Texturing and Inpainting a Complete Tubular 3D Object Reconstructed from Partial Views

Julien Fayer[a,b], Bastien Durix[a],
Simone Gasparini[a], Geraldine Morin[a]

[a]IRIT, University of Toulouse [b]SAS InnerSense

June 7, 2018

Introduction: context

- Texturing a 3D model generated from a limited set of images

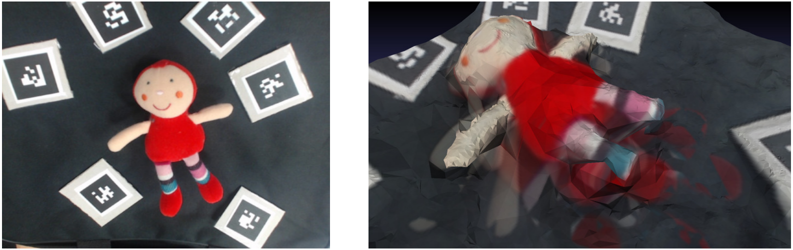

Figure: Reconstruction with the open-source state-of-the-art MVS pipeline AliceVision [AliceVision, 2017] with 4 images (right)

Introduction: context

- Texturing a 3D model generated from a limited set of images

Figure: Reconstruction with the open-source state-of-the-art MVS pipeline AliceVision [AliceVision, 2017] with 30 images (right)

Introduction: our proposal

- Goal: reconstruct a complete textured 3D object

- Hypothesis: object represented by a set of canal surfaces (*branches*) [Durix et al., 2016]

Introduction: our pipeline

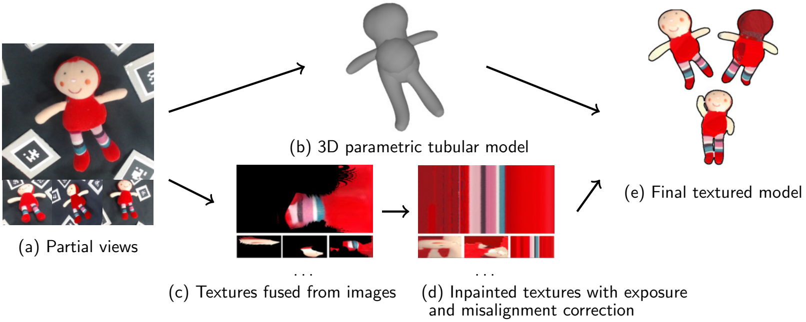

Figure: Pipeline

Outlines

1. Related works
   - Texturing 3D models generated by MVS
   - Texturing 3D models generated by other methods
2. Surface parametrization and texturing
   - Surface parametrization
   - Texturing a branch
3. Improving the textures
   - Exposure improvement
   - Alignment improvement
   - Texture completion
4. Experimental Results
   - Results
   - Limitations

1. Related works
   - Texturing 3D models generated by MVS
   - Texturing 3D models generated by other methods
2. Surface parametrization and texturing
   - Surface parametrization
   - Texturing a branch
3. Improving the textures
   - Exposure improvement
   - Alignment improvement
   - Texture completion
4. Experimental Results
   - Results
   - Limitations

3D Texturing
Single-view

- Only locally optimal → visible artifacts and discontinuities seams
- Solutions: Blending the images, labeling [Sinha et al., 2008], energy minimization [Waechter et al., 2014]

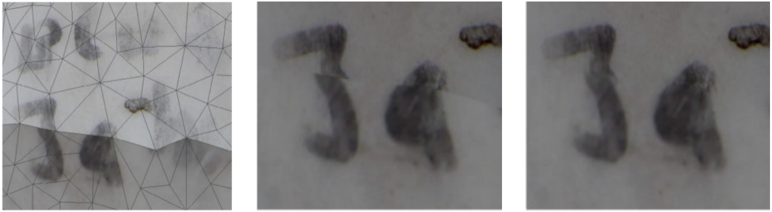

Figure: Solution of [Waechter et al., 2014]

3D Texturing
Multi-views

- Weighted blending of a subset of source images [Callieri et al., 2008] → Suffering of loss of quality and of details + generating artifacts
- Solution: patch based synthesis [Bi et al., 2017]

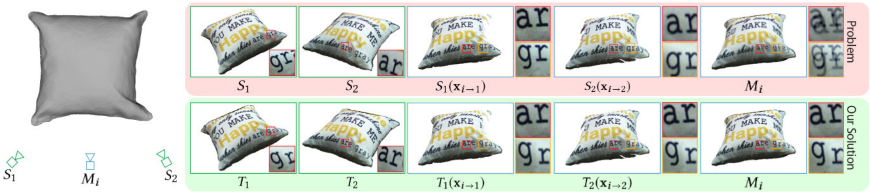

Figure: Patch based synthesis by [Bi et al., 2017]

Other reconstruction methods and texturing

- Shape from silhouette ([Carlos and Schmitt, 2004]): no parameterization → difficult to texture
- Reconstruction of tubular shapes:
  - 3D model from a single image with user interaction: texture recovered by back-projection [Chen et al., 2013]
  - Reconstruction by skeletons [Durix et al., 2016]: camera and geometry alignments issues

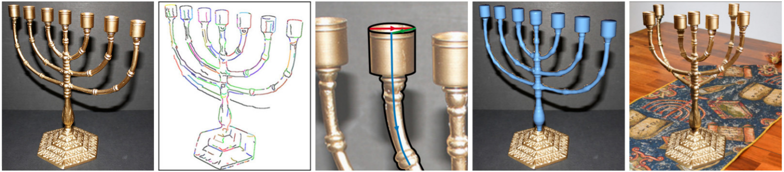

Figure: Results from [Chen et al., 2013]

Other reconstruction methods and texturing

- Advantage: complete reconstruction from partial views
- Contribution: texture reconstructed 3D model including occluded parts

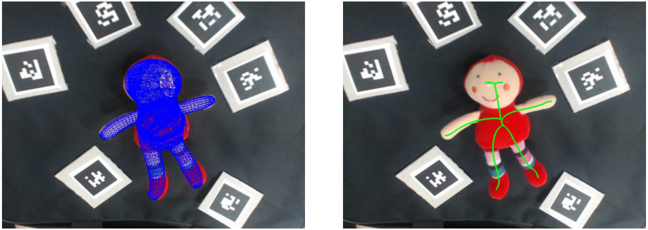

Figure: Estimated 3D skeleton [Durix et al., 2016]

1. Related works
    - Texturing 3D models generated by MVS
    - Texturing 3D models generated by other methods
2. Surface parametrization and texturing
    - Surface parametrization
    - Texturing a branch
3. Improving the textures
    - Exposure improvement
    - Alignment improvement
    - Texture completion
4. Experimental Results
    - Results
    - Limitations

Surface parametrization

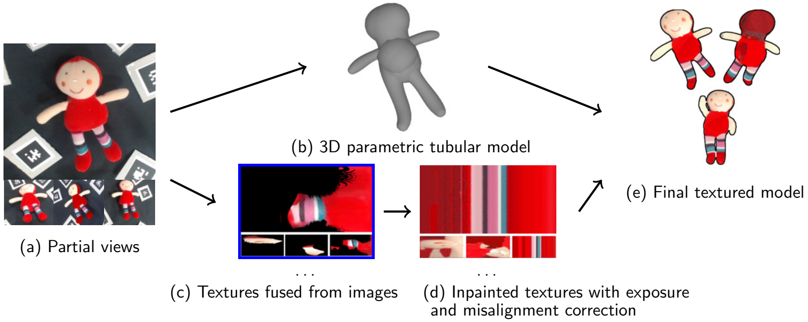

Figure: Pipeline

Texturing a branch from a single image

- Create a reference image texture I_b for each branch b
- Represented by a parametric canal surface $S(t, \theta)$ where $(t, \theta) \in A \times [0, 2\pi]$
- Occlusions handled by an adhoc z-buffer like method

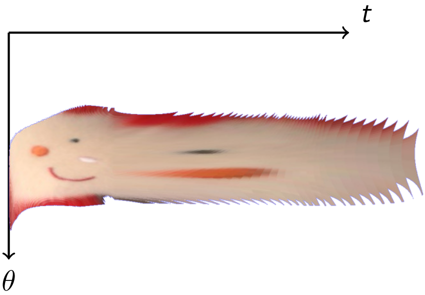

Figure: Surface parametrization

Texturing a branch from a single image

- Create a reference image texture I_b for each branch b
- Represented by a parametric canal surface $S(t, \theta)$ where $(t, \theta) \in A \times [0, 2\pi]$
- Occlusions handled by an adhoc z-buffer like method

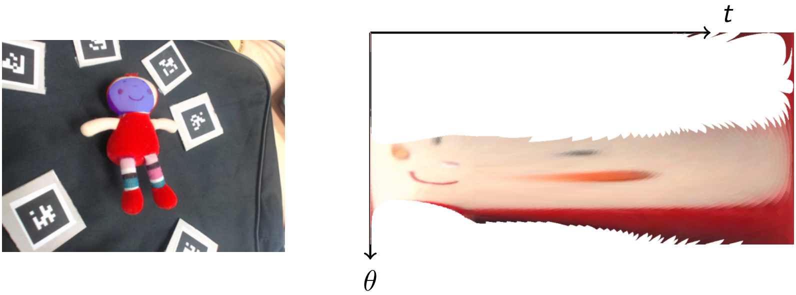

Figure: Surface parametrization

Texturing a branch from a single image

- Create a reference image texture I_b for each branch b
- Represented by a parametric canal surface $S(t, \theta)$ where $(t, \theta) \in A \times [0, 2\pi]$
- Occlusions handled by an adhoc z-buffer like method

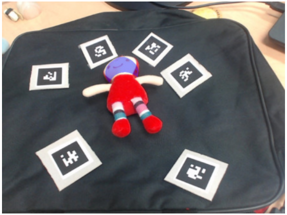

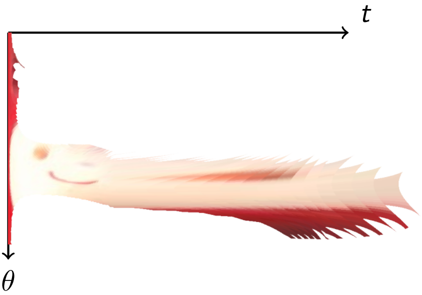

Figure: Surface parametrization

Texturing a branch from a single image

- Create a reference image texture I_b for each branch b
- Represented by a parametric canal surface $S(t, \theta)$ where $(t, \theta) \in A \times [0, 2\pi]$
- Occlusions handled by an adhoc z-buffer like method

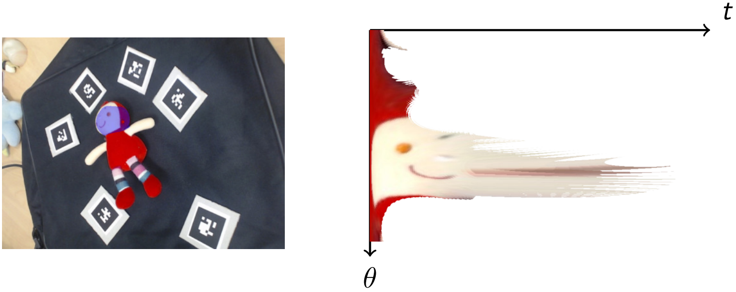

Figure: Surface parametrization

Texturing a branch from multiples images

- Apply the texture of the best image of a set V of images I_i
- Based on a radiometric confidence criterion of our previous work [Fayer et al., 2018]

Figure: Confidence of a reference image

Texturing a branch from multiples images

- Apply the texture of the best image of a set V of images I_i
- Based on a radiometric confidence criterion of our previous work [Fayer et al., 2018]

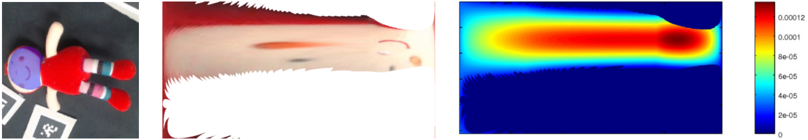

Figure: Confidence of a reference image

Surface parametrization and texturing

Texturing a branch from multiple images

- Each image I_i has a confidence map $\mathcal{C}_i$
- For each pixel p of the final reference image I_b

$$I_b(p) = I_i(p) \quad i = argmax_i(\mathcal{C}_i(p))$$

Figure: Merged reference image

Improving the textures

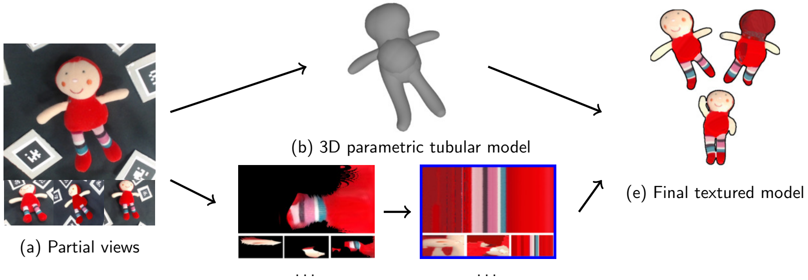

(a) Partial views
(b) 3D parametric tubular model
(c) Textures fused from images
(d) Inpainted textures with exposure and misalignment correction
(e) Final textured model

Figure: Pipeline

Improving the textures

- Remaining issues :
  - Exposure issues
  - Alignment issues
  - Completing the texture

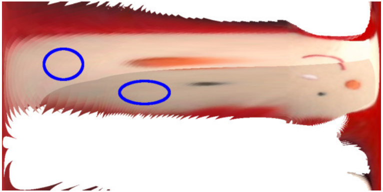

Figure: Exposure differences between two reference images

Improving the exposure

- Inspired from [Zhang et al., 2016]
- Each pixel p of a reference image I_b has a radiance value $r(p)$
- Each image I_i has a exposure coefficient α_i
- Find $(\alpha_i, r(p))$ for each pixel p and each image I_i^b generated from I_i ie:

$$\min_{\alpha_i, r(p)} \sum_{i,p} \left(I_i^b(p) - \alpha_i r(p) \right)^2$$

Improving the exposure

Figure: Exposure correction example

Improving the alignment

- Remaining issues :
  - Exposure issues
  - Alignment issues
  - Completing the texture

Figure: Alignment issues of merged reference image

Improving the alignment

- Inspired from [Arikan et al., 2014]
- Minimize the following energy:

$$E = \sum_{p \in I_b} E_d(p, l) + \lambda \sum_{(p_1, p_2) \in \mathcal{N}} E_r(p_1, p_2, l_1, l_2)$$

Figure: Labels map before and after the minimization

Improving alignments

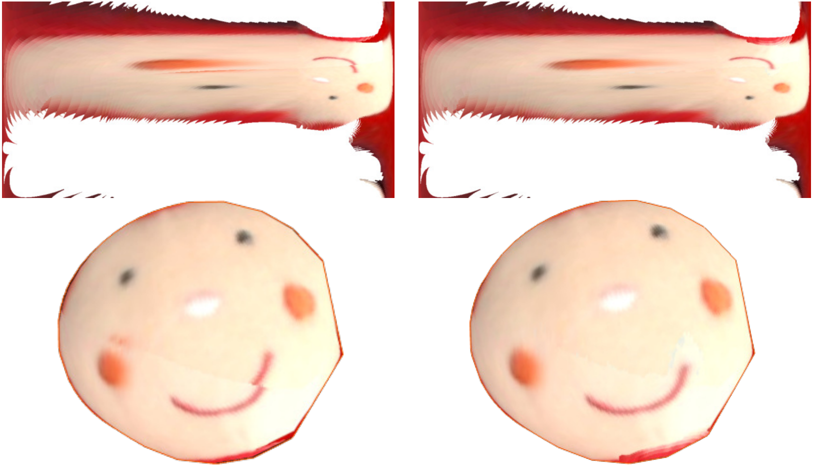

Figure: Alignment correction example

Completing the texture

- Remaining issues :
  - Exposure issues
  - Alignment issues
  - Completing the texture

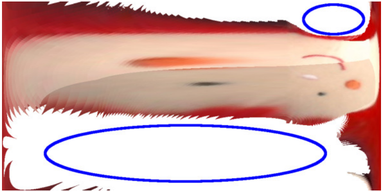

Figure: Missing regions in reference image

Completing the texture

- Two kinds of textures:
  - Regular textures
  - Circular textures: color of the pixel $p(t, \theta)$ is independent of θ
- Apply two completion methods according to the texture type

Figure: Examples of two textures: regular (left) and circular (right)

Inpainting: state-of-the-art

- Diffusion methods : Partial Differential Equations (PDEs) [Chan and Shen, 2001], Total Variation (TV)

Figure: Diffusion method example

Inpainting: state-of-the-art

- Patch based methods: PatchMatch [Barnes et al., 2009], statistical analysis [He and Sun, 2012]

Figure: Patch based method example

Completing the texture

- Create a structure image
- Generate a new texture image with the structure image and the partial texture image

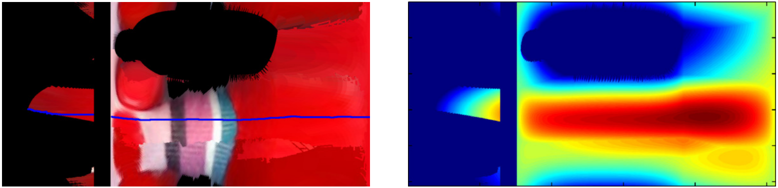

Figure: Circular texture completion

Completing the texture

- Create a structure image
- Generate a new texture image with the structure image and the partial texture image

Figure: Circular texture completion

Completing the texture

- Create a structure image
- Generate a new texture image with the structure image and the partial texture image

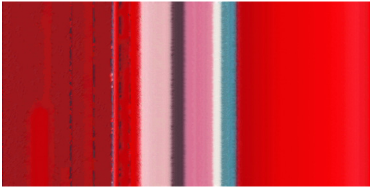

Figure: Circular texture completion

Completing the texture

- Create a structure image
- Generate a new texture image with the structure image and the partial texture image

Figure: Circular texture completion

Context

- Dataset: 5 plushes, 3 to 5 photos for each plush
  - https://github.com/Ibujah/dataset_plushes
- Good candidates: isotropic material, circular surfaces around the direction of the main axis
- Results website: https://sketchfab.com/jfayer/collections

Figure: Input plushes

Results

Figure: 3D models results

Results

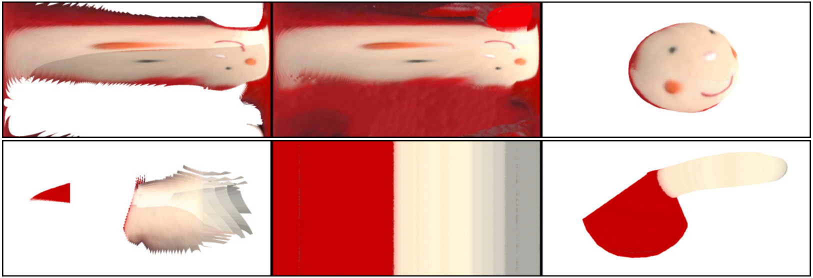

Figure: Mesh completion

Results

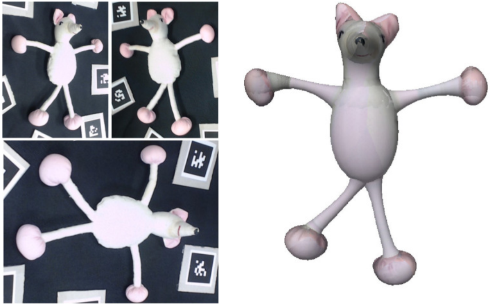

Figure: 3D models results

Results

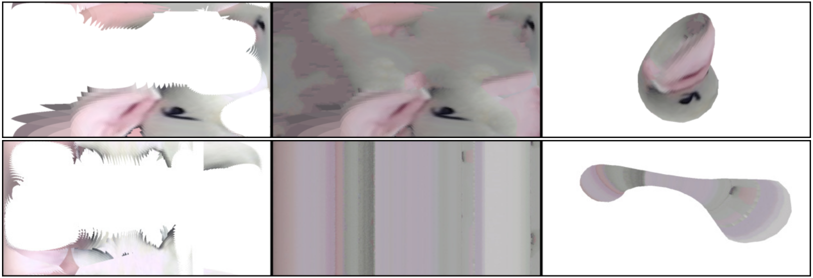

Figure: Mesh completion

Results

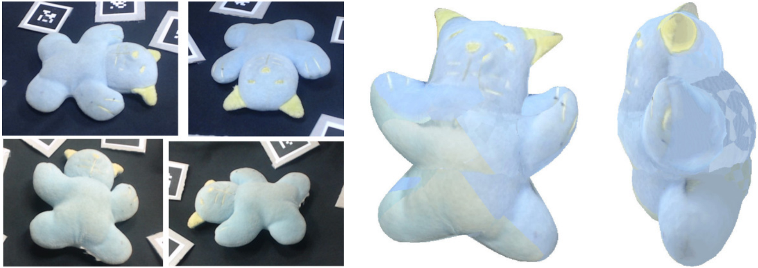

Figure: 3D models results

Results

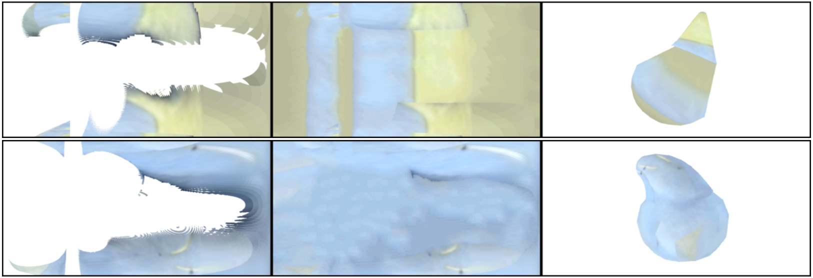

Figure: Mesh completion

Limitations

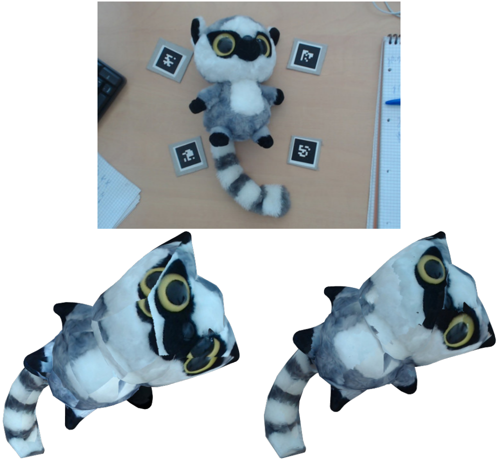

Figure: Limitations example

Conclusions

- Requiring only few images
- Texturing a complete 3D object
- Future directions:
  - Classify the texture into circular or regular
  - Generalize the work on exposure to smoothing of color
  - Handle shiny surfaces

AliceVision (2017).
Photogrammetric Computer Vision Framework.

Arikan, M., Preiner, R., Scheiblauer, C., Jeschke, S., and Wimmer, M. (2014).
Large-scale point-cloud visualization through localized textured surface reconstruction.
IEEE Transactions on Visualization and Computer Graphics, 20(9):1280–1292.

Barnes, C., Shechtman, E., Finkelstein, A., and Goldman, D. B. (2009).
PatchMatch: A randomized correspondence algorithm for structural image editing.
ACM Transactions on Graphics (Proc. SIGGRAPH), 28(3).

Bi, S., Kalantari, N. K., and Ramamoorthi, R. (2017).
Patch-based optimization for image-based texture mapping.
ACM Transactions on Graphics, 36(4):1–11.

Callieri, M., Cignoni, P., Corsini, M., and Scopigno, R. (2008).
Masked photo blending: Mapping dense photographic data set on high-resolution sampled 3D models.
Computers & Graphics, 32(4):464–473.

Carlos, H. and Schmitt, F. (2004).
Silhouette and stereo fusion for 3D object modeling.
Computer Vision and Image Understanding, 96(3):367–392.

Chan, T. F. and Shen, J. (2001).
Nontexture inpainting by curvature-driven diffusions.
Journal of Visual Communication and Image Representation, 12(4):436 – 449.

Chen, T., Zhu, Z., Shamir, A., Hu, S.-M., and Cohen-Or, D. (2013).
3sweep: Extracting editable objects from a single photo.
ACM Trans. Graph., 32(6):195:1–195:10.

Durix, B., Morin, G., Chambon, S., Roudet, C., and Garnier, L. (2016).
Skeleton-based Multiview Reconstruction.
In *Proceedings of the IEEE International Conference on Image Processing*, pages 4047–4051.

Fayer, J., Morin, G., Gasparini, S., Daisy, M., and Coudrin, B. (2018).
Radiometric confidence criterion for patch-based inpainting.
In *Proceedings of the International Conference on Pattern Recognition (ICPR 2018)*.
to appear.

He, K. and Sun, J. (2012).
Statistics of patch offsets for image completion.
In *Proceedings of the 12th European Conference on Computer Vision - Volume Part II*, ECCV'12, pages 16–29, Berlin, Heidelberg. Springer-Verlag.

Sinha, S. N., Steedly, D., Szeliski, R., Agrawala, M., and Pollefeys, M. (2008).
Interactive 3d architectural modeling from unordered photo collections.
ACM Trans. Graph., 27(5):159:1–159:10.

Waechter, M., Moehrle, N., and Goesele, M. (2014).
Let there be color! Large-scale texturing of 3D reconstructions.
In *Proceedings of the 2014 European Conference on Computer Vision (ECCV 2014)*, volume 8693 LNCS, pages 836–850.

Zhang, E., Cohen, M. F., and Curless, B. (2016).
Emptying, refurnishing, and relighting indoor spaces.
ACM Trans. Graph., 35(6):174:1–174:14.

Annexe: AliceVision Results

Figure: AliceVision [AliceVision, 2017] results

Annexe: AliceVision Results

Figure: AliceVision [AliceVision, 2017] results

Annexe: AliceVision Results

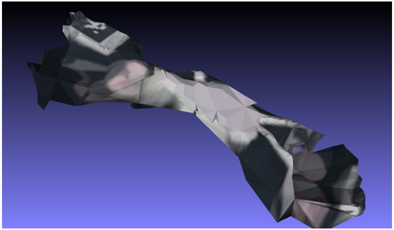

Figure: AliceVision [AliceVision, 2017] results

Annexe: AliceVision Results

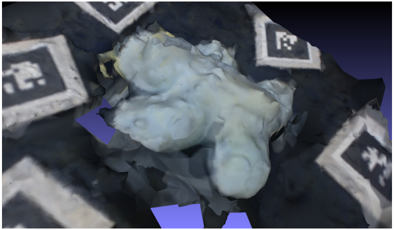

Figure: AliceVision [AliceVision, 2017] results

Annexe: Confidence criterion

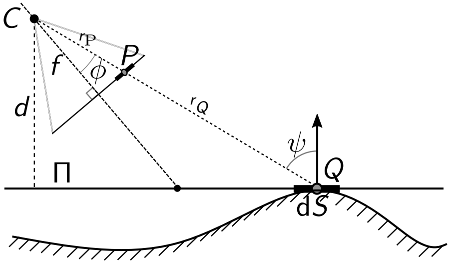

Figure: Confidence criterion applied to a viewpoint C and to a surface dS with tangent plane Π at point Q.

Annexe: Computation times

Plush	Blue	Mouse	Red	Bear	Rabbit
Perspective skeleton estimation	1.7 s	1.4 s	1.7 s	1.8 s	1.8 s
Triangulation	1.3 s	2.4 s	1.3 s	1.4 s	1.6 s
Z-buffering (CPU)	29 s	39 s	27 s	26 s	28 s
Texture extraction from images	9.9 s	15.1 s	9.6 s	8.4 s	9.8 s
Exposure correction	76 s	120 s	71 s	78 s	20 s
Registration correction	88 s	137 s	145 s	106 s	153 s
Texture completion	58 s	128 s	89 s	51 s	57 s
Total	264 s	443 s	345 s	273 s	271 s

Table: Computation times for some plushes