

Weakly Supervised Recognition of Surgical Gestures

Beatrice van Amsterdam¹, Hirenkumar Nakawala², Elena De Momi³, Danail Stoyanov¹

Abstract—Kinematic trajectories recorded from surgical robots contain information about surgical gestures and potentially encode cues about surgeon’s skill levels. Automatic segmentation of these trajectories into meaningful action units could help to develop new metrics for surgical skill assessment as well as to simplify surgical automation. State-of-the-art methods for action recognition relied on manual labelling of large datasets, which is time consuming and error prone. Unsupervised methods have been developed to overcome these limitations. However, they often rely on tedious parameter tuning and perform less well than supervised approaches, especially on data with high variability such as surgical trajectories. Hence, the potential of weak supervision could be to improve unsupervised learning while avoiding manual annotation of large datasets. In this paper, we used at a minimum one expert demonstration and its ground truth annotations to generate an appropriate initialization for a GMM-based algorithm for gesture recognition. We showed on real surgical demonstrations that the latter significantly outperforms standard task-agnostic initialization methods. We also demonstrated how to improve the recognition accuracy further by redefining the actions and optimising the inputs.

Index Terms—Classification, Gaussian Mixture Models, robotic surgery, kinematics, surgical gesture recognition

I. INTRODUCTION

Robot-Assisted Minimally Invasive Surgery (RAMIS) is an established practice across a range of surgical specialties, which helps to improve precision of the surgical manipulation and ergonomic comfort of the surgeon [1]. With RAMIS, a large dataset of video and kinematic trajectories of surgical interventions can be recorded from the robotic system, e.g. da Vinci surgical system (dVSS, Intuitive Surgical Inc., CA, USA). Surgical gesture recognition, i.e. segmentation and labelling of surgical action units, by analysing these datasets automatically can be used for multiple purposes, e.g. surgical skills assessment [2], [3] and automation [4], [5].

However, automatic gesture recognition is difficult to implement robustly due to the complexity of surgical tasks and the variability in users’ actions and patient-specific anatomy [7]. A number of approaches have been proposed to address this problem. Classical approaches are based on statistical models such as Gaussian Mixture Models (GMM) [8], Hidden Markov Models (HMM) [9], [10] and

This work was supported by the Wellcome/EPSCRC Centre for Interventional and Surgical Sciences (WEISS) (203145Z/16/Z) and the EPSCRC (EP/N027078/1, EP/P012841/1, EP/P027938/1, EP/R004080/1).

¹Beatrice van Amsterdam and Danail Stoyanov are with WEISS and UCL Robotics Institute, Department of Computer Science, UCL, London, UK. beatrice.amsterdam.18@ucl.ac.uk

²Hirenkumar Nakawala is with the Department of Computer Science, Università di Verona, Verona, Italy.

³Elena De Momi is with the Department of Electronics, Information and Bioengineering, Politecnico di Milano, Milan, Italy.




Fig. 1. Example of surgemes [6]: pushing needle through tissue (L1), transferring needle from right to left (L2), pulling suture with left hand (L3), transferring needle from left to right (L4).

Conditional Random Fields (CRF) [11], [12]. More recently, deep learning techniques have also been employed, providing the current state-of-the-art results [13], [14], [15]. The challenge with most of these methods is that they use manual annotations, which are costly, time consuming and subjective when generated by multiple participants. Due to subjectivity and smooth transitions between gestures, boundaries between consecutive gestures are often not clearly defined. Unsupervised methods, that automatically learn the segmentation criterion from the data, have been developed to overcome these limitations [8], [16], [17]. However, they typically perform less well than methods trained on labelled information, especially on data with high variability such as surgical trajectories. Hence, the potential of weak supervision could be to improve unsupervised learning while avoiding manual annotation of large datasets.

The aim of this paper is to propose a new weakly supervised approach for surgical gesture recognition, that allows to retain the amount of annotations of surgical gestures limited to very few demonstrations. In particular we used at a minimum one expert demonstration and its ground truth annotations to generate an appropriate initialization for a GMM-based unsupervised recognition algorithm, in order to improve upon standard task-agnostic initialization methods, such as random or K-means initialization [18]. We focused on recognition of *surgeme* units (Fig.1), the shortest “surgical motion unit with explicit semantic sense” [16] (e.g. *grasping the needle*, *pulling the suture*, etc.). We validated our algorithm on the JIGSAWS dataset [6], [19], featuring suturing demonstrations collected from eight surgeons with different skill level using the dVSS.

II. RELATED WORK

GMM-based methods: this group of works segment robot trajectories into action classes by fitting a GMM onto the available samples. In [20] the number of mixture components was chosen using the Bayesian Information Criterion (BIC) and the fitting was initialized using the K-means clustering algorithm. [8] presented a multi-level clustering approach for identification and pruning of segmentation points, which is based on a Dirichlet Process GMM (DPGMM), i.e. a mixture model where the number of clusters is determined by a DP. This work was extended integrating features extracted with deep neural networks from the video data [21], improving the recognition accuracy.

Heuristic initialization: a number of studies make use of heuristics to create an initial segmentation of the data. Examples include heuristics based on Zero Crossing Velocity [22], jerk profiles [23] or trajectory curvature [24]. However, the same heuristic could be not appropriate to explain every part of the data, and the output is often over-segmented [24].

Weakly supervised methods: only few weakly supervised approaches have been developed for surgical action recognition [25], [26], [27]. Some works, however, assume that the actions follow a pre-defined order, the goal is to find the action boundaries [26], [27]. In addition, these methods were only applied to recognise surgical phases, which represent high-level surgical states, using only the video data and lack the recognition of low-level surgeme units.

III. METHODS

Our approach relies on classical GMM clustering. Multiple reasons make GMM an appealing method for gesture recognition, such as performing simultaneous segmentation and classification, where one task does not rigidly influence the other, as in sequential approaches. GMM is intuitive because action classes are represented through independent means, covariance matrices and weights. Finally, the fuzziness at the segment boundaries is modelled through Gaussian intersections.

Notation: vectors are represented in bold lowercase letters (e.g. \mathbf{x}), matrices are represented in bold capital letters (e.g. \mathbf{A}) and scalars are represented in italic letters (e.g. t).

A. Data pre-processing

We used JIGSAWS [6], a public dataset comprising video and kinematic data captured at 30 Hz from the dVSS during multiple demonstrations of elementary surgical tasks, which were performed on phantoms by eight surgeons with different robotic surgical experience (expert, intermediate, novice). JIGSAWS also contains manual annotations describing the ground truth segmentation of each demonstration into action classes.

We tested our algorithm on the kinematic data recorded from the two Patient Side Manipulators (PSM1 and PSM2) of the dVSS [28]. The motion of each arm is described by a local frame attached at its end-effector using 19 kinematic variables, including Cartesian positions, a rotation matrix, linear velocities, angular velocities and a gripper angle.

TABLE I
KINEMATIC FEATURE VECTOR.

Indices	Description of variables
1–3	Right PSM tool tip position (xyz)
4–7	Right PSM tool tip orientation quaternion (\mathbf{q})
8–10	Right PSM tool tip linear velocity ($\dot{x}\dot{y}\dot{z}$)
11–13	Right PSM tool tip rotational velocity ($\dot{\alpha}\dot{\beta}\dot{\gamma}$)
14	Right PSM gripper angle (θ)
15–17	Left PSM tool tip position (xyz)
18–21	Left PSM tool tip orientation quaternion (\mathbf{q})
22–24	Left PSM tool tip linear velocity ($\dot{x}\dot{y}\dot{z}$)
25–27	Left PSM tool tip rotational velocity ($\dot{\alpha}\dot{\beta}\dot{\gamma}$)
28	Left PSM gripper angle (θ)
29–32	Euclidean distance signals (d_x, d_y, d_z, d)

The pre-processing pipeline of [16] was implemented:

- The rotation matrix \mathbf{R} describing the end-effector orientation with respect to the robot base is converted into a more compact quaternion representation \mathbf{q} , reducing the state vector to 14 variables for each arm.
- All the trajectories are smoothed with a low-pass filter with cut-off frequency $f_c = 1.5$ Hz in order to minimize the measurement noise.
- All the trajectories are normalized to zero mean and unit variance, in order to enable a fair comparison between signals with different unit of measure.

Additionally, four signals representing the distance between the two end-effectors along the three orthogonal axes (d_x, d_y, d_z) and their absolute Euclidean distance (d) are generated, in order to include information about the relationship between the two manipulators, resulting in a state vector $\mathbf{x}(t) \in R^p$ of $p = 32$ variables. Table I describes in detail the variables included in the kinematic feature vector.

Finally, the trajectories are subsampled from 30 Hz to 10 Hz for faster computation time.

B. Simultaneous action segmentation and recognition

We build upon the approach of [8], treating each demonstration $\mathbf{x}(t) \in R^p$, as a realization of a switched linear dynamical system with zero-mean Gaussian process noise $\mathbf{w}(t) \in R^p$:

$$\mathbf{x}(t+1) = \mathbf{A}_k \mathbf{x}(t) + \mathbf{w}(t), \quad \mathbf{A}_k \in \{\mathbf{A}_1, \dots, \mathbf{A}_N\} \quad (1)$$

where each different locally linear regime $\mathbf{A}_k \in R^{p \times p}$ corresponds to one of the N different surgemees composing the task. As explained in [8], under this hypothesis action recognition can be performed by fitting a GMM to the augmented state $\mathbf{n}(t)$, defined as:

$$\mathbf{n}(t) = [\mathbf{x}(t), \mathbf{x}(t+1), \dots, \mathbf{x}(t+W)] \in R^{pW}. \quad (2)$$

When $W = 1$, GMM fitting is indeed equivalent to solving multiple linear regression problems [29], one for each action class \mathbf{A}_k . After model fitting, each trajectory sample is assigned to its most likely mixture component, i.e. its most likely surgeme label.

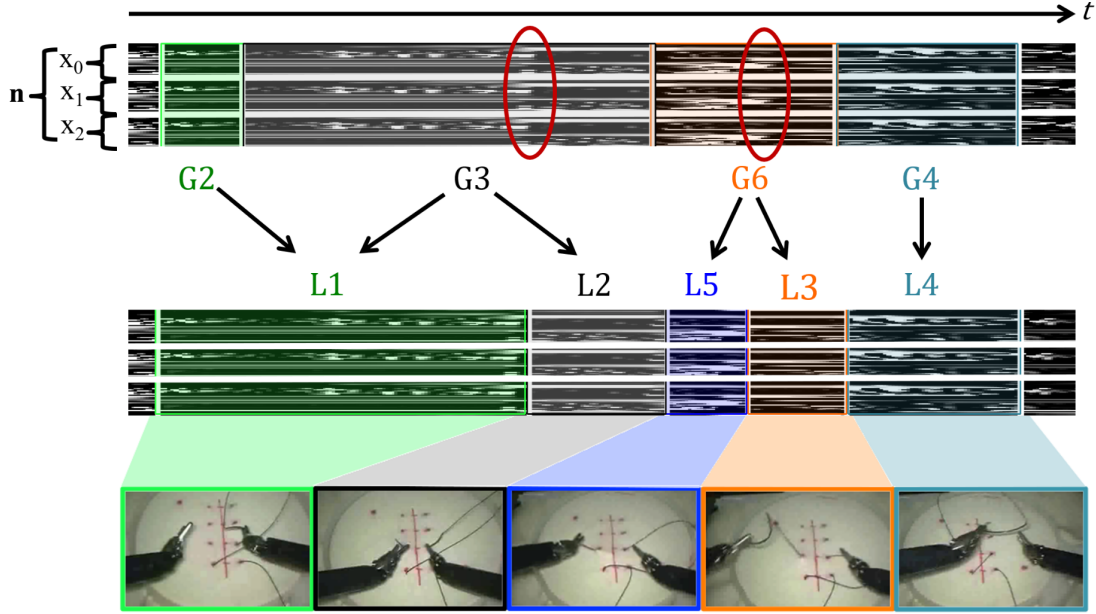


Fig. 2. The schematic shows the augmented state vector $\mathbf{n}(t) = [\mathbf{x}_0 \ \mathbf{x}_1 \ \mathbf{x}_2]$, $\mathbf{x}_0=\mathbf{x}(t)$, $\mathbf{x}_1=\mathbf{x}(t+1)$, $\mathbf{x}_2=\mathbf{x}(t+2)$, of a surgical demonstration example. Each pixel row corresponds to the value, mapped into gray levels (black = min value, white = max value), of a kinematic feature in time. An overlay of surgeme labels is shown above the raw kinematic values for both the original annotation (G) and the new annotation (L) we propose in this paper. The visual information at each surgeme is also shown below, although we do not use visual features explicitly within our approach.

C. Weakly supervised initialization

In order to initialize the GMM parameters (mean, covariance and weight of each mixture component), we use a small set of manually-segmented demonstrations. This set is composed of two demonstrations from expert users ($\mathbf{e}_1(t), \mathbf{e}_2(t)$) and one demonstration from an intermediate user ($\mathbf{e}_3(t)$), randomly selected among all demonstrations observed to be free from execution errors (such as needle dropping or multiple attempts of the same gesture). This set provides an exemplary execution of each possible action, which helps to generate a mixture with the correct number of components and appropriate shape. Exploiting the available ground truth annotations, we fit an initial GMM, denoted as GMM_0 , to the example demonstrations ($\mathbf{e}_1(t), \mathbf{e}_2(t), \mathbf{e}_3(t)$), thus obtaining the initial values of the mean vector ($\boldsymbol{\mu}_{0k}$), covariance matrix (\mathbf{C}_{0k}) and weight (w_{0k}) of each mixture component:

$$\mathbf{e}_1(t), \mathbf{e}_2(t), \mathbf{e}_3(t) \rightarrow \{\boldsymbol{\mu}_{0k}, \mathbf{C}_{0k}, w_{0k}\}_{k=1:N}$$

D. Trajectory segmentation

Once the initial mixture parameters have been generated, offline segmentation of the full dataset was performed by fitting another GMM onto the unlabelled demonstrations using the Expectation Maximization (EM) algorithm [30]. Each trajectory sample is assigned to its most likely mixture component, i.e. its most likely action label.

E. Ground Truth segmentation redefinition

As introduced in Section B, the action recognition method of this study relies on the hypothesis of local linearity

in demonstrations, where each locally linear regime corresponds to a different surgeme. Thus, we decided to analyse the suturing task in order to check if this hypothesis is approximately verified.

When observing the video records of the suturing demonstrations, we noticed the presence of sudden motion variations during the execution of some of the surgemes. By plotting (Fig. 2 top) the corresponding kinematic state vector $\mathbf{n}(t)$ as an image (where each pixel row corresponds to the value, mapped into gray levels, of a kinematic feature in time), overlaid with ground truth action boundaries (labels G), we noticed indeed the presence of sharp transitions of the kinematic pattern within those surgemes (e.g. see G3 and G6). In order to better satisfy the local linearity hypothesis required by our recognition algorithm, we therefore redefined our ground truth annotations (Fig. 2 bottom) in a way to avoid abrupt motion variations within surgemes.

Using the video feedback and the original annotations, all the trajectories have therefore been re-segmented according to the following criteria:

- Surgeme (G3) *pushing needle through the tissue* is split into (L1) *pushing needle through tissue* and (L2) *transferring needle from right to left*.
- Surgeme (G6) *pulling suture with left hand* is split into (L5) *extracting suture from tissue with left hand* and (L3) *pulling suture with left hand*.
- Surgeme (G11) *dropping suture and moving to end points* is, when necessary, split into (L7) *orienting needle*, (L9) *dropping suture*, and (L10) *moving to end points*.

- L1: pushing needle through tissue
- L2: transferring needle from right to left
- L3: pulling suture with left hand
- L4: transferring needle from left to right
- L5: extracting suture from tissue with left hand
- L6: grasping needle with right hand;
- L7: orienting needle
- L8: using right hand to help tighten suture
- L9: dropping suture
- L10: moving to end points

Fig. 3. Redefined action dictionary. Each surgeme is represented with a different colour for visualization purposes.

- Surgeme (G5) *moving to centre of workspace with needle in grip* is most of the times performed simultaneously with (L7) *orienting needle* or (G2) *positioning the tip of the needle*. We therefore include it in either of the two.
- Surgeme (G2) *positioning the tip of the needle* and (G3) *pushing needle through the tissue* are merged into (L1) *pushing needle through tissue* class, because there is no clear transition point between the two actions. Moreover, small needle repositioning motions are often performed when inserting the needle through the tissue.

The redefined action dictionary is presented in Fig. 3.

IV. EVALUATION METRICS

We evaluate our algorithm performance in a similar way to [21]. We use both extrinsic metrics, comparing the segmentation result to the ground truth annotations, and intrinsic metrics, measuring the compactness of the generated transition point clusters.

A. Extrinsic metrics

Accuracy: The accuracy represents the percentage of correctly labelled frames.

Normalized Mutual Information (NMI): The NMI measures the alignment between two sequences of labels (X and Y):

$$NMI(X, Y) = \frac{I(X, Y)}{\sqrt{H(X)H(Y)}}$$

where I is the mutual information and H the entropy [31]. This metric is independent of the absolute values of the labels, i.e. the score is not affected by permutations of cluster labels.

B. Intrinsic metrics

Silhouette Index (SI): The SI for a single sample is defined as:

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}$$

where $a(i)$ is the distance between that sample and the mean of the cluster it belongs to, while $b(i)$ is the distance between that sample and the mean of the nearest cluster it is not

part of. We employed the Euclidean distance metric. The Silhouette value is a measure of how similar an object is to its own cluster compared to other clusters. The individual Silhouette value ranges from -1 to +1. We normalised it between 0 and 1, as in [21]. A high value indicates that the object is well matched to its own cluster and poorly matched to neighbouring clusters. SI is the average Silhouette score over all N_s samples:

$$SI = \sum_{i=1}^{N_s} s(i)$$

High SI indicates that the clustering configuration is appropriate. We call SI_{GMM0} the SI computed on the clusters identified by our algorithm, and SI_{GT} the SI computed on the Ground Truth clusters.

V. EXPERIMENTS AND RESULTS

As described in Fig. 4, we conducted two sets of experiments, the first on a dataset comprising only expert demonstrations, and the second on a dataset comprising expert, intermediate and novice demonstrations.

JIGSAWS features 10 suturing demonstrations from expert surgeons, 10 suturing demonstrations from intermediate surgeons and other 20 from novice surgeons. Each demonstration has a different duration of approximately 1105 ± 432 frames. The expert demonstrations have generally the shortest duration.

A. Sliding window length

First of all, we conducted a preliminary test to select the optimal sliding window length (W), which defines the dimensions of the augmented feature state $\mathbf{n}(t)$. We applied our GMM_0 method on expert demonstrations for increasing values of W , and computed the corresponding accuracy score with redefined dictionary. As illustrated in Fig. 5, the trend is negative, as accuracy decreases when W increases. This could be explained by the corresponding increase in feature dimensionality, which jeopardises the clustering robustness




Fig. 4. We conducted a first set of experiments on expert demonstrations (in red colour), where we first performed a preliminary estimation of the optimal sliding window length (SW, in Section A). We then compared the performance of our initialization method with redefined dictionary to the performance with original Ground Truth annotations (GT, in Section B) and to the performance of K-means initialization method (INIT, in Section C). We finally performed input feature selection (FEAT, in Section D). We then tested, in a second set of experiments, the robustness of our method with redesigned dictionary and selected features on an extended set (EXT, in Section E) of expert, intermediate and novice demonstrations (in blue colour).

TABLE II
RESULTS ON EXPERT DEMONSTRATIONS.

Annotations		Original	Proposed				
Initialization		GMM_0	K-means	GMM_0			
Signals		All		No pose	No velocity	No distance	
Extrinsic metrics	Accuracy	58%	\	83%	55%	85%	77%
	NMI	0.58	0.58	0.72	0.41	0.73	0.67
Intrinsic metrics	SI_{GMM0}	0.55	0.52	0.57	0.48	0.60	0.55
	SI_{GT}	0.59	0.56	0.56	0.54	0.56	0.56




Fig. 5. Accuracy score as a function of the sliding window length W . The best recognition performance is obtained with $W=2$.

given the same amount of samples. We selected $W=2$, providing the best recognition performance.

B. Ground truth redefinition

We then analysed the performance of our GMM_0 method on both the original and the proposed ground truth annotations. The results are summarized in Table II. With the proposed annotations the extrinsic metrics show remarkable improvement with respect to the original ground truth, with accuracy score increasing of 25% and NMI score increasing of 14%. The SI_{GMM0} is also improved, advancing from 0.55 to 0.57.

In addition, we compared the recognition accuracy between groups of corresponding labels belonging to the original and the proposed action dictionaries. As shown in Fig. 6, the fusion of surgeme G2 and G3 into L1 and the separation of surgeme G3 into L1 and L2 give rise to action classes L1 and L2 which can be recognized more robustly, while the separation of G6 into L5 and L3 does not generate significant variations. Recognition accuracy of G11, split into L9 and L10, decreases with the proposed annotations, but the accuracy of all the other labels (G1, G4, G9, G8) is mostly improved. Overall, a more robust GMM distribution is generated when initialized with the proposed action dictionary. These results underline the influence of the action dictionary definition on action recognition performance.

C. Initialization technique

We compared the performance of our GMM_0 initialization method with respect to the commonly used K-means initialization method [32], with redefined dictionary (see




Fig. 6. Accuracy score comparison between groups of corresponding labels belonging to the original (in blue colour) and the proposed (in red colour) action dictionaries.

Table II). The number of K-means clusters is set as the number of action labels in the dictionary and the initial seeds are randomly sampled from the dataset. Being K-means algorithm fully unsupervised, no information about the identity of the generated clusters is available. For this reason the accuracy score was not computed.

GMM_0 initialization leads to 14% improvement of NMI, as well as increase of SI_{GMM0} . K-means indeed assumes equal prior probability for all K clusters (i.e. each cluster has roughly the same number of observations) [33] and it is randomly initialized. GMM_0 , on the other hand, exploits prior information to model the initial location, shape and size of the clusters, leading to more robust action identification.

D. Feature selection

We also studied the influence of different signals on the recognition accuracy. Specifically, we analysed the contribution of the pose (NO pose), velocity (NO velocity) and Euclidean distance (NO distance) signals by observing how the performance changes in their absence. The results in Table II suggest that the velocity signals should be discarded, as the recognition performance improves when they are excluded. Velocity signals are indeed a major source of within-cluster variability: not only users with different expertise level perform surgical tasks at different speeds, but even within the same demonstration velocity signals belonging to the same action class show high variability (see Fig. 7). The Euclidean distance signals we introduced, instead, give positive contribution to the classification accuracy, as the recognition performance degrades when they are excluded.




Fig. 7. Example of normalized position trajectory (top) and normalized linear velocity signal (bottom). Each colour represents a different surgeon, as described in Fig.3. Position segments having the same label show higher repeatability than the correspondent velocity segments. Velocity signals are indeed a major source of within-cluster variability.




Fig. 8. Example of segmentation output (bottom) and corresponding ground truth (top). Each colour represents a different surgeon, as described in Fig.3.

Fig. 8 shows an example of segmentation output, compared to its ground truth. Fig. 9 shows the 2D distribution, obtained with the t-SNE visualization technique [34], of the transition points identified by our algorithm on the expert set. Cluster compactness is visually comparable to the ground truth distribution.

E. Extended dataset

Finally we extended our algorithm validation to the full dataset, in order to test the robustness of our method, with redesigned dictionary and selected features, against increasing data variability and the presence of spurious motions. Specifically, we used the same initialization (GMM_0) as in the previous experiments, but we extended the unsupervised GMM fitting to all expert (E), intermediate (I) and novice (N) demonstrations. As summarized in Table III, the lower the expertise level of the surgeon, the lower the final accuracy,




Fig. 9. t-SNE representation of the transition point distribution identified by our algorithm (left), compared to the ground truth distribution (right). Each colour represents a different surgeon, as described in Fig.3.

TABLE III
RESULTS ON EXTENDED DATASET.

Metrics		E	E+I	E+I+N
<i>Extrinsic metrics</i>	<i>Accuracy</i>	85%	77%	59%
	<i>NMI</i>	0.73	0.66	0.46
<i>Intrinsic metrics</i>	SI_{GMM0}	0.60	0.53	0.50
	SI_{GT}	0.54	0.54	0.52

NMI and SI_{GMM0} scores. Simple GMM approaches do not exploit temporal constraints such as transition probabilities between actions. This constitutes a major limitation in the analysis of sequential information such as kinematic trajectories, resulting in limited performance as the data variability increases.

VI. CONCLUSION

This paper explored a new weakly supervised approach for surgical gesture recognition, that allows to retain the amount of annotations limited to very few demonstrations. We employed three demonstrations and their ground truth annotations to generate an appropriate initialization for a GMM-based recognition algorithm. Experimental results on real surgical kinematic trajectories during a training exercise confirm that weakly supervised initialization significantly outperforms standard task-agnostic initialization methods. We also demonstrated that recognition accuracy can be improved by carefully designing the optimal channel selection and the appropriate action granularity for the specific task at hand. We believe that inclusion of contextual and semantic information [35], [36] from video data would further boost the recognition performance [21].

However, manual redefinition of the action dictionary based on visual verification still involves a certain degree of subjectivity, and further validation should also be performed to assess the recognition performance for different sets of manually-segmented demonstrations employed in the initialization step. In addition, simple GMM approaches are not robust against increasing data variability. More complex GMM-based methods have been developed specifically for time series analysis, such as Gaussian-HMM [37] and GMM-HMM [38], [39]. These models introduce transition probabilities between different actions, thus generating a probabilistic action grammar that helps to improve the recognition accuracy. In future work, we will explore the effects of weak supervision on the initialization of transition and observation probability distributions in unsupervised HMM-based approaches. Finally, our experiments were conducted on a small-scale dataset, raising concerns about the generalization capability of our approach to surgical data modelling more broadly. More in-depth analysis will be performed on larger and more challenging datasets of robotic surgical demonstrations, e.g. [40].

REFERENCES

- [1] K. Moorthy, Y. Munz, A. Dosis, J. Hernandez, S. Martin, F. Bello, T. Rockall, and A. Darzi, "Dexterity enhancement with robotic surgery," *Surgical Endoscopy*, vol. 18, no. 5, pp. 790–795, 2004.
- [2] C. E. Reiley and G. D. Hager, "Task versus subtask surgical skill evaluation of robotic minimally invasive surgery," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 5761 LNCS, no. PART 1, pp. 435–442, 2009.
- [3] H. C. Lin, I. Shafran, D. Yuh, and G. D. Hager, "Towards automatic skill evaluation: Detection and segmentation of robot-assisted surgical motions," *Computer Aided Surgery*, vol. 11, no. 5, pp. 220–230, 2006.
- [4] N. Ettehad, S. Manaffam, and A. Behal, "Learning from demonstration: Generalization via task segmentation," in *IOP Conference Series: Materials Science and Engineering*, vol. 261, p. 012001, IOP Publishing, 2017.
- [5] R. Fox, S. Krishnan, I. Stoica, and K. Goldberg, "Multi-level discovery of deep options," *arXiv preprint arXiv:1703.08294*, 2017.
- [6] Y. Gao, S. S. Vedula, C. E. Reiley, N. Ahmidi, B. Varadarajan, H. C. Lin, L. Tao, L. Zappella, B. Béjar, D. D. Yuh, C. C. G. Chen, R. Vidal, S. Khudanpur, and G. D. Hager, "JHU-ISI Gesture and Skill Assessment Working Set (JIGSAWS): A Surgical Activity Dataset for Human Motion Modeling," *Modeling and Monitoring of Computer Assisted Interventions (M2CAI) – MICCAI Workshop*, pp. 1–10, 2014.
- [7] C. Cao, C. MacKenzie, and S. Payandeh, "Task and motion analyses in endoscopic surgery," in *Proceedings ASME Dynamic Systems and Control Division*, pp. 583–590, Citeseer, 1996.
- [8] 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.
- [9] L. Tao, E. Elhamifar, S. Khudanpur, G. D. Hager, and R. Vidal, "Sparse hidden markov models for surgical gesture classification and skill evaluation," in *International conference on information processing in computer-assisted interventions*, pp. 167–177, Springer, 2012.
- [10] B. Varadarajan, C. Reiley, H. Lin, S. Khudanpur, and G. Hager, "Data-derived models for segmentation with application to surgical assessment and training," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 5761 LNCS, no. PART 1, pp. 426–434, 2009.
- [11] L. Tao, L. Zappella, G. D. Hager, and R. Vidal, "Surgical gesture segmentation and recognition," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 8151 LNCS, no. PART 3, pp. 339–346, 2013.
- [12] E. Mavroudi, D. Bhaskara, S. Sefati, H. Ali, and R. Vidal, "End-to-end fine-grained action segmentation and recognition using conditional random field models and discriminative sparse coding," in *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1558–1567, IEEE, 2018.
- [13] A. P. Twinanda, S. Shehata, D. Mutter, J. Marescaux, M. De Mathelin, and N. Padoy, "EndoNet: A Deep Architecture for Recognition Tasks on Laparoscopic Videos," *IEEE Transactions on Medical Imaging*, vol. 36, no. 1, pp. 86–97, 2017.
- [14] C. Lea, R. Vidal, and G. D. Hager, "Learning convolutional action primitives for fine-grained action recognition," *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2016-June, pp. 1642–1649, 2016.
- [15] C. L. B. A. Reiter, and G. D. Hager, "Temporal Convolutional Networks: A Unified Approach to Action Segmentation," vol. 9915, pp. 47–54, 2016.
- [16] F. Despinoy, D. Bouget, G. Forestier, C. Penet, N. Zemiti, P. Poignet, and P. Jannin, "Unsupervised Trajectory Segmentation for Surgical Gesture Recognition in Robotic Training," *IEEE Transactions on Biomedical Engineering*, vol. 63, no. 6, pp. 1280–1291, 2016.
- [17] M. J. Fard, S. Ameri, R. B. Chinnam, and R. D. Ellis, "Soft Boundary Approach for Unsupervised Gesture Segmentation in Robotic-Assisted Surgery," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 171–178, 2017.
- [18] J. Blömer and K. Bujna, "Simple methods for initializing the em algorithm for gaussian mixture models," *CoRR*, 2013.
- [19] N. Ahmidi, L. Tao, S. Sefati, Y. Gao, C. Lea, B. B. Haro, L. Zappella, S. Khudanpur, R. Vidal, and G. D. Hager, "A Dataset and Benchmarks for Segmentation and Recognition of Gestures in Robotic Surgery," *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 9, pp. 2025–2041, 2017.
- [20] S. H. Lee, I. H. Suh, S. Calinon, and R. Johansson, "Autonomous framework for segmenting robot trajectories of manipulation task," *Autonomous Robots*, vol. 38, no. 2, pp. 107–141, 2014.
- [21] A. Murali, A. Garg, S. Krishnan, F. T. Pokorny, P. Abbeel, T. Darrell, and K. Goldberg, "TSC-DL: Unsupervised trajectory segmentation of multi-modal surgical demonstrations with Deep Learning," *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2016-June, pp. 4150–4157, 2016.
- [22] A. Fod, M. J. Matarić, and O. C. Jenkins, "Automated derivation of primitives for movement classification," *Autonomous robots*, vol. 12, no. 1, pp. 39–54, 2002.
- [23] B. Rohrer and N. Hogan, "Avoiding spurious submovement decompositions ii: a scattershot algorithm," *Biological cybernetics*, vol. 94, no. 5, pp. 409–414, 2006.
- [24] R. Lioutikov, G. Neumann, G. Maeda, and J. Peters, "Learning movement primitive libraries through probabilistic segmentation," *The International Journal of Robotics Research*, vol. 36, no. 8, pp. 879–894, 2017.
- [25] G. Quellec, K. Charrière, M. Lamard, Z. Droueche, C. Roux, B. Cochener, and G. Cazuguel, "Real-time recognition of surgical tasks in eye surgery videos," *Medical image analysis*, vol. 18, no. 3, pp. 579–590, 2014.
- [26] N. Padoy, T. Blum, S.-A. Ahmadi, H. Feussner, M.-O. Berger, and N. Navab, "Statistical modeling and recognition of surgical workflow," *Medical image analysis*, vol. 16, no. 3, pp. 632–641, 2012.
- [27] F. Lalys, L. Riffaud, D. Bouget, and P. Jannin, "An application-dependent framework for the recognition of high-level surgical tasks in the or," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 331–338, Springer, 2011.
- [28] P. Kazanzides, Z. Chen, A. Deguet, G. S. Fischer, R. H. Taylor, and S. P. DiMaio, "An open-source research kit for the da vinci® surgical system," in *2014 IEEE international conference on robotics and automation (ICRA)*, pp. 6434–6439, IEEE, 2014.
- [29] T. M. Moldovan, S. Levine, M. I. Jordan, and P. Abbeel, "Optimism-Driven Exploration for Nonlinear Systems," pp. 3239–3246, 2015.
- [30] T. K. Moon, "The expectation-maximization algorithm," *IEEE Signal processing magazine*, vol. 13, no. 6, pp. 47–60, 1996.
- [31] A. Strehl and J. Ghosh, "Cluster ensembles—a knowledge reuse framework for combining multiple partitions," *Journal of machine learning research*, vol. 3, no. Dec, pp. 583–617, 2002.
- [32] D. Arthur and S. Vassilvitskii, "k-means++: The advantages of careful seeding," in *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 1027–1035, Society for Industrial and Applied Mathematics, 2007.
- [33] P. S. Bradley and U. M. Fayyad, "Refining initial points for k-means clustering," in *ICML*, vol. 98, pp. 91–99, Citeseer, 1998.
- [34] L. V. D. Maaten and G. Hinton, "Visualizing Data using t-SNE," *Journal of Machine Learning Research 1*, vol. 620, no. 1, pp. 267–84, 2008.
- [35] X. Du, T. Kurmann, P.-L. Chang, M. Allan, S. Ourselin, R. Sznitman, J. D. Kelly, and D. Stoyanov, "Articulated multi-instrument 2-d pose estimation using fully convolutional networks," *IEEE transactions on medical imaging*, vol. 37, no. 5, pp. 1276–1287, 2018.
- [36] M. Allan, S. Ourselin, D. J. Hawkes, J. D. Kelly, and D. Stoyanov, "3-d pose estimation of articulated instruments in robotic minimally invasive surgery," *IEEE transactions on medical imaging*, vol. 37, no. 5, pp. 1204–1213, 2018.
- [37] S. Calinon, D. Florent, E. L. Sauser, D. G. Caldwell, and A. G. Billard, "An approach based on Hidden Markov Model and Gaussian Mixture Regression," *IEEE Robotics and Automation Magazine*, vol. 17, pp. 44–45, 2010.
- [38] H. Tang, M. Hasegawa-Johnson, and T. S. Huang, "Toward robust learning of the gaussian mixture state emission densities for hidden markov models," *Audio*, pp. 5242–5245, 2010.
- [39] C. Loukas and E. Georgiou, "Surgical workflow analysis with Gaussian mixture multivariate autoregressive (GMMAR) models: A simulation study," *Computer Aided Surgery*, vol. 18, no. 3-4, pp. 47–62, 2013.
- [40] D. Sarikaya, J. J. Corso, and K. A. Guru, "Detection and localization of robotic tools in robot-assisted surgery videos using deep neural networks for region proposal and detection," *IEEE Transactions on Medical Imaging*, vol. 36, pp. 1542–1549, July 2017.