

Analysing the Decorative Style of 3D Heritage Collections based on Shape Saliency

KARINA RODRIGUEZ ECHAVARRIA and RAN SONG, University of Brighton

As technologies for 3D acquisition become widely available, it is expected that 3D content documenting heritage artefacts will become increasingly popular. Nevertheless, to provide access and enable the creative use of this content, it is necessary to address challenges for its access. These include the automatic enrichment of 3D content with suitable metadata so that content does not get lost. To address these challenges, this paper presents research on developing technologies to support the organisation and discoverability of 3D content in the Cultural Heritage (CH) domain. This research takes advantage on the fact that heritage artefacts have been designed throughout the centuries with distinctive design styles. Hence, the shape and the decoration of an artefact can provide significant information on the history of the artefact. Hence, the main contributions of the paper include an ontology for documenting 3D representations of heritage artefacts decorated with ornament, such as architectural mouldings. In addition, the paper presents a complementary shape analysis method based on shape saliency to improve the automatic classification of the artefact semantic information based on its 3D shape. This method is tested on a collection of Regency ornament mouldings found in domestic interiors. This content provides a rich dataset on which to explore issues common to many CH artefacts, such as design styles and decorative ornament.

CCS Concepts: •**Computing methodologies** → **Shape analysis**; •**Information systems** → *Ontologies*; Multimedia and multimodal retrieval; •**Applied computing** → *Arts and humanities*; *Fine arts*;

Additional Key Words and Phrases: Shape saliency, Architectural Mouldings, Design style, 3D Patterns, 3D Motifs

ACM Reference Format:

Karina Rodriguez Echavarría and Ran Song. 2015. Analysing the Decorative Style of 3D Heritage Collections based on Shape Saliency *ACM J. Comput. Cult. Herit.* V, N, Article X (January YY), 17 pages.

1. INTRODUCTION

The assimilation of 3D technologies over the last few years is an undeniable fact, as demonstrated by the growing adoption of 3D acquisition, 3D entertainment, virtual reality and 3D printing technologies. As these technologies become widely available, it is expected that the demand for 3D content will increase. Thus, professionals in the Cultural Heritage (CH) domain are in a unique position to contribute 3D content documenting cultural heritage artefacts for a wider set of audiences. In turn, this will unleash opportunities for creative use of this content; underpinning the development of industries which will take advantage of the richness and tangibility of this content for diverse application domains, such as entertainment, conservation, design and manufacturing.

Nevertheless, to achieve this vision it is necessary to address some challenges. Currently, one major problem is the lack of underpinning technologies which can ensure the automatic enrichment of 3D

content with relevant historical metadata. To address this challenge, this paper presents research on developing technologies to support the organisation and discoverability of 3D content. Hence, the research contributes to address one of the major barriers for the ubiquitous use of 3D heritage content, which is the access to the content at the right place and time.

The research focuses on 3D content from the Cultural Heritage (CH) domain. Specifically, we use decorative Regency ornament mouldings found in domestic interiors. This content provides a rich dataset on which to explore issues surrounding design styles, such as decorative patterns and motifs. It also provides a suitable contribution to the Cultural Heritage domain, where many artefacts and architectural elements are decorated using styles relevant to the ideas and technologies defining a time period and place.

The main contributions of the paper include an ontology for documenting 3D representations of heritage artefacts decorated with ornament. This ontology is developed alongside a shape analysis method, based on shape salience, to improve the semantic information that is automatically extracted from a 3D shape. This will allow a variety of users to interrogate a repository using the underlying semantic connections within the content, rather than its low level geometric information (e.g. name or size of a 3D mesh).

The paper is organised as follows. Section 2 presents an overview of related work in the area of semantic analysis of 3D content. Section 3 provides background information in the use of ornament to decorate architectural elements, in particular during the Regency period in Great Britain. Moreover, section 4 describes the main elements of the ontology. Section 5 presents the saliency based shape retrieval method which is proposed to study the shape semantics of heritage artefacts, while section 6 shows the results of the current implementation. Finally, section 7 presents conclusions and further work.

2. RELATED WORK

Once artefacts are digitised using 3D acquisition technologies, such as photogrammetry or 3D scanning, the resulting content is usually kept stored in a computer, hard drive or uploaded into a repository accessible through the web. The latter option is usually the least popular option as 3D repositories are still in its infancy and existing repositories have mainly resulted from research projects, including the Digital Shape Workbench [VISIONAIR Consortium 2015], the Princeton University repository [Princeton University 2015], the Stanford Digital Forma Urbis Romae Project [The Stanford Digital Forma Urbis Romae Project 2015] and the 3D-COFORM repository [3D-COFORM Project 2015; Doerr et al. 2010].

Those heritage organisations which are currently undertaking the 3D digitisation of their artefacts do not yet have suitable solutions for storing, managing and providing suitable access to the 3D content. Hence, 3D content is usually stored in existing heritage organisations' digital assets libraries or collection management systems and the access which is provided via the organisation's website is usually hand-crafted. Furthermore, other portal services such as Europeana [The Europeana Foundation 2015] offer access to 3D models embedded on PDF documents or on websites from heritage organisations. With such a proliferation of different types of solutions, 3D content is at risk of not being found by typical searching mechanisms (e.g. text search) if there is no meaningful information associated with it. Although this problem is common to other content, such as text documents or images; 3D content has inherently richer semantic information associated with its shape than other content types. By semantic information we mean the collection of facts that describe an object within a given domain.

Artefacts from the cultural heritage domain can be described from multiple perspectives depending on the professional (e.g. curator, anthropologist and art historian) providing that description. As with other fields, different views on an artefact might be of particular interest in order to convey its most

important characteristics. According to [Hudson and Nancy 2001], artefacts are typically described in terms of their basic components (e.g. colour, shape, value, texture, space, time and motion, sound and smell), principles of design (e.g. balance, unity and variety, proportion and scale, and rhythm), medium (e.g. physical material used in the work), and style (e.g. a particular period style). It is important to note that only a few of these elements can be automatically extracted from a digitised 3D shape; although it is possible that others can be inferred.

Hence, semantic information related to a cultural heritage artefact might be intrinsic or extrinsic to the artefact's shape. While the creator or the value of an artefact is almost independent of its shape, the design style is related to the shape or visual appearance of the artefact and can relate an artefact to others in the same style. Hence, design style is often used as a descriptor to classify artefacts within a specific time period, place, art movement or school. The required knowledge to determine the design style of an artefact is usually stored in the head of art historians who are trained to observe the shape of artefacts, the fabrication techniques, the material as well as any decorative ornamentation and the types of forms used in that decoration. A further challenge when using design style to describe artefacts is that artefacts might present elements of more than one style. This is because design styles of previous generations have influenced the artistic expressions of later generations throughout history.

Existing solutions for understanding the semantics of 3D content rely on shape analysis methods. This includes methods to segment a give 3D shape into its more descriptive components, using segmentation techniques [Shamir 2008; Lavoué et al. 2012], as well as methods for labelling, classifying, and matching 3D shapes to others. Based on these existing solutions, there have been proposed solutions for semantic enrichment of cultural heritage artefacts using both manual or semi-automatic techniques [Shapira et al. 2010; Kalogerakis et al. 2010; Attene et al. 2007; Papaleo and De Floriani 2010; Hunter and Gerber 2010; Luca et al. 2011; Rodriguez Echavarria et al. 2009; Rodriguez Echavarria et al. 2012; Aliaga et al. 2011; Pittarello and Gatto 2011; Lo Buglio et al. 2013]. These solutions rely on different shape segmentation approaches as well as labelling or manually connecting the shape (or part of it) to a semantic network with relevant information on the knowledge domain. Most solutions attempt to understand what a shape represents, rather than study the shape to extract other relevant information which might help for its classification.

Relevant research in the study of style is the work of [Lun et al. 2015], who developed a method to evaluate the style of 3D objects based on similarly shaped, salient, geometric elements as a key indicator of stylistic similarity between structurally different objects. However, this method only allows for the classification of 3D content which presents discrete elements of style (e.g. not any more than one design style). Another study of interest is presented by [Doersch et al. 2012] who use a large repository of geotagged 2D imagery and seek to automatically find representative visual elements in the architecture of a place (e.g. windows, balconies, and street signs). Related to architecture is the research reported by [Lo Buglio et al. 2013], who proposed a bottom up approach to characterise the features of architectural columns. Their approach is based on the analysis of the low level geometric properties of a collection of columns. Our approach also looks at architectural elements of historic domestic interiors. However, our approach focuses on studying their decoration to extract semantic meaning of the design style of the 3D shape. Hence, the proposed automatic method classifies the 3D content according to artistic style and likely production methods.

3. DECORATIVE MOULDINGS FOR THE ORNAMENTATION OF HISTORIC DOMESTIC INTERIORS

[Meyer 1974] defines the term decoration as the art of applying ornament to beautify artefacts. The style of decoration is usually determined by the aim and material of the artefact to be decorated and, secondly, by the ideas ruling at different periods and among different nations. The idea of decorating artefacts has been used for many centuries. [Phillips and Bunce 1993] trace the use of ornament from

ancient and classical times through to the Gothic and Victorian periods, and culminating in the high-tech computer-aided designs we use today. Moreover, [Jones 1986] presents a comprehensive history of ornament along different periods and nations from antiquity to the medieval and the Renaissance in Europe, India, China, Pacific islands and the Islamic world. In this work, style in architecture can be defined as the form that human expression takes under the influence of climate, materials and fabrication processes at command. Hence, style can provide significant clues on the period and place where decorated artefacts have originated.

Furthermore, artefacts are usually decorated using *patterns*, which are defined as a design composed of one or more motifs, multiplied and arranged in an orderly sequence [Phillips and Bunce 1993]. *Motif* is defined as a unit with which the designer composes a pattern by repeating it at regular intervals over a 2D or a 3D surface [Phillips and Bunce 1993]. [Meyer 1974] makes a further distinction between organic (e.g. natural foliage, animals, human figures) and inorganic (e.g. geometrical elements) motifs. According to these definitions, artefacts decorated with ornament can be classified according to their style which provides information on their historical origin as well as by the type of imagery used for the elements of the motifs and patterns.

The research presented in this paper focuses on a specific type of decorated artefacts. These artefacts are the decorative mouldings which are used to ornament the interior of buildings. We have restricted the style of these artefacts to the Regency style and others trying to emulate this style. The Regency style originated in Britain and is identified at the time in which George IV acted as Regent (starting in 1811), in place of George III who was seriously ill. In 1820, George IV was proclaimed king of Great Britain. However, the Regency period is not constrained by these political events. Experts can trace Regency style expressions from as early as 1780 until 1840, when the Victorian style starts to be identified. Regency is also a unique style which was exported to the United States having a substantial influence on American design at the time of the Federal and Empire eras of architecture and decoration [Parissien 1992].

According to [Parissien 1992], the historical events of the Regency period reflected on the choice of materials and processes used for the architectural detailing of the Regency home. During this period, plaster and composition plaster became a popular choice for decorating the interior of buildings and households. Hence, decorative mouldings include a wide variety of artefacts which are known as plaster mouldings. These include coving, ceiling roses, corbels, panelling, pilaster and arches. Figure 1 shows examples of plaster ceiling roses with a Regency style. The city of Brighton and Hove, in the South East of the England, has many geographical areas with a distinctive architectural Regency style. As such, the interest in the use of decorative mouldings to decorate architectural elements both in interior and exterior of buildings remains high amongst heritage professionals and the public in general.

Historically, decorative mouldings have been produced using *plaster* materials by either running or casting. Running requires of a running mould with a sharp metal profile defining a cross section. This mould is then run through the plaster along its length to produce the moulding. Casting is a more complex process, as it requires of a mould in which *plaster* can be poured and retrieved when this has solidified. Moulds can be produced by different techniques. For instance, the pattern for a moulding's master or *buck* can be carved in wood and this can allow for the production of one or more moulds. The moulds are normally produced using resin or rubber which is poured over the *buck* to create an exact impression. Another method to produce a mould is to carve this directly into the wood with extraordinary delicacy and precision. This mould can then be used with specific *composition plaster* materials and the mould itself has a longer life than other resin or rubber moulds. *Composition plaster* is a commonly used material, and it differs from traditional plaster as it is produced by a mixture of ingredients, such as chalk, glue, linseed oil and resin [Thornton and Adair 1994]. Hence, composition plaster is poured directly into the wooden moulds to produce the mouldings.




Fig. 1. Regency style plaster decoration for building's interiors [Parissien 1992]

The vocabulary of motifs and patterns of Regency plasterers was influenced by Neo-Classical motifs. The more important the room, the heavier and more profuse the moulded ornamentation was. Some commonly used patterns include:

- Anthemion and palmette moulding: pattern comprising alternating palmette and anthemion motifs. The term anthemion comes from the Greek term for *flower*. The palmette is a motif resembling a stylized erect leaf divided into lobes, in the form of a fan or palm leaf, often supported by spirals.
- Egg and dart moulding: pattern of egg-shaped motif alternating with dart like motif.
- Bead and reel moulding: pattern with disks alternating singly or in pairs with oblong beads.
- Fret moulding: Patterns consisting of repeated, linear, geometrical shapes, usually angular, in a continuous band.

For this research we had access to the prestigious Jackson ornament collection provided by the Regency Town House [rth 2015]. This collection comprises of artefacts including architectural mouldings made of plaster and composite plaster, reverse cut moulds and *bucks* in hardwood and soft wood. Examples of artefacts in the collection are shown in figure 2. All of these mouldings are used for decorating wall and ceiling spaces. It was assembled by the ornamental firm George Jackson and Sons through the 19th and 20th centuries. The firm Jackson and Sons, founded in 1780 was key in revolutionising the way in which relief decoration was created. We have initially scanned a hundred artefacts from the Jackson collection using the Breuckmann smartscan device [Breuckmann 2015].



Fig. 2. Example of wooden moulds from the Jackson architectural ornament collection.

The resulting 3D models were stored in a 3D repository and basic metadata stored alongside. The next section describes in more detail the semantic information associated with the 3D content.

4. SEMANTIC ONTOLOGY TO DESCRIBE THE HERITAGE COLLECTION

The Jackson collection is complemented by a database with information such as a unique identifier, description (inc. historical marks), the artist who designed the master or mould and storage location. Also, the database contains information on the locations, when known, where copies of the mouldings are being used.

The conceptual organisation of the current database was mapped and enriched with additional information using the CIDOC Conceptual Reference Model (CRM). CIDOC-CRM [International Council of Museums 2011] provides a formal structure for describing the implicit and explicit concepts and relationships used in heritage artefacts' documentation. As such, the main objectives of this process was twofold: 1) to capture the information related to the processes, materials and visual elements used for this particular domain; and 2) to link to the digital files produced during the scanning process or other information related to the collection (including that existing in other databases).

Luca [Luca 2013] introduces the principle of points of view in order to structure semantically enriched 3D models representing architectural elements. The proposed ontology focuses mainly on the mouldings point of view of architectural elements. Hence, it does not attempt to describe the individual elements of a building but instead focuses on mouldings, wherever these happen to be located, and their production methods.

The resulting CIDOC-CRM model describes the information related to who, where and when the elements of the collection were produced as well as the information related to their production process and decoration style. The most relevant concepts of this conceptual model are shown in Figure 3, which include the moulding itself (E84 Information Carrier), the material (E57 Material), its origin (E52 Timespan and E40 LegalBody), the production process (E12 Production and E84 Information Carrier) and the style displayed by the moulding (E36 Visual Item). The model is generic as it can describe architectural mouldings with different production methods and using different decorative styles.

Therefore, as shown in Figure 3 we model an ‘E12 Production’ event linked to the ‘E84 Information Carrier’ concept, which describes the resulting moulding. In order to model how the moulding was produced, we use another ‘E84 Information Carrier’ concept to describe the mould, which in turn is linked to its own production event (‘E12 Production’). Moreover, the production of the mould (when required) is modelled through another ‘E12 Production’ concept which is further linked to an ‘E84 Information Carrier’ concept describing the *buck*.

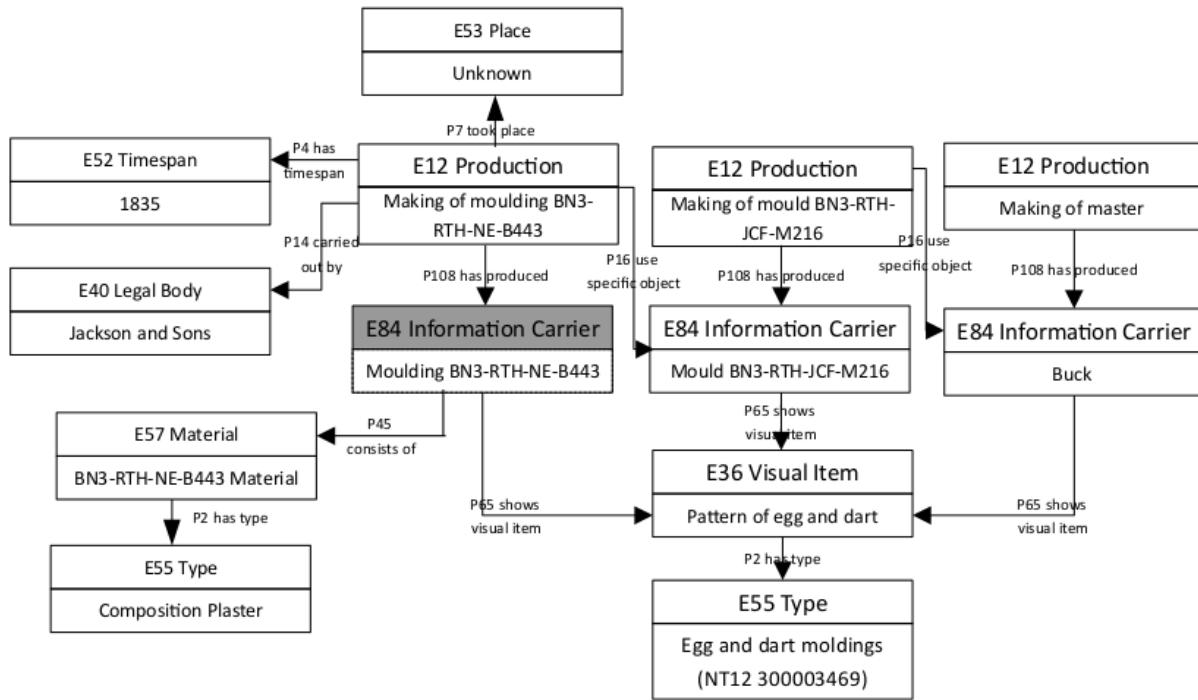


Fig. 3. Part of the ontology describing the plaster moulding production and the artistic style of the mouldings

In addition, all of the “E84 Information Carrier” concepts are linked to the “E36 Visual Item” concept. This concept is used to describe cross sections, ornament patterns and/or motifs shown in any of the mouldings, moulds or *bucks*. The “E36 Visual Item” concept is also linked to a particular style using the “E55 Type” concept. As such, we conceptualise artistic style by characterising similar physical features which are shared by all objects from the same style despite of when these objects were produced. As such, Regency style objects could have been produced in 1810 or today.

In order to have a vocabulary to describe the different types of cross sections, patterns and motifs, we use the Getty Art & Architecture Thesaurus (ATT) [Institute 2015]. This thesaurus includes elements commonly used in Regency mouldings, such as egg and dart, bead and reel as well as fret mouldings.

In addition, CRMdig [Doerr and Theodoridou 2011] is used to link the “E84 Information Carrier” to its corresponding scanned mesh (“D9 Data Object”). This link enables the shape analysis method to query those 3D meshes with known styles in order to enrich those shapes with unknown styles. As such, we can automatically enrich shapes with concepts such as pattern/motif and style. The shape retrieval method is described in the following section followed by initial results.

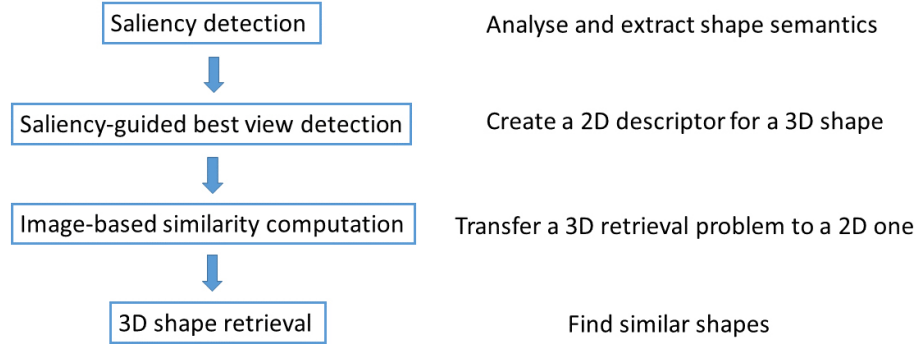


Fig. 4. The workflow of the proposed shape retrieval pipeline

5. SALIENCY-GUIDED ANALYSIS OF DECORATIVE STYLE

The analysis of decorative style of ornaments can be regarded as a retrieval process based on the 3D shape information. 3D shape retrieval is a fundamental topic in computer vision and graphics [Boyer et al. 2011; Bronstein et al. 2012]. In this paper, we present a novel shape retrieval method based on shape saliency. The workflow of the method is illustrated in Fig. 4 where we also explain the function of every step. Note that this pipeline differs from the method proposed in [Rodriguez Echavarria and Song 2015] from the second step. To make this journal version self-contained, we still describe and explain our method from the first step.

5.1 Shape saliency

3D shape saliency, first proposed in [Lee et al. 2005], is a measure of regional importance based on human perception. On one hand, the saliency map of a 3D shape highlights the most important geometric features on the surface. On the other hand, it is not just about shape geometry. Some experiments have been conducted to demonstrate that the computational model of shape saliency has better correlation with human eye fixations than a random model and a curvature-based model [Kim et al. 2010]. Thus in computer graphics, it is very promising for solving problems related to semantic analysis of 3D shapes. This has recently motivated many researchers [Shilane and Funkhouser 2007; Castellani et al. 2008; Song et al. 2012; Chen et al. 2012; Leifman et al. 2012; Hou and Qin 2013; Song et al. 2014] to investigate various computational saliency models.

We propose a novel algorithm for computing shape saliency based on *stochastic Laplacian*. Given a shape as a triangulated mesh M , we first calculate its Laplacian matrix. Typically, the Laplacian matrix of a mesh is based on the discretisation of a continuous Laplacian (e.g., the Laplace-Beltrami operator) defined (mathematically) for a smooth manifold using some weighted sum of adjacent vertices [Dey et al. 2010]. If a mesh M contains m vertices p_1, \dots, p_m , in its simplest form, the Laplacian matrix can be computed as:

$$L = D - A \quad (1)$$

where A is the adjacency matrix, given by

$$A(i, j) = \begin{cases} 1 & \text{if } p_i \text{ and } p_j \text{ are neighbours} \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

and D is a diagonal matrix in which D_{ii} is the degree of vertex p_i . This simplest computational model merely takes into account the combinatorics (connectivity). Ideally, a multiscale representation should

Algorithm 1: Saliency detection**Data:** A mesh M represented by a vertex matrix P and a face matrix T **Result:** A saliency map I **begin**Compute \hat{L} and then the stochastic Laplacian \mathcal{L} , input the number of scales as K and input a precision value λ ; $F^{(1)} = \mathcal{L}$;**for** $k \leftarrow 2$ **to** K **do** $F^{(k)} = F^{(k-1)} \hat{L}$; $\mathcal{F} = [F^{(k)}]_r$;**for** $i \leftarrow 1$ **to** m **do****for** $j \leftarrow 1$ **to** m **do****if** $\mathcal{F}(i, j) < \lambda$ **then** $\mathcal{F}(i, j) = 0$ **end****end****end** $F^{(k)} = \mathcal{F}$; $\mathcal{D}^{(k)} = |F^{(k)} - F^{(k-1)}|$;Calculate the difference map as $D^{(k)} = \sum_r \mathcal{D}^{(k)}$ where \sum_r denotes the row-based summation;**end** $\mathcal{I} = \sum_{k=2}^K D^{(k)}$; $I = \log(\mathcal{I})$ **end**

not only describe the combinatory structure of the mesh, but also encode its local geometric details. To incorporate local small-scale geometric information, the adjacency matrix is weighted by the distances between neighbouring vertices:

$$W(i, j) = \frac{1}{\|p_i - p_j\|^2} A(i, j); \quad (3)$$

This leads to the mesh Laplacian

$$L = D - W \quad (4)$$

We employ Eq. (3) other than the widely used cotangent because we find that it benefits the detection of salient features. The largest angle in an ill-shaped triangle has a negative cotangent but with a large absolute value. Since we need to compute the absolute values of the elements of the Laplacian to transfer it to a stochastic matrix, in the experiments we found that the cotangent Laplacian leads to high saliency values on ill-shaped triangles. The proposed mesh Laplacian is more meaningful in terms of salient features since usually most meshes retain more vertices around features (e.g., shape extremities). And we also observed that the cotangent Laplacian is not so efficient as the mesh Laplacian.

In this work, Eq. (4) is followed by two steps. First, we compute \hat{L} composed of the absolute values of the elements of the mesh Laplacian: $\hat{L}_{ij} = |L_{ij}|$. Second, we normalise \hat{L} subsequently so that the sum of each row is 1

$$\mathcal{L} = \left[\frac{\hat{L}}{\sum_r \hat{L}_{rj}} \right] \quad (5)$$




Fig. 5. Saliency detection on meshes. Top half: Original shapes; Bottom half: Saliency maps where warmer colour denotes salient regions.

where $[\cdot]_r$ denotes the row-based normalisation.

These two steps are very important since it actually turns the mesh Laplacian to the so-called *stochastic Laplacian*.

We first explicitly give our algorithm for computing shape saliency in Algorithm 1. It is based on the iterative updates of the stochastic Laplacian through matrix multiplication, and can be implemented much more easily and efficiently than previous methods which require eigendecomposition of the Laplacian or the design and dilation of wavelets. In the following, we analyse the algorithm and explain why it works.

The row based normalisation $\mathcal{F} = [F^{(k)}]_r$ in Algorithm 1 guarantees that the vertex matrix is updated by a stochastic matrix in each iteration, which in fact describes the transitions of a Markov chain. It is known that for a stochastic matrix \mathcal{F} , \mathcal{F}_{ij} denotes the probability of the one-step transition. Therefore, the k -th stochastic Laplacian produced through a series of $k - 1$ matrix multiplications gives the k -step transition probability. Since the matrix of stochastic Laplacian is sparse and here we also set a precision value λ to rule out small entries, most transitions are prohibitive (transition probability equals zero). And, only the transitions within a neighbourhood is available because it is constructed based on the adjacency matrix. Hence, the k -step transition actually defines a connected k -ring neighbourhood while all other connectivities/paths are prohibited. And when we perform the multiplication $F^{(k)} = F^{(k-1)}\hat{L}$ in Algorithm 1, we essentially perform a displacement in a k -ring neighbourhood for each vertex. Such a vertex displacement leads to the loss of local surface details and a hierarchy of meshes with multiple scales of details is formed. The more salient the vertices, the more the loss of corresponding local details. Therefore, we calculate the shape saliency as the aggregation of the difference maps at all scales excluding the first one in order to avoid being affected by noise perturbation. We finally output the saliency map I by computing the logarithm of \mathcal{I} merely for a better visualisation effect. Some results using shapes from the Jackson collection are shown in Fig. 5.

In Fig. 5, it can be seen that the detected salient regions are typically the most representative and informative features of the shape in a semantic sense. And more importantly, it shows some correspon-




Fig. 6. Image alignment via SIFT flow

dence. For instance, the egg and dart pattern is always detected as a salient region in different shapes. In general, shapes which can be categorised into the same semantic group could have highly correspondent salient features. This finding inspired us to develop an automatic shape retrieval system.

5.2 Shape retrieval

Typically, people understand the semantics of an architectural moulding by seeing its frontal view, which can be regarded as the best 2D view of the corresponding 3D model. In this work, we find the best 2D view of a 3D model as a descriptor which encodes the most important semantic information of the 3D model. This can be done by finding the viewpoint corresponding to the 2D view which maximises the observed saliency and the visibility of the 3D model. Given a viewpoint, the observed saliency with regards to it is defined as the sum of the saliency of all vertices that can be viewed from the viewpoint. The visibility of a 3D model with regards to a viewpoint is computed as the number of vertices that can be viewed from it. We use the algorithm proposed in [Rocha 2008] to check if a vertex is visible from a particular viewpoint.

In practice, we first create a set of 300 candidate viewpoints by uniformly sampling a sphere which bounds the 3D object. We use a sphere whose radius is twice as large as that of the tight bounding sphere. And then we calculate the observed saliency with regard to a viewpoint v as

$$X(v) = \sum_{p \in H(v)} I(p) \quad (6)$$

where $H(v)$ denotes the set of surface points visible from v and $I(p)$ is the saliency of a surface point. The visibility of v is computed as

$$Y(v) = |H(v)| \quad (7)$$

where $|H(v)|$ denotes the number of points in the point set $H(v)$.

Then both X and Y are normalised to be within the interval $[0, 1]$. The detection of the best 2D view follows two rules. First, the view should contain a large number of vertices. Second, the view should contain a large number of salient vertices. These two rules guarantee that the detected view is rich in information and thus potentially contains the most important semantics of the 3D model. So we find the viewpoint v^* which can maximise the sum of the observed saliency and the visibility

$$v^* = \arg \max_v (\mathcal{X} + \mathcal{Y}) \quad (8)$$




Fig. 7. Best 2D view detection for 3D models. Left: the original poses of input 3D models; Right: the detected best views of the 3D models

where \mathcal{X} and \mathcal{Y} are the normalised observed saliency and visibility respectively. Once we find the best viewpoint, we use the 2D view taken at the best viewpoint as a descriptor of the 3D model. Figure 7 shows a set of input 3D models and their corresponding best 2D views detected by our saliency-based method.

The best 2D view of a 3D shape is essentially a 2D descriptor of the 3D shape. Compared with the saliency histogram (which is essentially a 1D descriptor) used in [Rodriguez Echavarria and Song 2015], this newly proposed descriptor contains more information and thus is more discriminative. Consequently, it usually leads to more reliable retrieval results. This claim will be experimentally demonstrated in the next section.

After detecting the best views of the 3D models, we use these views, which are essentially 2D images, as descriptors to measure their similarity and directly use the measured similarity score to indicate the similarity between the corresponding 3D models. We employ the SIFT flow method [Liu et al. 2011] to measure the similarity between two views. SIFT flow detects dense correspondence between 2D images at scene level. The dense correspondence gives an alignment between a pair of images where corresponding regions from two different images are locally distorted to guarantee a good alignment between them as shown in Fig. 6. If the two images are identical, no distortion is needed to achieve a good alignment while if the two images are very different, a number of large local distortion are needed. Thus, the less the distortion, the more similar the two images. By measuring the amount of distortion which can be computed as the sum of the positional changes of all points before and after the alignment, we know how similar two images are quantitatively, which will be directly used to indicate the similarity between the two 3D models that these two images represent. In addition, since SIFT flow can handle images in different sizes, it is particularly suitable for measuring the similarity in this work. If one object only matches a small sub-region of the other one, their similarity score generated by SIFT flow is usually low.




Fig. 8. Saliency of shapes containing different numbers of vertices

6. RETRIEVAL RESULTS

We have tested the shape retrieval method described on the previous section on the scanned artefacts of the Jackson collection. On average, each shape is composed of 1.1 million vertices. It is very inefficient to analyse so many large shapes. Figure 8 demonstrates that the proposed saliency method is insensitive to simplification. In practice, we always employ QSlim [Garland and Heckbert 1997] to simplify the input shapes in order to speed up the retrieval.

Figure 9 shows the results of the proposed saliency-guided shape retrieval. In order to enrich semantically the 3D shapes, at least a few examples of each decorative pattern in mouldings in the database contain some semantic information regarding their production method and style (as described in section 4). These shapes will enable the classification of those 3D shapes that do not have semantic information linked to them. For the three queries in the first row in Fig. 9, our method successfully retrieves other shapes which are linked to an entry of ‘Pattern of egg and dart’ (E36 Visual Item) in their semantic network (see Fig. 3). In other words, the proposed saliency-guided method is capable of understanding the semantics corresponding to ‘Pattern of egg and dart’ which are typically independent of the overall size of the shape, the local geometry of the shape (e.g., the egg patterns in the query shapes are not exactly the same as those in the retrieved shapes) and the global structure of the shape (e.g., how many egg patterns it contains). This is desired as it is consistent with the classification of the semantic network.

We also note that the shapes which result from the shape retrieval method include mouldings, moulds and *bucks*. Shape-wise, there is no distinction between a moulding and a *buck*. This distinction is usually made by the material of the object rather than by its shape. In addition, moulds usually contain flat surface on the boundaries. Hence, the amount of boundary flat surfaces in some shapes in Fig. 9 is not relevant when determining the saliency of the shape. Nevertheless, these flat regions can potentially reveal some information regarding the type of object this shape represents (e.g. mould or moulding). This could be useful when there is limited information on the shape and it needs to be determined whether this is a moulding or a mould.

We have compared our method to the well-known D2-based shape retrieval [Osada et al. 2002]. We use all mesh vertices to compute the D2 descriptor. Increasing the number of sampled vertices does increase the performance of this descriptor since it becomes more discriminative, as expected. However, when the number is sufficiently large (e.g., 50,000), the improvement is very limited. The results of the comparison, shown in Fig. 9, demonstrate that the retrieval results are more semantically meaningful using our method. The results also demonstrate that saliency information can deliver an effective retrieval of 3D shapes which partially contain similar patterns and motifs.

Also, compared with the method proposed in [Rodriguez Echavarria and Song 2015] where we use the histogram of saliency as the shape descriptor, it can be seen that although the retrieval based on the histogram of saliency outperforms the D2 method, its performance is not as good as the method proposed in this paper. One interesting observation in Fig. 9 is that, except the first query shape which is the only one containing one and a half eggs in our database, for the other two query shapes, the shapes retrieved by our method contains roughly the same number of eggs and darts as the query



Fig. 9. Shape retrieval results. First row: query shapes (in yellow shading) selected from the Jackson collection shape database; Rows 2-5: the corresponding top four retrieved shapes (in cyan shading) using the proposed saliency-guided method; Rows 6-9: the corresponding top four retrieved shapes (in red shading) using the D2 method

shape. The reason is that in the measure of the similarity of two shapes where we use SIFT flow to compute the similarity of their corresponding best views, the number of the ‘egg and dart’ patterns does matter since it seeks dense point-wise correspondence between two images. By contrast, the shapes retrieved by using the histogram of saliency is not robust to the number of the ‘egg and dart’ patterns. This is because the histogram essentially reflects the proportion of the numbers of points with different

levels of saliency and thus not so discriminative for shapes containing different numbers of eggs and darts where the proportion of the numbers of points with different levels of saliency are roughly the same due to similar local geometry.

7. CONCLUSION

This paper presented research contributing towards the semantic enrichment of 3D repositories, in particular focusing on architectural mouldings. An ontology is used to describe these artefacts and it is complemented by a saliency-guided shape retrieval method which identifies similar mouldings. This mechanism allows the enrichment of the 3D models with information regarding their artistic style and production mechanism.

One of the advantages of our approach is that it can be applied to other heritage artefacts. In particular, it can be applied to those artefacts with some type of decorative ornament. The shape analysis algorithm could also be applied to a variety of heritage shapes to learn more about their semantic meaning. However, the main disadvantage of the current implementation is that it is unable to segment each individual pattern when more than one is used as shown in figure 1. This will require further work on segmentation of each pattern on a shape. As a limitation of the current work, we only scanned and reconstructed a set of 3D models of Regency style ornaments where most of them contain the ‘egg and dart’ patterns organised in a translational manner. The evaluation based on the existing database is thus not sufficiently solid due to the lack of a large-scale database containing a variety of decorative ornaments. Therefore, future work will involve including further design styles as well as a greater variety of patterns and motifs.

Acknowledgment

The authors would like to thank the Regency Town House for providing access to the Jackson ornament collection and relevant expertise for conducting the research.

REFERENCES

2015. The Regency Town House. (2015). <http://rth.org.uk/>
- 3D-COFORM Project. 2015. 3D Data Catalogue. <http://www.3dcoform.eu/x3domCatalogue/>. (2015).
- Daniel G. Aliaga, Elisa Bertino, and Stefano Valtolina. 2011. DECHO—A Framework for the Digital Exploration of Cultural Heritage Objects. *J. Comput. Cult. Herit.* 3, 3, Article 12 (Feb. 2011), 26 pages. DOI:<http://dx.doi.org/10.1145/1921614.1921619>
- Marco Attene, Francesco Robbiano, Michela Spagnuolo, and Bianca Falcidieno. 2007. Semantic annotation of 3D surface meshes based on feature characterization. In *Proceedings of the semantic and digital media technologies*. Springer-Verlag, Genoa, Italy, 126–139.
- Edmond Boyer, Alexander M Bronstein, Michael M Bronstein, Benjamin Bustos, Tal Darom, Radu Horaud, Ingrid Hotz, Yosi Keller, Johannes Keustermans, Artiom Kovnatsky, and others. 2011. SHREC 2011: robust feature detection and description benchmark. In *Proc. EG Workshop on 3D Object Retrieval* (2011).
- Breuckmann. 2015. SmartSCAN3D Series. (2015). <http://www.breuckmann.com>
- A Bronstein, M Bronstein, and Maks Ovsjanikov. 2012. 3D features, surface descriptors, and object descriptors. *3D Imaging, Analysis, and Applications* (2012).
- U. Castellani, M. Cristani, S. Fantoni, and V. Murino. 2008. Sparse points matching by combining 3D mesh saliency with statistical descriptors. In *Proc. Eurographics*. 643–652.
- Xiaobai Chen, Abulhair Saparov, Bill Pang, and Thomas Funkhouser. 2012. Schelling Points on 3D Surface Meshes. In *Proc. SIGGRAPH*.
- Tamal K Dey, Pawas Ranjan, and Yusu Wang. 2010. Convergence, stability, and discrete approximation of Laplace spectra. In *Proc. ACM-SIAM Symposium on Discrete Algorithms*. 650–663.
- Martin Doerr and Maria Theodoridou. 2011. CRMdig: A generic digital provenance model for scientific observation. In *Proceedings of 3rd USENIX Workshop on the Theory and Practice of Provenance*. Heraklion, Crete, Greece. http://www.usenix.org/event/tapp11/tech/final_files/Doerr.pdf

- Martin Doerr, Katerina Tzompanaki, Maria Theodoridou, Ch. Georgis, A. Axaridou, and Sven Havemann. 2010. A Repository for 3D Model Production and Interpretation in Culture and Beyond. In *Proceedings of the International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage*. 97–104. DOI: <http://dx.doi.org/10.2312/VAST/VAST10/097-104>
- Carl Doersch, Saurabh Singh, Abhinav Gupta, Josef Sivic, and Alexei A. Efros. 2012. What Makes Paris Look like Paris? *ACM Transactions on Graphics (SIGGRAPH)* 31, 4 (2012), 101:1–101:9.
- M. Garland and P.S. Heckbert. 1997. Surface simplification using quadric error metrics. In *Proc. SIGGRAPH*. ACM Press/Addison-Wesley Publishing Co., 209–216.
- Tingbo Hou and Hong Qin. 2013. Admissible Diffusion Wavelets and Their Applications in Space-Frequency Processing. *IEEE Transactions on Visualization and Computer Graphics* 19, 1 (2013), 3–15.
- Suzanne Hudson and Noonan-Morrissey Nancy. 2001. *The art of writing about art*. Cengage Learning.
- Jane Hunter and Anna Gerber. 2010. Harvesting community annotations on 3D models of museum artefacts to enhance knowledge, discovery and re-use. *Journal of Cultural Heritage* 11, 1 (2010), 81–90. DOI: <http://dx.doi.org/doi:10.1016/j.culher.2009.04.004>
- The Getty Research Institute. 2015. Art & Architecture Thesaurus®Online. (2015). <http://www.getty.edu/research/tools/vocabularies/aat/>
- International Council of Museums. 2011. CIDOC CRM Home page. (2011). <http://cidoc.ics.forth.gr/>.
- Owen Jones. 1986. *The Grammar of Ornament*. Studio Editions, London.
- Evangelos Kalogerakis, Aaron Hertzmann, and Karan Singh. 2010. Learning 3D Mesh Segmentation and Labeling. *ACM Transactions on Graphics* 29, 3 (2010).
- Y. Kim, A. Varshney, D.W. Jacobs, and F. Guimbretiere. 2010. Mesh saliency and human eye fixations. *ACM Transactions on Applied Perception* 7, 2 (2010), 12:1–12:13.
- G. Lavoué, J-P. Vandeborre, H. Benhabiles, M. Daoudi, K. Huebner, M. Mortara, and M. Spagnuolo. 2012. SHREC'12 Track: 3D Mesh Segmentation, M. Spagnuolo, M. Bronstein, A. Bronstein, and A. Ferreira (Eds.). Eurographics Association, Cagliari, Italy, 93–99. DOI: <http://dx.doi.org/10.2312/3DOR/3DOR12/093-099>
- C.H. Lee, A. Varshney, and D.W. Jacobs. 2005. Mesh saliency. In *Proc. SIGGRAPH*.
- George Leifman, Elizabeth Shtrom, and Ayellet Tal. 2012. Surface Regions of Interest for Viewpoint Selection. In *Proc. CVPR (oral)*.
- Ce Liu, Jenny Yuen, and Antonio Torralba. 2011. Sift flow: Dense correspondence across scenes and its applications. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* 33, 5 (2011), 978–994.
- D. Lo Buglio, V. Lardinois, and L. De Luca. 2013. Revealing shape semantics from morphological similarities of a collection of architectural elements: The case study of the columns of Saint-Michel de Cuxa. In *Digital Heritage International Congress (DigitalHeritage), 2013*, Vol. 1. 465–472. DOI: <http://dx.doi.org/10.1109/DigitalHeritage.2013.6743785>
- Livio De Luca. 2013. 3D Modelling and Semantic Enrichment in Cultural Heritage. In *54th Photogrammetric Week in Stuttgart*.
- Livio De Luca, Chawee Busayarat, Chiara Stefani, Philippe Véron, and Michel Florenzano. 2011. A semantic-based platform for the digital analysis of architectural heritage. *Computers & Graphics* 35, 2 (2011), 227–241. DOI: <http://dx.doi.org/doi:10.1016/j.cag.2010.11.009>
- Zhaoliang Lun, Evangelos Kalogerakis, and Alla Sheffer. 2015. Elements of Style: Learning Perceptual Shape Style Similarity. *ACM Trans. Graph.* 34, 4, Article 84 (July 2015), 14 pages. DOI: <http://dx.doi.org/10.1145/2766929>
- Franz S. Meyer. 1974. *A Handbook of Ornament*. Duckworth, Trowbridge & Esher.
- R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin. 2002. Shape distributions. *ACM Transactions on Graphics* 21, 4 (2002), 807–832.
- Laura Papaleo and Leila De Florian. 2010. Manual Segmentation and Semantic-based Hierarchical Tagging of 3D models. In *Eurographics Italian Chapter Conference*. 25–32.
- Steven. Parissien. 1992. *Regency style*. Phaidon, London.
- Peter Phillips and Gillian. Bunce. 1993. *Repeat patterns: a manual for designers, artists and architects*. Thames & Hudson, London. <http://capitadiscovery.co.uk/brighton-ac/items/44704>
- Fabio Pittarello and Ivano Gatto. 2011. ToBoA-3D: An Architecture for Managing Top-down and Bottom-up Annotated 3D Objects and Spaces on the Web. In *Proceedings of the 16th International Conference on 3D Web Technology (Web3D '11)*. ACM, New York, NY, USA, 57–65. DOI: <http://dx.doi.org/10.1145/2010425.2010436>
- Princeton University. 2015. Princeton Shape Benchmark. <http://shape.cs.princeton.edu/search.html>. (2015).
- Kelvin R Rocha. 2008. *A variational approach for viewpoint-based visibility maximization*. ProQuest.

- K. Rodriguez Echavarria, David Morris, and David Arnold. 2009. Web based presentation of semantically tagged 3D content for public sculptures and monuments in the UK. In *Proceedings International Conference on 3D Web Technology*. ACM, Darmstadt, Germany, 119–126.
- Karina Rodriguez Echavarria and Ran Song. 2015. Studying Shape Semantics of an Architectural Moulding Collection: Classifying Style Based on Shape Analysis Methods. In *Digital Heritage*. IEEE.
- K. Rodriguez Echavarria, Maria Theodoridou, Christos Georgis, David Arnold, Martin Doerr, Andre Stork, and Sebastian Pea Serna. 2012. Semantically Rich 3D Documentation for the Preservation of Tangible Heritage. In *Proceedings of the International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage*. Brighton, UK, 41–48. DOI:<http://dx.doi.org/10.2312/VAST/VAST12/041-048>
- Ariel Shamir. 2008. A survey on Mesh Segmentation Techniques. *Comput. Graph. Forum* 27, 6 (2008), 1539–1556.
- L. Shapira, S. Shalom, A. Shamir, D. Cohen-Or, and H. Zhang. 2010. Contextual Part Analogies in 3D Objects. *Int. J. Comput. Vision* 89, 2-3 (2010), 309–326.
- P. Shilane and T. Funkhouser. 2007. Distinctive regions of 3D surfaces. *ACM Transactions on Graphics* 26, 2 (2007), 7.
- R. Song, Y. Liu, R.R. Martin, and P.L. Rosin. 2012. Saliency-guided integration of multiple scans. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1474–1481.
- Ran Song, Yonghuai Liu, Ralph R Martin, and Paul L Rosin. 2014. Mesh saliency via spectral processing. *ACM Transactions on Graphics (TOG)* 33, 1 (2014), 6.
- The Europeana Foundation. 2015. Europeana. <http://www.europeana.eu>. (2015).
- The Stanford Digital Forma Urbis Romae Project. 2015. Stanford Digital Forma Urbis Romae Project. <http://formaurbis.stanford.edu/>. (2015).
- Jonathan Thornton and William Adair. 1994. Applied Decoration for Historic Interiors Preserving Composition Ornament. (1994). <http://www.nps.gov/tps/how-to-preserve/briefs/34-composition-ornament.htm>
- VISIONAIR Consortium. 2015. The Shape Repository of the Visualisation Virtual Services (VVS). <http://visionair.ge.imati.cnr.it/ontologies/shapes/>. (2015).