A Game-Theoretic Approach to Robust Inlier Selection

Andrea Torsello

Joint work with Andrea Albarelli, Emanuele Rodolà, Samuel Rota Bulò, and Marcello Pelillo

Matching Problem

The matching problem is one of finding correspondences within a set of elements, or features

Central to any recognition task where the object to be recognized is naturally divided into several parts

Several approaches (Too many to mention), but mostly based on global optimization (Quadratic Assignment Problem)

Approaches in the literature are generally rather greedy
(the more correspondences, the better)

Game-Theoretic Matching

Adopt a non-cooperative game-theoretic view to correspondence estimation

Correspondences are allowed to compete with one another in a *matching game*, a non-cooperative game where

- potential associations between the items to be matched correspond to strategies
- payoffs reflect the degree of compatibility between competing hypotheses

The solutions of the matching problem correspond to ESS's (dominant sets in the association space)

The framework can deal with general many-to-many matching problems even in the presence of asymmetric compatibilities.

Matching game

Let O1 and O2 be the two sets of features that we want to match and $A \subseteq O1 \times O2$ the set of feasible associations that satisfy the unary constraints. Each feasible association represents a possible matching hypothesis.

Let $C : A \times A \rightarrow R+$ be a set of pairwise compatibilities that measure the support that one association gives to the other.

A submatch (or simply a match) is a set of associations, which satisfies the pairwise feasibility constraints, and two additional criteria:

- High internal compatibility, i.e. the associations belonging to the match are mutually highly compatible
- low external compatibility, i.e. associations outside the match are scarcely compatible with those inside.

.

The proposed approach generalizes the association graph technique described by Barrow and Burstall to continuous structural constraints

Properties of Matching Games

Domain-specific information is confined to the definition of the compatibility function.

We are able to deal with many-to-many, one-to-many, many-to-one and one-to-one relations incorporating any hard binary constraints with the compatibilities (setting them to 0)

Theorem: Consider a matching-game with compatibilities $C = (c_{ij})$ with $c_{ij} \geq 0$ and $c_{ii} = 0$. If $x \in \Delta$ is an ESS then $c_{ij} > 0$ for all $i, j \in \sigma(x)$

Matching Examples

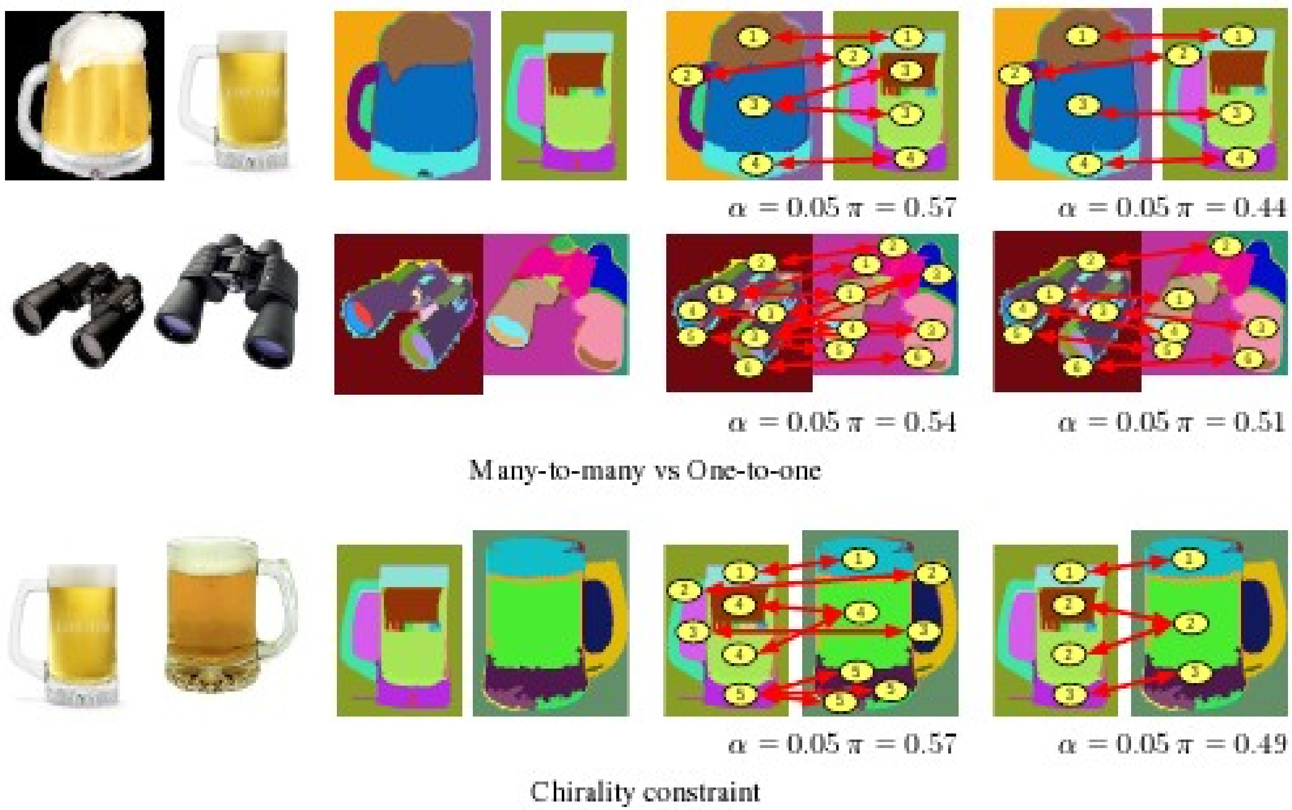

A. Albarelli, S. Rota-Bulò, A. Torsello, and M. Pelillo; ICCV 2009

Matching Examples

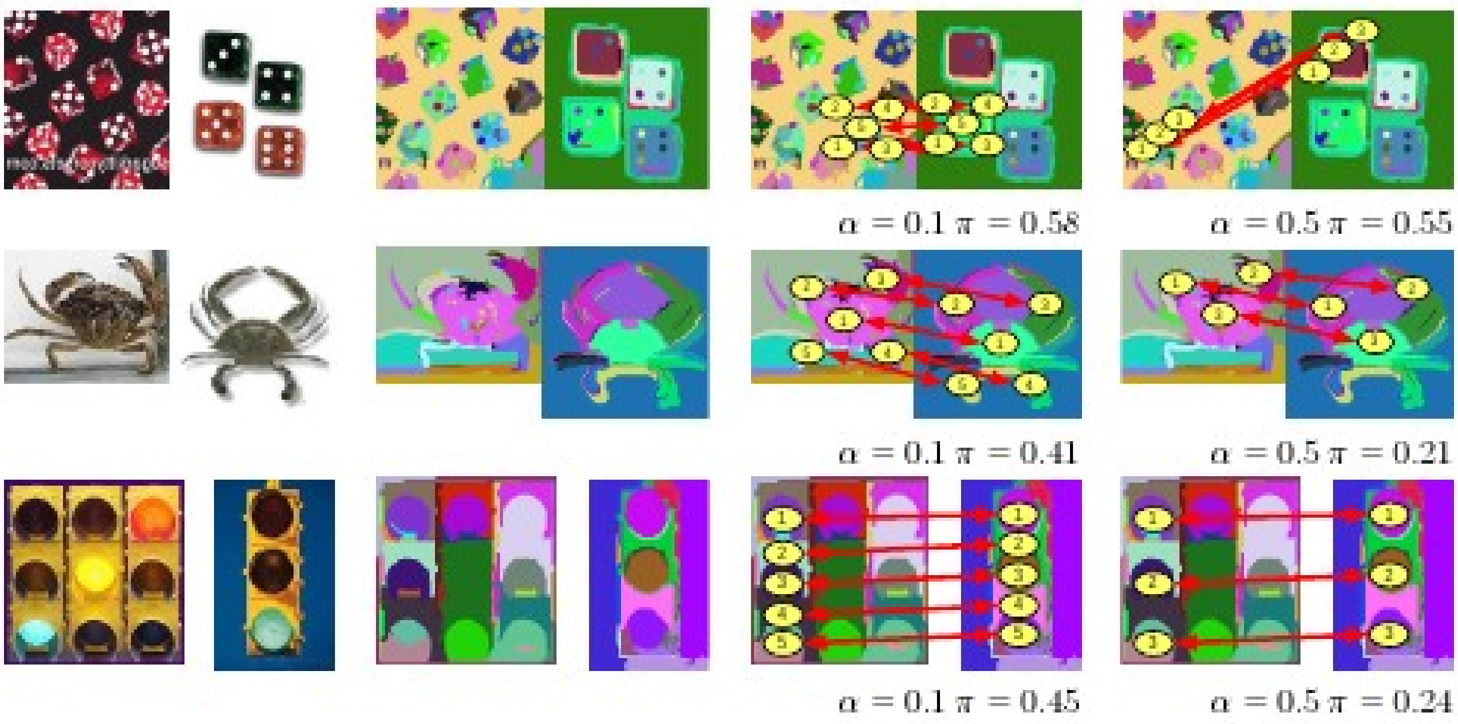

Effect of selectivity parameter α

A. Albarelli, S. Rota-Bulò, A. Torsello, and M. Pelillo; ICCV 2009

GT Matcher and Sparsity

The game-theoretic matcher deviates from the quadratic assignment tradition in that it is very selective: it limits to a cohesive set of association even if feasible associations might still be available

The matcher is tuned towards low false positives rather than low false negatives such as quadratic assignment

Quadratic assignment is greedy while the game theoretic matcher favours sparsity in the solutions

Matching and Inlier selection

There is a domain in which this property is particularly useful: Inlier selection

When estimating a transformation acting on some data, we often need to find correspondences between observations before and after the transformation

Inlier selection is the process of selecting correspondences that are consistent with a single global transformation to be estimated even in the presence of several outlier observations

Examples of problems include surface registration or point-feature matching

Matching and Inlier selection

Typical matching strategies are based on random selection (RANSAC) or the use of local information such as feature descriptors.

Global coherence checks are only introduced after a first estimation (filtering)

Filtering approaches are not very robust w.r.t. outliers (or structured noise)

The game theoretic approach drives the selection of correspondences that satisfy a global compatibility criterion

Estimation of Similarity Transformation

Input images SIFT features Lowe (RANSAC) GT matcher

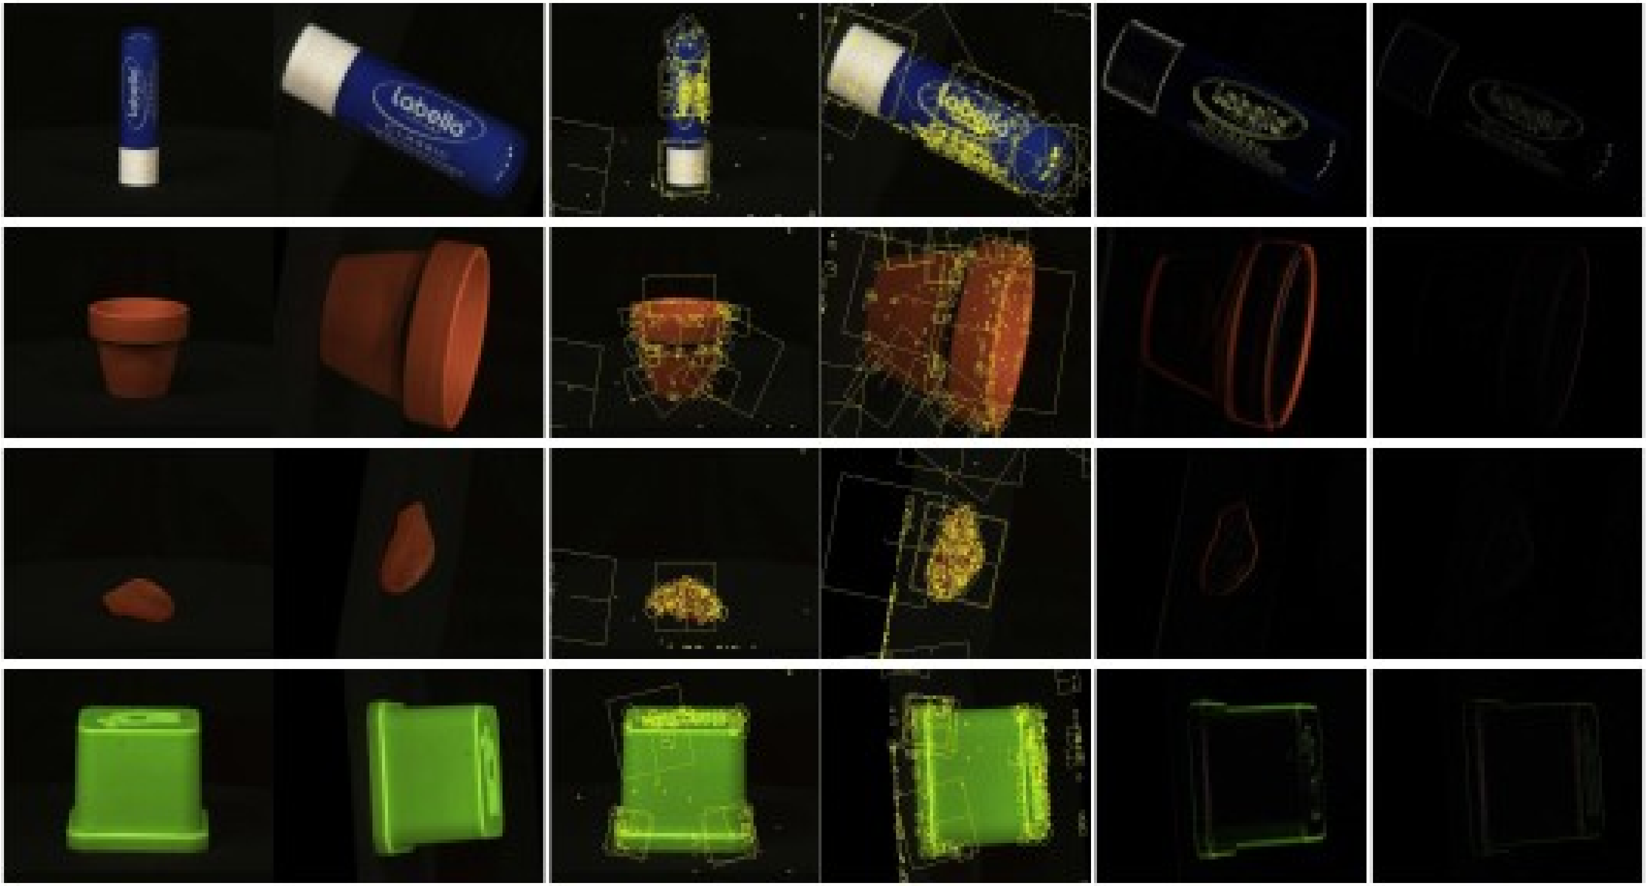

A. Albarelli, S. Rota-Bulò, A. Torsello, and M. Pelillo; ICCV 2009

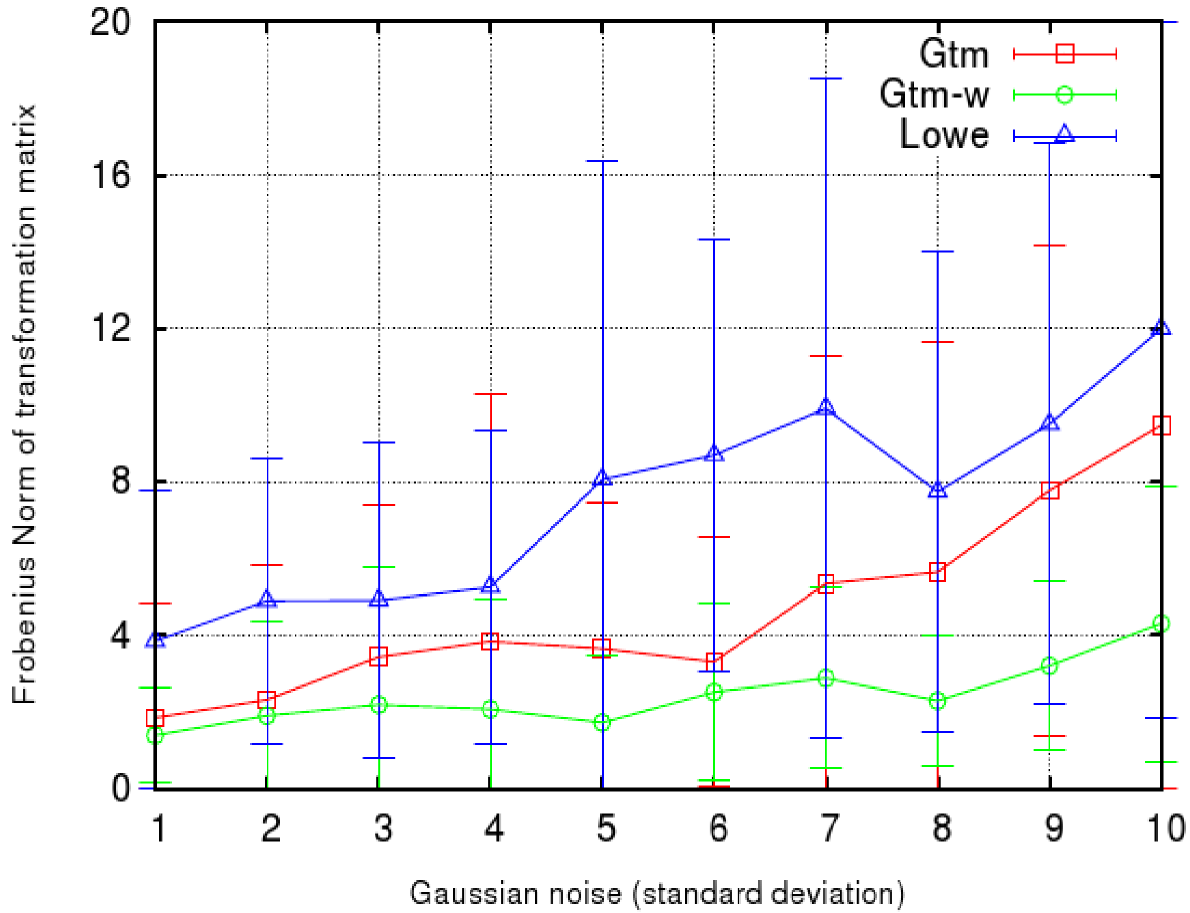

A. Albarelli, S. Rota-Bulò, A. Torsello, and M. Pelillo; ICCV 2009

Surface Registration

Descriptors are used just to reduce the set of feasible associations A

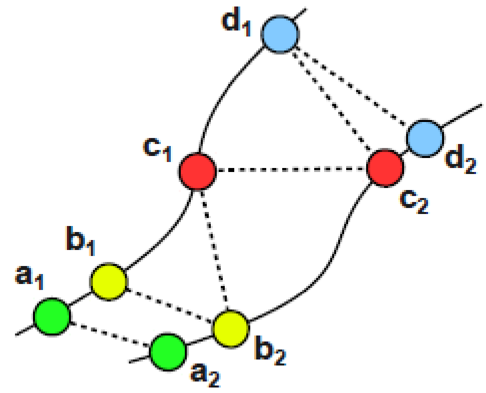

Π

	a_1a_2	b_1b_2	c_1b_2	c_1c_2	d_1c_2	d_1d_2
a_1a_2	0	1	0.1	0.1	0.7	0.9
b_1b_2	1	0	0	0.1	0.7	0.9
c_1b_2	0.1	0	0	0	0.6	0.4
c_1c_2	0.1	0.1	0	0	0	0.1
d_1c_2	0.7	0.7	0.6	0	0	0
d_1d_2	0.9	0.9	0.4	0.1	0	0

Compatibilities are related to rigidity constraints
(difference in distances between corresponding points)

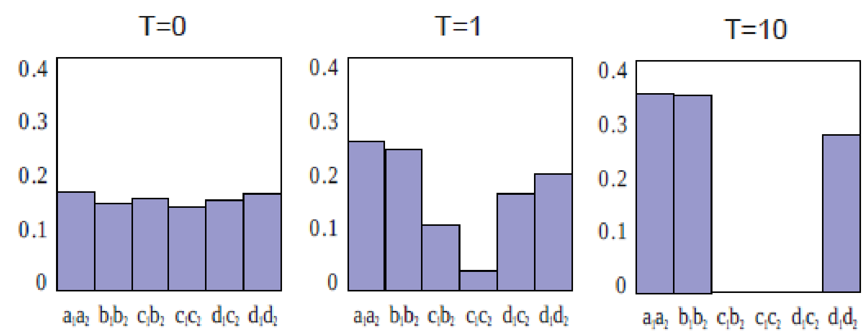

$$\pi((a_1, b_1), (a_2, b_2)) = \frac{min(|a_1 - a_2|, |b_1 - b_2|)}{max(|a_1 - a_2|, |b_1 - b_2|)}$$

Evolve using the replicator dynamics

$$\mathbf{x}_i(t+1) = \mathbf{x}_i(t) \frac{(C\mathbf{x}(t))_i}{\mathbf{x}(t)^T C \mathbf{x}(t)}$$

Surface Registration

DARCES Spin Images GT matcher

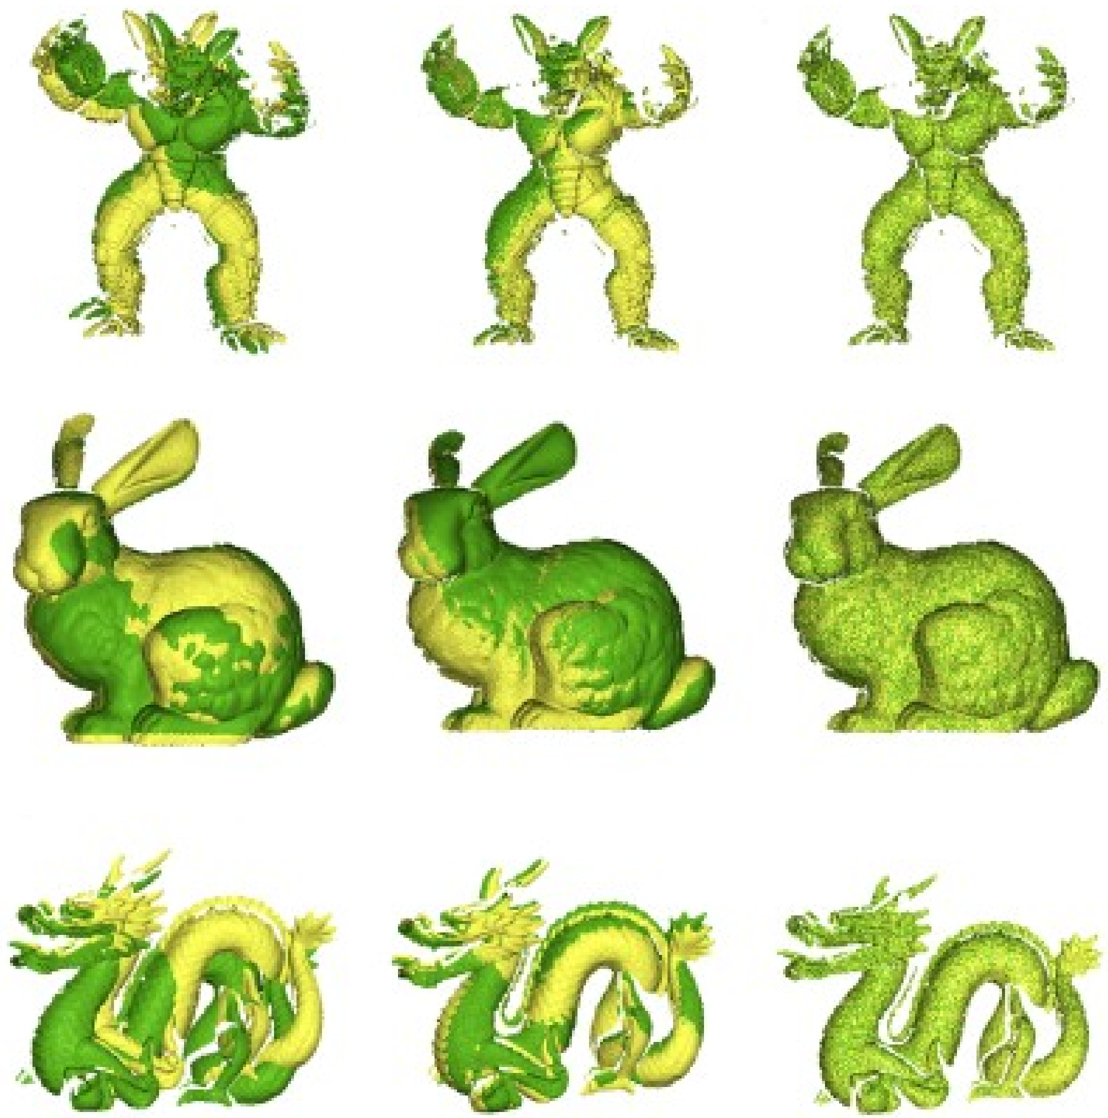

A. Albarelli, E.Rodolà, and A. Torsello; CVPR 2010

Surface Registration

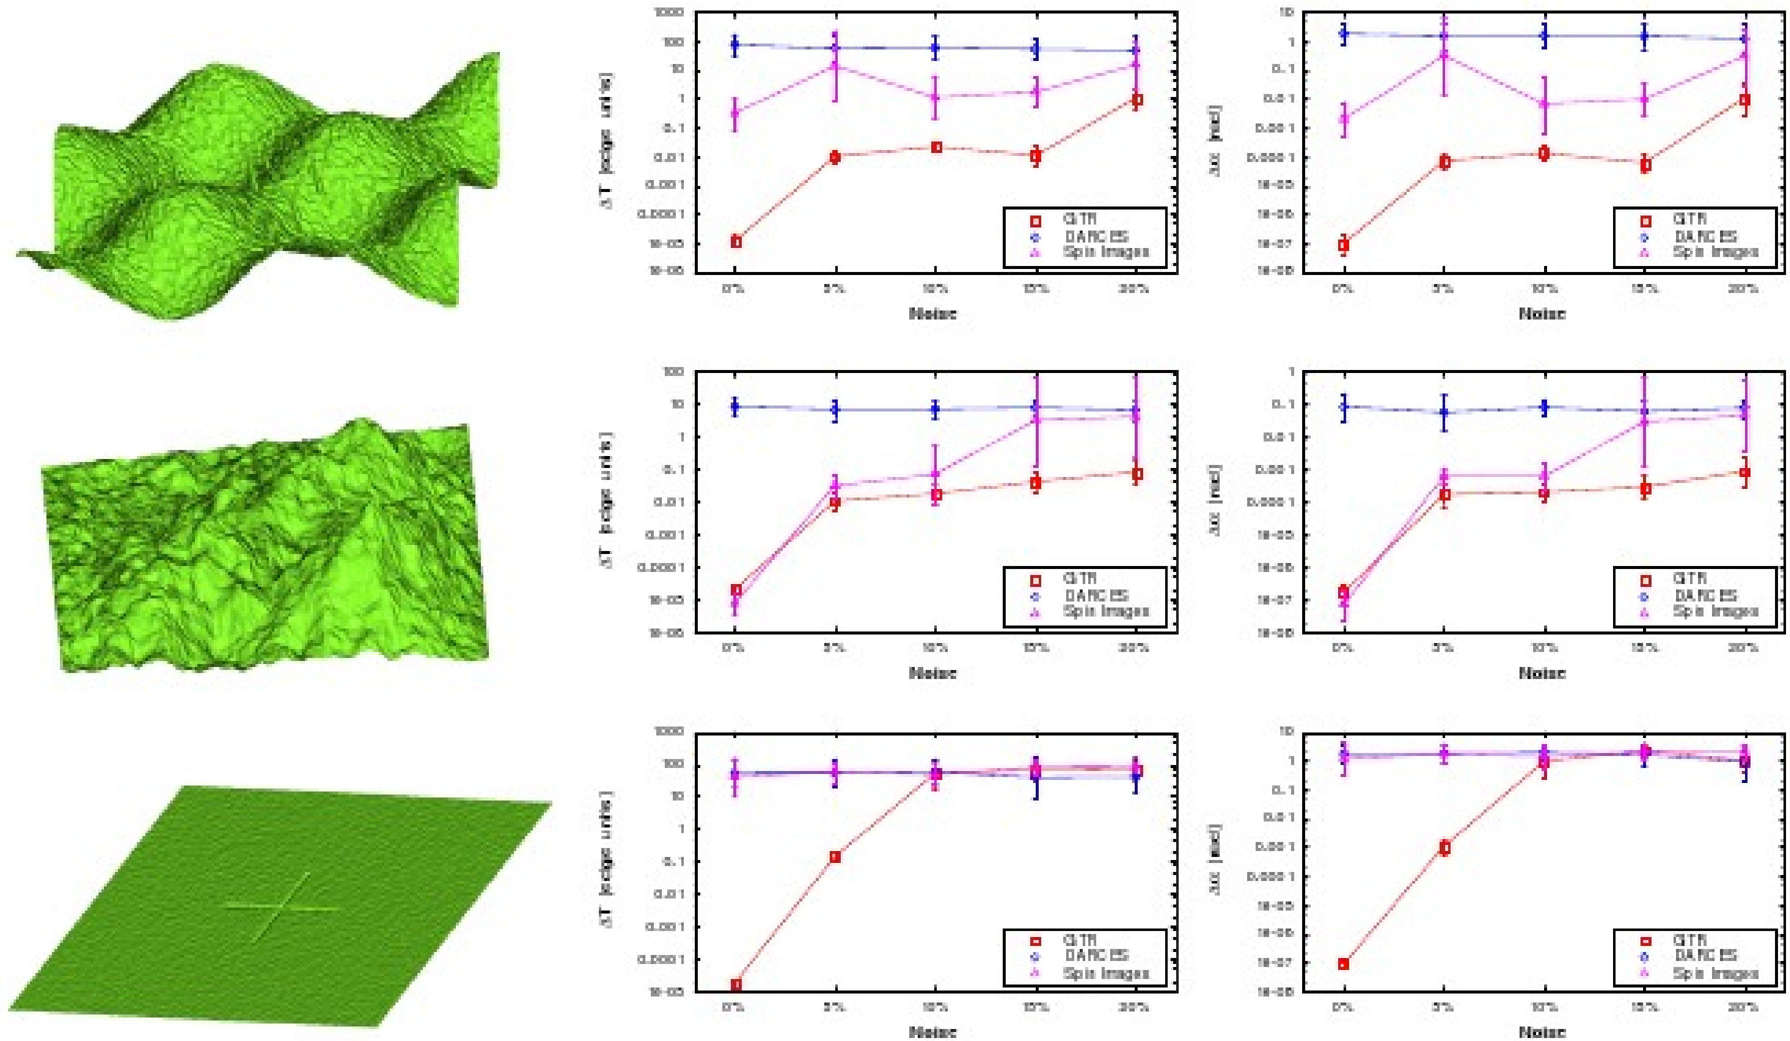

A. Albarelli, E.Rodolà, and A. Torsello; CVPR 2010

Surface Registration

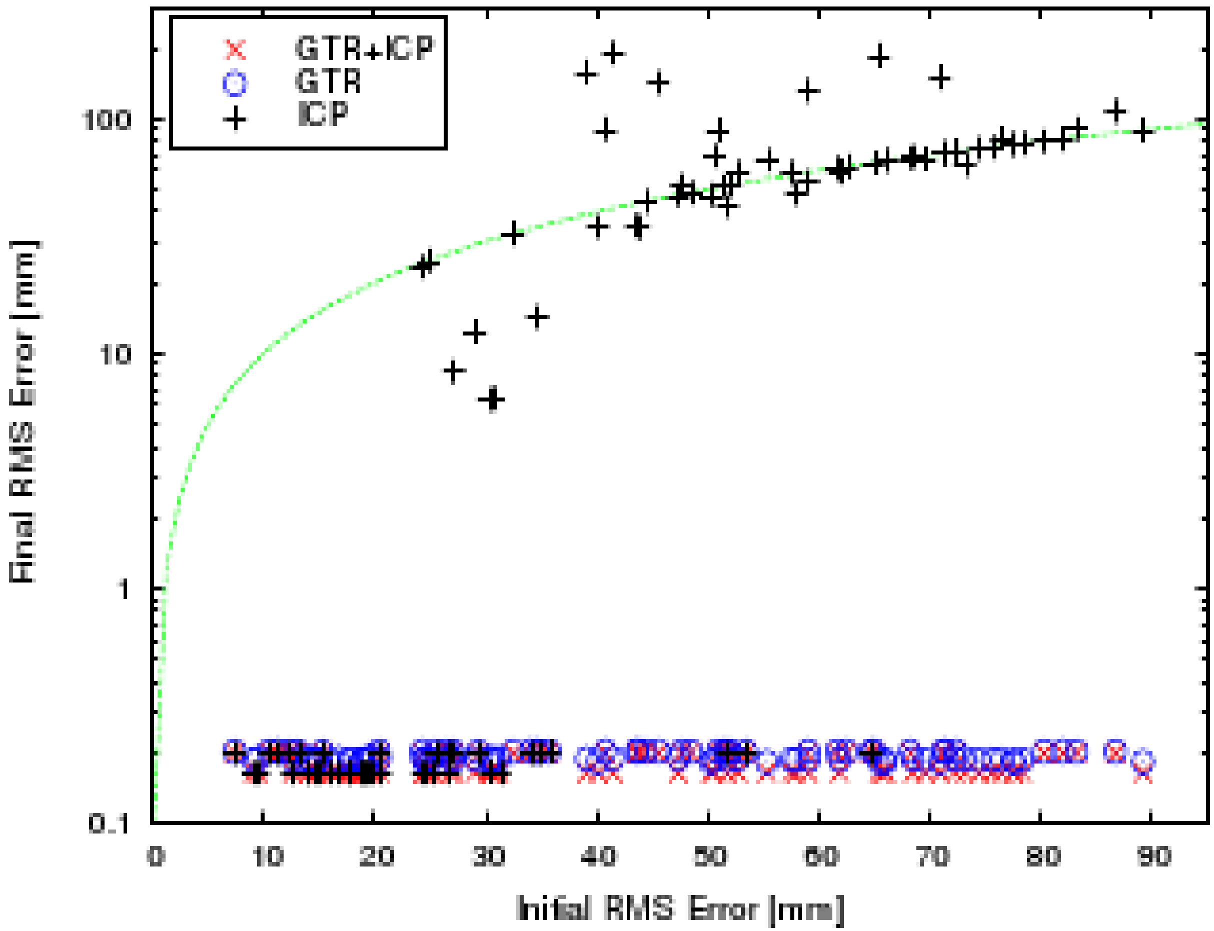

A. Albarelli, E.Rodolà, and A. Torsello; CVPR 2010

Point-Matching for Multi-View Bundle Adjustment

Define (local) compatibility between candidate correspondences through a weak (affine) camera model

We use the orientation and scale information in the feature descriptors to infer an affine transformation between the corresponding features Correspondence imply transformation

Two correspondences are compatible if they define similar transformations

$$\Pi((a_1, a_2), (b_1, b_2)) = e^{-\lambda \max(|a_2 - a_2'|, |b_2 - b_2'|)}$$

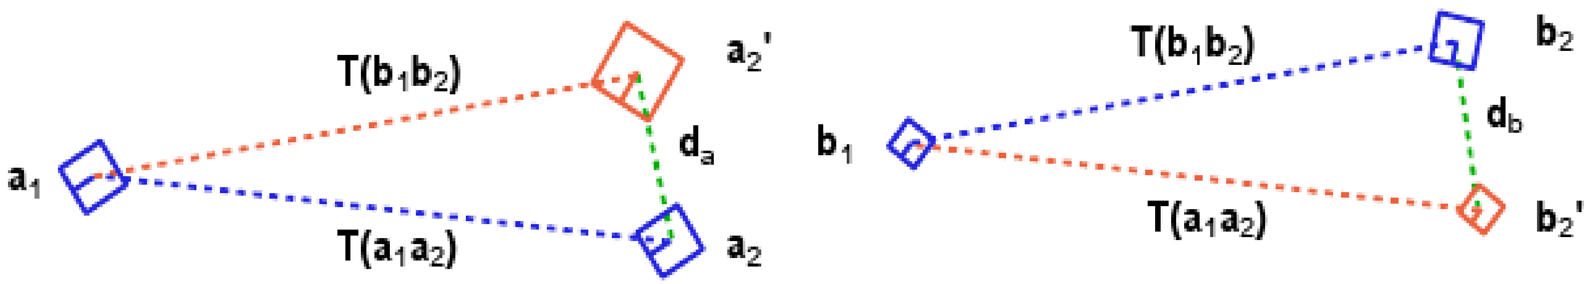

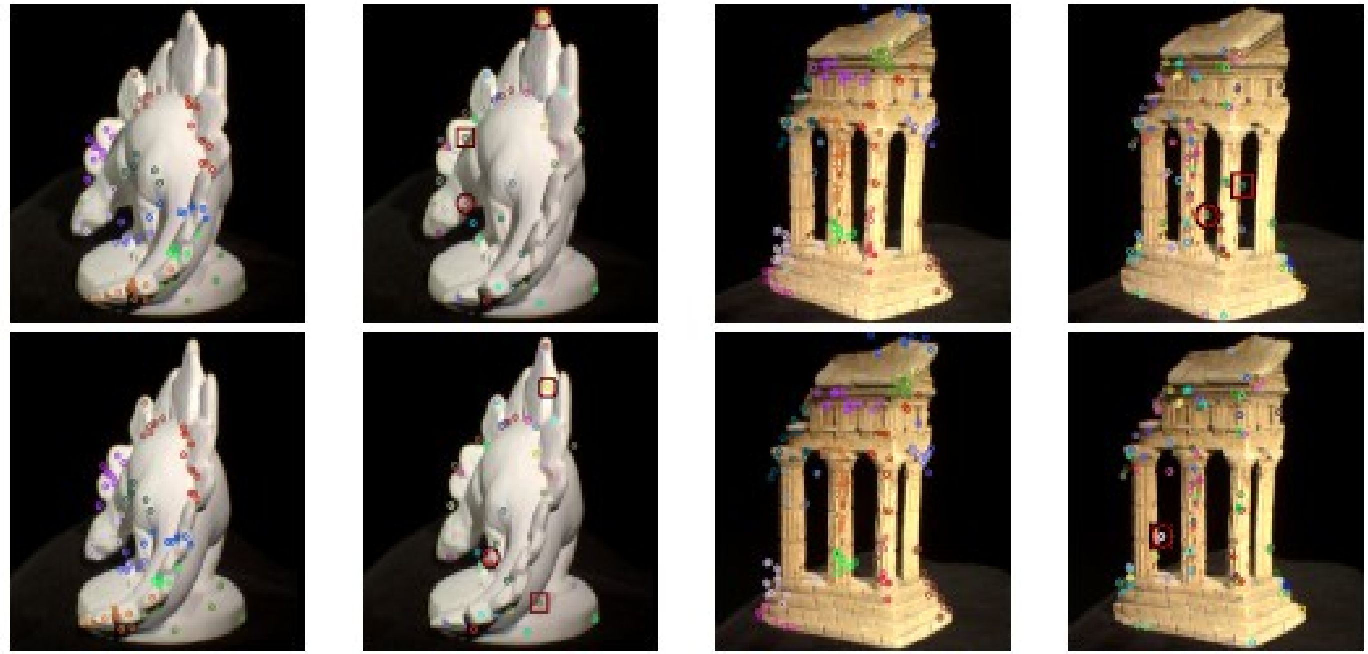

Dino sequence | Temple sequence

		Game-Theoretic	Bundler Keymatcher	Game-Theoretic	Bundler Keymatcher
Matches		14573	9245	25785	22317
ε	≤ 1 pix	24.83	6.49406	22.6049	24.6729
	≤ 5 pix	54.94	48.3659	62.7737	61.8957
	≥ 5 pix	20.21	45.1401	14.6214	13.4314
	Avg.	2.3086	4.5255	2.3577	2.3732
$\Delta\gamma$	Avg.	0.008313	0.009561	0.014050	0.014079
	S. dev.	0.002948	0.006738	0.000511	0.000825
	Max	0.013449	0.030661	0.014692	0.015442
Avg. levels		8.42	-	9.27	-

A. Albarelli, E.Rodolà, and A. Torsello; 3DPVT 2010

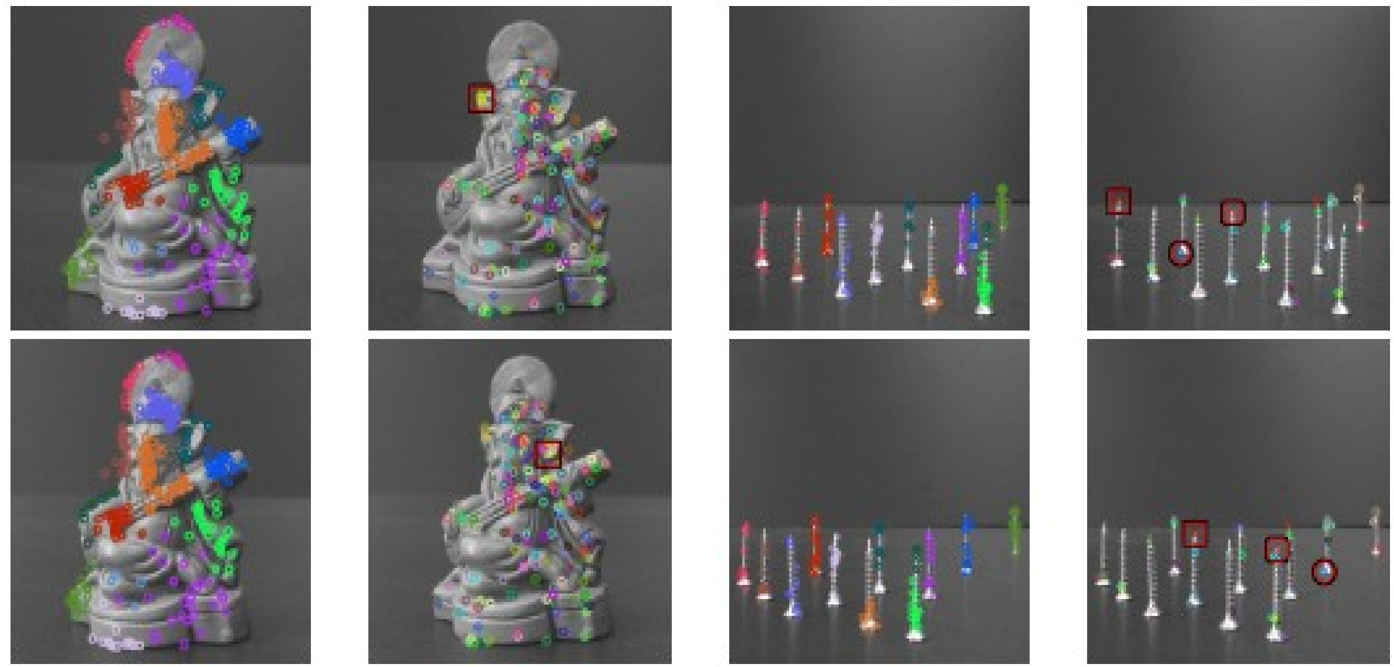

Ganesha stereo   Screws stereo

		Game-Theoretic	Bundler Keymatcher	Game-Theoretic	Bundler Keymatcher
Matches		280	200	211	46
ε	≤ 1 pix	98.2824	20	0	0
	≤ 5 pix	1.7175	80	34.7716	6.75676
	≥ 5 pix	0	0	65.2284	93.2432
	Avg.	0.321248	1.67583	5.86237	10.2208
$\Delta\alpha$		0.001014	0.007424	0.020822	0.030995
$\Delta\gamma$		0.048076	0.078715	0.106485	0.117885
Levels		14	-	12	-

A. Albarelli, E.Rodolà, and A. Torsello; 3DPVT 2010

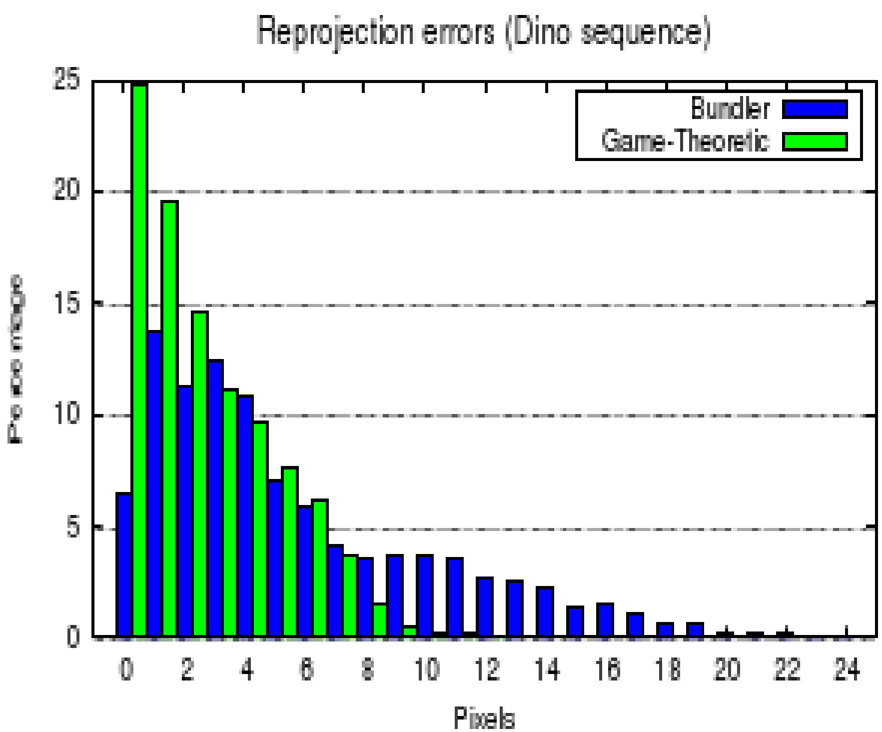

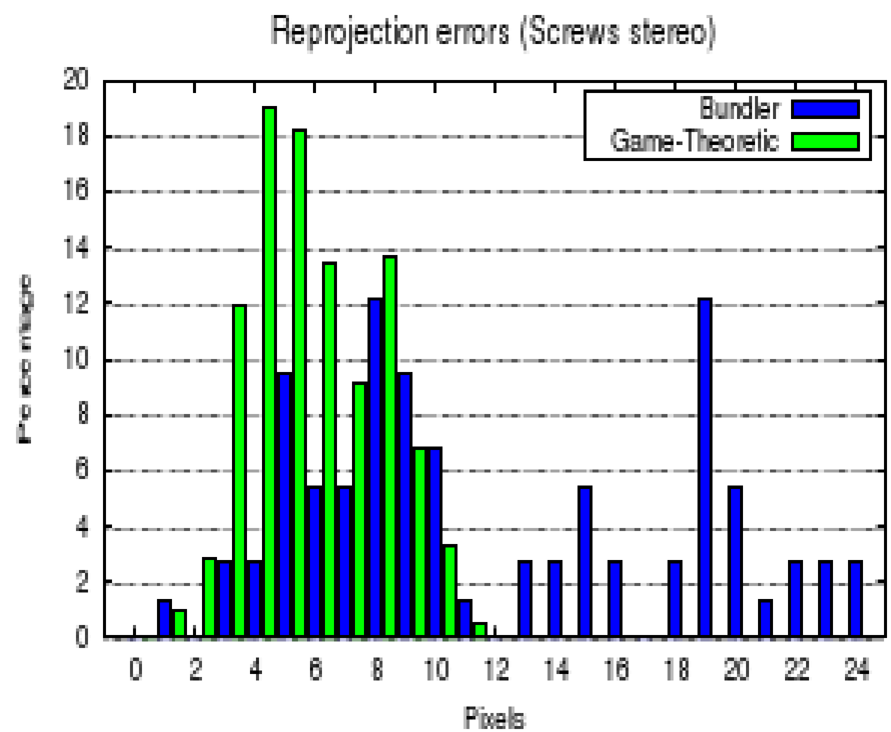

A. Albarelli, E.Rodolà, and A. Torsello; 3DPVT 2010

Conclusions

Presented a game theoretic approach to generic matching and inlier selection problems

Only highly compatible matches are enforced while incompatible correspondences are driven to extinction.

Robustness is achieved by driving the selection enforcing global consistencies in a pairwise setting

Experimental comparisons show the ability of our approach to obtain very accurate estimates in hard estimation problems with many outliers

References

A. Albarelli, S. Rota Bulò , A. Torsello, and M. Pelillo. Matching as a Non-Cooperative Game. ICCV 2009.

A. Albarelli, E. Rodolà, and A. Torsello. Robust Game-Theoretic Inlier Selection for Bundle Adjustment. 3DPVT 2010.

A. Albarelli, E. Rodolà, and A. Torsello. A Game-Theoretic Approach to Fine Surface Registration without Initial Motion Estimation. CVPR 2010.