Lecture 7: Thursday, 25 September 1997

Administrative:

Asst 3A (Object Modeling) status

How is it going ?

Next Wednesday, October 1^{st}

Guest Lecture: 7-9pm in 34-101
Paul Haeberli, SGI, “Media Synthesis”

By applying modern computer graphics to traditional drawing, painting, photography and manufacturing, we can synthesize a new class of creative media.

Several ongoing projects will be shown that explore parametric [modeling and] manufacturing, digital photography, abstract representations, and interaction.

(See `http://www.sgi.com/grafica`)

Today:

Perspective Transformation Wrap-Up

Local illumination models (H&B S14.2)

Perspective WrapUp

Perspective Transformation takes frustum to NDC:

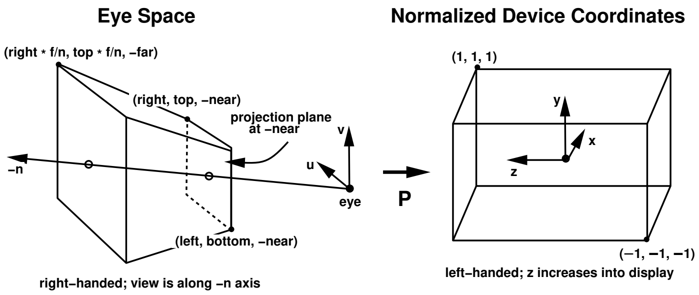

$$\mathbf{M} = \begin{pmatrix} \frac{2n}{r-l} & 0 & \frac{r+l}{r-l} & 0 \\ 0 & \frac{2n}{t-b} & \frac{t+b}{t-b} & 0 \\ 0 & 0 & -(\frac{f+n}{f-n}) & -(\frac{2fn}{f-n}) \\ 0 & 0 & -1 & 0 \end{pmatrix}$$

Why go through all this? Simplifies three things:

1) Clipping: clip against box, $x, y, z \in [-1..+1]$

2) Projection: simply drop z_{ndc}

3) Transformation to screen space (Tuesday, L8)

Amenable to hardware, software implementation

Types of Illumination Models

Local illumination

Also called *lighting, shading* models

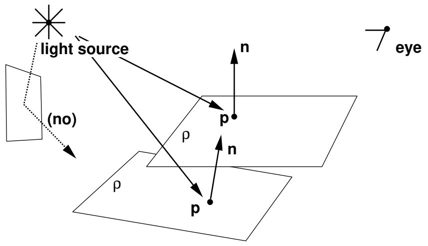

One or more “point” light sources

At infinity (parallel rays), or at finite point

No occlusion (other than back-face determination)

primary light sources only (no reflection)

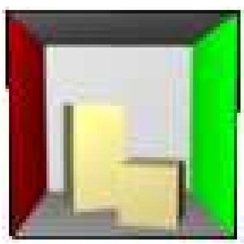

Types of Illumination Models

Semi-Local illumination

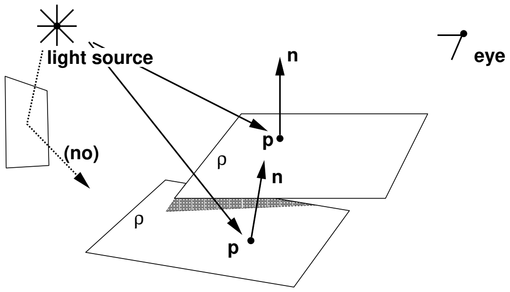

Shadow-tests for blocked *primary* sources

Point light sources yield *hard* shadows

No *secondary* light sources (reflection)

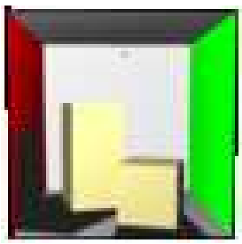

Global Illumination

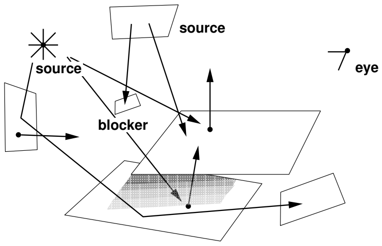

point, area, even volumetric light sources
area lights yield *soft* shadows (penumbrae)
secondary, tertiary, etc., light sources
comprehensive treatment of occlusion
view-independent (diffuse) effects

view-dependent (specular) effects

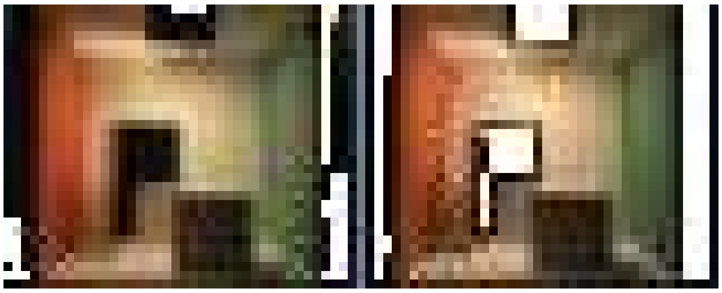

Cornell Global Illumination Research

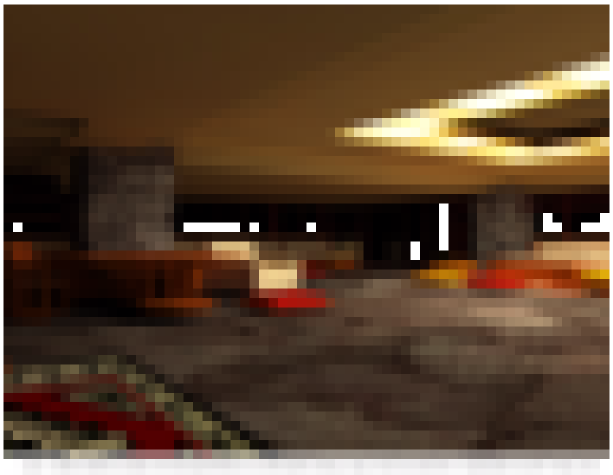

Local Illumination Models

Given point **p**, normal **N**, light **L**, viewer **V**
Shade **p** according to reflectance properties ρ

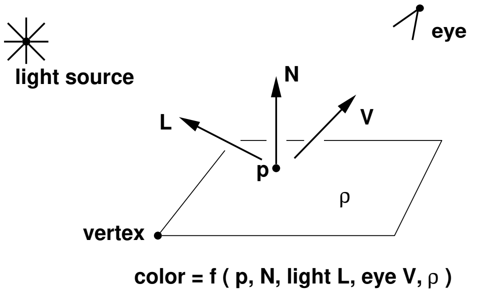

Emissive Objects

Simplest illumination model possible

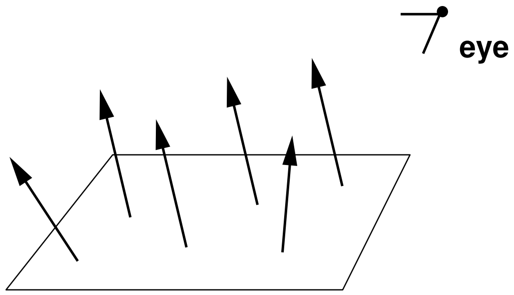

Every object emits light of some color, intensity

$$I = I_e$$

I_e is the surface's *emissive intensity*
I is the *reflected intensity* at shaded point
Examples:
Phosphorescent, bioluminescent materials

Ambient Illumination

Ambient (perfused) light often present due to reflection

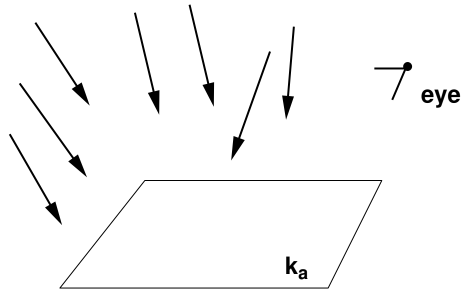

Quantitatively (I = Intensity; radiometric units later):

$$I = I_a k_a$$

k_a is the *ambient reflection coefficient*
I_a is the *ambient light intensity*
Examples:
Outside on a cloudy day
Outside, sunny day but in shadow
"Deep" building interior, no artificial lights

Ideal Diffuse Reflection: Qualitative

Consider illumination from distant source

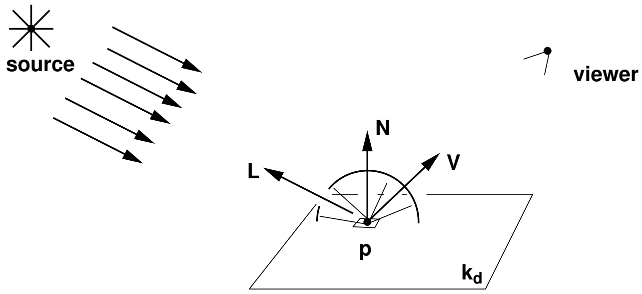

Ideal diffuse reflection occurs from perfect matte (unshiny) surfaces, not found in nature
reflection depends on angle between surface normal and direction **to light source**
apparent area depends on angle between surface normal and direction **to viewer**
geometry causes *viewer dependence* to cancel out

Fixed diffuse surface equally bright, regardless of direction from which it is viewed!

Ideal Diffuse Reflection: Quantitative

Lambert's Law (17th century physicist, mathematician)

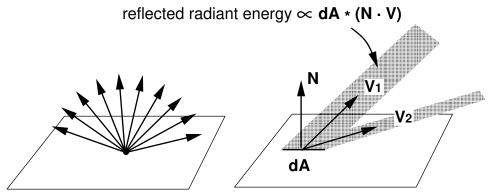

Light reflects isotropically from surface
Distributed equally to outgoing hemisphere
But: surface orientation creates preferred direction
Thus reflected **energy** is $\propto \mathbf{N} \cdot \mathbf{V}$
One way to understand: in terms of solid angle

Ideal Diffuse Reflection: Quantitative

Suppose a distant light source illuminates surface

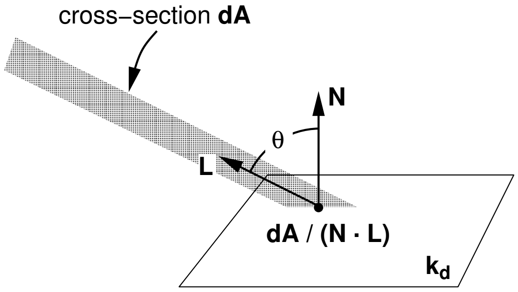

Beam with axis $-\mathbf{L}$, area dA, energy E impinges on surface element with normal $\mathbf{N}$

Beam distributes energy over area $dA/(\mathbf{N} \cdot \mathbf{L})$

Energy per unit area $\propto \mathbf{N} \cdot \mathbf{L}$

Lambert's Law

Combine Lambert's Law, with illumination geometry

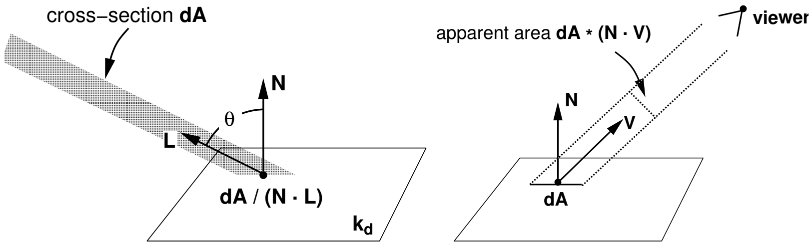

Diffuse reflected intensity $\propto \mathbf{N} \cdot \mathbf{V}$

But: *apparent area* $\propto \mathbf{N} \cdot \mathbf{V}$

Brightness is Energy per Apparent Area

Thus the observed brightness is:

$\propto \mathbf{N} \cdot \mathbf{L}$; $\quad \propto \mathbf{N} \cdot \mathbf{V}$; $\quad \propto 1/(\mathbf{N} \cdot \mathbf{V})$

In total, $\propto \mathbf{N} \cdot \mathbf{L}$

Brightness is **independent** of view direction $\mathbf{V}$!

Ideal Diffuse Reflection

Shading term:

$$I = I_l k_d \max(\mathbf{N} \cdot \mathbf{L}, 0)$$

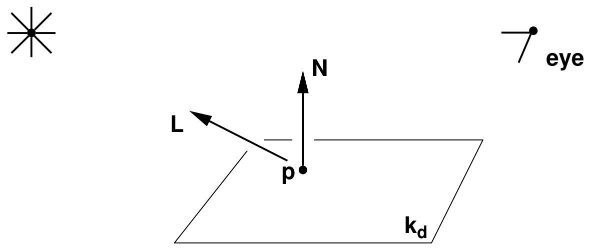

I_l is intensity of point light source

k_d is coefficient of diffuse reflection $0 \leq k_d < 1$

max prevents reflection from *backfacing* elements (underlying solid body opaque, thus "self-occluding")

$\mathbf{L}$ effectively constant for distant light sources

Ideal Specular Reflection

Occurs from ideal (perfect) mirror

Light along $\mathbf{L}$ bounces *only* in *reflected direction* $\mathbf{R}$

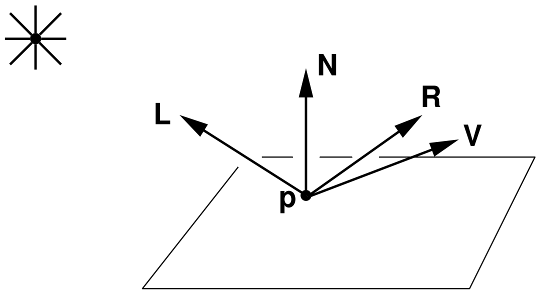

Quantitatively:

$$I = I_l \delta(R, V)$$

where $\delta(\mathbf{A}, \mathbf{B}) \equiv 1$ iff $\mathbf{A} = \mathbf{B}$, 0 otherwise

Uninteresting for most applications (non-physical)

Generalized Specular Reflection: Qualitative

Occurs from “mirror-like” surfaces

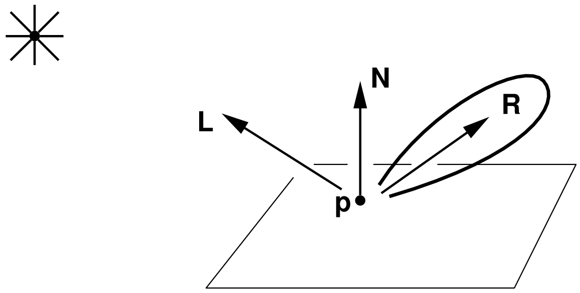

Highlights occur whose sharpness varies with surface shininess
Reflection *falls off* around *peak* reflection direction $\mathbf{R}$
Highlight *color* only weakly dependent on surface color

Specular Reflection: Quantitative

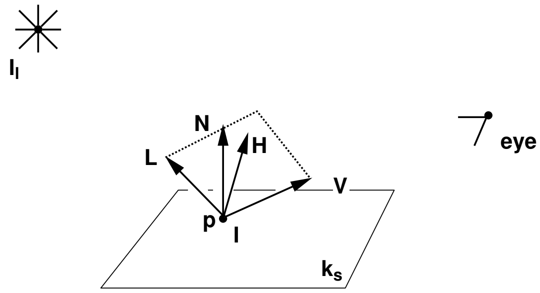

Usually written in terms of surface normal, half-vector

$$I = I_l k_s \max\left(\mathbf{N}\cdot\mathbf{H}, 0\right)^s$$

$\mathbf{H}$ = *half-vector* $(\mathbf{V}+\mathbf{L})/|\mathbf{V}+\mathbf{L}|$
$\mathbf{N}$ = surface normal
$\mathbf{V}$ = direction to viewer
$\mathbf{L}$ = direction to light source
k_s = coefficient of specular reflection
s = “shininess” parameter (sometimes e or n):
  $s = 1$ yields smooth, broad highlight
  $s = \infty$ yields perfect mirror (δ-function)
  s typically ranges from 1 to a few hundred

Phong Illumination Model (CACM 18(6), June 1975)

Diffuse term (independent of viewing direction)

$$I = I_l k_d \max(\mathbf{N}\cdot\mathbf{L}, 0)$$

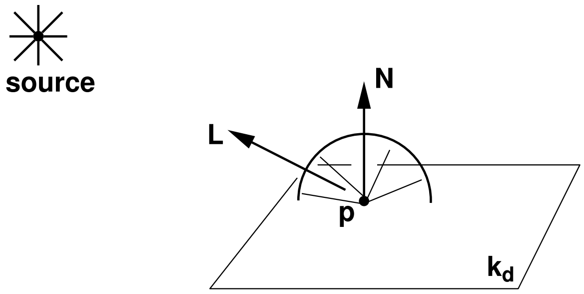

I = reflected intensity
k_d = coefficient of diffuse reflection
$\mathbf{N}$ = surface normal
$\mathbf{L}$ = direction to light source

Phong Illumination

Specular term (dependent on viewing direction)

$$I = I_l k_s \max\left(\mathbf{N}\cdot\mathbf{H}, 0\right)^s$$

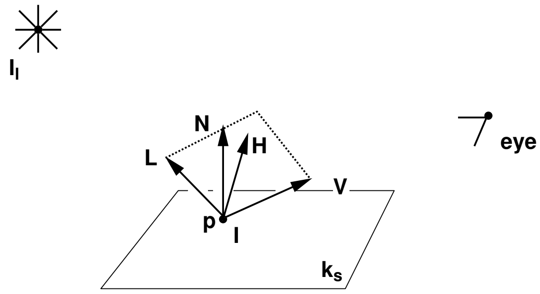

$\mathbf{H}$ = *half-vector* $(\mathbf{V}+\mathbf{L})/|\mathbf{V}+\mathbf{L}|$
$\mathbf{N}$ = surface normal
$\mathbf{V}$ = direction to viewer
$\mathbf{L}$ = direction to light source
k_s = coefficient of specular reflection
s = “shininess” parameter

Computing the reflection vector R

R is **L** mirrored about **N**

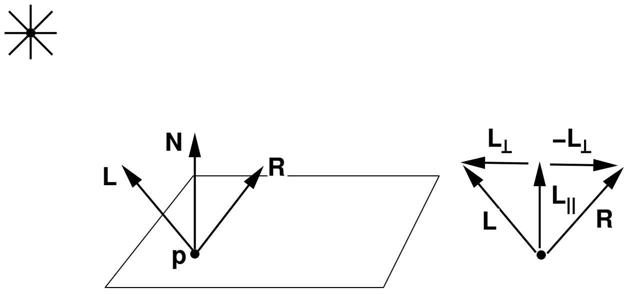

Partition **L** into components $\mathbf{L}_{\|}$ and $\mathbf{L}_{\perp}$; negate $\mathbf{L}_{\perp}$

$$\begin{aligned}\mathbf{R} &= \mathbf{L}_{\|} - \mathbf{L}_{\perp} \\ &= \mathbf{L}_{\|} - (\mathbf{L} - \mathbf{N}(\mathbf{L}\cdot\mathbf{N})) \\ &= \mathbf{N}(\mathbf{L}\cdot\mathbf{N}) - \mathbf{L} + \mathbf{N}(\mathbf{L}\cdot\mathbf{N}) \\ &= 2\mathbf{N}(\mathbf{L}\cdot\mathbf{N}) - \mathbf{L}\end{aligned}$$

L (and therefore $\mathbf{N}\cdot\mathbf{L}$) are constant
for distant sources (directional lights)

Computing the half-vector H

H is simply halfway between **L** and **V**

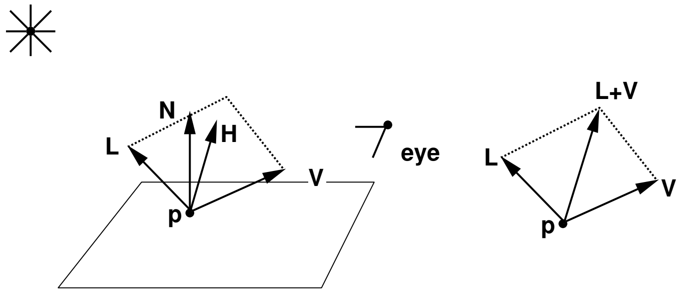

$$\mathbf{H} = \frac{\mathbf{L}+\mathbf{V}}{|\mathbf{L}+\mathbf{V}|}$$

H sometimes called "highlight vector"
hold **L**, **V** fixed, vary **N** to find maximal highlight

Distance Attenuation

Model falloff of intensity with distance
Problem: inverse quadratic falloff works poorly
at long distances: varies too slowly
at short distances: varies too quickly
due to non-physicality of point light model!
In practice, model attenuation with parameters c_k:

$$f_{att}(d_L) = \min\left(\frac{1}{c_1 + c_2 d_L + c_3 {d_L}^2}, 1\right)$$

d_L is distance from surface element **to light**
c_1: prevents blowup due to very close light sources
c_2: linear falloff due to large area sources
c_3: quadratic falloff due to small bright sources
min: prevents $f_{att}(d_L)$ from exceeding 1
Determining reasonable c_k can be problematic

Phong Illumination – Aggregate

Monochromatic lights and surfaces:

$$I = I_a k_a + f_{att}(d_L) I_l [k_d (\mathbf{N}\cdot\mathbf{L}) + k_s (\mathbf{N}\cdot\mathbf{H})^s]$$

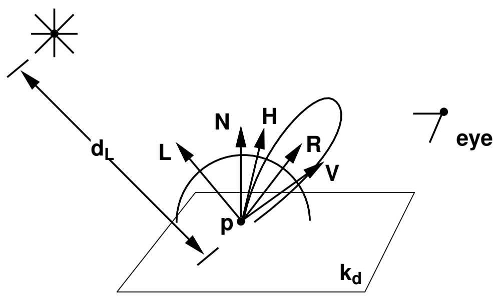

Colored lights and surfaces:

$$I_\lambda = I_{a\lambda} k_a C_{d\lambda} + f_{att}(d_L) I_{l\lambda} [k_d C_{d\lambda} (\mathbf{N}\cdot\mathbf{L}) + k_s (\mathbf{N}\cdot\mathbf{H})^s]$$

Typically, one equation for each *spectral component*
e.g., R ($\lambda = 620$nm), G ($\lambda = 550$nm), B ($\lambda = 440$nm)
Note: specular term independent of *material properties*
Approximation, with empirically derived parameters
CAUTION: Phong model is not physically correct!
Spectral integration; fall off; energy conservation

When to Apply Lighting Equation?

Apply in object space
 Adv: easy to compute normals
 Disadv: Must transform lights to object space
Apply after perspective transformation?
 No; lengths are not preserved
So: apply in world coordinates
 Adv: lights usually specified in world space
 This is most common practice, but:
  1) What about vertices outside frustum?
  2) How are obj. space normals handled?
Some systems shade pixel fragments at end of pipe!

When to Apply Lighting Equation?

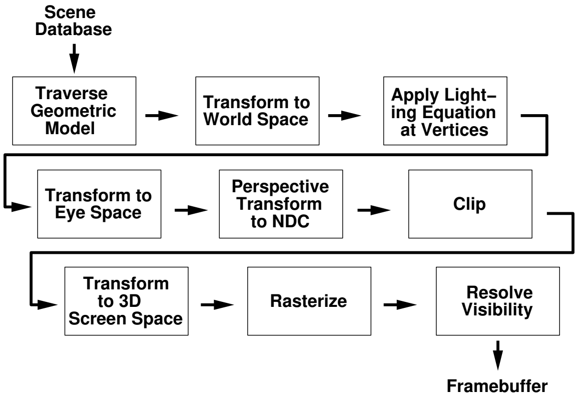

Transforming Normal Vectors

Suppose we apply a transformation $\mathbf{M}$ to an object
What are the normals of the transformed object?
Suppose we are shearing in x. Then

$$\mathbf{M} = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

Look at normals $(0, 1)$ and $(1, 0)$ of a square:

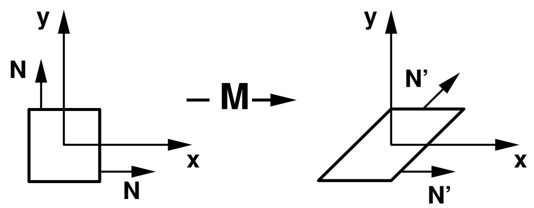

If we apply $\mathbf{M}$, we get $(1, 1)$ and $(1, 0)$!

Transforming Normal Vectors

Consider vector $\mathbf{v}$, and $\mathbf{n}$ such that $\mathbf{n} \perp \mathbf{v}$

$$\mathbf{n}^T \mathbf{v} = 0$$

for any (non-singular) transformation matrix $\mathbf{M}$

$$\mathbf{n}^T \mathbf{M}^{-1} \mathbf{M} \mathbf{v} = 0$$

which can be rewritten as

$$(\mathbf{n}^T \mathbf{M}^{-1})(\mathbf{M}\mathbf{v}) = 0$$

thus the *transpose* of the transformed normal is

$$\mathbf{n}^T \mathbf{M}^{-1}$$

and the transformed normal itself is

$$(\mathbf{n}^T \mathbf{M}^{-1})^T = \mathbf{M}^{-T} \mathbf{n}$$

Normal $\mathbf{n}$, under action of $\mathbf{M}$, becomes $\mathbf{M}^{-T} \mathbf{n}$

Example: Shear

For example above,

$$\mathbf{M} = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

Inverse $\mathbf{M^{-1}}$ is

$$\mathbf{M^{-1}} = \begin{pmatrix} 1 & -1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

Inverse Transpose $\mathbf{M^{-T}}$ is

$$\mathbf{M^{-T}} = \begin{pmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$$\begin{aligned} x' &= x \\ y' &= -x + y \end{aligned}$$

Now normals $(0, 1)$ and $(1, 0)$ map to:

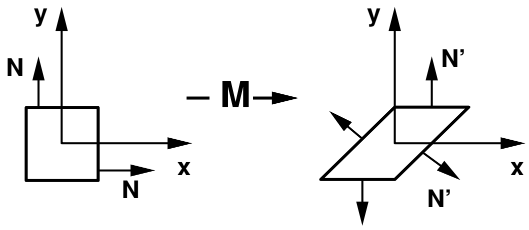

$(0, 1)$ and $(1, -1)$, respectively
Note length change! How do we fix this?

Lighting: Simplifying Assumptions

Incorrect modeling of *light source geometry*
  Can't get finite energy from zero volume!
Incorrect modeling of *absorptive media* (haze, etc.)
  Bogus intensity attenuation with distance
No modeling of *scattering* due to air, etc.
  Decreasing contrast, increasing blueness with distance
No dependence on wavelength λ, thin film, etc.
No dependence on polarization
No dependence on azimuthal angle of incidence
No dependence on time
Things are not so simple in reality!