

Consistently Orienting Facets in Polygon Meshes by Minimizing the Dirichlet Energy of Generalized Winding Numbers

Kenshi Takayama
ETH Zurich

Alec Jacobson
ETH Zurich

Ladislav Kavan
University of Pennsylvania

Olga Sorkine-Hornung
ETH Zurich

Abstract

Jacobson et al. [JKSH13] hypothesized that the local coherency of the generalized winding number function could be used to correctly determine consistent facet orientations in polygon meshes. We report on an approach to consistently orienting facets in polygon meshes by minimizing the Dirichlet energy of generalized winding numbers. While the energy can be concisely formulated and efficiently computed, we found that this approach is fundamentally flawed and is unfortunately not applicable for most handmade meshes shared on popular mesh repositories such as Google 3D Warehouse.

1 Introduction

The robust inside-outside segmentation method based on the generalized winding number proposed by Jacobson et al. [JKSH13] assumes that the polygonal facets in the input mesh are consistently oriented. If this assumption does not hold, which is often the case for hand-made meshes shared on the Internet (e.g., Google 3D Warehouse), it is necessary to first correct the facet orientations. In this technical report, we present an approach to correcting inconsistent facet orientations based on the Dirichlet energy of the generalized winding number. Our motivation came from the fact that if the mesh is watertight and its facets are consistently oriented, the generalized winding number is piecewise constant, and hence its Dirichlet energy is zero. Our intuition was that for a mesh with various imperfections such as holes and intersections, more consistently oriented facets would result in a smaller Dirichlet energy of the corresponding generalized winding number.

2 Method

2.1 Patch extraction

As a preprocessing step, facets that are topologically connected and forming a two-manifold surface are grouped together as a *patch*, similar to Borodin et al.’s method [BZK04]. If the mesh has non-manifold edges, the mesh is split around these edges, and multiple patches are generated. All facets belonging to the same patch share the same decision about whether they should be flipped or not in the output. Given n patches, our goal is to seek an assignment for n binary variables representing whether the i -th patch \mathcal{P}_i should be flipped or not ($1 \leq i \leq n$).

2.2 Generalized winding numbers

Let us first review the concept of generalized winding numbers [JKSH13]. The generalized winding number (hereafter we drop the word “generalized”) at a point $\mathbf{p} \in \mathbb{R}^3$ is defined as

$$w(\mathbf{p}) = \frac{1}{4\pi} \sum_{\tau \in \mathcal{M}} \theta_{\tau}(\mathbf{p}), \quad (1)$$

where \mathcal{M} is the input mesh (denoted as a set of facets), $\tau \in \mathcal{M}$ is a facet in \mathcal{M} , and θ_{τ} is the signed solid angle subtended by τ viewed from \mathbf{p} . Note that the sign of the angle is defined to be positive for points in the half space corresponding to the back side of a facet. In particular, the winding number of points inside and outside a watertight mesh is 1 and 0, respectively.

2.3 Dirichlet energy of the winding number

Since the winding number of a watertight and consistently oriented mesh is piecewise constant and hence its Dirichlet energy is zero, our intuition is that the Dirichlet energy of the winding number for an imperfect mesh should decrease as its facet orientations become more consistent. We give a mathematical formulation of the Dirichlet energy in the following.

For each i -th patch $\mathcal{P}_i \subseteq \mathcal{M}$ ($1 \leq i \leq n$), we want to determine a binary variable $s_i \in \{+1, -1\}$ called *sign*, where $s_i = -1$ means that all facets in \mathcal{P}_i should be flipped in the output. Now we can define the winding number with respect to the given sign $\mathbf{s} = (s_1, \dots, s_n)$ as

$$w_{\mathbf{s}}(\mathbf{p}) = \frac{1}{4\pi} \sum_{i=1}^n s_i \sum_{\tau \in \mathcal{P}_i} \theta_{\tau}(\mathbf{p}) . \quad (2)$$

Note that the orientation of each triangle τ is fixed to the original configuration when computing θ_{τ} . The Dirichlet energy of $w_{\mathbf{s}}$ denoted as $E(\mathbf{s})$ is a function of \mathbf{s} and is defined as the integral of the gradient magnitude squared of $w_{\mathbf{s}}$ over the volumetric domain $\Omega = \mathbb{R}^3 \setminus \mathcal{M}$:

$$E(\mathbf{s}) = \int_{\mathbf{p} \in \Omega} \|\nabla w_{\mathbf{s}}(\mathbf{p})\|^2 dV . \quad (3)$$

Reorganizing the integrand as

$$\|\nabla w_{\mathbf{s}}(\mathbf{p})\|^2 = \left\| \frac{1}{4\pi} \sum_{i=1}^n s_i \sum_{\tau \in \mathcal{P}_i} \nabla \theta_{\tau}(\mathbf{p}) \right\|^2 \quad (4)$$

$$= \frac{1}{16\pi^2} \left(\sum_{i=1}^n s_i \sum_{\tau \in \mathcal{P}_i} \nabla \theta_{\tau}(\mathbf{p}) \right) \cdot \left(\sum_{j=1}^n s_j \sum_{\mu \in \mathcal{P}_j} \nabla \theta_{\mu}(\mathbf{p}) \right) \quad (5)$$

$$= \frac{1}{16\pi^2} \sum_{i=1}^n \sum_{j=1}^n s_i s_j \left(\sum_{\tau \in \mathcal{P}_i} \nabla \theta_{\tau}(\mathbf{p}) \right) \cdot \left(\sum_{\mu \in \mathcal{P}_j} \nabla \theta_{\mu}(\mathbf{p}) \right) , \quad (6)$$

the energy can be written concisely in quadratic form as

$$E(\mathbf{s}) = \mathbf{s}^{\top} \mathbf{Q} \mathbf{s} , \quad (7)$$

where Q is an n -by- n matrix with elements

$$Q_{ij} = \frac{1}{16\pi^2} \sum_{\tau \in \mathcal{P}_i} \sum_{\mu \in \mathcal{P}_j} \int_{\mathbf{p} \in \Omega} \nabla \theta_{\tau}(\mathbf{p}) \cdot \nabla \theta_{\mu}(\mathbf{p}) dV . \quad (8)$$

Q is fixed per model independently of \mathbf{s} . Our goal is to find a set of facet orientations \mathbf{s}^* that gives the smallest Dirichlet energy:

$$\mathbf{s}^* = \arg \min_{\mathbf{s} \in \{-1, +1\}^n} E(\mathbf{s}) . \quad (9)$$

Since $s_i^2 = 1$, the energy can be further reorganized as

$$E(\mathbf{s}) = 2 \sum_{i=1}^{n-1} \sum_{j=i+1}^n s_i s_j Q_{ij} + \underbrace{\sum_{i=1}^n Q_{ii}}_{\text{const.}} . \quad (10)$$

Ignoring the diagonals of \mathbf{Q} , we can also consider

$$\tilde{E}(\mathbf{s}) = \sum_{i=1}^{n-1} \sum_{j=i+1}^n s_i s_j Q_{ij} \quad (11)$$

as the energy instead.

2.4 Computing the Dirichlet energy

We need to compute Q_{ij} in order to evaluate $E(\mathbf{s})$. A straightforward approach would be to tessellate the finite volume around the input mesh with sufficiently large volume and fine resolution to capture the variations of the winding numbers and then numerically compute the integral over this tessellation. However, this approach can be computationally expensive, especially when the mesh contains many small details and a fine tessellation is needed for achieving sufficient accuracy.

Fortunately, since the winding number function is harmonic (i.e., $\Delta w(\mathbf{p}) = 0$) [JKSH13], we can convert the volumetric integral over the domain Ω into a surface integral over the domain boundary $\partial\Omega$ using Green's first identity. For a once differentiable scalar function ψ and a twice differentiable scalar function ϕ , Green's first identity states

$$\int_{\Omega} \psi \Delta \phi dV + \int_{\Omega} \nabla \psi \cdot \nabla \phi dV = \int_{\partial\Omega} \psi (\nabla \phi \cdot \mathbf{n}) dS \quad (12)$$

where \mathbf{n} is an outgoing normal at the domain boundary (the positional variable \mathbf{p} is omitted here for brevity). By substituting $\psi = \theta_{\tau}$ and $\phi = \theta_{\mu}$ and recalling that winding numbers are harmonic, we get

$$\int_{\Omega} \theta_{\tau} \Delta \theta_{\mu} dV + \int_{\Omega} \nabla \theta_{\tau} \cdot \nabla \theta_{\mu} dV = \int_{\partial\Omega} \theta_{\tau} (\nabla \theta_{\mu} \cdot \mathbf{n}) dS , \quad (13)$$

which shows that the volumetric integral in (8) is now converted to a surface integral.

To compute this surface integral, we make the following notes on the domain and its boundary. The domain Ω consists of the entire 3D space \mathbb{R}^3 with a hole of infinitesimally thin slab carved out by each facet $\tau \in \mathcal{M}$, and the domain boundary $\partial\Omega$ consists of such holes. Each hole corresponding to τ has its front and back sides denoted by τ^+ and τ^- , respectively. Note that the domain's outgoing normals on τ^+ and τ^- are pointing toward the hole and are opposite to each other, see Figure 1. Now the domain boundary can be written as

$$\partial\Omega = \bigcup_{\tau \in \mathcal{M}} \tau^+ \cup \tau^- . \quad (14)$$

For a facet $\tau \in \mathcal{M}$ and a point on it $\mathbf{p} \in \tau$, let us denote the corresponding point on the front and back sides as $\mathbf{p}^+ \in \tau^+$ and $\mathbf{p}^- \in \tau^-$, respectively. Then, it is easy to see that the following properties hold:

$$\theta_\tau(\mathbf{p}^+) = -\theta_\tau(\mathbf{p}^-) = -2\pi \quad (15)$$

$$\nabla\theta_\tau(\mathbf{p}^+) = \nabla\theta_\tau(\mathbf{p}^-) , \quad (16)$$

where the latter is due to the derivative of the arctangent. Note that θ_τ and $\nabla\theta_\tau$ are undefined along the edges of τ ; we simply ignore the edge parts (i.e., parts connecting τ^+ and τ^-) in the surface integral, since the measure of all the edges vanishes (as opposed to the measure of the facets' interiors).

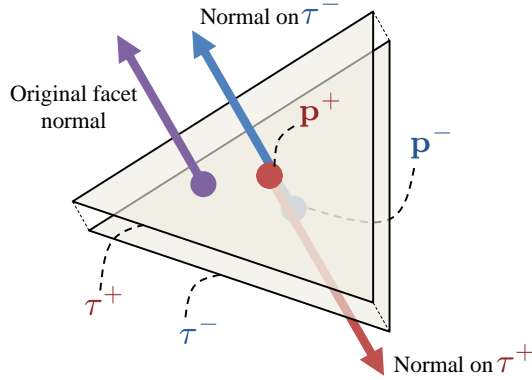


Figure 1: Notations for the domain boundary.

For any other third facet $\gamma \neq \tau$, θ_τ and $\nabla\theta_\mu$ take the same value on γ^+ and γ^- . Therefore,

$$\int_{\gamma^+ \cup \gamma^-} \theta_\tau(\nabla\theta_\mu \cdot \mathbf{n}) dS = \int_{\gamma^+} \theta_\tau(\nabla\theta_\mu \cdot \mathbf{n}) dS + \int_{\gamma^-} \theta_\tau(\nabla\theta_\mu \cdot \mathbf{n}) dS \quad (17)$$

$$= \int_{\gamma^+} \theta_\tau(\nabla\theta_\mu \cdot \mathbf{n}) dS + \int_{\gamma^+} \theta_\tau(\nabla\theta_\mu \cdot (-\mathbf{n})) dS \quad (18)$$

$$= 0 \quad (19)$$

where in (18) we replaced the integral over γ^- with an integral over γ^+ with a negated surface normal. Therefore, the integral over $\partial\Omega$ is reduced to an integral over $\tau^+ \cup \tau^-$, i.e.,

$$\int_{\partial\Omega} \theta_\tau(\nabla\theta_\mu \cdot \mathbf{n}) dS = \sum_{\gamma \in \mathcal{M}} \int_{\gamma^+ \cup \gamma^-} \theta_\tau(\nabla\theta_\mu \cdot \mathbf{n}) dS \quad (20)$$

$$= \int_{\tau^+ \cup \tau^-} \theta_\tau(\nabla\theta_\mu \cdot \mathbf{n}) dS . \quad (21)$$

Using the properties above, we arrive at:

$$\int_{\tau^+ \cup \tau^-} \theta_\tau(\nabla\theta_\mu \cdot \mathbf{n}) dS = \int_{\tau^+} (-2\pi)(\nabla\theta_\mu \cdot \mathbf{n}) dS + \int_{\tau^-} (2\pi)(\nabla\theta_\mu \cdot \mathbf{n}) dS \quad (22)$$

$$= -2\pi \int_{\tau^+} \nabla\theta_\mu \cdot \mathbf{n} dS + 2\pi \int_{\tau^+} \nabla\theta_\mu \cdot (-\mathbf{n}) dS \quad (23)$$

$$= -4\pi \int_{\tau^+} \nabla\theta_\mu \cdot \mathbf{n} dS . \quad (24)$$

In summary, Q_{ij} is computed as

$$Q_{ij} = \sum_{\tau \in \mathcal{P}_i} \sum_{\mu \in \mathcal{P}_j} F(\tau, \mu) \quad (25)$$

where

$$F(\tau, \mu) = -\frac{1}{4\pi} \int_{\mathbf{p} \in \tau^+} \nabla\theta_\mu(\mathbf{p}) \cdot \mathbf{n}(\mathbf{p}) dS . \quad (26)$$

2.5 Convergence of the Dirichlet energy

By deriving an analytic form of $F(\tau, \mu)$ using MAPLE, we found that it evaluates to infinity (either positive or negative depending on relative orientation of τ and μ) when some edges of τ and μ coincide geometrically. An intuitive explanation for this is that the magnitude of $\nabla\theta_\mu(\mathbf{p})$ becomes infinitely large when $\mathbf{p} \in \tau^+$ goes infinitesimally close to the edges of τ that are geometrically coinciding with the edges of μ . In particular, the diagonal entries Q_{ii} always evaluate to positive infinity. Although we can ignore the diagonals using (11), off-diagonal entries Q_{ij} can still evaluate to infinity if some boundary edges of \mathcal{P}_i and \mathcal{P}_j coincide geometrically. In our experiments, we numerically evaluate the energy with a sufficiently fine discretization of $\partial\Omega$ where the finite entries in \mathbf{Q} are converged.

2.6 Finding the minimizer of the Dirichlet energy

The minimization problem in (9) is referred to as Quadratic Unconstrained Binary Optimization (QUBO). QUBO is NP-hard, and various heuristic methods exist [BHT07]. In our experiments, we take a brute force approach and exhaustively search for the optimal solution, so that we can examine the nature of our approach with small toy examples.

3 Results

We implemented our approach in 3D as well as in 2D where the input consists of open or closed polylines (still referred to as “patches” in the following) and the same formulation applies. Although here we show our experiments only in 2D for the ease of visualization, the same argument also holds in 3D. Figure 2 shows cases where the input corresponds to the boundary of a solid shape with some missing pieces. In such cases, our algorithm produces the expected results, i.e., it orients the facets consistently and correctly identifies the shape’s interior. On the other hand, Figure 3 shows a case where our approach is fundamentally flawed: the input has some small patches attached to larger patches, and those small patches are oriented inside-out. The fundamental reason for this is that the Dirichlet energy only measures the variation of the function while ignoring the sign, and hence there is no penalty for orienting facets inside-out.




Figure 2: Examples where our approach produces reasonable results. The input polylines are represented by white lines with their orientations indicated by yellow arrows. The number below each image is the Dirichlet energy of the corresponding winding number function. The color maps represent the associated winding numbers, where higher winding numbers correspond to warmer colors (the actual mapping is slightly adjusted for each image to better visualize the distribution).




Figure 3: An example representing a torso with two small objects attached to the shoulder and the side. Our approach results in those small attachments oriented inside-out.

For an intuitive explanation of this behavior, it is useful to consider the *closure* of a patch \mathcal{P} denoted as $\bar{\mathcal{P}}$ [JKSH13], where the union of \mathcal{P} and $\bar{\mathcal{P}}$ bounds a finite volumetric domain $\Omega_{\mathcal{P}}$ (i.e., $\partial\Omega_{\mathcal{P}} = \mathcal{P} \cup \bar{\mathcal{P}}$). Denoting the winding number defined by a patch \mathcal{P} as $w_{\mathcal{P}}(\mathbf{p})$, the following holds:

$$w_{\mathcal{P}}(\mathbf{p}) + w_{\bar{\mathcal{P}}}(\mathbf{p}) = \begin{cases} 1, & \mathbf{p} \in \Omega_{\mathcal{P}} \\ 0, & \text{otherwise} \end{cases} . \quad (27)$$

Since the right hand side of (27) is piecewise-constant, the following also holds:

$$\nabla w_{\mathcal{P}}(\mathbf{p}) = -\nabla w_{\bar{\mathcal{P}}}(\mathbf{p}) . \quad (28)$$

Let us now consider a simple case with two patches \mathcal{P}_1 and \mathcal{P}_2 as shown in Figure 4. We can decompose the gradient and apply the above observation as:

$$\nabla w_{\mathcal{P}_1 \cup \mathcal{P}_2}(\mathbf{p}) = \nabla w_{\mathcal{P}_1}(\mathbf{p}) + \nabla w_{\mathcal{P}_2}(\mathbf{p}) \quad (29)$$

$$= -\nabla w_{\bar{\mathcal{P}}_1}(\mathbf{p}) - \nabla w_{\bar{\mathcal{P}}_2}(\mathbf{p}) . \quad (30)$$

It is clear in the figure that the current configuration of $\bar{\mathcal{P}}_1$ and $\bar{\mathcal{P}}_2$ yields smaller Dirichlet energy than the other configuration, and hence \mathcal{P}_2 is oriented this way in our approach.




Figure 4: An example showing the reason of why small attachments are oriented inside-out in our approach. Note again that the mapping between the winding number and the color is adjusted in each image to better visualize the function.

4 Conclusion and future work

We conclude that the presented approach has a fundamental flaw and is not applicable for most handmade meshes. One direction for further exploration is to apply a similar approach to surface reconstruction from an unoriented point cloud [MDGD⁺10, LW10], where only the positions of the points are given and finding good normals for the points is critical. In this case, the unknowns are the continuous normal directions, potentially making the problem easier. A key question, however, would be how to generalize the definition of winding numbers to point clouds.

Acknowledgements

We are grateful to Peter Kaufmann, Wenzel Jakob, Nobuyuki Umetani, Melina Skouras, Ilya Baran, and Ryan Schmidt for their advice and feedback. This work was supported in part by the ERC grant iModel

(StG-2012-306877), by an SNF award 200021_137879, by an NSF grant IIS-1350330 and by a gift from Adobe Research. Kenshi Takayama's work was funded by JSPS Postdoctoral Fellowships for Research Abroad. Alec Jacobson's work was supported by an Intel Doctoral Fellowship.

References

- [BHT07] Endre Boros, PeterL. Hammer, and Gabriel Tavares. Local search heuristics for quadratic unconstrained binary optimization (qubo). *Journal of Heuristics*, 13(2):99–132, 2007.
- [BZK04] Pavel Borodin, Gabriel Zachmann, and Reinhard Klein. Consistent normal orientation for polygonal meshes. In *Proceedings of the Computer Graphics International, CGI '04*, pages 18–25, 2004.
- [JKSH13] Alec Jacobson, Ladislav Kavan, and Olga Sorkine-Hornung. Robust inside-outside segmentation using generalized winding numbers. *ACM Trans. Graph.*, 32(4):33:1–33:12, 2013.
- [LW10] Shengjun Liu and Charlie C.L. Wang. Orienting unorganized points for surface reconstruction. *Computers & Graphics*, 34(3):209–218, 2010.
- [MDGD⁺10] Patrick Mullen, Fernando De Goes, Mathieu Desbrun, David Cohen-Steiner, and Pierre Alliez. Signing the unsigned: Robust surface reconstruction from raw pointsets. *Computer Graphics Forum*, 29(5):1733–1741, 2010.