GRAPES white paper

Based on deliverables D1.1-D3.1

December 2020

Grant Agreement number	860843
Project Acronym	GRAPES
Project title	Learning, processing and optimising shapes
Start date of project	01/12/2019
Duration of the project	54 months
Coordinator	I. Emiris (ATHENA)
Project website address	http://grapes-network.eu/
Members	ATHENA Research Center (GR, coordinator), U. Barcelona (ES), INRIA (FR), J. Kepler U. Linz (AT), RWTH Aachen University (DE), SINTEF AS (NO), U. Strathclyde (UK), Università della Svizzera italiana (CH), University of Rome Tor Vergata (IT), Vilnius University (LT), Geometry Factory SARL (FR).

Marie Skłodowska-Curie Actions

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 860843

Contents

1 Summary

GRAPES aims at considerably advancing the state of the art in a variety of fields ranging from Computational and Numerical Mathematics, to Geometric Modelling and CAD, up to Data Science and Machine Learning, in order to promote game changing approaches for generating, optimising, and learning 3D shapes. Research is articulated around 3 scientific work packages (WPs):

1. High-order methods and representations,
1. Algebraic & numeric tools in shape optimisation and analysis, and
2. Machine Learning for shapes.

Concrete applications include simulation and fabrication, design and visualisation, retrieval and mining, reconstruction, and urban planning.

Our 15 PhD candidates shall benefit from both top-notch research as well as a strong innovation component through a nexus of intersectoral secondments and Network-wide workshops.

Innovation and technology transfer rely on the active participation of SMEs, either as beneficiary, or as partner organisations hosting secondments.

This report summarises the objectives of the three scientific work packages, and the advances on the state-of-the-art within the first 12 months of the Project (1/12/2019 – 30/11/2020).

In WP1 we focus on new, alternative algebraic representations of complex shapes, namely high-order meshes and trivariate splines, so as to gain a compact and accurate representation of curved objects of arbitrary topology. We develop robust modellers, based on a toolbox for the fundamental primitives that must be executed extremely fast. One motivation stems from the need of converting raw measurements to powerful representations, in applications where simulations apply to physical objects or large-scale scenes. One subtask is to learn adaptive geometric priors to tackle the inherent ill-posed nature of reconstruction and to deal with defect-laden data.

Shape modification and optimisation under constraints for *Isogeometric analysis (IGA)* and tightly integrated *Design-through-Analysis* is highly desired in modern parametric modellers, targeting simulation and visualisation, computational mechanics and hydrodynamics, and propeller and ship-hull design. Our approach shall lead to reduced costs and turn-around times in product life cycle as well as higher fidelity in modern engineering systems. GRAPES also develops innovative spline technologies, including a locally refinable geometric framework, (not subject to the topological restrictions of NURBS), subdivision surfaces of arbitrary topology, and barycentric rational interpolation as well as ML for exploration of the associated design space. Another representation task is 3D reconstruction, given a set of 2D images, which is an algebraic version of the inverse vision problem.

The success of ML on notoriously difficult image analysis tasks has brought an overwhelming trend for applying such tools to 3D data. We devise new methods for learning 3D models relying on deep Neural Networks (NN), aimed at compact encodings for efficient generation, analysis, and search. We advance the state of the art in NN architectures aiming at data-driven geometric design, adaptable geometric representations in conjunction with curriculum learning, autoencoders, and generative adversarial networks (GANs). We introduce standard benchmarks for training and validation, based on an open 3D shape dataset.

At the time of writing of this report only two of the fellows have been recruited for a period longer than three months, and four are expected to complete the recruiting procedure within the first two months of 2021. The recruited fellows have taken relevant courses and seminars in their host institutions and have been familiarizing themselves to the scientific background and state-of-the-art needed for their individual projects. Their research has concentrated mostly on the fundamental aspects of the Project to prepare a sound basis to address the chosen applications fields.

2 Objectives and state-of-the-art

1.1 Work package 1: High-order methods and representations

Simplicial meshes have found wide adoption over the years to represent shapes and there is a rich literature for their generation, mainly because of the prevalence of linear basis functions in computer graphics and animation, as well as in physical simulation by means of low-order finite elements. However, properly capturing the boundaries of shapes, especially those that are highly curved, requires an inordinate number of simplices thus increasing computational complexity. By contrast, the use of high-order piecewise polynomial functions and elements leads to a drastic reduction of the number of elements needed to capture the domain boundaries and, ultimately, it allows to significantly lower runtime[1]. Therefore, high-order methods and representations are very promising and give rise to active research lines. WP1 leverages high-order (or curved) elements to develop new methods for mesh generation and reconstruction of 3D scenes from point clouds, with a view towards applications in geometry processing and design.

The focus of WP1 lies on the following 3 approaches:

- Meshes with curvilinear elements offer more geometric flexibility and improved numerical accuracy compared to more commonly used simplicial meshes. Their generation requires to simultaneously conform curved elements to a given boundary geometry, enforce a smooth and non-degenerate Jacobian everywhere, and keep the mesh as coarse as possible. Although several methods have been explored to accommodate all 3 properties, the generation of valid, practical curved meshes is still a challenge that we address. In architectural design and 3D printing, curved elements such as Dupin cyclides are tightly related to vertex-offset meshes which graciously simplify offset computations, the latter usually being a bottleneck; we extend such splines to arbitrary 3D topology.
- While state-of-the-art approaches in reconstruction are satisfactory at low level of detail, generating high-quality models (roof super-structures, detailed facades, free-form engineering objects, e.g., propellers) is an open problem. The problem's complexity comes from the variety of shapes, made up of linear and curved geometric primitives, as well as from the uncertainty, low sampling rate, and sparsity of input data. We explore 3D reconstruction of large-scale outdoor scenes from raw measurement data, while leveraging new methods in high-order meshing, extracting canonical curved primitives from sparse pointsets, and data-driven ML.

[1] L. Feng, P. Alliez, L. Busé, H. Delingette, M. Desbrun. *Curved optimal Delaunay triangulation.* ACM Tran. Graphics (SIGGRAPH), 37(4):61, 2018.

- Efficient geometry processing for design and manufacturing relies on a number of fundamental primitives, such as computing distances and intersections, that must be performed robustly, efficiently and accurately. We develop such a toolbox for high-order elements. Posed by our industrial partners, in particular RISC-SW for swept volumes in their CSG-like modeller, such questions today limit interactive modellers.

Task 1.1: Generation of valid high-order curved meshes (INRIA) (ESR5)
We aim to develop new mesh generation algorithms that are capable to produce valid high-order curved meshes made of Bézier elements. Based upon existing approaches, in particular the variational approach introduced by Feng et al.[1], we will devise new algebraic methods to have better control on the following key aspects: (1) the preservation of birationality of Bézier element parametrizations under deformation during an optimization process, (2) the detection of surface/surface intersections and self-intersection loci, by exploiting matrix-based representations and numerical linear algebra, in order to preserve sharp features and to control the regularity of the model and (3) refinement-based methods in order to split invalid curved elements during mesh generation.

Task 1.2: Extraction of geometric primitives from 3D point clouds (UB) (ESR3).
New algebraic methods will be explored for extracting optimal surfaces of higher degree and genus from a minimal point set. The extraction of cylinders and cones[2] and of tori from 3 points, two of them being oriented[3] have all been analysed. Extend these works to deal with more involved, but of practical interest, curved elements such as Dupin cyclides (Task 1.3), low-degree Bézier triangular and tensor product patches. Our new methods will be well suited for numerical computations with approximate data, as they are intended to serve the larger goal of improving speed and numerical accuracy in shape detection.

Task 1.3: Circular meshes and cyclidic splines of arbitrary topology (VU) (ESR14).
Regular circular quad (RCQ-) meshes are important in architectural design because of their vertex-offset property. It is well-known that RCQ-meshes produce smooth cyclidic splines composed of principal patches of Dupin cyclides with simply computable offsets. This property is crucial in CAD and 3D printing applications.

Also, 3D circular hexahedron meshes that generate the cyclidic splines composed of principal Dupin volumes will be explored. Locally they represent trivariate orthogonal systems for IGA. The classification of all possible Dupin volume singularities is planned.

Task 1.4: Piecewise smooth reconstruction of 3D scenes from raw point sets (GF) (ESR15).
Motivated by the need of adaptive and powerful representations for domain-specific applications such as urban planning, computational engineering, digital twins of production facilities or construction sites, safer transportation and autonomous driving, we explore the 3D reconstruction of large-scale indoor or outdoor scenes from raw measurement data (aerial or pedestrian LIDAR, dense photogrammetry). Departing from most approaches that reconstruct textured meshes mainly for visualization purposes, our focus is on “3D vector maps” which are semantic-aware, exhibit effective complexity-distortion tradeoffs and are the support of urban or simulation knowledge. Use cases are 3D levels of detail, in accordance with CityGML or other application-specific formats. We learn new error metrics between reconstructed surfaces and raw pointsets that are resilient to outliers and missing data.

[2] L. Busé, A. Galligo and J. Zhang. *Extraction of cylinders and cones from minimal point sets.* Graphical Models, 86(2016), 1-12.

[3] L. Busé, A. Galligo. *Extraction of tori from minimal point sets.* J. Computer-Aided Geometric Design, 58:1-7, 2017.

Task 1.5: Optimised predicate tools for geometric design and processing (ATH) (ESR1).
The main research directions are as follows: Develop tools of basic operations towards modelling and manufacturing, including ray shooting, surface-surface intersections, and computation of distance or tangents to an object. Treat various input representations, such as point clouds, meshes, matrix representation, implicit/parametric equations; the predicates are optimised for each case at a lower level, but representation-agnostic at a higher level. Leverage state-of-the-art algebraic tools to optimise time and space complexity, guarantee robustness, and handle degeneracies. Employ sparse polynomial elimination and the associated combinatorial techniques to exploit input sparseness, a domain in which Athena has strong expertise, e.g., structured resultants[4], matrix and determinantal resultant formulas, and Groebner bases.

1.2 Work package 2: Algebraic & numeric tools in shape optimisation and analysis

For more than 50 years, splines have been the standard representation for free form curves and surfaces in CAD, in particular in the car and aircraft industry, but also in Marine design and Architecture. Pioneered by Bézier at Renault, de Casteljau at Citroën, and de Boor at GM, in the 1960s, B-spline curves/surfaces captivate by their mathematical elegance, numerical efficiency, and intuitive design. Subdivision surfaces were discovered later in an effort to carry over the favourable properties of splines to handle surfaces with arbitrary topology and found wide adoption in the entertainment industry.

Both spline and subdivision surfaces have recently been combined with finite element analysis, leading to the new and powerful paradigm of Isogeometric Analysis (IGA), which allows to design, simulate, and analyse CAD models in a common framework, therefore avoiding geometry meshing and eventually saving costs; WP2 addresses crucial challenges in IGA and its applications. Despite their advantages, current technologies come with several shortcomings; for instance, the ISO STEP 10303 format and NURBS, the current industry standard, cannot guarantee that models be watertight, especially when creating meshes for FEA[5], nor enable local refinement in complex shape oprimisation problems. So, another goal of WP2 is to develop new shape representations that overcome these limitations.

The work-package approach is as follows:

- Local control and precision: In standard splines, refinement can be used to facilitate local control, and different methods have been proposed in the recent past[6] even in the multivariate setting. However, this does not increase local precision, needed in engineering; in geometric design, much better control over a local detail would be possible, if the control points were interpolated by the shape. Multi-degree splines and barycentric rational interpolants provide solutions to these problems but have been explored only in univariate setting. We will develop bivariate versions of these concepts and consider other techniques for improving local control and precision. We will investigate the properties of these representations and study their applicability in IGA and free-form design.

[4] I. Emiris, A. Mantzaflaris, E. Tsigaridas. *On the bit complexity of solving bilinear polynomial systems*. Proc. ACM ISSAC, pp.215-222, 2016.

[5] Sandia Labs reports that time spent fixing water-tigtness has increased from 73% (2005) up to levels as high as 90% (2017).

[6] T. Dokken, T. Lyche, K.F. Pettersen. *Polynomial splines over locally refined box-partitions.* Computer Aided Geom. Design 30(3):331-356, 2013.

- Continuity: Subdivision surfaces and solids are not C^2-continuous at extraordinary vertices and become numerically unstable in their vicinity, which limits their use in numerical computations, despite their favourable capability of describing shapes with arbitrary topology. To fully exploit their potential and make subdivision surfaces and solids suitable for IGA, we will explore new shape representations, based on polynomial and rational functions, that combine subdivision with other local refinement strategies, which guarantee the smoothness of the limit shape and can therefore be used in complex engineering systems, for example for the purpose of shape optimisation.
- The novel curves, surfaces, and solids under study, all have algebraic representations, and so do their projections in subspaces of lower dimension. A natural and scientifically challenging question regards the possibility of recovering these shapes from projections. We will develop algebraic algorithms for answering this question robustly and completely.

On a more technical level, our contributions and methodology include the following tasks:

Task 2.1: Modelling and simulation using analysis-suitable subdivision surfaces and solids (INRIA) (ESR6).
Unlike regular surface patches, irregular regions cannot be described using polynomial-based representations, such as B-splines or NURBS. Instead, the subdivision control net allows for irregular vertices which yield locally non-polynomial surfaces[7]. The approach consists in deriving concise, flexible, and polynomial-based representations that are compatible with subdivision surfaces and solids. Combine flexibility and smoothness towards conducting both modelling and simulation on a common engineering model.

Task 2.2: Shape optimisation via IGA, locally refinable parametric modellers, dimensionality reduction (UoS) (ESR10).

Recent developments in computational methodologies have facilitated designers to solve complex design problems with computational design tools. Nevertheless, performing shape optimisation in engineering context remains to be computationally intensive and challenging, mainly due to the intrinsic complexity of the underlying problems and the limitation of computational resources. Our work focuses on developing data-supported design tools, which help designers and engineers to:

- Overcome the computational burden of simulating complex engineering processes, and
- Generate appearance- and performance-driven designs at the conceptual design phase.

In this context, we have developed techniques like feature-to-feature learning, for reducing the complexity of problems involving many design parameters; introduced the novel concept of intra-sensitivity, for setting viable search spaces for shape optimisation; exploited geometric moments for sensitivity analyses, to aid the designer in the selection of most important parameters of design.

Task 2.3: Multi-degree spline technologies for Isogeometric analysis (UTV) (ESR13).
Extend univariate multi-degree splines to the multidimensional setting, beyond the tensor-product structure. Theoretical properties of multivariate multi-degree splines shall be investigated and proper manipulation algorithms have to be provided. A constructive approach shall be followed by investigating the potential of the Bézier extraction operator.

[7] B. Juettler, A. Mantzaflaris, R. Perl, M. Rumpf. *On numerical integration in Isogeometric subdivision methods for PDE on surfaces.* CMAME 302, 2016.

Task 2.4: Barycentric rational curves and surfaces (USI) (ESR12).
In order to unleash the favourable properties of rational functions in the context of design and analysis, extend the theory of barycentric rational interpolation to curves and surfaces. Devise algorithms for the evaluation, and together with SINTEF, for the local refinement of barycentric rational curves and surfaces and compare them to classical polynomial and spline curves and surfaces with respect to approximation order, stability, and efficiency. In cooperation with GF, develop a novel CGAL package for dissemination and to foster the use of these representations beyond GRAPES.

Task 2.5: Algebraic methods in multiview geometry (JKU) (ESR7).
Given one or several 2D projections, a fundamental problem is to compute the original 3D object: this is the inverse pinhole camera question. For a *smooth* algebraic surface given by its equation the question has been classically asked[8] and answered[9] in the wider context of abstract surfaces. We have obtained some results for non-smooth cases[10] for recovering the equation of the surface effectively and plan to employ discriminants, analysis of algebraic singularities, intersection theory, tropical and multiview geometry to settle the general problem: Decide if the given 2D representations are feasible and, if so, construct a 3D object without assumptions on dimension or topology, and identify the position of camera(s).

1.3 Work package 3: Machine Learning for shapes

In recent years, advances in ML, cheap computational resources, and the availability of big data have spurred the deep learning revolution in various domains. In particular, supervised learning in image analysis have led to superhuman performance in various tasks previously thought intractable, while unsupervised learning based on increasingly advanced generative models were applied to generate high-resolution synthetic images indistinguishable from real ones. Most of these successes, however, were achieved on rather simple data types such as images, videos, time series, which all have a Cartesian structure that trivially maps to tensor representations as used in Neural Networks (NN). On the other hand, there has been surprisingly little progress on 3D objects, despite their obvious importance in a wide spectrum of applications ranging from computational science and simulation, up to fabrication and urban planning, through medicine, augmented reality, and additive manufacturing. To be amenable to deep NNs, previous methods used view-based, or volumetric representations with 3D convolutional filters but cannot deal with deformable (or even rigid) transforms. Even recent techniques on point clouds (e.g., seminal models Pointnet[11] and Pointnet++) fail to exploit geometric structure, essentially ignoring the differential geometry properties. In WP3 we devise representations and architectures for deep learning to handle complex 3D shapes.

Our approach consists of the following 4 general goals:

- **Representational learning of 3D shapes.** The success of deep learning on some of the notoriously difficult image analysis tasks has brought in the past 5 years an overwhelming trend to abandon handcrafted constructions in favour of generic architectures able to learn

8 O. Chisini and C.F. Manara. *Sulla caratterizzazione delle curve di diramazione dei piani tripli.* Ann. Mat. Pura Appl., 25(1):255-265 (1946).

9 V.S. Kulikov. *On Chisini's conjecture.* Izv. Ross. Akad. Nauk Ser. Mat., 63(6):83-116 (1999).

10 M. Gallet, N. Lubbes, J. Schicho, J. Vrsek. *Reconstruction of surfaces with ordinary singularities.* Tech. Report: ArXiV:810.05559, 2018.

11 C.R. Qi, H. Su, K. Mo, L.J. Guibas. *Pointnet: Deep learning on point sets for 3D classification and segmentation.* In Proc. CVPR, 2017.

task-specific features from data. A key factor in the success of ML for non-Euclidean data (text, genes, audio) are the neural embeddings (e.g., Word2Vec, node2vec) mapping to Euclidean spaces, allowing for distance computations and vector arithmetic with a semantic interpretation. We study such encodings obtained via supervised learning, or through unsupervised learning from a large corpus of data. Geometric learning has advanced by introducing architectures for semi- and un-supervised learning (e.g., CapsuleNet), or for handling further representations, such as splines (e.g., spline CNN).

- **NNs for 3D shapes on alternative hardware.** Deep learning typically requires GPU or cloud-based resources, even for running trained models. With ML becoming ubiquitous and diversified in new applications such as geometry, it becomes indispensable to handle compact, efficient, specialised models on hardware such as mobile and hand-held devices, but also distributed systems (blockchain). We explore new NN architectures for shape generation and analysis, such as non-parametric architectures introducing degrees of freedom through reinforcement learning.
- **Generative modelling in geometric deep learning.** The nascent research field of geometric deep learning has obtained notable results in extending success from gridded to non-Euclidean data, such as graphs/meshes and manifolds, in particular for discriminative models such as classifiers. Moreover, interpolation and generation of new shapes seems now within reach by emerging deep learning models. Generative models with various input representations will be developed for geometric design, modelling, and processing.
- **Open 3D shape dataset.** As ML for 3D objects matures, it becomes critical to standardise training and benchmarking. All WP3 teams undertake the creation and curation of a reliable open 3D shape library. In conjunction with existing repositories (ModelNet, ShapeNet) we define new larger splits comprised of aligned/non-aligned objects, so as to manipulate more complex models in a measurable way. We consider weakly-supervised techniques and image retrieval benchmarking in terms of class splits, as well as part segmentation. We minimise additional annotation and, if needed, we follow semi-automated procedures.

On a more technical level, our contributions and methodology include the following tasks:

Task 3.1: Geometric modelling for evolutionary deep learning architectures (SINTEF) (ESR9).
This project explores the link between geometric shapes and neural networks. On one hand, this motivates to represent, generate and optimize shapes with neural networks to approach engineering problems typically governed by partial differential equations. On the other hand, it allows to investigate the behaviour of neural networks using well-studied concepts from geometry, such as discretization of a domain of interest.

Task 3.2: ML for geometric design (RWTH) (ESR8).
The initial research topic is to investigate neural network architectures that provide an interpretable latent representation which can be used for geometric modeling tasks („shape editing“). Through properly defined loss functions during the training phase, it should be enforced that components of the latent vector can be used as intuitive shape control handles. In order to approach this challenge, we started to implement and analyze the DeepMetaHandle architecture. Moreover, different geometry representations (such as meshes, point clouds, implicit functions) have been evaluated with respect to their suitability for DNN-based geometric modeling tasks.

Task 3.3: Deep learning of 3D shapes for retrieval (ATH) (ESR2).

The main research directions are the following: Design deep learning architectures for computing meaningful embeddings of 3D shapes in Euclidean spaces where shape similarity implies vector proximity for purposes of retrieval (and classification) in massive databases. Initially enhance existing architectures operating directly on point clouds, involving 3D space convolution, and enforcing rotation invariance using spherical coordinates. The establishment of benchmark datasets such as ShapeNet, ModelNet and ABC may lead us to re-position the initial goal of coordinating the creation of a 3D shape benchmark set.

Task 3.4: Geometric deep learning for shape analysis (USI) (ESR11).
Investigate the construction of intrinsic generative models (e.g., variational autoencoders or generative adversarial networks) for 3D shapes. Solve for the correspondence between the input and the generated shape, to account for the lack of canonical ordering in the non-Euclidean setting. In particular, study spectral and spatial formulations of intrinsic generative models, as well as topology generations.

Task 3.5: ML and interactive 3D visualisation of temporal point clouds for predicting morphological changes (UB) (ESR4).
Explore deep learning techniques for predicting temporal point clouds, define a model based on shapes, topology, and sampled properties, and study different ML approaches that use this model to predict morphological changes in two different approximations: voxel-based[12] [13] and point-based[1]. To help users better interpret morphological properties, provide a new information visualisation metaphor. To assist users in perceiving surface changes, analyse and study interactive visualisation to improve both the expert understanding and the performance of the ML techniques.

We shall work on the exploration of the state-of-the-art of deep learning techniques for predicting temporal point clouds. Note that the problem can be addressed from the point of view of classical ML approaches and from the point of view of deep learning algorithms. During this year, we have been initially working on the analysis and exploration of solutions based on classical ML approaches. In addition, we have started to analyse and work on different temporal point clouds datasets, extracted using LIDAR, from different types of terrains. We have also analysed the behavior of several machine learning algorithms in the new temporal point clouds datasets. Finally, we are currently preparing a journal article with the intention to submit it at the beginning of January 2021. The goal is to continue our work with the analysis of deep learning techniques and to compare with our previous findings.

3 Publications

The following publications have been made as of 30/11/2020. An updated list of publications can be found at http://grapes-network.eu/publications/.

- **Khan, S.**, Serani, A., Diez, M., Kaklis, P. (2020). *"Physics-Informed Feature-to-Feature Learning for Design-Space Dimensionality Reduction in Shape Optimisation"* in *American Institute of Aeronautics and Astronautics (AIAA), Scitech 2021 Forum*.
- R. Krasauskas, S. Zube, Kinematic interpretation of Darboux cyclides, Computer Aided Geometric Design, Vol. 83, November 2020. https://doi.org/10.1016/j.cagd.2020.101945

[12] A. Brock, T. Lim, J.M. Ritchie, N. Weston. *Generative and Discriminative Voxel Modeling with Convolutional Neural Networks*. CoRRabs/1608.04236 (2016).

[13] Y. Zhou, O. Tuzel. *VoxelNet: End-to-End learning for point cloud based 3D object detection*. In Proc. CVPR, pp.4490-4499, 2018.