Combinatorial Optimization in Computer Vision

Chapter 4: Integer Linear Programming

WS 2011/12

Ulrich Schlickewei
Computer Vision and Pattern Recognition Group
Technische Universität München

Plan for Today

Definition of Integer Linear Programs

LP Relaxation and Integral Polyhedra

Min-Cut Max-Flow as ILPs

3D Shape Matching via Integer Linear Programming

Definition of Integer Linear Program

Definition: An Integer Linear Program (ILP) is an optimization problem of the form

$$\begin{aligned} \min_x \quad & p^\top \cdot x \\ \text{s.t.} \quad & Ax = b \\ & x \geq 0 \\ & x \in \mathbb{Z}^n. \end{aligned}$$

Remark: Note that the only difference to LPs is by the additional constraint $x \in \mathbb{Z}^n$. While LPs are solvable in polynomial time, the general ILP is NP-hard.

Example: Linear Assignment Problem

Input: Set $S = \{s_1, \ldots, s_n\}, \;\; T = \{t_1, \ldots, t_n\}$ cost matrix C, with c_{ij} specifying the cost for transporting s_i to t_j

Wanted: Transport plan, that is a bijection $f : \{1, \ldots, n\} \rightarrow \{1, \ldots n\}$ which minimizes the cost $\sum_i c_{if(i)}$

ILP formulation:

$$\begin{aligned}
\min_x \quad & \sum_{i,j} c_{ij} x_{ij} \\
\text{s.t.} \quad & \sum_i x_{ij} = 1 \; \forall i \\
& \sum_j x_{ij} = 1 \; \forall j \\
& x \in \{0,1\}^{n^2}.
\end{aligned}$$

Plan for Today

Definition of Integer Linear Programs

LP Relaxation and Integral Polyhedra

Min-Cut Max-Flow as ILPs

3D Shape Matching via Integer Linear Programming

LP relaxation of ILP

The LP relaxation of an ILP

$$\begin{aligned} \min_x \quad & p^\top \cdot x \\ \text{s.t.} \quad & Ax = b \\ & x \geq 0 \\ & \cancel{x \in \mathbb{Z}^n.} \end{aligned}$$

is defined by dropping the integrality constraint.

The so-resulting LP is solvable in polynomial time and gives a lower bound for the ILP.

When is the LP relaxation equivalent?

Definition: Let $M = \{x \in \mathbb{R}^n \mid Ax = b,\ x \geq 0\}$ be a polyhedron.

- Then M is rational if A and b are rational.
- M is integral if all its vertices lie in $\mathbb{Z}^n$.

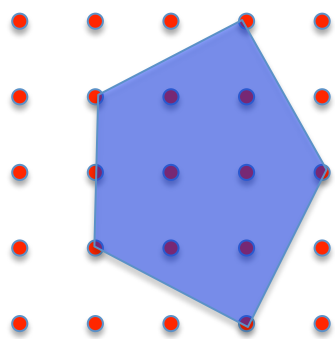

If a polyhedron is integral, then solving an ILP over M is equivalent to solving a LP over M.

Unimodular Matrices

Definition: Let $A \in \mathbb{Z}^{m \cdot n}$ be an integer-valued matrix.

- Assume that A has full row rank. Then A is unimodular if every non-singular $m \times m$ submatrix of A has determinant equal to 1 or -1.
- A is totally unimodular if every quadratic submatrix of A has determinant equal to 0, 1 or -1.

This very technical condition turns out to be initmately related to integral polyhedra. The condition on the determinant ensures by Cramer‘s rule that matrices are invertible over $\mathbb{Z}$.

Unimodular Matrices and Integer Polyhedra

Theorem: Let $A \in \mathbb{Z}^{mn}$ be an integer-valued matrix.

a) Assume that A has full row rank. Then A is unimodular if and only if for all $b \in \mathbb{Z}^n$ the polyhedron $M = \{x \in \mathbb{R}^n \mid Ax = b,\ x \geq 0\}$ is integral.

b) (Hoffmann-Kruskall) The matrix A is totally unimodular if and only of the polyhedron $M = \{x \in \mathbb{R}^n \mid Ax \leq b,\ x \geq 0\}$ is integral.

Note that unimodularity and total unimodularity is more than what is needed for particular ILPs to be solvable efficiently. Typically, one is only interested in the integrality of polyhedra for one sepcific value of $b \in \mathbb{Z}^n$.

Plan for Today

Definition of Integer Linear Programs

LP Relaxation and Integral Polyhedra

Min-Cut Max-Flow as ILPs

3D Shape Matching via Integer Linear Programming

Network Incidence Matrix

Consider a network $N = (V, E, s, t, c)$.

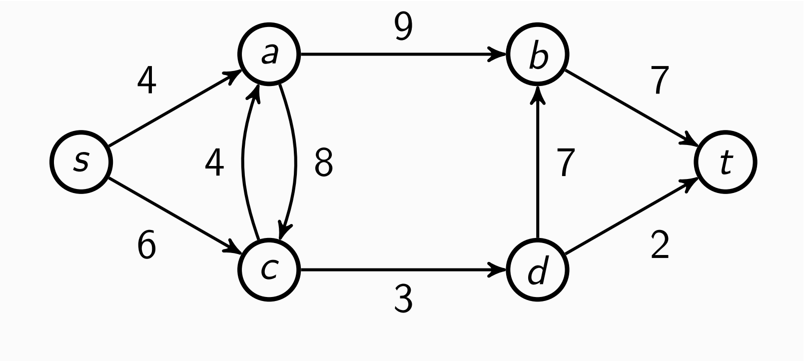

The incidence matrix $A = (a_{ve}) \in \{0, -1, 1\}^{|V| \times |E|}$ is defined by

$$a_{ve} = \begin{cases} 1 & \text{if edge } e \text{ ends at vertex } v \\ -1 & \text{if edge } e \text{ starts at vertex } v \\ 0 & \text{else.} \end{cases}$$

Incidence Matrix is Totally Unimodular

Lemma 1: The incidence matrix A is totally unimodular.

Proof: The Lemma follows from Lemma 2.

Lemma 2: Let $A \in \{0, -1, 1\}^{mn}$ be a matrix. Assume that each column of A contains at most one -1 and one +1. Then A is totally unimodular.

Proof: Induction on the size k of square submatrices B of A.

$k = 1$: okay.

$k > 1$: If there exists a column with at most one non-zero entry, develop the determinant by this column and use the induction hypothesis. Otherwise all columns sum up to zero and the determinant is zero.

Network Flows

Recall that a flow in $N = (V, E, s, t, c)$ is given by a function $f : E \to \mathbb{R}_{\geq 0}$ satisfying

$$f(e) \leq c(e) \quad \forall e \in E$$

$$\sum_{(v',v)\in E} f(v', v) - \sum_{(v,v'')\in E} f(v, v'') = 0 \qquad \forall v \in V.$$

The value of a flow is given by

$$|f| = \sum_{(s,v)\in E} f(s, v).$$

Max-Flow as Linear Program I

Let $w \in \{0,1\}^{|E|}$ be the vector defined by

$$w_e = \begin{cases} 1 & \text{if } e \text{ is incident to } s \\ 0 & \text{else.} \end{cases}$$

(in other words, w indicates outgoing edges from the source.
Note, that $-w$ is the s -th row of the incidence matrix A.)

Let A' be the reduced incidence matrix, defined by removing the rows corresponding to the source s and the sink t .

Max-Flow as Linear Program II

Then the max-flow problem can be formulated as the following Linear Program

$$
\begin{aligned}
\max_{y \in \mathbb{R}^{|E|}} \quad & w^\top \cdot y \\
\text{s.t.} \quad & A' y = 0 \\
& y \leq c \\
& y \geq 0.
\end{aligned}
$$

In order to dualize this LP, we transform it into standard inequality form:

$$
\begin{aligned}
\max_{y \in \mathbb{R}^{|E|}} \quad & w^\top \cdot y \\
\text{s.t.} \quad & \begin{pmatrix} A' \\ -A' \\ \mathbb{1} \end{pmatrix} \cdot y \leq \begin{pmatrix} 0 \\ 0 \\ c \end{pmatrix} \\
& y \geq 0.
\end{aligned}
$$

The Dual LP to Max-Flow I

By LP duality, the dual to the max-flow problem is given by

$$\min_{(u_1,u_2,x)\in\mathbb{R}^{2(|V|-2)+|E|}} \quad c^\top \cdot x$$

$$\text{s.t.} \qquad \begin{pmatrix}(A')^\top & -(A')^\top & \mathbb{1}\end{pmatrix} \cdot \begin{pmatrix}u_1\\u_2\\x\end{pmatrix} \geq w$$

$$\begin{pmatrix}u_1\\u_2\\x\end{pmatrix} \geq 0.$$

The Dual LP to Max-Flow II

This dual LP can be reformulated to the following equivalent LP

$$\begin{aligned} \min_{(u',x) \in \mathbb{R}^{|V|-2+|E|}} \quad & c^\top \cdot x \\ \text{s.t.} \quad & (A')^\top \cdot u' + x \geq w \\ & x \geq 0. \end{aligned}$$

Next, we recall that $-w$ coincides with the s -th row of A . In particular, if we extend u' by 1 in the s -th entry and by 0 in the t -th entry, then we obtain a vector $u \in \mathbb{R}^{|V|}$ with

$$Au + w = A'u'$$

The Dual LP to Max-Flow III

This allows to reformulate the dual LP to max-flow as follows:

$$\begin{aligned} \min_{(u,x)\in\mathbb{R}^{|V|+|E|}} \quad & c^\top \cdot x \\ \text{s.t.} \quad & A^\top \cdot u + x \geq 0 \\ & u_s = 1, \; u_t = 0 \\ & x \geq 0. \end{aligned}$$

We will now see that this dual LP is equivalent to the min-cut problem.

To start with, define for each feasible $(u, x) \in \mathbb{R}^{|V|+|E|}$ an s-t-cut by setting $S = \{v \in V \mid u_v > 0\}, \; T = \{v \in V \mid u_v \leq 0\}$. Then clearly $s \in S, \; t \in T$ and $V = S \sqcup T$, so that (S, T) is an s-t-cut.

The Dual LP to Max-Flow IV

Next, we note that by total unimodularity of A also the matrix $\begin{pmatrix} (A')^\top & -(A')^\top & \mathbb{1} \end{pmatrix}$ is totally unimodular. As a consequence, the polyhedron

$$M' = \left\{ \begin{pmatrix} u_1 \\ u_2 \\ x \end{pmatrix} \middle| \begin{pmatrix} (A')^\top & -(A')^\top & \mathbb{1} \end{pmatrix} \cdot \begin{pmatrix} u_1 \\ u_2 \\ x \end{pmatrix} \geq w, \;\; \begin{pmatrix} u_1 \\ u_2 \\ x \end{pmatrix} \geq 0 \right\}$$

is integral. This implies that the polyhedron

$$M = \left\{ \begin{pmatrix} u \\ x \end{pmatrix} \; \middle| A^\top u + x \geq 0, \;\; u_s = 1, \;\; u_t = 0, \;\; x \geq 0 \right\}$$

is integral.

The Dual LP to Max-Flow V

This implies that the dual LP

$$\begin{aligned}\min_{(u,x)\in\mathbb{R}^{|V|+|E|}} \quad & c^\top \cdot x \\ \text{s.t.} \quad & A^\top \cdot u + x \geq 0 \\ & u_s = 1,\ u_t = 0 \\ & x \geq 0.\end{aligned}$$

has an integral solution vector $\begin{pmatrix} u \\ x \end{pmatrix} \in M \cap \mathbb{Z}^{|V|+|E|}$.

Next, consider the e -th component of the constraint (where $e = (v_1, v_2) \in E$):

$$(A^\top \cdot u + x)_e = u_{v_2} - u_{v_1} + x_e \geq 0$$

The Dual LP to Max-Flow VI

Minimizing with respect to x and taking into account that $c \geq 0$ we see that for fixed u an optimal x is given by

$$x_{(v_1,v_2)} = \max(0, u_{v_1} - u_{v_2}).$$

Thus, we can interpret the dual LP as a labelling problem, where we search for a function $u : V \to \mathbb{Z}$, whose "weighted derivative"

$$\sum_{e=(v_1,v_2)\in E} c(v_1, v_2) \max(0, u(v_1) - u(v_2))$$

is minimal and which satisfies the "boundary condition" $u(s) = 1, u(t) = 0.$

The Dual LP to Max-Flow VII

This labelling problem is equivalent to the min-cut problem: It is certainly most economic to employ only the labels 0 and 1 and to change labels as rarely as possible. Moreover, if only these binary labels are used, for optimal x the cost $c^\top \cdot x$ is the cost of the cut induced by u .

Thus the LP dual to the max-flow problem

$$\begin{aligned} \min_{(u,x)\in\mathbb{R}^{|V|+|E|}} \quad & c^\top \cdot x \\ \text{s.t.} \quad & A^\top \cdot u + x \geq 0 \\ & u_s = 1,\ u_t = 0 \\ & x \geq 0. \end{aligned}$$

is the min-cut problem.

Min-Cut Max-Flow Duality

Theorem:

- The min-cut and the max-flow problems are dual linear programs. In particular, the cost of the minimum cut coincides with the value of the maximum flow.
- The polyhedron corresponding to the LP relaxation of the min-cut problem is integral.

Plan for Today

Definition of Integer Linear Programs

LP Relaxation and Integral Polyhedra

Min-Cut Max-Flow as ILPs

3D Shape Matching via Integer Linear Programming

The 3D Matching Problem

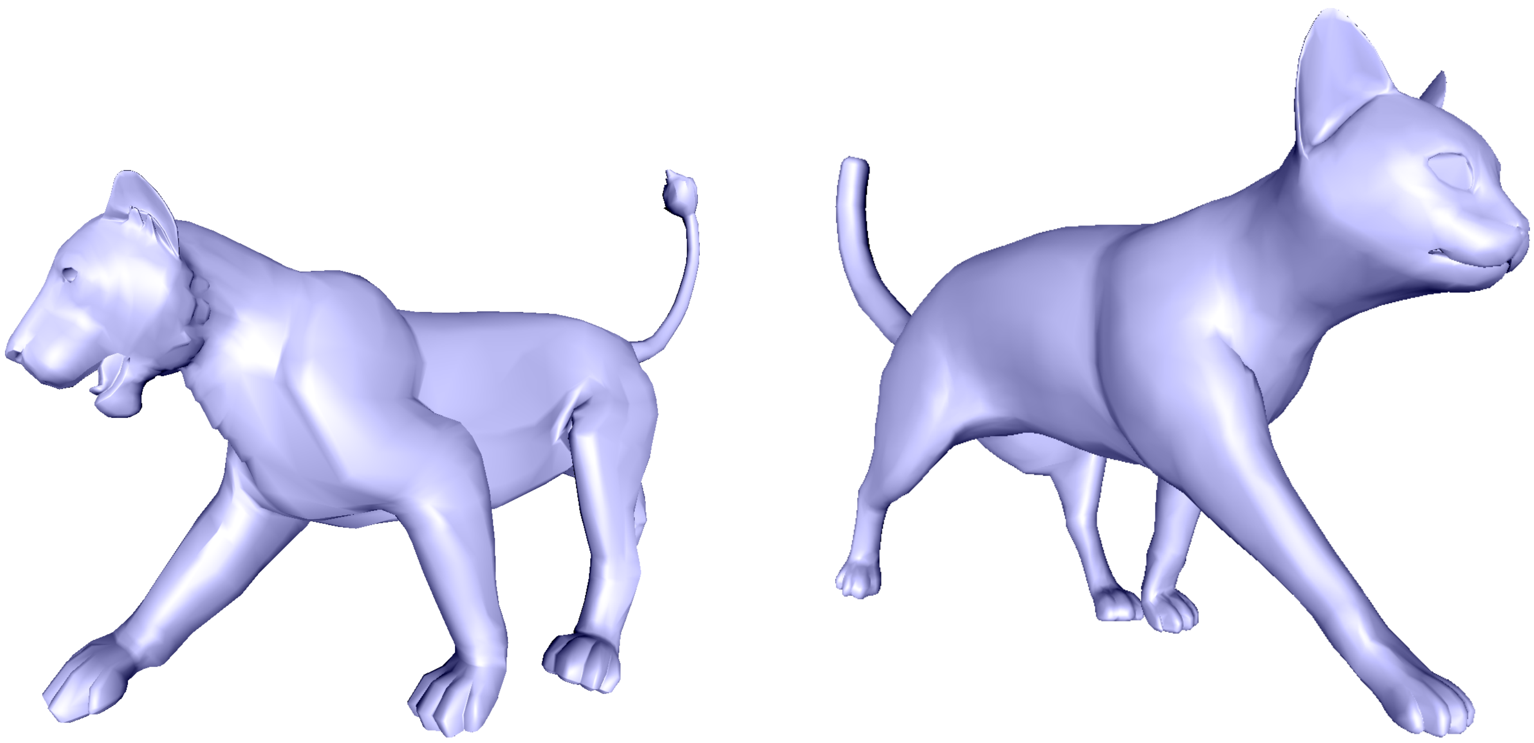

Problem: Find a correspondence between the two shapes, i.e. for each point on the lion find a corresponding point on the cat.

The 3D Matching Problem

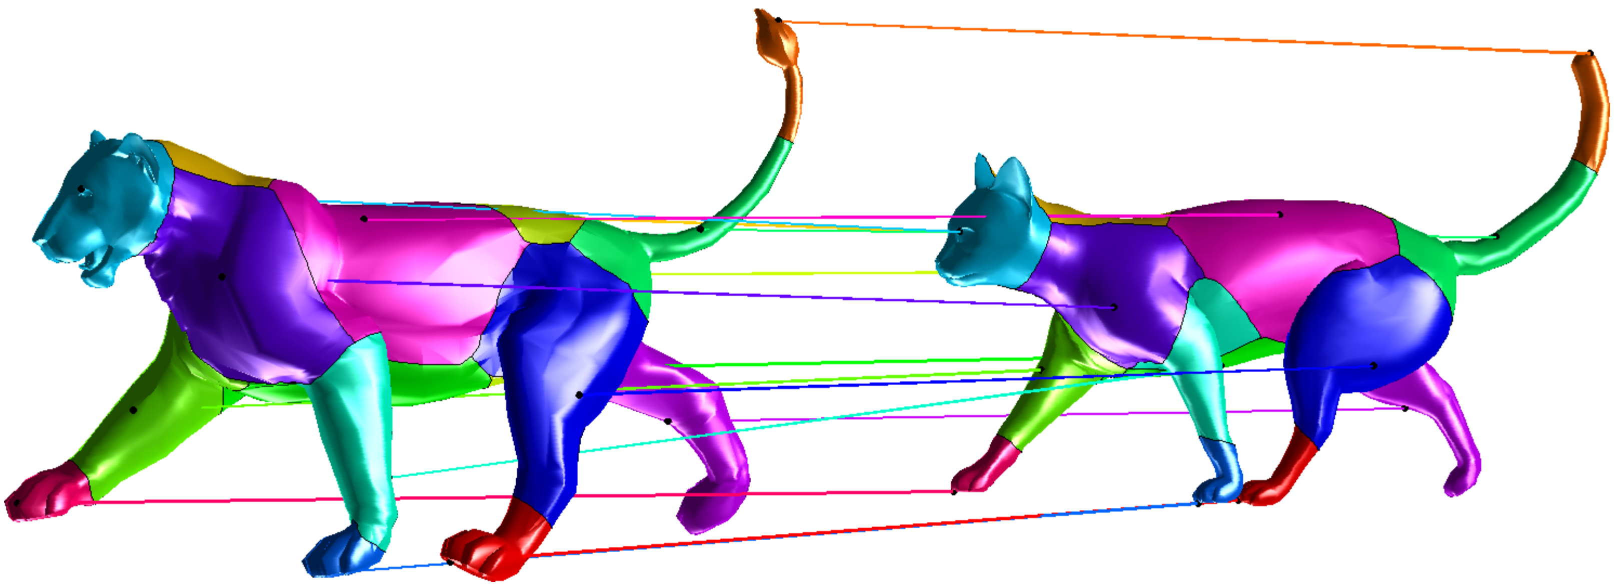

Desired properties:
- Preservation of orientation
- Meaningful correspodence

Mathematical Model I

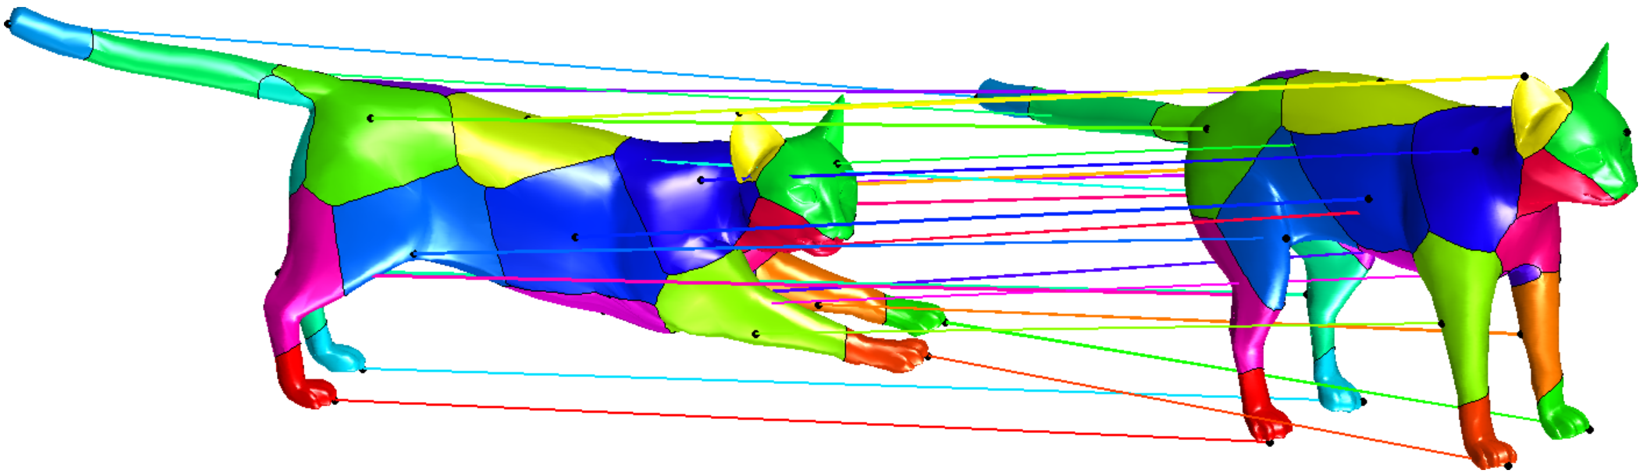

$$\varphi : \quad X \quad \rightarrow \quad Y$$

$$\varphi^* = \operatorname{argmin}_{\varphi \in \mathrm{Diff}^+(X,Y)} E(\varphi) + E(\varphi^{-1})$$

Mathematical Model II

- Mathematically, we search for an optimal orientation preserving diffeomorphism $\varphi : X \to Y$.
- We choose a physical model: We search for the deformation which causes minimal physical work. One can imagine this model as putting a glove over a hand – if all fingers are put in the right place, it costs less energy than if something goes wrong.
- The precise term for this energy, the thin-shell-energy, has been derived by Koiter in the 70's:

$$E(\varphi) = \underbrace{\int_X (\mathrm{tr}_{g_X} \mathbf{E}) + \mu \, \mathrm{tr}_{g_X} (\mathbf{E}^2)}_{E_{\mathrm{mem}}} + \lambda \underbrace{\int_X (H_X(x) - H_Y(\varphi(x))^2}_{E_{\mathrm{bend}}}$$

- Finally, we symmetrize the problem by considering

$$\varphi^* = \mathrm{argmin}_{\varphi \in \mathrm{Diff}^+(X,Y)} E(\varphi) + E(\varphi^{-1})$$

Representation of a Diffeomorphism by its Graph

Definition and Proposition: Let $\varphi : X \to Y$ be a diffeomorphism. The graph of φ is defined as the set $Z := \{(x, \varphi(x)) \mid x \in X\} \subset X \times Y.$ This set has the following properties:

i. Z is a differentiable, connected, closed surface in the product space $X \times Y$.
ii. The natural projections $\pi_X : Z \to X$ and $\pi_Y : Z \to Y$ are both diffeomorphisms.
iii. The two orientations which Z naturally inherits from X and Y coincide.

→ Shape Matching as Minimal Surface Problem

Visulization of Graph

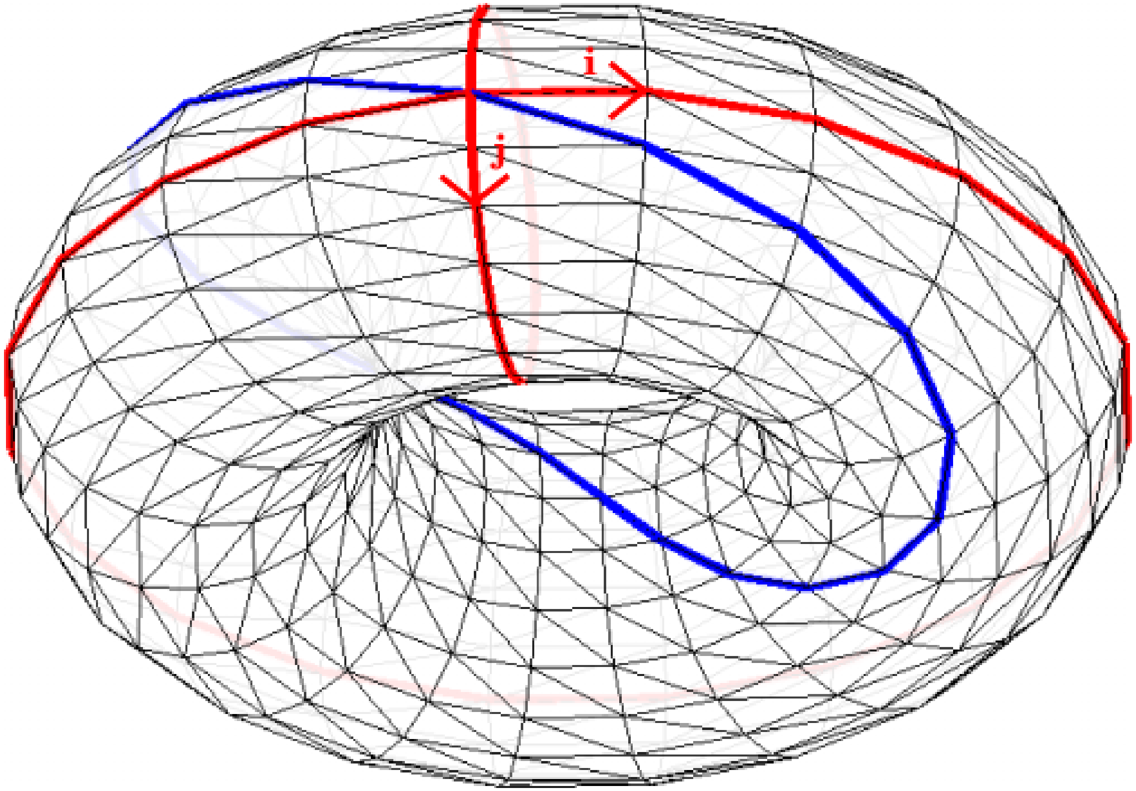

Example: Matching of the two red curves can be formulated as searching for the blue path.

Discrete Setting: Matching of Triangle Meshes

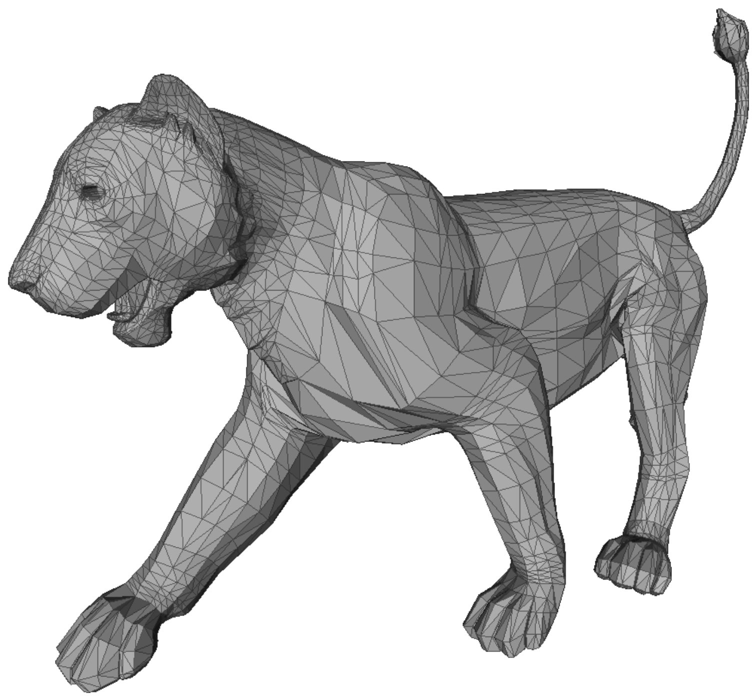

X

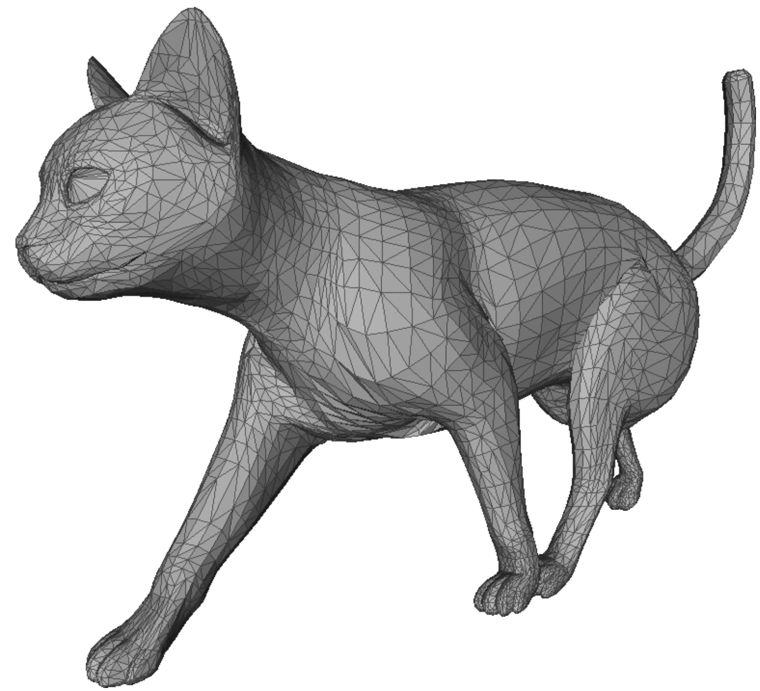

Y

Discrete Setting: Matching of Triangle Meshes

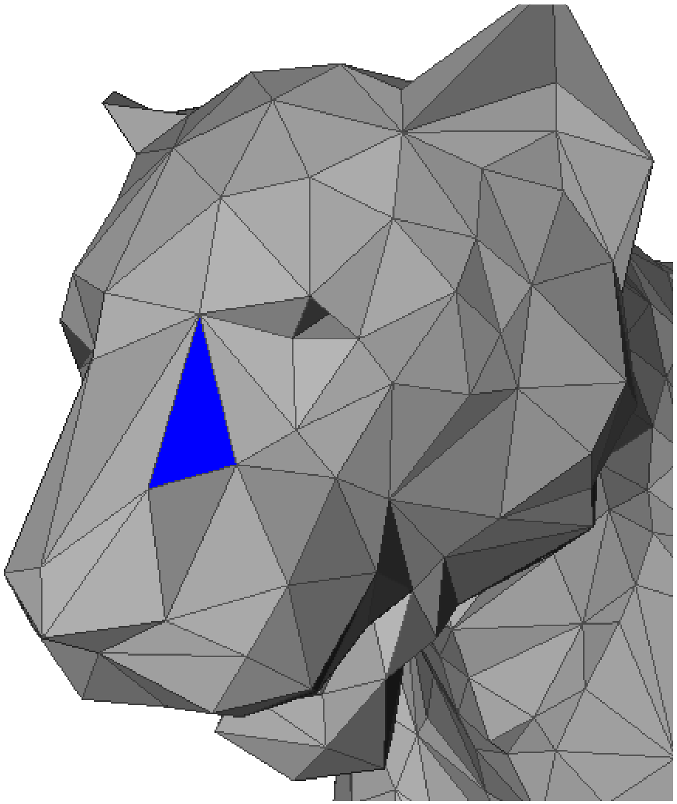

X

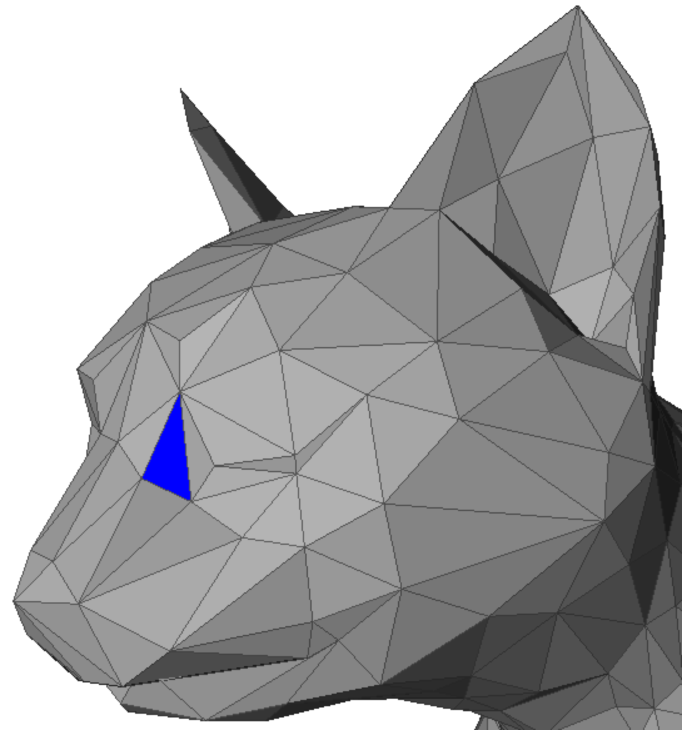

Y

Discrete Setting: Matching of Triangle Meshes

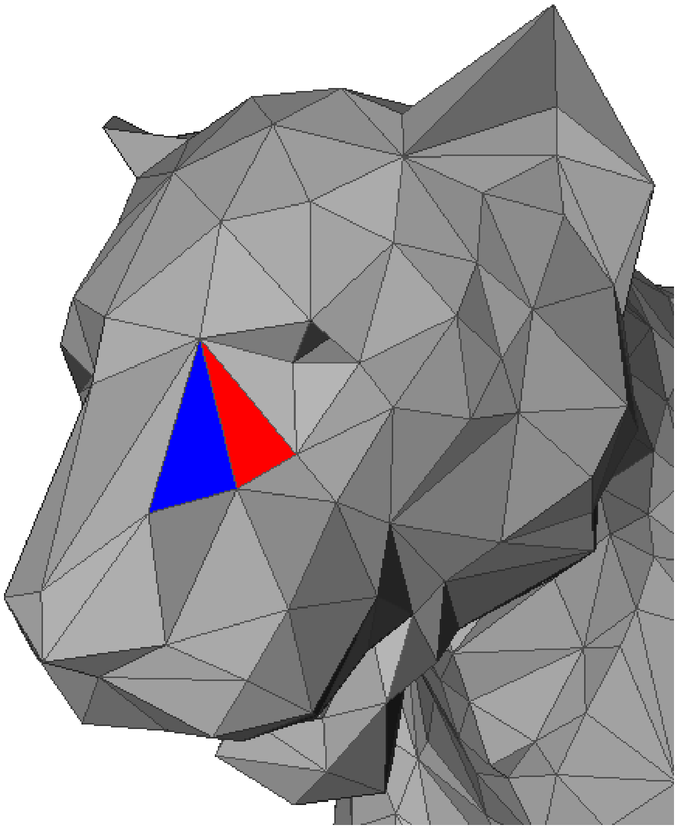

X

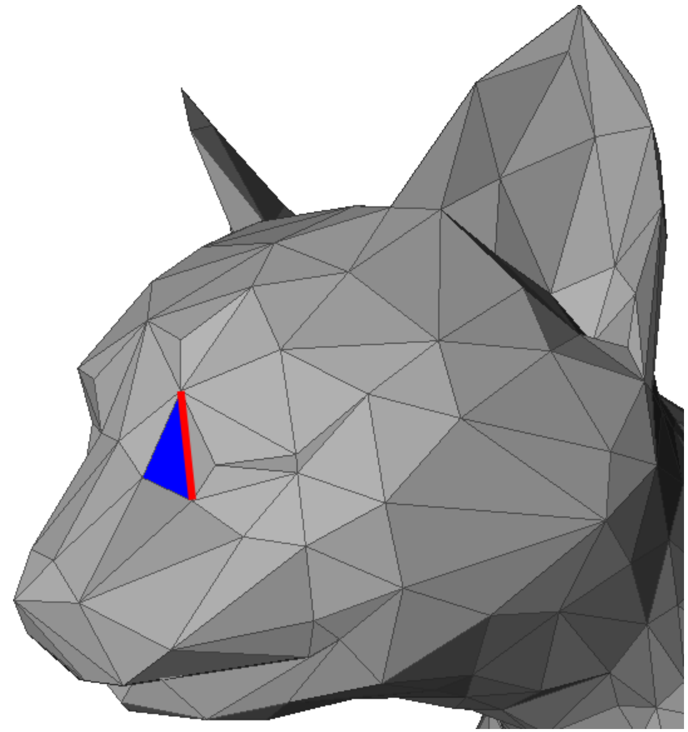

Y

Path to ILP Formulation

We will now discretize graph surfaces. Therefore we will see

- a set F of „product triangles" which are building blocks for discrete graph surfaces (similarly to the edges in the search of the blue path in the torus).
- discrete graph surfaces = indicator vectors $z \in \{0, 1\}^{|F|}$
- discrete versions of constraints i., ii., iii. Then the search space will be the set of subsets of F which satisfy the discrete constraints i., ii., iii.
- discrete energy vector $e \in \mathbb{R}^{|V|}$, which discretizes Koiter's thin-shell energy.

Product Triangles as Optimization Variables

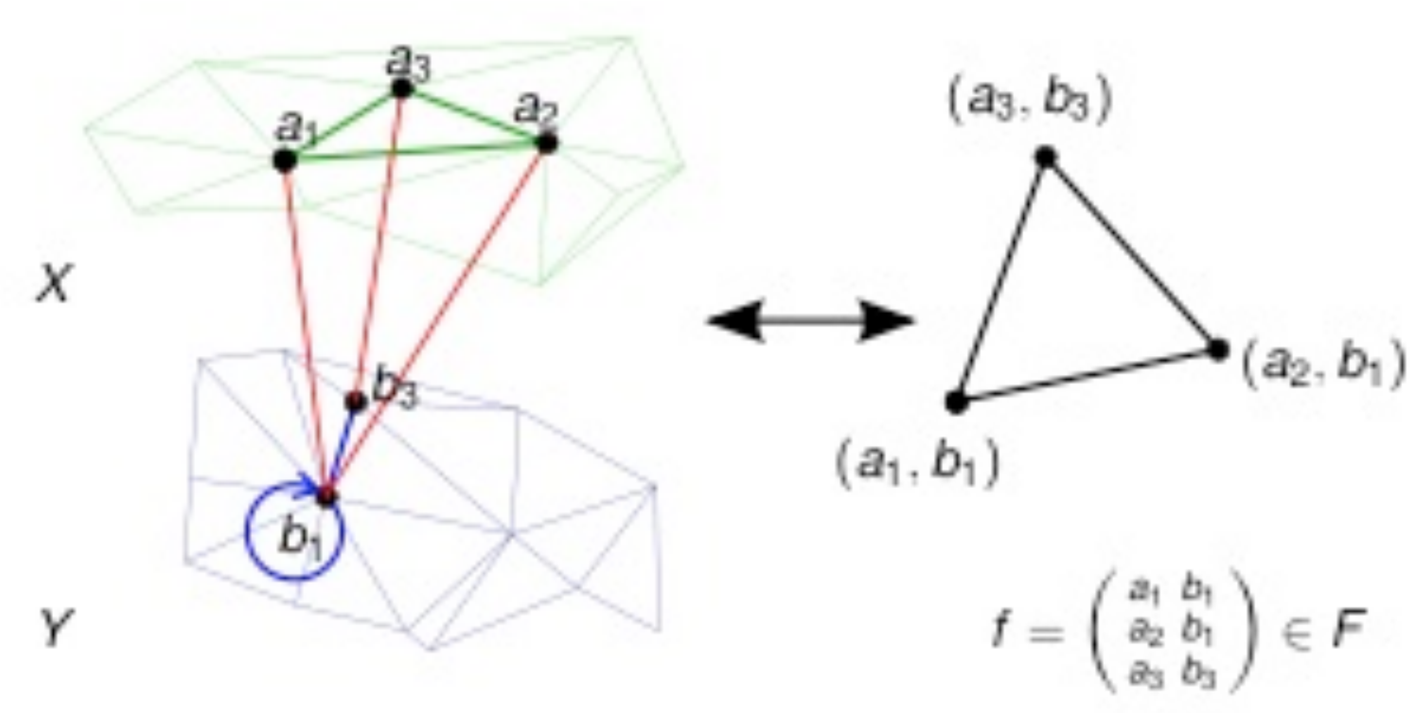

Product triangles as optimization variables. A product triangle f has the interpretation of setting into correspondence vertices a_i with b_i .

Stretching and Bending

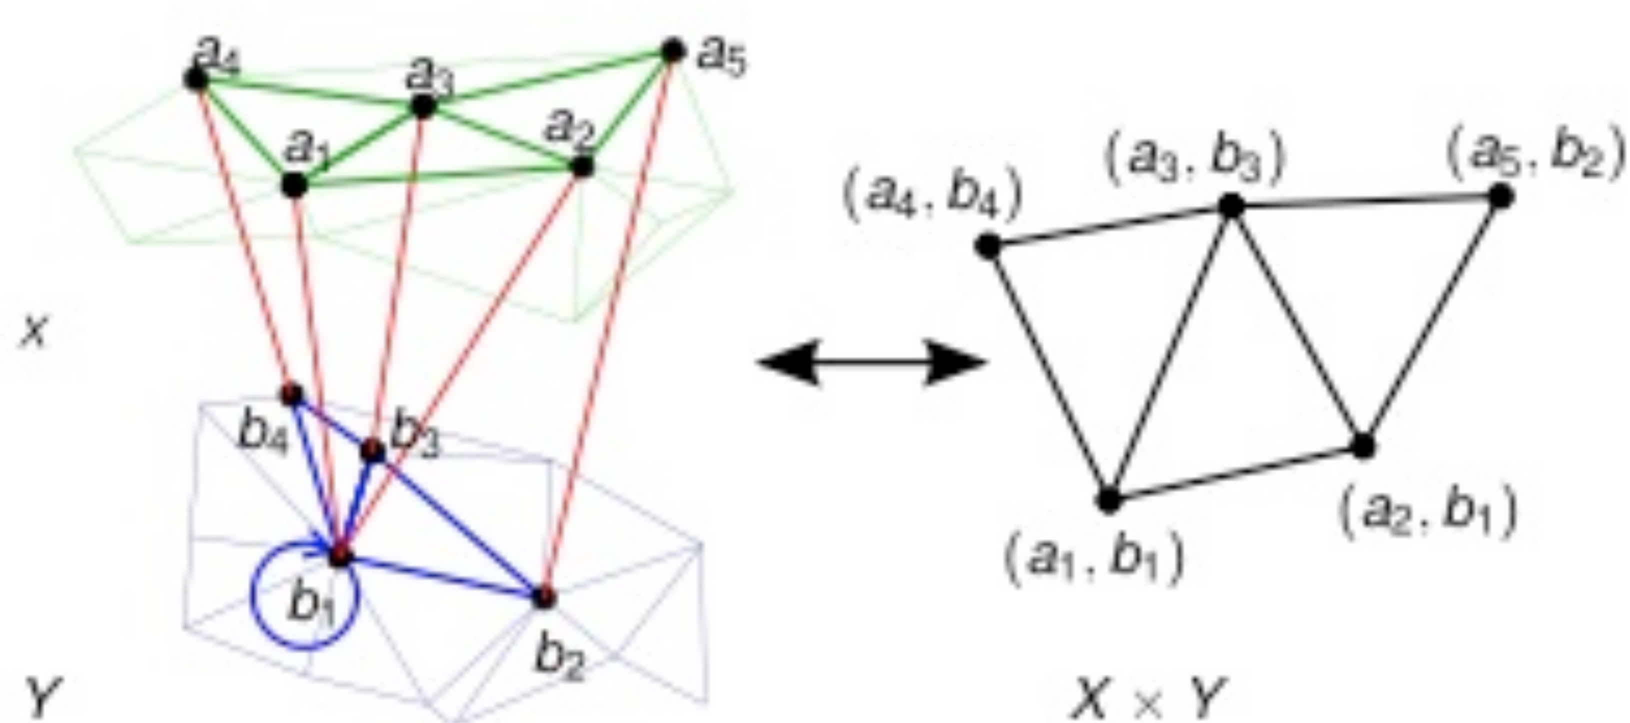

Boundary Operator

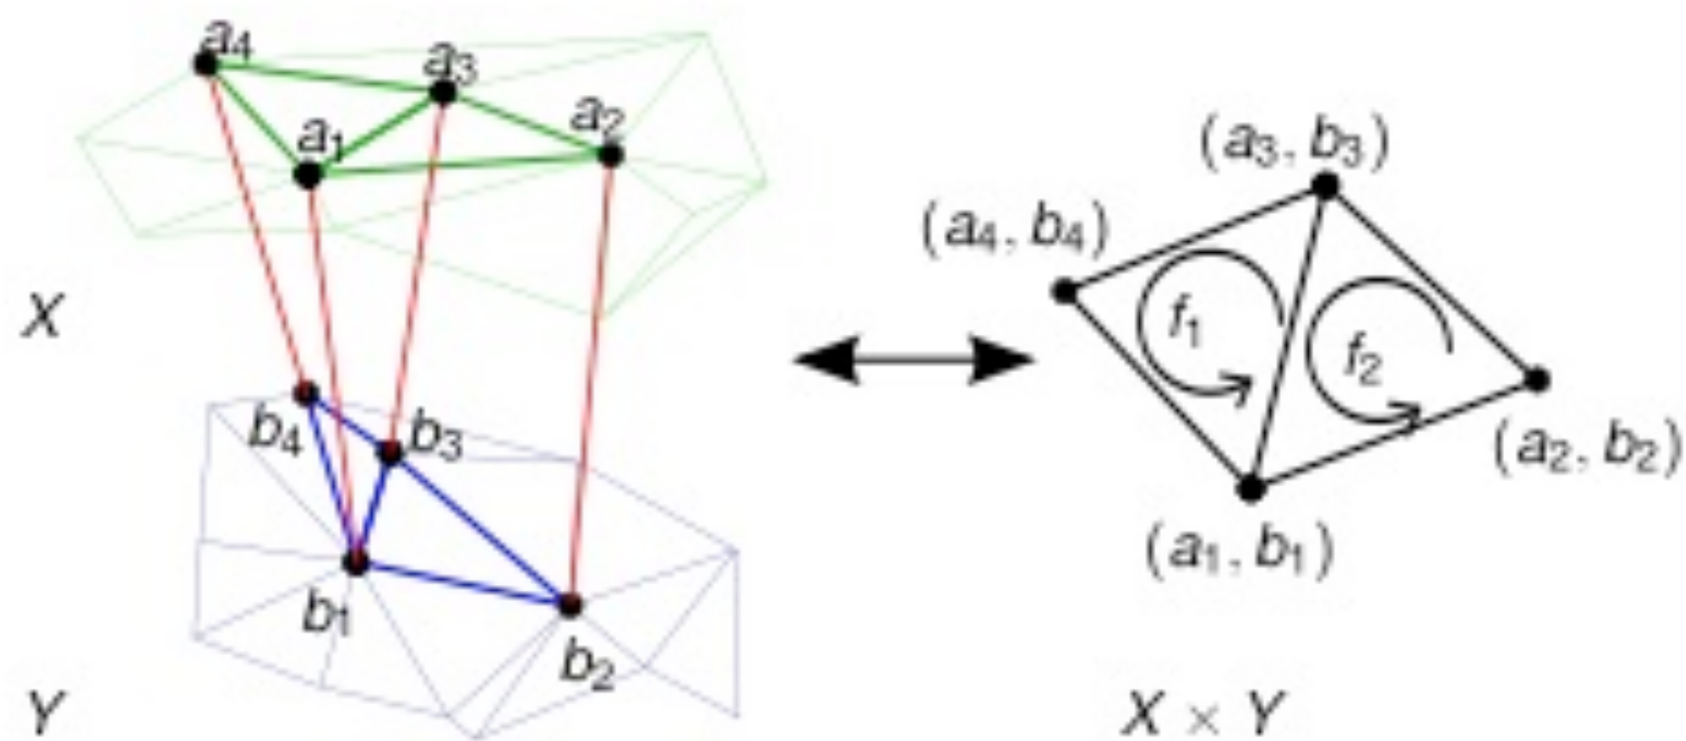

Boundary operator ensures geometric consistency. It is required that each product edge appears as often negatively oriented as it appears positively oriented in a discrete graph surface. The boundary constraint discretizes constraint i.

Projection Constraints and Discrete Energy

- Projection constraints ensure that each triangle on X and on Y is hit exactly once. These discretize constraint ii.
- Constraint iii. is built-in and does not need a discretization.
- Discrete thin-shell energy associates with each product triangle the physical deformation cost. This leads to a vector $e \in \mathbb{R}^{|V|}$

ILP Formulation for 3D Shape Matching

Alltogether, this approach rephrases 3D Shape Matching as the following Integer Linear Program

$$\begin{aligned} \min_{z} \quad & e^\top \cdot z \\ \text{s.t.} \quad & \begin{pmatrix} \partial \\ \pi_X \\ \pi_Y \end{pmatrix} \cdot z = \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix} \\ & \cancel{z \in \{0,1\}^{|F|}} \\ & z \in [0,1]^{|F|} \end{aligned}$$

We solve the relaxed version. Unfortunately, in this case the polyhedron is not integral.

Results

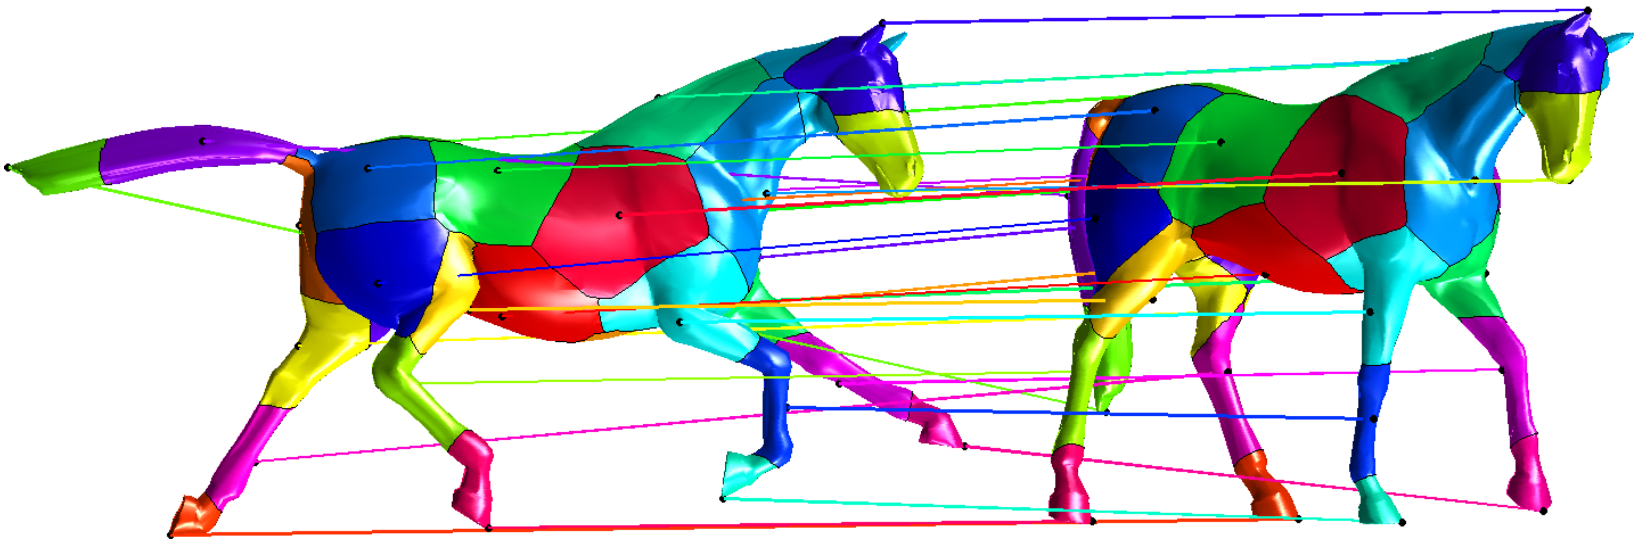

Articulation (Bending)

Results

Partial Matching

Summary

- Integer Linear Programs are Linear Programs with an additional integrality constraint.
- While LPs are solvable in polynomial time, the general ILP is NP-hard.
- If the constraint matrix is totally unimodular, the polyhedron is integral. In this case, LP relaxation is equivalent to ILP.
- The max-flow min-cut duality can be seen as dual LPs. The corresponding constraint matrix is totally unimodular.
- Geometrically consistent 3D shape matching can be casted as ILP.