CGI2012 manuscript No.
(will be inserted by the editor)

Transferring Characteristic Proportions to Modify the Artistic Style of Cartoons

Philip Buchanan · Michael Doggett

Abstract This paper presents a new automatic algorithm for extracting properties of an individual artist's style and applying those properties to other drawings. The algorithm can be used to ensure consistent artistic style across a project by automatically transferring a selected style to hand drawn images. We present an algorithm that uses multiple stages; subdivision of an image into selected geons, extraction of corresponding datasets from additional images, analysis and characterisation of the artist's style and finally transferring the artist's style to target images. When transferring the artist's style we use a complexity sensitive scaling algorithm to avoid artefacts. We apply this algorithm to line drawn cartoons and demonstrate that artistic style can be transferred between cartoon objects. Our results show that analysis of object proportions can be used practically to ensure consistency for objects with the same Geon structure but drawn by different artists.

Keywords Style Transfer · Artist Style · Cartoon Characters

1 Introduction

Content creation in modern entertainment is one of the most time consuming components of a project. Our research aims to reduce the workload of content creators by providing tools that allow them to easily work with artistic style. Our research deals specifically with cartoon style content, in media such as TV animation, movies, and video games.

Philip Buchanan
University of Canterbury

Michael Doggett
Lund University

TV series and 'AAA' video games typically have large art teams and enormous amounts of content. Ensuring stylistic consistency is difficult, and is something our algorithm aims to assist. In contrast, video games for the social space - such as facebook and the iPhone - typically have fast development cycles with small teams. Consistency within the team is rarely an issue, however content is often sourced from external locations. Providing tools to deal with artistic style will allow easier, faster and more coherent integration of generic art assets or user generated content.

In this paper, we propose a novel method for automatically extracting an individual artists style and applying it to other drawings. Our process centres on the characteristic proportions of an image. Unlike previous work which typically transfers artistic elements such as line style, we capture characteristics with spatial context and transfer those. When two different artists draw stylised versions of the same object, the exaggerations and shapes within the image are often significantly different. By measuring these properties across a large body of the artist's work, we can discover which proportions characterise the artist, and which are inherently part of the object.

Size, position, and shape differences can be measured in isolation and in relationship to each other. Characteristic proportions define the relationships within a set of similar shaped objects by the same artist. We can change the style of an object by altering its characteristic proportions to match the characteristic proportions of an image in our target style.

For example; a character with a large head and small body will have different characteristics to a character with a large head and large body, even if their overall height and width are the same. In this respect, objects of the same type drawn by the same artist have

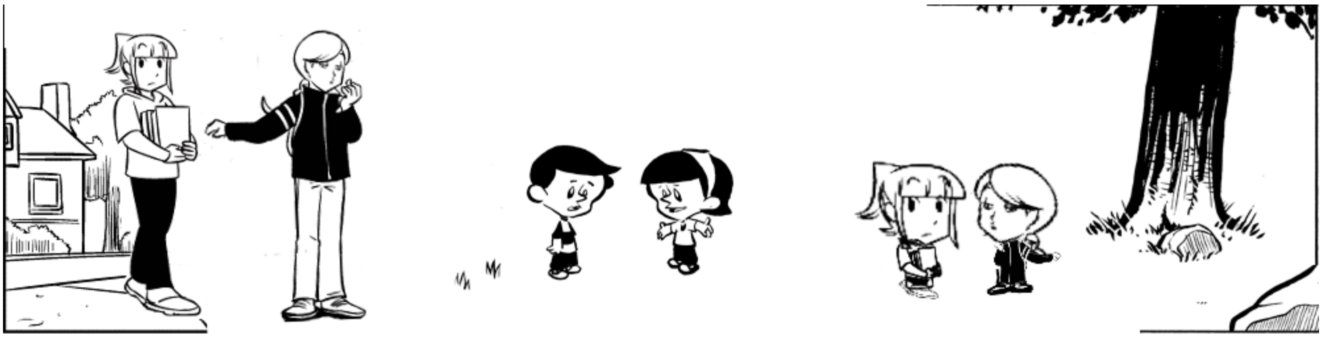

Fig. 1 Characters from the comic strip 'Sandra and Woo' have the artist style from 'Misery Loves Sherman' transferred to them. Left and Middle: Original characters. Our method automatically differentiates measurements that are specific to an artists' style from the measurements that belong to the object. Right: Left characters using the Middle character's style.

very similar characteristic proportions, whereas objects of the same type drawn by different artists often have different characteristic proportions.

Our algorithm automatically analyses the images, and creates a database with the characteristic proportions of a single artist. Given an input image and the appropriate geon subdivision, the algorithm finds the most likely target proportions and automatically resizes the source image to fit them. We also introduce a complexity sensitive scaling algorithm that prevents character details being distorted when scaling in one direction. Our algorithm does not deal with additional properties, and so constraining ourselves to line drawn cartoon images reduces the influence of rendering technique in the portrayal of style, and allows us to focus upon the proportions alone.

2 Previous Work

Non-photorealistic rendering research often deals with graphic style. Topics include rendering techniques and image analysis, and a common focus is replicating brush strokes or the typical shading of an artist. However, without altering the core image, this is seldom enough to replicate a specific artists style entirely. Given the focus upon specific rendering techniques, there is a surprising deficit of research linking styles to specific artists, or taking a broader overview of the components of visual style.

Modifying the artistic style of an image is difficult, and requires a plethora of techniques. Resources such as Freeman et al. [5] and Hertzmann et al. [7] showed that different line styles can be recorded and used to generate new images. They outlined the algorithms and techniques, and proved that for these limited cases, it is possible to programmatically redefine an aspect of visual style.

A brief study by Wayner [16] attempts to classify visual style using two descriptors; line length and line camber. It uses these descriptors to create a histogram of line lengths in an image, and then match this histogram with a database of authors. It was found that using only a line length classification, 95% accuracy could be achieved in the categorisation of black and white cartoon strips between 7 authors. A more thorough statistical analysis of artistic style was successfully performed by Shamir et al. [12]. Their research focuses upon painting technique and the classical school of art, and many of their metrics do not apply to digital media. However, these two studies are important because they show, definitively, that it is possible to mathematically classify and recognise artistic and visual style based upon simple metrics.

Two papers describe comprehensive techniques for the style translation of lines within 2D drawings. These outline relevant techniques and methods that may be important when implementing border style manipulation algorithms. Freeman et al. [5] outline the integration of two techniques used to transfer line style between a training data-set and a basic line drawing. 'Line style' in this paper refers to the width, curvature and corner data. A nearest neighbour algorithm was used to match lines in a database, which were then adjusted by the least-squares combination with a number of similar lines using a novel uneven point distribution. This was found to produce a translation algorithm that preserved the image and the visual fidelity of the style based upon percieved interest points.

Hertzmann et al. [7] explore the creation of 'curve analogies'. By using a texture-matching scale and rotation invariant transform, their algorithm can analyse a single line, stroke, or image, and transfer the style properties from this drawing onto a second drawing. One major advancement shown in this paper is the ability to perform this transplant between two very different images. While the paper talks about style in the context of visual style, it admits that this definition extends only as far as the shape variations within a single curve.

Both Freeman [5] and Hertzmann [7] use database-driven re-mapping techniques. The style mapping examples by Freeman et al. [5] are of high quality and show that this technique works well. However, the limitation of line style re-mapping shows in areas such as the human face - where modifications are stylistically consistent, but create unrealistic facial features. Our paper proposes a context aware algorithm using shape abstraction to avoid unrealistic modification of specific features.

3 Algorithm Overview

Our core algorithm analyses a database of images and detects consistent features across the dataset. It uses these consistent features to measure the properties of an image and perform a style transfer using measurements characteristic to the artist's style. Manually building this database would take considerable effort, and much of the work in this paper focuses upon automatically building datasets based upon unknown source images.

From an initial image, we create a structural description of the significant components. A combination of the bounding hull, image shape and complexity are used to subdivide the object into the basic geometric units (Geons) that will be used to measure characteristic proportions. This initial 'prototype' information is used to extract similar objects from a large number of source images. After the new objects are extracted and corresponding Geons fitted, the characteristic proportions are entered into the database.

These proportions are measured across a range of images from the same artist, and the distribution of measurements is used for categorisation. Each property is categorised depending upon whether it is part of the artists' characteristic style or unique to the specific input image and therefore irrelevant when performing the style modification. The style modification itself takes an image by the source artist and modifies the proportions to those characteristic to the target artist. Attention is needed to ensure the image content is preserved correctly when the properties are modified.

This process transforms an image into the same style as the target artist. The full algorithm is described in the following 4 sections, and the results displayed in section 8.

4 Prototype Subdivision

To extract our target proportions automatically from a large set of images, a basic prototype is supplied to demonstrate the object. Our algorithm divides these into unique shapes called geons that represent graphically significant components of the object. Geons are later used to extract the object from the background in target images, and then measure their rotation and scale.

Both Theobalt et al. [13] and Li et al. [9] present methods for decomposition of complex shapes into geometric units (Geons). Although both have different target applications, they each develop structural representations of context that would be ideal for our situation. However their internal structure is built from an analysis of the object surface in 3D, something that is difficult to extract or approximate from 2D images. With our focus upon a limited domain, a specialised subdivision algorithm is easier to implement and likely to function better in the 2D case.

Ullman et al. [14] find that Geons encompasing visual features of intermediate complexity are of best use in image classification. In the context of geon transfer, this should also provide the best result when remapping.

The object prototype is important because it heavily influences the extraction and measurement of source images that are used to compile the artist database.

We use a naive algorithm to do the separation by analysing the exaggeration, texture and internal lines present in most cartoon images. Based on a sample of common images, we assume that large variations in width or internal complexity indicate the boundaries between different parts of the object. A common case is a cartoon character, where we focus upon the proportions of head, body and legs. These form our three geons of intermediate complexity.

Figure 2 shows the subdivision process. First, a centreline is drawn through the longest axis of the image. If the image has significant bends, the centreline is also bent to account for these. If the image forks, then the most significant branch is chosen based upon its width and length. The complexity of an image at a point on the centreline is measured as the sum of derivitives at intervals along an orthogonal cross-section. A bounding hull is also drawn around the image. Because the area is not guaranteed to be closed, we impose a limit on the contraction and trade accuracy for a correct visual border. This bound is used to calculate the width in respect to the centreline

An evenly weighted combination of smallest widths and lowest complexity is used to find segmentations for

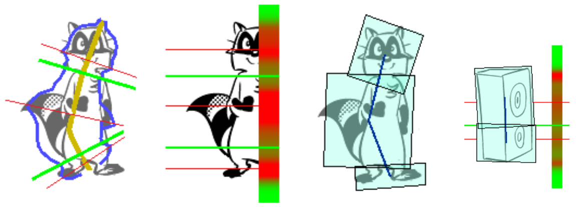

(a) (b) (c) (d)

Fig. 2 A cartoon prototype automatically subdivided into appropriate Geons. The yellow centreline and blue bounding hull line (a) are used to find appropriate subdivisions (green) based on width. Note the bounding hull is centered and therefore does not visually match the image borders. The image complexity (b) is shown untransformed, and green lines mark aproximate locations of low complexity. The character Fs (c) and object geons (d) are created based on these subdivision lines.

the body. The final number of sections is determined by the size of the image and the strength and placement of the potential cuts. Geons are created aligned to the centreline, using the height of the section and the maximum width of the section.

The width of the image is measured at boundaries of adjacent geons where the connectivity graph crosses. These measurements are later used to prevent discontinuities in the image when scaling along these boundaries. An aligned grid is used to additionally measure widths inside the geon, and later used to influence the contents of a Geon when scaling.

This method works without intervention for a range of common and simple cartoon objects, however it does not correctly subdivide complex or branching images. In this case the object can be subdivided and separate prototypes generated from the pieces. This is considered a viable solution, because the input we used primarily consisted of non-complex objects. If the topology is too complicated, it reduces the chance of finding a corresponding object in the database of the other artist. If an object cannot be subdivided, then geons can be placed manually as a last resort.

5 Image Extraction

To build a database with a large number of entries, we need to measure instances of the target object or character within a large input dataset. Occurrences of a character are often found in context, and so need to be extracted from the background before they can be used. This is especially challenging, because the character in the target image is usually in a different pose to the character provided as a prototype. These differences in pose cannot be ignored because they provide the data required to perform the style transfer. We therefore need to extract each geon individually.

Most widely used image recognition algorithms use a variation of SIFT [10] followed by a feature reduction and matching step often based on RANSAC [4]. However, cartoon content does not produce good feature points, and data within the geons is often not geometrically or linearly transformed. Two solutions are often used to overcome this problem. A system based on exemplar-based recognition is less susceptable to non-linear transforms and has been demonstrated to work on large datasets[6], however the heavy training requirement of such a system defeats the purpose of implementing one in this context. Graph matching [2]

One of the best solutions for image matching line drawings is to use a shape context descriptor [1]. This has high accuracy when matching line drawings and monochrome images, however, requires that target images have little background interference. Our requirement is for a robust image extraction algorithm that can find multiple divergent occurances in complex target images using only one source image. With image data restricted to line drawn cartoons, and the Geon structure providing aproximate locations, a custom search algorithm provided the best result for these specific requirements.

The algorithm shown in figure 3 is a brute-force approach that uses line directions and cartoon fill to match two images at a point. We start with the most complex geon to reduce the number of false matches. In the rare case where no matches are found, the process is repeated with a less complex Geon. This metric is determined using the same algorithm as outlined in Section 4.

Unlike photos, the images being matched have strong lines and large areas of flat colour. A feature map is generated for each image that equally emphasises these features. First the derivative of the image is taken to allow easier access to the line data, and then thresholded where $\Delta V < 0.5$ to remove insignificant slopes or noise, replacing this with the flat colour from the image. If the image instead has a pattern or noise, this is removed by averaging the colour. Blank backgrounds in prototype images are not matched, to ensure they do not interfere with background illustrations in the target image and cause false rejections.

The edge directions represented by the gradients in the derivative are then converted into a single normalised direction vector. This ensures better matching in cases where a line is offset and the opposing edge directions would have interefered with a correct match. A naive search is then performed by measuring the loose correlation of the source and target for every pixel in

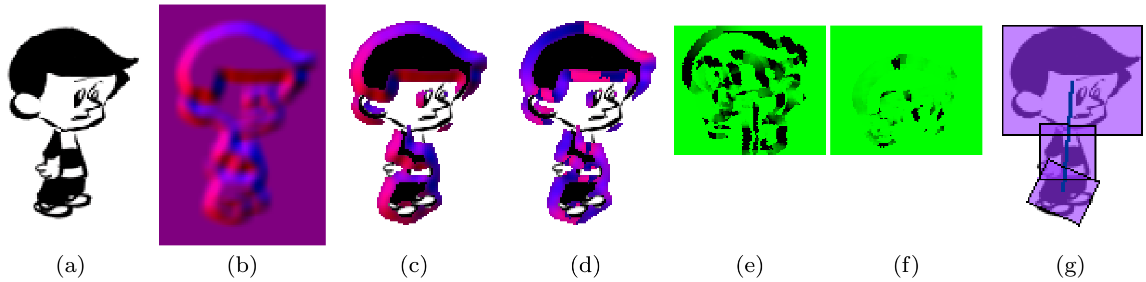

Fig. 3 Geon Extraction is peformed using an algorithm designed for this specific case. The derivatives of the source and target (a) images are combined with the flat colours (b), and vectors represent the line directions (c). A loose correlation is performed (d) at every pixel, where gradients and colours are matched within a radius. Correlation is performed across a range of rotations and scales, and the lowest difference (e) is considered to be a match. Subsequent Geons are matched within smaller areas (f) determined by the prototype. The final output (g) contains the size, rotation and position of each Geon in the target image.

the target image, with the source image rotated at increments of 1° and scaled by 1 pixel up to a range of 10%.

The loose correlation performed is similar to standard image correlation, except that overlapping pixels are also checked against nearby pixels to allow for differences in the drawings. This solves the situation where lines are mis-alinged by only a few pixels, yet do not match and hence distort the correlation result. Pixels are first compared based upon the vector direction if applicable, or based upon the flat colour if there is no vector. The best test radius for the sample dataset was determined to be 2% of the image size, but this needed to be increased for excessivly noisy or different images.

After the transformation of the first Geon is found, subsequent correlations do not need to be run across the entire image. The Geon structure in the prototype is used to determine the search areas for subsequent Geons, with the correlation being run in a radius of 200% around the location indicated.

This method worked sufficiently well for the sample dataset, requiring no training and needing calibration only in outlying cases. However, the naive search across all pixels in the image is an $O(n^3)$ algorithm and we leave optimization of this algorithm to future work. Additionally, this matching algorithm does not work on photographic images or images with high texture content that is liable to be miscategorised as lines. Given the specific nature of our dataset, these were acceptable tradeoffs.

6 Database Analysis

This algorithm provides the data needed to perform a proportion transfer. Within the database we have a set of objects, with their subcomponents (geons) in different poses. Some geon properties are important when representing an artists's style, while others are either irrelevant or apply only to a specific drawing of an object. By analysing the database for consistent configurations across the dataset, and discarding the properties that are drawing-specific, the Characteristic Proportions for a specific artist can be measured.

Prior reasearch by Freeman et al. [5] and Hertzmann et al. [7] performs style translation using line style. This type of context-free style translation is not uncommon, however it is unable to perform style translation using contextual clues from the image. The subdivision of images into geons and the knowledge about what they signify allows us to take measurements and analyse not just their individual properties, but also the relationship between them.

Figure 4 shows a database of extracted geons for a particular artist. The geon scale and position is measured relative to the prototype image, removing the scale bias that would otherwise be caused by different sized target images.

Further bias is removed by running the analysis process on a variety of objects of the same type, from the same artist. However running the process on characters often results in clear categorisations between different types of character within an artists' portfolio. Groupings such as 'Adult' and 'Child' or 'Good' and 'Bad' emerge. The differences in these groupings are more pronounced than the differences between artists. However the differences between, for example, the 'Child' groups from different artists is still significant enough to perform an effective style transfer.

Several measurements are recorded for each geon and compared to the other geons in the same class, in all other instances of the character. The size, angle, and distance to adjacent geon(s) are compared. In addition, the differences in angle and size between adjacent geons within a character are measured. For each property, we

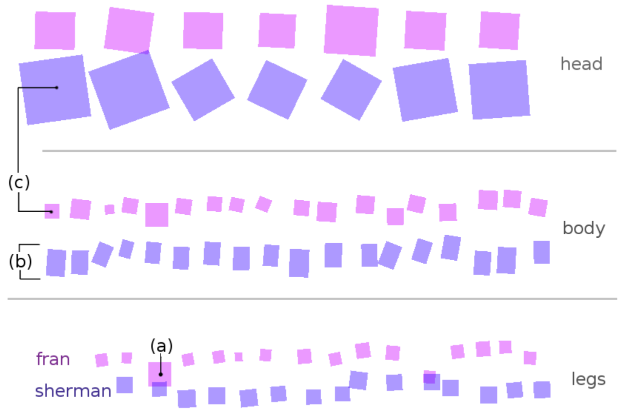

Fig. 4 This is a visualisation of the database showing geon measurements from two characters in 'Misery Loves Sherman' [3]. The rectangles visualise the size and rotation for each geon identified in the target images. Each colour denotes a character, blue representing Sherman and magenta representing Fran; and each row represents a specific class of geon — the head, body, and legs respectively. Three categories of geon are clearly visible. The size of (a) is unique to that geon, and needs to be recognised as an outlier; The height of the body geons (b) is consistently different for each character, and so is not a characteristic property for the artist; Lastly the ratio between the size of the geon classes marked by (c) is consistent across both characters, making this ratio characteristic to that artist. All three types of geon need to be recognised and categorised during the analysis stage.

measure the average (μ), the standard deviation (σ), and the total range in each direction (r). Exceptional scenes in the source image, or errors in the geon matching process can adversely affect the range, so outliers are discarded for this property.

To successfully transfer the characteristics between datasets, it is necessary to decide which properties are part of the characteristic style and which are specific to an individual character or image. If a property has a small variance and is similar across a wide range of images from the same artist, then it is likely to be part of their specific style. If a property has a large variance then it is probably dependent on the image or situation and therefore not of use as a style metric. This is difficult to measure, because most properties have different scales and there is no easy method to determine what a 'normal' variance is for each of them. Therefore, we compare a property from each class of geons.

The properties to be compared are outlined in Figure 5. It can be seen that the rotation of the head geon has a very wide range, as does the width of the leg geon. This is because the character was looking in different directions in the source images, and because the legs were drawn differently across a range of comics.

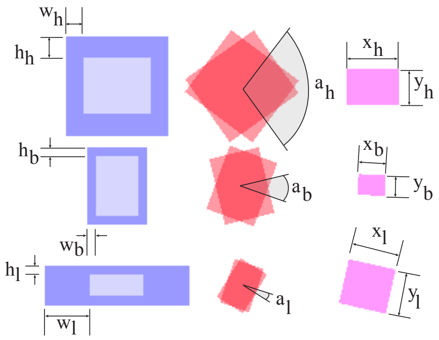

Fig. 5 This figure shows the possible range of values for a single set of character geons. Twelve instances of the character Lex from 'Not Just Nicky' [11] were analysed. The left column shows the sizes of the geons. The outer box shows the maximum size, and the inner box shows the minimum size. The middle column shows the range of rotations, and the right column shows the allowable range for each geons centre point.

The categorisation of these properties depends upon the target artist, but with a large range it is unlikely they are characteristic properties. In contrast, the rotation of the legs and the location of the body have very little deviation in their measurements and are likely to be characteristic properties. The properties are compared between artists to make that classification.

If a property shows a significant difference ($\mu > 0.5(\sigma_1 + \sigma_2)$) between two artists, then the property with the low deviation is counted as part of the style and the property with the high deviation is considered to be image specific. If there is no significant difference between artists, then we compare properties within one artist. Each property is compared to the equivalent in the other classes of geon. For example, the property measuring rotation of the head is compared to those measuring rotation of the body and the legs. The same test and categorisation is applied.

If there is still no significant difference found, then one of two situations has occurred. Either the property is a byproduct of the shapes and situation within the images and not characteristic to either artist—this is common for rotation, where the orientation of geons depends upon what's happening in the image; Or alternatively, both artists have this property as a strong component of their style—this is common for the ratio of sizes between adjoining geons. There is no ideal method for determining which of these options is correct, but the following formula gives reasonable results.

$$2\sigma \begin{cases} < r & \text{characteristic property} \\ > r & \text{image specific property} \end{cases} \tag{1}$$

Equation 1 relies upon the fact that our set of values is not evenly distributed. By measuring the ratio of the

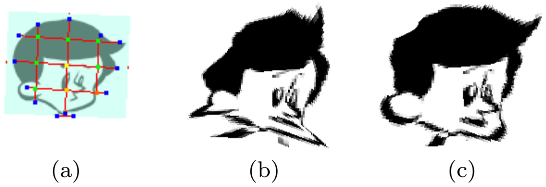

Fig. 6 In addition to scaling Geons, their contents are influenced using a weighted grid. The grid is extracted from the prototype image, including weightings based on complexity (a). If an image is simply scaled using the artist-significant gridlines, discontinuities are obvious (b). Weighting these values across the Geon produces a better result (c).

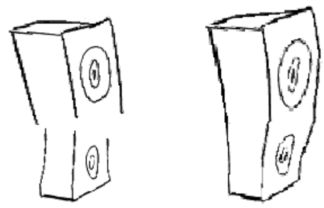

Fig. 7 When scaling, discontinuities are introduced at the borders of neighbouring Geons. Scaling based upon the ratio of the image at the seam ensures a better fit.

Fig. 8 Scaling lines linearly produces suboptimal results. Row (a) shows downscaling for a single geon and row (b) shows upscaling for a geon in context. The original image is presented in column 1, and column 2 shows a linear scaling. This produces artifacts such as flattening of the feet and a rotation of the waistband. Column 3 shows the results of scaling when curves are preserved and these artifacts no longer occur.

standard deviation to the range we can guess the type of distribution, and hence the type of the property. If σ is small compared to r, the distribution is likely to be clustered around the mean with only a few outliers, and the property is therefore likely to be part of the artists' style. If σ is large compared to r, the distribution is likely to have a wide spread. This means the property is unique to each image.

This categorisation method is also run on the weighted grid extracted during the prototype stage (Figure 6). The length of each gridline is placed in the database and initialised as an artist or image property. Smoothing is performed on adjacent lines to prevent large disconuities during the transfer.

7 Proportion Transfer

The final stage to complete a style transfer using these Characteristic Proportions is to transfer relevant geons in the source image to fit the characteristic proportions of the target style. Properties that measure the differences between geons depend on the properties of the geons themselves, and hence it is possible to end up in a situation where it's impossible to fullfill all the constraints. This means that it may not be possible to transfer all properties correctly, and therefore an order of precedence is established whereby the metrics with the most impact will be processed first. Four passes are run across the metrics, and properties within each pass are solved beginning with the geon that has the lowest standard deviation.

1. In the first pass, properties that are marked as characteristic for both the source image and the target image are processed first. This is given priority because removing the influence of the original artist and replacing it with the target artist creates the biggest impact.
2. Properties marked as characteristic for only the target image are processed second. If possible, the characteristic will be modified to take on the target image characteristics. However, this is only done if it does not interfere with the transformations performed in the first pass.
3. A special case arises where properties are marked as characteristic only for the source image. The style transfer is more effective if all original characetristic properties are removed, however there is no target characteristic value. Instead of completely transforming the geon, properties are adjusted by a random amount in the direction of the target mean.
4. Remaining properties aren't marked as characteristic for either artist, and are likely to be important to the image itself. These are left unmodified.

If transformed images are used in several locations, or if several different images are transformed, they should not have exactly the same proportions. To better reproduce the target characteristics, random values are generated for the target characteristic that fit with the measured average, distribution, and range.

The result of performing the transformation and introducing these random values is that discontinuities appear at the borders of neighbouring Geons. Image (a) in Figure 7 demonstrates a case where Geons have been scaled disproportionally and the discontinuities are highly visible. Both the image prototype and the image extraction steps record the actual width of the image at the borders of Geons. The connectivity information from the prototype generation is used to determine borders, and Equation 2 is used to scale borders to preserve line continuity as shown in Figure 7 (b).

$$R_{scale} = T_{border}/S_{border} * (T_{scale}/S_{scale}) \quad (2)$$

where R is the result image, T is the target image, and S is the source image.

After applying border scaling, the contents of the Geon are also scaled according to the cross-sectional grid using Equation 2. It was found that a subdivision of 6 gridlines produced the best visual results for the style transform.

Another common transformation for geons is an unequal scaling of the dimensions. This causes graphical artifacts because it modifies line curvature and significant points such as eyes or hands become stretched unrealistically. Stretching linearly often disrupts the perspective or distorts geometric objects. Note that this is not affected nor solved by the grid scaling, which runs independantly of the dimensions of the Geon. There are two common problem cases that are both handled with the same algorithm. Row (a) in Figure 8 shows the case where geons are reduced disproportionately in one direction, resulting in flattened shapes. Row (b) shows the extension of geons, often resulting in stretched shapes such as the flower on the chest and the waist band.

To solve this issue, images should be scaled only at places with minimal texture or features. The complexity along the image is recalculated based upon the scaled, border-aligned, content-shifted Geon, in the direction of the unequal scaling. A cross-section with low complexity is selected as close to the middle as possible. This is then stretched or removed entirely depending upon the direction of scaling. Directly removing a section of the image introduces the same discontinuity problem as at Geon borders, and so is solved using the image widths at that point with Equation 2.

Fig. 9 Source images used for the transfers in Figure 10

8 Results

Our algorithm was designed to extract an individual artist's style and apply it to other drawings. The process focuses on the proportions of an image, and in this respect it successfully transfers the characteristics from one artist to another. The process is entirely automated, an outcome that is important in contributing to the original goal of reducing artist workload. The effect of style translation using characteristic proportions can be seen clearly when the modified characters are viewed in context. In Figure 1, characters are shown from 'Sandra and Woo' [8] and 'Misery Loves Sherman' [3], as originally drawn by Puri Andini and Chris Eliopoulos respectively. Andini's characters have their characteristic properties modified to those of Eliopoulous' characters, and are then placed into a cartoon panel alongside the original. Due to the proportion transfer, the characters do not appear out of place and the result of style transfer allows Andini's characters to appear correctly in the same scene as Eliopoulos' characters.

The process has been run on a number of line drawn datasets with success. Three datasets are used in this paper, and where available, at least ten instances of each object are used to build the database. Using several different characters from four cartoons [8,3,15,11] shown in Figure 9, a target was applied and the results are shown in Figure 10.

In the introduction, we mentioned the difficulty of ensuring stylistic consistency across a large range of work. Although our algorithm would struggle to perceive the difference between two artists who strove to draw in the same style, it produces usable results in most other cases.

Due to the random influence in the proportion transfer, the results are never identical when running the transfer algorithm multiple times. Figure 11 shows five different transformations using Sandra from 'Sandra and Woo' [8] as the input image and 'Misery Loves Sherman' [3] as the target style.

One drawback of our approach is that an adequate amount of input data is required. One common scenario is a lack of source or target material, forcing the process to be run with only one side (or neither side) of the database populated. In the case where no database of source images exists, the single source image has all properties marked as artist-specific so that only characteristic properties are replaced on the target image. Figure 12 shows images produced when one or both databases are lacking. In the case where no database of target images exists, the process can be less effective but still produce usable results. Finally, running the style translation process is possible with no databases

Fig. 10 Several examples of the sytle transfer method. The left column is an example of the target style and the other columns have characters converted to that style.

at all, however the results can range from passable to unusable.

It is also possible to perform style translation on photos and images that are not line-drawn, although some of the processing stages must be skipped or performed manually. Figure 13 shows a situation where style transfer can be performed.

In addition to modifying images, the geon database can be used to identify the artist for existing images. By measuring how closely the new figures conformed to the characteristic properties of both artists a confidence value could be built and the character's artist identified. Figure 14 shows how this geon recognition produces correct results even for very similar datasets.

Fig. 11 Introducing a random element to the transform correctly reproduces the scatter of values in the target style. The algorithm is run on one source character five times to generate a range of results. It can be seen that each generated character conforms to the target style, yet has its own unique pose.

In Figure 14, three additional figures were chosen from the existing artists and correctly recognised.

Additional results can be seen at http://www.geons.tk/

9 Conclusion

This paper outlines the first results of transferring an artist's style from one image to another. The proportion transfer runs without requiring human judgement and produces acceptable results for a range of images.

Human perception has yet to be rivaled by computers, in large part due to the difficulty of contextual analysis. Our algorithm fills the gap between the functional context-free style analysis of Hertzmann et al. [7], and the non-functioning attempts at fully contextual systems. The spatially linked geons extracted from the image form a type of context that allows us to avoid the problems encountered by Freeman et al. [5]. Measuring

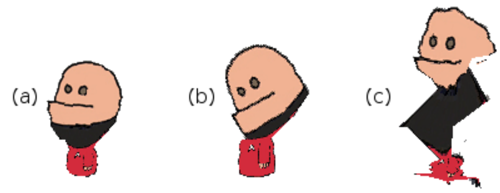

Fig. 12 A character from 'Yodablog' is transferred to the style of 'Misery Loves Sherman'. Transfer without a target database (a) produces acceptable results. Without the source database (b), the angle of the head is incorrectly transferred, resulting in stretching. Without any databse (c), the internal gridlines are directly adjusted to aproximate the shape of sherman - giving incorrect results.

Fig. 13 An example of transferring a cartoon style to a photographic image.

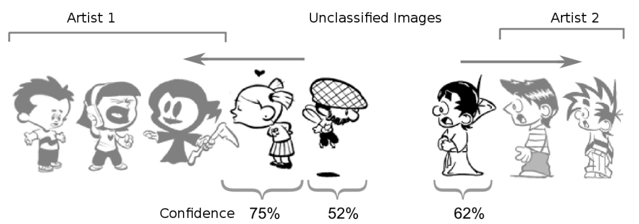

Fig. 14 Characteristic proportions can be used to classify unknown images. Three previously unused figures were compared with our analysed characteristic properties for each artist. The confidence of each categorisation is shown below the figure. The characters on either side were used to build the database of geons against which the matching was carried out.

geons allows spatial context to be taken into account without needing to classify the object itself.

The algorithm produces reasonable and predictable results when using line drawn cartoons. If both cartoons use similar types of line art, transformed images match the style of the target image accurately. However, two arbitrarily chosen artists are unlikely to have the same type of line art, and may draw in colour. In these cases, further processing would be needed to fully complete the style transfer. Algorithms exist to transfer the brush stroke style, however factors such as shading, colours, and image layout are also artist dependent. A similar analysis process could be applied to transform these.

Artistic style is a very broad topic, and relies in no small part upon human perception. Our process identifies geons and transforms them based upon a database of measurements. Although the results appear "correct" at a glance, it is difficult to evaluate whether a correct style transfer has actually been performed. Although this is a subjective judgment, a survey based evaluation of success could allow comparisons between different style transfer techniques and provide a benchmark to measure progress.

We believe this research opens up further opportunities in the area of style manipulation. This research begins to address the area of visual style, and shows that if the correct approach is found then it is possible to process art in ways that are often thought of as subjective.

Acknowledgements

Copyright of original images belong to the authors of the respective comic strips and are reproduced here with permission. Some images are creative commons.

References

1. S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(4):509 –522, apr 2002.
2. Praveen Dasigi and C.V.Jawahar. Efficient graph-based image matching for recognition and retrieval. In *Proceedings of National Conference on Computer Vision, Pattern recognition, Image Processing and Graphics (NCVPRIPG-2008)*, jan 2008.
3. Chris Eliopoulos. Misery Loves Sherman. `http://www.miserylovessherman.com/`, 2010.
4. Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24:381–395, June 1981.
5. William T. Freeman, Joshua B. Tenenbaum, and Egon C. Pasztor. Learning style translation for the lines of a drawing. *ACM Trans. Graph.*, 22, January 2003.
6. Andrea Lynn Frome. *Learning Local Distance Functions for Exemplar-Based Object Recognition.* PhD thesis, EECS Department, University of California, Berkeley, Aug 2007.
7. Aaron Hertzmann, Nuria Oliver, Brian Curless, and Steven M. Seitz. Curve analogies. In *Proceedings of the 13th Eurographics workshop on Rendering*, pages 233–246, 2002.
8. Oliver Knörzer and Puri Andini. Sandra and Woo. `http://www.sandraandwoo.com/`, 2010.
9. Xuetao Li, Tong Wing Woon, Tiow Seng Tan, and Zhiyong Huang. Decomposing polygon meshes for interactive applications. In *Proceedings of the 2001 symposium on Interactive 3D graphics*, I3D '01, pages 35–42, 2001.
10. David G. Lowe. Object recognition from local scale-invariant features. In *Proceedings of the International Conference on Computer Vision-Volume 2 - Volume 2*, ICCV '99, 1999.
11. Crowden Alexander Satz. Not Just Nicky. `http://www.nicky510.com/`, 2011.
12. Lior Shamir, Tomasz Macura, Nikita D. Orlov, Mark Eckley, and Ilya G. Goldberg. Impressionism, expressionism, surrealism : Automated recognition of painters and schools of art. *ACM Trans. Appl. Percept*, 2010.
13. Christian Theobalt, Christian Rössl, Edilson de Aguiar, and Hans-Peter Seidel. Animation collage. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics symposium on Computer animation*, SCA '07, pages 271–280, 2007.
14. Shimon Ullman, Michel Vidal-Naquet, and Erez Sali. Visual features of intermediate complexity and their use in classification. *Nature neuroscience*, 5(7):682–687, jul 2002.
15. Thierry Vivien. Yoda blog. `http://www.yodablog.net/`, 2011.
16. P. Wayner. Identification of artistic styles using a local statistical metric. In *Proceedings of IEEE Artificial Intelligence Applications*, pages 110–113, 1991.