



HAL
open science

Recognition of Laban Effort Qualities from Hand Motion

Maxime Garcia, Rémi Ronfard

► **To cite this version:**

Maxime Garcia, Rémi Ronfard. Recognition of Laban Effort Qualities from Hand Motion. MOCO'20 - 7th International Conference on Movement and Computing, Jul 2020, Jersey City/ Virtual, United States. pp.1-8, 10.1145/3401956.3404227 . hal-02899999

HAL Id: hal-02899999

<https://inria.hal.science/hal-02899999v1>

Submitted on 15 Jul 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Recognition of Laban Effort Qualities from Hand Motion

Maxime Garcia

Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LJK
38000 Grenoble, France
maxime.garcia1@ensimag.grenoble-inp.fr

Rémi Ronfard

Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LJK
38000 Grenoble, France
remi.ronfard@inria.fr

ABSTRACT

In this paper, we conduct a study for recognizing motion qualities in hand gestures using virtual reality trackers attached to the hand. From this 6D signal, we extract Euclidean, equi-affine and moving frame features and compare their effectiveness in the task of recognizing Laban Effort qualities. Our experimental results reveal that equi-affine features are highly discriminant features for this task. We also compare two classification methods on this task. In the first method, we trained separate HMM models for the 6 Laban Effort qualities (light, strong, sudden, sustained, direct, indirect). In the second method, we trained separate HMM models for the 8 Laban motion verbs (dab, glide, float, flick, thrust, press, wring, slash) and combined them to recognize individual qualities. In our experiments, the second method gives improved results. Together, those findings suggest that low-dimensional signals from VR trackers can be used to predict motion qualities with reasonable precision.

CCS CONCEPTS

• **Computing methodologies** → **Maximum likelihood modeling; Classification and regression trees;**

KEYWORDS

Motion Qualities Recognition, Laban Movement Analysis

ACM Reference Format:

Maxime Garcia and Rémi Ronfard. 2020. Recognition of Laban Effort Qualities from Hand Motion. In *7th International Conference on Movement and Computing (MOCO '20)*, July 15–17, 2020, Jersey City/ Virtual, NJ, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3401956.3404227>

1 INTRODUCTION

Immersive virtual reality systems such as HTC Vive or Oculus Rift allow users to use their own body as a user interface. More specifically, such systems provide precise and reliable measurements of the six dimensional motion of trackers attached to the user's head and hands in real time. This has shown to be useful in applications such as 3D drawing, painting and sculpting, among others.

Previous work has looked at various methods for recognizing hand gestures for such purposes. In this paper, we are interested in measuring the Laban Effort qualities associated with such gestures. Recognizing motion qualities is particularly useful for understanding motion intents and can be used in interactive systems. The range of

application that use Laban Effort qualities as input is quite diverse and includes live performances interactions [9] as well as content stylization, especially 3D character animations [2]. More generally, Laban Motion Analysis provides a descriptive vocabulary which can be used as a sub-representation of human motions.

In our study, we ask users to draw 3D shapes with various motion qualities, typically lighter or stronger, quicker or more sustained, more or less direct. Using the subjective motion qualities as our ground truth, we created a small data set of 11 gestures performed by 3 users, including two Laban experts, in 8 motion qualities.

In our study, we were interested to answer the following questions:

- Can we automatically recognize motion qualities of gestures, i.e. decide whether a given gesture was performed lightly or strongly, suddenly or sustainedly, directly or indirectly ?
- Can we automatically recognize Laban Effort Verbs, i.e. combinations of Laban Efforts in the Space, Time and Weight dimensions simultaneously?
- Which geometric features are most useful in recognizing motion qualities in either case?

Through our study, we brought quantitative results answering those three challenges. More particularly, we compared two Hidden Markov Model (HMM) based classifiers, one trained on Laban Elements while the other was trained to recognize Effort Verbs (also sometime called Laban Action Drives), and compared their advantages and weaknesses. In addition, we conducted several classification studies using three different feature sets extracted from the input gestures and showed the relevance of natural motion descriptors such as equi-affine speed and acceleration on the task of Laban Effort recognition.

2 LABAN MOVEMENT ANALYSIS

Laban Movement Analysis (LMA) [17] is a method for characterizing quality in human movements in terms of timing, space used, posture and intention. The full LMA uses five description categories, namely Effort, Shape, Phrasing, Body and Space, to classify motion. In this work, we focus on the Effort category and more precisely, we are interested in recognizing Effort qualities from hand motions only.

Laban Effort describes the dynamic quality of motion and is based on the 4 following descriptors (with the associated extreme values): Space (direct - indirect), Time (sustained - sudden), Weight (light - strong) and Flow (bound - free) making the Effort category isomorphic to $[-1; 1]^4$.

Expert studies on Laban Effort expression [5] observed that Space, Time and Weight are often sufficient to describe actions, with either two or three of them being active at the same time. Motions involving one descriptor are defined as Factors (characterized by two opposed

MOCO '20, July 15–17, 2020, Jersey City/ Virtual, NJ, USA

© 2020 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *7th International Conference on Movement and Computing (MOCO '20)*, July 15–17, 2020, Jersey City/ Virtual, NJ, USA, <https://doi.org/10.1145/3401956.3404227>.

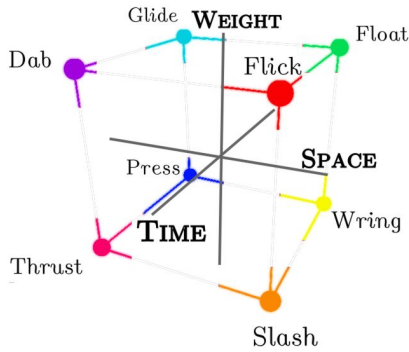


Figure 1: Laban’s Effort cube: the eight Effort Verbs are represented at the corners of the cube. The three Effort Factors of Space, Time and Weight, are represented as axes of the cube. The six Effort Elements (light, strong, sudden, sustained, direct, indirect) are the faces of the cube.

Elements), those involving two descriptors are defined as States and those involving three of them as Drives. In this study, we omit the Flow dimension and focus on the so-called Action Drives which involve the three dimensions of Space, Time and Weight. The Action Drives are also referred as the 8 Effort Verbs : dab, glide, float, flick, thrust, press, wring, slash. Figure 1 shows a representation of the 8 Verbs positioned in the Laban Effort cube. In table 1, we detail the decomposition of each Effort Verb into its corresponding Elements.

Table 1: Laban Effort Verbs decomposition into their corresponding Elements. Each Effort Verb contains three Elements. Each Element is contained in four Effort Verbs.

Verbs \ Elements	Direct	Indirect	Sudden	Sustained	Light	Strong
Dab	✓		✓		✓	
Flick		✓	✓		✓	
Slash		✓	✓			✓
Thrust	✓		✓			✓
Glide	✓			✓	✓	
Float		✓		✓	✓	
Press	✓			✓		✓
Wring		✓		✓		✓

3 RELATED WORK

3.1 Spatial interaction

Our research is motivated by the recent popularization of 3D spatial interaction devices [20] in 3D shape design and visual arts. In the near future, we predict that those same devices will also be used to create 3D motions and animations. Indeed, using tangible interaction devices for creating short animation sequences is a good practice allowing even non experts to create motion. As tangible interaction enhances interactivity, various puppetry interfaces have been proposed in order to control characters in real time. While numerous

systems producing either simple rigid motions [16] or more elaborate articulated animations [25] have been studied, none of them aimed at expressive animation. In this work, we present experimental results obtained with the HTC Vive controller and Lighthouse tracking system, which provide a complete 6D motion signal at 60 fps.

3.2 Geometry of natural motion

Researchers have noticed intriguing properties of spontaneous human hand movements, especially during writing and drawing. More particularly, [27] highlighted the importance of two-dimensional equi-affine velocity for qualifying uniform drawing motion.

Based on those observations, we conjecture that three-dimensional equi-affine velocity and its derivatives (acceleration and jerk) may play a role in expressing and recognizing expressive qualities of hand motions.

3.3 Gesture recognition

Human gesture recognition is a process which is now integrated in many applications, especially for entertaining purposes as it can be found in video games or interactive performances (theater, concerts,...). This particular problem has been studied both from videos and human motion data [20]. Moreover, gesture recognition from motion data is typically done using machine learning algorithm taking either 3D trajectories and their geometrical properties such as speed and acceleration as input. As 3D trajectories are temporal sequences, many studies for gesture recognition have been performed using Hidden Markov Model [31]. Another class of methods performing pattern matching in a auxiliary space also tackle this problem. This was studied for recognizing 2D gestures [30] as well as 3D gestures [7].

Hidden Markov Models have been used extensively for both gesture segmentation and recognition in the past (see [28] for a comprehensive survey). HMMs have also been used to segment and recognize sign language from hand motion [11, 33] and this has been generalized to parallel HMMs for two-hand signing [34]. In related work, Green used HMMs to recognize full body gestures by segmenting actions into dynemes, which are a generalization of gestures to full body motions [14].

3.4 Motion quality recognition

Motion qualities are recognizable features of a given motion which convey the intent of the performer, describing how the motion is executed. Motion qualities have been researched extensively in the 1940s by Rudolf Laban for the purpose of dance notation [17] and previous work in the computer graphics community has attempted to recognize Laban motion qualities from natural human motion [6, 22] and to generate 3D character animation with recognizable Laban motion qualities [5, 10], or both [2]. In recent years, Laban motion analysis has gained recognition in other fields including computer human interaction [1] and robotics [19]. As a result, an important topic of research has been to recognize Laban motion qualities from various motion signals using machine learning.

Camurri et al. classify hand gestures according to the Space and Time dimensions of Laban’s Theory of Effort [8]. Silang Maranan

et al. recognize Laban Effort qualities from a single wearable accelerometer in real time [29]. Bernardet et al. report subjective experiments evaluating the ability of expert movement analysts to recognize motion qualities in videos of the pantomimed movements of knocking at a door or giving directions [4].

Extracting Laban Effort qualities from low dimensional motion was explored in the context of wrist motion data [12] using multi-modal input, including wrist position, speed, acceleration and muscle activation. In contrast, we take as input a single rigid body motion and extract a rich set of geometric and kinematic features, such as Euclidean and equi-affine velocity, acceleration, jerk, curvature and torsion.

4 OVERVIEW

In this paper, we present a study focusing on Laban Effort classification from hand motions recorded with an HTC Vive controller held by users. More precisely, we are interested in finding relevant geometric features that describe motion qualities in terms of Space, Time and Weight efforts. To do so, we designed three sets of features and evaluated how they influence classification results. We also performed Laban Effort classification using two different sets of HMMs. The first set contains 8 Verb HMMs trained on Verbs. The second set contains 6 Element HMMs trained on Elements. We compare the three feature sets and two HMM sets on the two tasks of classifying motions into Verbs and Elements.

In the next sections, we present our data acquisition setup as well as our experimental protocol (section 5). We then introduce the different feature sets that we extracted from the input 6D trajectories (section 6) which are then used for training and testing our two classifiers described in section 7. Finally, experimental results are presented in section 8 before concluding and discussing.

5 DATA SET ACQUISITION

We collected our data set of rigid hand motions by tracking the moving frame of a manipulated object, in our case the the HTC Vive Controller (see Figure 2). Other similar methods have used a variety of available devices, including video and Kinect cameras, inertial measurement units (IMU), magnetic sensors or optical tracking systems [15, 20].

Our data set is composed of 6-D input curves from which we compute geometrical properties as explained in Section 6. We registered data from 3 users including 2 LMA experts and one user familiar with LMA. A total of 104 curves was recorded in which 88 are predefined patterns performed in the 8 Effort Verbs, and 16 are free style gestures also performed in the 8 Effort Verbs. Using patterns for training is especially useful to dissociate motion qualities within a single pattern while freestyle trajectories were introduced to avoid over-fitting and bias. The eleven predefined patterns are presented in Figure 3. Every curve of the data set was then re-sampled such that each curve contains 512 points. For classification results, we randomly split our curve data set into 56 curves (7 per Verb) for training and 48 curves (6 per Verb) for testing. In addition, we augmented our training set by sub-sampling each of the 56 training curve 15 times, totaling 840 curves and including 105 examples per Verb. In total, the data set is thus composed of 888 curves divided into 840 training curves and 48 testing curves.



Figure 2: User recording 3D patterns in different Laban Effort qualities.

6 EUCLIDEAN, AFFINE AND MOVING FRAME FEATURES

In this section, we describe the geometric features we extract from user’s hand motion which will be used as input for our Laban Effort classifiers (see section 7). We do not attempt to recognize motion qualities directly from the raw 6D hand motion trajectories (end-to-end recognition) but instead choose features from different geometries, which have been found to be important in describing natural human motion in previous work [3].

Similar to previous work like [24] and [12], we based our feature choice on *a priori* assumptions relating geometric features and their influence in Effort quality variations. Contrary to previous work, we do not record any physiologic data from users, which makes our task more difficult. Fdili-Alaoui et al. tracked the muscular activity of the wrist and emphasized the importance of this physiological measurement for recognizing motion qualities [12]. Instead, we propose to study a larger set of geometric features which we relate to individual Laban Elements.

We relate Torsion τ and Curvature κ to the Space Element. Direct motions are should have constant τ and κ while Indirect motions often deviate from the neutral motion direction and implies variations of those two features. This can be illustrated in the 2D example presented in Figure 4.

As the Time Element opposes Sustained to Sudden motion, we include speed $V_e = (v_e, \omega)$, containing linear and angular velocity information as well as its two successive derivatives $A_e = (\dot{v}_e, \dot{\omega})$ and $J_e = (\ddot{v}_e, \ddot{\omega})$ computed using finite differences. Those features are expressed in the global world space. In addition, we also compute those the linear velocity and its derivative feature in the moving frame of the HTC Vive controller v, \dot{v}, \ddot{v} (with $v = R_i(p_i - p_{i-1})$, p_i being the 3D coordinates of the i^{th} point of the curve and R_i the




Figure 3: 6D Patterns used both for training and testing.

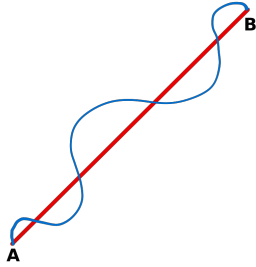


Figure 4: Space Element 2D representation using two paths from A to B. Red: Direct path. Blue: Indirect path.

rotation matrix associated with the moving frame at point i).

Studying Weight, which opposes Strong to Light motions, might seem difficult as we do not have the physiologic data which was shown to be useful in [12]. However, we notice that uniformity of motion is an important aspect differentiating Strong from Light gestures. Strong motions are irregular and wide while Light motions are continuous and fluid. It turns out that equi-affine velocity provides a suitable measure of the perceptual uniformity of motion, both in 2D [18, 21, 27, 32] and 3D [26]. In 2D, the equi-affine velocity of a point moving along a planar trajectory is defined as $V_a = |v_e \times \dot{v}_e|^{\frac{1}{3}}$ which can also be written as $V_a = v_e \kappa^{\frac{1}{3}}$. In their work, [27] shows that humans draw spontaneously at a constant equi-affine speed which is known as the 1/3 power law.

More generally, equi-affine velocity was identified as being highly correlated with our perception of speed of movement in drawing gestures [26]. Indeed, experimental psychologists have shown that human subjects describe movements performed with constant equi-affine velocity as uniform. By analogy, we hypothesize that spatial equi-affine velocity should also play an important role for measuring the perceived velocity of 3D hand motion.

In 3D, equi-affine velocity can be computed efficiently using the triple scalar product $V_a = |\dot{r}, \ddot{r}, \ddot{\ddot{r}}|^{\frac{1}{6}}$ which is again related to the Euclidean velocity with the formula $V_a = V_e \kappa^{\frac{1}{3}} \tau^{\frac{1}{6}}$ where κ is the curvature of the stroke as before and $\tau = \frac{|\dot{r}, \ddot{r}, \ddot{\ddot{r}}|}{\|\dot{r} \times \ddot{r}\|^2}$ is the torsion of the stroke.

Similarly to the 2D case, recent work has shown that human subjects describe spatial movements performed with equi-affine velocity as uniform [13, 23, 26] and this is known as the 1/6 power law. Those findings make spatial equi-affine velocity a likely candidate to measure subjective motion qualities and we therefore include it in our feature set. We also compute equi-affine acceleration as the first derivative of V_a .

In Figure 5 we show an example of V_e and V_a profiles for Light and Strong motions. In addition, we also show, in Figure 6, the linear and angular velocities profiles in the moving frame attached to the hand for the same examples. Those two figures highlight the uniform aspect of Light motions against Strong motions especially in terms of equi-affine speed behavior (more constant for Light motions). We can also notice high speed and rotational speed variations for Strong motions.

In total, we measure 18 moving frame features and 7 world frame features (Euclidean velocity, acceleration and jerk, curvature, torsion, equi-affine velocity, and equi-affine acceleration). We split those feature into three different sets presented in table 2.

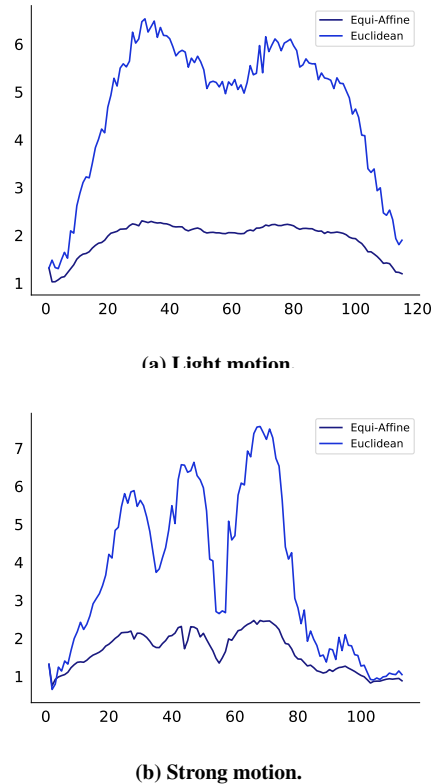


Figure 5: Euclidean V_e and Equi-affine V_a velocity profiles of Loop gesture in Light and Strong motion qualities.

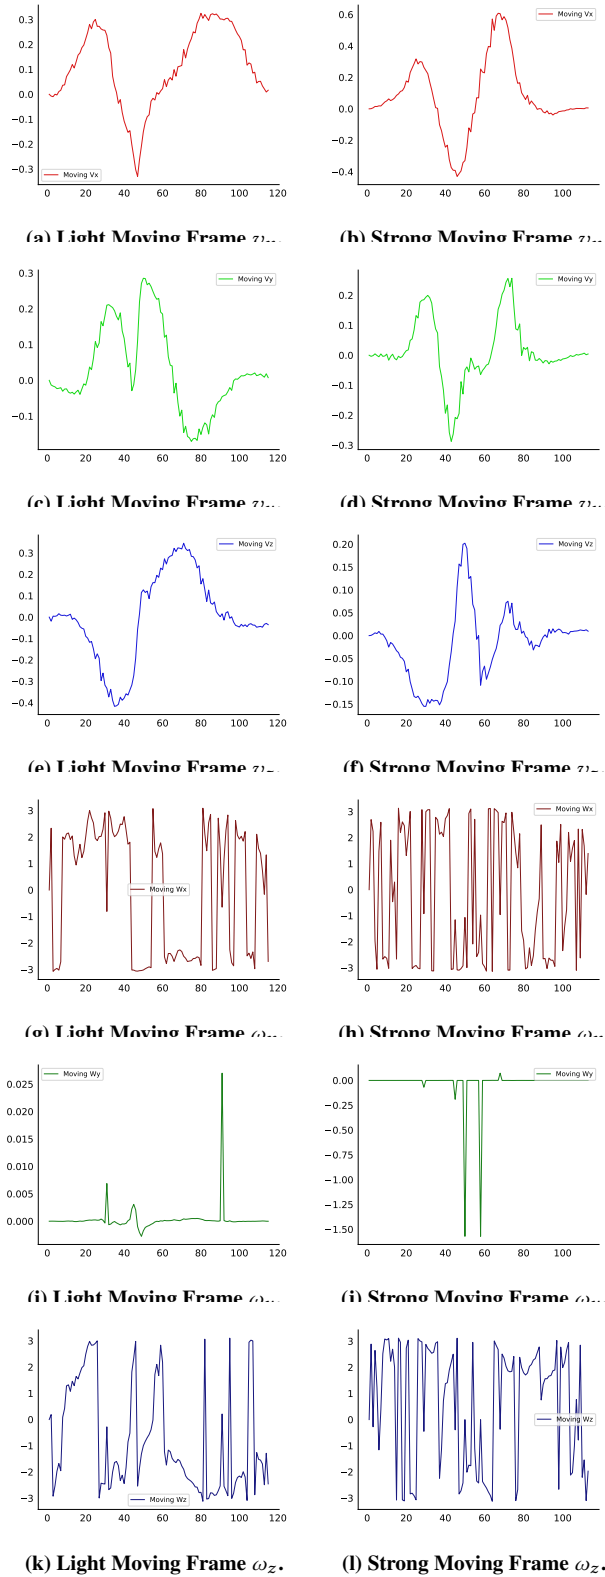


Figure 6: Linear and angular velocities in the moving frame of Loop gesture in Light and Strong motion qualities.

Table 2: Feature sets used for Laban Effort classification.

Euclidean	Equi-Affine	Moving Frame
V_e	V_a	v, ω
γ_e	γ_a	$\dot{v}, \dot{\omega}$
J_e	κ, τ	$\ddot{v}, \ddot{\omega}$

7 LABAN EFFORT CLASSIFICATION

We study two classification tasks in this paper. In the Verb classification task, our goal is to recognize instances of the eight Effort Verbs from observations. In the Factor classification task, our goal is to classify observations along the separate dimensions of Time (is the motion sudden or sustained?), Space (is the motion direct or indirect?) and Weight (is the motion light or strong?).

We approach each classification task with two different methods. In the Verb HMM method, we train one HMM per Effort Verb using all available training examples for that Verb. On the Effort Verb classification task, we compute the 8 Verb likelihoods directly from the trained Verb HMMs and return the verb with maximum likelihood. On the Factor classification task, we simply compute the most likely Verb and return its component Elements in each Factor, as shown in Table 1).

In the Element HMM method, we train one HMM per Element using all available training examples from that Element. On the Factor classification task, we compute the 6 Element likelihoods directly from the trained Element HMMs and return the Element with maximum likelihood in each Factor. On the Verb classification task, we compute the Verb likelihoods indirectly, as a combination of the likelihoods for the three elements contained in the verb (as shown in Table 1), and we return the most likely Verb.

For both methods, we used left-to-right HMMs with 6 states. The states were trained independently for each of the 14 HMMs from all available training examples. The number of states was chosen empirically after extensive experiments using between one and twelve states because it gave the best results.

8 EXPERIMENTAL RESULTS

In this section, we present classification results performed on the test set described in Section 5 using the two classifiers described in Section 7 and the three feature sets presented in Section 6. We first study the Verb classification task for both methods and compare results using different feature set combinations. Then, we study the Factor classification task and compare the two methods Factor-wise. Finally, we study the special case of Time and Weight classification, where the space Dimension is ignored.

8.1 Verb classification

For each classifier, we compared the precision for Verb classification using different geometrical features among the feature sets described in table 2. We ran three classification processes in total. The first process only includes Euclidean features. The per Verb classifier performed 25% of successful classification rate against 18% for the Element HMM method. Adding equi-affine properties to our set of feature for the second process, we observed significantly improved

classification results with 43% success rate for the per Verb method against 25% for the Element one. Finally, best results were achieved when also adding Motion Frame related properties to our feature set, achieving 44% rate for the per Verb method and 25% for the composed method, even if improvements are subtle at the end of this process. Verb classification results are summarized in Figure 7.

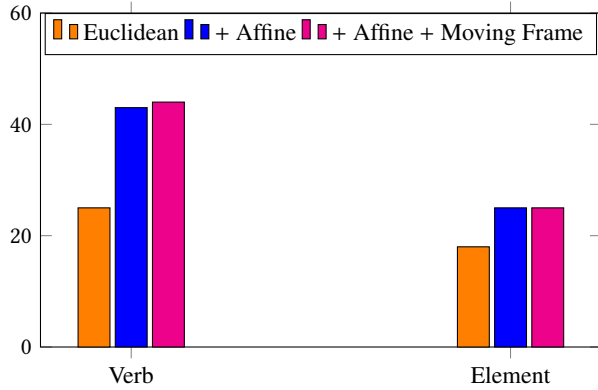


Figure 7: Laban verb classification accuracy in percentage using the verbs HMM classifier and the Element HMM classifier.

Those first results reveal that the Verb HMMs outperform the Element HMMs in the task of Verb classification. In addition, we notice that equi-affine features greatly improve classification results for both classifiers while it is not the case for moving frame features. This confirms that equi-affine features are discriminant features for the task of Verb classification.

As confusion still occurs in the overall classification results, we first computed the confusion matrices of the two methods and gathered the results in Figure 8.

These matrices especially highlight that the two classifiers have very different confusion categories. Moreover, we can notice that the Element HMM method seems to get confused over categories which share the same Time and Weight Effort but which differs in terms of Space like this is the case for Dab/Flick, Slash/Thrust and Glide/Float. In the same manner, the Verb HMM method is not confused by the Time Effort. But it is less accurate in categories with different Weight and Space Elements at the same time.

8.2 Factor classification

We also evaluated how our classifiers performed on individual Factor classification in order to determine more precisely the advantages and weaknesses of the two methods. We summarized the results in Figure 9 and 10

We first observe that Time is highly recognizable in both methods with classification results above 85% when using all features. Interestingly, we notice that the Element HMM method is more accurate than the Verb HMM method for Space and Time classification regardless of the feature sets we used.

On the contrary, the Verb HMM method is more accurate in terms of Weight classification as well as Verb classification. Overall, it is not surprising that the Element HMM method performs better at

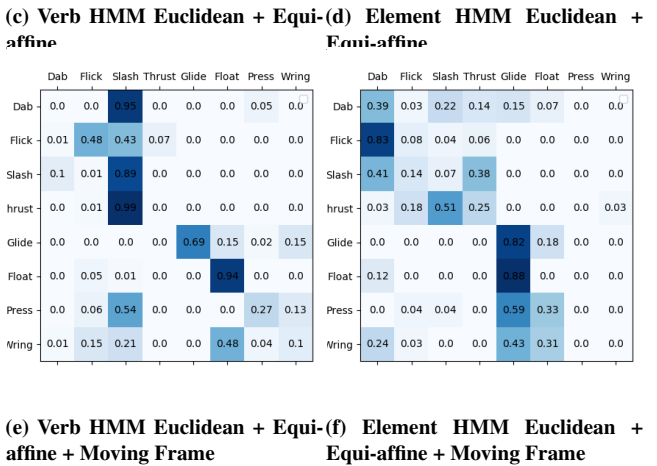
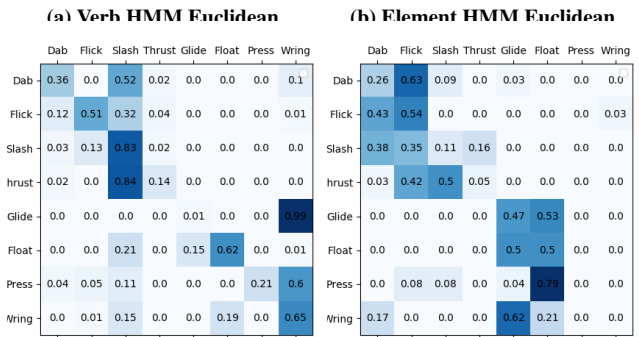
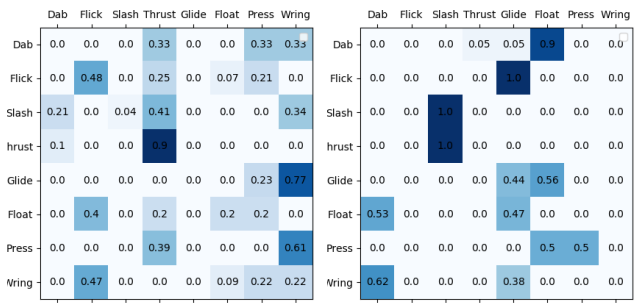


Figure 8: Confusion matrices. Left: Classification confusion using one HMM per Verb. Right: Classification confusion using one HMM per Effort Element.

individual classification. However, its accuracy in terms of individual Element classification is only slightly better than the Verb HMM, yet it performs poorly in terms of Verb classification. Thus, these results highlight that learning over individual Elements misses out correlations between Elements when combining classification results while learning over the Verbs directly is usually more robust.

We also notice that the Verb HMM method achieves better results than [12] who used one HMM per Element on the task of Factor classification. While we did not study classification results on their data (which may not be well suited for our study) we achieve 44%

precision against 26% as reported in [12]. Those results are not directly comparable as the testing sets differs. In future experiments, it would be useful to compare classification results using common testing sets.

For both classifiers, we notice that Space and Weight Factors are more difficult to recognize and cause a significant drop in Verb recognition rates. In the case of the Verb HMM method we also notice that the Space Factor is the most difficult dimension to classify.

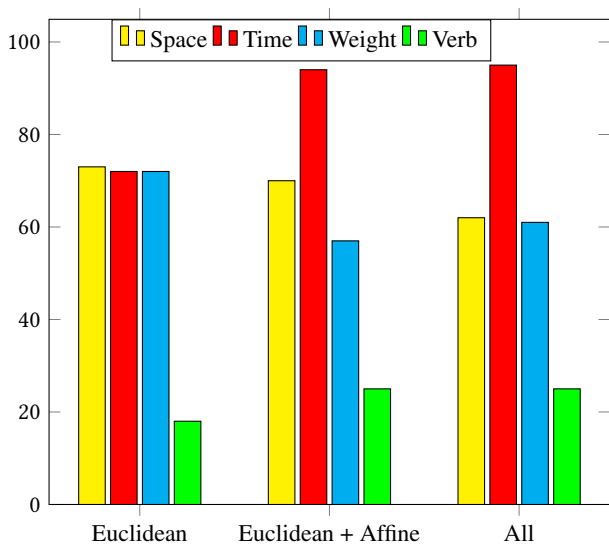


Figure 9: Factor and Verb classification results using the Element HMM classifier.

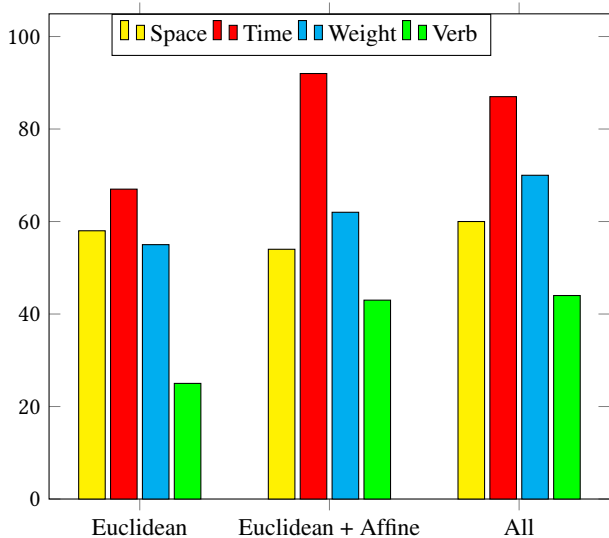


Figure 10: Factor and Verb classification results using the Verb HMM classifier.

8.3 Time/Weight classification

To better evaluate the impact of the Space Factor in our classification results we removed it from the classification results and computed the Time/Weight State classification precision rate using the Verb HMM and the Element HMM methods. The two methods give almost identical results when considering only the Time and Weight Factors. More precisely, the best results are obtained by combining the Euclidean and Equi-affine feature sets in both cases with 63% of successful recognition for the Verb HMM method and 62% for the Element HMM. We observe a small drop of successful classification rate when considering all feature sets for the two methods (61% for the Verb HMM and 58% for the Element HMM) indicating that adding the moving frame feature set brings confusion which we can observe for the Time Factor in Figure 9 and 10.

In addition, the similarity of the classification results between the two methods reveals that the Element HMM method fails to classify curves in terms of Space when it correctly classifies those curves in terms of Time and Weight which is less the case for the Verb HMM method.

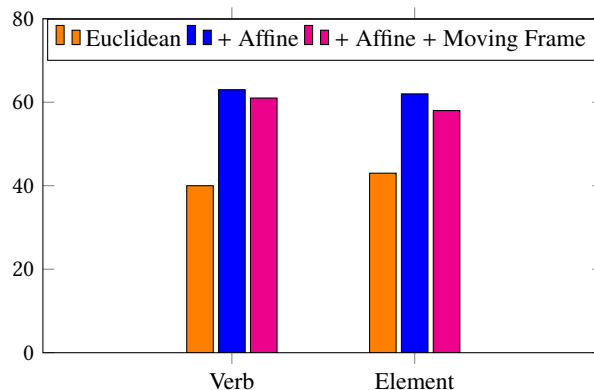


Figure 11: Laban State classification accuracy in percentage using the Verb HMM classifier and the Element HMM classifier.

9 LIMITATIONS AND FUTURE WORK

Our study and our methods come with some limitations. First, Verb classification from hand motion remains a difficult problem as we achieve 44% of successful classification rate at best. More precisely, while we obtain high recognition accuracy for Time Element recognition, Space and Weight classification still remains a hard task. Indeed, hand 6D trajectory might not be sufficient to describe users' full intent, it could be useful to also capture more body parts motion like the elbow and the shoulder. Besides, we did not properly evaluate the role of rotations in the classification process as we allowed users to hold the HTC Vive controllers in any way they wished, which may have introduced rotation bias on our database (the role of rotation in classification can be seen in Figure 6).

For future work, it would be interesting to study Laban Effort classification on all Laban Motion States, including Space/Time and Space/Weight classifications to better understand where confusion between elements comes from. We also would like to perform subjective user experiments to measure how well the intended motion

qualities are perceived by naive users, and compare those subjective assessments with our classification results. Finally, we would like to further enlarge our training and testing data set and also include physiological features like captured by [12] and compare classification results. Moreover, it would be interesting to evaluate properly if the geometric features we used in this paper can substitute to physiological features.

10 CONCLUSION

We presented a new study for recognizing Laban Efforts from hand motion characterized by a 6D trajectory. Our study reveals that the choice of classification structures and geometric features can make a significant difference. Our experimental results suggest that training HMMs directly on Effort Verbs can achieve better results than training HMMs on Effort Elements. We also demonstrated the role of equi-affine features which greatly enhance Verb classification results. While future work is needed to confirm our results with larger data sets, our initial findings suggest that low-cost VR tracking systems may be used to recognize motion qualities in a variety of hand gestures, with possible applications to computer human interactions and expressive computer animation.

REFERENCES

- [1] Sarah Fdili Alaoui, Baptiste Caramiaux, Marcos Serrano, and Frédéric Bevilacqua. 2012. Movement Qualities as Interaction Modality. In *Proceedings of the Designing Interactive Systems Conference*. 761–769.
- [2] Andreas Aristidou, Qiong Zeng, Efstathios Stavarakis, KangKang Yin, Daniel Cohen-Or, Yiorgos Chrysanthou, and Baoquan Chen. 2017. Emotion Control of Unstructured Dance Movements. In *Proceedings of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA '17)*.
- [3] Daniel Bennequin, Ronit Fuchs, Alain Berthoz, and Tamar Flash. 2009. Movement Timing and Invariance Arise from Several Geometries. *PLoS Computational Biology* 5, 7 (2009).
- [4] U. Bernardet, S. Fdili Alaoui, K. Studd, K. Bradley, P. Pasquier, and T. Schiphorst. 2019. Assessing the reliability of the Laban Movement Analysis system. *PLoS ONE* 14, 6 (2019).
- [5] Leslie Bishko. 2014. Animation Principles and Laban Movement Analysis: Movement Frameworks For Creating Empathic Character Performances. ETC Press, Pittsburgh, PA, USA, 177–203.
- [6] D. Bouchard and N. Badler. 2007. Semantic Segmentation of Motion Capture Using Laban Movement Analysis. In *Proceedings of the 7th International Conference on Intelligent Virtual Agents*, Vol. 4722. 37–44.
- [7] Ernesto Bribeasa. 2007. Classification and Generation of 3D Discrete Curves. *Applied Mathematical Sciences* 57 (2007), 2805 – 2825.
- [8] Antonio Camurri, Corrado Canepa, Simone Ghisio, and Gualtiero Volpe. 2009. Automatic Classification of Expressive Hand Gestures on Tangible Acoustic Interfaces According to Laban’s Theory of Effort. In *Gesture-Based Human-Computer Interaction and Simulation*. 151–162.
- [9] Antonio Camurri, Shuji Hashimoto, Matteo Ricchetti, Andrea Ricci, Kenji Suzuki, Riccardo Trocca, and Gualtiero Volpe. 2000. EyesWeb: Toward Gesture and Affect Recognition in Interactive Dance and Music Systems. *Computer Music Journal* 24 (04 2000), 57–69.
- [10] Diane Chi, Monica Costa, Liwei Zhao, and Norman Badler. 2000. The EMOTE Model for Effort and Shape. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '00)*. ACM Press/Addison-Wesley Publishing Co., 173–182.
- [11] Dominik Endres, Yaron Meirovitch, Tamar Flash, and Martin Giese. 2013. Segmenting sign language into motor primitives with Bayesian binning. *Frontiers in Computational Neuroscience* 7 (2013), 68.
- [12] Sarah Fdili Alaoui, Jules Françoise, Thecla Schiphorst, Karen Studd, and Frederic Bevilacqua. 2017. Seeing, Sensing and Recognizing Laban Movement Qualities. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. 4009–4020.
- [13] Tamar Flash and Amir A. Handzel. 2007. Affine differential geometry analysis of human arm movements. *Biological Cybernetics* 96, 6 (2007), 577–601.
- [14] Richard Green. 2004. Extending the Perceptual User Interface to Recognise Movement. In *Computer Human Interaction*, Masood Masoodian, Steve Jones, and Bill Rogers (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 121–132.
- [15] Saikat Gupta, Sujin Jang, and Karthik Ramani. 2014. PuppetX: A Framework for Gestural Interactions with User Constructed Playthings. In *Proceedings of the 2014 International Working Conference on Advanced Visual Interfaces (AVI '14)*. 73–80.
- [16] Robert Held, Ankit Gupta, Brian Curless, and Maneesh Agrawala. 2012. 3D Puppetry: A Kinect-based Interface for 3D Animation. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology (UIST '12)*. ACM, 423–434.
- [17] Rudolf Laban. 1950. *The mastery of movement*. Macdonald & Evans, London.
- [18] Francesco Lacquaniti, Carlo Terzuolo, and Paolo Viviani. 1983. The law relating the kinematic and figural aspects of drawing movements. *Acta Psychologica* 54, 1 (1983), 115 – 130.
- [19] Jean-Paul Laumond and Naoko Abe (Eds.). 2016. *Dance Notations and Robot Motion*. Springer.
- [20] Joseph J. LaViola and Daniel F. Keefe. 2011. 3D Spatial Interaction: Applications for Art, Design, and Science. In *ACM SIGGRAPH 2011 Courses (SIGGRAPH '11)*. Article 1, 75 pages.
- [21] Sergey Lebedev, Wai Hon Tsui, and Peter Van Gelder. 2001. Drawing Movements as an Outcome of the Principle of Least Action. *J. Math. Psychol.* 45, 1 (Feb. 2001), 43–52.
- [22] Chris Bregler Lorenzo Torresani, Peggy Hackney. 2007. Learning Motion Style Synthesis from Perceptual Observations. In *Proc. of Neural Information Processing Systems (NIPS)*.
- [23] Uri Maoz, Alain Berthoz, and Tamar Flash. 2009. Complex unconstrained three-dimensional hand movement and constant equi-affine speed. *Journal of Neurophysiology* 101, 2 (February 2009), 1002–1015.
- [24] Megumi Masuda, Shohei Kato, and Hidenori Itoh. 2010. A Laban-Based Approach to Emotional Motion Rendering for Human-Robot Interaction. In *Entertainment Computing - ICEC 2010 (Lecture Notes in Computer Science)*. Springer, Berlin, Heidelberg.
- [25] Sageev Oore, Demetri Terzopoulos, and Geoffrey Hinton. 2002. A Desktop Input Device and Interface for Interactive 3D Character Animation. In *Graphics Interface*.
- [26] Frank E. Pollick, Uri Maoz, Amir A. Handzel, Peter J. Gibling, Guillermo Sapiro, and Tamar Flash. 2009. Three-dimensional arm movements at constant equi-affine speed. *Cortex* 45, 3 (2009), 325 – 339. Special Issue on Cognitive Neuroscience of Drawing.
- [27] F. E. Pollick and G. Sapiro. 1997. Constant affine velocity predicts the 1/3 power law of planar motion perception and generation. *Vision Research* 37, 3 (Feb. 1997), 347–353.
- [28] Siddharth S. Rautaray and Anupam Agrawal. 2015. Vision Based Hand Gesture Recognition for Human Computer Interaction: A Survey. *Artif. Intell. Rev.* 43, 1 (Jan. 2015), 1–54.
- [29] Diego Silang Maranan, Sarah Fdili Alaoui, Thecla Schiphorst, Philippe Pasquier, Pattarawat Subyen, and Lyn Bartram. 2014. Designing for Movement: Evaluating Computational Models Using LMA Effort Qualities. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 991–1000.
- [30] Matthew Thorne, David Burke, and Michiel van de Panne. 2004. Motion Doodles: An Interface for Sketching Character Motion (SIGGRAPH '04). ACM, 424–431.
- [31] Arthur Truong and Titus Zaharia. 2017. Laban movement analysis and hidden Markov models for dynamic 3D gesture recognition. *EURASIP Journal on Image and Video Processing* 2017, 1 (Aug. 2017), 52–1 – 52–16.
- [32] P. Viviani and M. Cenzato. 1985. Segmentation and coupling in complex movements. *Journal of Experimental Psychology: Human Perception and Performance* 11, 6 (December 1985), 828–845.
- [33] Christian Vogler and Dimitris Metaxas. 1998. ASL Recognition Based on a Coupling Between HMMs and 3D Motion Analysis. In *Proceedings of the Sixth International Conference on Computer Vision*.
- [34] Christian Vogler and Dimitris Metaxas. 2001. A Framework for Recognizing the Simultaneous Aspects of American Sign Language. *Computer Vision and Image Understanding* 81, 3 (2001), 358 – 384.