

Geometry Super-Resolution by Example

Thales Vieira Alex Bordignon Thomas Lewiner
Matmídia Laboratory
Department of Mathematics, PUC-Rio. Rio de Janeiro, Brazil
www.matmídia.mat.puc-rio.br/{thalesv,alexlaier,tomlew}

Luiz Velho
Visgraf Laboratory
IMPA. Rio de Janeiro, Brazil
www.impa.br/~lvelho

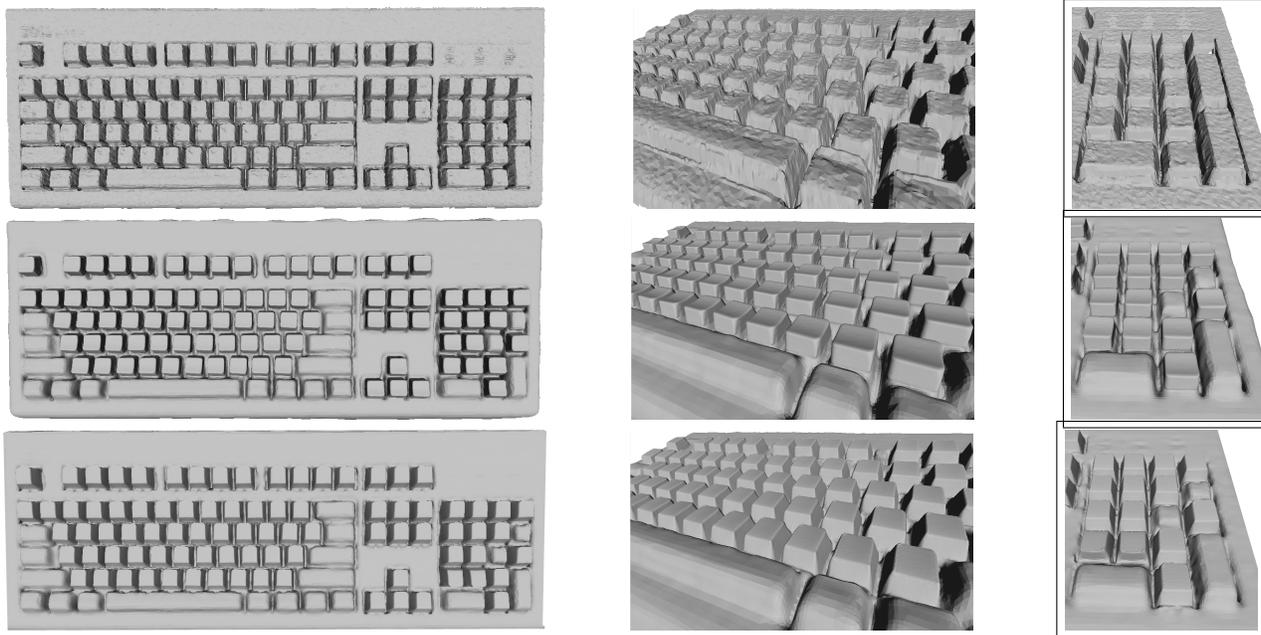


Figure 1. Super-resolution of a scanned keyboard with a high-resolution scan of one of its keys: original model (top), reconstructed with the detail exemplar of a key scanned at high resolution (middle) and reconstructed with the exemplar key generated by low resolution accumulation (bottom).

Abstract—The acquisition of high-resolution 3D models still requires delicate and time-consuming processes. In particular, each detail of the object should be scanned separately, although they may be similar. This can be simplified by copying a small set of details at different places of the model, synthesizing high geometric resolution from details exemplars, as introduced in this paper for three different contexts : when the detail exemplars are scanned separately at high resolution, when they are synthesized or edited from other models, or when they are obtained by accumulating repeated instances of the detail in the low-resolution scan. The main challenge here is to correctly register the high-resolution details with the low resolution model. To address this issue, this work proposes a careful resolution manipulation of 3D scans at each step of an automatic registration pipeline, combined with a robust selection of alignments. This results in a fully automatic process for geometry super-resolution by example. Experiments on synthetic and real data sets show applicability in different contexts, including resolution increase, noise removal by example and geometric texture insertion.

Keywords-Super-resolution; 3D Registration; 3D Scanning; Surface Reconstruction;

I. INTRODUCTION

With the advent of cheaper scanning devices, 3D model acquisition became more popular. However, the scanning of high-resolution model remains a complex and time-consuming task. Regardless of the equipment used, the 3D scanning hardware imposes a limit in the highest achievable resolution of a single scan. Additionally, the acquired data is contaminated by sensor noise, further compromising the effective model resolution.

There is an inherent trade-off between scanning time and model resolution: While for digitizing an object in low resolution only a few broad scans may suffice ; To scan the same object in high resolution many close scans would be needed to capture each fine detail, even if they are repeated in the model. This trade-off can be improved if we reuse a single exemplar of each fine detail, turning the scan process simpler, faster and still producing models with high resolution. This goal can be accomplished using *super-resolution by example!*

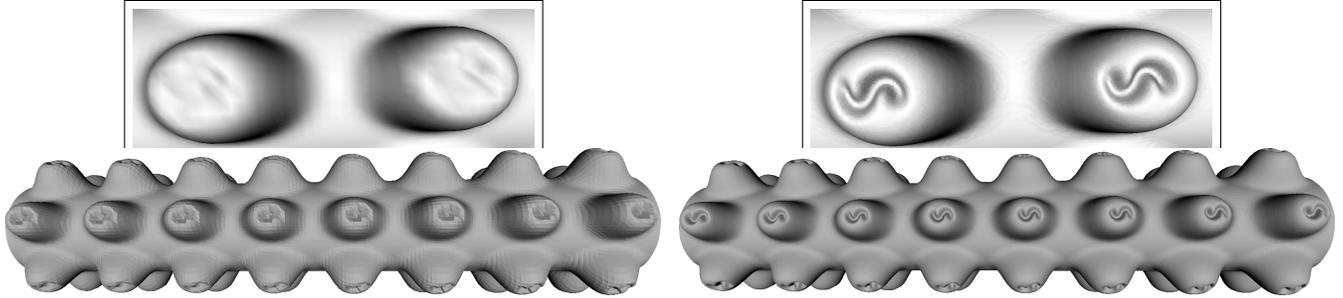


Figure 2. Intelligent geometry insertion: substituting the exemplar detail for a synthetic texture, a new object can be obtained by super-resolution, incorporating the texture automatically at the right places.

The main idea consists in exploiting the fact that most human-made objects have a very coherent local structure that can be summarized by few exemplar features. In this context, an efficient strategy for high-resolution scanning is to quickly digitize the entire object in low-resolution with a small number of scans, and to subsequently scan or extract the few characteristic features of the object in high resolution (Figure 1). The scanning process stops here. As a post-process, the super-resolution method reconstructs the model in high resolution by extrapolating the global geometry based on the local exemplars.

Contributions: Super-resolution is a powerful method that has already been applied to images and video. However, its adaptation to geometric models is non-trivial due to several technical difficulties, among which registration with different resolutions and the robust identification of many detail occurrences in large low-resolution models. The contribution of this paper is the development of a fully automatic method for geometric super-resolution by example that overcomes these issues. To do so, we propose a careful resolution manipulation of 3D scans at each step of an automatic registration pipeline, combined with a robust selection of valid alignments.

This approach actually applies to slightly different contexts than classical super-resolution, when the resolution of the exemplar detail is higher than the whole model. Depending on the context, the exemplar can be obtained from a separate scan at high resolution (Figure 1), when the user has access to the scanning process, or synthesized artificially for intelligent geometry texture insertion (Figure 2). The exemplar can also be retrieved by accumulating the repeated low-resolution instances of a selected detail, identified by the same mixed-resolution registration (Figure 11).

II. RELATED WORK

There are two basic approaches for super-resolution: removing noise from the accumulation of samples and detail synthesis by example.

Super-resolution from multiple samples exploits redundancy in replicated data sets of the same scene in order to

filter out noise and enhance the signal, effectively recovering information beyond the sampling frequency [1]. This approach suits well for video sequences, as proposed for example by Schultz and Stevenson [2] or Irani and Peleg [3]. It has been used for geometry super-resolution in the work of Kil *et al.* [4]. By scanning several times the same object from slightly displaced viewpoints, the authors obtain an accumulation of geometry that can be used to remove noise. However, this slows down the scanning process.

Super-resolution by example exploits the global coherency of a scene. More precisely, a small set of example features generally contains a summary of information about the multi-scale structure of the scene. By properly analyzing these example features, they can be used to predict and extrapolate high-resolution details of a coarse resolution version of the scene. This approach has been first introduced from fractal-based image compression [5] and then extended to high-resolution synthesis [6]. Similar ideas have been developed for 3D models [7], [8]. They reconstruct a surface from a noisy point cloud and prior shapes. Pieces of these shapes are compared with the point cloud geometry descriptors, and a Bayesian process validates such match. In their work, the prior shapes and the point cloud have the same resolution, allowing the inclusion of non-local geometrical contexts for matching. In [9], geometric moments are used as descriptors to match patches from high quality prior shapes to the point cloud. Then, the point cloud is augmented by a variation of the MLS projection procedure.

Our approach mainly follows the super-resolution by example paradigm. However, differently from these approaches, we match noisy detail at high resolution entirely onto a large low-resolution model: no geometry context is used for matching. We also borrow from the multiple samples approach when the high-resolution exemplar is not available for scan, in order to generate it by accumulating its repeated occurrences, approaching structure extraction applications[10]. In that sense, this work can be used as geometry texture synthesis [11], [12], allowing to correctly position texture elements from low resolution information.

III. REVIEW OF SINGLE RESOLUTION SCAN REGISTRATION USING SPIN IMAGES

One of the main difficulties of adapting super-resolution methods originally conceived for images to 3D geometry is that an image has a trivial topology and also possesses a regular structure while surfaces generally do not. In order to overcome this difficulty it is necessary to resort to invariant feature analysis that are coordinate independent and also employ registration methods for point sets. Spin images [13] constitute one of the most successful feature types for the analysis of geometric data. In our technique we use spin images to identify correspondences between parts of the complete coarse resolution model and exemplar structures containing high-resolution details. This section details how to use such local descriptors for registration [14], [15].

Scan registration pipeline: The registration process tries to identify the matching parts of models and estimate the spatial transformations that align them. When restricting to rigid transformations, the identification can be done by associating to some significant points of the model a rigid-invariant descriptor. The significant points are usually chosen as high mean-curvature points in each range image. The matching descriptors then correspond to matching parts of the models. Since the matching may be imperfect, these correspondences must be filtered, and the alignment refined, typically using iterative closest points algorithms [16].

Spin images: Given a reference point p , its spin image is bi-dimensional density histogram of the neighbors q of p (Figure 3). To be rigid invariant, this histogram is represented in radial coordinates: the two coordinates of q are its distance to the tangent plane at p and its distance to p when projected in the tangent plane. This histogram is normalized to cope with slight variations of the sampling rate in the two models.

The spin-image comparison is performed using linear correlation coefficients and specific statistical measures [13]. In particular, it counts only the number of correct matchings between the histograms, and does not penalize unmatched parts. Two parameters define the spin image: the histogram size and the extension of the neighborhood. For usual scan registration, the histogram size varies from 10 to 20, and the extension is set to twice the point density to ensure enough points in each class of the histogram [13], turning the spin image less sensitive to noise.

Correspondence grouping: The robustness of this descriptor is improved by geometric consistency filtering. For each reference point in the first model, there should be a corresponding point in the second model with similar spin image. Groups of such correspondence should be coherent, in particular the distances between points should be preserved across correspondences. This geometric consistency test would require checking every combination of correspondences. To optimize this, a greedy strategy prioritizing correspondences that are far apart is generally used, since far apart

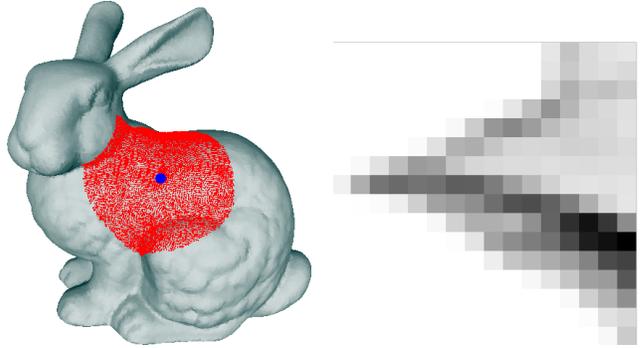


Figure 3. Spin images at the central point (right) are rigid-invariant geometry descriptors for a neighborhood of a reference point (left).

correspondences improve the transformation accuracy. Each group of correspondences generates a rigid transformation, and the greedy search stops when the transformed models have a sufficient overlapping.

This procedure *as is* would be ineffective for super-resolution purposes, since it does not cope with different resolutions and eventual multiple valid transformations.

IV. MIXED-RESOLUTION REGISTRATION

In this work, we automatically identify and replace all instances of a given high resolution *detail* in a low resolution 3d *model*, which can be a triangle mesh or an oriented point cloud. The proposed method can be decomposed into five steps, each of them suffering significant modifications for super-resolution. Initially, we decompose the model and the detail in two resolutions: the original resolution, which will be used in the final result, and a low resolution, used for registration purposes. We choose the resolution carefully to preserve as much information as possible while reducing both the model and the detail to a common ground (Section IV-A). For super-resolution from multiple samples, the user selects the detail example from the model, so that this phase is unnecessary. We further compute the spin-image parameters for small details matching (Section IV-B). We then filter and group correspondences from different occurrences of the detail (Section IV-C). After that, we identify the unique occurrences of the detail in the model by detecting and removing duplicate or symmetric transformations (Section IV-D). Finally, we incorporate the new geometry information in the model (Section IV-E).

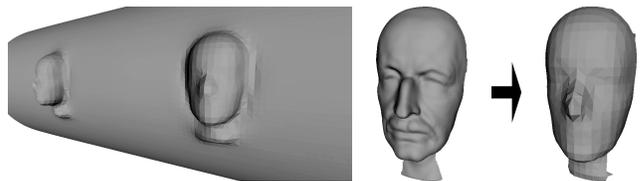


Figure 4. Scaling the model (left) and the detail (middle) to a common resolution (right) by reconstructing them at the same octree depth.

A. Scaling

In order to identify information from the detail to the model, we scale them to a common resolution (Figure 4). We estimate the resolution r of a triangle mesh as the median of the edges’ lengths for meshes, and using the k nearest neighbors for point sets, as described by Pauly *et al.* [17]. For noisy triangle meshes, we reduce the estimated resolution by multiplying r with the inverse of the mean curvature variance, following the noise estimation of Vieira and Shimada [18].

The scaling must maintain a reasonable amount of feature points to serve as references for the spin image. To ensure that, the low resolution model is simply filtered by a two-step normal smoothing / vertex fitting method [19]. To scale the high-resolution detail, we use a Poisson reconstruction [20] at low resolution, setting the maximal depth of its octree to $\log(r)$. The Poisson scheme is mainly based on normals, which help in preserving the locus of high-curvatures in the detail. This way, the scaling process removes high frequencies from both the model and the detail, keeping the common high curvatures, and naturally eliminates random noise.

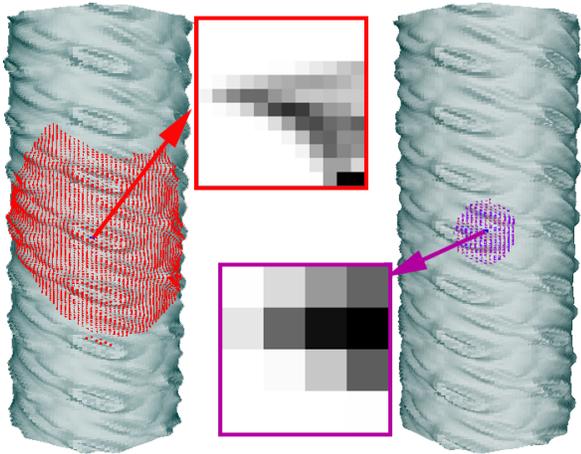


Figure 5. As opposed to scan registration (left), the extension of spin-image for super-resolution must be adapted (right). It is automatically computed from the resolution of the model and the size of the detail.

B. Local feature matching

The classical spin-image design and matching described at Section III would fail for our super-resolution application. In particular usual spin-images use large neighborhoods, since in that case we expect large overlapping around few very sharp feature points. In our context, we will prefer smaller neighborhoods and more feature points since the detail is small and its high curvatures have been smoothed in the scaling process (Figure 5).

More precisely, we automatically set the neighborhood extension for the spin images to $n = \alpha d$, where d is the

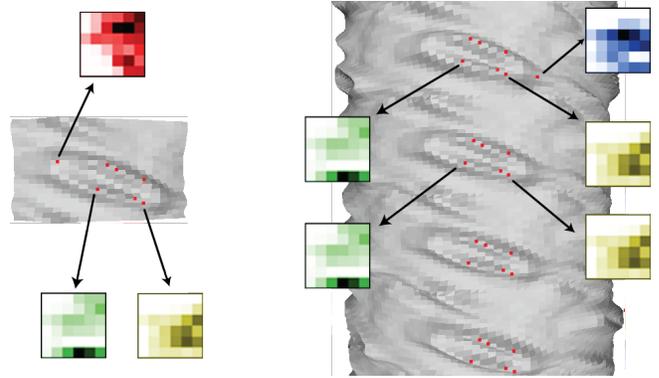


Figure 6. Matching between the spin images of reference points (red dots) in the detail (left) and the model (right).

length of the detail’s bounding box diagonal. For the spin-image reference points, we select the top β highest curvature points in the detail, or at least 10 points to guarantee matching robustness (Figure 6). In our experiments, we obtained good results with $\alpha = \frac{1}{4}$ and $\beta = 10\%$. To assure that matching reference points in the model will be selected, we use the interval defined by the smallest and highest curvature of the selected points in the detail. Moreover, we avoid creating spin-images with neighborhoods overflowing the detail. To do so, we only select reference points in the detail at distance less than n to its center.

Finally, when comparing two spin-images, we take into account every entry of the histogram, as opposed to the classical matching, where only the overlapping entries are considered. Given two spin-images P and Q , the standard linear correlation coefficient R_{PQ} is defined as

$$R_{PQ} = \frac{1}{\sigma_P \sigma_Q} \left(N \sum_i p_i q_i - \sum_i p_i \sum_i q_i \right),$$

where σ_P is the standard deviation of histogram P and N is the number of bins. In the classical spin-images matching, the similarity measure is defined as $C_{PQ} = \text{atanh}^2(R_{PQ}) - \frac{\lambda}{N-3}$. C is then a loss function that returns higher values for highly correlated images that have a large number of overlapping bins (N). The λ parameter is given by the expected number of overlapping pixels entries. In particular, this strategy is well suited when considering geometry that do not fully overlap, like classical 3d scan registration. For mixed resolution, we look for full overlap, so we consider only the linear correlation coefficient R_{PQ} .

C. Correspondences filtering

Every matching pair of spin-images corresponds to a positioning of the detail in the model. The valid occurrences are obtained by selecting coherent groups of such correspondences. Such group is coherent if they induce a rigid transformation that maps the detail onto a single occurrence in the model. The rigid transformation criterion is similar

to the classical case (Section III). The additional single occurrence criterion is specific to our super-resolution. It requires the correspondences inside a coherent group to be distributed in the model on an area smaller than the detail surface. In particular, this criterion greatly accelerates the process, since only a small number of correspondences combinations must be tested.

D. Identification of the detail occurrences

Each coherent group of at least three correspondences induces a transformation T . However, T may not map the detail onto one of its occurrences. Moreover, different groups can lead to similar transformations, in particular since the model may contain only approximate copies of the detail (Figure 7). Therefore, we first filter each transformation T according to the overlap between the detail mapped by T and the model. We are looking *a priori* for fully overlapping detail occurrences, i.e. 100% of the transformed detail points. Observe that this is significantly higher than usual scan registration thresholds. However, as the model is prone to noise and detail occurrences are not exactly identical, we validate transformations whose resultant overlapping points are higher than $k\%$ of the number of detail points. In our experiments, we set k to 90%.

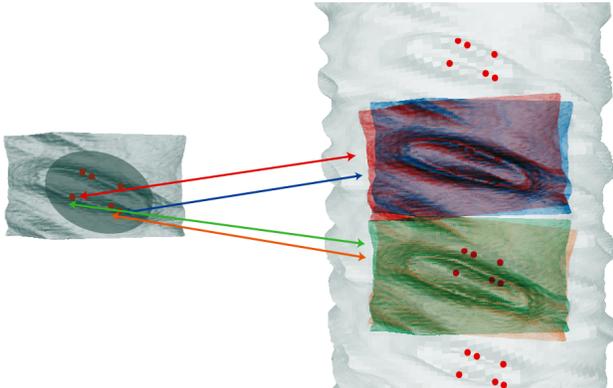


Figure 7. Duplicate transformations must be identified in a symmetry-proof manner.

Secondly, we eliminate duplicate transformations by measuring their difference in a symmetry-proof manner. To do so, we define a pseudometric M between transformations T_1 and T_2 , which considers the axis-aligned bounding boxes of $T_1(detail)$ and $T_2(detail)$. The distance between the resultant bounding boxes lower-left-front points p_1 and p_2 , and upper-right-back points q_1 and q_2 is measured by

$$M(T_1, T_2) = \frac{1}{2} \|p_1 - p_2\| + \frac{1}{2} \|q_1 - q_2\| .$$

When $M(T_1, T_2) \leq \frac{1}{2} \|p_1 - q_1\|$, the transformed details considerably overlap each other, eventually with symmetry. In this case of overlapping, we keep only the transformation that results in the highest overlap with the model.

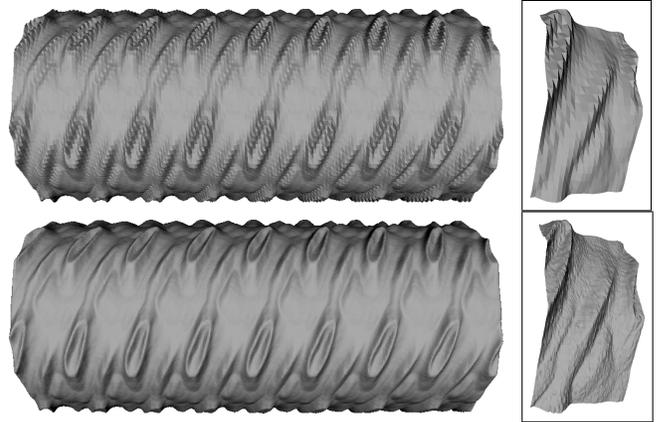


Figure 8. The geometry is merged by replacing the overlap region (top) with the detail (bottom).

E. Detail insertion

Given the set of unique transformations, we refine these transformations by a simple iterative closest point algorithm [16]. Aside from improving the registration accuracy between the transformed detail and the model, this algorithm detects the regions of the model that the detail overlaps. To insert the high resolution into the model, it is necessary to remove this overlapping part. In particular, this is crucial when using this technique for example-based noise removal or geometry texture insertion (Figure 8). We finally merge the geometry by replacing the overlap region in the model by the high-resolution detail geometry.

V. ON THE EXEMPLAR DETAILS

The exemplar details can be obtained through three different means. The best one is to acquire it at high resolution during the scanning process. However, this is limited to applications where the user has access to the scanning device, and when this device allows for zooming until the desired resolution, which is the case of laser or structured light on specific materials. The second way is to synthesize the detail or to retrieve it from an external data base. This case is useful for geometry texture insertion.

In the case where the user has access only to the low-resolution scan, it is possible to gradually generate an ex-



Figure 9. Filtering of the accumulated geometry for the keyboard example: from the aligned detail occurrences (left) are projected in a smooth manifold surface (right).

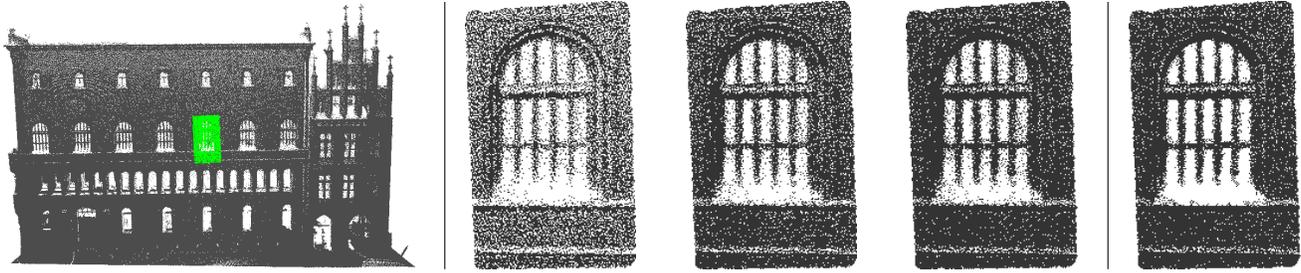


Figure 10. Generation of a high-resolution exemplar by accumulation: from an initial user selection of a window (left), the registration process identifies other occurrences of windows and copy their geometry back onto the initial detail (on the illustration with 2, 4 and 6 windows copied). The accumulated geometry is then filtered using a variation of MLS (right).

exemplar, borrowing ideas from super-resolution from multiple samples: Instead of using a super-resolution detail scanned in high quality, which is not always available, we can use the detail occurrences in the model as multiple samples to reconstruct a super-resolution detail.

The user can select part of the low-resolution scan as input, creating the initial detail. We register this detail with the model, forcing a different location, and copy the geometry of the overlapped region. We generate a new detail by merging the point sets of the initial detail and the copied geometry, aligned back onto the detail. We smooth the merged point set using an MLS variation similar to the work of Weyrich *et al.* [21] (Figure 9). We repeat the operation, accumulating geometry to increase the resolution of the initial detail (Figure 10).

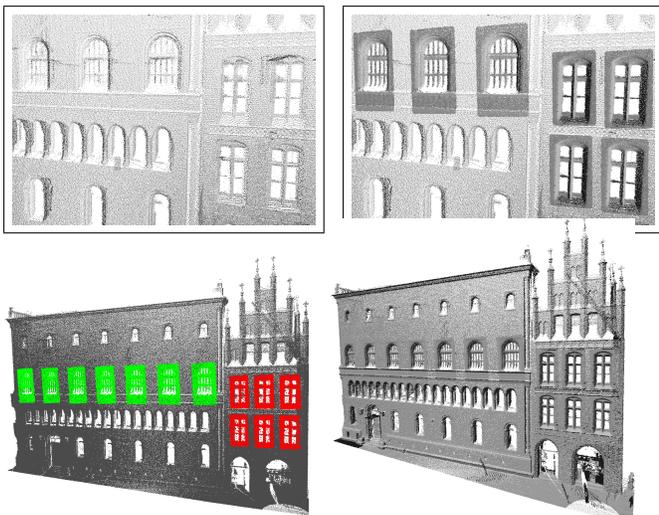


Figure 11. Superresolution on two groups of windows from the rathaus model: identified windows at low resolution (left) and the obtained super-resolution model (right), displayed as a raw point set.

VI. RESULTS

We implemented the above procedure into a fully automatic algorithm that takes two or three group of meshes

or point sets as input: the low resolution models, the high-resolution details and eventually geometric textures to merge instead of the details. We used a kd-tree for neighborhood searching for the spin-image neighborhoods, the correspondence groups creations and to compute overlapping regions. The mean curvature used for the noise estimation and the reference points selection was computed according to the work of Meyer *et al.* [22] for the mesh case, and using the work of Pauly *et al.* [17] for point sets. Finally, we use again Poisson reconstruction [20] on the merged points to display our results as surfaces.

A. Experiments

We apply our method on synthetic and real objects. The synthetic models are cylinders with small meshes manually placed on it, and reconstructed at low resolution (Figures 2, 8 and 5). The details are small perturbations of the small meshes. The real case comes from a laser scan of a keyboard, with one of its key scanned at high resolution (Figures 1 and 14). For this experiment, we used a Minolta VIVID 910 scanner with a telephoto lens. The complete keyboard was scanned at a distance of 2.5m, while the single key detail at a distance of 0.5m.

We further test our technique in other context: geometry texture insertion on the real scan of the keyboard, changing the key detail by an edited one (Figure 13). We also tested it on real data with no scanned high-resolution exemplar, but generating this super-resolution by geometric accumulation. On the keyboard of Figure 1(bottom), we manually selected one key as initial detail. We can observe that the mixed registration process worked as well as with the high-resolution scan, e.g. failing on the same keys of the numeric keypad.

We applied the same technique on the rathaus¹. In that case we selected two different windows as initial details, and detected all their occurrences in the models (Figure 11). Observe that this is a point set model. The timings of the super-resolution steps are reported in Table I.

¹Altes Rathaus Hannover from the Institute of Cartography and Geoinformatics of the Leibniz University of Hannover

Model	detail occurr.	low model points $\times 10^3$	detail ref. points $\times 10^3$	total points $\times 10^3$	spin image	filtering	validation	insertion	total
					corresp. secs.	secs.	secs.	secs.	secs.
max planck	3	53.4	19.3	110.4	2.7	1.9	1.7	1.4	6.2
textured cylinder	48	25.3	9.2	442.3	1.5	0.9	0.9	14.0	15.5
s cylinder	48	25.3	9.2	442.3	1.9	1.4	1.4	13.8	18.7
keyboard	86	88.1	11.6	1,237.8	28.9	15.4	84.3	39.2	128.7
+ flower keyboard	86	88.1	23.6	2,087.3	0.0	0.0	0.0	177.2	177.2
rathaus window #1	7	1,063	17	1,460	256.7	25.3	10.4	30.1	322.5
rathaus window #2	6	1,460	14	1,821	190.5	22.1	9.7	32.6	254.9

Table I

SOME EXPERIMENTS ON SYNTHETIC MODELS AND ON REAL OBJECTS (KEYBOARD, RATHAUS). THE TIMINGS OF EACH STEP OF OUR METHOD REPORTED WITH THE SIZE OF THE ORIGINAL MODEL, THE NUMBER OF DETAIL OCCURRENCES AND THE NUMBER OF INSERTED POINTS AFTER REGISTRATION. EXCEPT FOR THE FLOWER KEYBOARD, WHICH CORRESPONDS ONLY TO THE TEXTURE INSERTION, MOST OF THE TIME IS SPENT ONTO THE SPIN IMAGE CORRESPONDENCE, SINCE THE MIXED RESOLUTION IS MORE DELICATE AT THIS STEP.

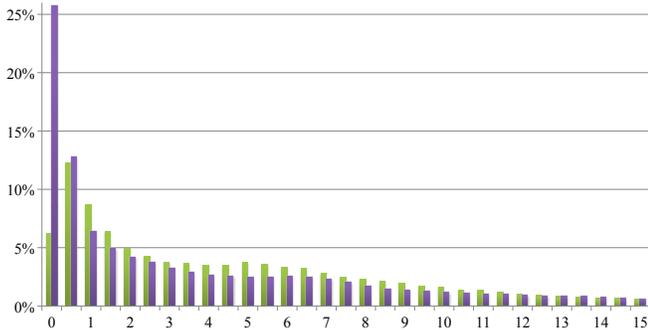


Figure 12. Curvature distributions of the original keyboard model (left green bars) and of the super-resolution keyboard after Poisson Reconstruction (right purple bars): The method removes the noisy keys, represented mostly by middle-valued curvature points in the histogram, and inserts keys with flat regions, which are represented mostly by very low curvature points in the histogram.

B. Discussion

We evaluated the method in terms of robustness and final detail level. For all models, the method was robust in registering all detail occurrences. On synthetic data, it correctly identified all the occurrences of the detail, and aligned them with a very high precision even at low resolution (Figure 8 and Table II). The textured cylinder of Figure 5 shows that the method works well in the border of the details, even when there is no deformation between the detail occurrences.

In the 86 buttons keyboard case (Figure 1), only one extra invalid registration occurred in a bigger button from the numerical pad. This limitation is due to the similar features of the lower corners of the key. Looking more carefully, the scan noise leads to small misalignments of each key, which the simple iterative closest point (ICP) algorithm we used is not able to adjust. These limitations can be reduced by using advanced ICP algorithms, in particular incorporating non-rigid transformations [23].

We finally checked that the proposed super-resolution method actually increases the resolution of the final model, even after the reconstruction method used to generate the final result. To do so, we computed the mean curvature

Resolution	$M(T_1)$	$M(T_2)$	$M(T_3)$	average
L 6	0.41 %	0.28 %	0.25 %	0.31 %
L 7	0.07 %	0.19 %	0.33 %	0.19 %
L 8	0.16 %	0.12 %	0.21 %	0.16 %
L 9	0.15 %	0.11 %	0.20 %	0.15 %

Table II

ACCURACY OF THE 3 DETAIL ALIGNMENTS IN THE SYNTHETIC MODEL OF FIGURE 4. THE RESOLUTION CORRESPONDS TO THE OCTREE LEVEL USED FOR THE MODEL CREATION, AND THE ERRORS ARE RELATIVE TO THE TRANSFORMATION APPLIED FOR THE SYNTHESIS, IN PERCENTAGE OF THE DETAIL BOUNDING BOX DIAGONAL.

distribution before and after the super-resolution process on the keyboard model. We can check on Figure 12 that the super-resolution did remove the noisy keys with middle-valued curvature.

VII. CONCLUSIONS

We proposed a geometry super-resolution by example for generating high quality meshes out of a rough model and exemplars of its details. Our method works by matching medium frequencies between the model and the detail, and then incorporating high frequencies geometric information. This method is effective to speed up the process of high-resolution scanning, and may also serve in other contexts as example-based noise removal and geometry texture insertion. The proposed method is limited to models which have repeated occurrences of a shape, and restrict the resolution increase to the regions of those occurrences. Increasing the resolution of other parts would require inpainting-like tools to extrapolate the geometry [24], together with a super-resolution scheme as an extension of the one proposed here.

ACKNOWLEDGMENTS

The authors would like to thank CNPq, CAPES and FAPERJ for their support during the preparation of this work, and the Institute of Cartography and Geoinformatics of the Leibniz University of Hannover for the Altes Rathaus Hannover data.



Figure 13. The transformations obtained by our mixed-resolution registration can be used for inserting geometric textures by simply changing the exemplar before the insertion process, here on the keyboard of Figure 1.

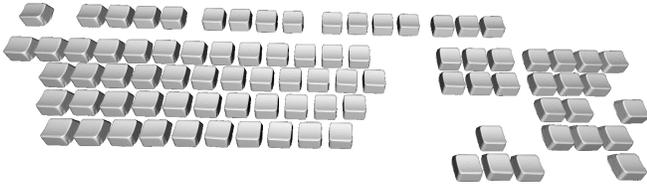


Figure 14. All the 86 keys are correctly matched, and an extra key (“enter” of the numerical pad) has been inserted, although smaller than the real key.

REFERENCES

- [1] M. Bertero and P. Boccacci, “Super-resolution in computational imaging,” *Micron*, vol. 34, pp. 265 – 273, 2003.
- [2] R. Schultz and R. Stevenson, “Extraction of high-resolution frames from video sequences,” *Transactions on Image Processing*, vol. 5, no. 6, pp. 996–1011, 1996.
- [3] M. Irani and S. Peleg, “Super resolution from image sequences,” in *Pattern Recognition*, vol. 90. IEEE, 1990, pp. 115–120.
- [4] Y. J. Kil, N. Amenta, and B. Mederos, “Laser scanner super-resolution,” in *Point Based Graphics*, 2006.
- [5] A. Pentland and B. Horowitz, “A practical approach to fractal-based image compression,” in *Digital images and human vision*. MIT Press, 1993, pp. 53–59.
- [6] W. T. Freeman, T. R. Jones, and E. C. Pasztor, “Example-based super-resolution,” *Computer Graphics and Applications*, vol. 22, no. 2, pp. 56–65, 2002.
- [7] M. Pauly, N. Mitra, J. Giesen, M. Gross, and L. Guibas, “Example-based 3d scan completion,” in *Symposium on Geometry Processing*, 2005, pp. 23–32.
- [8] R. Schnabel, P. Degener, and R. Klein, “Completion and Reconstruction with Primitive Shapes,” in *Computer Graphics Forum*, vol. 28, no. 2, 2009, pp. 503–512.
- [9] R. Gal, A. Shamir, T. Hassner, M. Pauly, and D. Cohen-Or, “Surface reconstruction using local shape priors,” in *Symposium on Geometry Processing*, 2007, pp. 253–262.
- [10] M. Pauly, N. Mitra, J. Wallner, H. Pottmann, and L. Guibas, “Discovering structural regularity in 3D geometry,” *Transactions on Graphics*, vol. 27, no. 3, pp. 43–43, 2008.
- [11] G. Turk, “Texture synthesis on surfaces,” in *Siggraph*. ACM, 2001, pp. 347–354.
- [12] L. Ying, A. Hertzmann, H. Biermann, and D. Zorin, “Texture and shape synthesis on surfaces,” in *Workshop on Rendering Techniques*, 2001, pp. 301–312.
- [13] A. Johnson, “Spin-images: A representation for 3-d surface matching,” Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, 1997.
- [14] T. Vieira, “Registro automático de superfícies usando spin-images,” Master’s thesis, Department of Mathematics, UFAL, 2007, in portuguese.
- [15] T. Vieira, A. Peixoto, L. Velho, and T. Lewiner, “An iterative framework for registration with reconstruction,” in *VMV*, 2007, pp. 101–108.
- [16] P. Besl and N. McKay, “A Method for Registration of 3D Shapes,” *Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992.
- [17] M. Pauly, M. Gross, and L. Kobbelt, “Efficient Simplification of Point-Sampled Surfaces,” in *Visualization*. IEEE, 2002, pp. 163–170.
- [18] M. Vieira and K. Shimada, “Surface mesh segmentation and smooth surface extraction through region growing,” *Computer Aided Geometric Design*, vol. 22, no. 8, pp. 771–792, 2005.
- [19] M. Callieri, P. Cignoni, F. Ganovelli, C. Montani, and P. Pingi, “VCLab’s tools for 3D range data processing,” in *Virtual Reality, Archaeology and Intelligent Cultural Heritage*, 2003, pp. 9–18.
- [20] M. Kazhdan, M. Bolitho, and H. Hoppe, “Poisson Surface Reconstruction,” in *Symposium on Geometry Processing*, 2006, pp. 61–70.
- [21] T. Weyrich, M. Pauly, R. Keiser, S. Heinzle, S. Scandella, and M. Gross, “Post-processing of Scanned 3D Surface Data,” in *Point Based Graphics*, 2004, pp. 85–94.
- [22] M. Meyer, M. Desbrun, P. Schröder, and A. Barr, “Discrete differential geometry operators for triangulated 2-manifolds,” in *Visualization and Mathematics*, 2002.
- [23] B. J. Brown and S. Rusinkiewicz, “Global non-rigid alignment of 3-d scans,” *Transactions on Graphics*, vol. 26, no. 3, p. 21, 2007.
- [24] A. Sharf, M. Alexa, and D. Cohen-Or, “Context-based surface completion,” *Siggraph*, vol. 23, no. 3, pp. 878–887, 2004.
- [25] R. Destobbeli and L. Velho, “Super-resolution,” IMPA, Tech. Rep. TR-2002-08, 2002.
- [26] M. Alexa, J. Behr, D. Cohen-Or, S. Fleishman, D. Levin, and C. Silva, “Point Set Surfaces,” in *VIS ’01: Proceedings of the conference on Visualization ’01*. Washington, DC, USA: IEEE Computer Society, 2001, pp. 21–28.