Motion Compression using Principal Geodesic Analysis

Maxime Tournier[1] Xiaomao Wu[1] Nicolas Courty[2]
Elise Arnaud[1] Lionel Revéret[1]

[1]Laboratoire Jean Kuntzmann

Université de Grenoble

INRIA Rhône-Alpes

France

[2]Université Bretagne-Sud, France

March 30, 2009

Motion Capture

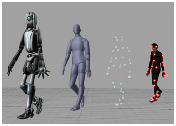  

- Ubiquitous technique (entertainment, biomechanics, ...)
- Quality data, at the expense of the size:
  - 120Hz
  - 30-60 DOFs

Size does matter

Storage / transmission are difficult:

- Storage size is limited
- Bandwidth is expensive

Motion Edition

To improve immersion: provide the user with diversity

- Pose selection from a database
- Motion synthesis
  - Blending of existing clips *Motion Graphs* [KSG02]
  - Learning on a database *Style-Based IK* [GMHP04]
- Need big quantities of data

Problem Formulation

We want a motion representation that is both:

- Compact
- Easily editable

Motion Capture Compression

- Temporal coherence / Spatial coherence

  Splines, wavelets — reduce DOFs, PCA

- Motion database [Ari06] vs Single motion [LM06][BPP07]

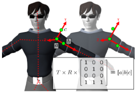

- Joints orientations [BPP07] vs Marker positions [Ari06][LM06]

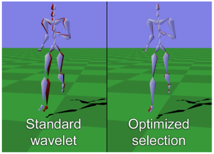

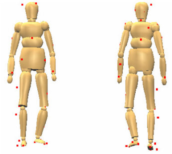

Euler angles, errors accumulation — bones' rigidity lost: needs fitting

What's Important

- Avoid skeleton fitting: use orientations
- Exploit spatial coherence to build pose models
- IK always needed to fix artifacts: integrate it into the pipeline
  - Possibly exploiting spatial coherence

- Our point of view: compression as motion re-synthesis

Motion Synthesis

- Exploit spatial coherence
  - Fewer DOFs, encoding statistical correlations
- Ease expensive optimizations
  - Physically plausible motion synthesis [SHP04]
- Style-based Inverse Kinematics
  - Learning using probabilistic inference [GMHP04]

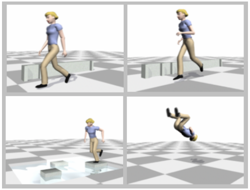

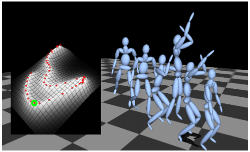

Statistical Models for Rotation Data

- $SO(3)$: not a vector space !
  - Linear methods are not suitable
  - Deviations, degenerate results
- Work with charts *(=linearizations)*
  - Only valid locally, introduce distortions
  - Troublesome for statistical analysis
  - Resort to *expensive* non-linear methods
- Need for an *intrinsic* method
  - Takes the *geometry* of $SO(3)$ into account
  - Principal Geodesic Analysis (PGA) [FLPJ04]

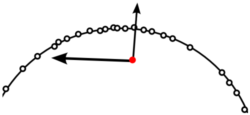

PCA of data on a circle

Overview of PGA

In a nutshell:

- Projects data over *geodesics*
  - Locally length-minimizing curves
  - = *vectorial lines in the linear case*
- Maximizes projected variance
- No analytical expression for geodesic directions in the general case
  - Optimization
  - Approximation in a "good" chart

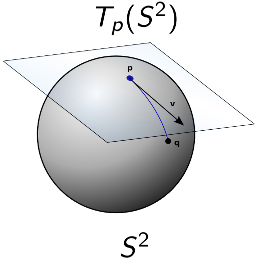

Analogy with PCA

PCA	PGA
Data: $x_i \in \mathbb{R}^{3n}$	Data: $x_i \in M = SO(3)^n$

Analogy with PCA

PCA

Data: $x_i \in \mathbb{R}^{3n}$

Data mean: $\mu = \frac{1}{m} \sum_i x_i$

PGA

Data: $x_i \in M = SO(3)^n$

Intrinsic mean: $\mu = \underset{y \in M}{argmin} \sum_i d(y, x_i)^2$ [Pen06]

Analogy with PCA

PCA	PGA
Data: $x_i \in \mathbb{R}^{3n}$	Data: $x_i \in M = SO(3)^n$
Data mean: $\mu = \frac{1}{m}\sum_i x_i$	*Intrinsic* mean: $\mu = \underset{y \in M}{argmin} \sum_i d(y, x_i)^2$ [Pen06]
k principal components $v_i \in \mathbb{R}^{3n}$	k geodesic directions $v_i \in T_\mu(M)$

Analogy with PCA

PCA	PGA
Data: $x_i \in \mathbb{R}^{3n}$	Data: $x_i \in M = SO(3)^n$
Data mean: $\mu = \frac{1}{m} \sum_i x_i$	*Intrinsic* mean: $\mu = \underset{y \in M}{argmin} \sum_i d(y, x_i)^2$ [Pen06]
k principal components $v_i \in \mathbb{R}^{3n}$	k geodesic directions $v_i \in T_\mu(M)$
Reconstruction: $y = \mu + \sum_{j=1}^{j=k} \alpha_j . v_j$	Reconstruction: $y = \mu . \prod_{j=1}^{j=k} e^{\alpha_j . v_j}$

$(\alpha_j)_{j \leq k}$ = coordinates along principal component/geodesics

Our approach

Most of the pose information is encoded by the end-effector positions

1. Store only a few marker trajectories as controllers of the pose
2. Parametrize the poses of one motion using PGA
3. Derive an IK algorithm using the PGA pose model
4. Exploit temporal coherence in the marker trajectories

Definitions - Notations

- Skeleton: n joints + root
- Pose: $p \in SO(3)^n$
  - + position/orientation of skeleton in space
- Animation: m frames

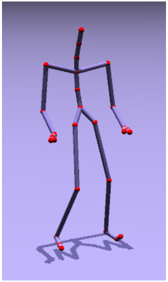

Pose Space Parametrization

- Apply PGA to the motion poses: $M = SO(3)^n$
- Keep only the first k geodesics
  - k chosen to account for 99% variance of the input data
  - In practice, $5 \leq k \leq 20$ most of the time
- Reduced coordinates: $(\alpha_j)_{j \leq k}$
  - Coordinates along the principal geodesics
  - $\approx$ weight of each *eigen pose*

Some examples of geodesics

IK by Searching the Pose Space

Optimize $f : (\alpha_j)_{j \leq k} \mapsto$ end-effector positions

- Non-linear least squares
  - Levenberg-Marquardt solver
- *Analytical Jacobian* thanks to the geodesics formulation
  - Real-time
- Good optimization behavior
  - Pose extrapolation
  - No local extrema reported

Compression of End-effector Trajectories

- Spline interpolation
  - Recursive insertion of p control points (usually, $p \leq \frac{m}{8}$)
  - Error threshold
  - Only store p control points, possibly modified
- Or any other temporal-coherence-based compression scheme
- Also compress root absolute position/orientation

Quick Recap

Size Analysis

$$m \overbrace{\begin{pmatrix} \\ \\ \text{original data} \\ \\ \\ \end{pmatrix}}^{3\times(n+2)} \overset{\text{Compression}}{\Longrightarrow}$$

Size Analysis

$$m\left(\overbrace{\quad\text{original data}\quad}^{3\times(n+2)}\right) \xrightarrow{\text{Compression}} \begin{array}{l} k\left(\overbrace{\quad geodesics \quad}^{3\times n}\right) \\ 1\left(\quad mean \quad\right) \end{array}$$

Size Analysis

$3\times(n+2)$

m (original data)

Compression $\Rightarrow$

$3\times n$

k (*geodesics*)

1 (*mean*)

p (end-effectors) (root pos/orient)

3×5 3×2

Typical Example

- Walking motion (CMU subject #17, trial #8)
  - $n = 30, m = 6179$
- Kept $k = 12$ geodesics, $p = \frac{m}{32}$
- Ratio $r = \frac{(30+2).m}{30.(k+1)+7.\frac{m}{32}} \approx 111$

Decompression

Straightforward:

1. Decompress root joint positions/orientations
2. Decompress end-effector trajectories, expressed in root frame
3. Use PGA-based IK to recover each pose

Distortion metric

- For comparison, we use the same distortion rate d as [LM06]
- $d = 100\frac{\|A - \tilde{A}\|}{\|A - E(A)\|}$
  - A: original absolute positions
  - $\tilde{A}$ reconstructed absolute positions
  - $E(A)$ mean positions of A repeated

Example Results

Walking:
$k = 12$, $p = n/64$, $d = 0.3\%$

Boxing:
$k = 12$, $p = n/16$, $d = 0.49\%$

Breakdance:
$k = 15$, $p = n/32$, $d = 0.56\%$

Limitations/Future Work

- Online optimization
  - Still real-time
  - Implementation could be faster
- Separate encoding of sharp features
  - Splines: smoothes important C^1 discontinuities
  - Error quantizing...
- Motion segments for long, diverse clips

Conclusion

- High compression ratios
  - Up to 1:100, without 16 bits quantization
- Compact, editable data

Thank You for Your Attention !

Questions ?

Okan Arikan.
Compression of motion capture databases.
ACM Trans. Graph., 25(3):890–897, 2006.

Philippe Beaudoin, Pierre Poulin, and Michiel van de Panne.
Adapting wavelet compression to human motion capture clips.
In *GI '07: Proceedings of Graphics Interface 2007*, pages 313–318, New York, NY, USA, 2007. ACM.

P. Thomas Fletcher, Conglin Lu, Stephen M. Pizer, and Sarang C. Joshi.
Principal geodesic analysis for the study of nonlinear statistics of shape.
IEEE Transactions on Medical Imaging, 23(8):995, 2004.

Keith Grochow, Steven L. Martin, Aaron Hertzmann, and Zoran Popovic.
Style-based inverse kinematics.

In *ACM SIGGRAPH 2004 Papers on - SIGGRAPH 04 SIGGRAPH 04*, page 522, New York, NY, USA, 2004. ACM Press.

Lucas Kovar, John Schreiner, and Michael Gleicher.
Footskate cleanup for motion capture editing.
In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation - SCA 02 SCA 02*, page 97, New York, NY, USA, 2002. ACM Press.

Guodong Liu and Leonard McMillan.
Segment-based human motion compression.
In *SCA '06: Proceedings of the 2006 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 127–135, Aire-la-Ville, Switzerland, Switzerland, 2006. Eurographics Association.

Xavier Pennec.

Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements.
Journal of Mathematical Imaging and Vision, 25(1):127, 2006.

Alla Safonova, Jessica K. Hodgins, and Nancy S. Pollard.
Synthesizing physically realistic human motion in low-dimensional, behavior-specific spaces.
In *ACM SIGGRAPH 2004 Papers on - SIGGRAPH 04 SIGGRAPH 04*, page 514, New York, NY, USA, 2004. ACM Press.

Liu et al. 2006:

Sequence	Compression	Distortion rate d	Decompr. time (msec/pose)
Jumps, bends, lift up	1:55.2	5.1	0.7
Long breakdance sequence	1:18.4	7.1	0.7
Walf, stretch, punches, drink	1:61.7	5.1	0.7
Walk, stretch, punches, kicks	1:56.0	5.4	0.7

Our results:

Subject/Trial	Description	# poses	Rate	Distortion	msec/pose
09/06	Running, short	141	1:18	0.36	7.88
17/08	Walking, long	6179	1:182	0.049	16.2
15/04	Mixed, dancing, boxing	22948	1:69	1.55	30.6
85/12	Breakdance	4499	1:97	0.56	20.42
17/10	Boxe	2783	1:61	0.49	15.97

Arikan et al. 2006

Exploite les redondances au sein d'une base de données

- Repères fictifs
  - Interpolation spline des trajectoires
  - ACP par paquets sur les points de contôle
- Moindres carrés pour recaler le squelette
- Encodage séparé des empreintes de pied

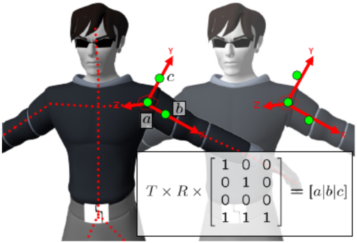

Liu et McMillan 2006

Compression des positions des marqueurs, un seul mouvement

- Modèles linéaires sur portions de mouvement
- Segmentation automatique
- Réduction par ACP, interpolation spline des coordonnées
- Recalage du squelette

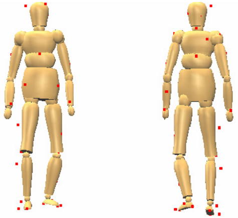

Beaudoin et al. 2007

Compression par ondelettes des angles d'Euler, un seul mouvement

- Cumul hiérarchique des erreurs
  - Séléction automatique de la base
- Problèmes classiques angles d'Euler
  - Pas plus court chemin, Gimbal Lock
- Correction par IK

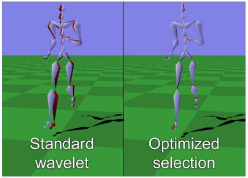

Comparaisons - Mouvements base CMU

Comparaison avec [LM06] (meilleurs résultats)

- Métrique de distortion: $d = 100\frac{\|A-\tilde{A}\|}{\|A-E(A)\|}$
  - Renseigne très peu sur la qualité réelle des animations...
  - Aucune métrique de perception satisfaisante
- Note: [LM06] quantifient les données sur 16 bits

Exponential Map

A special chart: the exponential map

Exponential Map

A special chart: the exponential map

- *Tangent directions* $\leftrightarrow$ *geodesics* starting at $x \in M$
  - $exp_x : T_x(M) \longrightarrow M$
  - $log_x : M \longrightarrow T_x(M)$

Exponential Map

A special chart: the exponential map

- *Tangent directions* $\leftrightarrow$ *geodesics* starting at $x \in M$
  - $exp_x : T_x(M) \longrightarrow M$
  - $log_x : M \longrightarrow T_x(M)$

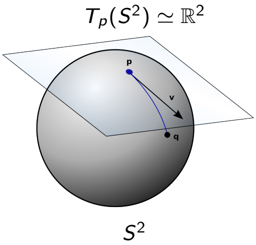

Exponential Map

A special chart: the exponential map

- *Tangent directions* $\leftrightarrow$ *geodesics* starting at $x \in M$
  - $exp_x : T_x(M) \rightarrow M$
  - $log_x : M \rightarrow T_x(M)$

$T_p(S^2) \simeq \mathbb{R}^2$

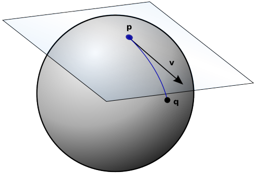

S^2

$exp_p(v) = q$

$log_p(q) = v$

Exponential Map

A special chart: the exponential map

- *Tangent directions* $\leftrightarrow$ *geodesics* starting at $x \in M$
  - $exp_x : T_x(M) \rightarrow M$
  - $log_x : M \rightarrow T_x(M)$
- Easily computed for $SO(3)$

$T_p(S^2) \simeq \mathbb{R}^2$

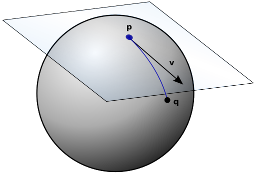

S^2

$exp_p(v) = q$

$log_p(q) = v$

Exponential Map

A special chart: the exponential map

- *Tangent directions* $\leftrightarrow$ *geodesics* starting at $x \in M$
  - $exp_x : T_x(M) \rightarrow M$
  - $log_x : M \rightarrow T_x(M)$
- Easily computed for $SO(3)$
- A way to express *geodesic distance* $d(x, y)$
  (hence variance)

$T_p(S^2) \simeq \mathbb{R}^2$

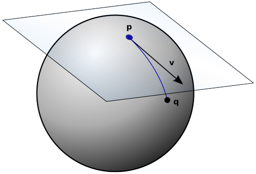

S^2

$exp_p(v) = q$

$log_p(q) = v$

Approximate Principal Geodesics Computation

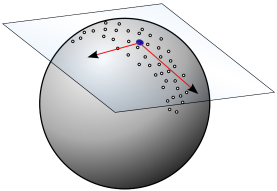

Approximate Principal Geodesics Computation

1. Linearize the data in $T_\mu(M)$
   - By using log_μ
   - Best chart w.r.t the data

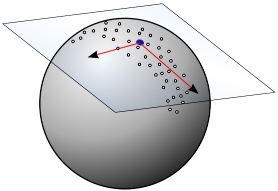

Approximate Principal Geodesics Computation

1. Linearize the data in $T_\mu(M)$
   - By using log_μ
   - Best chart w.r.t the data
2. PCA of tangent data:
   use principal components as geodesics directions