# *Compact and Efficient Generation of Radiance Transfer for Dynamically Articulated Characters*

D. Nowrouzezahrai,

P. Simari, E. Kalogerakis, K. Singh, E. Fiume

dgp

Dynamic Graphics Project
University of Toronto
*www.dgp.toronto.edu*

# Goal, Challenges & assumptions

- Self-shadowing of a *dynamically* articulating character mesh
  - Under varying external illumination
  - High-performance (> 100 Hz)

- Per-vertex visibility is prohibitively expensive to calculate every frame

- Diffuse-only characters
- Direct-illumination
- Low-frequency distant lighting and shading

# Previous Work

- Kirk and Arikan, Kontkanen and Aila
  - AO as a linear function of joint angles
- Ren et al. abstract animating geometry with spheres and use SH log space accumulation
- Zhou et al. precompute shadow fields for rigid motion

# Main Idea

- Given a training animation, **f** frames long, of a mesh with **v** vertices:
  1. Precompute the per-vertex/per-frame cosine-weighted visibility functions (expensive)
  2. SH project these spherical functions, yielding **c** coefficients (per-vertex, per-frame = **fvc** floats in a training set)
- Can we fit a linear model of the joint angles to this dataset?
- Is more aggressive compression possible/useful?

# Is a Linear Model Sufficient?

- We show empirically that a linear model of the joint angles <u>sufficiently fits the SH projected cosine-weighted visibility data</u>

- Statistically speaking:
  - We show training and test error
  - Residual distribution

# Why the Linear Model is Good

- Sufficient reconstruction of the dataset
- Optimally solvable in closed form
- Smooth by construction
  - Consistent, plausible results
- Compact representation (as small as 1MB)
- Very fast to evaluate (100 Hz, including shading)

# Regularization

- Our technique can be used solely as a compression method for animated shadowing data
- But, in general, we want to train on a small (representative) dataset and *generalize* well to unseen poses
- We apply regularization to get sensible (actually, very sensible) reconstruction error for dynamically articulated animations

# Empirical Proof of Linearity

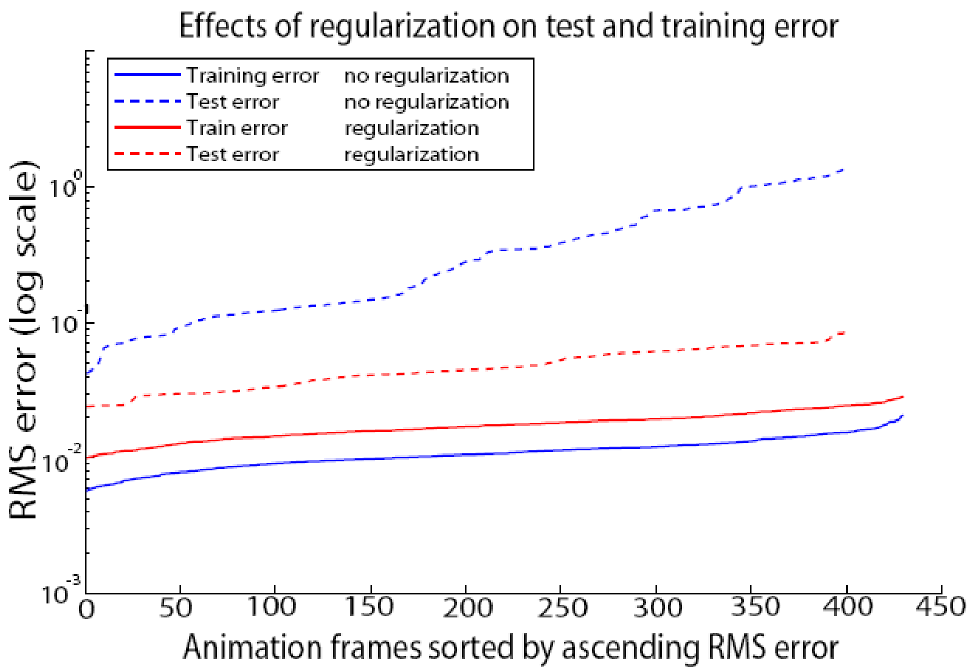


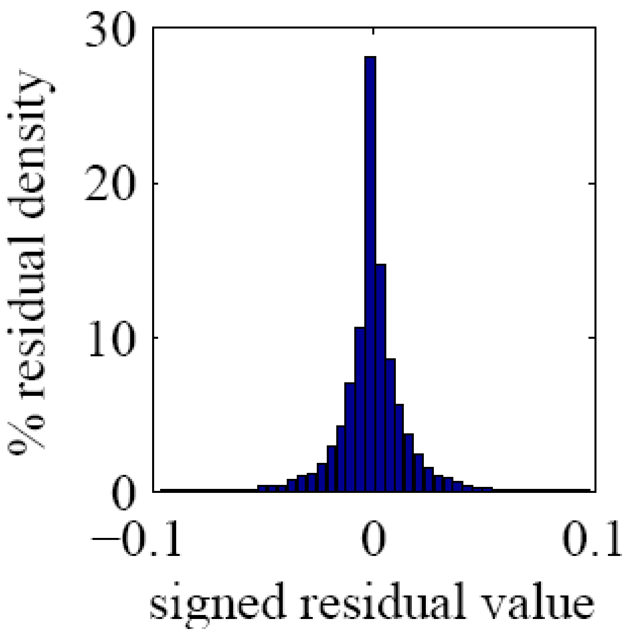

# Dimensionality Reduction

- Given an articulated character with **a** joint angles:
  - for each of the **c** coefficients, we have a linear mapping, **X** of size **a** x **v**
- Observations:
  1. PCA over joint angles shows high variance in only the first few eigenvectors
  2. PCA over each SH coefficient at vertices also yields a similar, yet less pronounced, drop-off of variance

# Visual Artifacts due to PCA

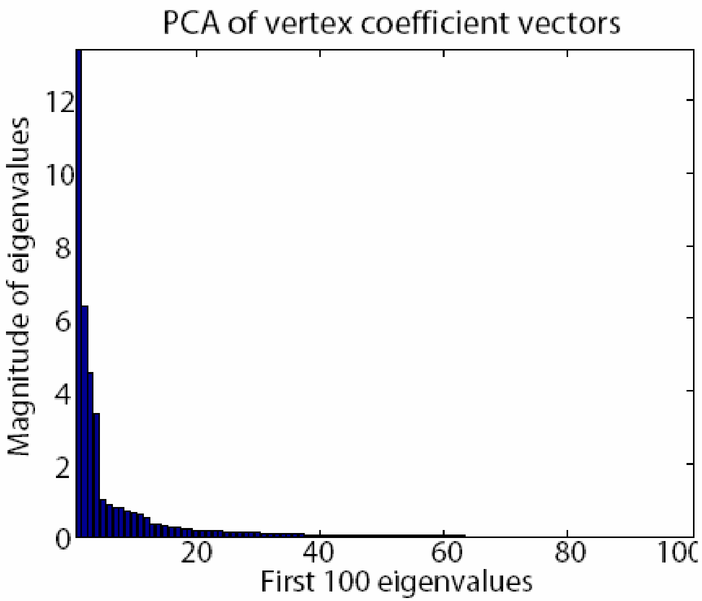


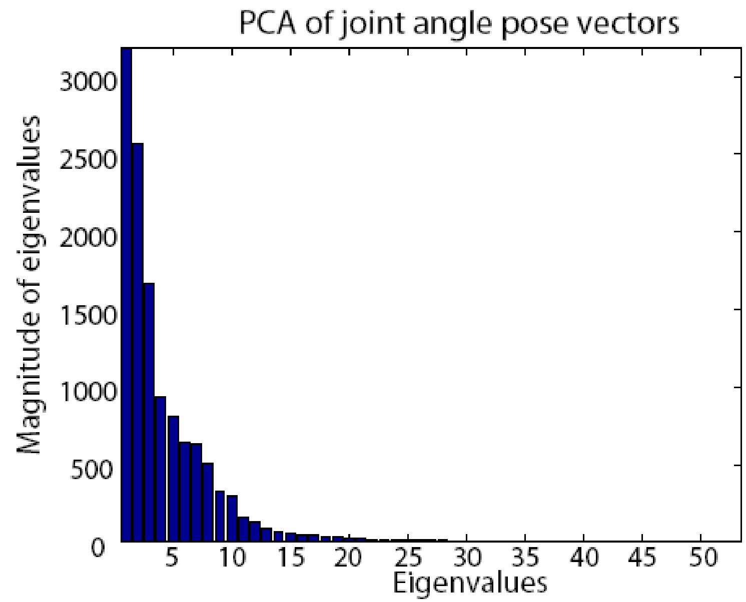



22 eigenvalues (90% variance)


500 eigenvalues (99% variance)


1000 eigenvalues (99.9% variance)

# Run-time Reconstruction

No Dimensionality Reduction:

1. Query joint angles
2. Perform **c** vector/matrix multiplications with linear models
   - → This yields the **c** SH cosine-weighted visibility coefficients
3. Dot the currently generated per-vertex SH coefficients with the external lighting coefficients → Final shade

# Run-time Reconstruction cont…

With Dimensionality Reduction:

1. Query joint angles
2. Project joint angle vector to lower-dimensional joint angle vector
3. Perform **c** vector/matrix multiplications with <u>reduced</u> linear models
4. Reproject low-dimensional coefficients to original number of vertices
5. Dot the currently generated per-vertex SH coefficients with the external lighting coefficients → Final shade

# Storage Requirements

| | Master Pai (11.5k vertices & 54 joint angles) Trained on a 430 frame animation. | | |
|---|---|---|---|
| | Memory: Absolute | Memory: % of total data | Memory: % of 1 mesh's data |
| Linear Model (No PCA reduction) | 87.5MB | 12.56% | 5400% |
| PCA (15,500) | 23.6MB | 3.39% | 1457% |
| PCA (8,22) | **1.0MB** | **0.146%** | **62.9%** |

| | Armadillo (24.9k vertices & 25 joint angles) Trained on a 250 frame animation. | | |
|---|---|---|---|
| | Memory: Absolute | Memory: % of total data | Memory: % of 1 mesh's data |
| Linear Model (No PCA reduction) | 85.4MB | 10.0% | 2500% |
| PCA (15,500) | 48.6MB | 5.68% | 1421.13% |
| PCA (8,22) | **2.1MB** | **0.248%** | **61.93%** |

# Conclusions

- A linear function of joint angles captures low-frequency shadowing effects on:
  - Training animation sequences, and
  - Dynamically generated (test) animation

- Dimensionality reduction allows to further trade reconstruction accuracy with storage requirements

# Thank you! Questions?