–3412.
- [9] M. Buchin, A. Driemel, M. J. van Kreveld, and V. S. Adinolfi, "Segmenting trajectories: A framework and algorithms using spatiotemporal criteria," *Journal of Spatial Information Science*, vol. 3, pp. 33–63, 2011.
- [10] Z. Izakian, M. S. Mesgari, and R. Weibel, "A feature extraction based trajectory segmentation approach based on multiple movement parameters," *Engineering Applications of Artificial Intelligence*, vol. 88, p. 103394, 2020.
- [11] J. De Schutter, "Invariant description of rigid body motion trajectories," *Journal of Mechanisms and Robotics*, vol. 2, no. 1, pp. 011004/1–9, 2010.
- [12] K. Lynch and F. Park, *Modern Robotics*. Cambridge University Press, 2017.
- [13] H. Bruyninckx, "Robot kinematics and dynamics," <https://u0011821.pages.gitlab.kuleuven.be/pdf/2009-HermanBruyninckx-robotics-textbook.pdf>, November 2021, course notes, Department of Mechanical Engineering, KU Leuven.
- [14] M. Chasles, "Note sur les propriétés générales du système de deux corps semblables entr'eux et placés d'une manière quelconque dans l'espace; et sur le déplacement fini ou infiniment petit d'un corps solide libre," in *Bulletin des Sciences Mathématiques, Astronomiques, Physiques et Chimiques*, vol. 14, 1830, pp. 321–326.
- [15] M. Ceccarelli, "Screw axis defined by giulio mozzi in 1763 and early studies on helicoidal motion," *Mechanism and Machine Theory*, vol. 35, no. 6, pp. 761–770, 2000.
- [16] G. Chirikjian, "Partial bi-invariance of $se(3)$ metrics 1," *Journal of Computing and Information Science in Engineering*, vol. 15, p. 011008, 2015.
- [17] F. C. Park, "Distance Metrics on the Rigid-Body Motions with Applications to Mechanism Design," *Journal of Mechanical Design*, vol. 117, no. 1, pp. 48–54, 1995.
- [18] L. Kavan, S. Collins, J. Žára, and C. O'Sullivan, "Geometric skinning with approximate dual quaternion blending," *ACM Transactions on Graphics (TOG)*, vol. 27, no. 4, pp. 1–23, 2008.
- [19] Y. Guo, Y. Li, and Z. Shao, "Dsrf: A flexible trajectory descriptor for articulated human action recognition," *Pattern Recognition*, vol. 76, pp. 137–148, 2018.
- [20] B. Roth, "Finding geometric invariants from time-based invariants for spherical and spatial motions," *Journal of Mechanical Design*, vol. 127, no. 2, pp. 227–231, 2005.
- [21] "Self-supervised segmentation of temporal rigid body trajectories," https://gitlab.kuleuven.be/robotgenskill/public_code/segmentation_paper_code.