5 3D Objects

5.1 Surface Representations

As with 2D objects, we can represent 3D objects in **parametric** and **implicit** forms. (There are also explicit forms for 3D surfaces — sometimes called "height fields" — but we will not cover them here).

5.2 Planes

- **Implicit:** $(\bar{p} - \bar{p}_0) \cdot \vec{n} = 0$, where $\bar{p}_0$ is a point in $\mathbb{R}^3$ on the plane, and $\vec{n}$ is a normal vector perpendicular to the plane.

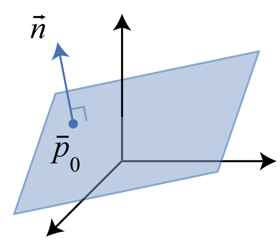

A plane can be defined uniquely by three non-colinear points $\bar{p}_1, \bar{p}_2, \bar{p}_3$. Let $\vec{a} = \bar{p}_2 - \bar{p}_1$ and $\vec{b} = \bar{p}_3 - \bar{p}_1$, so $\vec{a}$ and $\vec{b}$ are vectors in the plane. Then $\vec{n} = \vec{a} \times \vec{b}$. Since the points are not colinear, $\|\vec{n}\| \neq 0$.

- **Parametric:** $\bar{s}(\alpha, \beta) = \bar{p}_0 + \alpha\vec{a} + \beta\vec{b}$, for $\alpha, \beta \in \mathbb{R}$.

Note:
This is similar to the parametric form of a line: $\bar{l}(\alpha) = \bar{p}_0 + \alpha\vec{a}$.

A **planar patch** is a parallelogram defined by bounds on α and β.

Example:
Let $0 \leq \alpha \leq 1$ and $0 \leq \beta \leq 1$:

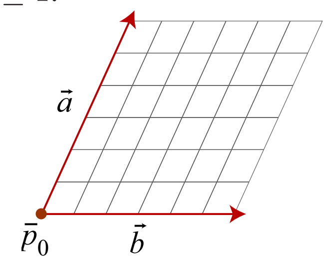

Copyright © 2005 David Fleet and Aaron Hertzmann

5.3 Surface Tangents and Normals

The **tangent** to a curve at $\bar{p}$ is the instantaneous direction of the curve at $\bar{p}$.

The **tangent plane** to a surface at $\bar{p}$ is analogous. It is defined as the plane containing tangent vectors to all curves on the surface that go through $\bar{p}$.

A **surface normal** at a point $\bar{p}$ is a vector perpendicular to a tangent plane.

5.3.1 Curves on Surfaces

The parametric form $\bar{p}(\alpha, \beta)$ of a surface defines a mapping from 2D points to 3D points: every 2D point (α, β) in $\mathbb{R}^2$ corresponds to a 3D point $\bar{p}$ in $\mathbb{R}^3$. Moreover, consider a curve $\bar{l}(\lambda) = (\alpha(\lambda), \beta(\lambda))$ in 2D — there is a corresponding curve in 3D contained within the surface: $\bar{l}^*(\lambda) = \bar{p}(\bar{l}(\lambda))$.

5.3.2 Parametric Form

For a curve $\bar{c}(\lambda) = (x(\lambda), y(\lambda), z(\lambda))^T$ in 3D, the tangent is

$$\frac{d\bar{c}(\lambda)}{d\lambda} = \left(\frac{dx(\lambda)}{d\lambda}, \frac{dy(\lambda)}{d\lambda}, \frac{dz(\lambda)}{d\lambda}\right). \tag{1}$$

For a surface point $\bar{s}(\alpha, \beta)$, two tangent vectors can be computed:

$$\frac{\partial \bar{s}}{\partial \alpha} \text{ and } \frac{\partial \bar{s}}{\partial \beta}. \tag{2}$$

Derivation:
Consider a point (α_0, β_0) in 2D which corresponds to a 3D point $\bar{s}(\alpha_0, \beta_0)$. Define two straight lines in 2D:

$$\bar{d}(\lambda_1) = (\lambda_1, \beta_0)^T \tag{3}$$
$$\bar{e}(\lambda_2) = (\alpha_0, \lambda_2)^T \tag{4}$$

These lines correspond to curves in 3D:

$$\bar{d}^*(\lambda_1) = \bar{s}(\bar{d}(\lambda_1)) \tag{5}$$
$$\bar{e}^*(\lambda_2) = \bar{s}(\bar{d}(\lambda_2)) \tag{6}$$

Copyright © 2005 David Fleet and Aaron Hertzmann

Using the chain rule for vector functions, the tangents of these curves are:

$$\frac{\partial \bar{d}^*}{\partial \lambda_1} = \frac{\partial \bar{s}}{\partial \alpha}\frac{\partial \bar{d}_\alpha}{\partial \lambda_1} + \frac{\partial \bar{s}}{\partial \beta}\frac{\partial \bar{d}_\beta}{\partial \lambda_1} \quad = \quad \frac{\partial \bar{s}}{\partial \alpha} \tag{7}$$

$$\frac{\partial \bar{e}^*}{\partial \lambda_2} = \frac{\partial \bar{s}}{\partial \alpha}\frac{\partial \bar{e}_\alpha}{\partial \lambda_2} + \frac{\partial \bar{s}}{\partial \beta}\frac{\partial \bar{e}_\beta}{\partial \lambda_2} \quad = \quad \frac{\partial \bar{s}}{\partial \beta} \tag{8}$$

The normal of $\bar{s}$ at $\alpha = \alpha_0$, $\beta = \beta_0$ is

$$\vec{n}(\alpha_0, \beta_0) \quad = \quad \left(\left. \frac{\partial \bar{s}}{\partial \alpha} \right|_{\alpha_0, \beta_0} \right) \times \left(\left. \frac{\partial \bar{s}}{\partial \beta} \right|_{\alpha_0, \beta_0} \right). \tag{9}$$

The tangent plane is a plane containing the surface at $\bar{s}(\alpha_0, \beta_0)$ with normal vector equal to the surface normal. The equation for the tangent plane is:

$$\vec{n}(\alpha_0, \beta_0) \cdot (\bar{p} - \bar{s}(\alpha_0, \beta_0)) \quad = \quad 0. \tag{10}$$

What if we used different curves in 2D to define the tangent plane? It can be shown that we get the same tangent plane; in other words, tangent vectors of all 2D curves through a given surface point are contained within a single tangent plane. (Try this as an exercise).

Note:
The normal vector is not unique. If $\vec{n}$ is a normal vector, then any vector $\alpha\vec{n}$ is also normal to the surface, for $\alpha \in \mathbb{R}$. What this means is that the normal can be scaled, and the direction can be reversed.

5.3.3 Implicit Form

In the implicit form, a surface is defined as the set of points $\bar{p}$ that satisfy $f(\bar{p}) = 0$ for some function f. A normal is given by the gradient of f,

$$\vec{n}(\bar{p}) = \nabla f(\bar{p})|_{\bar{p}} \tag{11}$$

where $\nabla f = \left(\frac{\partial f(\bar{p})}{\partial x}, \frac{\partial f(\bar{p})}{\partial y}, \frac{\partial f(\bar{p})}{\partial z} \right)$.

Derivation:
Consider a 3D curve $\bar{c}(\lambda)$ that is contained within the 3D surface, and that passes through $\bar{p}_0$ at λ_0. In other words, $\bar{c}(\lambda_0) = \bar{p}_0$ and

$$f(\bar{c}(\lambda)) \quad = \quad 0 \tag{12}$$

Copyright © 2005 David Fleet and Aaron Hertzmann

for all λ. Differentiating both sides gives:

$$\frac{\partial f}{\partial \lambda} = 0 \tag{13}$$

Expanding the left-hand side, we see:

$$\frac{\partial f}{\partial \lambda} = \frac{\partial f}{\partial x}\frac{\partial \bar{c}_x}{\partial \lambda} + \frac{\partial f}{\partial y}\frac{\partial \bar{c}_y}{\partial \lambda} + \frac{\partial f}{\partial z}\frac{\partial \bar{c}_z}{\partial \lambda} \tag{14}$$

$$= \nabla f(\bar{p})|_{\bar{p}} \cdot \frac{d\bar{c}}{d\lambda} = 0 \tag{15}$$

This last line states that the gradient is perpendicular to the curve tangent, which is the definition of the normal vector.

Example:
The implicit form of a sphere is: $f(\bar{p}) = ||\bar{p} - \bar{c}||^2 - R^2 = 0$. The normal at a point $\bar{p}$ is: $\nabla f = 2(\bar{p} - \bar{c})$.

Exercise: show that the normal computed for a plane is the same, regardless of whether it is computed using the parametric or implicit forms. (This was done in class). Try it for another surface.

5.4 Parametric Surfaces

5.4.1 Bilinear Patch

A **bilinear patch** is defined by four points, no three of which are colinear.

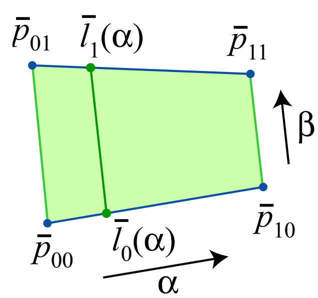

Given $\bar{p}_{00}, \bar{p}_{01}, \bar{p}_{10}, \bar{p}_{11}$, define

$$\bar{l}_0(\alpha) = (1-\alpha)\bar{p}_{00} + \alpha\bar{p}_{10},$$
$$\bar{l}_1(\alpha) = (1-\alpha)\bar{p}_{01} + \alpha\bar{p}_{11}.$$

Copyright © 2005 David Fleet and Aaron Hertzmann

Then connect $\bar{l}_0(\alpha)$ and $\bar{l}_1(\alpha)$ with a line:

$$\bar{p}(\alpha, \beta) = (1 - \beta)\bar{l}_0(\alpha) + \beta\bar{l}_1(\alpha),$$

for $0 \leq \alpha \leq 1$ and $0 \leq \beta \leq 1$.

Question: when is a bilinear patch not equivalent to a planar patch? Hint: a planar patch is defined by 3 points, but a bilinear patch is defined by 4.

5.4.2 Cylinder

A **cylinder** is constructed by moving a point on a line l along a planar curve $p_0(\alpha)$ such that the direction of the line is held constant.

If the direction of the line l is $\vec{d}$, the cylinder is defined as

$$\bar{p}(\alpha, \beta) = p_0(\alpha) + \beta\vec{d}.$$

A **right cylinder** has $\vec{d}$ perpendicular to the plane containing $p_0(\alpha)$.

A **circular cylinder** is a cylinder where $p_0(\alpha)$ is a circle.

Example:
A right circular cylinder can be defined by $p_0(\alpha) = (r\cos(\alpha), r\sin(\alpha), 0)$, for $0 \leq \alpha < 2\pi$, and $\vec{d} = (0, 0, 1)$.

So $p_0(\alpha, \beta) = (r\cos(\alpha), r\sin(\alpha), \beta)$, for $0 \leq \beta \leq 1$.

To find the normal at a point on this cylinder, we can use the implicit form $f(x, y, z) = x^2 + y^2 - r^2 = 0$ to find $\nabla f = 2(x, y, 0)$.

Using the parametric form directly to find the normal, we have

$$\frac{\partial \bar{p}}{\partial \alpha} = r(-\sin(\alpha), \cos(\alpha), 0), \text{ and } \frac{\partial \bar{p}}{\partial \beta} = (0, 0, 1), \text{ so}$$

$$\frac{\partial \bar{p}}{\partial \alpha} \times \frac{\partial \bar{p}}{\partial \beta} = (r\cos(\alpha)r\sin(\alpha), 0).$$

Note:
The cross product of two vectors $\vec{a} = (a_1, a_2, a_3)$ and $\vec{b} = (b_1, b_2, b_3)$ can

be found by taking the determinant of the matrix,

$$\begin{bmatrix} i & j & k \\ a_1 & a_2 & a_3 \\ b_1 & b_2 & b_3 \end{bmatrix}.$$

5.4.3 Surface of Revolution

To form a **surface of revolution**, we revolve a curve in the x-z plane, $\bar{c}(\beta) = (x(\beta), 0, z(\beta))$, about the z-axis.

Hence, each point on $\bar{c}$ traces out a circle parallel to the x-y plane with radius $|x(\beta)|$. Circles then have the form $(r\cos(\alpha), r\sin(\alpha))$, where α is the parameter of revolution. So the rotated surface has the parametric form

$$\bar{s}(\alpha, \beta) = (x(\beta)\cos(\alpha), x(\beta)\sin(\alpha), z(\beta)).$$

Example:
If $\bar{c}(\beta)$ is a line perpendicular to the x-axis, we have a right circular cylinder.

A torus is a surface of revolution:

$$\bar{c}(\beta) = (d + r\cos(\beta), 0, r\sin(\beta)).$$

5.4.4 Quadric

A **quadric** is a generalization of a conic section to 3D. The implicit form of a quadric in the standard position is

$$\begin{aligned} ax^2 + by^2 + cz^2 + d &= 0, \\ ax^2 + by^2 + ez &= 0, \end{aligned}$$

for $a, b, c, d, e \in \mathbb{R}$. There are six basic types of quadric surfaces, which depend on the signs of the parameters.
They are the ellipsoid, hyperboloid of one sheet, hyperboloid of two sheets, elliptic cone, elliptic paraboloid, and hyperbolic paraboloid (saddle). All but the hyperbolic paraboloid may be expressed as a surface of revolution.

Copyright © 2005 David Fleet and Aaron Hertzmann

Example:
An ellipsoid has the implicit form

$$\frac{x^2}{a^2} + \frac{y^2}{b^2} + \frac{z^2}{c^2} - 1 = 0.$$

In parametric form, this is

$$\bar{s}(\alpha, \beta) = (a \sin(\beta) \cos(\alpha), b \sin(\beta) \sin(\alpha), c \cos(\beta)),$$

for $\beta \in [0, \pi]$ and $\alpha \in (-\pi, \pi]$.

5.4.5 Polygonal Mesh

A **polygonal mesh** is a collection of polygons (vertices, edges, and faces). As polygons may be used to approximate curves, a polygonal mesh may be used to approximate a surface.

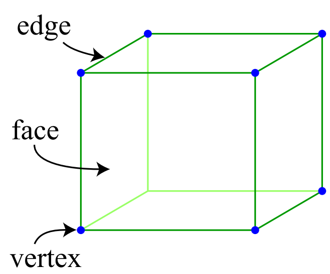

A **polyhedron** is a closed, connected polygonal mesh. Each edge must be shared by two faces.

A **face** refers to a planar polygonal patch within a mesh.

A mesh is **simple** when its topology is equivalent to that of a sphere. That is, it has no holes.

Given a parametric surface, $\bar{s}(\alpha, \beta)$, we can sample values of α and β to generate a polygonal mesh approximating $\bar{s}$.

5.5 3D Affine Transformations

Three dimensional transformations are used for many different purposes, such as coordinate transforms, shape modeling, animation, and camera modeling.

Copyright © 2005 David Fleet and Aaron Hertzmann

An affine transform in 3D looks the same as in 2D: $F(\bar{p}) = A\bar{p} + \vec{t}$ for $A \in \mathbb{R}^{3\times3}$, $\bar{p}, \vec{t} \in \mathbb{R}^3$. A homogeneous affine transformation is

$$\hat{F}(\hat{p}) = \hat{M}\hat{p}, \text{ where } \hat{p} = \begin{bmatrix} \bar{p} \\ 1 \end{bmatrix}, \hat{M} = \begin{bmatrix} A & \vec{t} \\ \vec{0}^T & 1 \end{bmatrix}.$$

Translation: $A = I$, $\vec{t} = (t_x, t_y, t_z)$.

Scaling: $A = \text{diag}(s_x, s_y, s_z)$, $\vec{t} = \vec{0}$.

Rotation: $A = R$, $\vec{t} = \vec{0}$, and $\det(R) = 1$.

3D rotations are much more complex than 2D rotations, so we will consider only elementary rotations about the x, y, and z axes.

For a rotation about the z-axis, the z coordinate remains unchanged, and the rotation occurs in the x-y plane. So if $\bar{q} = R\bar{p}$, then $q_z = p_z$. That is,

$$\begin{bmatrix} q_x \\ q_y \end{bmatrix} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} \begin{bmatrix} p_x \\ p_y \end{bmatrix}.$$

Including the z coordinate, this becomes

$$R_z(\theta) = \begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Similarly, rotation about the x-axis is

$$R_x(\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta) & -\sin(\theta) \\ 0 & \sin(\theta) & \cos(\theta) \end{bmatrix}.$$

For rotation about the y-axis,

$$R_y(\theta) = \begin{bmatrix} \cos(\theta) & 0 & \sin(\theta) \\ 0 & 1 & 0 \\ -\sin(\theta) & 0 & \cos(\theta) \end{bmatrix}.$$

Copyright © 2005 David Fleet and Aaron Hertzmann

5.6 Spherical Coordinates

Any three dimensional vector $\vec{u} = (u_x, u_y, u_z)$ may be represented in **spherical coordinates**. By computing a polar angle ϕ counterclockwise about the y-axis from the z-axis and an azimuthal angle θ counterclockwise about the z-axis from the x-axis, we can define a vector in the appropriate direction. Then it is only a matter of scaling this vector to the correct length $(u_x^2 + u_y^2 + u_z^2)^{-1/2}$ to match $\vec{u}$.

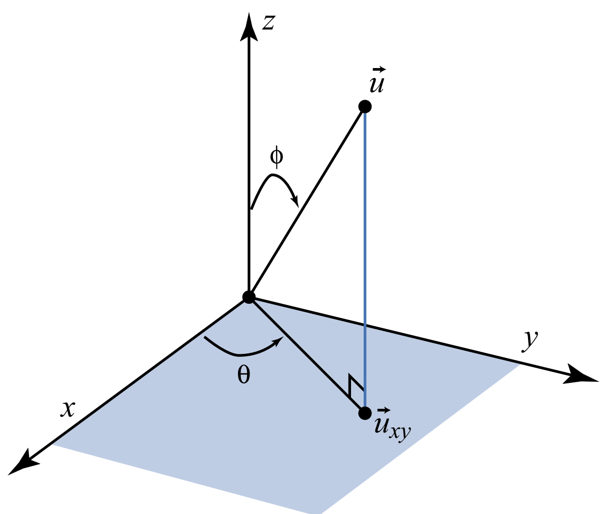

Given angles ϕ and θ, we can find a unit vector as $\vec{u} = (\cos(\theta)\sin(\phi), \sin(\theta)\sin(\phi), cos(\phi))$.

Given a vector $\vec{u}$, its azimuthal angle is given by $\theta = \arctan\left(\frac{u_y}{u_x}\right)$ and its polar angle is $\phi = \arctan\left(\frac{(u_x^2+u_y^2)^{1/2}}{u_z}\right)$. This formula does not require that $\vec{u}$ be a unit vector.

5.6.1 Rotation of a Point About a Line

Spherical coordinates are useful in finding the rotation of a point about an arbitrary line. Let $\bar{l}(\lambda) = \lambda\vec{u}$ with $\|\vec{u}\| = 1$, and $\vec{u}$ having azimuthal angle θ and polar angle ϕ. We may compose elementary rotations to get the effect of rotating a point $\bar{p}$ about $\bar{l}(\lambda)$ by a counterclockwise angle ρ:

1. Align $\vec{u}$ with the z-axis.

   - Rotate by $-\theta$ about the z-axis so $\vec{u}$ goes to the xz-plane.
   - Rotate up to the z-axis by rotating by $-\phi$ about the y-axis.

   Hence, $\bar{q} = R_y(-\phi)R_z(-\theta)\bar{p}$

2. Apply a rotation by ρ about the z-axis: $R_z(\rho)$.

Copyright © 2005 David Fleet and Aaron Hertzmann

3. Invert the first step to move the z-axis back to $\vec{u}$: $R_z(\theta)R_y(\phi) = (R_y(-\phi)R_z(-\theta))^{-1}$.

Finally, our formula is $\bar{q} = R_{\vec{u}}(\rho)\bar{p} = R_z(\theta)R_y(\phi)R_z(\rho)R_y(-\phi)R_z(-\theta)\bar{p}$.

5.7 Nonlinear Transformations

Affine transformations are a first-order model of shape deformation. With affine transformations, scaling and shear are the simplest nonrigid deformations. Common higher-order deformations include tapering, twisting, and bending.

Example:
To create a nonlinear taper, instead of constantly scaling in x and y for all z, as in

$$\bar{q} = \begin{bmatrix} a & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & 1 \end{bmatrix} \bar{p},$$

let a and b be functions of z, so

$$\bar{q} = \begin{bmatrix} a(\bar{p}_z) & 0 & 0 \\ 0 & b(\bar{p}_z) & 0 \\ 0 & 0 & 1 \end{bmatrix} \bar{p}.$$

A linear taper looks like $a(z) = \alpha_0 + \alpha_1 z$.
A quadratic taper would be $a(z) = \alpha_0 + \alpha_1 z + \alpha_2 z^2$.

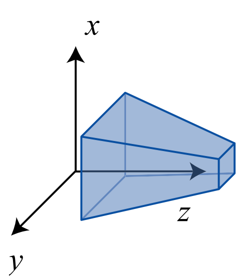

(a) Linear taper

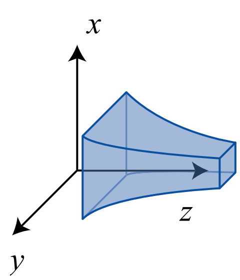

(b) Nonlinear taper

5.8 Representing Triangle Meshes

A triangle mesh is often represented with a list of vertices and a list of triangle faces. Each vertex consists of three floating point values for the x, y, and z positions, and a face consists of three

Copyright © 2005 David Fleet and Aaron Hertzmann

indices of vertices in the vertex list. Representing a mesh this way reduces memory use, since each vertex needs to be stored once, rather than once for every face it is on; and this gives us connectivity information, since it is possible to determine which faces share a common vertex. This can easily be extended to represent polygons with an arbitrary number of vertices, but any polygon can be decomposed into triangles. A tetrahedron can be represented with the following lists:

Vertex index	x	y	z
0	0	0	0
1	1	0	0
2	0	1	0
3	0	0	1

Face index	Vertices
0	0, 1, 2
1	0, 3, 1
2	1, 3, 2
3	2, 3, 0

Notice that vertices are specified in a counter-clockwise order, so that the front of the face and back can be distinguished. This is the default behavior for OpenGL, although it can also be set to take face vertices in clockwise order. Lists of normals and texture coordinates can also be specified, with each face then associated with a list of vertices and corresponding normals and texture coordinates.

5.9 Generating Triangle Meshes

As stated earlier, a parametric surface can be sampled to generate a polygonal mesh. Consider the surface of revolution

$$\bar{S}(\alpha, \beta) = [x(\alpha)\cos\beta, x(\alpha)\sin\beta, z(\alpha)]^T$$

with the profile $\bar{C}(\alpha) = [x(\alpha), 0, z(\alpha)]^T$ and $\beta \in [0, 2\pi]$.

To take a uniform sampling, we can use

$$\Delta\alpha = \frac{\alpha_1 - \alpha_0}{m}, \text{ and } \Delta\beta = \frac{2\pi}{n},$$

where m is the number of patches to take along the z-axis, and n is the number of patches to take around the z-axis.

Each patch would consist of four vertices as follows:

$$S_{ij} = \begin{pmatrix} \bar{S}(i\Delta\alpha, j\Delta\beta) \\ \bar{S}((i+1)\Delta\alpha, j\Delta\beta) \\ \bar{S}((i+1)\Delta\alpha, (j+1)\Delta\beta) \\ \bar{S}(i\Delta\alpha, (j+1)\Delta\beta) \end{pmatrix} = \begin{pmatrix} \bar{S}_{i,j} \\ \bar{S}_{i+1,j} \\ \bar{S}_{i+1,j+1} \\ \bar{S}_{i,j+1} \end{pmatrix}, \quad \text{for} \quad \begin{array}{l} i \in [0, m-1], \\ j \in [0, n-1] \end{array}$$

To render this as a triangle mesh, we must *tesselate* the sampled quads into triangles. This is accomplished by defining triangles P_{ij} and Q_{ij} given S_{ij} as follows:

$$P_{ij} = (\bar{S}_{i,j}, \bar{S}_{i+1,j}, \bar{S}_{i+1,j+1}), \text{ and } Q_{ij} = (\bar{S}_{i,j}, \bar{S}_{i+1,j+1}, \bar{S}_{i,j+1})$$

Copyright © 2005 David Fleet and Aaron Hertzmann