Shadow and Specularity Priors for Intrinsic Light Field Decomposition

EMMCVPR Submission #5 - supplementary material

Abstract

In this supplemental material, we show quantitative and qualitative results of our method on three more synthetic scenes and compare them to previous work. For three more Lytro light fields and one real-world data set, we show only decomposition results for a qualitative evaluation, since there is no ground truth available. We perform comparison with four other approaches proposed by Alperovich and Goldluecke [1], Chen and Koltun [2], Jeon et al.[5] and Garces et al.[3]. All results are presented for the center view, with the specular component scaled for better visibility. For light field methods: ours, [1] and [3] we recorded videos of the decomposition for a cross-hair shaped subset of 13 views. We obtain those views from the 9×9 light field by excluding edge views and extracting a cross-hair shaped subset. Also, we provide videos of the shadow score β. Although in [3] authors provide decomposition result for the whole light field, we record videos only for the cross-hair shaped subset for better comparison. Note: text appearing in red is a link and can be used to navigate the document. Text appearing in blue is a link to a video file.

Contents

References

[1] A. Alperovich and B. Goldluecke. A variational model for intrinsic light field decomposition. In *Proc. ACCV*, 2016.

[2] Q. Chen and V. Koltun. A simple model for intrinsic image decomposition with depth cues. In *Proc. ICCV*, 2013.

[3] E. Garces, J. I. Echevarria, W. Zhang, H. Wu, K. Zhou, and D. Gutierrez. Intrinsic light field images. *Computer Graphics Forum*, 2017.

[4] R. Grosse, M. K. Johnson, E. H. Adelson, and W. T. Freeman. Ground truth dataset and baseline evaluations for intrinsic image algorithm. In *Proc. ICCV*, 2009.

[5] J. Jeon, S. Cho, X. Tong, and S. Lee. Intrinsic image decomposition using structure-texture separation and surface normals. In *Proc. ECCV*, 2014.

Light field *Koala*

Type:	Real world data set captured with the Lytro Illum plenoptic camera
Size:	9×9 views
Resolution:	434×625 pixels
Disparity range:	$[-1.5 \;\; 1.5]$

Scene description: We present a real world data set that we used in the main paper. The scene contains two objects, one is an almost Lambertian koala toy, the other a saxophone, which is highly specular. This scene compares how the algorithms perform on data taken with Lytro camera, and how they deal with complex specular objects.

center view

estimated disparity

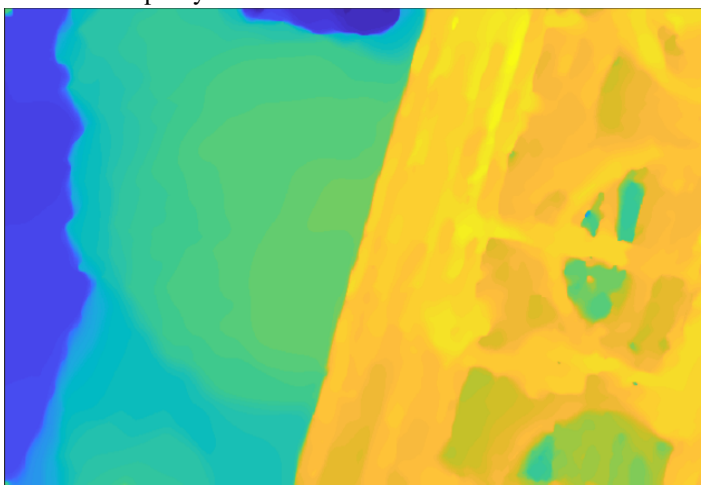

center view

albedo

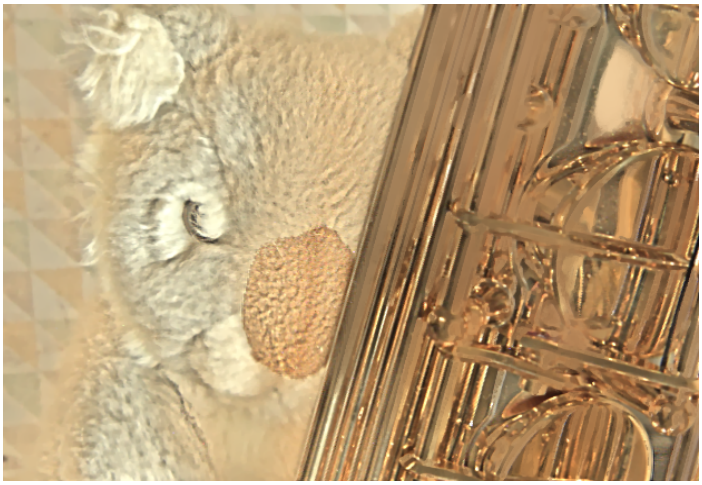

specularity

shading

direct shading

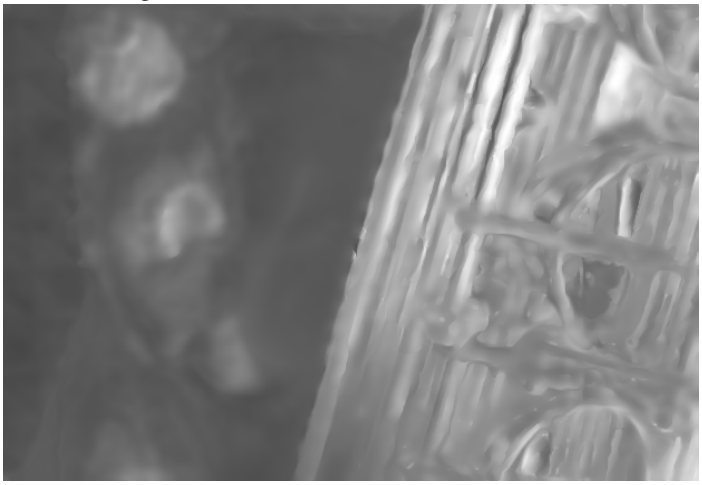

indirect shading

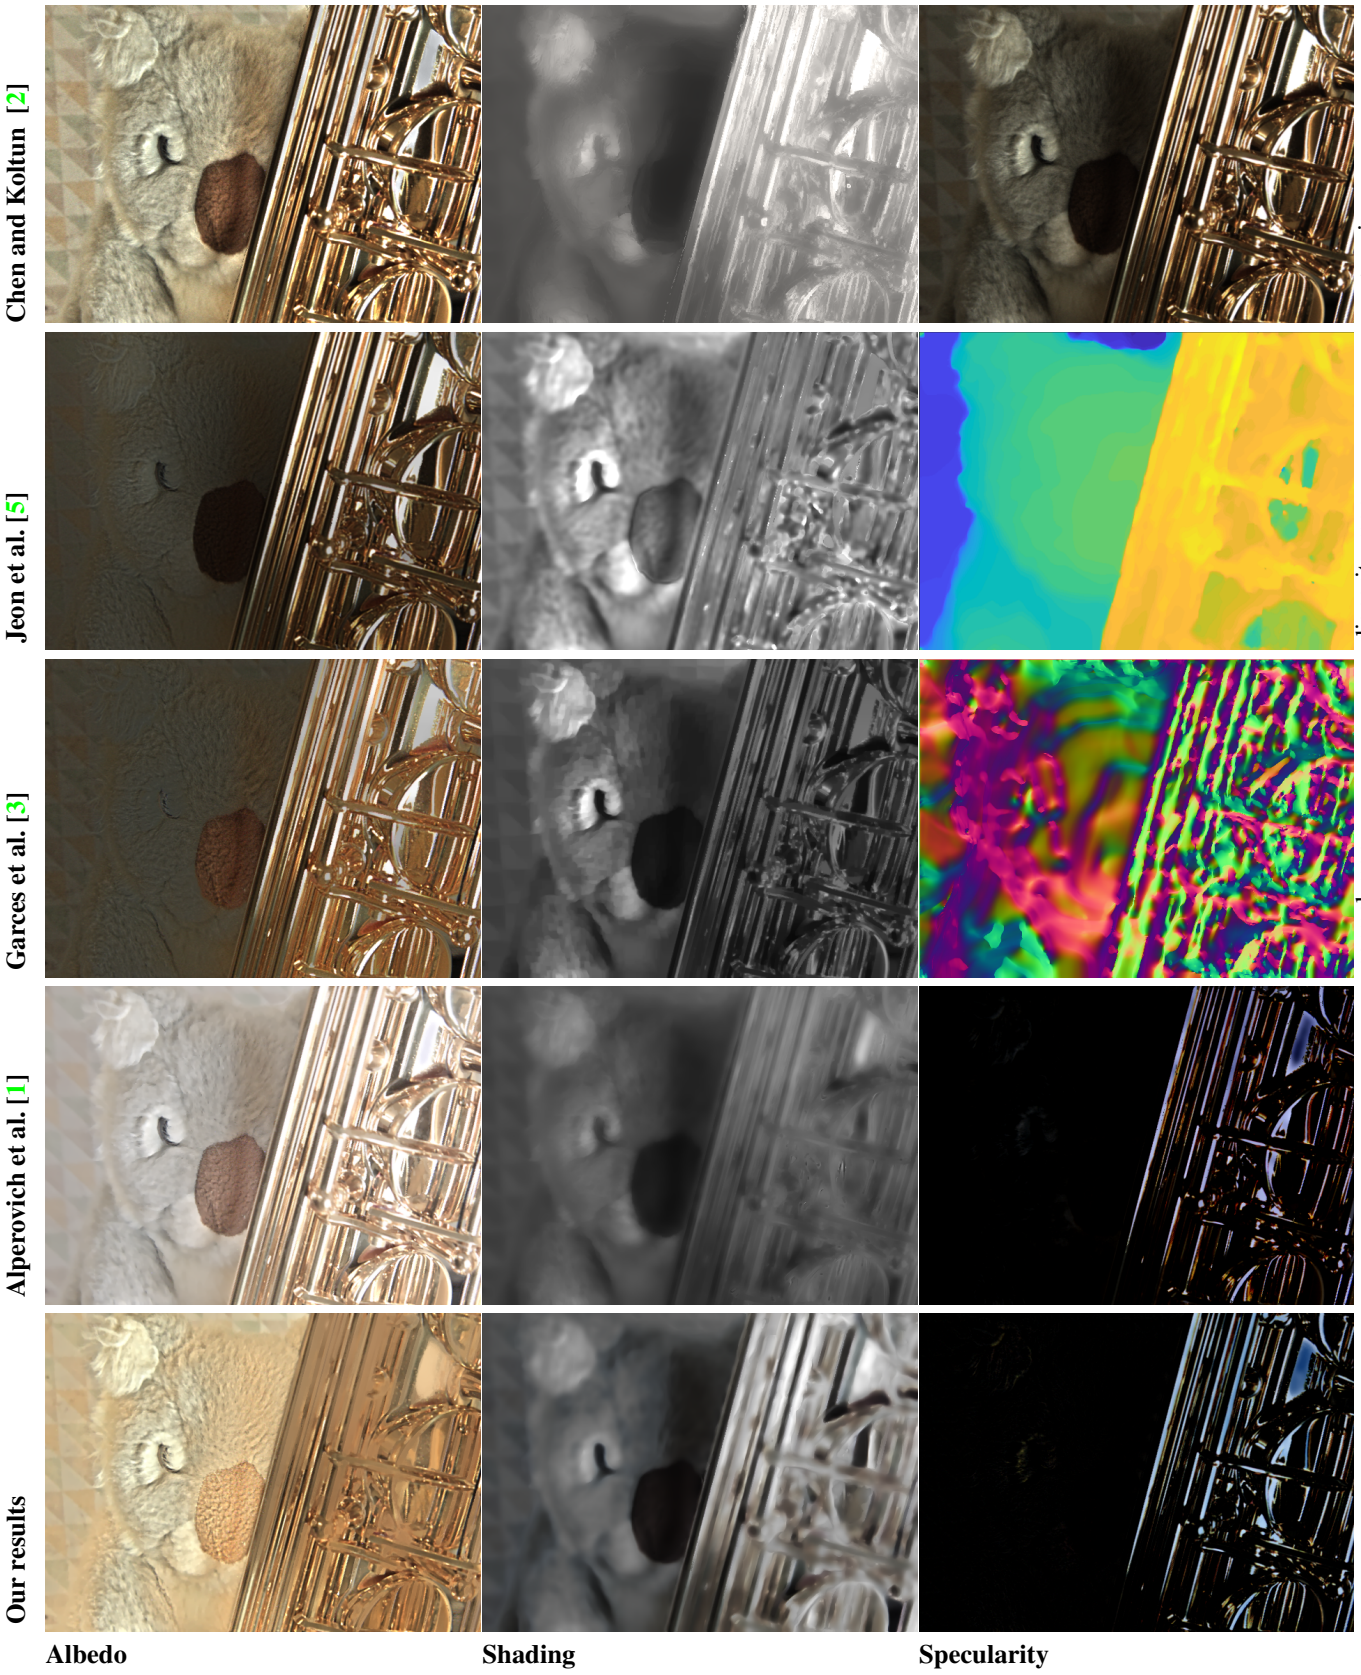
Chen and Koltun [2]
Jeon et al. [5]
Garces et al. [3]
Alperovich et al. [1]
Our results
center view
disparity map
normal map
Albedo
Shading
Specularity

specular mask, h

estimated diffuse chromaticity, χ

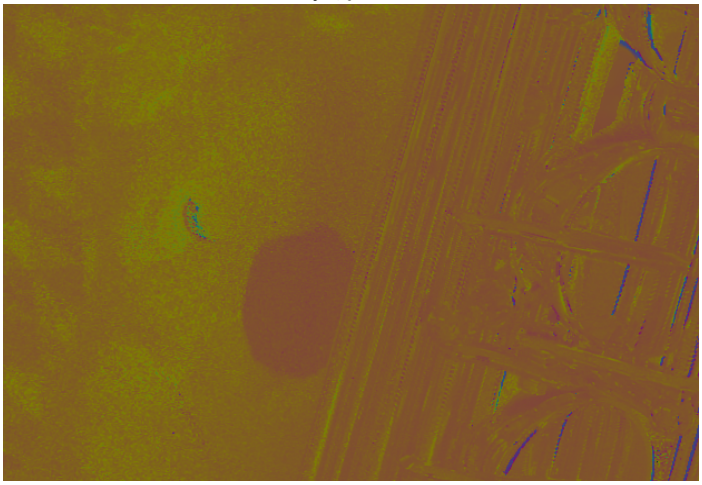

point-wise shadow confidence, β_{pw}

resulting shadow score, β

shadow boundaries score, δ_β

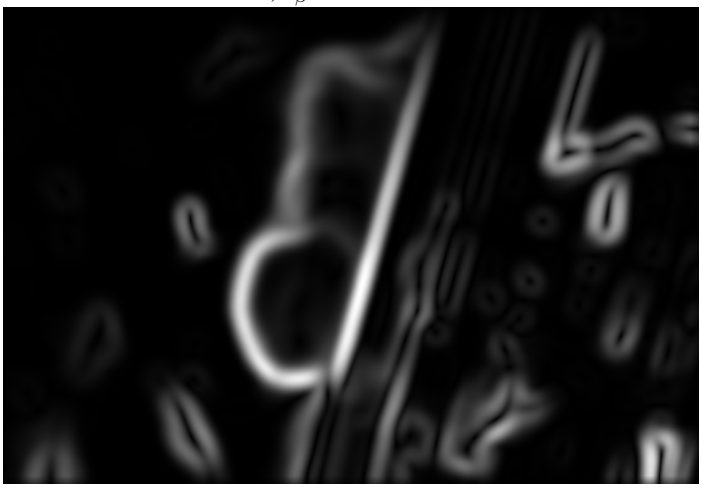

- Input light filed
- **Our results**
  - Albedo
  - Shading
  - Specularity
  - Direct shading
  - Indirect shading
  - Shadow score
- **Results by Alperovich and Goldluecke [1]**
  - Albedo
  - Shading
  - Specularity
- **Results by Garces et al. [3]**
  - Albedo
  - Shading

Light field *Origami*

Type:	Real world data set captured with the Lytro Illum plenoptic camera
Size:	9×9 views
Resolution:	434×625 pixels
Disparity range:	$[-1.5 \;\; 0.8]$

Scene description: We present a real world data set with two objects, an origami cat and owl. Both objects cast shadows on the floor, the owl has a strong highlight. With this example, we want to evaluate the algorithms on a data set with cast shadows.

center view

estimated disparity

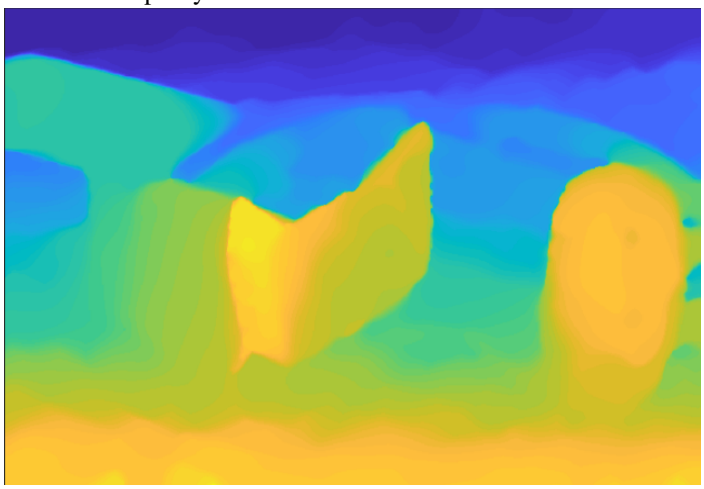

center view

albedo

specularity

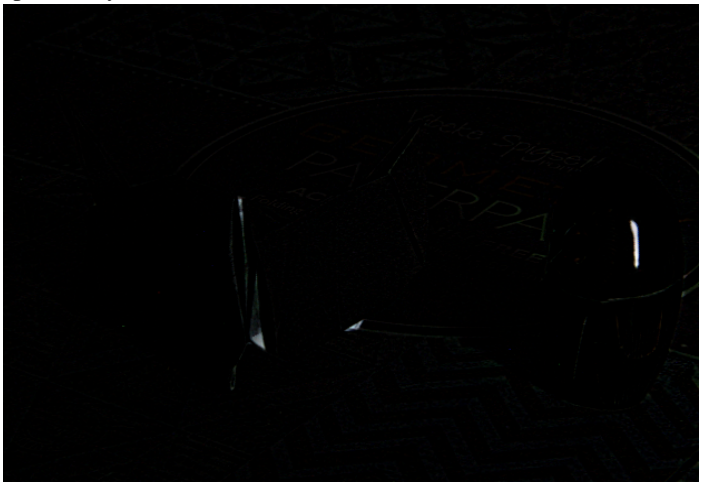

shading

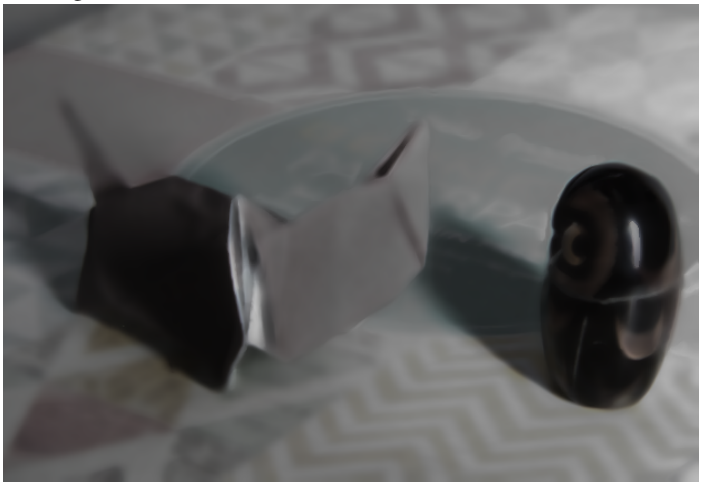

direct shading

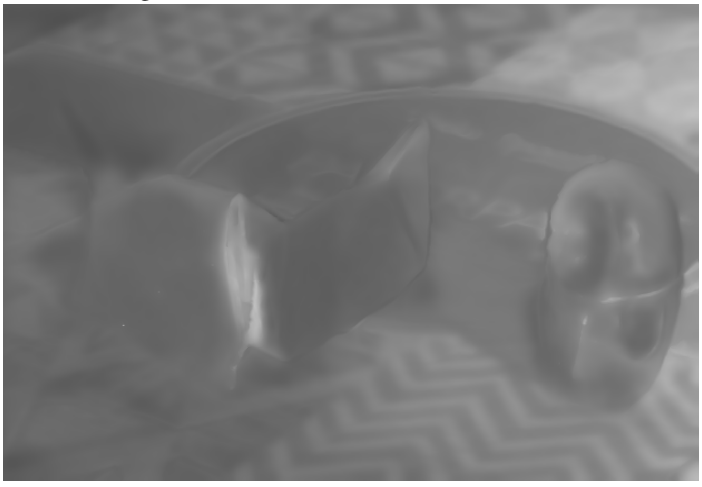

indirect shading

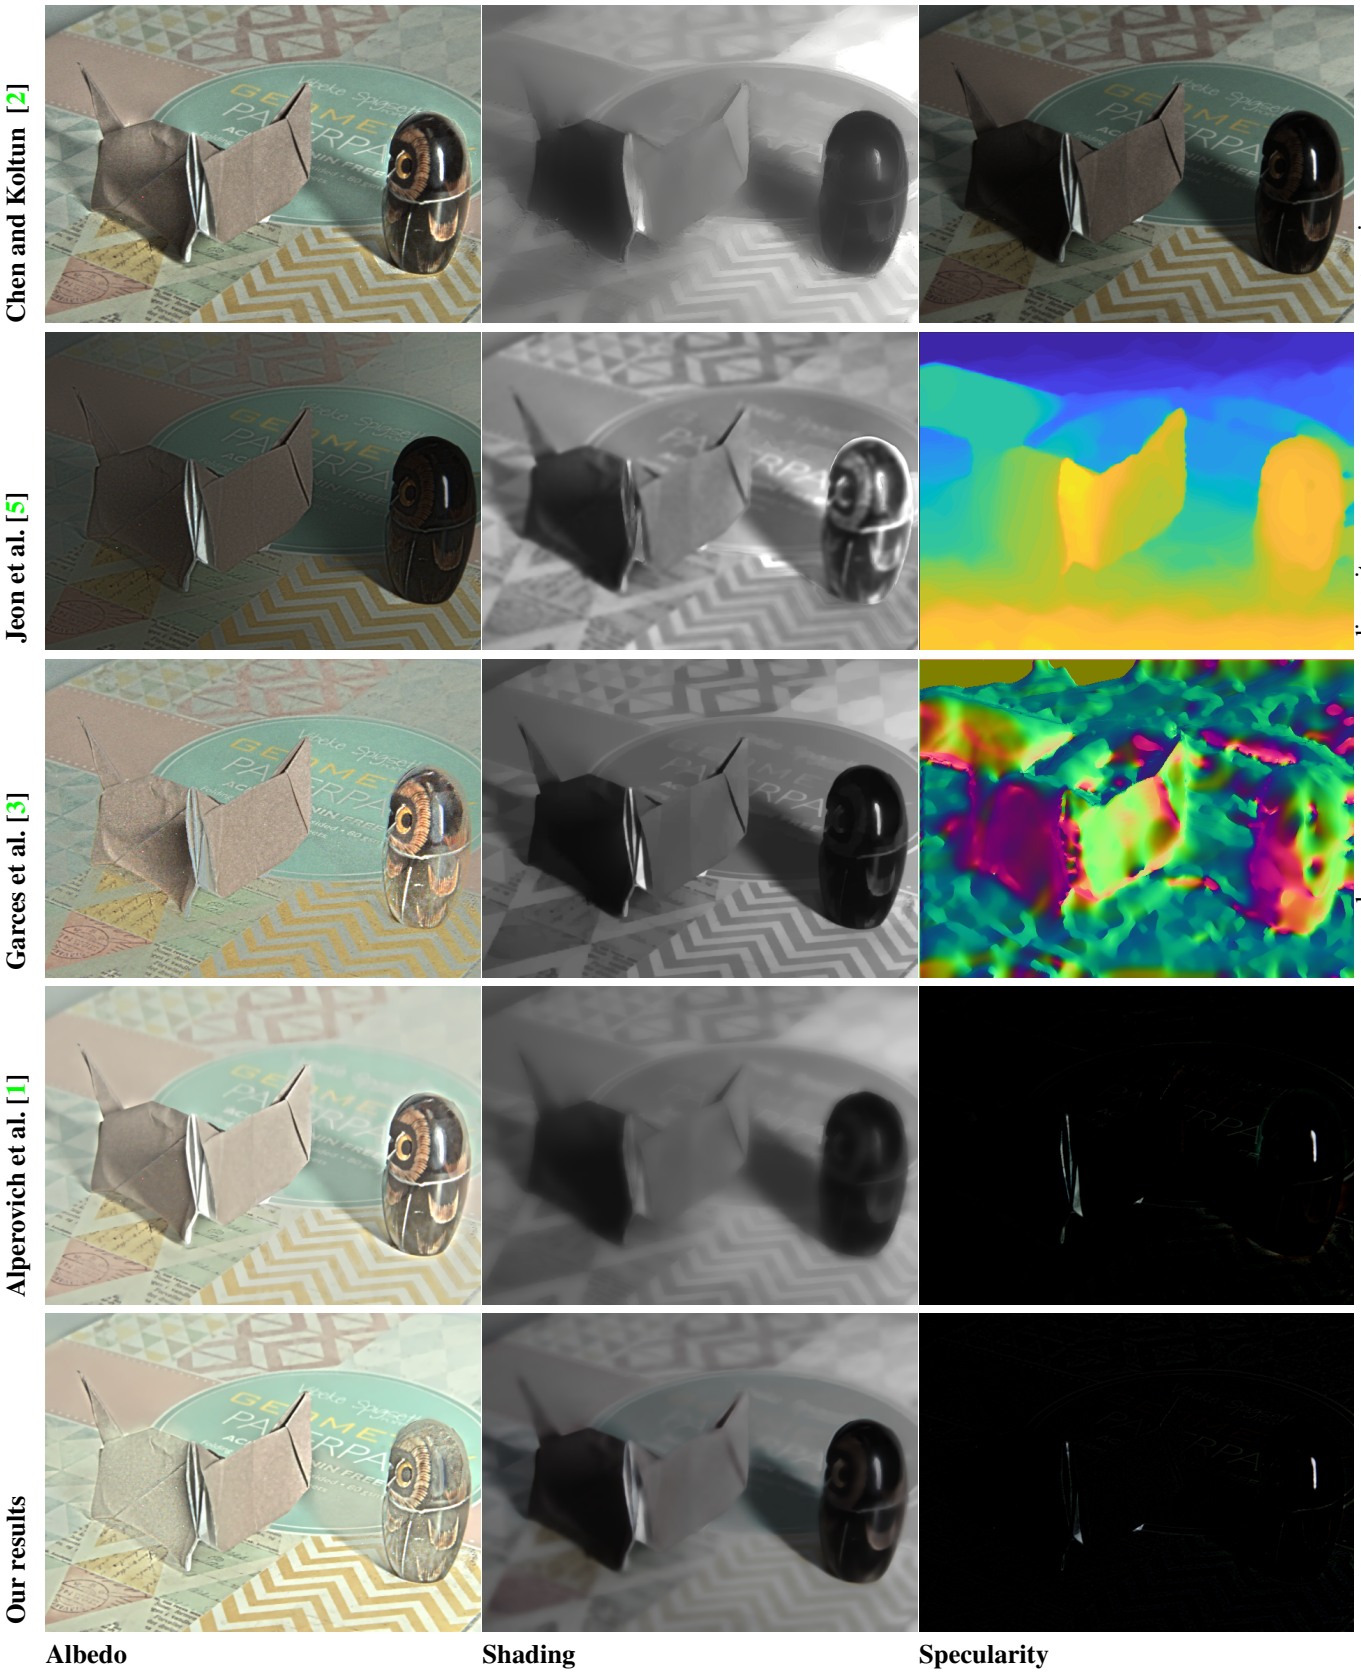
Chen and Koltun [2]
Jeon et al. [5]
Garces et al. [3]
Alperovich et al. [1]
Our results
center view
disparity map
normal map
Albedo
Shading
Specularity

specular mask, h

estimated diffuse chromaticity, χ

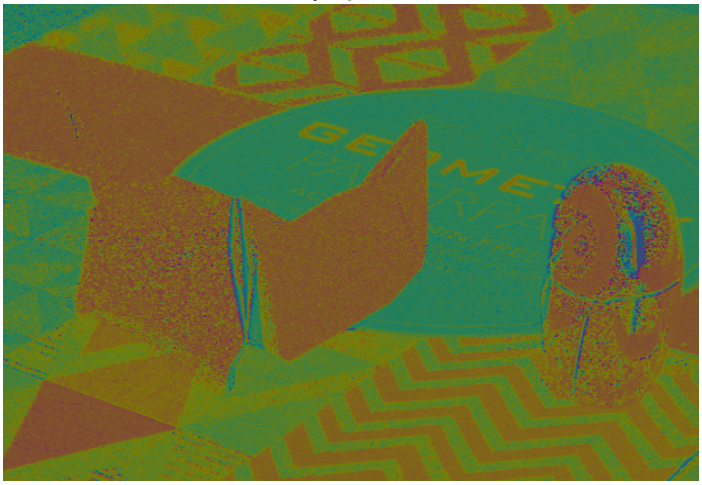

point-wise shadow confidence, β_{pw}

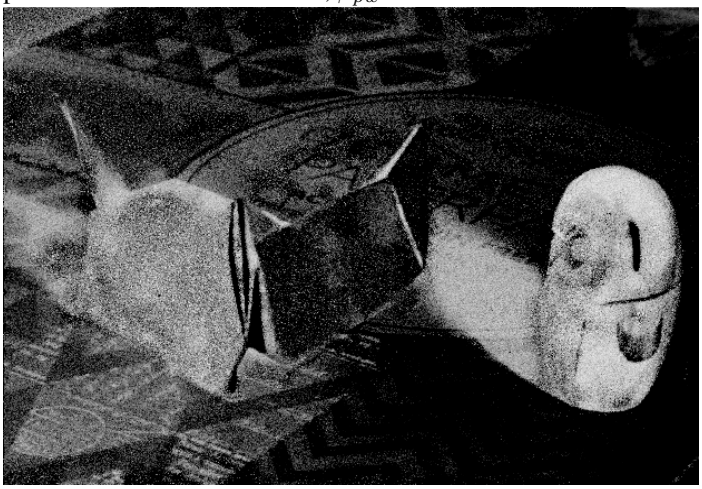

resulting shadow score, β

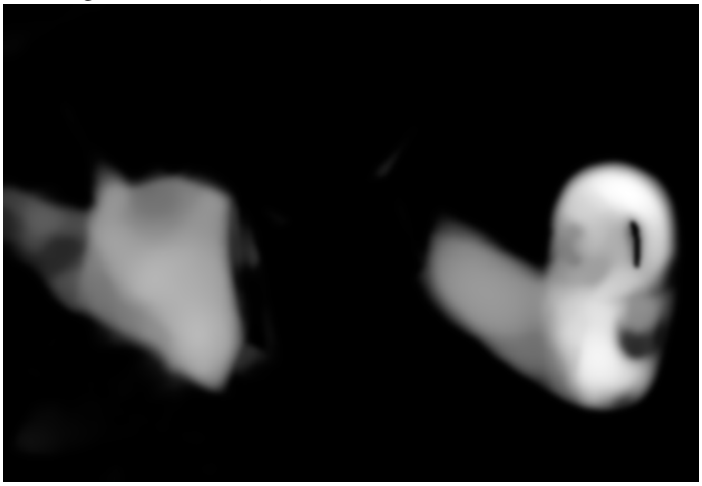

shadow boundaries score, δ_β

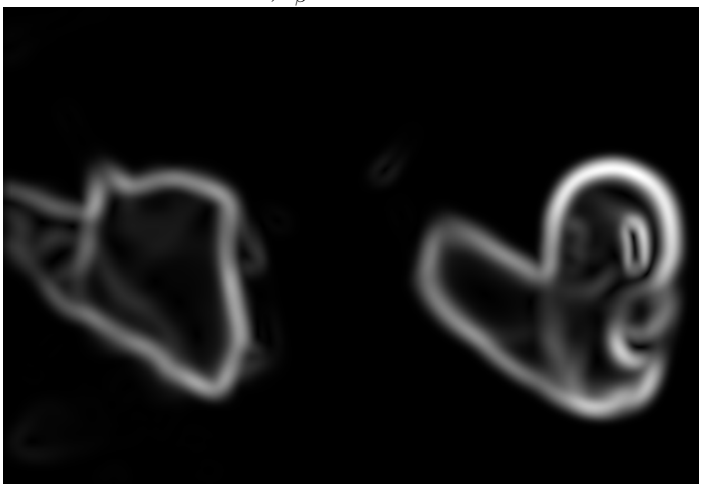

- Input light filed
- **Our results**
  - Albedo
  - Shading
  - Specularity
  - Direct shading
  - Indirect shading
  - Shadow score
- **Results by Alperovich and Goldluecke [1]**
  - Albedo
  - Shading
  - Specularity
- **Results by Garces et al. [3]**
  - Albedo
  - Shading

Light field *Koala1*

Type:	Real world data set captured with the Lytro Illum plenoptic camera
Size:	9×9 views
Resolution:	434×625 pixels
Disparity range:	$[-1.5 \;\; 1.4]$

Scene description: We present one more scene with koala and saxophone. This scene contains more shadows and less specularity. Also, cast shadows have stronger boundaries.

center view

estimated disparity

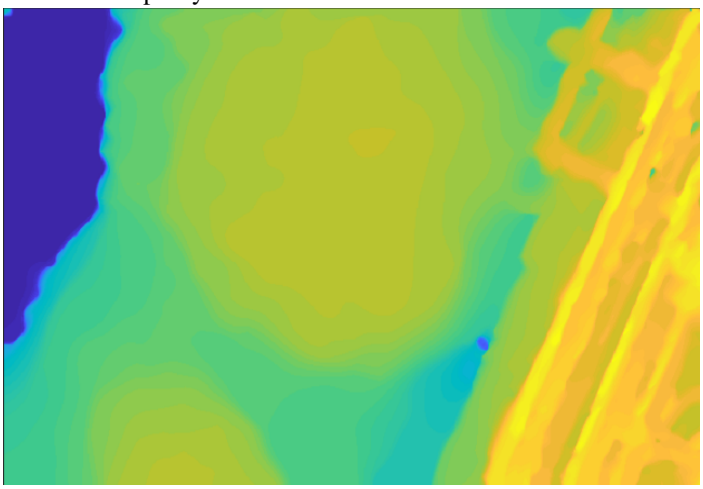

center view

albedo

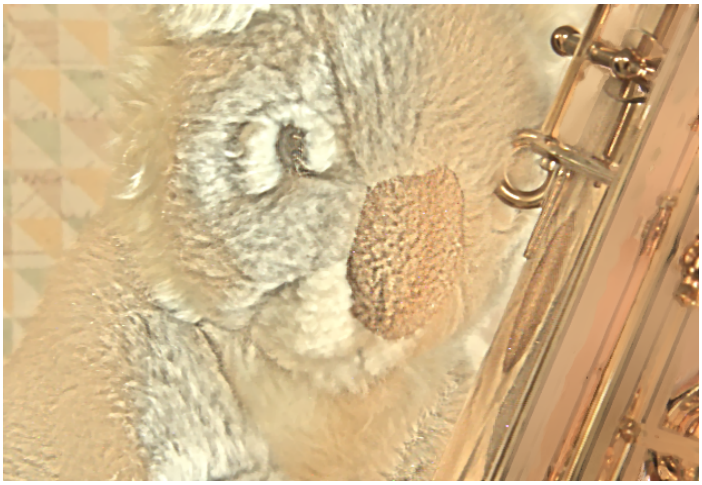

specularity

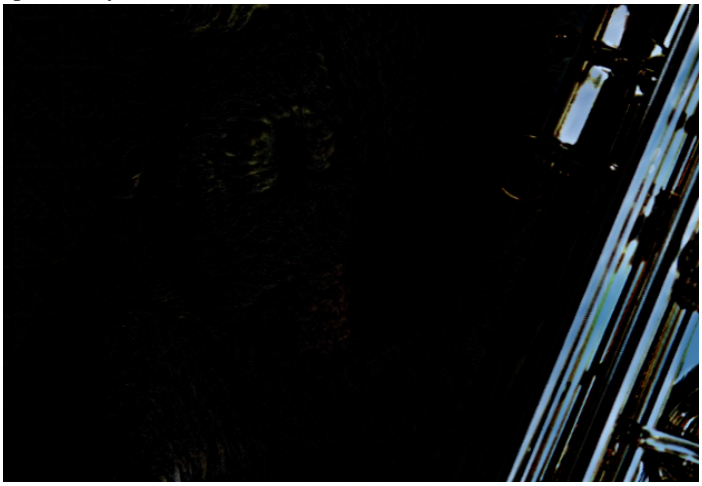

shading

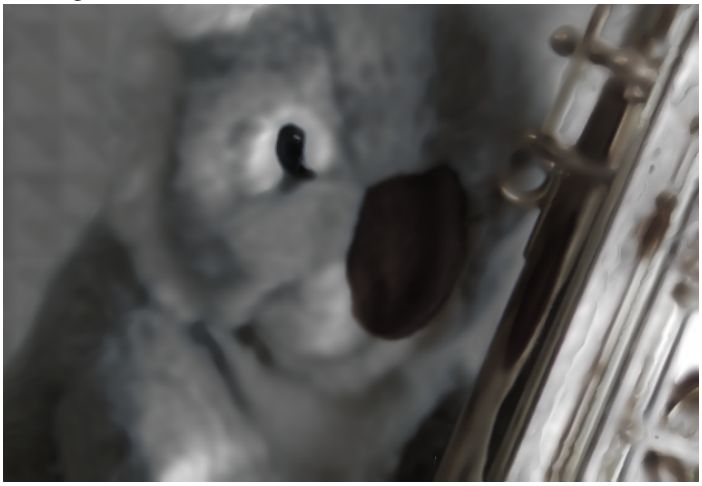

direct shading

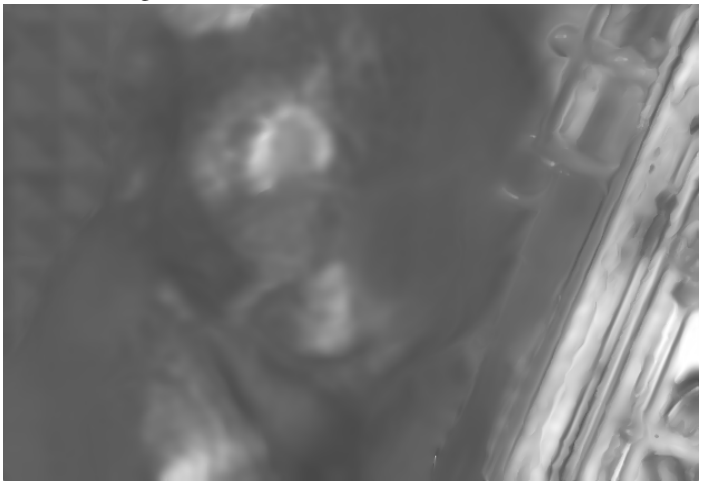

indirect shading

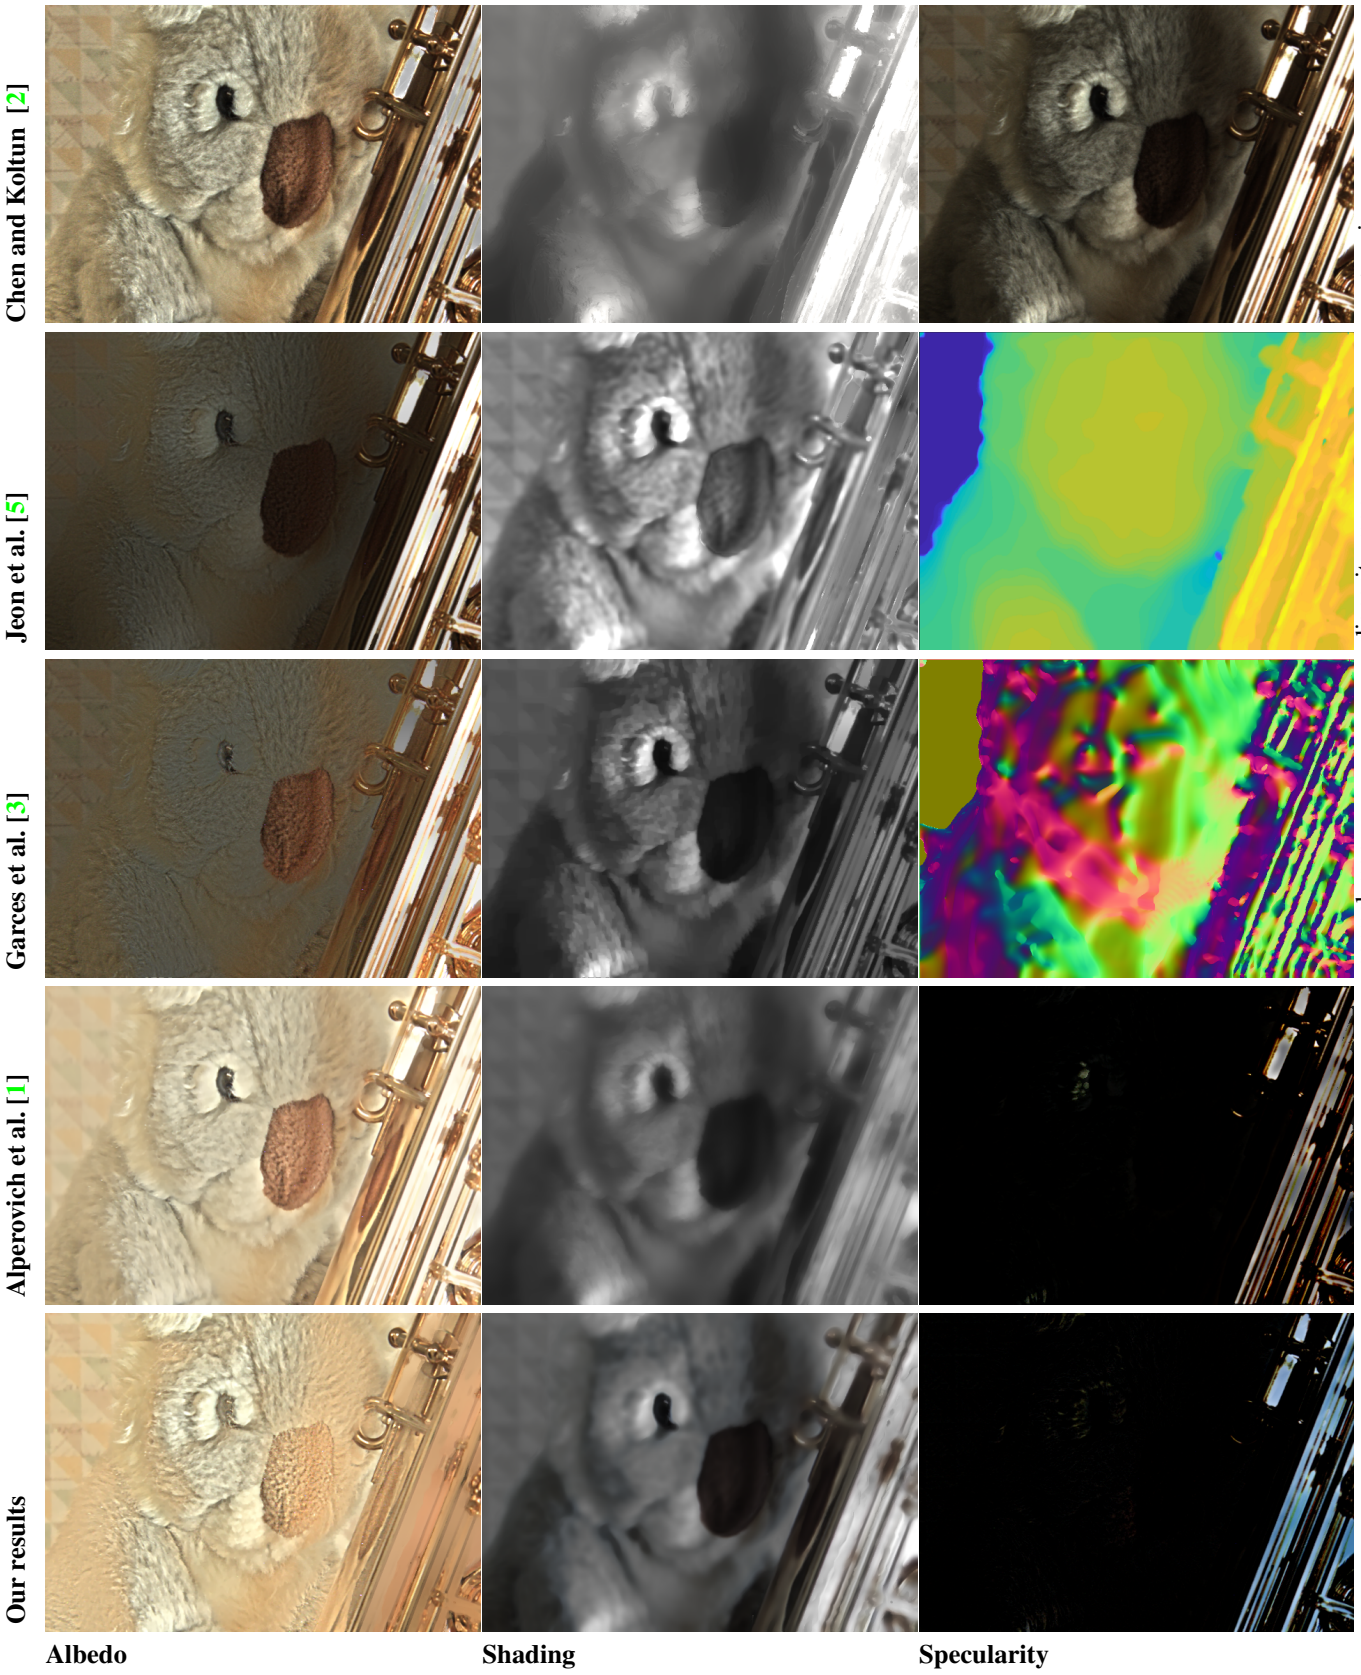
Chen and Koltun [2]
Jeon et al. [5]
Garces et al. [3]
Alperovich et al. [1]
Our results
center view
disparity map
normal map
Albedo
Shading
Specularity

specular mask, h

estimated diffuse chromaticity, χ

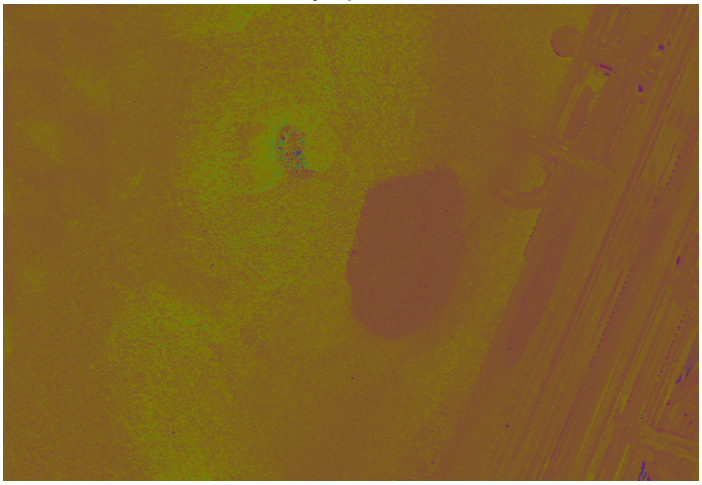

point-wise shadow confidence, β_{pw}

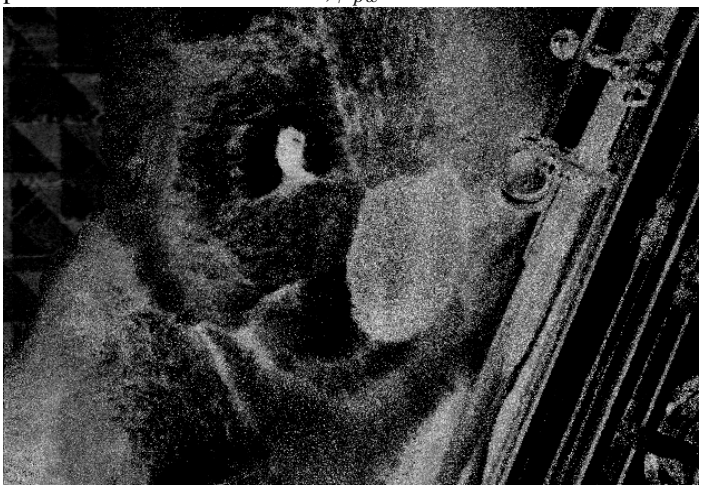

resulting shadow score, β

shadow boundaries score, δ_β

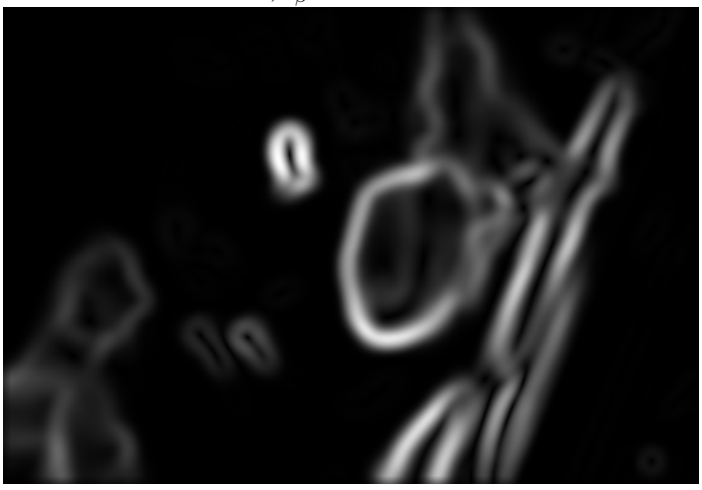

- Input light filed
- **Our results**
  - Albedo
  - Shading
  - Specularity
  - Direct shading
  - Indirect shading
  - Shadow score
- **Results by Alperovich and Goldluecke [1]**
  - Albedo
  - Shading
  - Specularity
- **Results by Garces et al. [3]**
  - Albedo
  - Shading

Light field *Antonius diffuse*

Type:	Synthetic data set generated with Blender using Cycles rendering engine
Size:	9×9 views
Resolution:	512×512 pixels
Disparity range:	$[-1.5 \;\; 1.3]$

Scene description: The following two datasets have identical scene geometry. However, the first scene contains an object with pure Lambertian reflection, the second an object with specular reflection. Here, we present the Lambertian one. With this example, we want to evaluate how algorithms perform on a scene without specularity. For our algorithm and [1], we assign zeros to the specularity confidence mask h, so that they in effect run without specularity component detection. Note that the proposed method is the only one which almost completely removes the cast shadow.

center view

estimated disparity

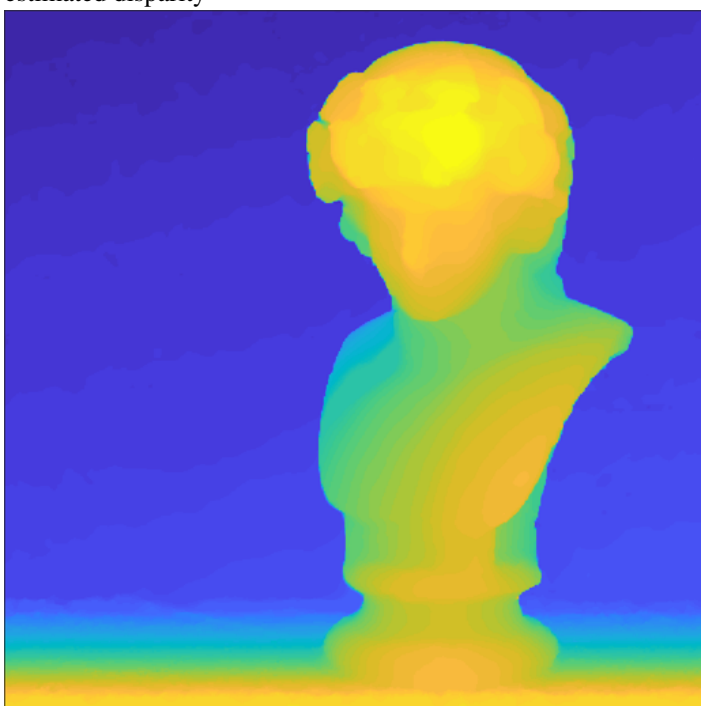

center view

albedo

specularity

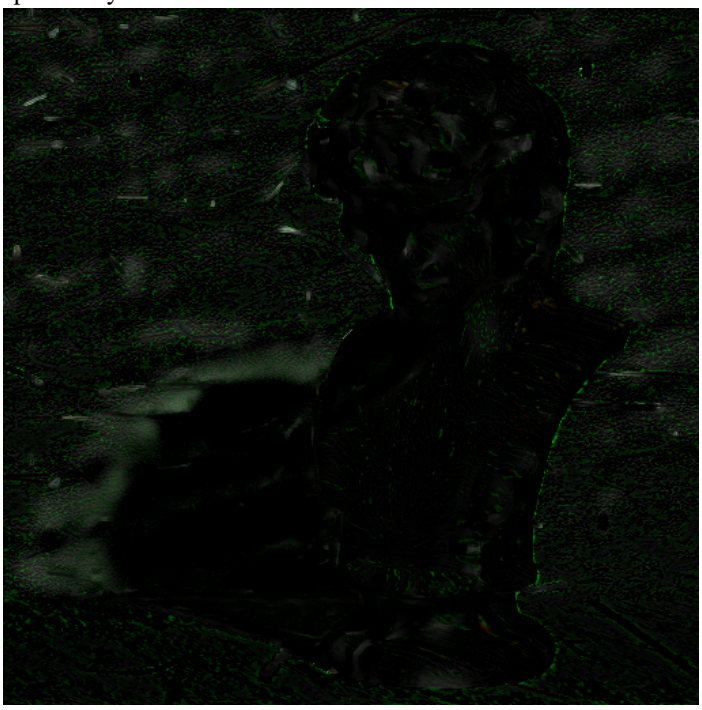

shading

direct shading

indirect shading

Ground truth

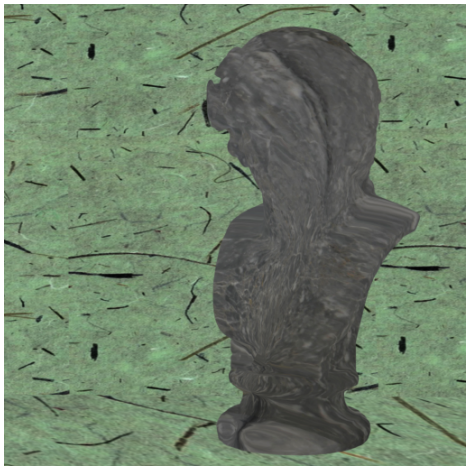

results by Alperovich and Goldluecke [1]


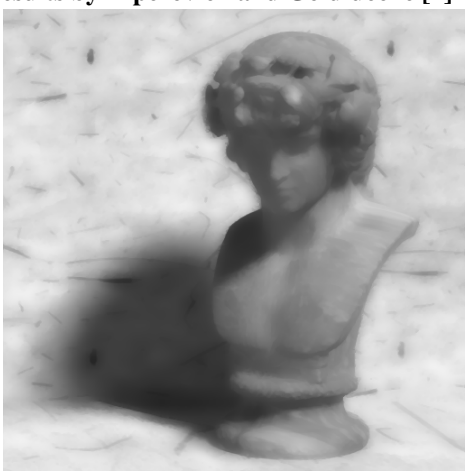
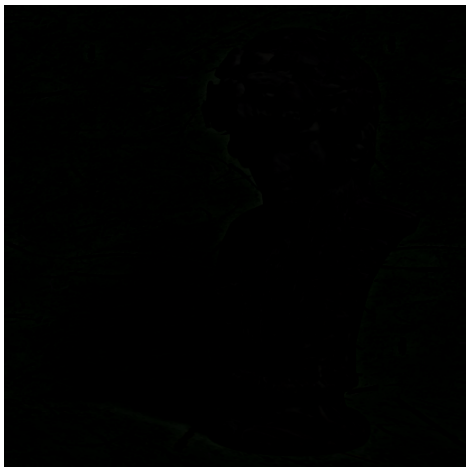

Our results



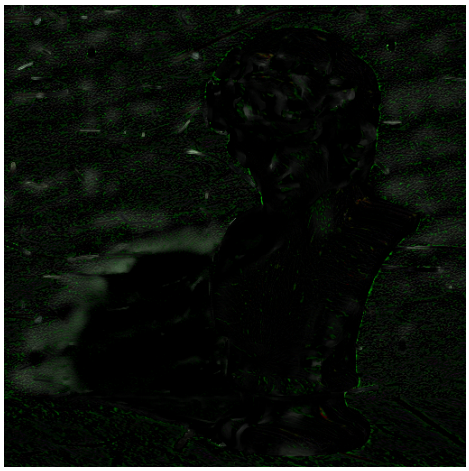

Albedo **Shading** **Specularity**

results by Garces et al.[3]

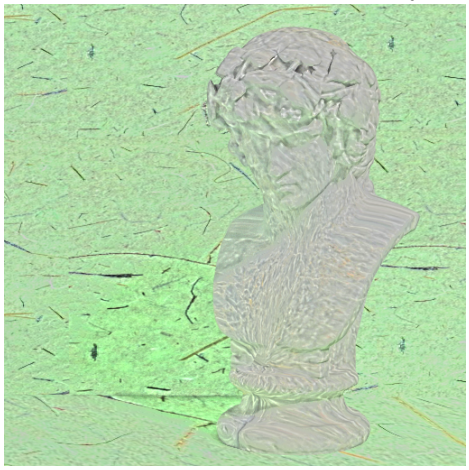

normal map

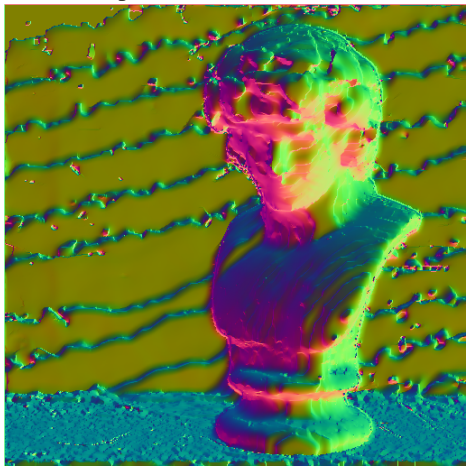

results by Chen and Koltun [2]

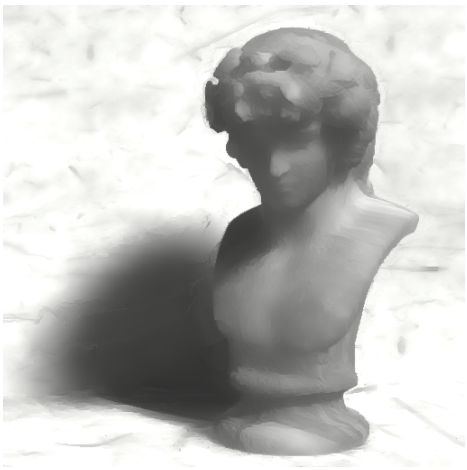

center view

results by Jeon et al.[5]

Albedo

Shading

disparity map

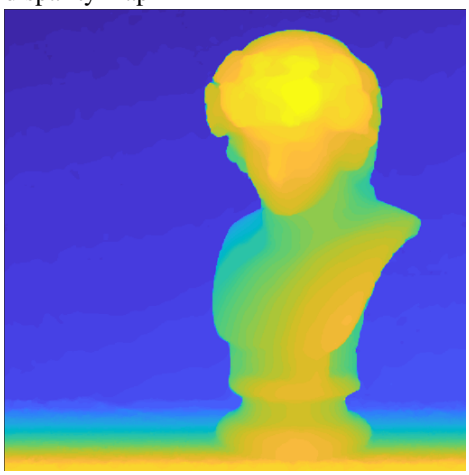

Table 1: Quantitative evaluation of the center view albedo, shading and specularity

	LMSE			GMSE		
	A	S	H	A	S	H
Chen and Koltun	0.0033	0.0262		0.0078	0.0682	
Jeon et al.	0.0053	0.0377		0.0085	0.0975	
Garces et al.	0.0041	**0.0208**		0.0082	0.0636	
Alperovich and Goldluecke	0.0038	0.0231	0	0.0082	0.0624	0
Our	**0.0032**	0.0213	**0**	**0.0076**	**0.0607**	**0**

specular confidence, h

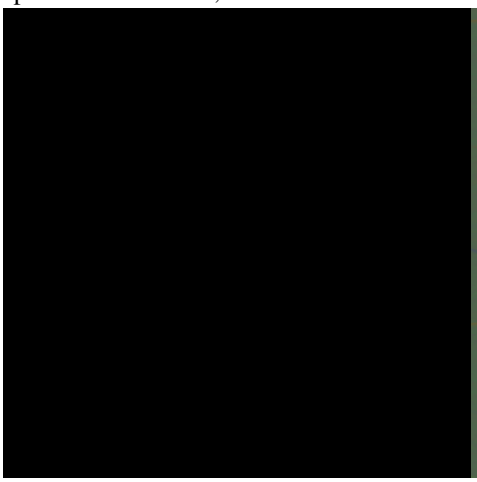

estimated diffuse chromaticity, χ

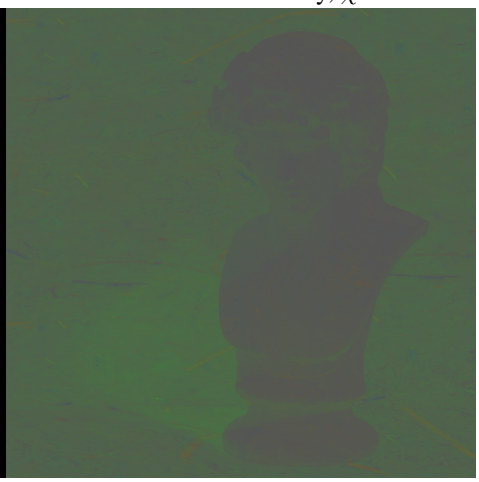

point-wise shadow confidence, β_{pw}

resulting shadow score, β

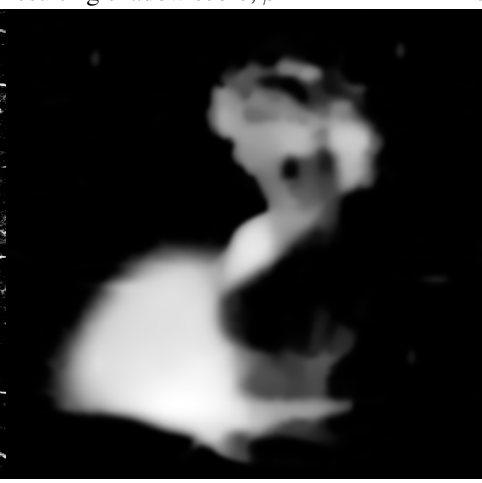

shadow boundaries score, δ_β

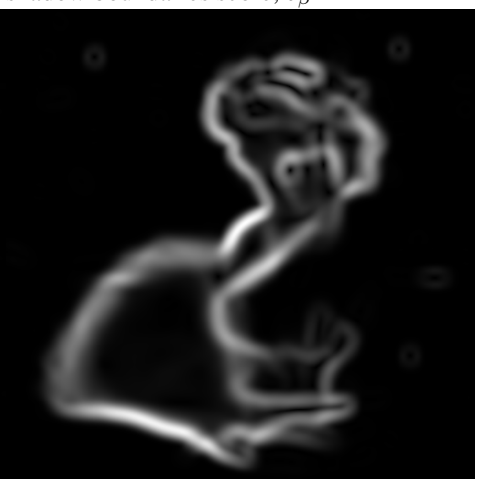

- Input light filed
- **Our results**
  - Albedo
  - Shading
  - Specularity
  - Direct shading
  - Indirect shading
  - Shadow score
- **Results by Alperovich and Goldluecke [1]**
  - Albedo
  - Shading
  - Specularity
- **Results by Garces et al. [3]**
  - Albedo
  - Shading

Light field *Antonius specular*

Type:	Synthetic data set generated with Blender using Cycles rendering engine
Size:	9×9 views
Resolution:	512×512 pixels
Disparity range:	$[-1.5 \;\; 1.3]$

Scene description: Geometrically, the same as the previous, but the object now has not only complicated geometry, but also and specular reflection. Again, we perform much better on removing the cast shadow.

center view

estimated disparity

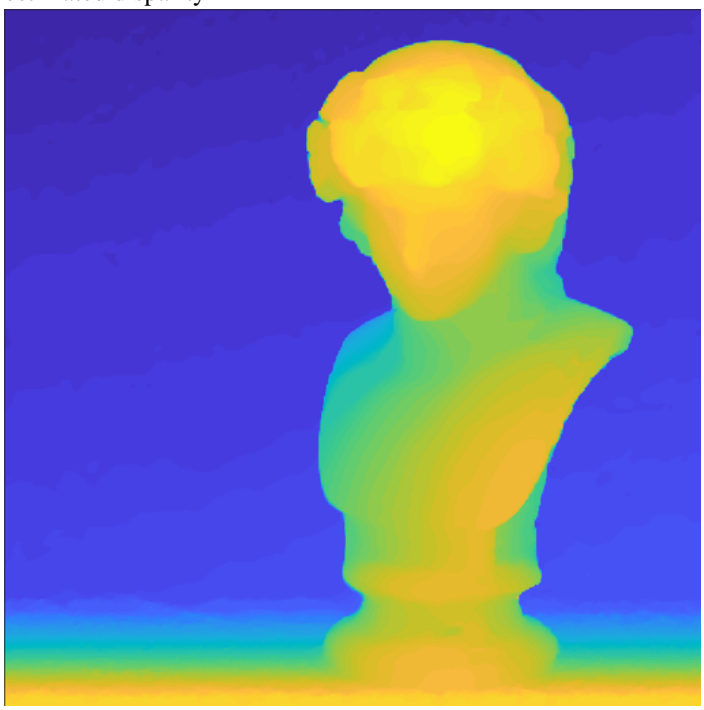

center view

albedo

specularity

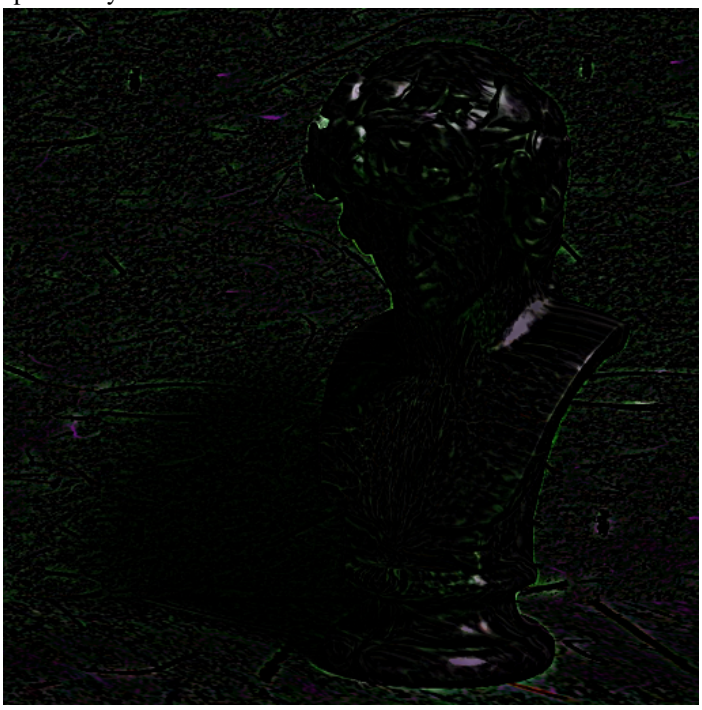

shading

direct shading

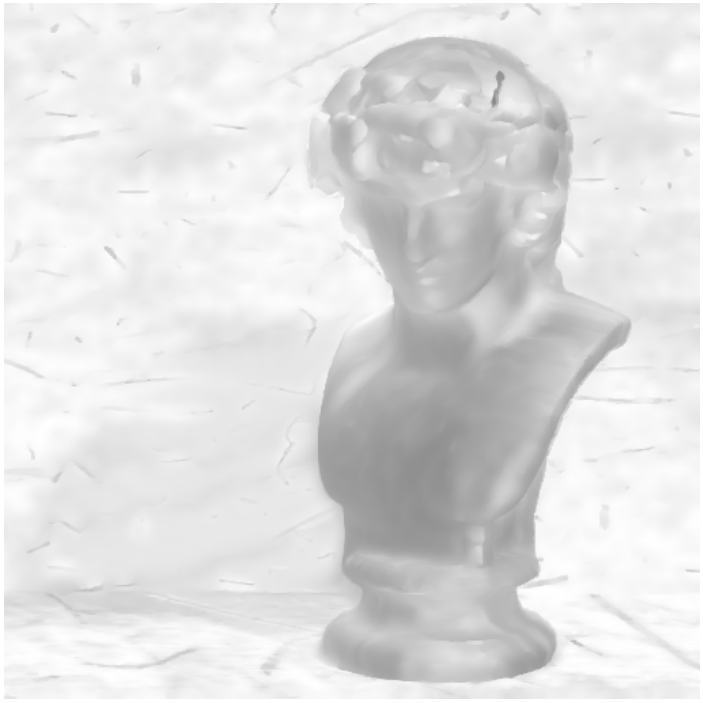

indirect shading

Ground truth

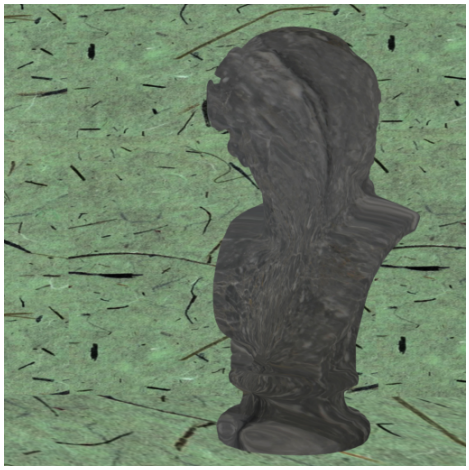

results by Alperovich and Goldluecke [1]


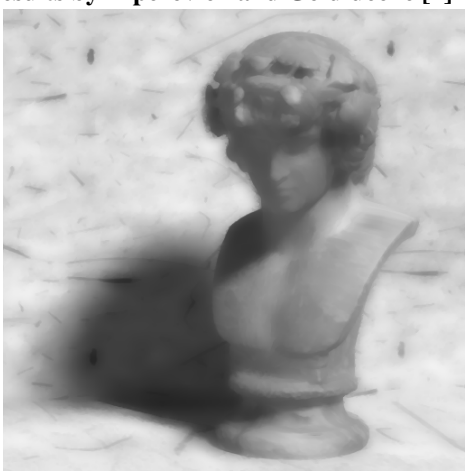


Our results



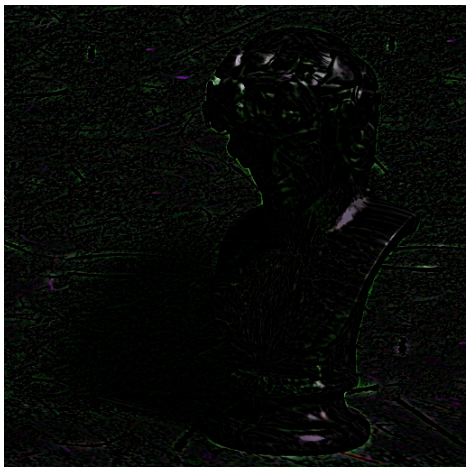

Albedo **Shading** **Specularity**

results by Garces et al.[3]

normal map

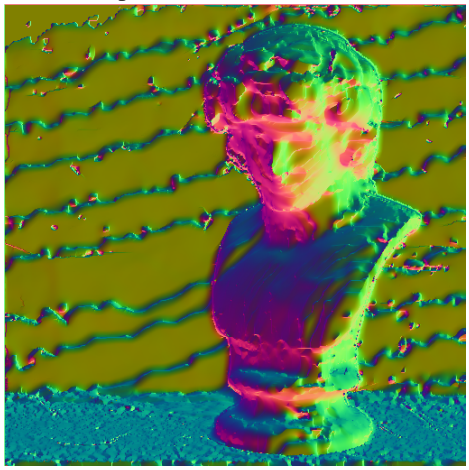

results by Chen and Koltun [2]

center view

results by Jeon et al.[5]

Albedo

Shading

disparity map

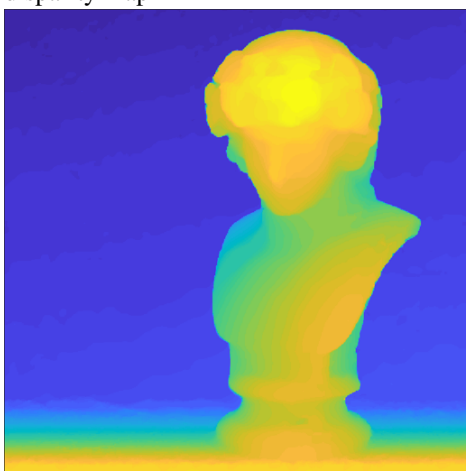

Table 2: Quantitative evaluation of the center view albedo, shading and specularity

	LMSE			GMSE		
	A	S	H	A	S	H
Chen and Koltun	**0.0041**	0.0247		0.0104	0.0582	
Jeon et al.	0.0057	0.0372		**0.0102**	0.0968	
Garces et al.	0.0049	0.0189		0.0107	0.0541	
Alperovich and Goldluecke	0.0045	0.0218	0.0008	0.0105	0.0538	**0.0018**
Our	0.0046	**0.0183**	**0.0008**	0.0108	**0.0507**	0.002

specular confidence, h

estimated diffuse chromaticity, χ

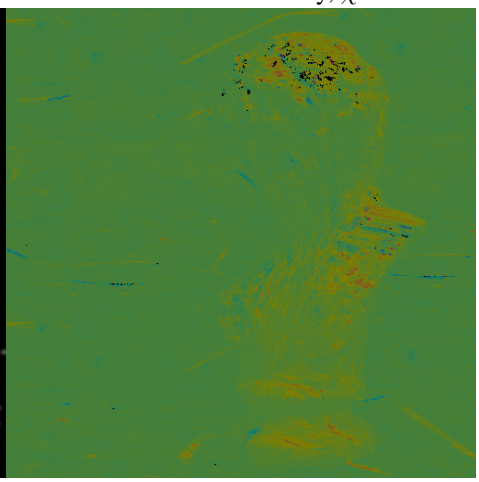

point-wise shadow confidence, β_{pw}

resulting shadow score, β

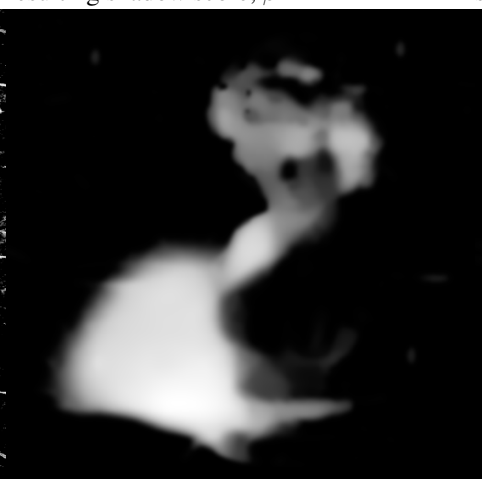

shadow boundaries score, δ_β

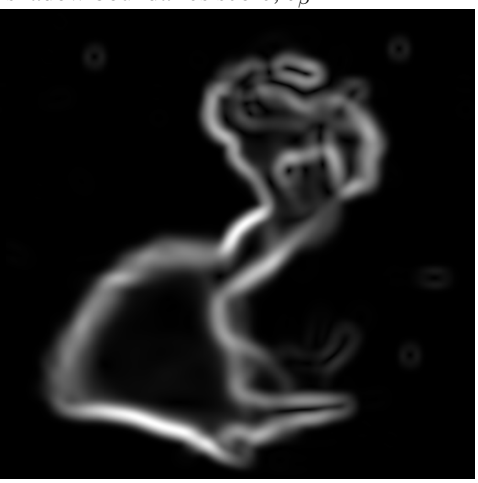

- Input light filed
- **Our results**
  - Albedo
  - Shading
  - Specularity
  - Direct shading
  - Indirect shading
  - Shadow score
- **Results by Alperovich and Goldluecke [1]**
  - Albedo
  - Shading
  - Specularity
- **Results by Garces et al. [3]**
  - Albedo
  - Shading

Light field *Monkey*

Type:	Synthetic data set generated with Blender using Cycles rendering engine
Size:	9×9 views
Resolution:	512×512 pixels
Disparity range:	$[-1.5 \;\; 1.5]$

Scene description: We present another data from the main paper, but with more information and results included. The scene contains two specular objects, one of them with complicated geometry. Both objects cast strong shadows on the floor. With this data set, we evaluate how the algorithms deal with complicated geometry, cast shadows, interreflections, and specularity.

center view

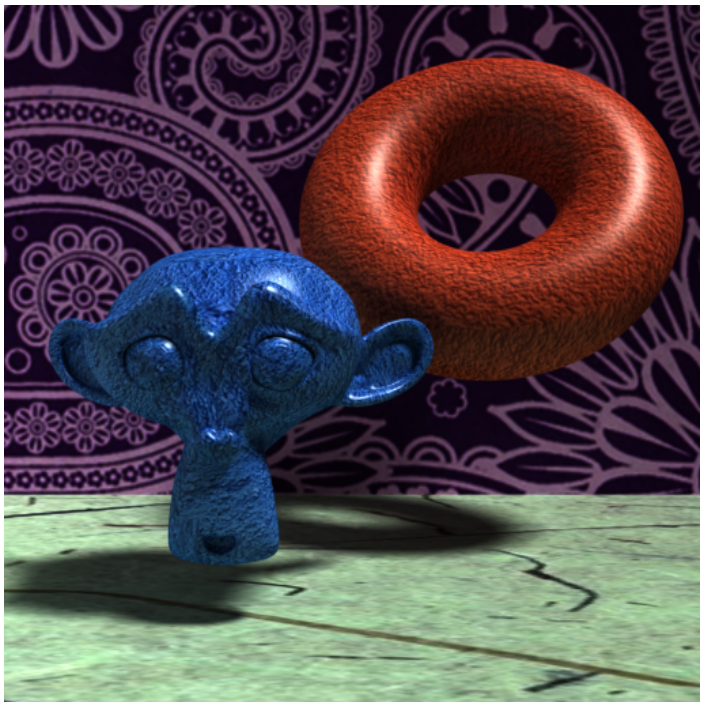

estimated disparity

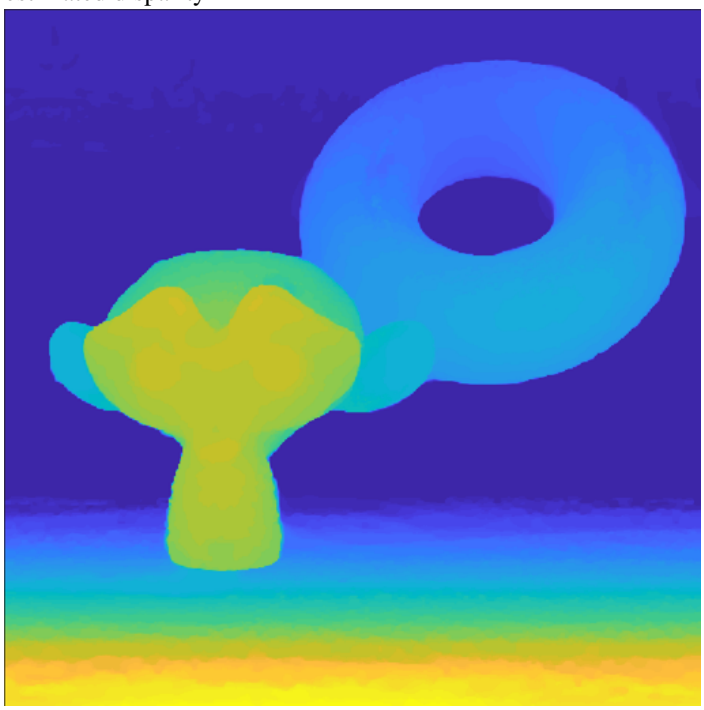

center view

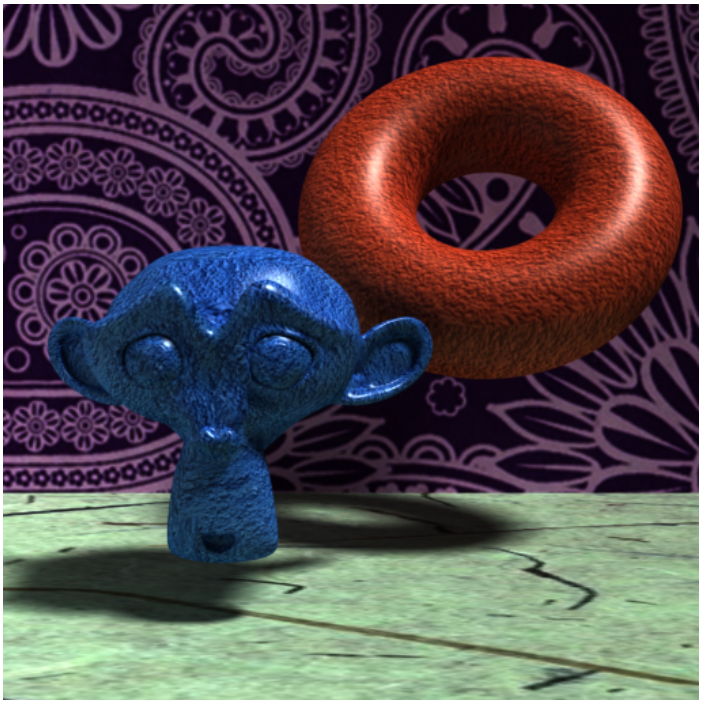

albedo

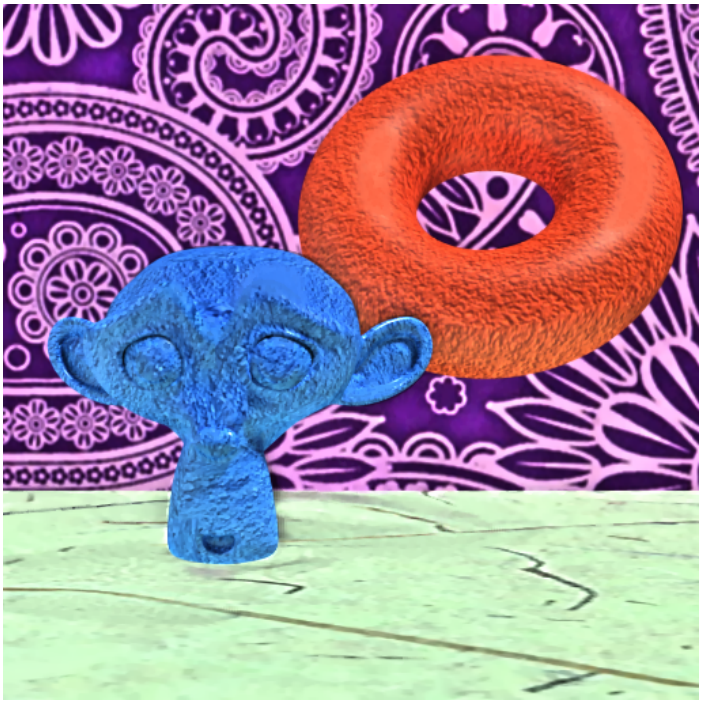

specularity

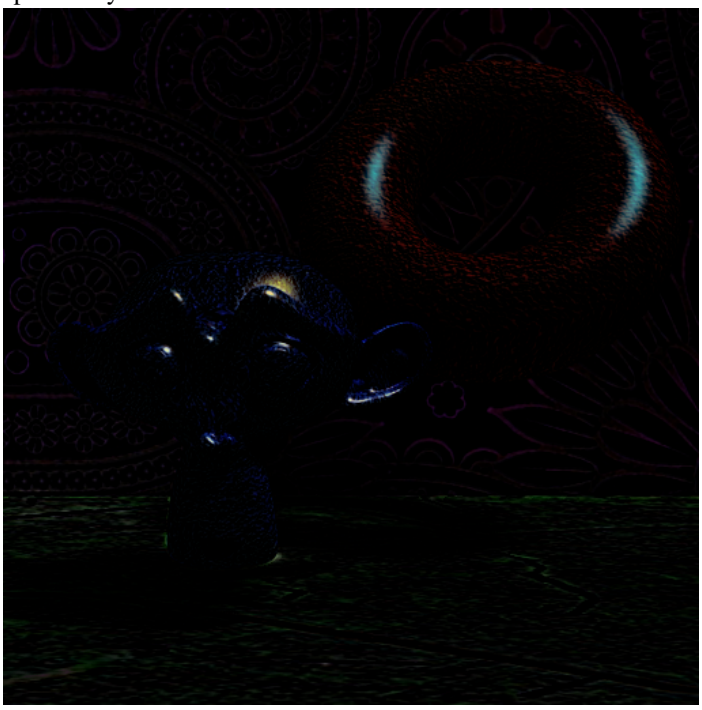

shading

direct shading

indirect shading

Ground truth

results by Alperovich and Goldluecke [1]

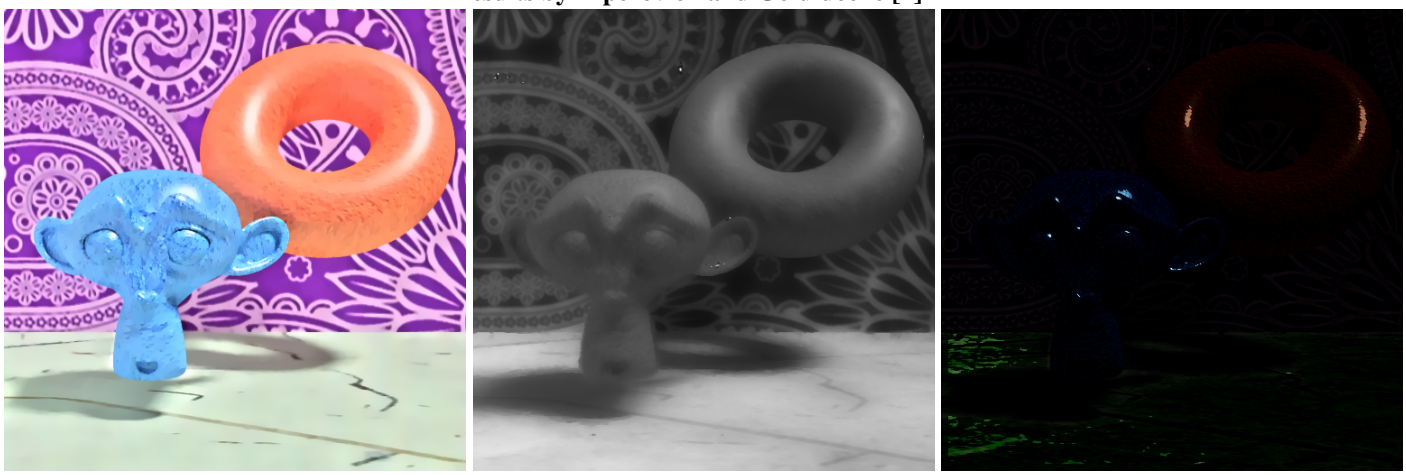

Our results

Albedo **Shading** **Specularity**

results by Garces et al.[3]

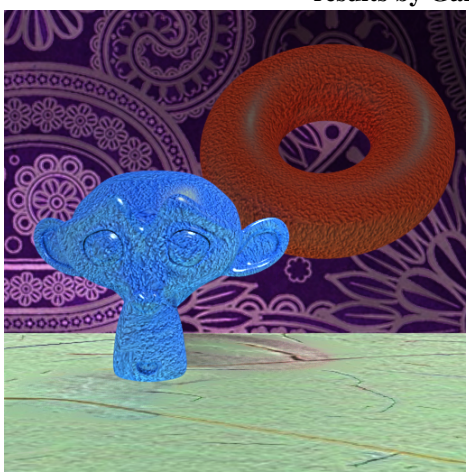

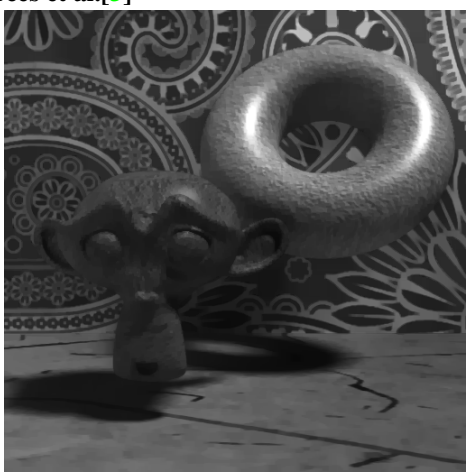

normal map

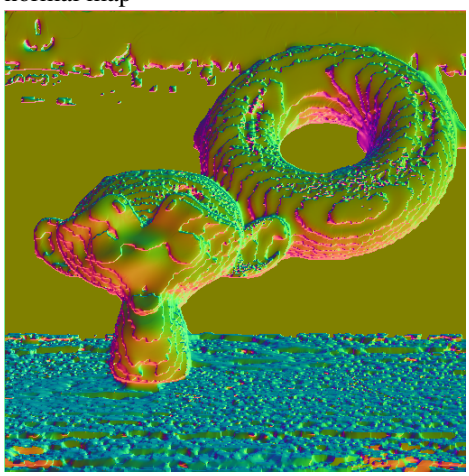

results by Chen and Koltun [2]

center view

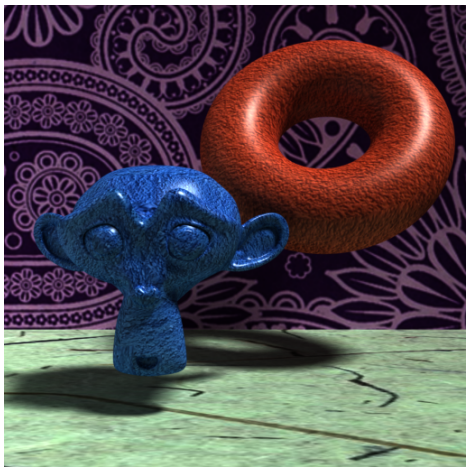

results by Jeon et al.[5]

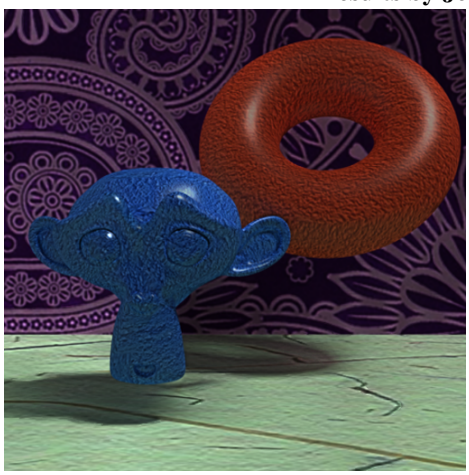

Albedo

Shading

disparity map

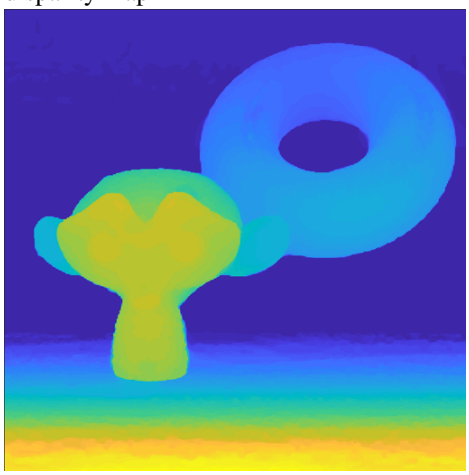

Table 3: Quantitative evaluation of the center view albedo, shading and specularity

	LMSE			GMSE		
	A	S	H	A	S	H
Chen and Koltun	0.0065	0.0152		0.0141	0.0325	
Jeon et al.	0.0189	0.0290		0.0272	0.0940	
Garces et al.	0.0076	0.0192		0.0154	0.0846	
Alperovich and Goldluecke	0.0074	0.0116	0.0015	0.0118	0.0262	0.0019
Our	**0.0022**	**0.0064**	**0.0006**	**0.0025**	**0.0110**	**0.0014**

specular confidence, h

estimated diffuse chromaticity, χ

point-wise shadow confidence, β_{pw}

resulting shadow score, β

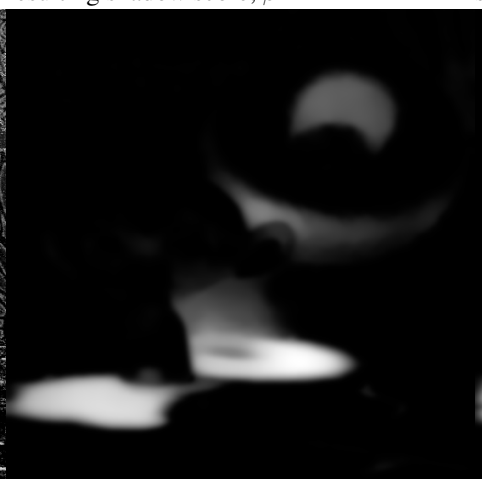

shadow boundaries score, δ_β

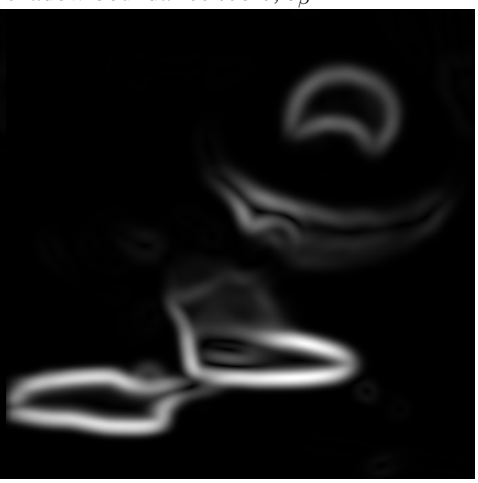

- Input light filed
- **Our results**
  - Albedo
  - Shading
  - Specularity
  - Direct shading
  - Indirect shading
  - Shadow score
- **Results by Alperovich and Goldluecke [1]**
  - Albedo
  - Shading
  - Specularity
- **Results by Garces et al. [3]**
  - Albedo
  - Shading

Light field *Chicken*

Type:	Real world data set captured with industrial camera mounted on gantry
Size:	9×9 views
Resolution:	497×710 pixels
Disparity range:	$[-1.5 \;\; 1.0]$

Scene description: We present a real world data set with non-Lambertian object, illuminated with approximately white light. This scene was captured with the industrial camera mounted on a gantry.

center view

estimated disparity

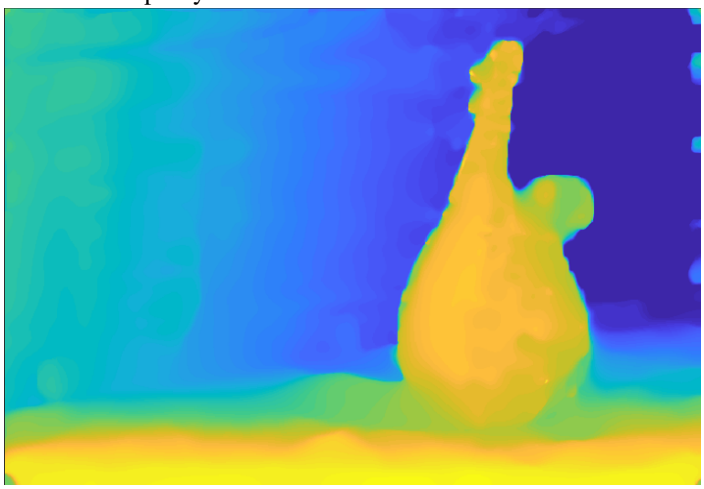

center view

albedo

specularity

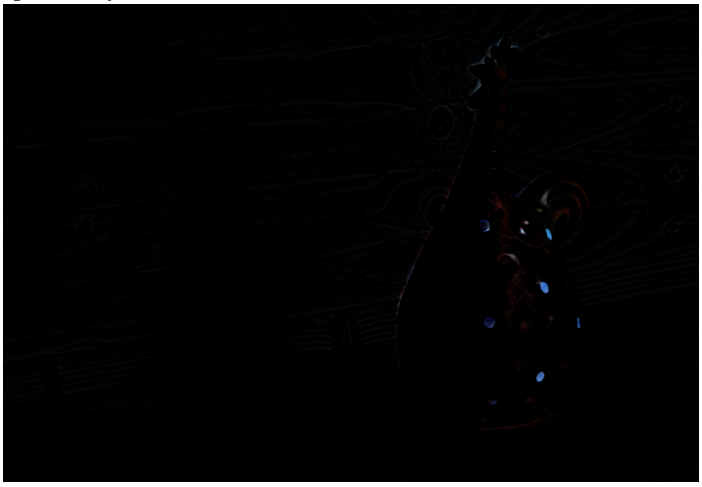

shading

direct shading

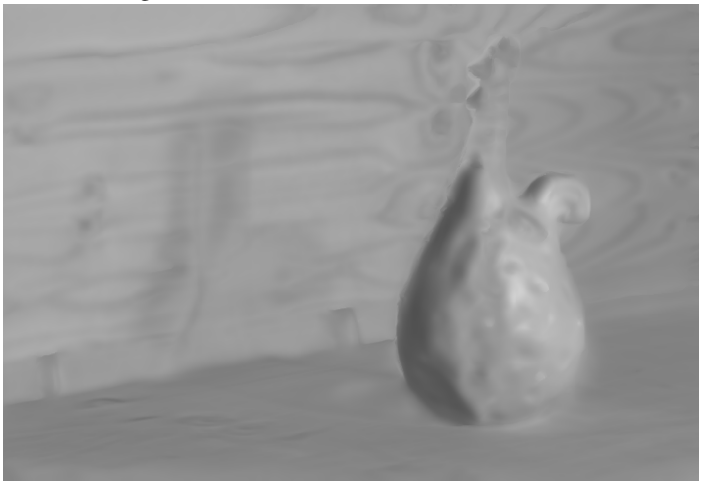

indirect shading

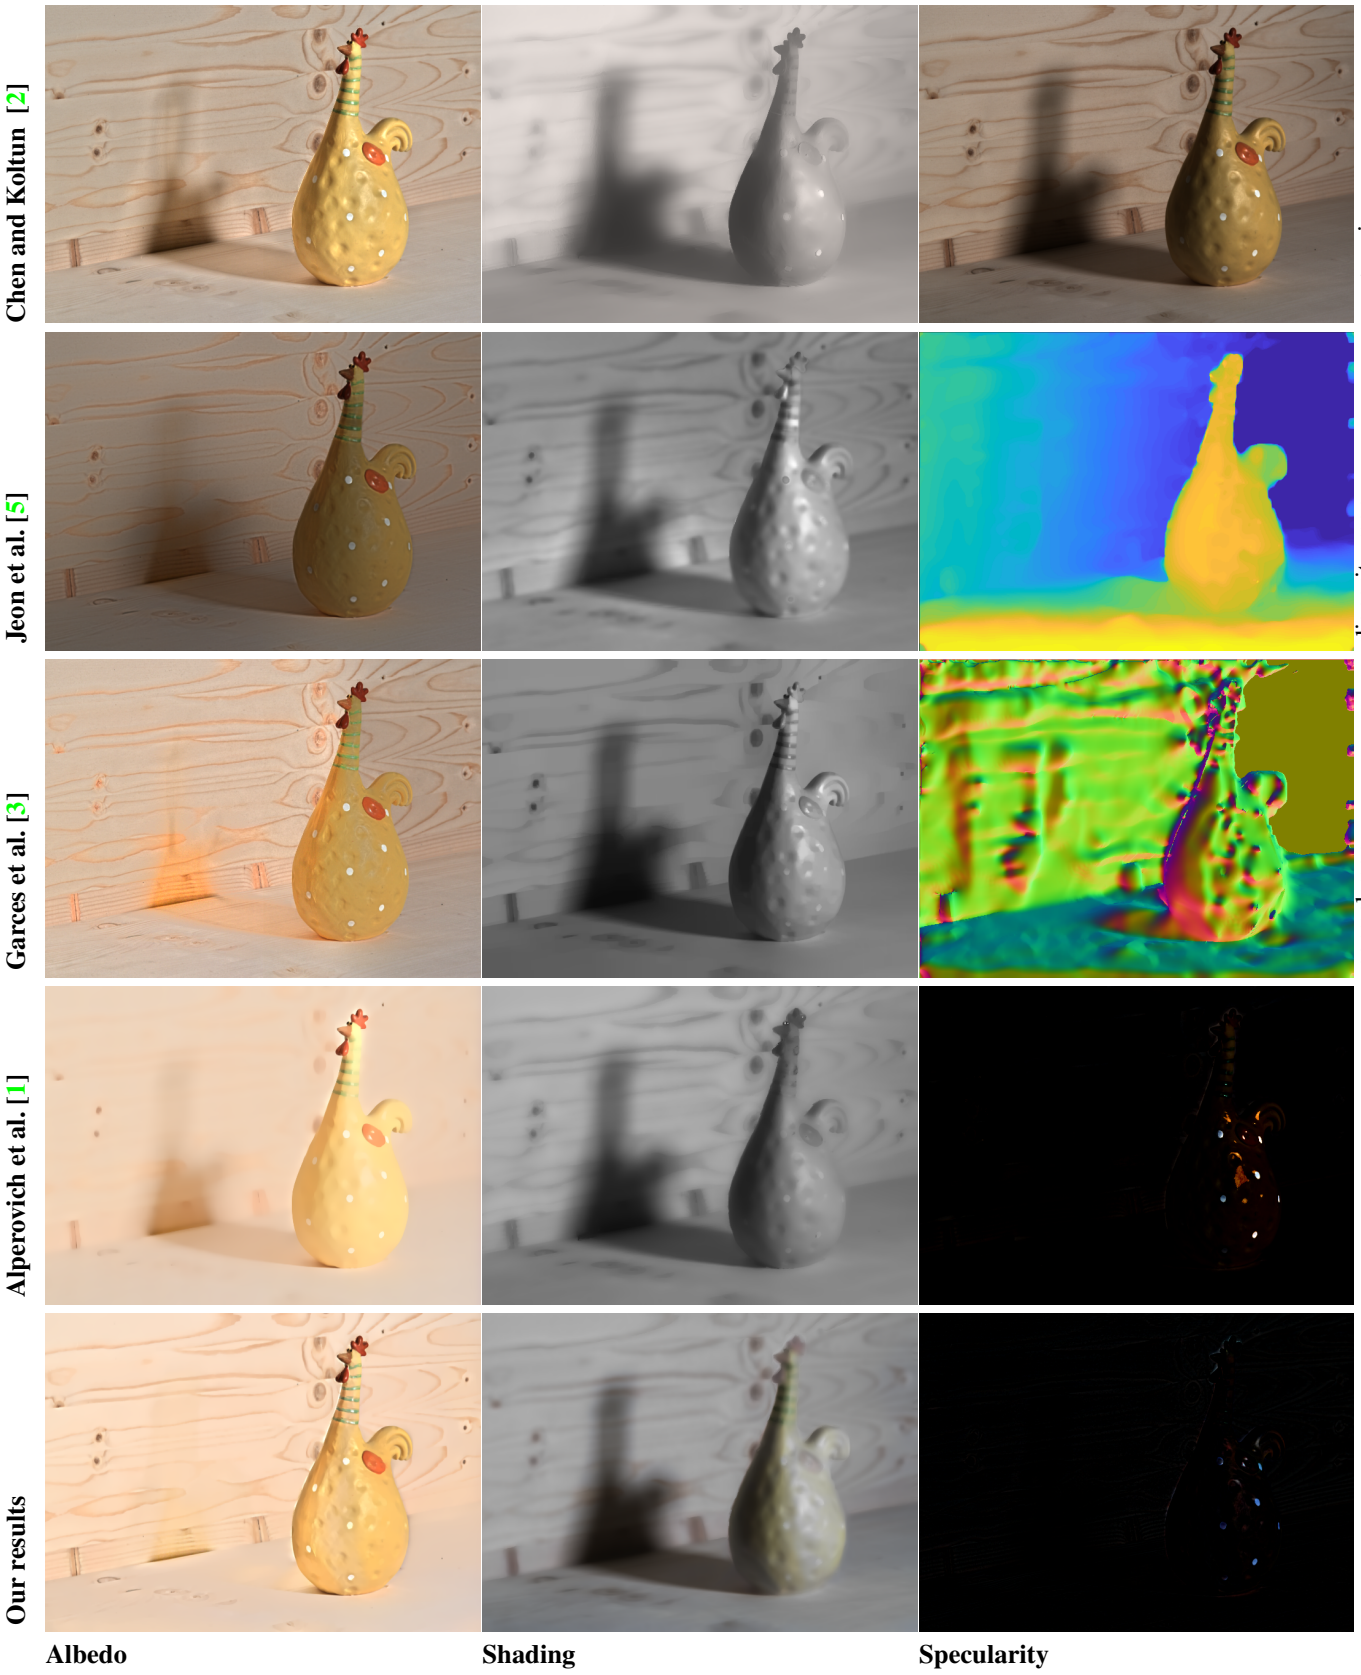
Chen and Koltun [2]
Jeon et al. [5]
Garces et al. [3]
Alperovich et al. [1]
Our results
center view
disparity map
normal map
Albedo
Shading
Specularity

specular mask, h

estimated diffuse chromaticity, χ

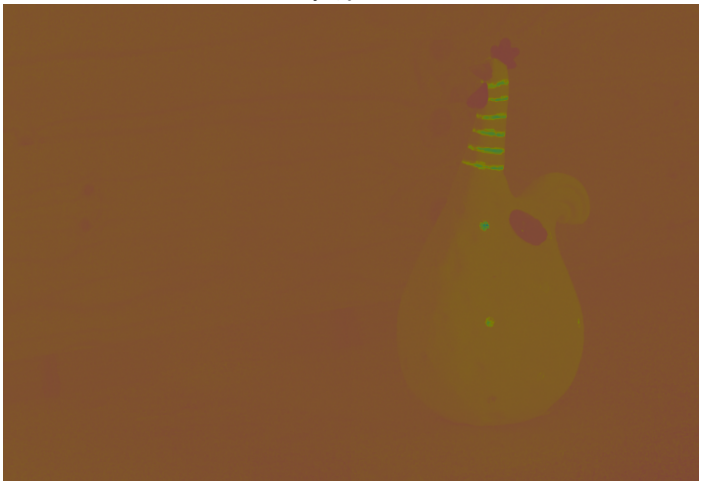

point-wise shadow confidence, β_{pw}

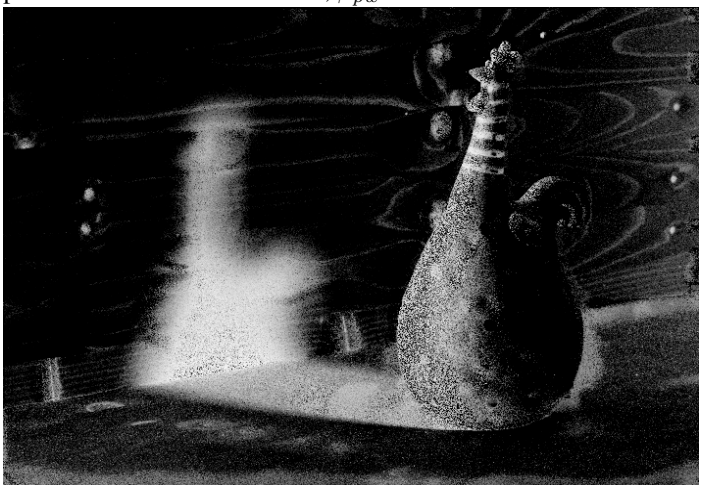

resulting shadow score, β

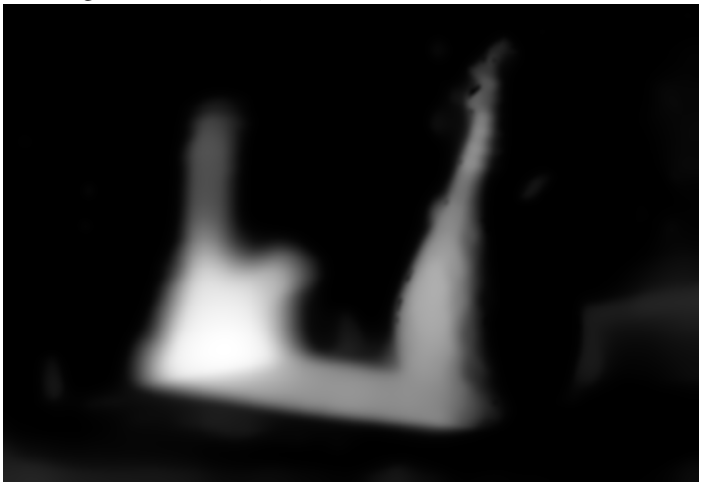

shadow boundaries score, δ_β

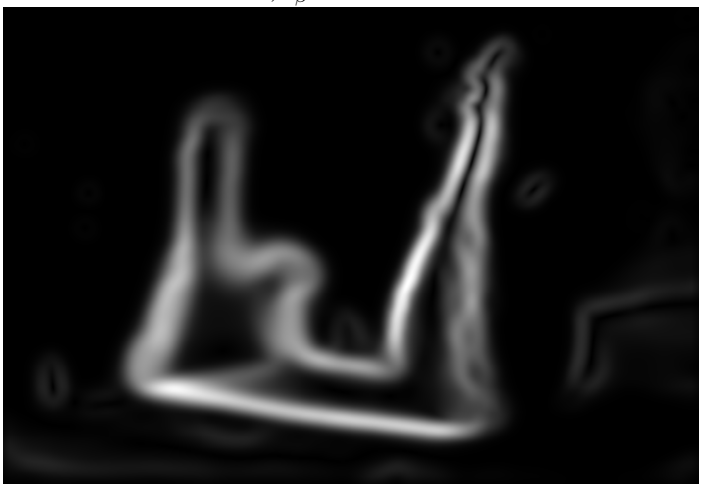

- Input light filed
- **Our results**
  - Albedo
  - Shading
  - Specularity
  - Direct shading
  - Indirect shading
  - Shadow score
- **Results by Alperovich and Goldluecke [1]**
  - Albedo
  - Shading
  - Specularity
- **Results by Garces et al. [3]**
  - Albedo
  - Shading