NEWSLETTER | Issue #2 | November 2018

WhoLoDancE

WHOLE-BODY INTERACTION LEARNING FOR DANCE EDUCATION

STRIDING TOWARDS THE GOAL

Issue **#2** | November 2018

NEWSLETTER INFORMATION

Contributors
Katerina El Raheb - *Athena Research Centre*
Rosamaria Cisneros, Ruth Gibson, Sarah Whatley, Karen Wood - *Coventry University*
Amalia Markatzi - *Lykeion Ton Hellenidon*
Mirko De Maldè, Edwin Morley-Fletcher, Anna Rizzo, Vana Tsimpopolou - *Lynkeus*
Oshri Even-Zohar - *Motek Entertainment*
Vladimir Viro - *Peachnote GmbH*
Michele Buccoli, Augusto Sarti, Massimiliano Zanoni - *Politecnico di Milano*
Pablo Palacio - *Stocos*
Antonio Camurri, Stefano Piana - *Università degli Studi di Genova*
Jean Marc Matos - *K. Danse*

Cover
Moments of the dance performance by Instituto Stocos during the dissemination event in Madrid (*Naves Matadero - Centro Internacional de Artes Vivas*, 17-18 April 2018). Photo credits: Pablo Palacio.

Layout and design
Alessandro Ingrosso, Stefania Servidio
graficheria.it

Print edition
Triograf Srl Roma

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 688865

FOREWORD

Dear Readers,
Welcome to the second issue of the annual WhoLoDancE newsletter, which gives an account of the work carried out in the second year of the project, concluded in December 2017, and provides some insights on the new research directions and dissemination initiatives scheduled for this last year of activity.
As a general view, the second year has seen the consolidation and further enrichment of the digital usage of the dance movements collected through the motion capture sessions held in the first year. All these outcomes can now be accessed on-line through the *WhoLoDancE Movement Library* (page 8). Here, a newly implemented *synchronization technique* (page 24) has allowed the automatic alignment of motion capture, video and audio tracks for all performances in the database.
Additionally, the project has seen the fulfilment and release of the first prototypes of innovations, which were some of the key goals of WhoLoDancE: the *blending engine* (page 21), allowing to assemble movement segments from different genres in new diverse choreographies; the *similarity search engine* (page 12), enabling to search and retrieve movements similar to the one of interest, in regard to user-defined criteria; the *Choreomorphy* tool (page 17), offering users a virtual reality-based experience for visualising their body movements in real-time through a range of different avatars, by the use of motion capture devices; the *movement sketching* tool (page 19), which aims to provide the possibility to record, display and analyse users' dance movements with basic low-budget devices, with special regard to dance students and practitioners. All applications are currently being integrated into a comprehensive system, namely the *WhoLoDancE Framework* (page 31), providing both a web-based access for basic tools as well as an integration based on the Unity gaming engine, for real-time interactive applications.
Worthwhile mentioning, the newsletter reports also on the important work of elaboration of *movement principles* (page 27), i.e., a set of qualities recognising the dynamic and expressive properties of dance movements, together with the *use-case scenarios* identified for the WhoLoDancE tools.

^ A dance student experimenting with the Choreomorphy tool during the WhoLoDancE Seminar in Madrid (*Naves Matadero - Centro Internacional de Artes Vivas*, 17-18 April 2018).
Photo credits: Amalia Markatzi.

As final contribution, the newsletter provides an outline of the *exploitation strategy* envisaged for the further expansion of the project (page 33): the creation of a digital platform for learning as well as entertainment, creative and didactic purposes, namely *WholoMove*, where artists, teachers and students could share their work, learn, interact and receive credits for their contributions by having recourse to a special token as a new form of digital currency usable within the project.
In the last pages, you can find an overview of the upcoming *dissemination events* organised within WhoLoDancE (page 36), where our partners will present and discuss the project achievements with the relevant reference scientific and practice communities, or set up interactive workshops and exhibitions for both professionals and non-experts, where spectators will have the opportunity to attend partners' dance performances and have a preview experience of the prototype tools under development.

Hope you enjoy the reading!

Anna Rizzo
Chief Editor

Moments of the dance performance by Instituto Stocos during the dissemination event in Madrid (*Naves Matadero - Centro Internacional de Artes Vivas*, 17-18 April 2018). Photo credits: Pablo Palacio.

INDEX

WhoLoDancE: an overview

ANNA RIZZO
Lynkeus

The WhoLoDancE *(Whole-Body Interaction Learning for Dance Education)* project has been conceived as an unprecedented approach to dance, leveraging breakthrough technologies to digitalise, preserve and convey the European dance cultural patrimony to bring a profound innovation in the way this longstanding art is traditionally created, learnt and taught. This passed through the recording of a wide range of dance movements with motion capture techniques, their annotation and enrichment, and the development of cutting edge algorithms to explore, analyse and re-elaborate these data at the service of dance learning and choreography, together with the implementation of virtual body exploration experiences able to offer dance practitioners a new way of self-reflection-based experimentation.

To achieve these goals, the project workplan has been articulated in three phases: *Phase I* (Jan-Dec 2016), dedicated to the acquisition of preliminary knowledge and movement data; *Phase II* (Jan-Dec 2017), directed to the definition of emotional and music-dance representation models and the preliminary deployment of data-driven and model-driven analysis software; *Phase III* (Jan-Dec 2018), aimed at the final delivery of data-driven models and tools, together with visual interactive user interfaces.

Λ Figure 1. Rosamaria Cisneros (c-DaRE, Coventry University) performing flamenco during the third motion capture session (Schram Studios, Amsterdam, 10-20 July 2016).

Motion capture, knowledge acquisition and semantic models: laying the basis of WhoLoDancE

The first year of activity has been pivotal to create the ground knowledge and data on which to build our data-driven tools. The Consortium planned three motion capture sessions, the former held at *Casa Paganini Infomus* research centre of University of Genova (March 2016), the latter two taking place at *Motek Entertainment studios* (May 2016) and *Schram* Studios (July 2016), Amsterdam. These led to the production of a substantial volume of kinetic material representative of the four selected dance genres, i.e., ballet, contemporary, flamenco and folk Greek dances. This year has also been fundamental for the acquisition of preliminary knowledge from end-users through evaluative sessions and interviews with dance experts inside and outside WhoLoDancE, necessary for a first definition of different users' profiles and use-case scenarios.

Beyond *Phase II*: prototype finalisation, definition of new representation and learning models

The second year of activity has been dedicated to the delineation of models for emotional and music-dance representation, as well as data-driven and model-driven analysis software. Particularly, the first half of *Phase II* has seen the finalisation of a web-based user interface for the data repository, namely the *WhoLoDancE Library of Movements* (WML), further supplemented by the *annotation* tool, designed for annotation and segmentation of raw mocap data by dance partners. This manual process, which consists in the *characterization of captured dance sequences in respect to movement qualities and their segmentation in correlation with rhythm*, has been essential for a variety of applications, including similarity search, movement sketching and multi-modal rendering, and constituted the base for the training of machine learning-based algorithms for automatic annotation and segmentation, which will support the enrichment of further movement data uploaded in the future. Additionally, a second improved version of the *blending engine*, dedicated to *assembly of movements in new sequences*, was released thanks to the collaborative testing work of technical partners and precious users' insights from dance partners. The research and development work on the *similarity search engine*, allowing to *retrieve segments similar to a given one of interest*, has led to a further refinement of the algorithm and its extension to high-level features involving movement expression. Regarding volumetric displays for virtual body interaction experiences, *Phase II* has seen the implementation of the *Choreomorphy* interface, which allows users to *watch themselves within different avatars in real-time*. In addition, the Consortium successfully released the first prototype of the *movement sketching* software, meant to provide non-professional users with access to the library from low-budget devices, enabling to create new dance sequences and query the repository in search of similar segments. Lastly, *Phase II* brought the definition of *dance movement principles* and *dance learning scenarios*, which have significantly contributed to define innovative approaches to dance teaching and learning. Combining information and ideas derived from technical developments of the first half of *Phase II*, a number of interesting use-case scenarios of educational and creative activities were also developed, highlighting where each tool can represent an added-value for the teaching and learning of dance. *Phase II* also witnessed the conception of an ambitious exploitation plan, envisaging the creation of a multi-sided platform where to extend the WhoLoDancE proof-of-concept to other types of dance genres and adjacent activity areas, while a working group dedicated to related IPR issues is currently examining the proposal of an automated, blockchain-based management of copyrights and contractual agreements regarding the digital use of movements and related technologies.

Λ Figure 2. Moments of the performance by K. Danse during the Metabody-WhoLoDancE dissemination event (*Centre Culturel Bellegarde*, Toulouse, 15-16 December 2017). Photo credits: Fabien Leprieult.

Striding towards the goal: a challenging year underway

The last project year is seeing WhoLoDancE partners finalising their efforts towards the refinement, validation and integration of the developed tools into a unique, seamless comprehensive framework, where users will be able to take the best of what has been attained for professional or personal use. The platform, namely *WhoLoDancE Framework*, will rely on two different access modalities: on one side, a web-based front-end dedicated to distributed, low-end tools accessible from any personal device; on the other, a remote access modality based on the renowned Unity gaming engine, allowing performance of real-time, interactive and immersive applications. In the last months, the Consortium has been intensively busy preparing some final interactive exhibitions, composed as a mixture of technology-based experiences and artistic performances, to be presented during upcoming dissemination events including *Romaeuropa Festival*, *Festival della Scienza* and EUROMED.

WHO WE ARE

Lynkeus is an independent strategy consultancy specialised in the conception and management of EU-funded projects, with a specific expertise on eHealth, data analytics, data privacy and security. Founded in 2000, it works to identify and promote the best cutting-edge technological solutions to complex socio-economic issues, ranging from innovative applications in the field of eHealth and precision medicine to the adoption of blockchain and smart contract technologies for a number of scenarios. Within WhoLoDancE, Lynkeus plays the role of Project Coordinator, and is responsible for science communication and dissemination activities, as well as exploitation and IPR management.

The **Centre for Dance Research (C-DaRE) at Coventry University** is one of Europe's leading research centers for dance. Its research is rooted in the creation, analysis and publication of diverse dance practices. Its researchers collaborate on an international basis with artists and research organizations, with funding from national and international research councils, trusts and the EU. The team from Coventry University contributes with expertise in dance knowledge and artistic practice and leads the evaluation of the project tools and the development of the relevant use-case scenarios, thanks to the broad and varied connection with the dance community at national and international level.

Motek Entertainment innovates 3D animation and performance capture by developing digital assets, content, formats, custom pipelines and solutions to technical and creative challenges. It excels in real-time interactive environments, immersive experiences, mapping and digital holography. Motek is one of the pioneers in the field of motion capture, with a history of new production technologies and integration of techniques from the world of research, medicine and robotics into affordable, broadcast-ready solutions. The role of Motek in WhoLoDancE involves the creation of the motion capture repository, the development of a standalone dance-related movement blending engine, and the implementation of holographic volumetric display technologies.

K. Danse is a dance company involved in numerous creative projects mixing contemporary dance and digital arts. The choreographic approach of K. Danse develops a contemporary movement language by the dialectic confrontation between the physical body (lived, experienced) and the visual body (seen, virtual). K. Danse develops work in four main areas of activity: performance making, mediation, art-science research, and organization of events. In WhoLoDancE, K.Danse contributed with choreographic knowledge and expertise to the preparation and rehearsal of motion capture sessions, the conception of dance movement principles, the follow-up on motion captured data with annotation and movement segmentation and the organization of several dissemination events.

The **Lykeion Ton Hellenidon** founded in 1911, has played an important role in the preservation of Greek cultural traditions and folk customs but also in the domain of the defence of women's rights. The group presents artistic performances on Greek traditional themes, participating in prestigious cultural events, and constitutes a Consultative Organization of the UNESCO Intergovernmental Committee for the Safeguarding of the Intangible Cultural Heritage. As the other dance partners, LCGW has been contributing to WhoLoDancE by recording movements for the motion repository, segmenting and annotating dance sequences, evaluating project tools and supporting their development, particularly machine learning-based processing methods (annotation, segmentation, beat tracking, etc.).

The **Athena Research Center in Information, Communication and Knowledge Technologies (Athena RC)** was founded in 2003 and aims at the development of scientific and technological research in the sectors of information technology, knowledge, communication and automating production processes. It focuses on the implementation and exploitation results aiming at the production of laboratory prototypes, products and services applied to a variety of sectors such as Cultural Heritage, Health and Earth Sciences, Education and other. Within WhoLoDancE, the Athena RC team is responsible for the drafting or user requirements, the definition of learning scenarios and conceptual framework, the implementation of data management platform, as well as the creation of end-user interfaces including desktop, and whole-body interaction technologies.

Instituto Stocos is a project based on the transfer of concepts and abstractions taken from scientific disciplines into performative contexts, focusing on the analysis, research and development of the interaction between body gesture, music and interactive visual imagery. Instituto Stocos also organises events and publications related to these areas. In WhoLoDancE, Instituto Stocos contributed in the elaboration of movement principles and choreographic strategies, in the recording of motion capture sessions for classic ballet and contemporary dance, developing music and dance representation models and in the organisation of dissemination events and performances.

InfoMus/CasaPaganini is a research centre of the Department of Informatics, Bioengineering, Robotics and System Engineering (DIBRIS) of the University of Genoa. Founded in 1984, its research method is characterized by cross-fertilisation between scientific, humanistic and artistic contributions. The centre participates in many international research and education projects, develops multimedia systems, platforms and applications for industry partners, takes part in artistic productions and museum projects where systems and research results are applied and validated. In WhoLoDancE, InfoMus/CasaPaganini contributes to the development of movement analysis techniques and novel applications to access movement datasets, capture movements with low-cost devices and get relevant information in the database.

Peachnote develops technologies and applications making it easier to enjoy music more, both analytically and creatively. It collects large datasets and applies to them efficient algorithms, data mining and machine learning to make music history and practice more accessible and novel real-time multimedia applications possible. Peachnote publishes web and mobile applications for novice and advanced musicians and participates in research collaborations with academic partners. In WhoLoDancE, Peachnote provides the search and similarity engine for the collected motion data and designs real-time motion-based applications.

Ranked among the top 50 universities in the world (QS ranking) in all areas of science and technology, **Politecnico di Milano (PoliMi)** is the largest technical university in Italy and Europe. The Image and Sound Processing Group (ISPG) has a long-standing experience in the fields of multi-view analysis for 3D reconstruction and behavioral analysis, pattern recognition for video and sound analysis, sound production and processing. In WhoLoDancE, PoliMi is involved in the development of advanced analytics and machine learning-based algorithms arming the platform, such as similarity search, automatic segmentation, beat tracking and annotation.

The WhoLoDancE library of movements: a priceless resource for preservation, research and innovation in dance

KATERINA EL RAHEB
Athena Research Center

MASSIMILIANO ZANONI
Politecnico di Milano

AMALIA MARKATZI
Lykeion
Ton Hellenidon

The WhoLoDancE Movement Library represents a web-based interface designed to leverage the large collection of dance movements recorded during motion capture sessions. The system allows to browse recordings by genre, search for specific performances by movement descriptors, interact with recordings and annotate movement segments. Movement annotation and segmentation by dance experts have been flanked by automatic tools which "learn" from their work, to provide extensive and reliable data to future end-users.

One of the very first steps of the WhoLoDancE project has been to record a huge set of dance movements, performed by professional dancers from K. Danse, Stocos and Lykeion Ton Hellenidon, through a motion-capture system: this has led to the creation of a diverse and articulated archive of movement recordings from classical ballet, contemporary dance, flamenco and Greek folk dance. Besides representing an important patrimony for the preservation of European cultural heritage, the repository was conceived to become a fundamental resource for dance teaching, research and innovation, also serving as a source of inspiration for the creation of new movements in choreography.

At that point, making this motion capture data available to dance practitioners has been a primary project asset. In the course of the second year, the Consortium developed an associated web-based application, namely the *WhoLoDancE Movement Library (WML)*, a user-friendly and versatile interface to access the dance movement database and explore it with a wide range of functionalities.

Browsing, searching, visualizing: the WML application

In the database, each recording is associated with a range of metadata information, including *title, genre, performer* and *date of recording*. Through the application, users can browse recordings sorted by *genre* (Figure 1), but also search for specific performances making use of relevant keywords (*tags*) related to metadata. After searching or browsing, users are redirected to the results page, where they can select the recording(s) of interest. Filtering, sorting and pagination is also supported. A special player allows the synchronized playback of a video together with the corresponding motion capture (MoCap) file, along with a timeline showing the recording progression (Figure 2). Also, users can interact independently with the MoCap and video, e.g., by rotating the 3D avatar in the MoCap to watch simultaneously the back of the performer (from the MoCap) the front of the performer (from the video). The first release version of the application was evaluated by external dance experts during Users' Board Sessions (Milan, 5-6 December 2016; London, 26 June 2017) and diverse audiences from numerous

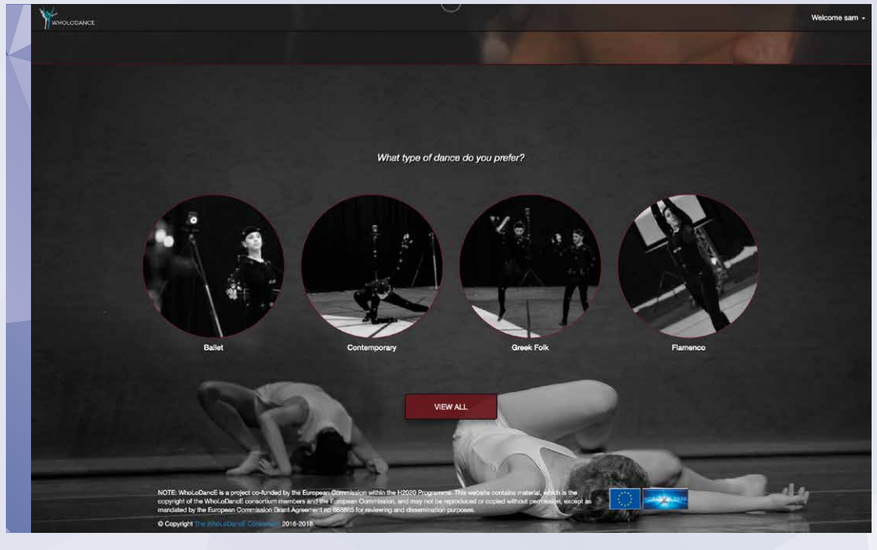

< Figure 1. Web interface view of the movement library entry page, with the possibility to browse motions by dance genre.

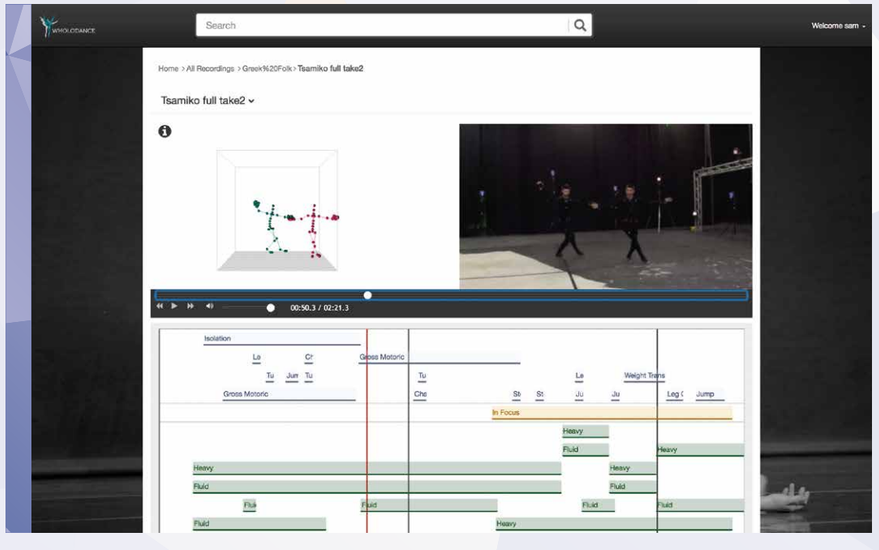

< Figure 2. Web interface view of the movement library with interactive visualisation of a chosen dance performance, with synchronized display of the video and MoCap file, along with the relevant timeline and movement annotation.

dissemination events, receiving fairly positive and useful feedback. The second version, currently in progress, will include additional functionalities, such as the possibility to create private profiles, save recordings of interest and create playlists, in view of enhancing the educational nature of the platform and providing a more personalised experience.

Annotation as the base for applications development

The work of dance researchers is widely based on *dance movement description*, the characterization of dance movements based on a series of dance-related concepts (*semantic descriptors*), and the process of describing dance performances using descriptors is called *annotation*. A dedicated tool was developed within the WML to enable annotation of performances by Consortium dance experts and, in later stage, any WML user. A tabular and a timeline view have been arranged to facilitate viewing, adding, editing and deleting annotations. The tool played a fundamental role in collecting "ground -truth" data from dance experts, creating a dataset of annotations which served as base for further applications. Technology and dance experts worked together to define the best methodology for annotation within one conceptual framework, helping to reach a common understanding of how the movement should be described. As a result, the database has constituted the base for implementation of computing algorithms such as *feature extraction* (i.e., automatic extraction of movement features), *similarity search* (page 12), or *automatic segmentation*, currently arming the platform, as well as functionalities such as *search-by-descriptors*, allowing to search performances by the use of tags.

Machine-learning based annotation: when humans train algorithms

As mentioned above, the WML application allows manual annotation of performances or even single segments. However, in several situations manual annotation is unfeasible or misleading: this is the case of repositories containing huge amounts of recordings, or real-time performance analysis, where annotation must be done just after motion recording. For these reasons, in WhoLoDancE we planned the development of software-based strategies to perform *automatic annotation* of dance performances. In the project we employ three types of descriptors, namely *(1) movement principles* (i.e., higher level concepts such as directionality or coordination), *(2) movement qualities* (i.e., related to how a movement is performed e.g., fluid or fragmented) and *(3) actions* (e.g., step, turn, jump etc.). Among descriptors, it is possible to highlight two main categories, those that are objective and little interpretable, and those that are more abstract and subjective.

Objective descriptors are easily referable to clear and basic elements of the physics or geometry of movement or posture: this is the case of the movement principle *balance* (i.e., a measure of how the line of gravity of the body falls within the line of the supporting limbs) or *symmetry* (i.e., the ability to compute the same movement simultaneously using both sides of the body). Since objective descriptors can be referred to physical or geometrical characteristics, they can be modelled through mathematical formulations, called *rule-based algorithms* or *rule-based models*.

Subjective descriptors refer to more abstract concepts. For most of them, it is hard to provide an objective definition: this is the case, for instance, for movement qualities such as *fluidity* or *heaviness*. Despite the difficulty in providing a formal definition, human can use them to describe movement properties: this brings the idea that a link to some physical or geometrical characteristics exists but resides in the human mind and is too complex to be formalized through a mathematical model. However, it is possible to develop algorithms able to discover this connection by "learning" from given examples: these algorithms are called *artificial intelligence* or *machine learning* techniques. In the initial phase, namely *training* stage, machine learning algorithms are provided with examples that are highly representative of the concept under investigation, and a large set of physics-relates and geometrical-related features are extracted from each example. Then, the algorithm attempts to discover characteristic patterns in the data that can be associated to this concept. The result of this "learning stage" is a model able to automatically annotate previously unseen dance movements.

Given a concept to model, the selection of the set of representative examples is crucial for the quality of the *learned* model. For this reason, it was mandatory to produce a wide set of dance movements annotated by experts. Through the WML, experts annotated dance performances in relation with the concept to train with numeric values, from 1 to 10. If more than one expert was involved in the procedure, a final consensus value was computed with the use of statistical indicators.

As time passes by, the algorithm is further optimized by reiterative training with new human-annotated examples to improve its accuracy. In this way, the machine learning algorithm becomes able to describe a dance performance, in regard to a given set of descriptors, in the very same way a dance expert would do. Without the need of constant time-consuming and tedious work, the database is thus enriched with highly informative annotations, easing user search and constituting a solid base for dance analysis.

The added value of the WML for dance teaching and analysis

Throughout the interdisciplinary collaboration for designing, developing and testing the WML, the Consortium and external experts had the chance to reflect on its added value in dance education and research. For example, the possibility to watch simultaneously a performance from different perspectives emerged as a big advantage in comparison to regular 2D videos and one perspective cameras, especially for dances like Greek folk, which are performed in a circle, usually with long and heavy costumes. For all dance genres, movement recordings have followed a systematic schedule and an associated shot-list, aimed to create a valuable learning context: this has been particularly relevant for Greek traditional dances, where basic steps, variations and improvisations of the leader dancer have been captured for more than 50 dances from all over Greece, allowing to appreciate local differences and nuances. Worthwhile mentioning, the continuous discussion between technologists and dance experts opened interesting research issues, such as subjectivity and objectivity of movement description, and raised new perspectives on the teaching of traditional dances, usually subjected to more conventional approaches.

^ Figure 3. Amalia Markatzi from Lykeion Ton Hellenidon explaining traditional Greek folk dances at the Metabody-WhoLoDancE dissemination event (*Centre Culturel Bellegarde*, Toulouse, 18-19 December 2017). Photo credits: Fabien Leprieult.

v Figure 4. Amalia Markatzi from Lykeion Ton Hellenidon engaging the audience into traditional Greek folk dances during the WhoLoDancE Seminar in Madrid (*Naves Matadero - Centro Internacional de Artes Vivas*, 17-18 April 2018). Photo credits: Pablo Palacio.

A first user interface for similarity search in dance

With hundreds of motion capture recordings, the WhoLoDancE Movement Library is a highly valuable resource for dancers and choreographers, but raises the need of easy tools to navigate it. PoliMi, UniGe and Peachnote worked together on an interface for querying a movement segment in the library and automatically retrieving similar ones, with respect to different criteria, from physical aspects to qualities expressed in the dance.

MICHELE BUCCOLI, Politecnico di Milano | **VLADIMIR VIRO,** Peachnote GmbH

The *WhoLoDancE Movement Library* (WML, page 8), developed through motion capture recording sessions and now organised into a web-based service, offers the rare opportunity to consult a remarkably wide and diverse range of dance motions from different genres (ballet, contemporary dance, flamenco and Greek folk dance at present, and possibly many more in the future). However, such high number of recordings makes extremely hard to navigate and find a specific performance, requiring specific strategies to enable a straightforward and user-friendly browsing.

Similarity search as a natural way to identify common patterns in dance

The WML is armed with a wide range of utilities: among these, the possibility to sort performances according to metadata categories, such as *dance genre*, *company*, *performer*, or *performance name* (Figure 1). Thanks to the annotation tool, it is also possible to search a specific item through relevant terms (*tags*) annotated by project dance experts. Such tools, however, rely on a number of assumptions: the metadata are correct (*no typos*) and unique (*no ambiguity*), the user already knows which performance is looking for, the query can be expressed with tags, and the dataset is fully and properly annotated.

When navigating large multimedia libraries, however, commonly users have no or hardly an idea of what they're looking for. Often they find an item of interest simply by chance, and then search for similar ones. This is mostly the case for *learning scenarios in dance (page 27)*, when a student aims to find different ways to perform the same movement, corresponding to a series of recordings very similar to each other. Investigating the concept of similarity was therefore a crucial task in WhoLoDancE, having in mind to achieve an intuitive way to navigate the WML.

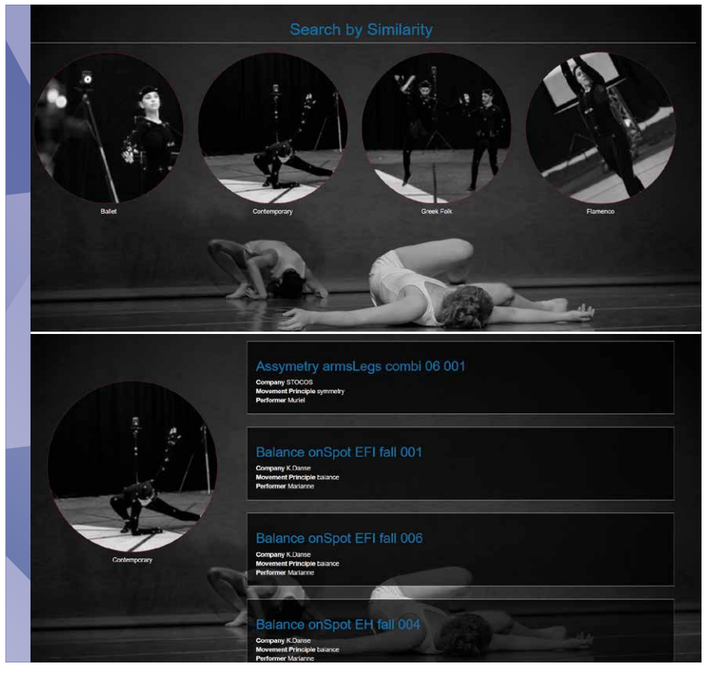

V Figure 1. *Top:* selection of performances by dance genre in the WhoLoDancE Movement Library (WML). *Bottom:* list of performances for a selected dance genre in the WML interface.

The concept of similarity in dance

The concept of similarity has been widely investigated and represents the core of many applications, from search engines to recommender systems. Similarity is usually modelled as a *metric*: given two movements, we assign a number describing *how much* they are similar to each other; the more similar the movements, the higher their similarity value.

On this basis, our task has been to design an algorithm automatically computing the similarity value of two given movements. We may think of similarity as an *absolute* concept: as a matter of fact, if we ask a person how much two movements are similar, the answer would probably be «it depends». In fact, evaluating the similarity of two movements implies choosing some of their properties and evaluating the degree of similarity in regard to such properties. For instance, two recordings of a jump in ballet and contemporary dance are similar in their physical act, but rather different in the style. For these reasons, similarity search requires two subsequent steps: in first place, the definition and automatic extraction of these properties; then, the design of an algorithm to compare such properties and compute the similarity.

Step 1: movement feature extraction

Researchers at *Politecnico di Milano* (PoliMi) and University of Genova (UniGe) designed a set of techniques to capture and numerically estimate some of these movement properties. Their values, called *features*, can be seen as the evolution of properties over time. For instance, from a movement recording we are able to extract the dancer's position on the stage (Figure 2). Using the feature, is then possible to derive the dancer's path throughout the performance. Some features extracted by PoliMi include limb movement *direction, velocity and acceleration*, as well as *angles* between limbs (e.g., between arm and forearm). Such features are usually referred as *low-level features*, since they quantify properties related to the physical (low) level of movements.

UniGe, instead, focused on *mid-* and *high-level* features, related to abstract movement qualities such as *energy, lightness, smoothness, symmetry, coordination* between left and right part of the body, and between feet or hands. Once these features are extracted from the WML motion recordings, they are collected, processed and uploaded to the similarity engine for computing similarity values.

Step 2: computation of similarity scores

Peachnote GmbH developed the core of

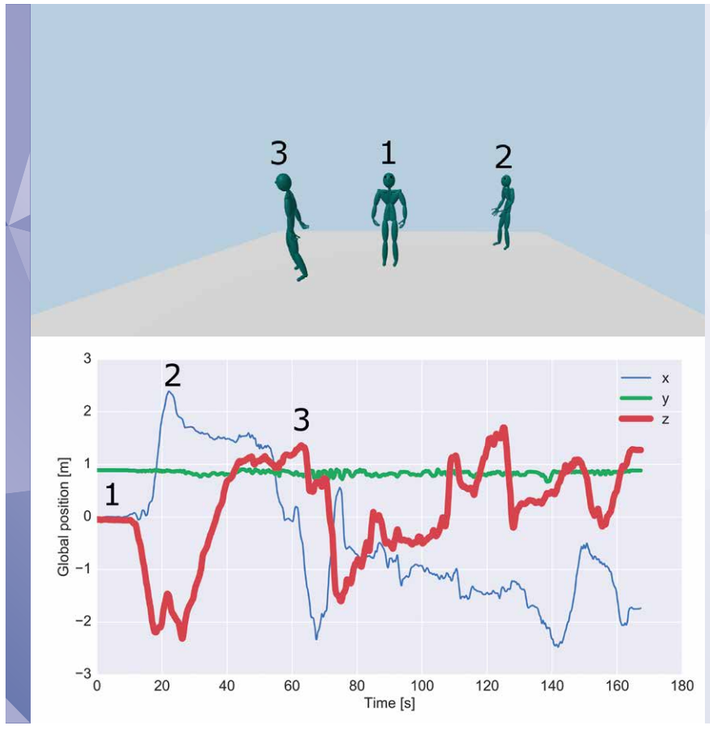

V Figure 2. *Top:* three moments (frames) of a flamenco performance: the dancer's position changes from the centre of the stage (1) to the top-right corner (2) to the mid-left position (3). *Bottom:* features extracted from the same flamenco performance to represent the global position of the dancer on the stage: from the centre (x=z=0 m), to the top-right (x=2 m, z=-2 m), to the mid-left (x=0 m, z=1.2 m). Note that y remains the same, as the dancer does not jump.

this system, the similarity engine, which uses algorithms to find similar patterns based on movement features. The engine can operate on a variable number of features, depending on how many are selected and the desired degree of detail; for example, the finger motions can be omitted if irrelevant to the dance genre.
Each feature is represented by a sequence of numbers (*real-valued time series*). When movement recordings are indexed in the database, similar sub-sequences are grouped together and organized in a hierarchical way, within a tree-like data structure (*clustering tree*). Once a similarity query comes in, the engine walks through the hierarchy of stored sequences (*nodes*) from top to bottom. At each level, the engine compares the query sequence with all nodes and decides which one contains the closest sequence. Then, only the most promising branch of the tree is inspected further, while the others are ignored. In this way, the engine can prune most of distant sequences without comparing them directly, avoiding a lot of unnecessary computations.
For every feature, the system assigns a similarity score to each recording. For example, given a query movement and two features (*global smoothness, global lightness*), the algorithm computes a list of recordings with similar smoothness value, and another one with close lightness values. Then, the engine combines the lists to prioritize those recordings which are similar in as many features as possible (Figure 3). In our example, the engine combines the *smoothness* and *lightness* list and computes a global similarity score for each recording.
The various features may have a different relevance in estimating global similarity. For example, for flamenco sequences, finger motions are extremely important, while they are less relevant for ballet. To address the issue, the similarity engine allows to specify a *weighted template*, reporting how to weigh similarity of different features when combining them. Back to the previous example, we can retrieve a list with the global similarity being composed of 75% smoothness and 25% lightness. The definition of weighted templates is crucial in many scenarios, since it allows *partial-body similarity search*: to search for sequences that share only a similar upper-body motion, we simply set the weight of all features related to feet and legs to 0, effectively disregarding them. A series of weighted templates is available by default, but users can define a customised one when issuing a *query*. Thanks to an optimised algorithm and efficient indexing, the similarity engine can retrieve most similar recordings in a few seconds from the request.

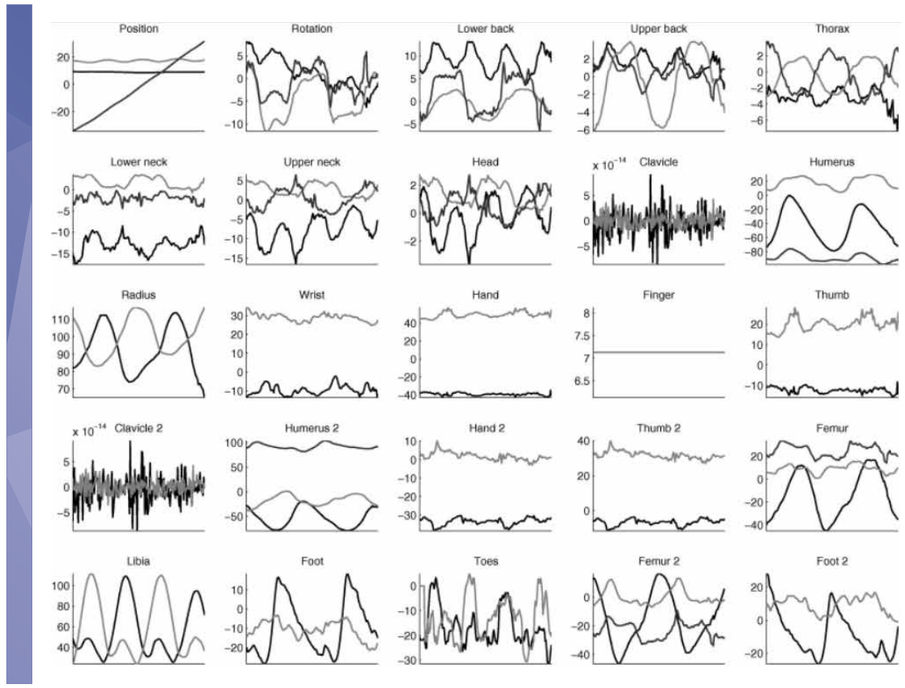

^ Figure 3. The similarity engine analyses and compares different features separately, before combining them into a global similarity value.

The final system: a web-based prototype for similarity search
Once feature extraction and design of the similarity engine had been completed, PoliMi developed a web-based prototype of a user interface for similarity search. The graphic design is based on the WML, with the possibility of filtering performances by genre and accessing to the performances metadata.
To execute the query, the user is asked to select a 5-second movement excerpt from the input performance and a weighted template of features of interest, and then simply click on the search button (Figure 4, top). The similarity engine computes a similarity value for each performance in the database and sorts the performances by similarity value. Then, the interface shows the list of results according to similarity values, from the highest to the lowest. The user is now able to watch the original performance (the query segment) and relevant output results played simultaneously side-by-side, to facilitate the comparison (Figure 4, bottom).

Use-case scenarios: dance learning, creation and research
Although the prototype was initially conceived to evaluate the combination of feature extraction and similarity search, this system can be applied to many possible scenarios, from dance learning and training, to choreography and dance research. For instance, the similarity search tool could be used in combination with the *blending engine* (page 12), to support choreographers in finding new movement combinations similar to the ones just created with the engine. When applied to dance research, similarity search could be possibly employed for finding similarities across different dance (sub)genres, as for Greek folk dance (Figure 5), where several variants, sharing common features, coexist in as many regions of the country. Future versions of the WML interface will include similarity search functionality too, with a recommender system automatically suggesting movements similar to the one the user is watching.

Showcasing the tool: the MOCO 2017 conference
The current prototype has been tested by project dance partners (Stocos, K. Danse and Lykeion Ton Hellenidon), and publicly presented during the demo session of the *4th International Conference on Movement Computing* (MOCO 2017), held in London on 28-30 June 2017. Here, most users were stuck by the its precision and flexibility, and some of them suggested it as a potential source of inspiration. Strongest point of the system is, in fact, the possibility to search within the entire WML across multiple dance genres. This allows to find, for instance, ballet steps similar to flamenco or Greek folk movements, bringing unexpected insights to dance professionals, to be possibly exploited for innovating and diversifying choreographies in unprecedented ways.

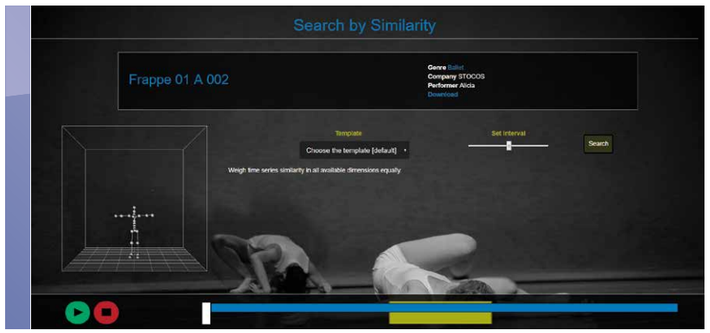

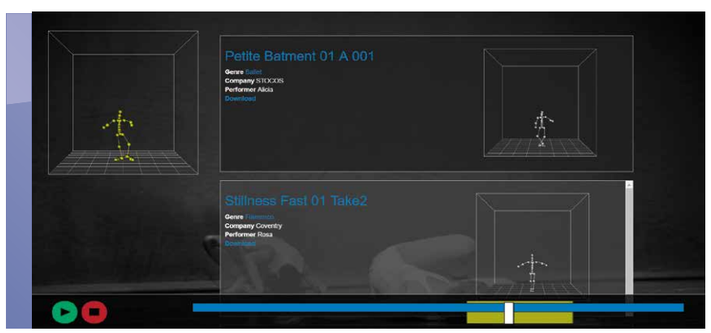

v Figure 4. *Top:* A performance is selected from the similarity search interface, where relevant details are shown. The user can now select the template and the excerpt of interest to apply the similarity search function. *Bottom:* similarity search tool interface with query excerpt (on the left) and two similarity search results excerpts. The player executes the query and the results' excerpts simultaneously, to improve comparability.

Future objectives: automatic segmentation and user customisation

One of remaining flaws of the current prototype is the manual selection of the excerpt, a rather tedious and time-consuming task for the user. As recordings often present one clear movement, the system should be able to automatically identify and set it as the query excerpt.

With this aim in mind, PoliMi and UniGe are investigating the development of automatic segmentation. Its implementation involves the same similarity search development steps: extraction of features and design of an algorithm to break down movements into their salient components. Regarding the former, they have identified and extracted hundreds of possible features, by combining several processing techniques in different ways. These include feature *selection strategies*, used to identify properties with the highest descriptive power, and *aggregation strategies*, employed to combine a wide range of features into few general ones. For the latter, they have are evaluating different algorithms, including those used to segment musical pieces. Various combinations of tentative features, processing techniques and segmentation algorithms have been tested so far, yielding quite different results. As it is hard for non-expert to assess which combinations perform at best, the ball has passed to dance partners. Our dance experts are going to manually segment and annotate a small subset of dance motions, enabling technologists to compare their work with software performance and choose combinations bearing the closest results. In future versions of the similarity search system, users will be able to choose whether to manually select the excerpt, or simply let the system do the work.

A further direction of research involves the concept of similarity. When automatically estimating similarity, we assume a general consensus on the value the majority of users would assign. According to psychology studies, however, we know that similarity is rather a perceptual and subjective quality. A future goal for similarity search will be to achieve a *customised* engine, where the retrieving function will be computed *according to the user's preferences*.

v Figure 5. A dance performance manually segmented by project dance partner Amalia Markatzi.

ROSAMARIA CISNEROS,
Coventry University

KATERINA EL RAHEB,
Athena Research Center

OSHRI EVEN-ZOHAR,
Motek Entertainment

CHOREOMORPHY:

an interactive system for real-time movement display and improvisation

The Choreomorphy tool offers the rare chance of improvising movements and visualising them in real-time in a variety of different avatars and settings, so as to enhance the user's self-reflection and experimentation. Its prototype version has been thoroughly tested by dance experts and practitioners, revealing a high potential in innovating the way dance movement is created and taught.

WhoLoDancE brings together a team of expert technologists, artists and multi-dimensional model designers to develop creative digital tools to facilitate and innovate learning and teaching of dance.

Among them, Choreomorphy offers the unique experience of engaging with dance through a digital platform, while experiencing improvisation through 3D avatar visualisations.

Choreomorphy is a novel interactive system that supports reflective dance improvisation in both the educational and choreographic context and is addressed to dance experts- as well as *amateurs*-users. When motion capture devices are applied all over the dancers' body, the Choreomorphy interface allows to visualize their movements in real time, as well as to record and display them later on. Most importantly, Choreomorphy gives the option of customizing the visualization of the dancer's body and movement, by switching among different avatars and settings in real-time, facilitating self-reflection and experimentation. This enables users to watch themselves in different entities and shapes, as well as focus on specific aspects of their movement such as traces, trails or volumetric space.

^ Figure 1. A dancer trying the MoCap devices on during the Metabody-WhoLoDancE dissemination event (*Centre Culturel Bellegarde*, Toulouse, 18-19 December 2017). Photo credits: Fabien Leprieult. v

Figure 2. Amalia Markatzi from Lykeion Ton Hellenidon experiencing with different avatars during Metabody-WhoLoDancE dissemination event (*Centre Culturel Bellegarde*, Toulouse, 18-19 December 2017). Photo credits: Fabien Leprieult.

Evaluating the tool through the dancers' perspective

The system was evaluated through lab sessions with experienced dance practitioners, as well as in a series of public dissemination events, designed to serve as a formative evaluation of users' experience and interaction, as well as to provide insight about desired features for the interface and avatars (e.g., human vs. non-human shape, human vs. non-human articulation, face characteristics). Also, we aimed at exploring the potential impact of such a tool in an artistic, creative, and educational context for dance practitioners that wish to experiment with visual metaphors and imagery.

Working with dance experts revealed to what extent the experience of dance practitioners seeing their own movement in different real time visualizations can have a deep effect on their practice and cannot be compared to seeing recorded movements through a single avatar.

The contribution of the system in the area of digital applications for dance, according to practitioners, is two-fold.

Firstly, Choreomorphy goes beyond the limits of existing dance education systems, which are based on a mimesis educational approach, suitable mostly for beginner dancers. Choreomorphy, by contrast, is based on imagery techniques to promote a reflexive approach, addressed to more advanced practitioners. The rich controls for fine tuning the avatar, scene and movement visualization characteristics during the session in an uninterrupted way is, according to users, the second main asset of the tool. The system provides an intuitive interface that can be used by practitioners themselves or their assistants (choreographer, teacher, co-dancer) to experiment in one session with different combinations of visual effects.

Unlike most available digital dance learning systems, Choreomorphy is not a system that asks the dancer to imitate specific movements and provide a score. It rather aims at providing a variety of choices in visualizing movement based on imagery examples used in dance practice. In addition, taking into account the importance of imagery in dance practice, Choreomorphy can be applied to digital improvisation sessions to advance research on dance movement, by exploring the impact of body and movement representation in digital environments and draft requirements regarding the avatar characteristics (anthropomorphism, textures, shape, gender, etc.) and their influence on movement improvisation.

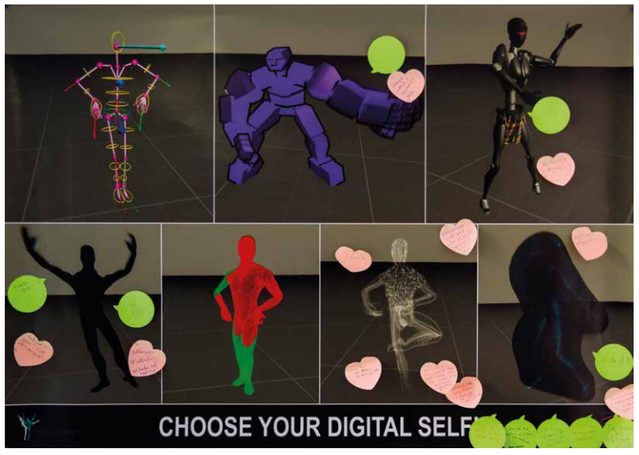

Figure 3. WhoLoDancE poster showing different Choreomorphy avatars, with users' feedbacks, at Metabody-WhoLoDancE dissemination event (*Centre Culturel Bellegarde*, Toulouse, 18-19 December 2017). Photo credits: Fabien Leprieult.

A running prototype of the "movement sketching" tool

STEFANO PIANA,
Università degli Studi di Genova

ANTONIO CAMURRI,
Università degli Studi di Genova

The movement sketching tool represents a way for less equipped users, such as dance students or non professional amateurs, to record, display and analyse their dance movements, accessing the project learning tools from low-cost devices. The tool will in future be able to upload onto the WhoLoDancE library, enriching the repository with possibly infinite kinds of dance genres and styles.

The "movement sketching" paradigm has been designed as a way to make the WhoLoDancE technology accessible to everyone, giving users lacking professional equipment (e.g., motion capture systems, virtual and augmented reality headsets), the opportunity to explore the possibilities offered by the various tools developed within WhoLoDancE, such as consulting the repository of movements, or making use of the blending engine to assemble short dance sequences.

Keeping this in mind, researchers at University of Genova (UniGe) developed a first working prototype of a "movement sketching" tool that, relying on low-cost motion capture devices (e.g., Microsoft Kinect[1] and other RGB-D cameras, Notch Sensors[2] and accelerometers, even those integrated in smartphones), allows users to capture short dance sequences and re-play them in 3D on conventional displays (smartphones, tablets) or low-cost virtual and augmented reality devices. The virtual reality representation makes possible to observe movements from different positions, to evaluate gestures as well as technical and artistic aspects.

Many functionalities, minimum equipment

The sketching tool can be interfaced with a variety of input devices, from RGB-D cameras (e.g., Kinect) to IMU-based capture systems (e.g., Notch sensors). These, in particular, are quite affordable and provide good quality full-body motion captures, enabling to record movement sequences on a mobile device wearing a set of 11 coin-sized sensors (Figure 1).

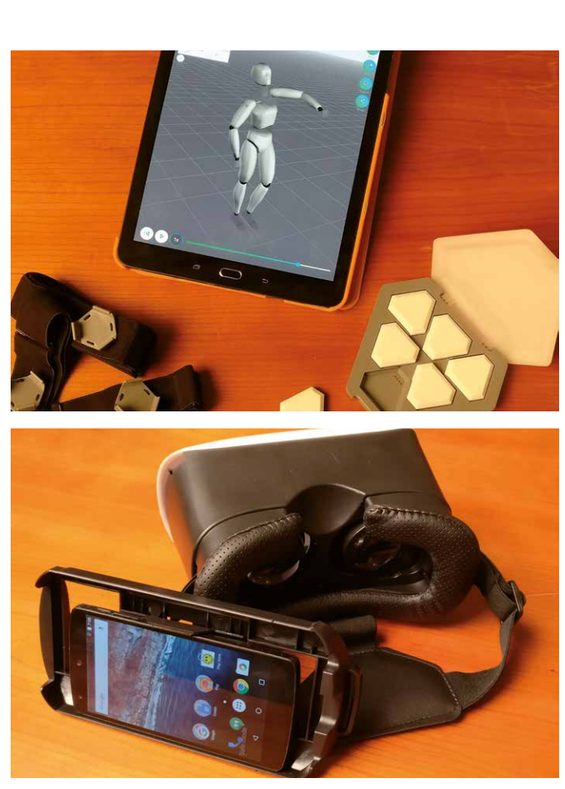

[1] www.developer.microsoft.com/en-us/windows/kinect
[2] www.wearnotch.com

Figure 1. *Top*: one of the supported low-cost capture devices, the Notch sensors. *Bottom*: a low-cost CardBoard VR headset.

One important functionality of the sketching tool is the automatic analysis of movement qualities: once a sequence is recorded, it automatically extracts a set of features, ranging from low-level (e.g., limb energy, speed and acceleration) to more complex high-level qualities (e.g., fluidity, limb coordination, balance). Each of these can be displayed as a graphic (Figure 2, *top*), enabling users to assess the evolution of their performances in regard to different aspects of interest and analyse them individually over time. Additionally, the extracted qualities can be employed for querying the WhoLoDancE Movement Library (WML): making use of the *similarity search engine* (page 21), the tool is in fact able to identify the closest performances among all recordings available in the repository. The kind and number of qualities can be manually selected, allowing to focus on different movement aspects (e.g., technical, expressive or choreographic). At this point, the user can compare recorded segments with the ones retrieved from the repository, with particular focus on the automatically extracted dimensions. An example of this process is shown in Figure 2, *top*: on the left, the recorded sequence; on the right, the most similar segment retrieved from the repository. At the bottom is the temporal evolution of one of the movement qualities extracted from a recording (*green*), and a similar sequence found in the online repository (*red*).

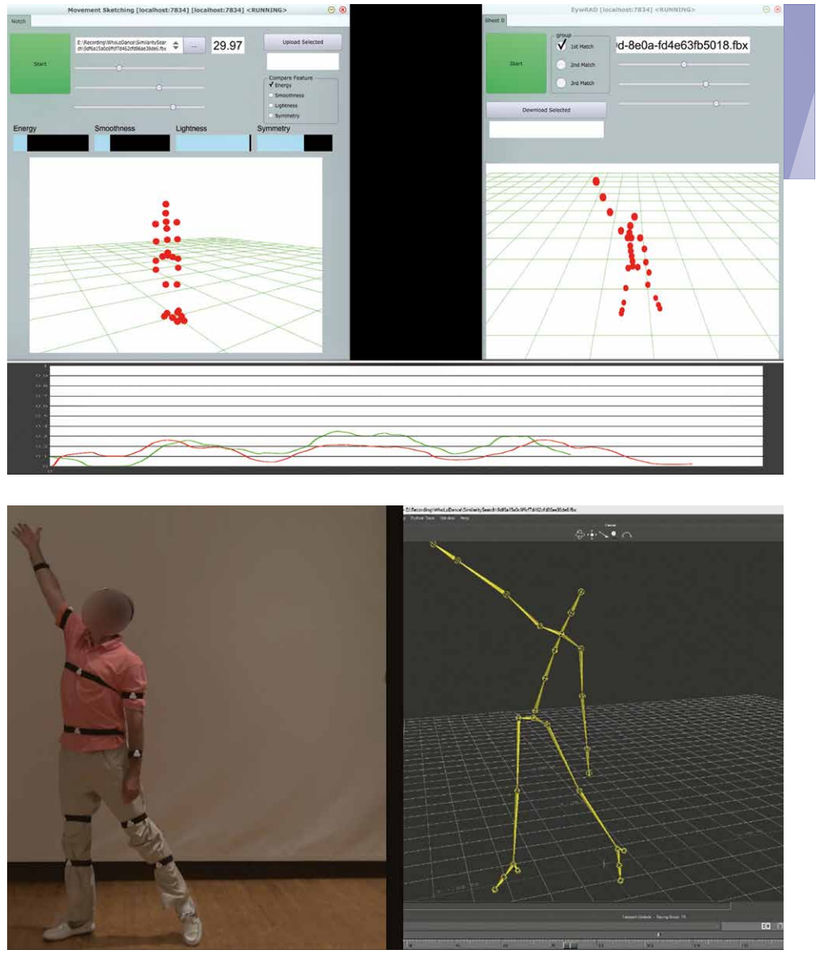

∧ Figure 2. *Top*: first prototype interface of the sketching tool. *Bottom*: recording movements wearing Notch sensors. *Bottom right*: 3D rendering of the Virtual Reality interface.

Diverse applications and innovative perspectives

This tool has the potential to help dance students compare their style and technique with professional dancers', support professionals in refining their style, movements and expressivity and even offering choreographers new ways of creating performances. In the next months, the prototype will be extended to include new quality extraction algorithms, enriched with additional functionalities, such as real-time analysis and feedback, and improved augmented-reality and virtual-reality displays. Finally, as last step, the tool will be integrated with the movement library: in this way, it will be possible to upload recorded sequences onto the repository, enriching the catalogue with new genres and styles.

Our work on "movement sketching" and dance analysis is bringing up new fascinating aspects we are currently exploring. Among the topics under investigation is the possibility to record new movement qualities, particularly expressive ones (e.g., emotional content), as well as expand to other contexts related to movement (e.g., sport, martial arts, fitness, instrument playing).

Motion blending engine on the go

OSHRI EVEN ZOHAR,
Motek Entertainment

ROSAMARIA CISNEROS,
Centre for Dance Research, Coventry University

The WhoLoDancE motion blending engine, firstly released in July 2016, constitutes a fundamental tool for composition of multiple-blended movement sequences. After further work, Motek has come to an improved model with streamlined operations, user-defined layouts, additional blending types and playback options. Enabling to load repositories, assemble full choreographies, create new blends and save results, it now represents a precious asset for choreographers, dance professionals, students and teachers.

The WhoLoDancE motion blending engine has been conceived as tool to combine ('blend') different dance movements into new composite motion sequences, to be applied in teaching and choreography. Its first prototype, presented in *Issue 1*, allowed for the interactive blending and composition of movement sequences from the motion capture (MoCap) data of the WhoLoDancE library, as well as the ex-novo creation of new movements based on the existing motion segments.

Overcoming technical challenges and refining users' features

While moving forward, however, some of the initial expectations had to be challenged. We indeed realized that the sheer amount of MoCap data sets gathered into the digital repository implied fairly long loading times, posing considerable difficulties to the real-time blending of movements in an interactive manner. Enabling blends of every sequence to every other sequence, while allowing separation of each body part, required the optimization of the current software and the creation of pre-cached blends for every blend possibility (i.e., the combination of each body part with each motion sequence). On the users' side, WhoLoDancE artistic partners provided very valuable feedback, helping to make interface colour and layout schemes more intuitive for non-technical people. Following these inputs, we tackled the capability of allowing users to customize the interface to their preferred layout and colour scheme. To maximize interoperability, we also implemented the capacity to export resulting blends as FBX files, enabling users to visualize and treat those as new "building blocks data" for creation of new blends.

The last year of the project will be devoted to streamline the entire pipeline user operations, from the loading of the repository, through the assembly of full choreographies, to the creation of new blends

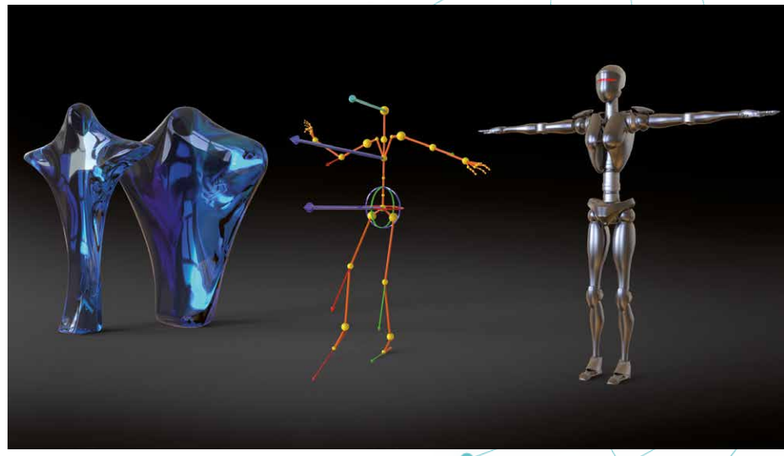

∨ Figure 1. Visualization of avatars developed in WhoLoDancE.

< Figure 2. An example of complex multi sequence blending.

and saving of results. Besides, we aim to enable additional functionalities, such as the streaming of data from the blending engine to the platform integrating the different tools implemented in WhoLoDancE (page 31).

Evaluating the tool: the dancer's insights

To gather further insights, the system has been also evaluated by a range of dance professionals during dedicated lab sessions, to serve as a formative evaluation of the user experience and explore the potential impact of the tool in creative and educational contexts. Also, in early November 2017 the Centre for Dance Research of Coventry University organised a drop-in session with key stakeholders in London considering a variety of perspectives, including the commercial potential of the tool. Both usability tests and interviews to potential users identified there is an actual need for such a tool, particularly in the choreographic context. On that regard, participants highlighted how the blending engine may have a clear effect on how people view cultural dance and may open the possibilities of choreographing with dance genres they might not have considered before, also thanks to the incredible innovative value of the repository of MoCap data. The impact of the evaluative session offered the Consortium clear feedback on the direction the blending engine tool should take in the future, particularly in the fields of machine learning, human-computer interaction, the MoCap community and the performing arts.

A promising market to come: the entertainment industry

For these reasons, we also invested some efforts in defining a potential market for the blending engine, starting from exploring the Motek client network in the entertainment, gaming and education industry, given the conspicuous increase in the demand for motion capture-based assets in those fields (20%-30% on average, up to 80% in the gaming industry) reported since 2004[1]. Apparently, the hottest demand concerns digital effects based on motion capture, such as new techniques for animating computer-generated objects in major film productions, with an estimated global budget of 14.4 billion for 2016, 35% of which for motion capture-based effects[2].

Despite a significant interest in motion capture, this technique is still prohibitively expensive for all but the largest productions. The appearance of lower-end but reasonable-quality inertial motion capture systems is starting to change this trend, but the moment where relevant motion capture repositories will be widely available is still a long way in the future.

In this framework, the dance motion repository generated within WhoLoDancE can be seen as the first batch of many more to come, which may involve additional dance genres, different professions, sports, martial arts and social behaviours. Mining raw motion data and streamlining sequence production involves selecting motions from thousands of different capture sequences, some containing up to 10.000 frames. The selection could be based on the optimal sequence length, motion complexity and sector relevance. On these premises, we believe that the blending engine will position WhoLoDancE at the forefront of MoCap-based digital assets creation and marketplace, successfully extending it to additional motion-driven industries.

[1] ibisworld.com
[2] hollywoodreporter.com

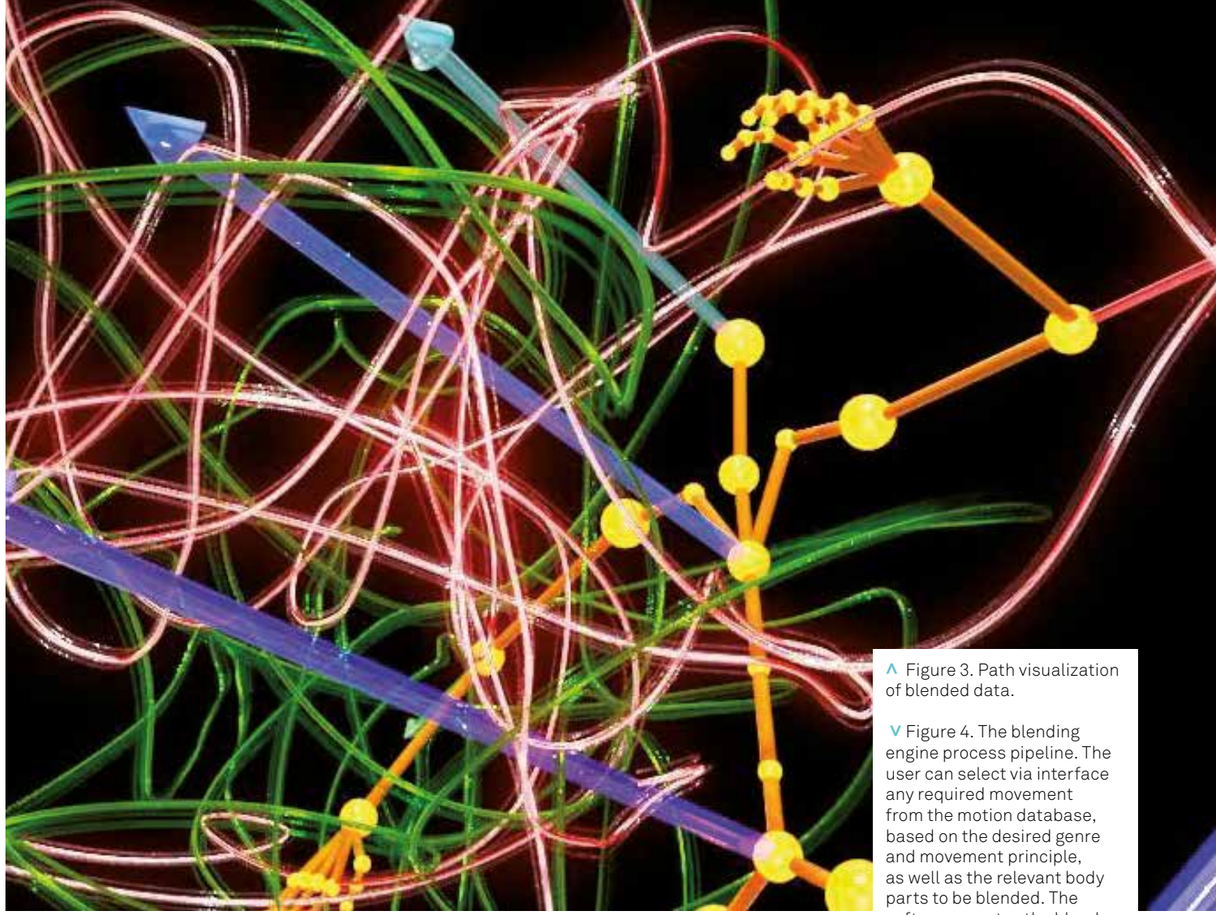

^ Figure 3. Path visualization of blended data.

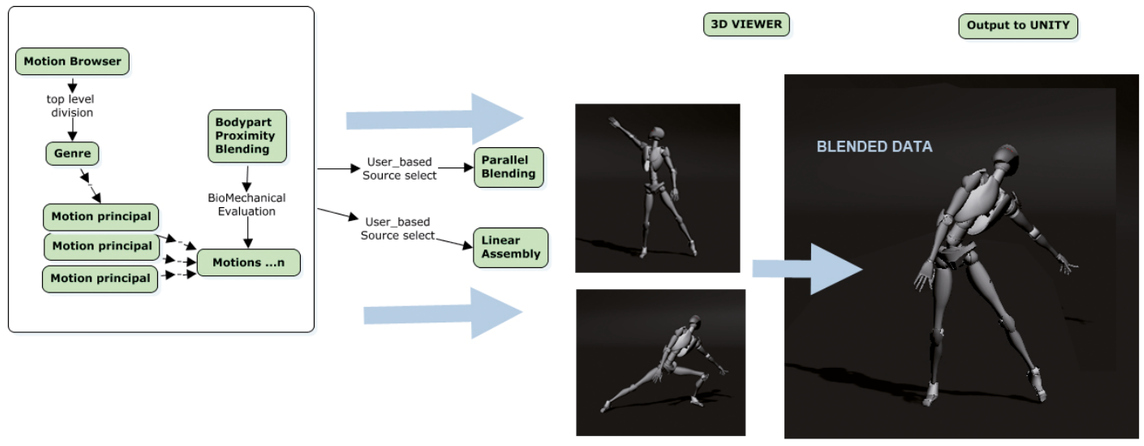

V Figure 4. The blending engine process pipeline. The user can select via interface any required movement from the motion database, based on the desired genre and movement principle, as well as the relevant body parts to be blended. The software creates the blends via the proximity module and evaluates the correct biomechanical relationship. The user can choose to create a linear blend (where sequences will be assembled one after each other) or a parallel blend (where sequences occupy the same timeline). The resulting blend can be displayed in the 3D viewer module or streamed to the Unity environment.

Multi-modality and synchronicity in the spotlight

AUGUSTO SARTI, Politecnico di Milano
MASSIMILIANO ZANONI, Politecnico di Milano
MICHELE BUCCOLI, Politecnico di Milano

The analysis of dance performances is a complex procedure that involves evaluation of motion tracking and music. To capture such complex information, we make use of different acquisition modalities, such as video, motion capture and audio, but the joint analysis of multimodal signals requires them to be aligned in time. In WhoLoDancE, we developed an innovative approach to automatically estimate the correct alignment of video, motion capture and music tracks related to the same dance session.

Motion tracking and analysis have recently received a great deal of attention in the dance art. Here, they are being widely employed for the automatic annotation of performances and the definition of similarity metrics among steps, procedures that represent precious supports to dance teachers and students, as well as incredibly powerful tools for the work of choreographers.

Many approaches to motion analysis are based on video signals, which are easy to capture and have proved to bring good results. However, videos remain a poor representation to grasp the complexity of dance performances, which involve both movement expression and space: for instance, video lack information on depth. To achieve this, we resort to motion capture (MoCap), a more powerful representation of the body pattern in space. Meanwhile, dance movement cannot be disjointed from music: across centuries, music has played a key role in dance performance in several styles and contexts, providing the melodic and rhythmic structure movements interact with. This is the case for flamenco and Greek folk music, where the melody flow and the pattern of beat sequences build the base for choreography. With these considerations in mind, we believe the analysis of different types of recordings (audio, video, MoCap) of the same performance (*multimodal analysis*) allows to better capture the complexity of dance practice, thanks to the variety of information it is possible to extract.

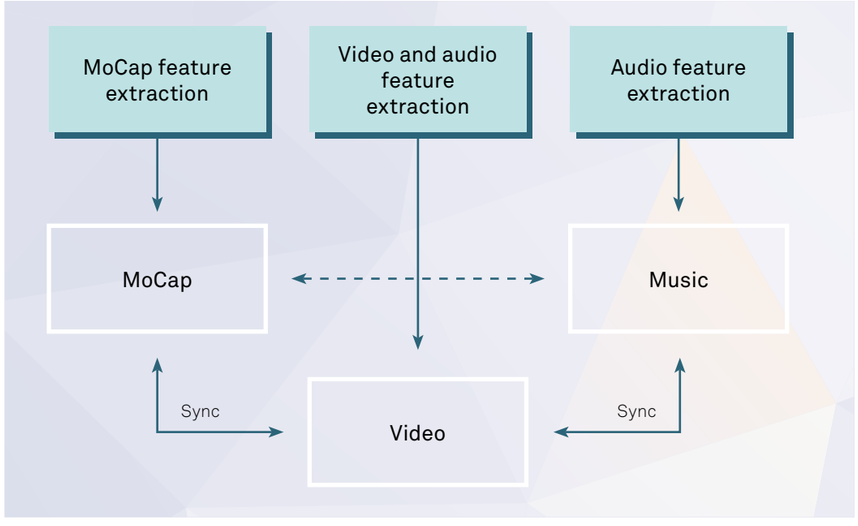

< Figure 1. Global scheme of the multimodal stream alignment. The scheme shows the two-layer procedure we adopted: feature extraction and stream synchronization.

Synchronisation/alignment of multimodal signals

To perform multimodal analysis, it is crucial for signals to be *synchronized/aligned*. Often synchronization is done at the time of acquisition through a trigger clock (e.g., SMPTE), able to automatically synchronize recording machines (cameras, sound and MoCap recorders). Unfortunately, some devices do not support an external clock, requiring a post-recording manual alignment. This can become a cognitively heavy task when the number of recordings is high, making necessary to adopt an automatic technique to estimate the correct offset for time alignment.

Video, MoCap and audio are quite different media, capturing diverse aspects of a performance with specific data formats (*signals*). Nonetheless, it is possible to identify some temporal points common to all three signals, the identification of which constitutes the base for the alignment procedure. Firstly, we find a *representation model* for each format, i.e., a set of temporal descriptors highly representative of the media; then, we execute the alignment of these points of interest. In our method, we split the task into two sub-problems: MoCap-video alignment and MoCap-music alignment. An overall scheme of the method is shown in Figure 1.

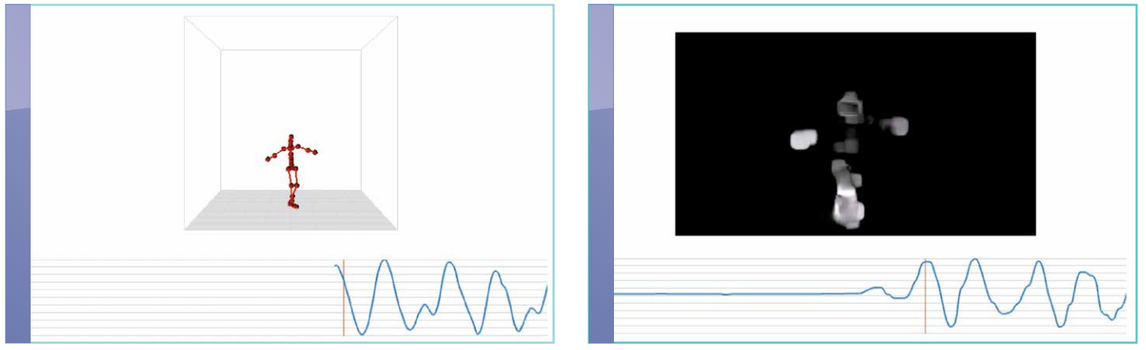

^ Figure 2. *Left*: example of a velocity-based MoCap feature related to the vertical direction. *Right*: example of a velocity-based video feature related to the vertical direction.

MoCap/video alignment

This method is based on the correlation of features extracted from video and MoCap signal. The correlation provides the likelihood of the two signals to belong to the same session (*matching*) at a certain time offset (*alignment*). This task involves sev-

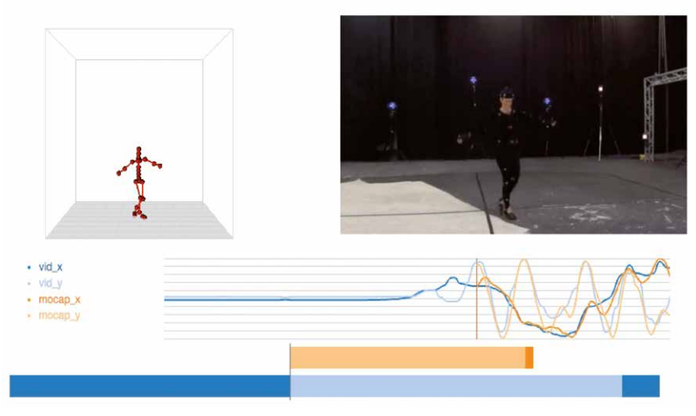

V Figure 3. Example of the four velocity-based descriptors related to the dance session extracted for the alignment procedure.

eral steps. First, we need to identify the most feasible descriptors for the representation of signals. A MoCap and a video related to the same session share an analogous number of movements in the scene. For this reason, we use velocity-based descriptors in horizontal and vertical direction, which are comparable between MoCap and video. Velocity-based descriptors in MoCap are extracted, for each instant of the stream, as the average of all joint velocities (Figure 2, *left*), whereas in videos are calculated as the average velocity of each pixel corresponding to the dancer (Figure 2, *right*). Along the performance, the descriptors generate two curves (i.e., horizontal velocity and vertical velocity) for the MoCap and as many for the video. The alignment is then performed by computing the correlation between the two pairs of curves (Figure 3).

Music/video alignment

During the performance, videos and environmental sound, including performance music, are acquired as different tracks within the same recording. As music is extrapolated from the video recording, the audio track tends to have rather low quality, being degraded by background noises such as audience talking, reverberation of the room and dancers' steps and movements. Given the audio track from a video recording, the audio synchronization aims at finding *which* music piece was played during the performance and its correct *time alignment*.

Before the compass was invented, sailors could navigate through the sea at night just by examining stars and their positions. This is one of the reasons why constellations were invented: to define patterns to help sailors memorize these positions. With much less poetry, we can think of a starred sky as a binary map, filled with 1 where there is a star and 0 elsewhere. This is the first step of audio synchronization: we extract a binary map (*constellation map*) from the audio track and music piece, and then compare these maps instead of the actual signals. The extraction of the binary map is made on the signal *spectrogram*, which reflects the evolution of audio frequencies over time. A variety of signal processing techniques are applied to derive a robust *descriptor* for audio signals, i.e., a number "describing" the signal pattern, which remains stable even if some noise is introduced. As a result, the constellation map of an audio track is very similar to the corresponding music piece. Once constellation maps are extracted, they are stored and indexed as a set of *fingerprints*, unique representations that make extremely difficult to confuse one audio signal with the other. Each fingerprint is stored together with its corresponding instant: back to the sailor example, we take a picture of the sky and store constellations there, with a note about the portion of sky where these constellations are found.

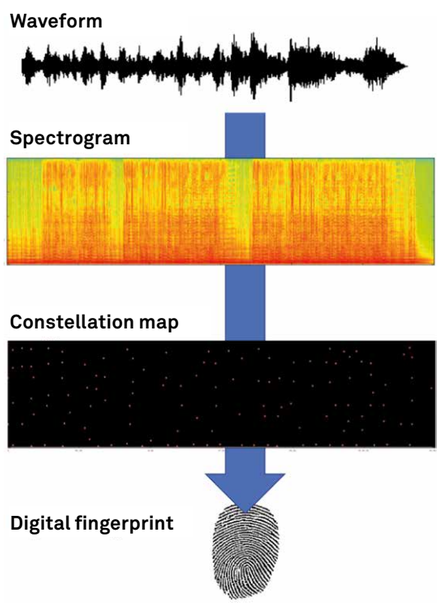

v Figure 4. Outline of the procedure for fingerprint extraction from an audio stream.

In WhoLoDancE, we first collected all music pieces played during recording sessions, extracted constellation maps and corresponding fingerprints, and stored them in the database. Then, we computed a constellation map and relevant fingerprints for each audio track and queried the database in search of music pieces sharing analogous fingerprints. Among all candidate pieces and time alignments, we selected the ones with maximum common fingerprints. By performing this operation over all audio tracks, we were able to synchronize all video recordings with the corresponding music tracks and, in turn, all music tracks with the corresponding MoCap signals (Figure 4).

Towards cutting edge dance learning models: the WhoLoDancE movement principles

SARAH WHATLEY,
Centre for Dance Research, Coventry University

KATERINA EL RAHEB,
Athena Research Center

At base, the WhoLoDancE project is focused on how we might build tools to support, augment and enhance the teaching, learning and making of dance. Each tool is opening up valuable questions about the nature of dance, about how dance is taught, learnt and created, what is particular about different dance genres and where there are common features. The identification of high level movement principles, which underpin all dance genres, and which then lead to a schema of movement qualities that are important for recognising the dynamic and expressive properties of dance movement, forms the backbone of the project.

The aim of WhoLoDancE is to develop advanced technology-based tools for dance learning, to be designed upon the principles of contemporary teaching models. Beside traditional mimicking, these include enhancing the dancers' movement literacy, increasing their ability to analyse their own motions, as well as enriching their movement skills feeding the learners' creativity and musicality, leading to a broader dance experience. To reach this goal, it has been essential to identify a set of properties, namely the *movement principles*, which summarise the embodied skills to be improved in the dance learning process, regardless of dance genre and style. An interdisciplinary focus group, composed of information technology researchers of the Athena Research Centre and dance researchers from the Centre for Dance Research of Coventry University (C-DaRE), has been set up to face the issue. Following a preliminary investigation, the team defined a first list of ten movement principles.

Preliminary movement principles

Symmetry: the ability to perform with both right and left side of the body, arm or leg the exact same movement, simultaneously or sequentially, both in position and in motion. As for the other principles, it also includes the opposite, the capacity to play with asymmetry and isometry.

Directionality: the awareness of body orientation in space. Usually this is derived from the position of hips and torso, but interesting postures might derive from various directions of each

^ Figure 1. *Balance* represents the ability to stand and move in balance, but also out of balance, depending on whether the line of gravity falls within the line of your supporting limb(s) or not, and relies on the awareness of the different vector forces on your body.

< Figure 2. One of the most important skills in every kind of dancing, *coordination* represents the ability to synchronise (or not) different parts of the body, that can move in the same or separate tempos.

body part in relation to a specific space, e.g., the audience, camera or studio.
Balance: the ability to stand and move in balance, but also out of balance, depending on whether the line of gravity falls within the line of your supporting limb(s), or not. It relies on the awareness of the different vector forces on your body.
Alignment – posture stability: the awareness of the geometry of the body (e.g., the sagittal, horizontal, vertical axes) and planes, and how the relation of different body parts and joints create "lines" in the body shape.
Weight bearing vs gesturing: this principle concerns the capacity to distinguish between movements implying bearing weights (e.g., weight transference, stepping) and gestures which simply involve an analogous intention/expression.
Gross vs fine motoric/isolation/articulation: the ability to distinguish small movements executed by specific body parts (e.g., hand, hip, shoulder) without moving the rest of the body, from those moving larger parts of the body as a whole.
Coordination: one of the most important skills in every kind of dancing, it represents the ability to synchronise (or not) different parts of the body, that can move in the same or separate tempos. *Motion through space*: the capacity of progressing through space, towards particular directions, paths etc., vs. dancing on the spot. Also, the use of body as a moving point in space, or as a continuously changing moving volume.
Rhythm and phrasing: the ability to move in particular (predefined or improvised) rhythms. This principle also implies how the dancer's movement is related (or not) to the music and its rhythmical aspects (tempo, time signature, rhythmic patterns etc.).
Stillness: while movement seems to be the essence of dance, a dancer needs to improve her/his ability to remain still, either if it is part of a choreography or interpretation of rhythmical pauses, or an exercise for balance and isolation of body parts. Stillness is related to 'presence', i.e., the ability for the dancer to be bodily aware, and has been investigated in previous interdisciplinary work.

Seeking Consortium's feedback
After conceiving the principles, the team prepared a short questionnaire in view of collecting the relevant partners' feedback, to gain insight from both dance and technical experts' point of view. This was composed of a number of key questions:

1. Define a set of movements/movement sequences that are connected with the movement principle (*dance partners*).
2. How important is this principle for your dance genre-style? (*dance partners*).
3. Give two examples-learning scenarios of teaching the particular principle. Which of the learning principles would you choose? (*dance partners*).
4. Do you see the connection between the movement principle and the low-level features we should capture and analyse? Is there a challenge/obstacle? (*technical partners*).
5. Is there any related work in your institution in capturing-analysing this particular movement principle? (*technical partners*).

Main objectives of the questionnaire were to validate if the preliminary movement principles were appropriate for the different dance genres. Also, as dance partners were asked to provide video shots on movement principles correlating their answers, the questionnaire also served to collect a first base of kinetic material, allowing the motion capture experts to create a relevant shot list.

Learning scenarios for the WhoLoDancE tools
As the tools have developed, each of which offers a different way for dance learners to consider questions about where movement initiates, how movement propagates, how movement 'habits' generate or limit dance, and wider concepts that are fundamental to dance, such as corporeality and perception, the project team has given attention to how the tools could be applied in practice. The C-DaRE, working closely with Athena, has carried out extensive discussions, surveys and interviews with our potential user communities to find out about current practice, needs and curiosities and any perceived barriers to changing practices in the domain of dance. Based on feedback and the team's own knowledge as experts in dance teaching and dance making, a number of learning scenarios have been developed to show the potential application of the tools. These examples are indicative of how the tools may be valuable to different user groups.

Using the similarity search tool
A Greek folk dance group are interested in examining their dance form. They go to the WhoLoDancE library and search for examples. They are intrigued by the motion capture sequences that reveal information about the skeleton moving; they discuss the new insights they gain about the physical properties of their practice. They do another search to compare the 'same' sequences performed by different dancers and ask: *What does this tell us about regional variation? How does different geographical location change the dance? Gender specificity? What new insights do we gain about the movement principles of our dance?* They search by visual similarity to discover other sequences that share similar properties and find connections between Greek folk dance and other dance genres. How does this give them information about the way that dance circulates and develops over time and across boundaries?

^ Figure 4. Traditional Greek folk dance performance. Photo credits: Lykeion ton Hellenidon.

Using the movement library and annotator
An expert flamenco dancer is refining rhythmic footwork and wants to examine her own performance through close analysis of the relationship between the biomechanics of the footwork, the underlying 'narrative' and metaphorical properties of the dance. She goes to the movement library and searches relevant Flamenco sequences by keywords from repository and selects three most appropriate for her exploration. She selects three sequences and annotates these to develop a close analysis of action/quality/dynamic relationship, identifying the movement principles that are relevant for the sequences. She reconstructs the sequences in the studio and focuses on the movement qualities and actions that she highlighted, then returns to edit her annotations, refining following her time in the studio. She is able to uncover new information provided by the chain of movement observation to recreation in the studio; from narrative source to image to movement pattern to performance, and what is particular to her own individual dancing.

Using the blending engine tool
An emerging contemporary dance choreographer is looking for creative inspiration and ideas for how to deepen his knowledge about his own movement ideas to support him working with his group of young performers. He browses the movement library and is drawn to the blending tool as a way of experimenting with novel movement invention. He uses the blending tool to devise four sequences combining actions from ballet and contemporary dance. He is particularly interested in interrupted

v Figure 3. "*Flamenco Dance*" (CC BY 2.0) by Dance Photographer - Brendan Lally.

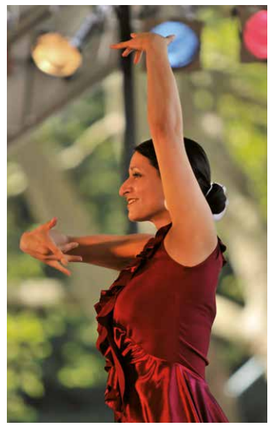

v Figure 5. "*Gregory Catellier*" (CC BY 2.0) by BurnAway

movement, in subverting movement flow and examining how those movement sequences generate new imagery and metaphors that can generate an idiosyncratic movement vocabulary. The movement principles aid him in noticing how he can vary his habitual movement patterns. He introduces those sequences to his dancers who reconstruct and together they devise choreographic sequences and discuss the qualities that are emerging. They search for more sequences using the same search terms – and continue to work iteratively in the studio and on the site.

Using the low-end movement sketching tool

A pre-professional ballet dancer is wanting to improve her upper and lower body co-ordination and fluidity in an adage sequence. She wants to do some private study, away from the class, in preparation for an audition. She goes to the library to find an example of an adage phrase performed by a professional. She wears the Notch motion capture tool to imitate the movement (the shape and upper/lower body coordination) and the tool provides a score that tells her how close her capture is to the original movement and how fluid the movement is in comparison. She is interested in finding other examples of similar adage phrases to explore similarities and differences in movement qualities as related to the movement principles, and how that relates to co-ordination. She records her movement imitating the new phrases using the Notch and compares again her dancing with that of the professional. She uses this information to refine her dancing.

< Figure 7. "*Jenny Hackwell wearing Bloch Leeba Leota*" (CC BY 2.0) by Dancewear Central

Using virtual reality, choosing the avatar and wearing a holographic device

A dance tutor sets a project on his students that focuses on learning differences between dance genres and quality of movement by dancing in close relationship with an avatar. The students choose an avatar (e.g. mechanical, blob, trails) and go to the library to search for a movement sequence that they think is 'typical' of a chosen dance genre. They go into the studio and reconstruct the sequence. Using a holographic device, they then dance 'with' the chosen avatar projection of the sequence to work on embodying the movement qualities, rhythmic structure, actions and movement through space. They are asked to self-assess their own level of 'accuracy' in comparison with the avatar and how the different avatars affect their perception of movement dynamics in relation to their body and space. From both an objective and embodied perspective, the students are asked to analyse the movement principles of their chosen sequence and discuss differences and connections of embodying the different avatars.

Each scenario illustrates the potential for the tools to support learning at different levels and can enable learning to happen in the time needed by and determined by the learner, who might be using the tools in their own time, in their own space, learning alongside and in support of the learning with a teacher/choreographer. The introduction of multi-modal tools and technologies can prompt new questions about dance whilst producing robust information about the complexities of dance, how dance is created, learnt and taught.

v Figure 6. "*dance_120*" (CC BY 2.0) by victorio + the camera

An integrated platform for digital resources: the WhoLoDancE Framework

MASSIMILIANO ZANONI, Politecnico di Milano
OSHRI EVEN-ZOHAR, Motek Entertainment
ANNA RIZZO, Lynkeus

WhoLoDancE developed a variety of applications with different technical needs and performances. To provide users with a unique solution leveraging all available functionalities, the Consortium implemented a three-layered comprehensive framework, including low-end, web-based applications as well as high-end, interactive applications taking the Unity 3D™ game engine as its core base, to cover all current application scenarios and possibly expand them in the future.

With the project termination approaching, the WhoLoDancE Consortium deemed necessary to develop a unified solution where to integrate all digital resources under implementation, to ultimately provide final users with a comprehensive resource where to search for dance sequences, segment and annotate movements, assemble new choreographies, perform avatar-based analysis of users' movements: in one word, the *WhoLoDancE Framework*. This resource is designed to serve all envisaged learning scenarios and dance practices, as well as be further expandable to additional scenarios in the future, by covering different state-of-the-art technologies, easily expandable and re-usable in different contexts.

A layered architecture

With this in mind, the framework (Figure 1) has been designed in three layers: (1) *the back-end*, (2) *the front-*

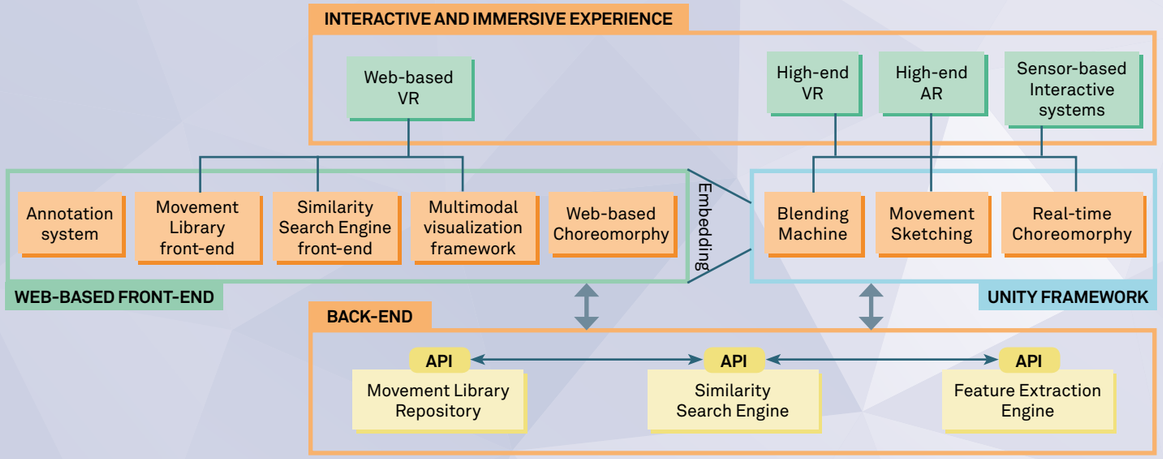

> Figure 1. The architecture of the WhoLoDancE Framework. The *back-end* (yellow) represents the foundation layer, while the *front-end* (orange) includes all applications and tools made available to users, including low-demanding web-based tools (movement library, annotation, similarity search, multimodal visualisation) and high-performance, real-time Unity-based technologies (blending machine, movement sketching). The upper *interactive and immersive layer* (green) incorporates applications for interactive and immersive experience (VR, AR, sensor-based systems).

end, and (3) the upper, *interactive immersive layer*. The *back-end* represents the foundation layer, where most of WhoLoDancE applications are based. It implements a number of services including the storage layer, the feature extraction and the similarity search engines. To be flexible and expandable, tools in the back-end layer are cross-platform and easily accessible from all applications in the layers above, and provided as Internet services, each with a dedicated server and access programming interface (API). The second layer is the *front-end*, an interface covering all applications that will available to users, including the movement library front-end and specific learning-scenarios applications, such as the blending machine. Due to the number of different scenarios to cover, we implemented the front-end considering two main technologies: web-based and Unity-based technologies. The first is appropriate for distributed, low-demanding (low-end) tools not requiring installation but having to run on several platforms: these include the movement library, the annotation, similarity search, multimodal visualisation tools, as well as the web-based version of Choreomorphy. Whereas the second has been adopted for real-time, interactive and immersive applications (*high-end*), such as the blending machine, the movement sketching tool and the real-time version of Choreomorphy. Since the project mostly concerns real-time, interactive applications, the *interactive and immersive layer* stays at the top level. The layer includes all the engines suitable for interactive and immersive experiences, such as VR and AR engines, as well as sensor-based interactive systems.

Unity: a versatile engine for high-performance applications

Unity 3D™ is foremost a game engine, widely adopted as a platform for the creation and distribution of games. Unity is built as a cross-platform engine to develop 2D and 3D videogames and simulations for computers, consoles and mobile devices, and is highly versatile, supporting a wide range of devices (e.g., game consoles, VR headsets and mobiles) and most operating systems (Microsoft, Mac) available on the market, and capable of either reading web-based data from external servers or deploying fully embedded web-based content. Another strong point is represented by its extensive editing capacity and flexibility: in fact, the Unity editor features multiple tools enabling rapid editing and iteration in the development cycles, including play mode for quick previews in real-time. Also, it includes a range of user-friendly tools for designing immersive experiences and game worlds, as well as a strong suite of developer tools for implementing game logic and high-performance gameplay. Unity also supports features for specific needs across genres and includes a navigation system for creation of non-player characters (NPCs) that can intelligently move around the game world. The system uses navigation meshes that are created automatically from scene geometry, or even dynamic obstacles, to alter the navigation of the characters at runtime. Unity is the most largely used of available engines, with a growing community and a wide range of cross-sector applications, and comes at several cost models, depending on the application and development path. In WhoLoDancE, we adopted the Unity Personal license, which is free to use. The different applications developed by Consortium partners, which are partly web-based, partly standalone desktop applications and partly already Unity-embedded, will be integrated with different modalities. One type of integration will involve the full embedding into the Unity platform (as for Choreomorphy), another type will use intermittent web-based communication to the platform (as for annotation and similarity search tool).

< Figure 2. "Gone Home - TV Room" The Fullbright Company (CC BY-SA 3.0), via Wikimedia Commons. *Gone Home* is a first-person adventure exploration video game developed and published by The Fullbright Company, made with the Unity game engine.

An ICO for WholoMove

EDWIN MORLEY-FLETCHER,
Lynkeus

An ambitious collective exploitation opportunity, inspired by the multi-sided platform (MSP) concept, as developed by the 2014 Nobel Laureate for Economics, Jean Tirole, is being examined on behalf of all its partners by Lynkeus, the coordinator of WhoLoDancE.

Given its technological outcomes, moving from proof-of-concept system to becoming a tested operational reality for dance, the WhoLoDancE exploitation idea entails that a similar approach can be applied to other types of dance genres (as well as social dances) and to adjacent areas, like wellness, sports, martial arts, crafts and work ergonomics, body movements for rehabilitation, and even robotics. WhoLoDancE has demonstrated the capacity of recording, reconstructing, and preserving the representation and heritage of priceless movement skills, providing cutting edge digital technology by which a growing community of users will be able to acquire relevant in-depth knowledge and experiment with new ways of learning and teaching, as well as of designing new anatomically sound movements and choreographies.

WhoLoDancE partners have converged on an integrated web-based and Unity-based platform, leveraging the technical developments already accomplished and having recourse to Unity as cross-platform engine of choice. Additionally, it was also agreed that enabling a wider use of less costly low-end devices would be useful, hence, deciding to go for a multi-layered software licensing approach. This now allows to implement a freemium general access policy, providing browsing of and some limited access to the platform for free, in addition to a fee-based access to all WhoLoDancE functionalities. These fees are expected to trigger a feedback mechanism to support further acquisition of motion captures, broadening the content database and ensuring its sustainability, while providing, through blockchain and smart contract applications, an economic return not only to the technology developers but also to the artists contributing to the motion capture sessions. On the premise that the internet has been forcing virtually all industries to upgrade their economic models

v Figure 1. Proposed WholoMove logo and original sculpture. Patrick Steen, "Dansende Golven" (2008), Oostende, Belgium. Photo credits: Reimund Schmidt-De Caluwe (sdc-photo) 2015.

and to transform assets that were not traditionally exchanged into economic goods or services, WhoLoDancE partners have agreed to entrust Lynkeus with the goal of devising how to issue digital tokens, launching an Initial Coin Offering (ICO).

The idea is to follow the trend in which ICOs are becoming the financial foundation of choice for many platforms, where participants exchange useful services paying for them through specially issued electronic coins (tokens), tied to a specific purpose. The assumption is that spreading currency issuance spreads also value creation. Tokens provide a specialised means of exchange and mediation capable of incentivising demand and supply of new digital goods, implementing through a blockchain concertation layer a precise traceability of all transactions, thus securing attribution and property of both the token and the transacted asset. This in turn makes them both scalable and secure, attracting the gathering of further data, generating demand, and lowering distribution costs. By giving permanence to all pertinent digital information in the platform economy, well-designed tokenomics not only enforce full transparency by default, but also allow the alignment of incentives of various players operating on the platform and contributing to its growth. The ambition is to establish the first scalable, decentralised library of annotated motion capture files along with tools to edit, blend, sell, and distribute real world body motions and the tools to leverage them into an algorithmic marketplace. The shared hope is that this platform will allow artists, craftsmen, athletes or anyone else able to perform skillful movements to capture their performance in local motion-capture environments, to upload these contents into a quality controlled, integrated environment, and to share them with a global audience.

The suggestion put forward by the coordinator has been to call the wider project multi-sided platform WholoMove, highlighting by this the determination to go way beyond the dance world alone.
The shared goal is that WholoMove will represent a "first" in the dynamic exploitation of a successful EU-funded H2020 project. It is also meant to act as a highly visible testimony of the financial attractiveness that the outcome of a project may eventually have, given the strong scientific and entrepreneurial background guaranteed by the highly competitive EU selection process, such as the one that WhoLoDancE went through, for being funded in the first place, and being evaluated during its unrolling. This overall technical evaluation has been received as a significant encouragement by WhoLoDancE partners in furthering their project's integration and progress. Now, here is a list of 13 tools which are ready to undergo a process of community evaluation, by which they will be made temporarily available to a wide selection of dance experts among those who will candidate as evaluators. The WholoMove platform will have the ambition to serve all industries concerned with, or making use of, human movements in digital formats, including technology developers and service providers working in those markets.

A selected blockchain system (leveraging the experience accrued in parallel by running the MyHealthMyData EU-funded H2020 project) will be used as the ledger of reference for transacting WholoMove services and executing smart contracts. Through the latter, content creators will define terms and conditions under which their work can be purchased, used, and distributed, while data fingerprinting, time-stamping and other data protection technologies will enforce proper and legal use of content files. Motion-capture services will be advertised through the platform by qualified providers whose technologies will be directly compatible with the WholoMove system, where movements correlated with metadata will allow to capture authorship, content descriptions, file specifications, etc., and provide rich classifications and easy navigation.

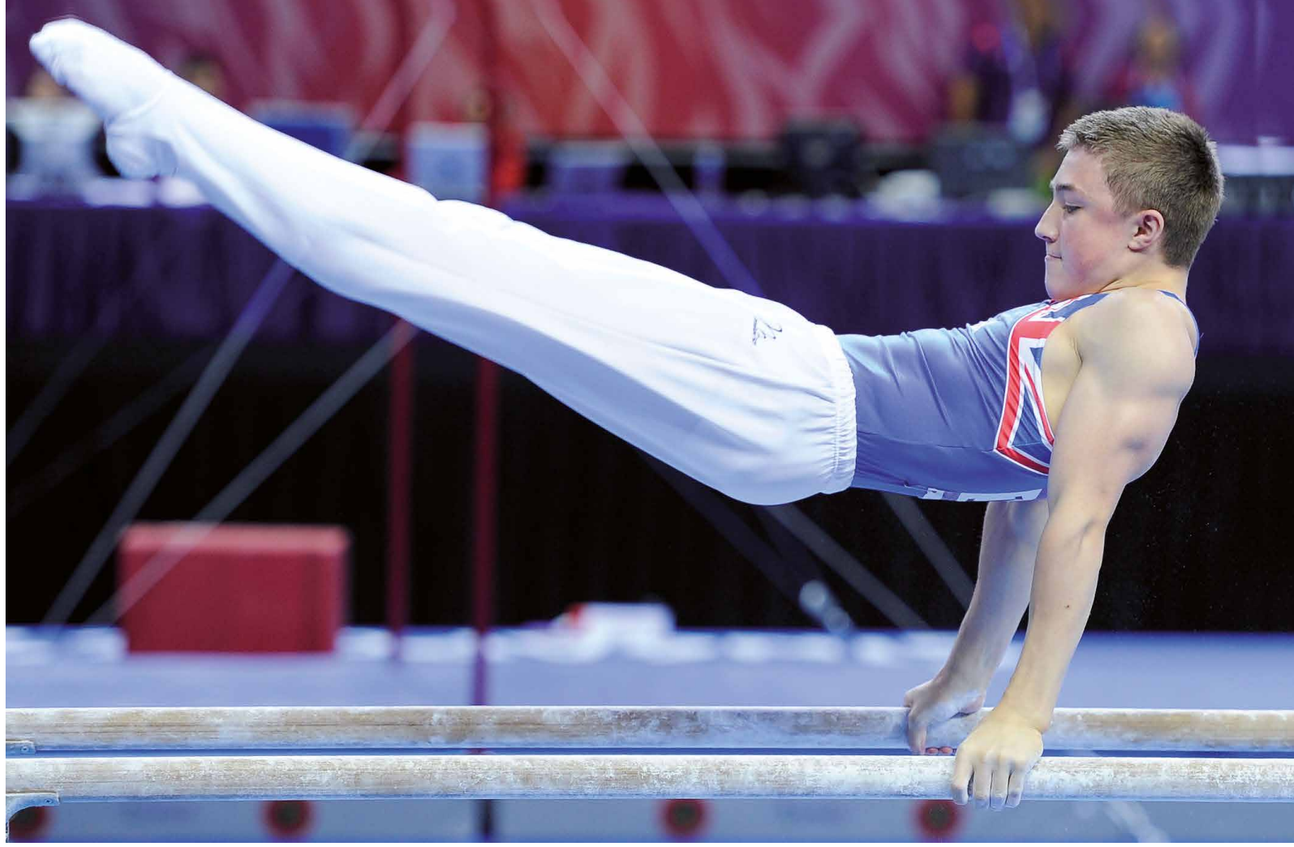

Λ Figure 2. "*Day 8 Gymnastics (22 Aug 2010)*" (CC BY-NC 2.0) by Singapore 2010 Youth Olympic Games. Once the WhoLoDancE proof-of-concept system will be established for dance, a similar approach could be possibly applied to other types of dance genres and to adjacent areas, like fitness, sports, martial arts, crafts and work ergonomics, and even body movements for rehabilitation.

WholoMove will aim at establishing for the first time a marketplace to share motion-capture files in creative and scientific contexts, and at the same time to foster the creation of new content by providing users with tools to search, compare, select and finally blend human movements in novel dynamic sequences. Generating new, anatomically sound movements and forms will in turn activate a content-based network effect increasing the value of the user experience. Musculoskeletal constraints frameworks, encoded in the blending engine, guarantee the creation of sound body dynamics with the highest creative flexibility for users to experiment and, in the process, enrich the content repository. While a detailed business plan will soon be available for clarifying how WholoMove is going to proceed further, a general approach can already be anticipated, which will be based on the "bundle of contracts" theory of the firm. On this basis, Lynkeus is looking into how to launch the ICO, including with it with a fully compliant and transparent token issuing process and the underlying blockchain system.

Dissemination events

WHOLODANCE EXPERIENCE // LAB & PERFORMANCE

ROMAEUROPA FESTIVAL
7 OCTOBER 2018 - Rome, Italy

The event was hosted in the prestigious context of the **Romaeuropa Festival**, among the top four in Europe for the promotion and diffusion of contemporary arts, theatre, dance and music.

It was articulated into a series of interconnected experiences:

- **the interactive laboratory**
  An 'experiential journey' exploring the Consortium-developed digital technologies and their applications for teaching, learning, improvisation and movement creation;

- **the demo performance**
  ("Two Pandoras: the thin membrane between resiliency and fragility"), showing the creative possibilities of motion analysis tools and sound synthesis through the combination of choreographic and musical composition methods;

- **a public debate**
  on future applications in the field of dance and beyond, including the potential development of a technological platform dedicated to the analysis, training and creative exploration of movement-related activities.

WHOLODANCE PERFORMATIVE WORKSHOP

FESTIVAL DELLA SCIENZA
27 OCTOBER 2018 - Genoa, Italy

Accounting for about 200.000 visitors, 300 events and 350 speakers each year, the **Festival della Scienza** of Genoa is the leading science festival in Italy and one of the biggest in Europe, constituting a point of reference to explore science through exhibitions, laboratories, shows and conferences.

The 2018 edition, entitled "Change", hosted a one-day event dedicated to WhoLoDancE, which took the form of a hands-on laboratory where to discover the four dance genres and participate in the various ways its digital technologies can be integrated into new pedagogical processes for teaching, learning and analysis of dance movements, such as motion capture, annotation of movement qualities, segmentation, similarity search, virtual reality or movement sketching.

The event concluded with a performance showing the creative potential of the tools for sonification of movement qualities for the creation of sound-driven choreographies.

wholodance.eu
info@wholodance.eu