Improving High-Speed Scanning Systems by Photometric Stereo

VAST 2011: the International Symposium on Virtual Reality, Archaeology and Cultural Heritage

Frédéric Larue, Matteo Dellepiane and Roberto Scopigno

ISTI-CNR, Pisa, Italy

October 18-21, 2011

Introduction
3D scanning for CH

3D measurement tools already widespread for:

- Digital archiving
- Digital inspection
- Restoration of artifacts
- Communication purposes

Introduction
3D scanning for CH

3D measurement tools already widespread for:

- Digital archiving
- Digital inspection
- Restoration of artifacts
- Communication purposes

Why?

New inspection/processing possibilities, more DOF (interactive virtual tour, fragments re-assembly, ...)

Introduction

3D scanning for CH

3D measurement tools already widespread for:

- Digital archiving
- Digital inspection
- Restoration of artifacts
- Communication purposes

Why?

New inspection/processing possibilities, more DOF
(interactive virtual tour, fragments re-assembly, ...)

However...

Tedious, time consuming, often requires expertise

Introduction

High-speed scanning systems

Existing real-time scanners

RUSINKIEWICZ S., HALL-HOLT O., LEVOY M.:
Real-time 3D model acquisition.
ACM Trans. Graph., 2002

ZHANG L., CURLESS B., SEITZ S.:
Spacetime stereo: shape recovery for dynamic scenes.
IEEE CVPR, 2003

WEISE T., LEIBE B., VAN GOOL L.:
Fast 3D scanning with automatic motion compensation.
IEEE CVPR, 2007

Introduction

High-speed scanning systems

Advantages wrt. traditional technologies

- Digital copies produced in a few minutes
  - enables to digitize faster large collections
- Interactive feedback
  - intermediate results can be inspected on-the-fly
- No post-processing/manual intervention
  - acquisitions can be performed by non-expert users

Introduction

High-speed scanning systems

Advantages wrt. traditional technologies

- Digital copies produced in a few minutes
  - enables to digitize faster large collections
- Interactive feedback
  - intermediate results can be inspected on-the-fly
- No post-processing/manual intervention
  - acquisitions can be performed by non-expert users

Cost to pay for rapid acquisition

Loss of accuracy

Introduction
Proposed approach

Common hardware configuration

- Based on structured light
  - embeds a camera and a projector

Introduction

Proposed approach

Common hardware configuration

- Based on structured light
  - embeds a camera and a projector

Produced data

- Acquired geometry
- Calibration information
- Video flow
  - illumination: scanner's projector
  - viewpoint: scanner's acquisition camera

Introduction
Proposed approach

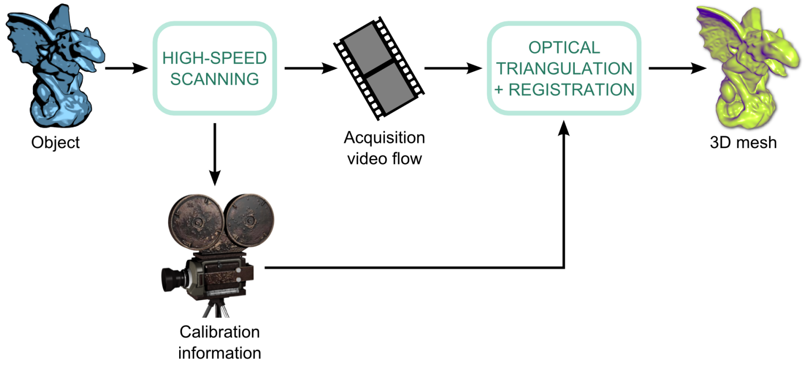

Introduction

Proposed approach

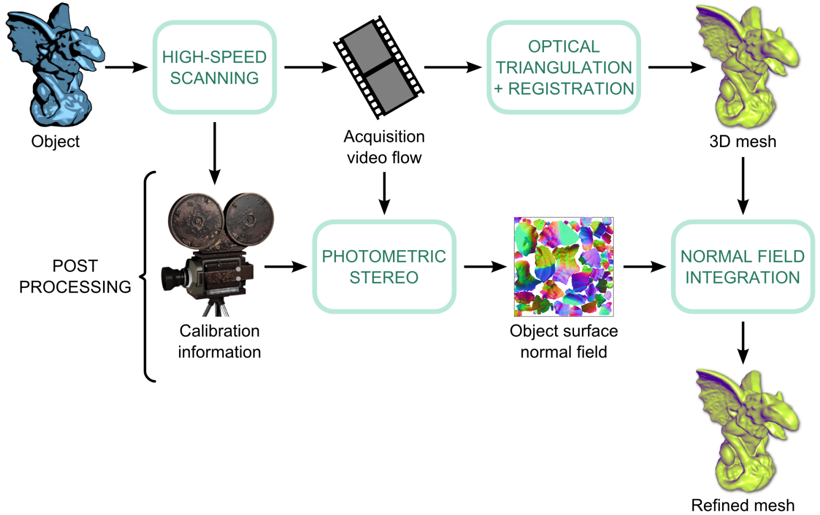

Introduction

Proposed approach

Similar approach

Nehab D., Rusinkiewicz S., Davis J., Ramamoorthi R.:
Efficiently combining positions and normals for precise 3D geometry.
ACM Trans. Graph., 2005

- Combines phase-shifting and photometry to improve shape measurement

Introduction

Proposed approach

Similar approach

NEHAB D., RUSINKIEWICZ S., DAVIS J., RAMAMOORTHI R.:
Efficiently combining positions and normals for precise 3D geometry.
ACM Trans. Graph., 2005

- Combines phase-shifting and photometry to improve shape measurement

Cons

- Involves a scanner made on purpose, not common systems
- Not intended to real-time scanning

Geometry Refinement

Basic estimation of normal vectors

Color / normal relation for Lambertian surfaces

The color c_i observed in a single video frame F_i at a surface point p is given by:

$$c_i = \frac{\rho_d}{{d_i}^2} \left\langle \vec{n}_p, \vec{l}_i \right\rangle$$

with:

- $\vec{n}_p$: unit normal vector at p
- $\vec{l}_i$: unit light direction at p for frame F_i
- d_i: light distance from p for frame F_i
- ρ_d: depends on light intensity, surface diffuse albedo, camera transfer function

Geometry Refinement

Basic estimation of normal vectors

When p is visible in several frames $\{F_i\}_{1 \leq i \leq N}$

A solution for $\vec{n}_p$ can be found by *least square fitting*:

$$\rho_d \ \vec{n}_p = \arg\min_X \ (LX - C)^2$$

with:

$$L = \begin{bmatrix} l_{1,x} & l_{1,y} & l_{1,z} \\ & \vdots & \\ l_{N,x} & l_{N,y} & l_{N,z} \end{bmatrix} \quad \text{and} \quad C = \begin{bmatrix} c_1 \ {d_1}^2 \\ \vdots \\ c_N \ {d_N}^2 \end{bmatrix}$$

Geometry Refinement

Basic estimation of normal vectors

When p is visible in several frames $\{F_i\}_{1 \leq i \leq N}$

A solution for $\vec{n}_p$ can be found by *least square fitting*:

$$\rho_d \ \vec{n}_p = \arg\min_X \ (LX - C)^2$$

with:

$$L = \begin{bmatrix} l_{1,x} & l_{1,y} & l_{1,z} \\ & \vdots & \\ l_{N,x} & l_{N,y} & l_{N,z} \end{bmatrix} \quad \text{and} \quad C = \begin{bmatrix} c_1 \ {d_1}^2 \\ \vdots \\ c_N \ {d_N}^2 \end{bmatrix}$$

Drawback

Solution quality depends on the sampling distribution

Geometry Refinement
Basic estimation of normal vectors

Case of high-speed scanners

- Distribution depends on the scanning trajectory
  - highly uneven

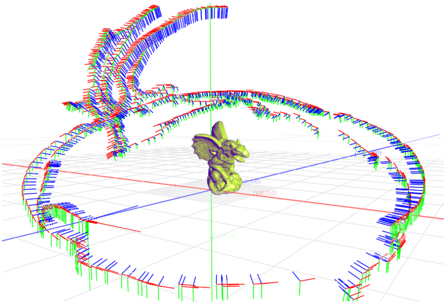

Geometry Refinement

Basic estimation of normal vectors

Case of high-speed scanners

- Distribution depends on the scanning trajectory
  - highly uneven
- Impact on the fitting:
  - well constrained along the trajectory
  - imprecise everywhere else

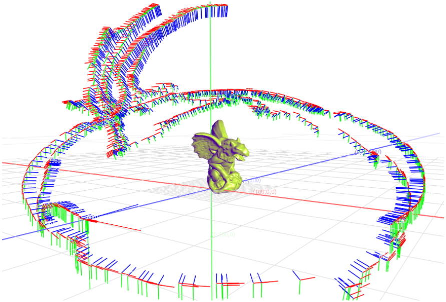

Geometry Refinement

Basic estimation of normal vectors

Case of high-speed scanners

- Distribution depends on the scanning trajectory
  - highly uneven
- Impact on the fitting:
  - well constrained along the trajectory
  - imprecise everywhere else

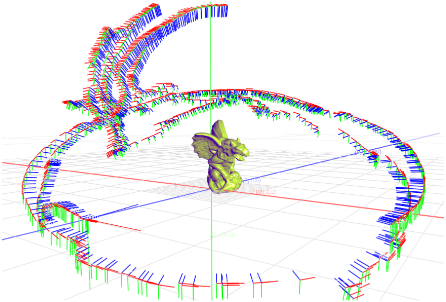

Our solution

Sampling distribution analysis to constrain the least square solution

Geometry Refinement

Light sampling analysis for improved normal fitting

Sampling distribution analysis

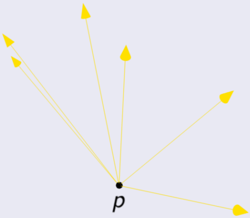

Consider a point p and its sampling of light directions

Geometry Refinement

Light sampling analysis for improved normal fitting

Sampling distribution analysis

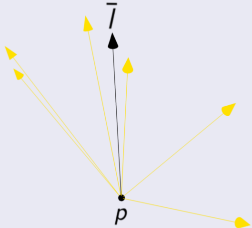

Compute the mean light direction $\bar{l}$

Geometry Refinement

Light sampling analysis for improved normal fitting

Sampling distribution analysis

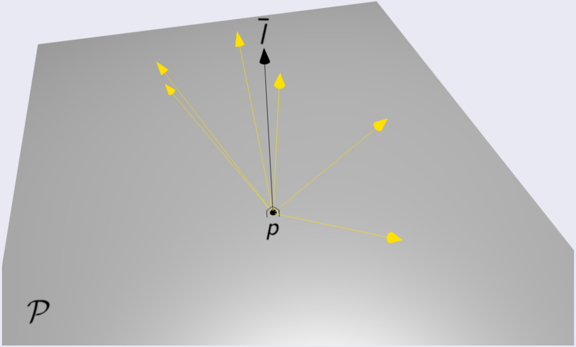

Create the plane $\mathcal{P}$ incident to p and orthogonal to $\bar{l}$

Geometry Refinement

Light sampling analysis for improved normal fitting

Sampling distribution analysis

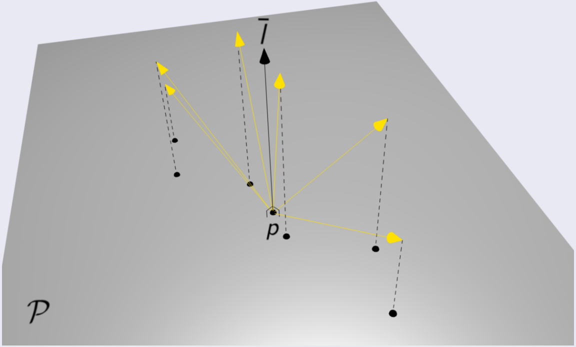

Project light samples onto $\mathcal{P}$

Geometry Refinement

Light sampling analysis for improved normal fitting

Sampling distribution analysis

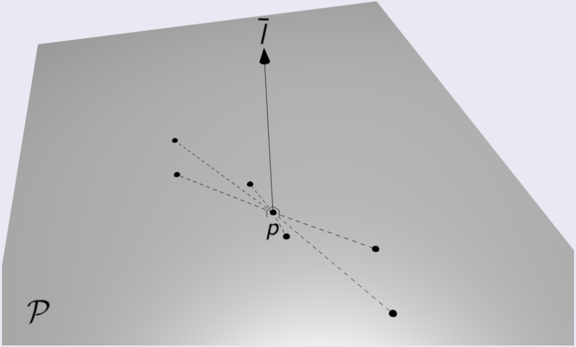

Project light samples onto $\mathcal{P}$

Geometry Refinement

Light sampling analysis for improved normal fitting

Sampling distribution analysis

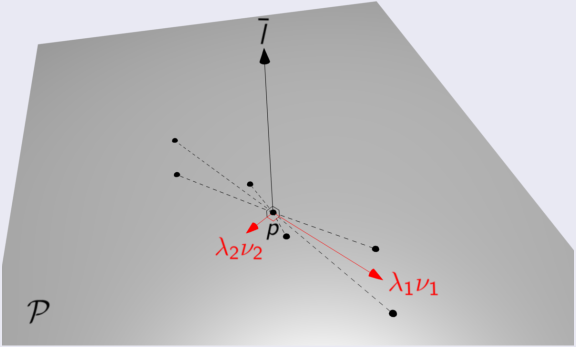

Perform a PCA on the projected samples

Leads to eigenvectors ν_1, ν_2 and eigenvalues λ_1, λ_2

Geometry Refinement

Light sampling analysis for improved normal fitting

Sampling distribution analysis

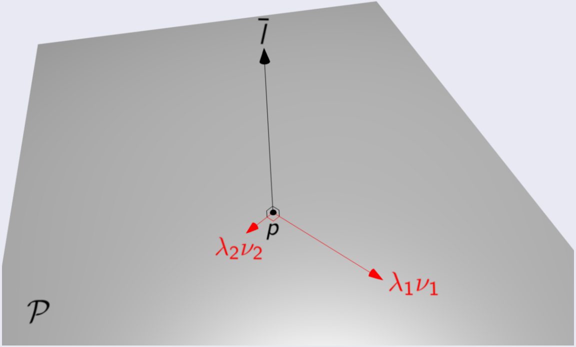

Along ν_1, the sampling distribution is the best

Along ν_2, it is the poorest

Geometry Refinement

Light sampling analysis for improved normal fitting

Adaptive correction of least square solution

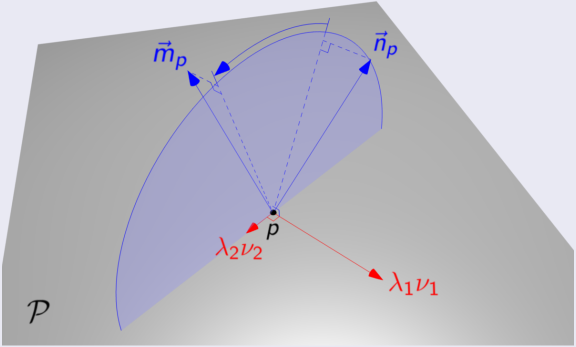

The fitted normal $\vec{n}_p$ is moved closer to the initial mesh normal $\vec{m}_p$ by a rotation around ν_1

Geometry Refinement

Light sampling analysis for improved normal fitting

Adaptive correction of least square solution

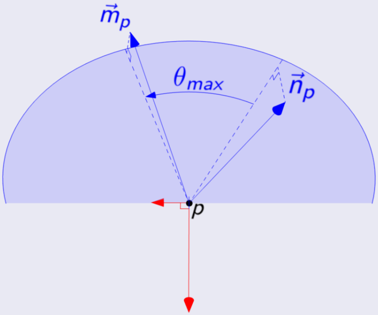

$$\text{Rotation amplitude} = \left(1 - \frac{\lambda_2}{\lambda_1}\right)\theta_{max}$$

where θ_{max} is the maximum correction angle

Geometry Refinement

Light sampling analysis for improved normal fitting

Adaptive correction of least square solution

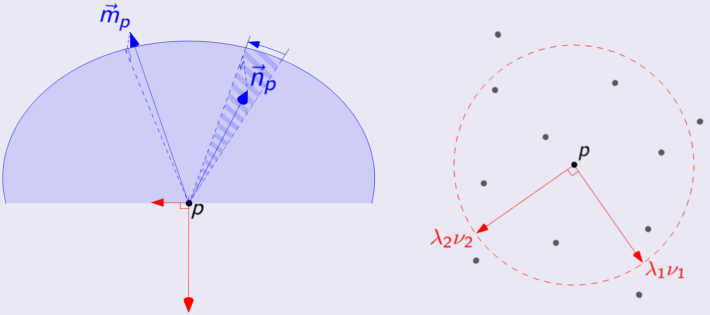

Even distribution: $\left(1 - \frac{\lambda_2}{\lambda_1}\right) \approx 0$

Smaller correction

Geometry Refinement

Light sampling analysis for improved normal fitting

Adaptive correction of least square solution

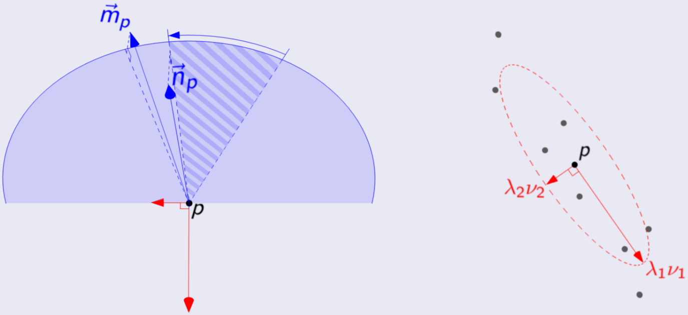

Uneven distribution: $\left(1 - \frac{\lambda_2}{\lambda_1}\right) \approx 1$

Larger correction

Geometry Refinement

Integration of the normal field

Iterative integration

- Each point p is moved along its initial normal $\vec{m}_p$ in the mesh of an amount Δ computed by:

$$\Delta = \left(\underbrace{A(p)}_{\substack{\text{push } p \text{ to follow} \\ \text{the normal field} \\ \text{curvature}}} + \underbrace{B(p)}_{\substack{\text{enforce } p \text{ to stay} \\ \text{close to the initial} \\ \text{surface}}} \right)$$

- Repeated until convergence

Results

Basic normal fitting *vs.* Correction by sampling analysis

Case 1 – Light directions *evenly* distributed

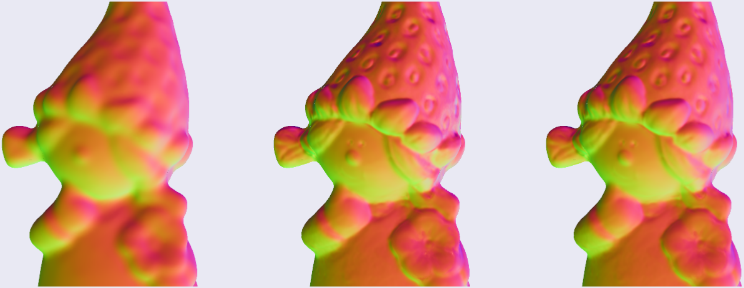

Normals from initial mesh | Normals from least square | Corrected least square

- both fittings behave similarly

Results

Basic normal fitting *vs.* Correction by sampling analysis

Case 2 – Light directions *unevenly* distributed

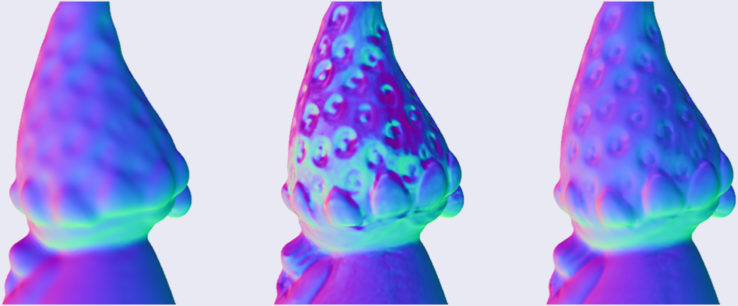

Normals from initial mesh — Normals from least square — Corrected least square

- incoherences in the normal field are corrected
- features are still recovered

Results
Mesh refined by normal integration

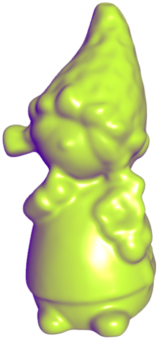

Initial mesh

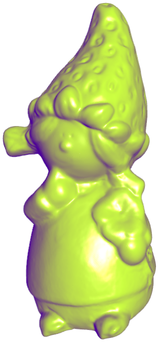

Refined mesh

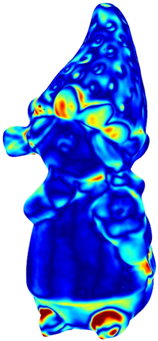

Hausdorff distance

Results

Mesh refined by normal integration

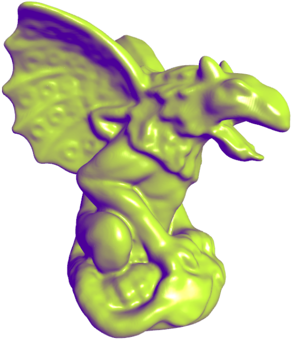

Initial mesh

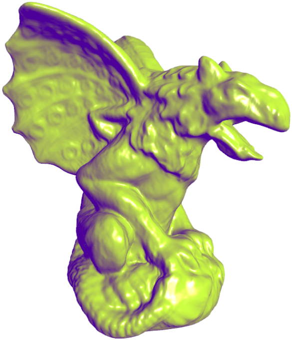

Refined mesh

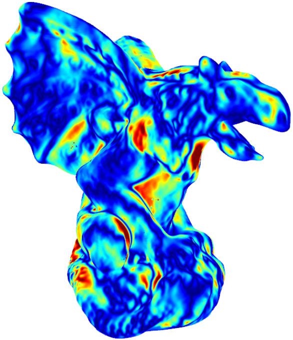

Hausdorff distance

Results

Examples using only texture and normal map

Results

Examples using only texture and normal map

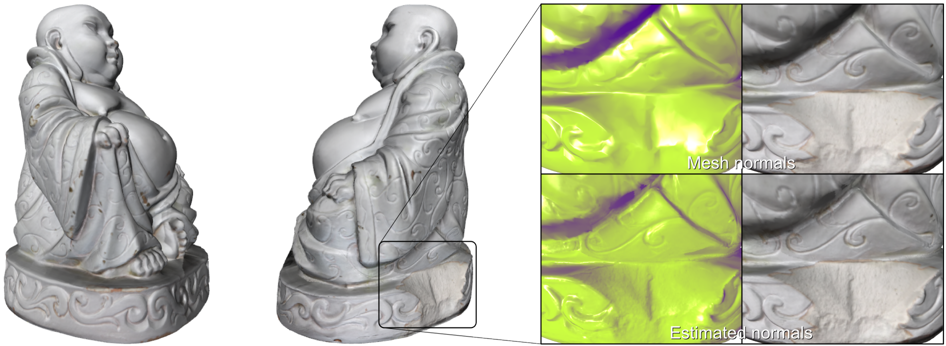

Conclusion

Pros

- Uses only data explicitly provided by the scanner
- Exploit high data density for refinement
- Accounts for a possible limited coverage of lighting directions
- Integrable in the classical pipeline of high-speed scanners

Limitations

- Fitting may fail for highly specular materials
- Technology mostly limited to small objects

Conclusion

Future work

- Formulate correction as an additional constraint for LS fitting
- Evaluation of results' accuracy
  - which ground truth to compare with?
- Get rid of the Lambertian material assumption
  - weighting accounting for deviation between lighting samples and perfect Lambertian model
- Online refinement
  - progressively updated while measurement is going on

Thank you for your attention

Thanks to ETH Zurich for datasets.

Thank you for your attention

Thanks to ETH Zurich for datasets.

And before going back home, make a researcher happy...

Ask a question

Tech Slide 1

Geometry integration

Iterative integration

- Each point p is moved along its initial normal $\vec{m}_p$ in the mesh of an amount Δ computed from its nearest neighbors $\mathcal{N}_p$:

$$\Delta = \frac{1}{2|\mathcal{N}_p|}\left(\underbrace{\sum_{q\in\mathcal{N}_p}\langle\vec{m}_p,\vec{n}_q\rangle\langle\vec{n}_q,\vec{pq}\rangle}_{\substack{\text{push } p \text{ to follow}\\ \text{the normal field}\\ \text{curvature}}} + \underbrace{\sum_{q\in\mathcal{N}_p}\langle\vec{m}_p,\vec{pq}\rangle}_{\substack{\text{enforce } p \text{ to stay}\\ \text{close to the initial}\\ \text{surface}}}\right)$$

- Repeated until convergence

Tech Slide 2

Performances

Dataset	*Video length (# frames)*	*Processing time (in seconds)*		
		Nor. fitting	*Integr.*	**Total**
Gargoyle	1070	198	5	**203**
Dwarf	2110	277	6	**283**

Table: Processing statistics for our test datasets.