

Gesture Recognition in Robotic Surgery: a Review

Beatrice van Amsterdam¹, Matthew J. Clarkson¹, Danail Stoyanov¹

Abstract—Objective: Surgical activity recognition is a fundamental step in computer-assisted interventions. This paper reviews the state-of-the-art in methods for automatic recognition of fine-grained gestures in robotic surgery focusing on recent data-driven approaches and outlines the open questions and future research directions. **Methods:** An article search was performed on 5 bibliographic databases with combinations of the following search terms: robotic, robot-assisted, JIGSAWS, surgery, surgical, gesture, fine-grained, surgeme, action, trajectory, segmentation, recognition, parsing. Selected articles were classified based on the level of supervision required for training and divided into different groups representing major frameworks for time series analysis and data modelling. **Results:** A total of 52 articles were reviewed. The research field is showing rapid expansion, with the majority of articles published in the last 4 years. Deep-learning-based temporal models with discriminative feature extraction and multi-modal data integration have demonstrated promising results on small surgical datasets. Currently, unsupervised methods perform significantly less well than the supervised approaches. **Conclusion:** The development of large and diverse open-source datasets of annotated demonstrations is essential for development and validation of robust solutions for surgical gesture recognition. While new strategies for discriminative feature extraction and knowledge transfer, or unsupervised and semi-supervised approaches, can mitigate the need for data and labels, they have not yet been demonstrated to achieve comparable performance. Important future research directions include detection and forecast of gesture-specific errors and anomalies. **Significance:** This paper is a comprehensive and structured analysis of surgical gesture recognition methods aiming to summarize the status of this rapidly evolving field.

Index Terms—gesture recognition, surgical data science, robotic surgery.

I. INTRODUCTION

Robot-assisted minimally invasive surgery (RMIS) is now well established in clinical practice as a means of enhancing surgical instrumentation and ergonomics with respect to conventional laparoscopic approaches [1]. Surgical specialisations like gynecology and urology have driven the adoption of RMIS, but diffusion of the robotic approach is increasing as hospitals become equipped with surgical robotic systems, surgeons trained to use them and both the capital and running costs reduce [2]. In addition to better surgical instrumentation, surgical robots can potentially be platforms for enabling data driven solutions for better OR capabilities [3].

The work was supported by the Wellcome/EPSCRC Centre for Interventional and Surgical Sciences (WEISS) [203145Z/16/Z]; the Engineering and Physical Sciences Research Council (EPSRC) [EP/P027938/1, EP/R004080/1, EP/P012841/1]; the Royal Academy of Engineering Chair in Emerging Technologies scheme; the EPSRC i4Health CDT (EP/S021930/1).

¹Beatrice van Amsterdam, Matthew J. Clarkson and Danail Stoyanov are with the Wellcome/EPSCRC Centre for Interventional and Surgical Sciences (WEISS), University College London, UK. beatrice.amsterdam.18@ucl.ac.uk




Fig. 1. Surgeme examples for the suturing task: (a) positioning needle tip on insertion point, (b) pushing needle through the tissue, (c) pulling needle out of tissue, (d) transferring needle from left hand to right hand.

Besides digital video, surgical robots can capture quantitative instrument motion trajectories during surgical interventions enabling analysis of surgical activity that is not possible with traditional instrumentation. These new data streams have paved the way for data-driven computational models in computer assisted interventions (CAI) and surgical data science (SDS) [4] to facilitate better procedural understanding and intra-operative support, including the possibility of linking back to the robot control loop to achieve surgical automation. The resultant data has also allowed research into automated surgical competence and skill assessment [5, 6, 7] and adaptive surgical training [8], aiming at quantifying surgical process without expert monitoring.

Understanding surgical process has often been approached by decomposition into pre-defined procedural segments at different granularity levels. Starting from the finest level, surgical processes have been hierarchically decomposed into *dexemes*, *surgemes* (generally called *gestures*, Fig. 1), *activities*, *phases*, *procedures* and *states* [9, 10]. Unfortunately, the lack of a universal taxonomy and of clear segmentation criteria has led researchers to propose different decomposition strategies (e.g. semantic [11] vs event-based [12] labelling) and alternative motion dictionaries [13, 11, 14], resulting in conceptual inconsistencies and a lack of standardised terminology.

Segmentation at different levels highlights distinct aspects of the surgical procedure, leading to a range of applications. For example, higher-level situation-awareness (state, phase or activity recognition) is well suited for workflow optimization, scheduling and resource management [15]. Recognition of fine-grained motion (*dexemes*, *surgemes*) is instead better suited for low-level analysis, assessment and reproduction

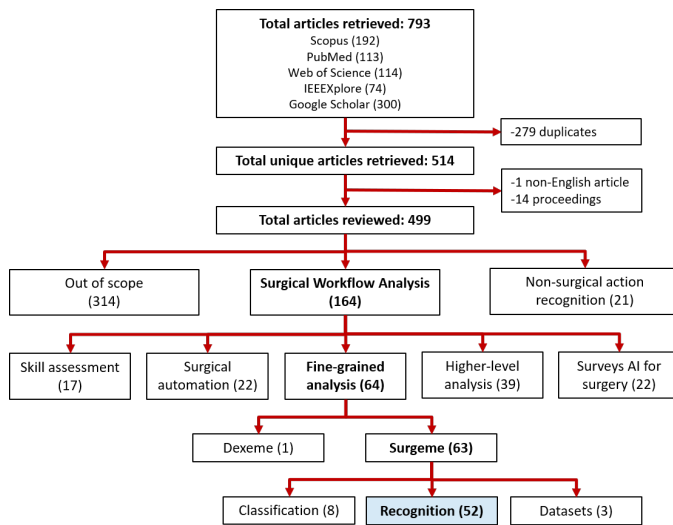


Fig. 2. Selection strategy. We identified a total of 52 articles on automatic recognition of fine-grained gestures in robot-assisted surgery, which is primarily laparoscopic and technically challenging for recognition due to environment deformations, camera motion and illumination changes.

of dexterous motion. While dexemes describe short motion segments devoid of medical sense (e.g. turning left) [10], surgemes represent surgical gestures made with a specific purpose (e.g. grabbing the needle) and are the focus of this review. Kinematic trajectory segmentation into gestures has been shown to have potential to describe technical skill development [13, 11, 14] in a more quantitative manner than relying on global performance measures as time to completion or total path length [16]. Gesture recognition also finds application in surgical automation because short motion segments are less complex and therefore easier to learn and generalize than long surgical tasks, and can be regarded as modular blocks of motion to compose and reuse [17].

Fine-grained analysis of surgical demonstrations is however non-trivial and presents a number of challenges. When generated manually, gesture annotations are costly, time consuming and subject to personal interpretation of multiple participants. Due to subjectivity and smooth transitions between gestures, boundaries between consecutive segments are often not clearly defined and it is difficult, even for a single annotator, to generate consistent segmentations as data grow over time. For this reason significant research effort is aimed at automating the annotation process utilizing both video and instrument kinematics, either individually or in combination [18, 19, 20]. Automation remains very challenging due to the complexity of surgical tasks and the presence of different variability factors, including surgeon’s skill level, operative style, type of procedure and patient-specific anatomy, but could help to efficiently generate the large amount of training data required for robust CAI implementations. Beyond off-line generation of training datasets for surgical automation and training, automatic recognition can be exploited for online applications where context-awareness is essential, such as systems for surgical process monitoring, error detection and intra-operative guidance and assistance of different kinds (e.g. robotic, visual, haptic) [21, 22]. This will have a potentially high impact

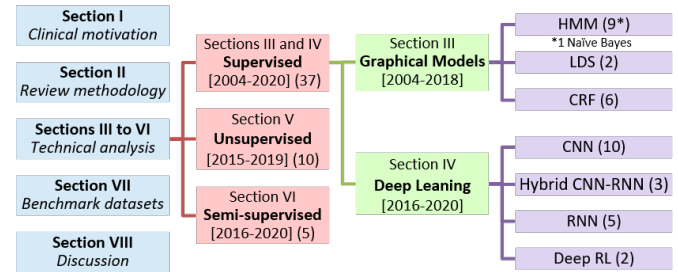


Fig. 3. Review structure. The selected methods were classified based on the learning paradigm (red boxes) and recognition technique (green and purple boxes).

in robotic surgery, where demands for increased safety and effectiveness are particularly strong [23].

In this paper, we review existing methods for automatic recognition of fine-grained gestures in RMIS, illustrating how technical challenges have been addressed, what problems still need to be investigated and what could be the future research directions in the field. Comprehensive reviews on higher-level surgical workflow recognition have recently been reported and are outside the scope of this study [24, 10, 9].

II. METHODOLOGY AND STRUCTURE

We used Google Scholar, Scopus, PubMed, Web of Science and IEEEExplore for searching the literature titles, abstracts and keywords with the query: “(robotic OR robot-assisted OR JIGSAWS) AND (surgery OR surgical) AND (gesture OR fine-grained OR surgeme OR action OR trajectory) AND (segmentation OR recognition OR parsing)”, where JIGSAWS represents the first open source dataset for surgical gesture recognition (see Section VII-A). As Google Scholar only allows full-text searches, the 300 most relevant publications were extracted from the 47,100 retrieved items. Integrating the 5 databases and discarding all duplicates, non-English texts and full conference proceedings, our query delivered a total of 499 matches in October 2020.

We adopt the term “recognition” to refer to the joint problem of “segmentation” and “classification” of surgical gestures, i.e. the simultaneous segmentation of surgical demonstration into distinct motion units and the classification of these units into meaningful action categories. While computer-vision methods for classifying videos with known action boundaries have been successful, temporal segmentation of multiple action instances is still a hard-won problem. Classification of trimmed videos also finds limited applicability in surgical scenarios due to gesture annotation costs, thus research has been fairly limited (only 8 studies were retrieved with our query) and is not presented in this review. We also excluded articles focused on coarser (“higher-level”) granularity and studies that use non-novel recognition methods for skill assessment or surgical automation. After a selection process as illustrated in Fig. 2, we identified a total of 52 publications to be reviewed.

We use the field decomposition shown in Fig. 3 to organise our analysis. Supervised approaches are classified into two groups representing major frameworks for time series analysis and data modelling: probabilistic graphical models (Section

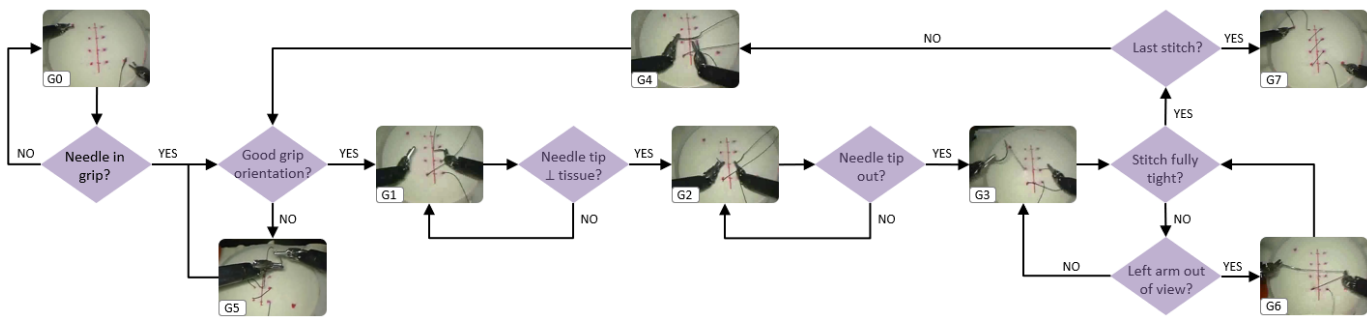


Fig. 4. Example of grammar graph for a suturing demonstration performed with the right hand. Gesture labels: G0 = reaching needle with right hand (START), G1 = positioning needle with right hand, G2 = pushing needle in tissue with right hand, G3 = pulling suture with left hand, G4 = transferring needle from left hand to right hand, G5 = orienting needle on right hand, G6 = tightening suture with right hand, G7 = dropping suture (END). The transition between gestures depends on specific conditions, here represented in purple. Images from JIGSAWS [25].

III) and deep learning (Section IV). We then describe efforts to mitigate the need for manual annotations (Sections V and VI). A subdivision of the methods by data modality is instead highlighted in Tables II, III, IV and V of Section VII.

III. SUPERVISED LEARNING - GRAPHICAL MODELS

Inspired by time series analysis for speech recognition, structured temporal models such as probabilistic graphical models of temporal processes have been extensively used for gesture recognition in robotic surgery [25]. Similarly to the natural language, where syntactic rules govern the generation of words and sentences, the surgical language can be decomposed into surgical motion units at different granularity levels, and probabilistic grammars can be identified to describe the structure of higher-level surgical tasks (Fig. 4).

There is often significant flexibility in choosing the structure and parameterization of a graphical model. Given a sequence of observations and corresponding output labels, generative models such as Hidden Markov Models (HMM) and Linear Dynamical Systems (LDS) explicitly attempt to capture the joint probability distribution over inputs and output. Discriminative models like Conditional Random Fields (CRF), on the other hand, attempt to model the conditional distribution of the output labels given the input data. The advantage with respect to generative models is that dependencies in the input features don't need to be modelled, so that conditional models can have simpler structure [26].

A. Hidden Markov Models

HMMs were first used to classify motion segments obtained by thresholding the kinematic signals, which had limited recognition rate but could be significantly improved with appropriate data transformation including normalization to zero-mean and unit-variance, and Linear Discriminant Analysis (LDA) [31]. Normalization compensates for different units of measure in the kinematic features, while LDA reduces the data dimensionality and model complexity, which is important for robust parameter estimation when only little data is available. LDA's ability to enhance the separation between classes is another desirable property for fine-grained gesture recognition,


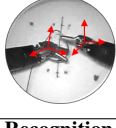
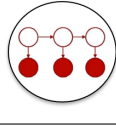
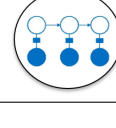
where different classes exhibit similar motion and visual appearance.

Highly discriminative feature extraction is key to the recognition problem, so that even simple frame-wise classifiers, such as the Naïve Bayes classifier, can provide high recognition accuracy [32]. Following [31], LDA was successfully applied to the high-dimensional kinematic signals recorded from the da Vinci surgical robot in combination with a continuous HMM for gesture recognition, composed of basic HMMs representing different surgemes and trained individually [33, 41]. Recognition rate of such models further improves promoting discrimination between sub-gestures (dexeme) rather than entire gestures (surgeme), as this helps to capture the evolution and internal variability of each segment [34]. Automatic derivation of the optimal HMM topology (i.e. the optimal number of hidden states) can also help to discover skill-related and context-specific sub-gestures. While yielding accurate recognition on new trials of seen users, data-derived HMMs may also introduce user-specific sub-gestures and therefore overfit on observed subjects [34].

Alternative data representation and reduction techniques investigated for recognition include Factor-Analyzed Hidden Markov Models (FA-HMM), which yield better recognition compared to standard HMMs with LDA [39]. Sparse Hidden Markov Models (SHMM) were also applied to the recognition of surgical gestures [35, 36]. Available implementations of SHMM represent each surgeme with a different hidden state, whose observations are modelled as sparse linear combinations of atomic motions. The advantage with respect to conventional discrete [7] or Gaussian [6, 32, 34] observation models relies in the sparsity property, which guarantees more robust learning with high-dimensional signals such as the robot kinematics. While each surgical gesture can be associated to a different dictionary that is learnt independently [35], more compact solutions propose unique motion dictionaries shared among all gestures [36].

Finally, interaction forces measured between the robotic tool tip and the environment have been employed as complementary data modality to the robot kinematics [38]. Force information allows one to assign physical interpretation to HMM state observations.

TABLE I
PROBABILISTIC GRAPHICAL MODELS FOR SURGEME RECOGNITION.

Feature extraction	Pros	Cons
 Video Bag-of-features [18] Semantic features [27, 28] CNN [29, 30]	Spatio-temporal information Semantic information Spatio-temporal information, data-derived	Hand-crafted Hand-crafted, localized information Short attention span
 Kinematics LDA [31, 32, 33, 34] Sparse Codings [35, 36, 37] Convolutional filters [28]	Enhanced (sub-)gesture separation Expressive and compact data representation Sub-gesture analysis, temporal information	No temporal information No temporal information Short attention span
Recognition model	Pros	Cons
 Generative Naïve Bayes [32] HMM [31, 33, 34, 35, 36, 38] LDS [39, 40]	Simple model Transition model between (sub-)gestures Continuous hidden state	Frame-wise classification Markov assumption, short attention span Local linearity hypothesis, short attention span
 Discriminative CRF [18, 27, 37, 29, 28, 30]	Segment-to-segment dependencies (MsM-CRF, SC-CRF)	No temporal regularization of adjacent samples (SC-CRF), lack of longer-range dependencies

B. Linear Dynamical Systems

A fundamental drawback of the HMM is its reliance on discrete switching states to model temporal dependencies in the input signals, thus failing to capture more complex correlations in the continuous kinematic and video data. Switched Linear Dynamical Systems (S-LDS) were then proposed to model temporal dependencies between adjacent observations [39]. S-LDSs approximate non-linear systems, such as demonstrations of complex surgical tasks, using a suitable number of linear systems. The data distribution within each gesture is thus modelled through the linear evolution of a continuous hidden state, while the transition between different regimes relies on a discrete switching hidden state. Thanks to improved dynamics modelling, recognition with S-LDS outperformed FA-HMM significantly [39].

S-LDS models can be further extended to capture the dynamics of past observations samples, thus becoming Switched Vector Auto-Regressive models (S-VAR). By imposing specific structure on the model parameters, with certain parameters depending on the discrete latent state and others which are time-invariant, S-VAR models can learn both gesture-dependent and global dynamics in the kinematic signals, resulting in improved recognition accuracy [40].

C. Conditional Random Fields

CRFs were first used to recognize surgical gestures in combination with video data. Inspired by related work on video-based surgeme classification [44], bag of spatio-temporal visual features extracted around space-time interest points or dense trajectories were combined with the kinematic signals in input to a Markov/semi-Markov recognition model (MsM-CRF), which is able to capture frame-level and segment-level dependencies in the input signals [18].

Long-range temporal dependencies were also modelled with the Skip-Chain conditional random field (SC-CRF) [27], incorporating a new skip-length data potential into the CRF energy function to capture sample dependencies at d frames of distance, where d is a skip-length hyper-parameter. In addition, a set of semantic features were extracted from the laparoscopic videos based on the distance of each surgical tool to the closest object in the surgical environment. The authors argued that in constrained environments such as surgical training stations, where the set of objects is known, semantic features encoding object-based information could help to discriminate surgical gestures better than abstract visual features as in [18].

As for HMMs, efforts were spent on improving data representation and feature extraction for CRFs. Inspired by [36], kinematic signals were modelled as sparse linear combinations of atomic motions from an over-complete motion dictionary shared among all gestures [37]. The motion dictionary was learnt jointly with a SC-CRF to capture temporal dependencies and generate more discriminative signal representations, which is valuable to fine-grained analysis. Similarly to [34], surgical kinematics were also analysed at the sub-gesture level. Temporal convolutional filters were used to model the evolution of “latent action primitives” composing each surgical gesture and capture variability within surgemes [28]. The real kinematic data were then replaced with virtual kinematic signals estimated from the video data through a spatial CNN, addressing the situation where kinematic information is available during training, but only video data are available for testing [30]. Convolutional filters were later employed in a similar but more powerful multi-stage framework based on deep convolutional neural networks for spatio-temporal feature extraction from the raw video data (ST-CNN) [29]. Nevertheless, feature extraction strategies as in [28] and [29] are only able to capture




Fig. 5. Deep learning for surgeme recognition: (a) temporal convolutional, (b) recurrent and (c) reinforcement learning strategies. Input kinematic or video sequences ($t = 1 : T, T = \text{sequence length}$) are represented at the bottom and corresponding predictions at the top. Highlighted in green are the input samples that influence the current system predictions (red circles). a: TCN schematic (figure reproduced from [42]). All input samples are processed simultaneously through layers of temporal convolutional filters to deliver predictions over the entire sequence. b: Generic multi-layer RNN ($l = 1 : L, L = \text{number of layers}$). The input samples (x_t) are processed sequentially in time, with each prediction (y_t^l) depending in principle on all the previous samples through memory of previous hidden states (h_{t-1}^l). c: Gesture recognition with reinforcement learning [43]. At each time stamp, predictions are generated over K frames ahead.

local temporal correlations in the observations. Modelling of long-range dependencies is entirely entrusted to the high-level temporal classifier (SC-CRF and sM-CRF respectively). A summary of the presented graphical models for surgeme recognition is reported in Table I.

IV. SUPERVISED LEARNING - DEEP LEARNING

Recently, Deep Neural Networks (DNN) have emerged as the go-to choice for powerful feature extractors that can be learnt automatically from the data, in contrast with hand-crafted filters based on domain-specific prior knowledge [45]. In DNNs, layers of computational units are stacked to generate features at different semantic levels and with increased spatio-temporal receptive fields. This is advantageous for action recognition because long-range spatial and temporal dependencies encoded in the input data can be captured already in early processing stages, easing the subsequent recognition task.

A. Convolutional Neural Networks

Temporal Convolutional Network (TCN) is one of the first deep convolutional models proposed for fine-grained surgeme recognition [19]. TCN uses a hierarchy of temporal convolutions and pooling layers to capture long-range temporal dependencies in the input data (Fig. 5a), thus learning action dependencies in surgical demonstrations better than competing methods and producing well structured predictions. Integration of bilinear pooling with learnable weights further boosts recognition performance incorporating second-order statistics adaptable to the data [46]. As opposed to RNNs, TCN predictions are computed simultaneously for each time stamp and training is much faster.

Built on a similar encoder-decoder backbone, more recent solutions rely instead on two-stream processing: one stream to capture multiscale contextual information and action dependencies in the input sequence, the other to extract low-level information for precise action boundary identification

[47, 48]. Basic temporal convolutions were replaced with a set of Deformable Temporal Residual Modules (TDRM) to capture the temporal variability of human actions and merge the two streams at multiple processing levels [47]. Alternatively, atrous (aka dilated) temporal convolutions and pyramid pooling were used to explicitly generate multi-scale encoding of the input data, that were merged with the low-level local features just before the decoding phase [48]. Both architectures were able to improve consistently both frame-wise and segmental evaluation scores (see Section VII for a description and explanation of the most common evaluation metrics) upon the competing methods.

While pooling operations help to increase the temporal receptive field of a network, they are also responsible for partial loss of fine-grained information and less precise identification of the gesture boundaries. Stacking multiple layers of dilated convolution with increasing dilation factor was proposed as alternative strategy to model long range temporal dependencies between surgical gestures [49]. Gesture predictions can be further refined in a multi-task, multi-stage framework where each stack of dilated convolutions is applied to the output of the previous stage, and the whole system is trained on the sum of all stage losses, as well as on the auxiliary task of surgical skill score prediction [50]. Alternatively, two convolutional stages can be bridged with a self-attention kernel to encode local and global temporal dependencies in the input signals, returning highly accurate predictions with little training cost [51].

All aforementioned architectures represent high-level temporal models that can process efficiently only low-dimensional signals such as kinematic data or pre-encoded visual features. The latter are often derived from the original video frames through spatial encoding techniques [29], albeit unable to capture the dynamics in surgical demonstrations. To address this issue, spatio-temporal feature extraction strategies from short video snippets were proposed using 3D CNNs [52] or Spa-

tial Temporal Graph Convolutional Networks (ST-GCN) [53]. While 3D CNNs operate directly on the video data, achieving superior recognition performance to related spatial [29] and spatio-temporal [29] models, ST-GCN are applied on spatio-temporal skeleton representations of the surgical tools, which are automatically estimated from the video clips. Skeleton-based features are independent of the background and thus robust to scene and context variations, making them suitable for recognition across different surgical tasks or datasets. Clip-based methods like 3D CNNs and ST-CNNs have however limited span of attention, which prevents them from learning long-range temporal dependencies, yielding suboptimal segmental scores. Integration with higher-level convolutional or recurrent models could further boost segmentation performance.

Finally, efforts were spent on improving recognition robustness to inter-surgeon data variability, associated with heterogeneity factors such as surgical style and expertise level, which could hinder generalization performance on unobserved subjects [25] (see Sections VII-C and VIII-B for more details). Robustness to the subject's identity was obtained via adaptive feature normalization [54], where the feature statistics are estimated over data sub-sets recorded from the same surgeon aiming to reduce inter-surgeon data variability and ease the recognition process.

B. Recurrent Neural Networks

Recurrent Neural Networks (RNN) have been used to capture long-term non-linear dynamics in surgical kinematic data [20]. In-depth analysis and comparison of different architectures (simple Bidirectional RNNs, Bidirectional LSTMs, Bidirectional GRUs and Bidirectional Mixed History RNNs) showed that LSTMs and GRUs, which were conceived to alleviate the vanishing gradient problem and can be trained more robustly, are less sensitive to hyper-parameter choice and achieve the best recognition performance [55].

Since predictions are computed sequentially in time, RNNs can naturally handle input signals of different duration and can operate in real time. They also maintain memory mechanisms so that predictions depend in principle on all the previous time stamps (Fig. 5b). Their temporal span of attention, however, is in practice much shorter and they fail to capture multi-scale temporal dependencies in the observations [56]. To address this issue, Multi-Scale Recurrent Neural Network (MS-RNN) [56] extracts multi-scale temporal information from the kinematic signals by means of hierarchical convolutions with wavelets, which are then given as additional inputs to a LSTM for gesture recognition. MS-RNN is trained end-to-end to generate compact descriptors and lessen LSTM susceptibility to overfitting.

Alternative approaches have been recently proposed to regularize training of recurrent architectures and avoid overfitting, such as using data augmentation to improve the network robustness to rotation [57] or multi-task learning to jointly recognize surgical gestures and estimate the progress of the surgical task, in order to introduce ordering relationships and temporal context explicitly into the recognition process [58].

C. Hybrid Convolutional-Recurrent Networks

Temporal Convolutional and Recurrent Network (TricorNet) [59] represents the first hybrid network combining convolutional and recurrent paradigms to recognise surgical gestures. TricorNet features a temporal convolutional encoder similar to [19] to learn local motion evolution, followed by a hierarchy of bidirectional LSTMs to capture long-range temporal dependencies and decode the prediction sequence. Since recurrent cells are in principle able to carry long-term memory of the past observations, the authors argued that hybrid solutions could better model high-level action dependencies than purely convolutional systems with fixed-size local receptive field.

Rather than combining different models sequentially, Fusion-KV [60] performs multi-modal learning with multiple recognition models (CNN-TCN, TCN, LSTM, Random Forest, Support Vector Machine) operating in parallel on three different data sources: kinematics, video and system events collected from the robotic platform (e.g. camera follow, instrument follow, surgeon head in/out of the console, etc). Predictions from all streams are fused through a weighted voting scheme that takes into account recognition performance of each individual stream on the training data. The strengths of different methods and data are thus combined to deliver improved recognition performance. Fusion-KV was later incorporated in a multi-task sequence-to-sequence framework for surgical gesture and instrument trajectory prediction, where gesture labels and instrument positions are estimated over multiple steps in the future given a past observation window [61]. Compact features are first extracted from multiple data sources (full video frames, region-of-interest around surgical tools and kinematic data) and concatenated into a single multi-modal feature vector. Multi-modal features for an observation window in the past and corresponding gesture labels estimated with Fusion-KV are then given as input to an attention-based sequence-to-sequence decoder for gesture prediction, showing 10-step prediction rates comparable to state-of-the-art recognition (0-step prediction) models.

D. Deep Reinforcement Learning

Surgeme recognition can alternatively be modelled as a sequential decision-making process that can be learnt with Reinforcement Learning (RL), as recently proposed by [43]. In their work, an intelligent agent observes each surgical sequence from the beginning. At each time stamp, it selects an appropriate step size and gesture label according to a specific policy, and then moves step-size frames ahead, classifying all the frames in-between as the selected gesture (Fig. 5c). The reward function for policy learning is driven by both classification error and chosen step size, introducing a bias towards larger steps to promote prediction smoothness. While the system was able to achieve competitive results with respect to several state-of-the-art techniques, it sometimes delivered uncertain predictions, especially for rare gestures or at the segmentation boundaries. A tree-search component was therefore introduced to refine the low-confidence predictions from the policy network by looking into the future, improving recognition of the action boundaries [62].

V. UNSUPERVISED LEARNING

Because of the lack of labelled examples, unsupervised approaches are mostly focused on the segmentation sub-problem, i.e. the identification of motion unit boundaries within surgical demonstrations.

Transition State Clustering (TSC) [63], a multilevel generative model based on hierarchical Gaussian Mixture Model (GMM), was applied to surgical kinematic data to identify segmentation points pruning spurious segments generated by inconsistent motion, under the assumption that action order is partially consistent across surgical demonstrations. The kinematic signals were then integrated with visual features extracted with a pre-trained CNN, improving segmentation performance [64]. TCN methodology was later extended using Dense Convolutional Encoder-Decoder Network (DCED-Net) for unsupervised feature extraction from surgical videos [65]. Trained on the task of image reconstruction, DCED-Net extracts discriminative visual features using specific convolutional blocks, called Dense Blocks, to reduce information loss during dimensionality reduction. Segmentation results were further improved with a majority vote strategy to eliminate spurious transition points exploiting both kinematic and visual information, and an iterative algorithm to merge similar adjacent segments [65].

More complex generative models were later proposed to describe the joint probability distribution of observed kinematic data and latent action sequences. Surgical demonstrations were for instance modelled as Markov decision processes learnt through imitation learning with options [66], where complex policies are broken down into simpler low-level primitives called options. The model parameters were estimated with a policy-gradient algorithm aimed to maximize the likelihood of the available data and simultaneously segment them into different motion primitives. Alternative solutions such PRISM (PRocedure Identification with a Segmental Mixture model) [67] were developed under the assumption that all demonstrations are described by a common latent sequence of actions. While generalizable to non-Markov processes, PRISM performance degrades significantly with increased action ordering variability.

Other clustering techniques group the data samples based on feature similarities and temporal constraints. Spectral clustering was tailored to temporal segmentation through specific regularization terms considering local dependencies in the motion data, where consecutive samples are likely to belong to the same action class [68]. Good results were obtained with bottom-up clustering, where the sequential nature of motion signals can be exploited by merging neighbouring segments according to various merging and stopping criteria [69, 70]. In [69], surgical demonstrations were first segmented into a set of very fine-grained motion primitives (dexeme) based on the persistence and dissimilarity of consecutive segments between critical points in the kinematic trajectories. A descriptive signature was then created to represent each dexeme and used for classification into surgeme categories. Starting instead from the finest possible segmentation of the kinematic signals [70], neighboring segments were merged

iteratively using multiple compaction scores based on different distance measures (similarity between PCA results, distance between segment centres, dynamic time warping distance) and fuzzy membership scores to model gradual transitions between surgical gestures.

Rather than grouping similar data samples or segments together, opposite approaches search for optimal segmentation criteria. Spatio-temporal and variance properties of the kinematic data can be exploited to identify candidate segmentation points around the peaks of distance and variance profiles computed from the original signals [71]. The method's disadvantage, however, is in the tedious tuning of threshold values.

Finally, zero-shot learning (ZSL) approaches have also shown potential for surgical gesture recognition [72]. ZSL consists in learning a set of high-level semantic attributes describing instances of specific action classes and exploiting them to recognize instances of new classes, provided with their attribute description (signature) but no labelled examples. [72] modelled surgical gestures with dynamic signatures represented by a set of attributes that change in time according to action-specific rules. Given a surgical demonstration, the presence or absence of those attributes were detected on each video frame and compared to the set of signatures for action matching, achieving better results than fully supervised unstructured baseline models.

VI. SEMI-SUPERVISED LEARNING

While unsupervised approaches mitigate significantly labelling costs and issues, as labels are only necessary for model testing, they still show a significant gap in performance with respect to the supervised methods. A midway solution is represented by semi-supervised learning [73, 74, 75, 76], where a small pool of annotated demonstrations is used for model training. Such demonstrations can be used to transfer gesture labels to a large set of unlabelled observations previously aligned with DTW [73], or to initialize clustering of unlabelled data and avoid tedious parameter tuning [74]. These systems are however highly sensitive to large data variability, showing rapid drop in recognition accuracy with increased heterogeneity in terms of gesture ordering or surgical skill level respectively.

Promising results were obtained with RNNs within a two-phase pipeline comprising unsupervised representation learning and semi-supervised recognition [75]. An RNN-based generative model was first trained to learn the full joint distribution over the kinematic signals, an auxiliary task that needs no additional annotations but can extract informative features from the input data. These features were then used by a bidirectional LSTM for gesture recognition, which was only trained on a very small number of labelled demonstrations. Decoupling between representation learning and gesture recognition was later solved via semi-supervised interleaved training [77]. Meaningful visual embeddings were learnt using triplet loss to pull together images in the same action segment and push away samples from different segments, where action pseudo-labels for unlabeled training sequences were iteratively



Fig. 6. Snapshots from suturing, knot-tying and needle passing demonstrations in JIGSAWS [25].

predicted by the current estimate of a recurrent network trained via cross-entropy loss. Thanks to robust representation learning, the gap in performance with respect to equivalent fully supervised systems was limited.

Semi-supervised learning was also explored in combination with transfer learning between surgical tasks [76]. One-dimensional projections of the kinematic signals derived from their Self-Similarity Matrix (SSM), which preserves motion patterns in the data, were used to learn segmentation policies of specific surgical tasks (e.g. knot-tying) from annotated demonstrations of other tasks (e.g. suturing). In contrast to [73, 74, 75], the class of the generated segments was not predicted.

VII. BENCHMARK DATASETS AND EVALUATION PROTOCOLS

Early techniques for surgical gesture segmentation and skill assessment were validated on study-specific data using different evaluation metrics, hindering comparative research and progress tracking in the field. To address this issue, the JHU-ISI Gesture and Skill Assessment Working Set (JIGSAWS) [13, 25] was released as first public dataset containing video and kinematic recordings of robotic surgical demonstrations, as well as gesture and skill annotations.

A. The JIGSAWS dataset

JIGSAWS contains synchronized video and kinematic data captured from the da Vinci Surgical System (dVSS) during the execution of elementary surgical tasks, including 39 suturing, 36 knot-tying and 20 needle-passing demonstrations (Fig. 6). All trials are performed on a training phantom by eight surgeons with different robotic surgical experience: experts reported more than 100 hours of training, intermediate between 10 and 100, and novice less than 10 [13].

Video data: Stereo videos from the da Vinci endoscopic cameras are recorded at 30Hz and at a resolution of 640 x 480. Demonstrations are all observed from the same point of view, as camera movement and zoom are completely absent. Calibration parameters are not provided.

Kinematic data: Kinematic trajectories of the two master tool manipulators (MTMs) and patient side manipulators (PSMs) are recorded in sync with the video data. The pose of each arm is represented by a local reference frame linked to the end-effector, and their motion is described by 19 kinematic variables that include Cartesian positions, a rotation matrix, linear velocities, angular velocities and a gripper angle. Kinematic features of different manipulators are all referred to a common coordinate system.

Annotations: A common dictionary comprised of 15 action classes is used to describe all three surgical tasks in the dataset. Ground truth annotations for each trial include the sequence of gestures, boundary timestamps and the demonstrator's expertise level.

B. Experimental setup

Two standard cross-validation setups have been used to analyze performance and generalization capability of the available recognition systems:

- *Leave-one-supertrial-out (LOSO):* five validation folds are created taking one different trial from each user. The LOSO scheme can be used to evaluate model generalization to new trials performed by known surgeons.
- *Leave-one-user-out (LOUO):* eight validation folds are created taking all trials performed by the same user. The LOUO scheme can be used to evaluate model generalization to new and unknown surgeons, which is more challenging due to subject-specific style variability.

Notably, deep learning schemes typically limit extensive cross-validation due to high computational requirements and long training time and most recent studies have used only LOUO to discriminate different methods. LOUO is considered as the gold standard because it does not allow approaches to overfit to surgeon-specific features, which is also important for surgical skill assessment where networks should not learn to classify based on the surgeon but on the execution quality.

C. Evaluation metrics

Evaluation metrics for action recognition can be broadly divided into three categories highlighting different strengths and weaknesses of the proposed methods:

Frame-wise metrics: standard classification metrics such as average accuracy, precision, recall and F1-score, have been widely used to assess frame-wise recognition performance. These were however conceived for non-sequential data and are not sufficient for an exhaustive analysis of temporal predictions and robust method comparison. Predictions with similar accuracy might in fact show large qualitative differences in terms of action ordering and over-segmentation (i.e. prediction of many insignificant action boundaries).

Segmental metrics: the two most common segmental metrics, which have been specifically designed to evaluate the temporal structure of action predictions, are the *edit score* and the *segmental F1 score with threshold k/100 (F1@k)*. The edit score evaluates action ordering but not their timing, penalizing out-of-order predictions and over-segmentation. It is computed as the Levenshtein distance between true and predicted label sequences, either normalized at the dataset level (Edit), where the common normalization factor is the maximum number of segments in any ground-truth sequence, or at the instance level (Edit*), where the normalization factor for each sequence is the maximum number of segments in either the ground-truth or the prediction [29]. F1@k also penalizes over-segmentation, but it additionally evaluates the temporal overlap between predictions and ground truth segments with reduced sensitivity

TABLE II
AVERAGE CROSS-VALIDATION SCORES (ACCURACY AND EDIT*) ON THE SUTURING DEMONSTRATIONS OF JIGSAWS - GRAPHICAL MODELS.
GREEN = KINEMATICS, YELLOW = VIDEO, RED = KINEMATICS AND VIDEO.

Method	Year	LOSO		LOOU		Data
		Acc	Edit*	Acc	Edit*	
<i>Composite-HMM</i> [34, 25]	2009	82.2	-	74.0	-	PSMs - temporal concatenation, LDA
<i>KSVD-SHMM</i> [35, 25]	2012	83.4	-	73.5	-	PSMs
<i>SDSDL</i> [36]	2015	86.3	-	78.7	-	PSMs & MTMs - mean-variance normalization, PCA
<i>SDL+SC-CRF</i> [37]	2018	86.2	75.5	78.5	58.0	PSMs & MTMs - mean-variance normalization, PCA
<i>MsM-CRF</i> [18, 25]	2013	82.0	-	67.8	-	PSMs & MTMs - SVM scores (unary and pairwise)
		84.4	-	77.3	-	RGB frames - BoF around dense trajectories, SVM scores (unary and pairwise)
		85.1	-	79.1	-	kinematics: PSMs & MTMs - SVM scores (pairwise) video: RGB frames - BoF around dense trajectories, SVM scores (unary)
<i>STCNN+sM-CRF</i> [29]	2016	-	-	74.2	66.6	RGB frames & Motion History Images
<i>SC-CRF</i> [27, 25]	2015	85.2	-	81.7	-	PSMs (positions, linear velocities, gripper angle)
		85.0	-	81.6	-	kinematics: PSMs (positions, linear velocities, gripper angle) video: Semantic features (distance tool-object)
<i>LC-SC-CRF</i> [28]	2016	-	-	83.5	76.9	kinematics: PSMs (positions, linear velocities, gripper angle) video: Semantic features (distance tool-object)
<i>CNN + LC-SC-CRF</i> [30]	2016	-	-	76.6	-	RGB frames

PSMs = Patient Side Manipulators, MTMs = Master Tool Manipulators, LDA = Linear Discriminative Analysis, PCA = Principal Component Analysis, SVM = Support Vector Machine, BoF = Bag of Features

to slight temporal shifts, compensating for annotation noise around the segment boundaries. F1@k is obtained by computing the intersection over union (IoU) overlap score between each predicted segment and the corresponding ground truth segment of the same class. That prediction is considered as true positive (TP) if the IoU is above a threshold $\tau = k/100$, otherwise it is a false positive (FP). TPs and FPs are then used to compute the final F1 score [42].

Normalized mutual information (NMI) has been used less frequently to measure alignment between two action sequences in unsupervised settings [64], as this score is not affected by permutations of cluster labels.

Unsupervised metrics: instead of comparing predictions to their ground truth, unsupervised metrics such as the silhouette score (SS) have been used to measure the compactness of action clusters generated by unsupervised recognition methods [64].

Other metrics used in the reviewed articles include Segmentation Accuracy (Seg-Acc) [65], Temporal Structure Score (TSS) [67] and detection scores based on temporal tolerance windows [66, 71]. Formulas and pseudocode of the most common evaluation metrics are reported in Appendix A.

D. Results

Results reported on the suturing demonstrations of JIGSAWS are grouped into supervised graphical models (Table II), supervised deep learning (Table III), semi-supervised (Table V) and unsupervised (Table IV) methods. Despite referring to the same dataset, attention should be paid to the input

data when comparing different methods, especially when the performance gap is limited. Improvements in segmental scores, for example, could be partially attributed to the adoption of lower sampling rate rather than to the method itself. Small accuracy gaps could also be linked to specific data pre-processing (filtering, normalization, dimensionality reduction) and kinematic and visual channel selection. Given the small dataset size and similarity between certain results, more rigorous comparative analysis could be supported by statistical testing to highlight significant performance gaps. Evaluation on larger surgical datasets should however be the main goal of future researches.

VIII. DISCUSSION AND OPEN QUESTIONS

Automatic recognition of surgical gestures is a complex task and a number of considerations affect the current state of the art in available methods as detailed below.

A. Robust temporal modelling

Our comparative analysis on the JIGSAWS dataset revealed that recent deep learning models are able to capture complex temporal dependencies in surgical motion, leading to state-of-the-art recognition performance. In particular, hierarchical temporal convolutions and recurrent modules can model multi-scale temporal information for robust classification and boundary identification, especially when integrated with temporal attention mechanisms [51, 61]. Only a few studies, however, have considered the continuous nature of surgical motion and explicitly modelled gradual transitions between action

TABLE III
AVERAGE LOUO CROSS-VALIDATION SCORES (ACCURACY, EDIT*, EDIT AND F1@10) ON THE SUTURING TRIALS OF JIGSAWS - DEEP LEARNING.
GREEN = KINEMATICS, YELLOW = VIDEO, RED = KINEMATICS AND VIDEO.

Method	Year	Acc	Edit*	Edit	F1@10	Data
<i>LSTM</i> [20, 19]	2016	80.5	75.3	80.2	-	PSMs (positions, linear velocities, gripper angle) - 5Hz, mean-variance norm
<i>BiRNN</i> [55]	2019	82.1	78.9	82.7	-	PSMs (linear velocities, angular velocities, gripper angle) - 5Hz
<i>BiMIST</i> [55]	2019	84.7	76.4	74.3	-	PSMs (linear velocities, angular velocities, gripper angle) - 5Hz
<i>BiLSTM</i> [55]	2019	84.7	88.1	91.6	-	PSMs (linear velocities, angular velocities, gripper angle) - 5Hz
<i>BiGRU</i> [55]	2019	84.8	88.5	91.6	-	PSMs (linear velocities, angular velocities, gripper angle) - 5Hz
<i>APc</i> [58]	2020	85.5	85.3	-	-	PSMs (positions, linear velocities, gripper angle) - 5Hz, mean-variance norm
<i>MS-RNN</i> [56]	2019	90.2	-	89.5	-	PSMs & MTMs - 5Hz
<i>ST-GCN</i> [53]	2020	67.9	-	-	-	RGB frames - 90-frame snippets, ResNet50 pose estimation
<i>S-CNN</i> [29, 19]	2016	74.0	37.7	-	-	RGB frames & Difference Images
<i>DenseNet + Adaptive Norm</i> [54]	2020	75.8	-	-	-	PSMs & MTMs
<i>ST-CNN</i> [29, 19]	2016	77.7	68.0	-	-	RGB frames & Difference Images
<i>TCN</i> [19]	2016	79.6	85.8	-	-	PSMs (positions, linear velocities, gripper angle) - 10Hz
		81.4	83.1	-	-	RGB frames & Difference Images - 10Hz, S-CNN [29] feature extraction
<i>TDNN</i> [49]	2019	80.4	-	-	-	PSMs & MTMs
<i>TCN + RL</i> [43]	2016	82.1	87.9	-	91.1	PSMs (positions, linear velocities, gripper angle)
		81.4	88.0	-	92.0	RGB frames & Difference Images - 10Hz, S-CNN [29] feature extraction
<i>TCN + RL + Tree Search</i> [62]	2020	81.7	88.5	-	92.7	RGB frames & Difference Images - 10Hz, S-CNN [29] feature extraction
<i>MTL-VF</i> [50]	2020	82.1	86.6	-	90.6	RGB frames - C3D [78] feature extraction
<i>TCED-Bd</i> [46]	2019	82.2	87.7	-	91.4	RGB frames & Difference Images - 10Hz, S-CNN [29] feature extraction
<i>TricorNet</i> [59]	2017	82.9	86.8	-	-	RGB frames & Difference Images - 10Hz, S-CNN [29] feature extraction
<i>3D-CNN</i> [52]	2019	84.0	80.7	-	87.2	RGB frames - 5Hz, 16-frame snippets, scale jittering, corner cropping
<i>TDRN</i> [47]	2018	84.6	90.2	-	92.9	RGB frames & Difference Images - 10Hz, S-CNN [29] feature extraction
<i>AT-Net</i> [48]	2019	-	91.5	-	94.6	RGB frames & Difference Images - 10Hz, S-CNN [29] feature extraction
<i>Symm dilation + attention</i> [51]	2020	90.1	89.9	-	92.5	RGB frames & Difference Images - 10Hz, S-CNN [29] feature extraction
<i>Fusion-KV</i> [60]	2020	86.3	87.2	-	-	kinematics: PSMs (positions, linear velocities, Euler angles, angular velocities) video: RGB frames & Difference Images - VGG16 feature extraction
<i>Fusion-KV+LSTM (pred @Is)</i> [61]	2020	84.3	-	-	-	kinematics: PSMs for past observation window video: RGB frames for past observation window

PSMs = Patient Side Manipulators, MTMs = Master Tool Manipulators

segments [70, 66, 76, 71]. Transition handling has been alternatively addressed by under-penalizing classification errors to take into account uncertainties in ground truth annotations [73]. Promising results have also been achieved with reinforcement learning, but related work in this area is still very limited and requires deeper analysis.

As for the lookahead time, the majority of reported methods relies on a variable number of future frames for robust recognition. While real-time performance or ability to predict future samples are not always needed and pursued, further efforts in this direction could broaden the applicability of such algorithms within the surgical theatre, supporting systems for surgical process monitoring and anomaly detection and prevention [21, 61].

B. Discriminative feature extraction

Robust analysis and understanding of fine-grained surgical gestures is especially complicated by the high variability of data with same action class. Gestures represent generic blocks of motion which can be shared across different tasks, phases and procedures, and significant variation comes from such context including the instrument types and the patient-specific environment. Other variability factors comprise the surgeon's skill level, experience and individual style, which can significantly alter the duration, kinematics and order of actions in a subject-specific manner [25]. For example, novice surgeons tend to make different mistakes and perform multiple gesture attempts and adjustment motions, while expert demonstrations are generally faster and smoother. Even among

TABLE IV
AVERAGE SCORES ON THE SUTURING DEMONSTRATIONS OF JIGSAWS - UNSUPERVISED METHODS.
GREEN = KINEMATICS, YELLOW = VIDEO, RED = KINEMATICS AND VIDEO.
**LEAVE-ONE-OUT ESTIMATE.

Method	Year	Scores	Metrics	Data
<i>TSSC-LR</i> [68]	2017	49.7 / 32.8	Acc / Edit*	PSMs & MTMs - 6 Hz, unit-norm normalization
<i>TSC-DL</i> [64]	2016	0.52** / 0.31**	SS / NMI	PSMs & MTMs - 10 Hz, temporal concatenation
		0.58** / 0.16**	SS / NMI	RGB frames - 10 Hz, cropping, rescaling, channel norm, VGG feature extraction, PCA, temporal concatenation
		0.73** / 0.63**	SS / NMI	kinematics: PSMs & MTMs - 10 Hz, temporal concatenation video: RGB frames - 10 Hz, cropping, rescaling, channel norm, VGG feature extraction, PCA, temporal concatenation
<i>TSC-DCED-Net</i> [65]	2018	0.57** / 61.9**	NMI / Seg-Acc	kinematics: PSMs & MTMs - wavelet transform video: RGB frames - DCED-Net feature extraction
<i>DDO</i> [66]	2017	72.0	Fraction of correct boundaries (300ms tolerance)	kinematics: PSMs (positions, gripper angle) video: RGB frames - rescaling
<i>PRISM</i> [67]	2019	46.7 / 33.3	TSS / NMI (expert only)	MTMs - 10 Hz, PCA, temporal concat, mean-variance norm
<i>Soft-UGS</i> [70]	2017	0.74 / 0.71 / 0.72 / 0.65	Recall / Precision / F1 / SS	PSMs
<i>SpaceTime + Variance</i> [71]	2019	65.2 / 92.0 / 75.4	Recall / Precision / F1	PSMs (positions, orientations) - smoothing
<i>Zero-shot</i> [72]	2019	56.6 / 61.7	Acc / Edit*	RGB frames

PSMs = Patient Side Manipulators, MTMs = Master Tool Manipulators, PCA = Principal Component Analysis, Acc = accuracy, SS = Silhouette Score, NMI = Normalized Mutual Information, Seg-Acc = Segmentation Accuracy, TSS = Temporal Structure Score

TABLE V
AVERAGE SCORES ON THE SUTURING DEMONSTRATIONS OF JIGSAWS - SEMI-SUPERVISED METHODS.
GT = THE NUMBER OF LABELLED TRIALS IN THE TRAINING SET.
(*) AVERAGE OVER 10 RUNS. (**) AVERAGE OVER EXHAUSTIVE SPLITS.

Method	Year	GT	Scores	Metrics	Data
<i>GMM0</i> [74]	2019	3	0.59 / 0.46 / 0.50	Acc / NMI / SS (revised annotations)	PSMs (positions, orientations, gripper angle, euclidean distance between PSMs) - 10 Hz, smoothing, mean-variance normalization, temporal concatenation
<i>DTW + label transfer</i> [73]	2016	5	0.67 / 0.59 / 0.69 / 0.65 (*)	Acc / Precision / Recall / F1 (1s tolerance)	PSMs & MTMs - temporal concatenation, SDAE feature extraction, temporal derivative, smoothing, downsampling
		20	0.68 / 0.63 / 0.79 / 0.70 (*)	Acc / Precision / Recall / F1 (1s tolerance)	PSMs & MTMs - temporal concatenation, SDAE feature extraction, temporal derivative, smoothing, downsampling
<i>LSTM</i> [75]	2019	1	70.4 / 78.2 (**)	Acc / Edit	PSMs (linear velocities, angular velocities, gripper angle) - 5Hz
		7	82.4 / 89.0 (**)	Acc / Edit	PSMs (linear velocities, angular velocities, gripper angle) - 5Hz

PSMs = Patient Side Manipulators, MTMs = Master Tool Manipulators, SDAE = Stacked Denoising Autoencoder, Acc = accuracy, NMI = Normalized Mutual Information, SS = Silhouette Score

experts, differences might be notable in terms of surgical style, which are then reported among novice surgeons mimicking their tutors.

Similarity of data with different action class is also problematic because the amount of available information to discriminate between fine-grained gestures is limited. Considering the suturing task, for example, the same surgical tools and objects (a needle and a thread) are operated within the same environment during the whole task execution, conferring similar visual appearance to all the constituent segments (Fig. 4). Even the kinematic patterns of different gestures are sometimes very similar (e.g. when a surgeon attempts to push a needle through the tissue multiple times, the gestures “positioning needle tip

on insertion point” and “pushing needle through the tissue” can be easily confused). Additional data such as tool usage information and video of the whole operating room, which can help to discriminate high-level surgical phases [79], are instead redundant for fine-grained analysis.

Discriminative feature extraction from surgical video and kinematics has thus been an important research focus area. Most of the reviewed methods, however, analyze all data samples independently or consider local neighborhoods in time. Only few studies integrate discriminative feature extraction and high-level temporal modelling through end-to-end training [37], thus embedding long-range temporal information in the low-level features. Nevertheless, end-to-end training

of large architectures such as CNN-LSTMs, often applied to the analysis of high-level surgical workflow [80, 81], is problematic on very small datasets like JIGSAWS.

Discriminative information could also derive from optical flow motion features, which have aided surgical phase recognition [82] and action classification [83, 84], or from semantic analysis of the video data [27]. Semantic visual features could be potentially used to represent information about the surgical environment and manipulated objects that are not explicitly captured by abstract video representations or by the kinematic data, such as identification and localization of anatomical structures, measures of tissue deformation, pose or state of manipulated objects, presence of bleeding, smoke or other surgical tools operated by a surgical assistant. In addition, semantic information could also be gained indirectly through multi-task or transfer learning with appropriate auxiliary tasks. Multi-task learning has been extensively explored at coarse granularity levels, where systems detecting surgical tools or predicting the remaining time of surgery have also shown improved phase recognition capabilities [79, 85, 86]. Gesture recognition networks have only recently been enhanced with parallel branches for progress [58], skill scores [50] or surgical tool trajectory [61] prediction, and other auxiliary tasks could potentially be conceived to aid recognition at fine granularity (e.g. tasks based on semantic image segmentation, object tracking or anomaly detection).

C. Multi-modal data integration

The integration of multi-modal data could also play an important role in improving the performance of available recognition systems, as different sensor modalities often measure complementary information. Kinematic data, for example, describe poses and velocities of the surgical tools in the three-dimensional space, whereas visual features carry implicit complementary information about the surgical environment and instrument-tissue interactions. The combination of kinematic and visual features has in fact yielded consistent improvement over the two individual modalities [18, 27, 64, 60]. Force and torque signatures have also been used for motion decomposition and skill assessment [38, 5, 6], even though this information may not be readily available from system manufacturers. This lack of data access has also meant that multi-modal data integration for surgerme recognition has not been investigated thoroughly.

Kinematic and visual information have been combined at the input level [64, 65, 27], at intermediate level [61] or in the final prediction [60] and pruning stages [64, 65]. Concatenation of multi-modal data at low processing levels might be suboptimal and inefficient, as the two sensing modalities are not directly comparable due to differences in cardinality, semantics and stochasticity [64] (e.g. visual features are typically much larger than kinematic samples and more noisy). A splicing weight between kinematic and visual features has been used to find appropriate importance levels in different surgical tasks [65], but it entails extensive tuning and can not adapt dynamically to environmental changes. Class-specific weights for multi-modal late fusion [60] are also unable to capture intra-gesture

dynamic properties. Improved integration could be achieved with flexible systems that seek timely multi-modal cooperation at multiple processing levels, with adaptive weighting or knowledge transfer from the strongest to the weakest modality.

D. Translational research

Despite recent advances, the field and applicability to real surgical scenarios have been greatly hindered by the lack of large and diverse datasets of annotated demonstrations, which are essential for robust training of modern recognition systems based on machine learning and deep learning. Current methods need testing on varied surgical tasks and procedures, as well as on data from real interventions with complex environments, blood, specularities, camera motions, illumination changes, occlusions and higher variability in motion and gesture ordering [60, 87, 88]. International collaborations for clinical data collection and sharing would not only accelerate the generation of such datasets, but also allow the representation of variations in human physiology linked to ethnicity as well as a wider spectrum of surgical techniques. Development of a common fine-grained surgical language with standardized decomposition criteria for different procedures and target anatomical structures is then needed in order to integrate information from multiple data sources.

It is however not trivial to define a comprehensive gesture dictionary and meaningful segmentation criteria which are consistent in the presence of unconstrained movement and generalizable to a variety of anatomies and surgical styles. The characterization of “surgical gesture” is in fact an open problem itself, in particular for bimanual operations. Datasets like JIGSAWS regard fine-grained gestures as short action segments performed by either of the two arms, while the motion of the other arm is ignored. This strategy allows for faster labelling, but suffers from ambiguities on complex demonstrations where the two arms might cooperate or perform different actions at the same time. Such ambiguity can be resolved with finer-grained analysis and multi-label segmentation with separate classification of right and left arm motions. Multi-label analysis is burdensome but more precise and can additionally account for action combinations [14].

The optimal decomposition strategy could also depend on application requirements. A robotic collaborative assistant might not need to recognize highly detailed surgical motions and localize precise action boundaries in order to provide the surgeon with timely support. On the other side, robust understanding of task workflow, bimanual cooperation and action transitions are required to achieve surgical automation and successfully execute complex action sequences in real case scenarios. As for surgical skill assessment, the proper decomposition strategy could depend on the annotators’ evaluation technique. Identification of known skill-related gestures and failure modes while ignoring style-related variability might be beneficial.

Even when the action dictionary is well defined, however, manual labelling of surgical gestures is extremely costly, time consuming, prone to errors and inconsistencies. To mitigate data requirements, the recognition problem can be tackled

in a semi-supervised or unsupervised manner, where labels are only necessary for model testing. Even if not yet able to outperform supervised methods, these have a lot of potential to capitalise on large amounts of unlabelled data [89].

When collection of clinical records is problematic, one of the most viable strategies to generate large-scale surgical data is represented by virtual reality simulators [90], which can nowadays reproduce real surgical environments with impressive realism, as well as provide complementary information (e.g. surgical tool kinematics or semantic segmentation of the 3D scene) to aid recognition. Further efforts could then be directed towards new strategies for knowledge transfer between surgical environments (simulation vs real surgery), as well as between surgical tasks [76, 30] or surgical robots [91], in order to exploit cross-domain information to improve generalization performance.

Finally, an emerging and closely related research sub-field is represented by the detection and forecast of surgical errors in fine-grained gestures. It has been recently shown that common gesture-specific errors can be identified in real-time if provided with corresponding real-time gesture labels [21]. Public availability of error annotations [21] will hopefully stimulate future research in this relatively unexplored field, which is fundamental for several CAI implementations including workflow monitoring and surgical automation.

IX. CONCLUSION

In this paper, we have presented a comprehensive analysis and critical review on the recognition of surgical gestures from intraoperatively collected information. Automatic recognition of surgical gestures is challenging due to the variability, complexity and fine-grained structure of surgical motion. The digitization of surgical video and increased availability of sensor information in robotic instrumentation has driven the emergence and evolution data-driven recognition models. While increasing the robustness of automatic recognition, deep-learning-based models still face many challenges compounded by the limited availability and small size of surgical training datasets that do not capture the domain variability of clinical practice across diverse and sometimes procedure specific activities. The field would also benefit from standardized definition of the segmentation criteria for different applications and surgical procedures to support comparative analysis and research across multiple datasets. Technically, unsupervised and semi-supervised approaches, which do not require large amounts of ground truth labels, are important directions for future research in order to enable scaling of methods across surgical procedures and allow adaptability to the evolution of surgical technique and instrumentation.

APPENDIX A

FORMULAS OF THE MOST COMMON EVALUATION METRICS

Accuracy (Acc): Given a sequence of length N ,

$$Acc = \frac{Nc}{N} * 100, \quad 0 < Acc < 100$$

where Nc is the number of correctly labelled frames.

Edit score (Edit): Given a sequence of labels τ and corresponding ground truth γ :

$$Edit = 1 - \frac{D(\tau, \gamma)}{N} * 100, \quad 0 < Edit < 100$$

where D is the Levenshtein distance between τ and γ , N is either the maximum number of segments in τ or γ (instance level normalization), or in any ground truth sequence (dataset level normalization).

Segmental F1 score with threshold $k/100$ (F1@k): Given a sequence of predicted segments PS and corresponding ground truth segments TS :

$$TP = FP = 0$$

$$N_{used} = 0$$

for each predicted segment PS_j in PS :

c = label of PS_j

TS^c = list of TS segments of class c

$IoU_{all} = IoU(PS_j, TS^c)$

$IoU_i = \max(IoU_{all})$

TS_i^c = segment in TS^c corresponding to $\max(IoU_{all})$

if $IoU_i > k/100$ and TS_i^c not used:

$$TP = TP + 1$$

Flag TS_i^c as used

$$N_{used} = N_{used} + 1$$

else:

$$FP = FP + 1$$

$$FN = \text{length}(TS) - N_{used}$$

$$\text{precision} = TP / (TP + FP)$$

$$\text{recall} = TP / (TP + FN)$$

$$F1@k = 2 * (\text{precision} * \text{recall}) / (\text{precision} + \text{recall}) * 100$$

where IoU is the intersection over union overlap score, TP are true positives, FP false positives and FN false negatives. Detailed implementation is publicly available [92].

Normalized Mutual Information (NMI): Given two sequences of labels τ and γ :

$$NMI(\tau, \gamma) = \frac{I(\tau, \gamma)}{\sqrt{H(\tau)H(\gamma)}}, \quad 0 < NMI(\tau, \gamma) < 1$$

where I is the mutual information and H the entropy.

Silhouette Score (SS): The SS for a single data sample i is defined as:

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}, \quad -1 < s(i) < 1$$

where $a(i)$ is the distance (e.g. Euclidean distance) between that sample and the mean of its own cluster, while $b(i)$ is the distance between that sample and the mean of the nearest cluster it is not a part of. SS is the average Silhouette score over all N_s samples:

$$SS = \sum_{i=1}^{N_s} s(i), \quad -1 < SS < 1$$

SS can be normalized to $0 < SS < 1$.

REFERENCES

- [1] K. Moorthy et al. Dexterity enhancement with robotic surgery. *Surgical Endoscopy and Other Interventional Techniques*, 2004.
- [2] K. H. Sheetz et al. Trends in the adoption of robotic surgery for common surgical procedures. *JAMA Network Open*, 2020.
- [3] F. Chadebecq et al. Computer vision in the surgical operating room. *Visceral Medicine*, 2020.
- [4] L. Maier-Hein et al. Surgical data science for next-generation interventions. *Nature Biomedical Engineering*, 2017.
- [5] J. Rosen et al. Markov Modeling of Minimally Invasive Surgery Based on Tool / Tissue Interaction and Force / Torque Signatures for Evaluating Surgical Skills. *IEEE transactions on Biomedical Engineering*, 2001.
- [6] J. Rosen et al. Task Decomposition of Laparoscopic Surgery for Objective Evaluation of Surgical Residents' Learning Curve Using Hidden Markov Model. *Computer Aided Surgery*, 2002.
- [7] C. E. Reiley and G. D. Hager. Task versus Subtask Surgical Skill Evaluation of Robotic Minimally Invasive Surgery. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2009.
- [8] N. Vaughan et al. An overview of self-adaptive technologies within virtual reality training. *Computer Science Review*, 2016.
- [9] F. Lalys and P. Jannin. Surgical process modelling: A review. *International Journal of Computer Assisted Radiology and Surgery*, 2014.
- [10] O. Dergachyova. *Knowledge-based support for surgical workflow analysis and recognition*. PhD thesis, 2017.
- [11] S. S. Vedula et al. Analysis of the Structure of Surgical Activity for a Suturing and Knot-Tying Task. *PLOS ONE*, 2016.
- [12] G. De Novi et al. Event-driven Surgical Gesture Segmentation and Task Recognition for Ocular Trauma Simulation. *Intelligent Environments (Workshops)*, 2012.
- [13] Y. Gao et al. 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*, 2014.
- [14] J. Chen et al. Use of automated performance metrics to measure surgeon performance during robotic vesicourethral anastomosis and methodical development of a training tutorial. *The Journal of urology*, 2018.
- [15] M. Maktabi and T. Neumuth. Online time and resource management based on surgical workflow time series analysis. *International Journal of Computer Assisted Radiology and Surgery*, 2017.
- [16] T. N. Judkins et al. Objective evaluation of expert and novice performance during robotic surgical training tasks. *Surgical Endoscopy*, 2009.
- [17] D. Nagy. A DVRK-based Framework for Surgical Subtask Automation. *Acta Polytechnica Hungarica*, 2019.
- [18] L. Tao et al. Surgical gesture segmentation and recognition. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2013.
- [19] C. Lea et al. Temporal convolutional networks: A unified approach to action segmentation. *European Conference on Computer Vision (ECCV)*, 2016.
- [20] R. DiPietro et al. Recognizing Surgical Activities with Recurrent Neural Networks. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2016.
- [21] M. S. Yasar and H. Alemzadeh. Real-time context-aware detection of unsafe events in robot-assisted surgery. *arXiv preprint arXiv:2005.03611*, 2020.
- [22] G. De Rossi et al. Cognitive robotic architecture for semi-autonomous execution of manipulation tasks in a surgical environment. *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [23] R. H. Bell. Why Johnny cannot operate. *Surgery*, 2009.
- [24] N. Padoy. Machine and deep learning for workflow recognition during surgery. *Minimally Invasive Therapy and Allied Technologies*, 2019.
- [25] N. Ahmidi et al. A Dataset and Benchmarks for Segmentation and Recognition of Gestures in Robotic Surgery. *IEEE Transactions on Biomedical Engineering*, 2017.
- [26] C. Sutton and A. McCallum. An introduction to conditional random fields. *Foundations and Trends in Machine Learning*, 2011.
- [27] C. Lea et al. An Improved Model for Segmentation and Recognition of Fine-Grained Activities with Application to Surgical Training Tasks. *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2015.
- [28] C. Lea et al. Learning convolutional action primitives for fine-grained action recognition. *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.
- [29] C. Lea et al. Segmental spatiotemporal CNNs for fine-grained action segmentation. *European Conference on Computer Vision (ECCV)*, 2016.
- [30] C. Rupprecht et al. Sensor substitution for video-based action recognition. *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2016.
- [31] T. E. Murphy. Towards Objective Surgical Skill Evaluation with Hidden Markov Model-based Motion Recognition. page 116, 2004.
- [32] H. C. Lin et al. Towards automatic skill evaluation : Detection and segmentation of robot-assisted surgical motions. *Computer Aided Surgery*, 2006.
- [33] C. E. Reiley et al. Automatic recognition of surgical motions using statistical modeling for capturing variability. *Studies in Health Technology and Informatics*, 2008.
- [34] B. Varadarajan et al. Data-Derived Models for Segmentation with Application to Surgical Assessment and Training. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2009.
- [35] L. Tao et al. Sparse hidden Markov models for surgical gesture classification and skill evaluation. *International conference on information processing in computer-assisted interventions*, 2012.
- [36] S. Sefati et al. Learning Shared, Discriminative Dictionaries for Surgical Gesture Segmentation and Classification. *Modeling and Monitoring of Computer Assisted Interventions (M2CAI) - MICCAI Workshop*, 2015.
- [37] E. Mavroudi et al. End-to-end fine-grained action segmentation and recognition using conditional random field models and discriminative sparse coding. *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2018.
- [38] M. Selvaggio et al. Physics-based task classification of da Vinci robot surgical procedures. 2018.
- [39] B. Varadarajan. Learning and Inference Algorithms for Dynamical System Models of Dextrous Motion. 2011.
- [40] B. Varadarajan and S. Khudanpur. Learning and inference algorithms for partially observed structured switching vector autoregressive models. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2011.
- [41] F. Grigoli. A method for simultaneous task recognition and gesture classification in surgical robotics. 2018.
- [42] C. Lea et al. Temporal convolutional networks for action segmentation and detection. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [43] D. Liu and T. Jiang. Deep Reinforcement Learning for Surgical Gesture Segmentation and Classification. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2018.
- [44] B. Béjar Haro et al. Surgical Gesture Classification from Video Data. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2012.
- [45] N. O'Mahony et al. Deep learning vs. traditional computer vision. *Science and Information Conference*, 2019.
- [46] Y. Zhang et al. Local temporal bilinear pooling for fine-grained action parsing. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [47] P. Lei and S. Todorovic. Temporal Deformable Residual Networks for Action Segmentation in Videos. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [48] J. Wang et al. Atrous temporal convolutional network for video action segmentation. *IEEE International Conference on Image Processing (ICIP)*, 2019.
- [49] G. Menegozzo et al. Surgical gesture recognition with time delay neural network based on kinematic data. *International Symposium on Medical Robotics (ISMR)*, 2019.
- [50] T. Wang et al. Towards accurate and interpretable surgical skill assessment: A video-based method incorporating recognized surgical gestures and skill levels. *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2020.
- [51] J. Zhang et al. Symmetric dilated convolution for surgical gesture recognition. *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2020.
- [52] I. Funke et al. Using 3D Convolutional Neural Networks to Learn Spatiotemporal Features for Automatic Surgical Gesture Recognition in Video. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2019.
- [53] D. Sarikaya and P. Jannin. Towards Generalizable Surgical Activity Recognition Using Spatial Temporal Graph Convolutional Networks. *arXiv preprint arXiv:2001.03728*, 2020.
- [54] A. Kaku et al. Be Like Water: Robustness to Extraneous Variables Via Adaptive Feature Normalization. *arXiv preprint arXiv:2002.04019*, 2020.

- [55] R. DiPietro et al. Segmenting and classifying activities in robot-assisted surgery with recurrent neural networks. *International Journal of Computer Assisted Radiology and Surgery*, 2019.
- [56] I. Gurcan and H. Van Nguyen. Surgical Activities Recognition Using Multi-scale Recurrent Networks. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [57] D. Itzkovich et al. Using Augmentation to Improve the Robustness to Rotation of Deep Learning Segmentation in Robotic-Assisted Surgical Data. *IEEE International Conference on Robotics and Automation (ICRA)*, 2019.
- [58] B. van Amsterdam et al. Multi-Task Recurrent Neural Network for Surgical Gesture Recognition and Progress Prediction. *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [59] L. Ding and C. Xu. TricorNet: A Hybrid Temporal Convolutional and Recurrent Network for Video Action Segmentation. *arXiv preprint arXiv:1705.07818*, 2017.
- [60] Y. Qin et al. Temporal Segmentation of Surgical Sub-tasks through Deep Learning with Multiple Data Sources. *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [61] Y. Qin et al. davincinet: Joint prediction of motion and surgical state in robot-assisted surgery. *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [62] X. Gao et al. Automatic Gesture Recognition in Robot-assisted Surgery with Reinforcement Learning and Tree Search. *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [63] S. Krishnan et al. Transition state clustering: Unsupervised surgical trajectory segmentation for robot learning. *The International Journal of Robotics Research*, 2015.
- [64] A. Murali et al. TSC-DL: Unsupervised trajectory segmentation of multi-modal surgical demonstrations with Deep Learning. *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.
- [65] H. Zhao et al. A fast unsupervised approach for multi-modality surgical trajectory segmentation. *IEEE Access*, 2018.
- [66] R. Fox et al. Multi-level discovery of deep options. *arXiv preprint arXiv:1703.08294*, 2017.
- [67] K. Goel and E. Brunskill. Learning procedural abstractions and evaluating discrete latent temporal structure. *International Conference on Learning Representations (ICLR)*, 2019.
- [68] L. Clopton et al. Temporal Subspace Clustering for Unsupervised Action Segmentation. *Computational Sensing and Medical Robotics (CSMR) REU*, 2017.
- [69] F. Despinoy et al. Unsupervised Trajectory Segmentation for Surgical Gesture Recognition in Robotic Training. *IEEE Transactions on Biomedical Engineering*, 2016.
- [70] M. J. Fard et al. Soft Boundary Approach for Unsupervised Gesture Segmentation in Robotic-Assisted Surgery. *IEEE Robotics and Automation Letters*, 2017.
- [71] Y. Tsai et al. Unsupervised Task Segmentation Approach for Bimanual Surgical Tasks using Spatiotemporal and Variance Properties. *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [72] J. D. Jones et al. Zero-shot Recognition of Complex Action Sequences. *arXiv preprint arXiv:1912.03613*, 2019.
- [73] Y. Gao et al. Unsupervised surgical data alignment with application to automatic activity annotation. *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.
- [74] B. van Amsterdam et al. Weakly Supervised Recognition of Surgical Gestures. *IEEE International Conference on Robotics and Automation (ICRA)*, 2019.
- [75] R. DiPietro and G. D. Hager. Automated Surgical Activity Recognition with One Labeled Sequence. *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2019.
- [76] Y. Tsai et al. Transfer learning for surgical task segmentation. *IEEE International Conference on Robotics and Automation (ICRA)*, 2019.
- [77] A. Tanwani et al. Motion2vec: Semi-supervised representation learning from surgical videos. *IEEE International Conference on Robotics and Automation (ICRA)*, 2020.
- [78] D. Tran et al. Learning spatiotemporal features with 3d convolutional networks. *IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [79] A. P. Twinanda et al. EndoNet: A Deep Architecture for Recognition Tasks on Laparoscopic Videos. *IEEE Transactions on Medical Imaging*, 2017.
- [80] T. Yu et al. Learning from a tiny dataset of manual annotations: a teacher/student approach for surgical phase recognition. *arXiv preprint arXiv:1812.00033*, 2018.
- [81] G. Yengera et al. Less is More: Surgical Phase Recognition with Less Annotations through Self-Supervised Pre-training of CNN-LSTM Networks. *arXiv preprint arXiv:1805.08569*, 2018.
- [82] G. Quellec et al. Real-time recognition of surgical tasks in eye surgery videos. *Medical Image Analysis*, 2014.
- [83] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations (ICLR)*, 2015.
- [84] D. Sarikaya and P. Jannin. Surgical Gesture Recognition with Optical Flow only. *arXiv preprint arXiv:1904.01143*, 2019.
- [85] S. S. Mondal et al. Multitask Learning of Temporal Connectionism in Convolutional Networks using a Joint Distribution Loss Function to Simultaneously Identify Tools and Phase in Surgical Videos. *arXiv preprint arXiv:1905.08315*, 2019.
- [86] X. Li et al. Progress Estimation and Phase Detection for Sequential Processes. *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2017.
- [87] F. Luongo et al. Deep learning-based computer vision to recognize and classify suturing gestures in robot-assisted surgery. *Surgery*, 2020.
- [88] V. Bawa et al. Esad: Endoscopic surgeon action detection dataset. *arXiv preprint arXiv:2006.07164*, 2020.
- [89] Y. Fu et al. More unlabelled data or label more data? a study on semi-supervised laparoscopic image segmentation. *MICCAI Workshop on Domain Adaptation and Representation Transfer*, 2019.
- [90] F. Bovo et al. Surgical robot simulation with BBZ console. *Journal of Visualized Surgery*, 2017.
- [91] N. Madapana et al. DESK: A Robotic Activity Dataset for Dexterous Surgical Skills Transfer to Medical Robots. *arXiv preprint arXiv:1903.00959*, 2020.
- [92] C. Lea. Temporal convolutional networks. <https://github.com/colincsl/TemporalConvolutionalNetworks>, 2017.