Welcome to:

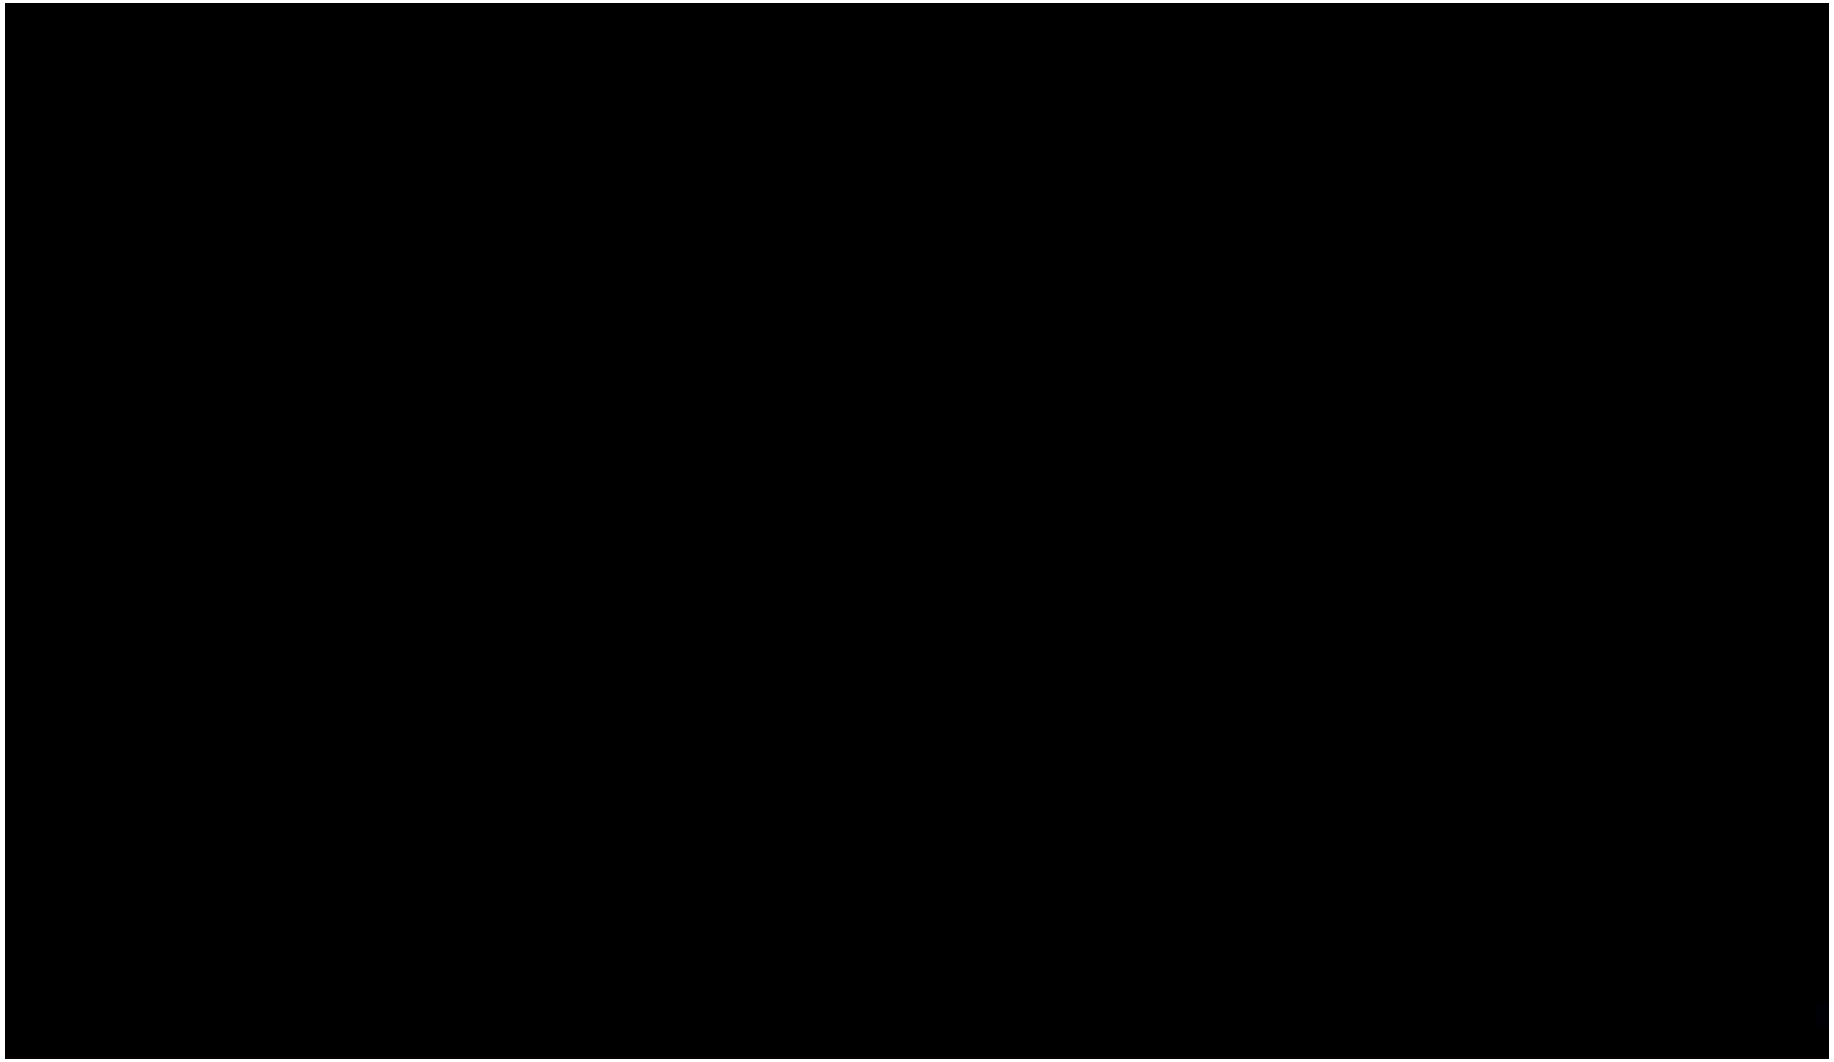

Welcome to:

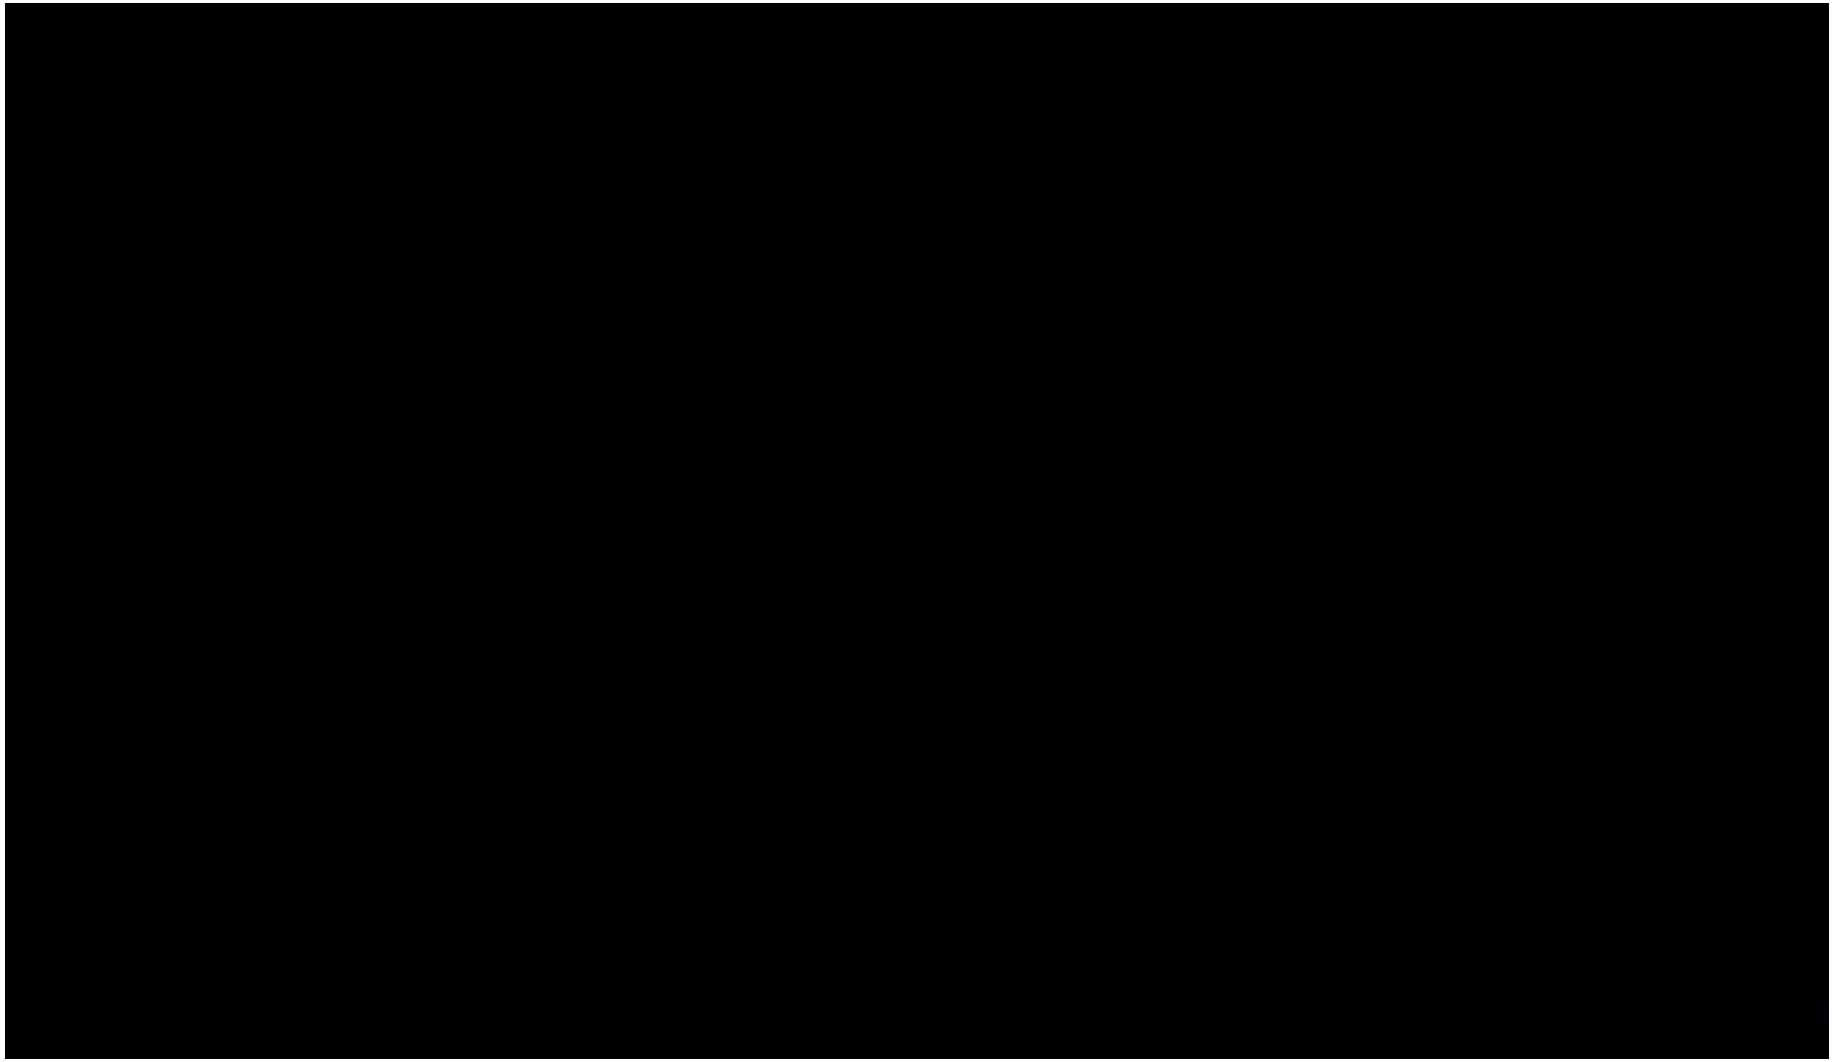

Welcome to:

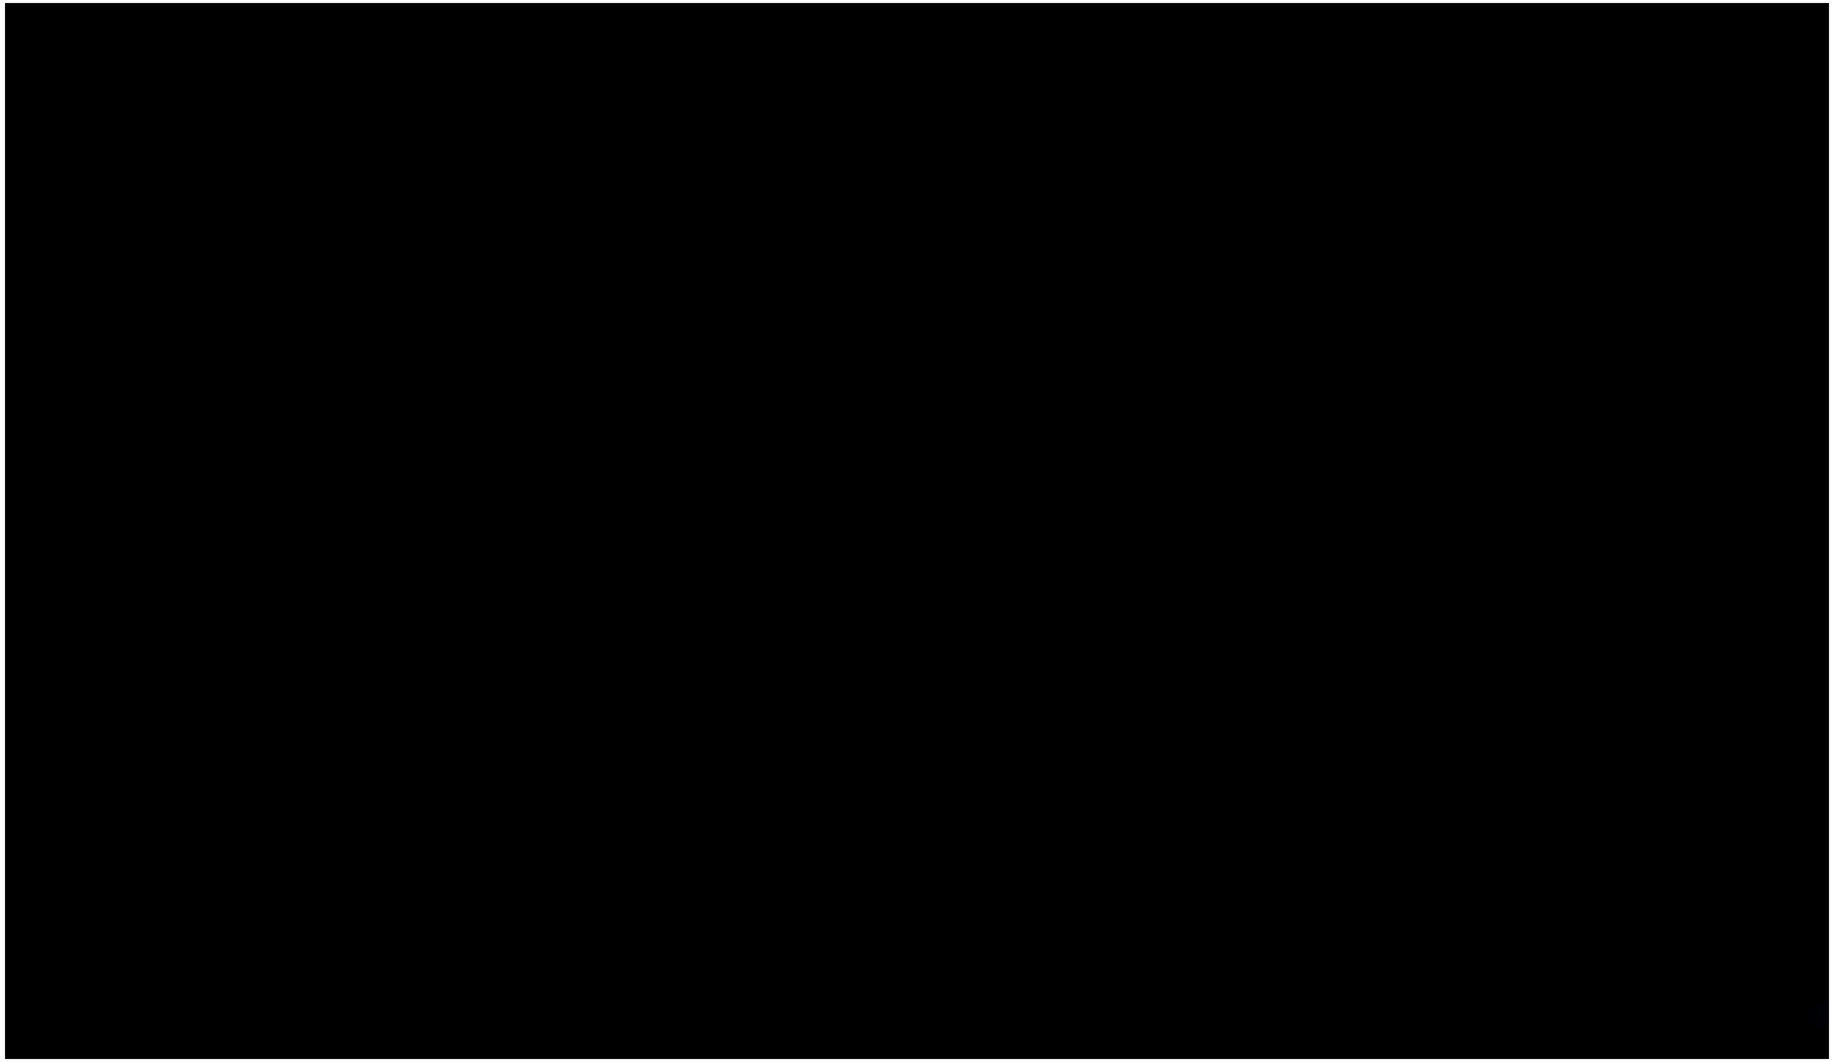

I’m Assaf, nice to meet you ☺

I’m Assaf, nice to meet you ☺

I’m Assaf, nice to meet you ☺

t

I’m Assaf, nice to meet you ☺

I’m Assaf, nice to meet you ☺

I'm Assaf, nice to meet you ☺

I’m Assaf, nice to meet you ☺

t

I'm Assaf, nice to meet you ☺

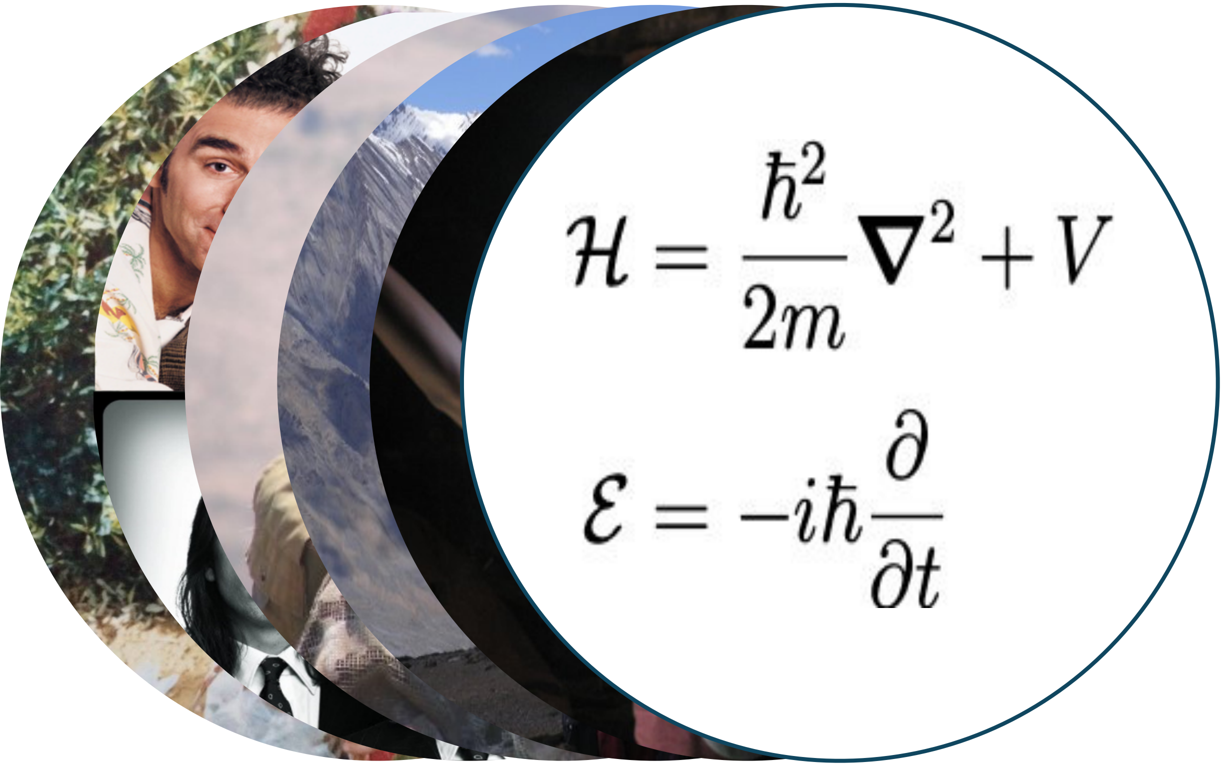

$$\mathcal{H} = \frac{\hbar^2}{2m}\nabla^2 + V$$

$$\mathcal{E} = -i\hbar\frac{\partial}{\partial t}$$

t

I’m Assaf, nice to meet you ☺

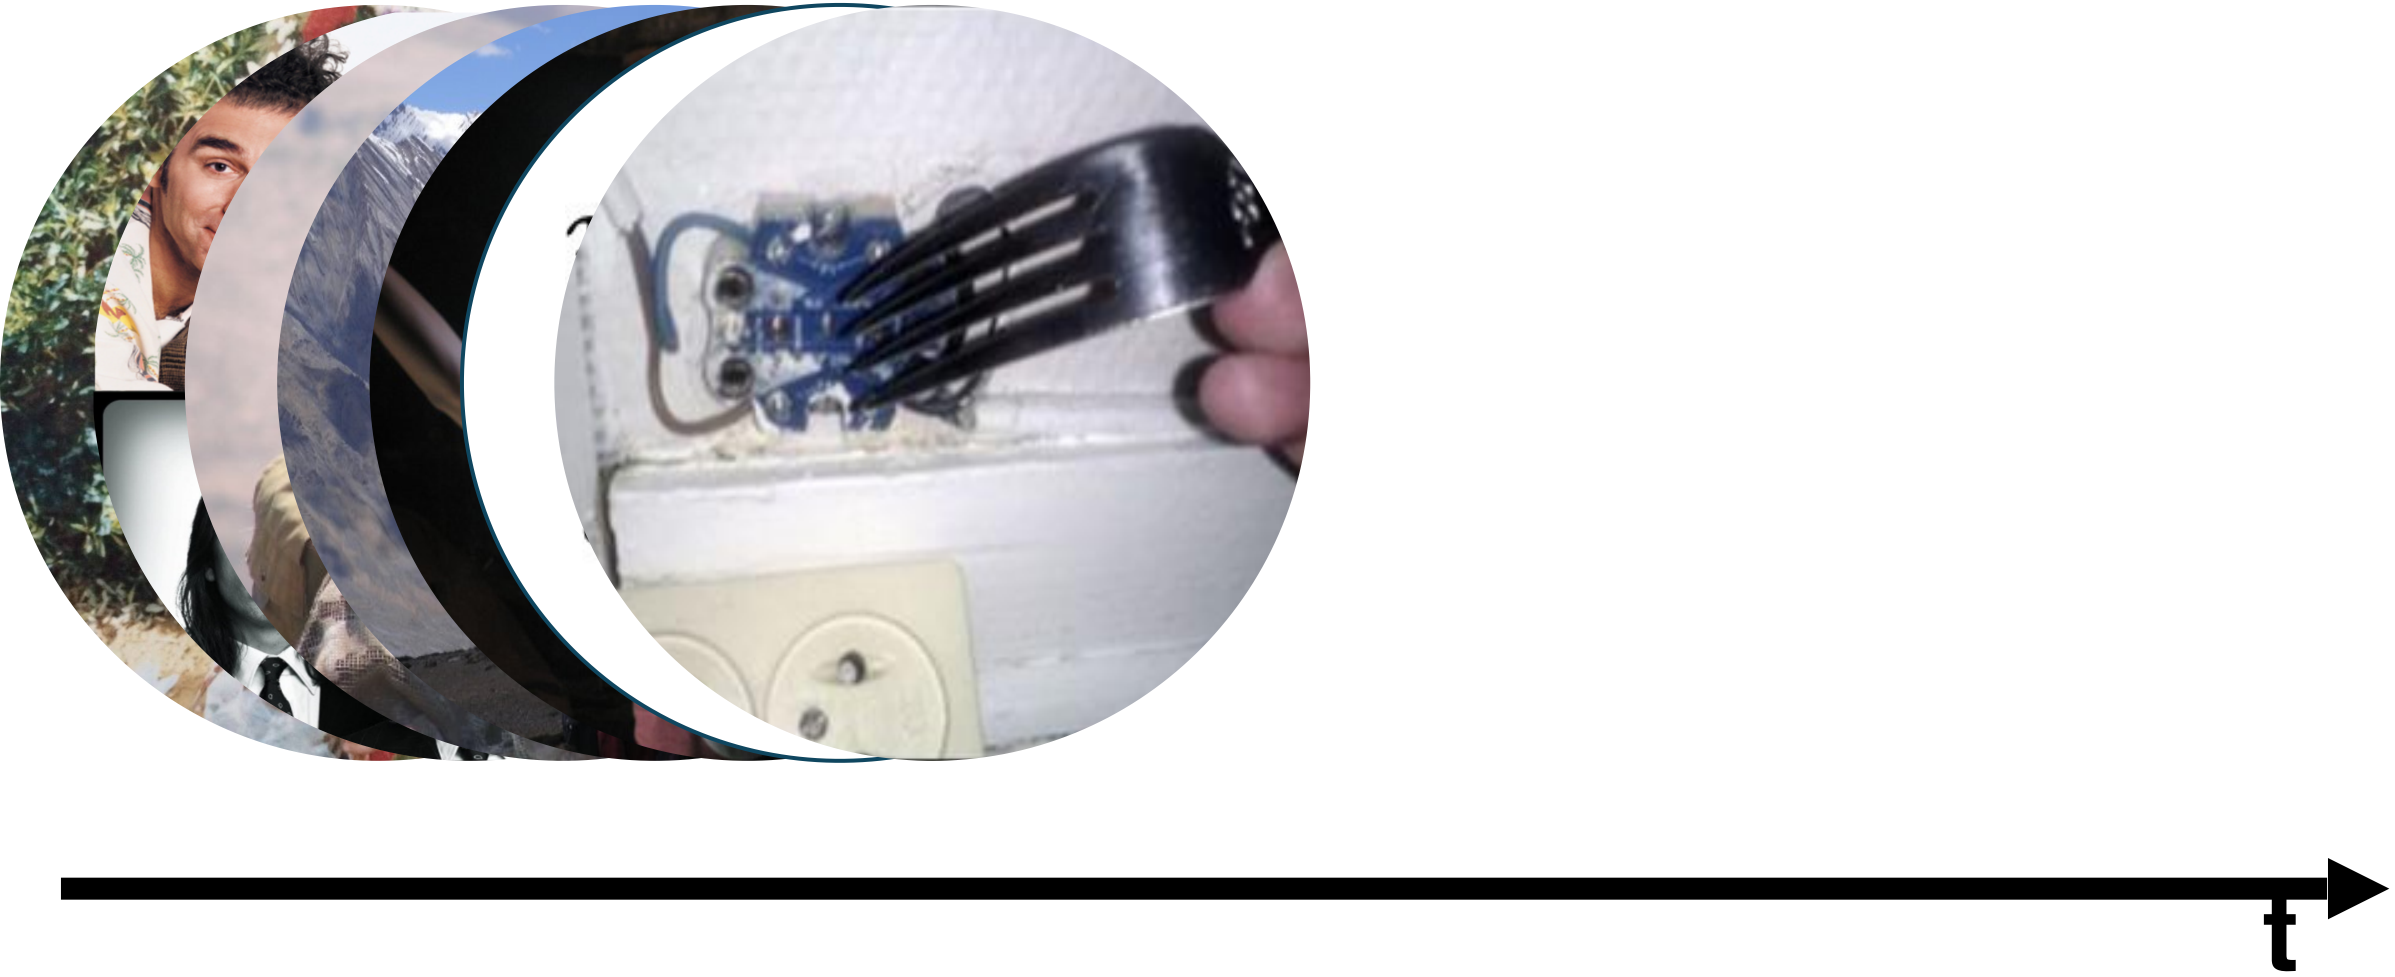

I’m Assaf, nice to meet you ☺

t

I’m Assaf, nice to meet you ☺

t

I’m Assaf, nice to meet you ☺

I’m Assaf, nice to meet you ☺

t

I’m Assaf, nice to meet you ☺

t

I’m Assaf, nice to meet you ☺

I’m Assaf, nice to meet you ☺

I’m Assaf, nice to meet you ☺

I'm Assaf, nice to meet you ☺

t

I'm Assaf, nice to meet you ☺

I’m Assaf, nice to meet you ☺

I’m Assaf, nice to meet you ☺

I'm Assaf, nice to meet you ☺

I’m Assaf, nice to meet you ☺

TAs:

Amit Shmidov

Oren Chikli

Today:

00:00) Course administration and context

00:10) History and motivation

00:20) The physics of light

00:30) Image formation and basic optics

01:00) ---- Break ----

01:15) Convolution

02:00) ---- Break ----

02:15) The digital image pipeline

02:45) ---- Break ----

03:00) Tutorial (**Amit**): Basic PyTorch and point ops

Administration

MCV Course Website

Administration

- Everything is on the course website.

MCV Course Website

Administration

- Everything is on the course website.
- In-person (when no war is going on).

MCV Course Website

Administration

- Everything is on the course website.
- In-person (when no war is going on).
- 3.5 credit points.

MCV Course Website

Administration

- Everything is on the course website.
- In-person (when no war is going on).
- 3.5 credit points.
- 3 hrs. Lecture + 1 hr. tutorial
  (Tutorial covers new material).

MCV Course Website

Administration

- Everything is on the course website.
- In-person (when no war is going on).
- 3.5 credit points.
- 3 hrs. Lecture + 1 hr. tutorial
  (Tutorial covers new material).
- 5 homework assignments + Final Project.
  Homework is demanding, but worth it.

MCV Course Website

Administration

- Everything is on the course website.
- In-person (when no war is going on).
- 3.5 credit points.
- 3 hrs. Lecture + 1 hr. tutorial
  (Tutorial covers new material).
- 5 homework assignments + Final Project.
  Homework is demanding, but worth it.
- Total 7 days late HW submission for the semester.

MCV Course Website

Administration

- Everything is on the course website.
- In-person (when no war is going on).
- 3.5 credit points.
- 3 hrs. Lecture + 1 hr. tutorial
  (Tutorial covers new material).
- 5 homework assignments + Final Project.
  Homework is demanding, but worth it.
- Total 7 days late HW submission for the semester.

MCV Course Website

”If we have seen further, it is by standing on the shoulders of Giants"

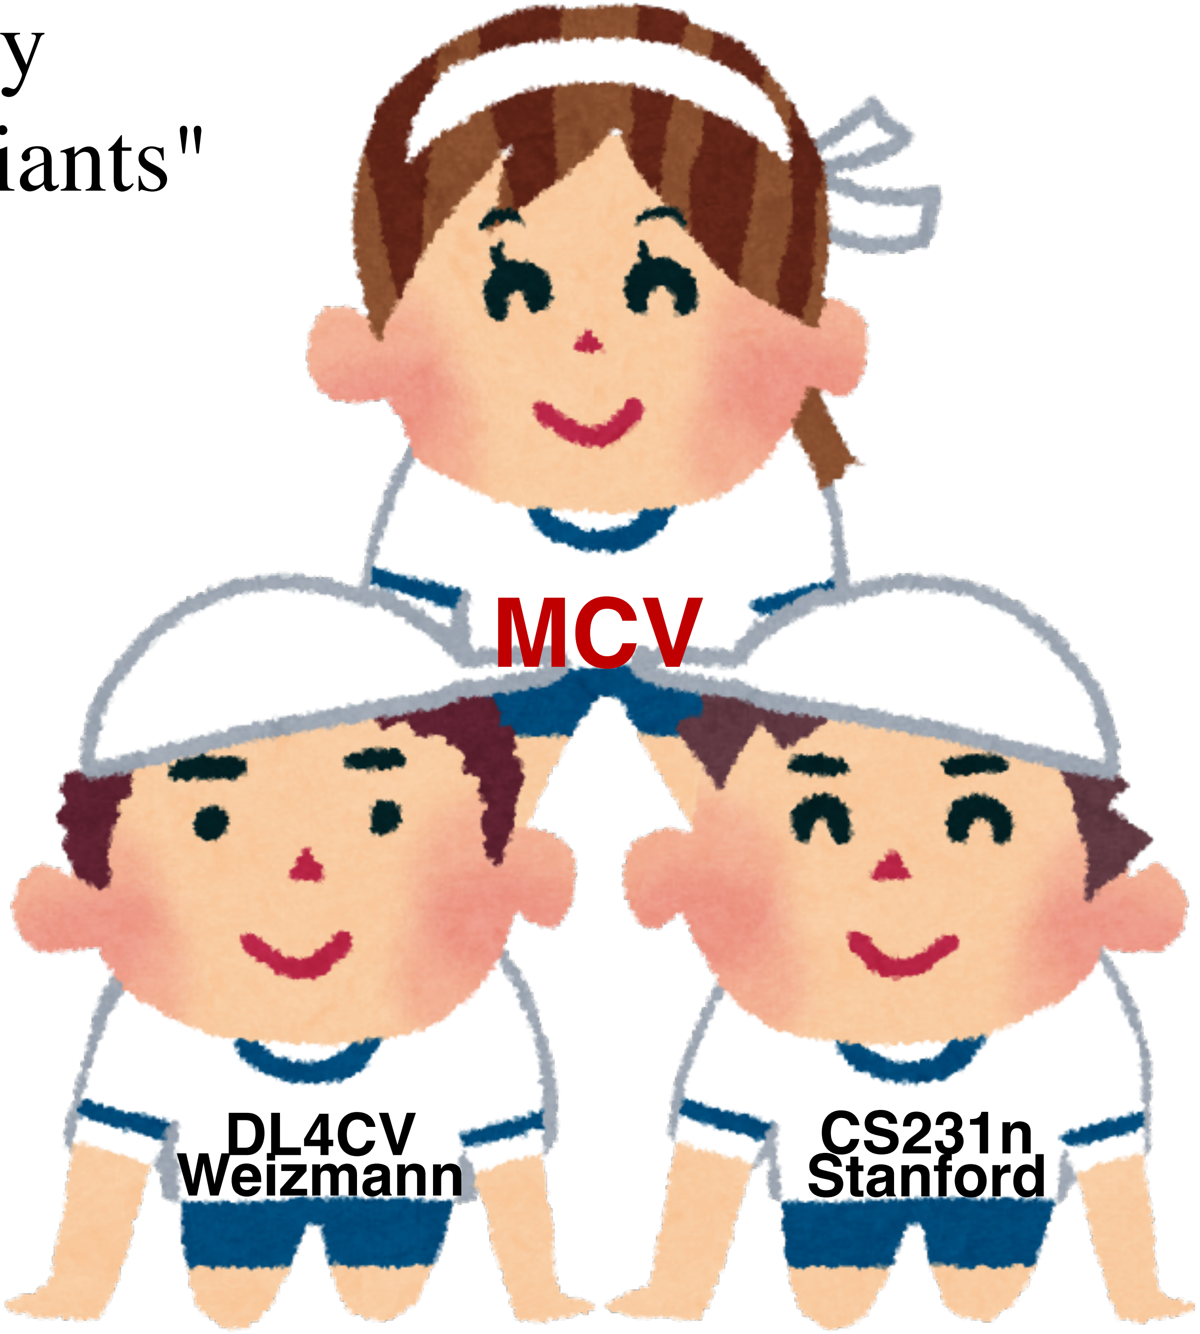

"If we have seen further, it is by standing on the shoulders of Giants"

- From basics to most recent SotA

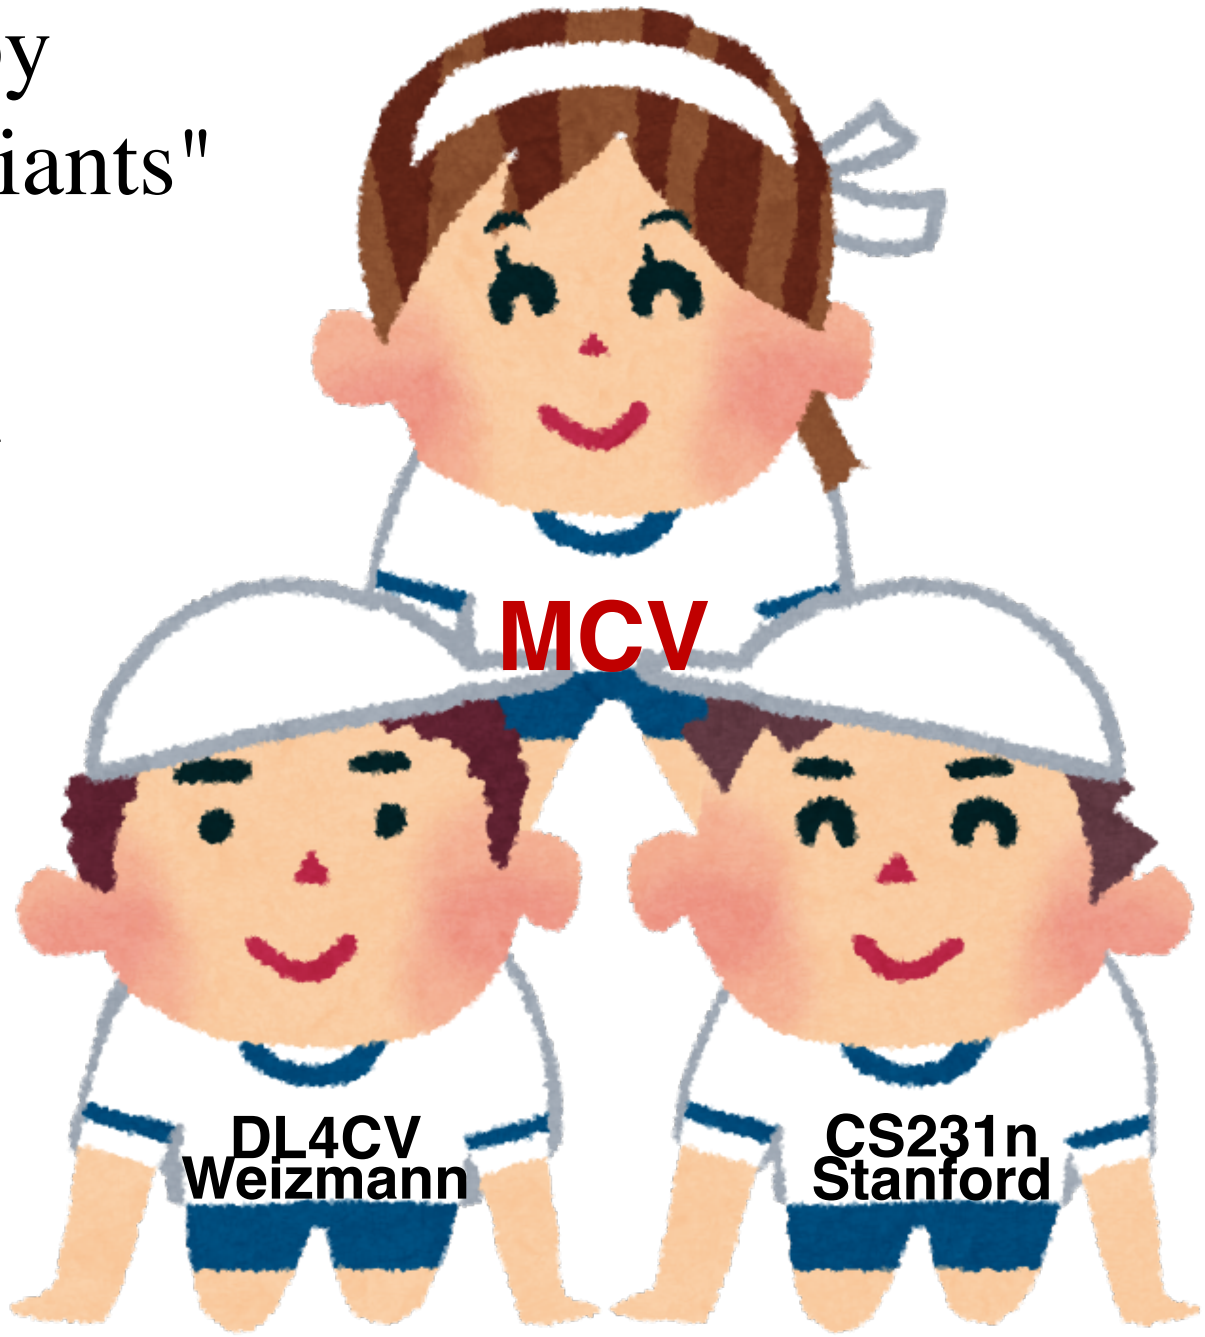

”If we have seen further, it is by standing on the shoulders of Giants"

- From basics to most recent SotA
- Intuition

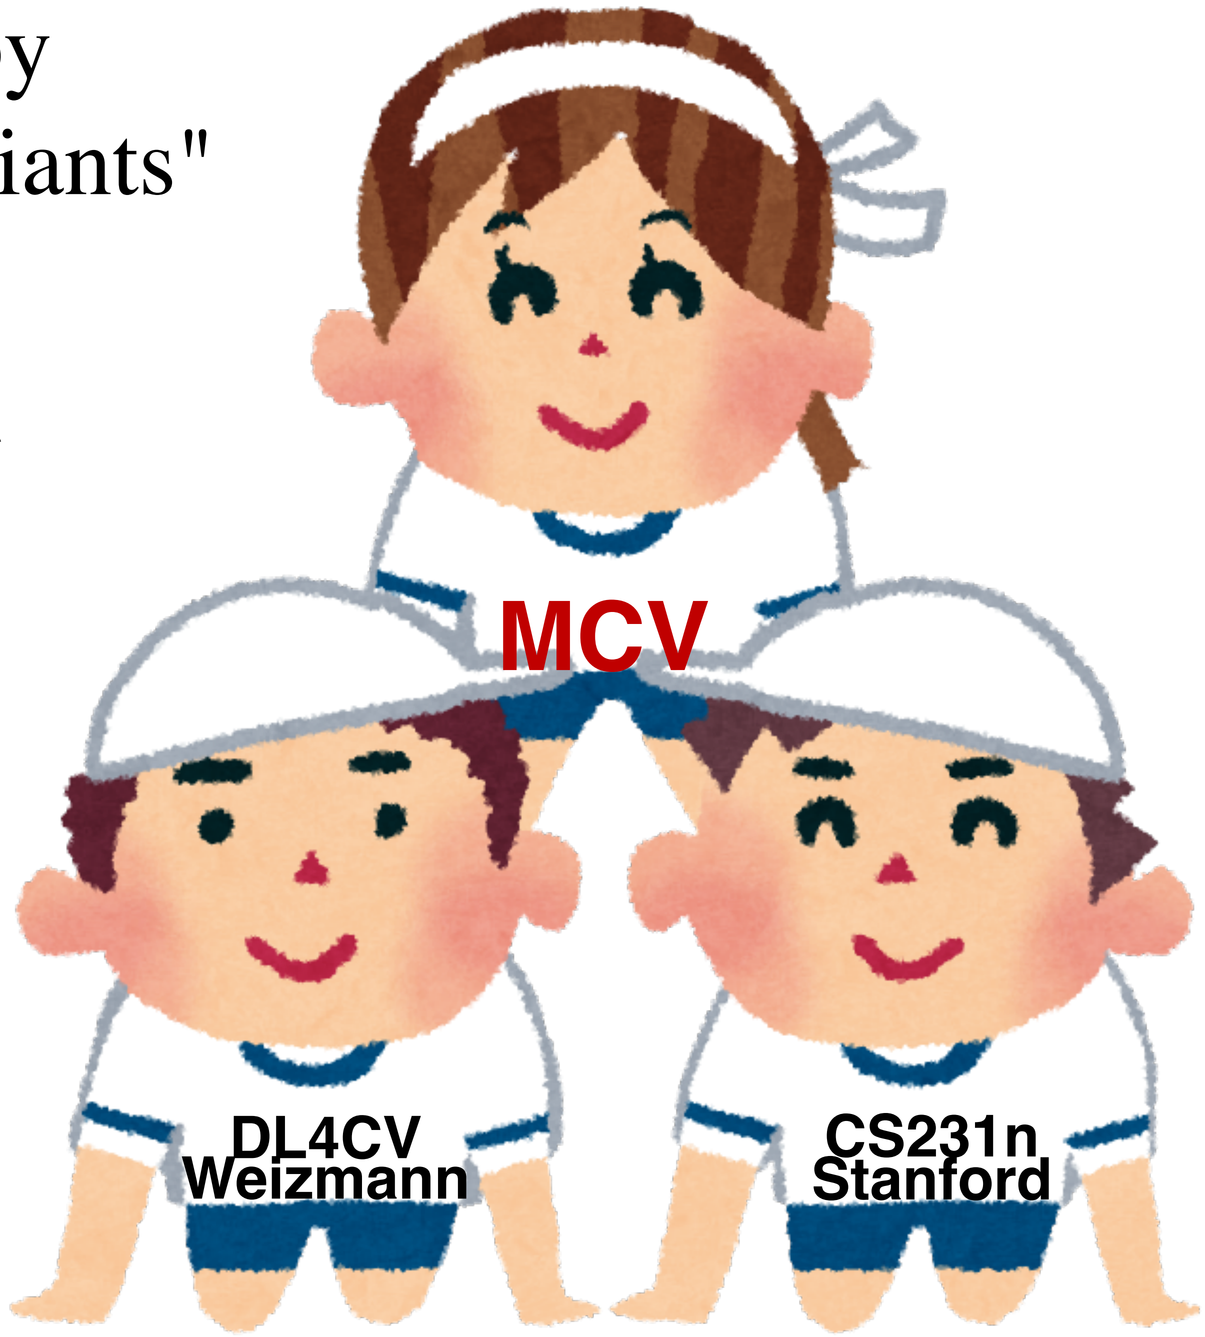

”If we have seen further, it is by standing on the shoulders of Giants"

- From basics to most recent SotA
- Intuition
- Hands-on

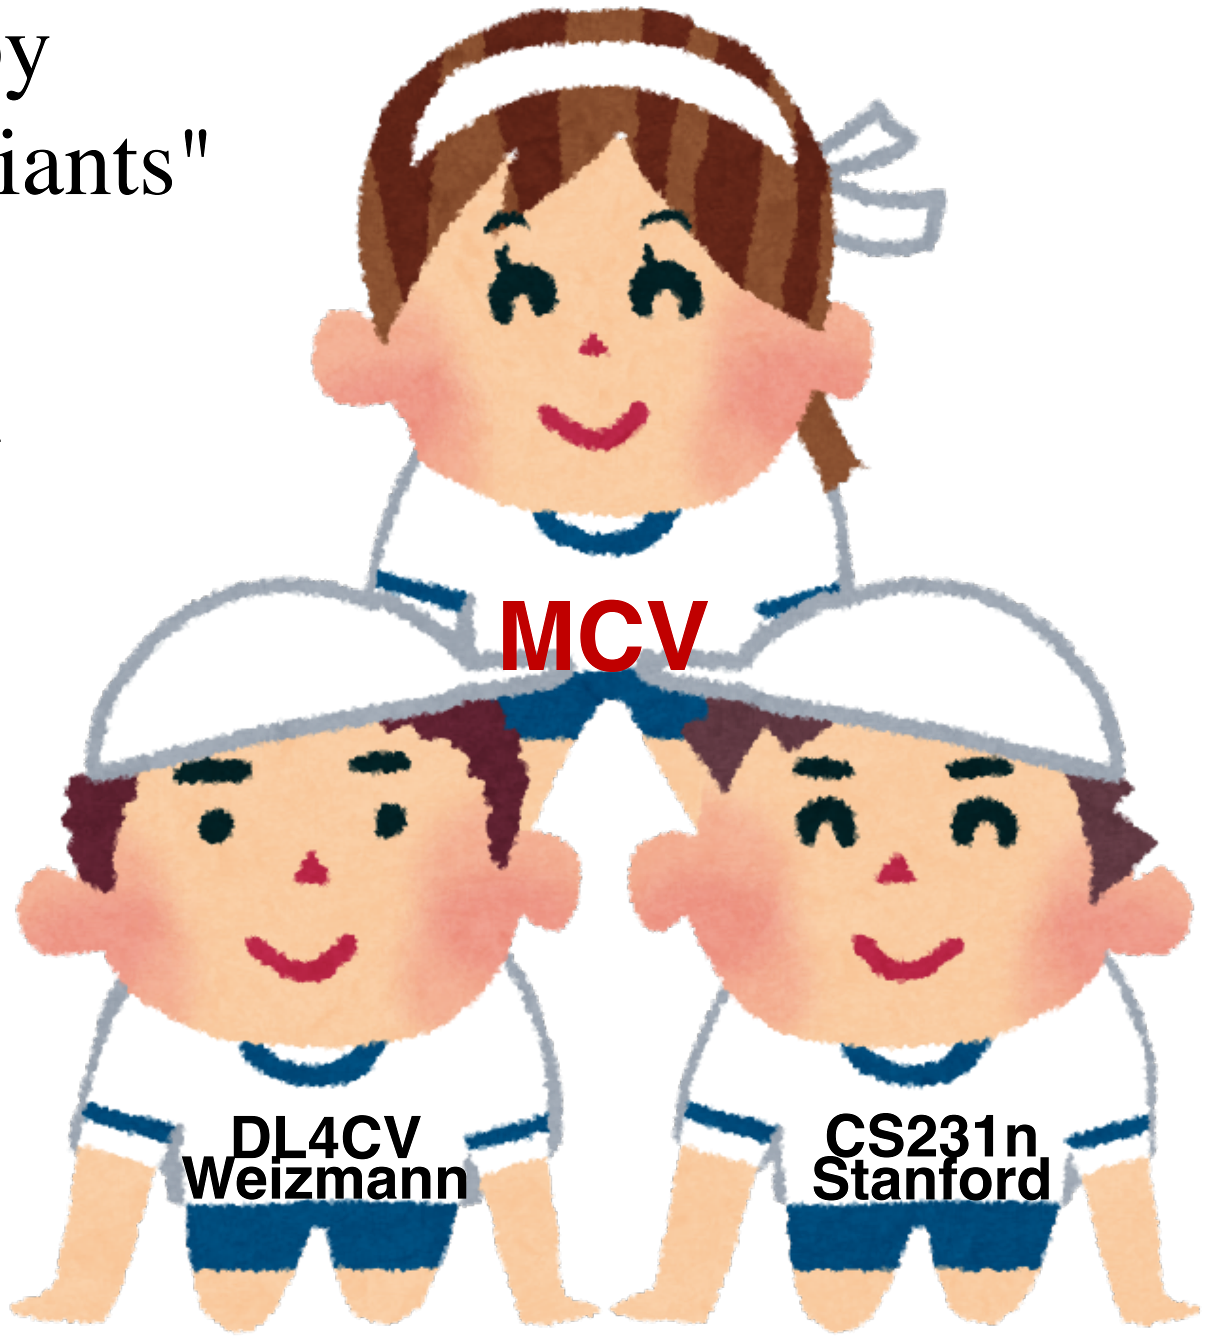

”If we have seen further, it is by standing on the shoulders of Giants"

- From basics to most recent SotA
- Intuition
- Hands-on
- We get to the bottom of things

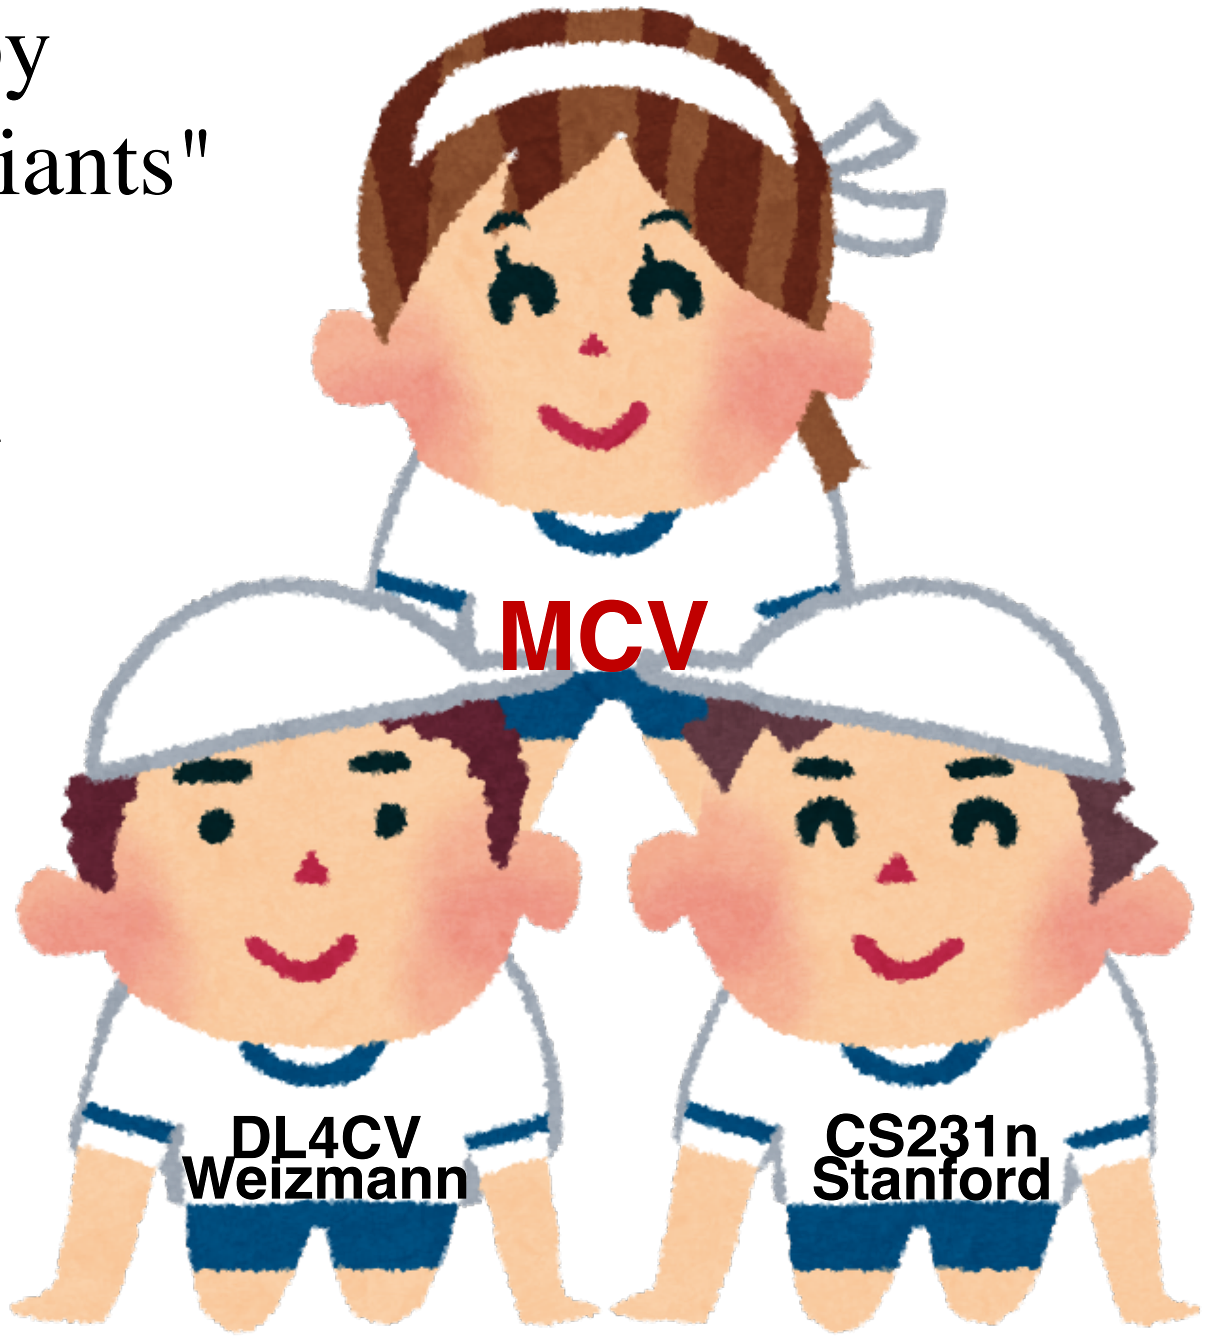

”If we have seen further, it is by standing on the shoulders of Giants"

- From basics to most recent SotA
- Intuition
- Hands-on
- We get to the bottom of things
- Research oriented

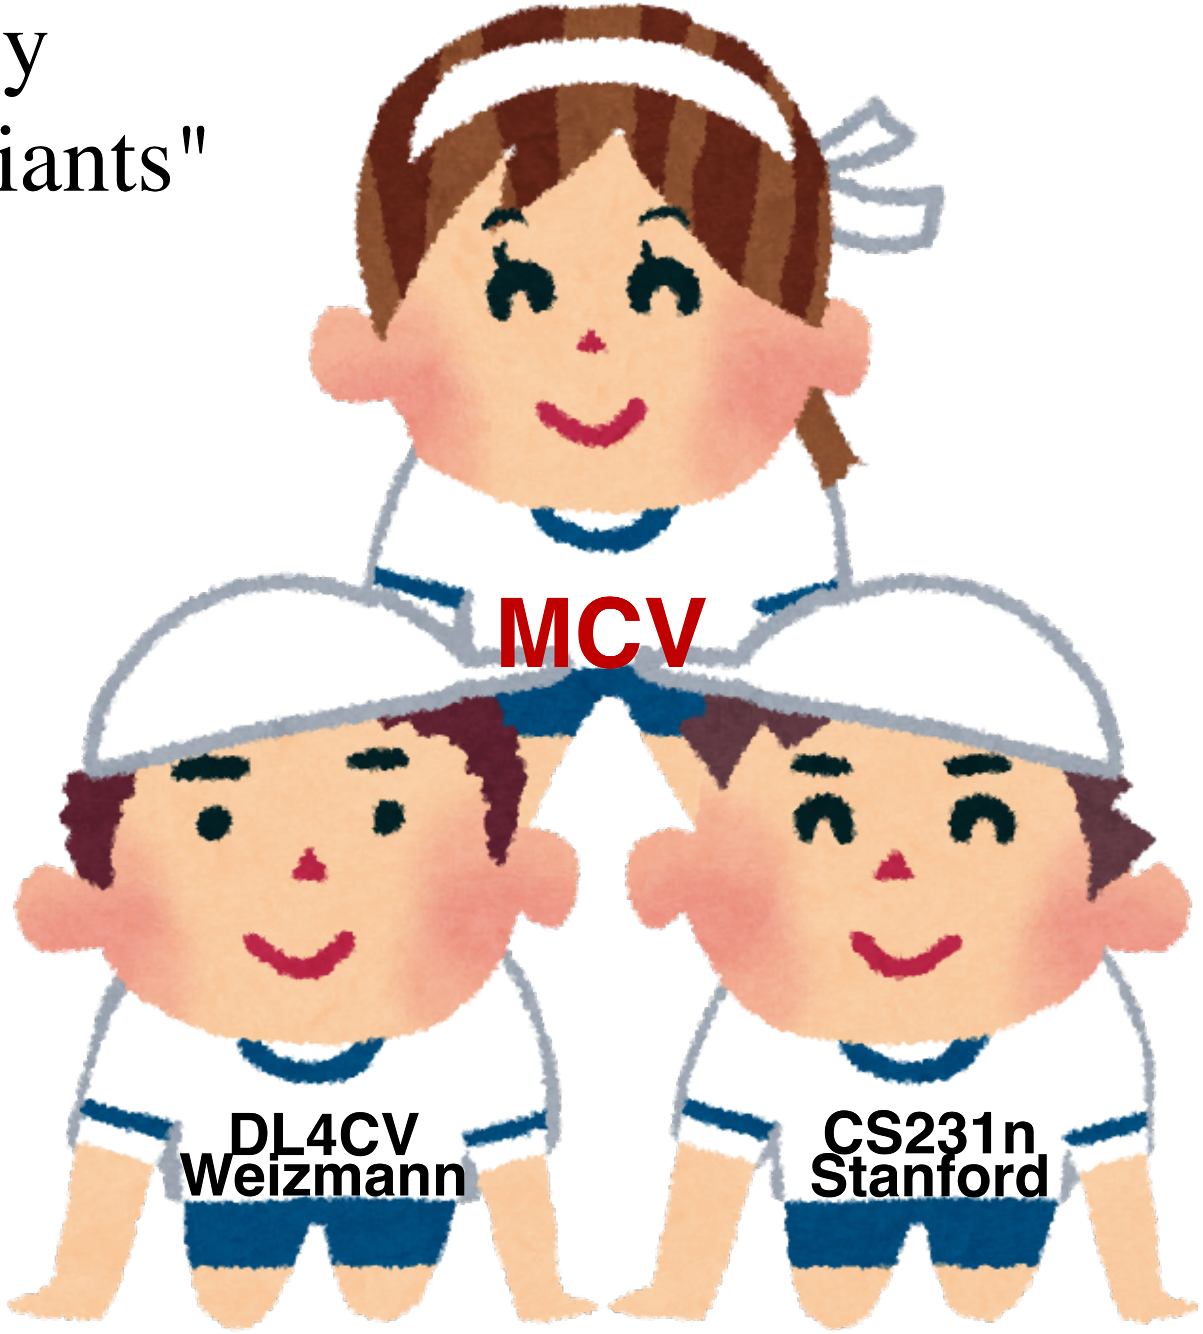

”If we have seen further, it is by standing on the shoulders of Giants"

- From basics to most recent SotA
- Intuition
- Hands-on
- We get to the bottom of things
- Research oriented
- Openness

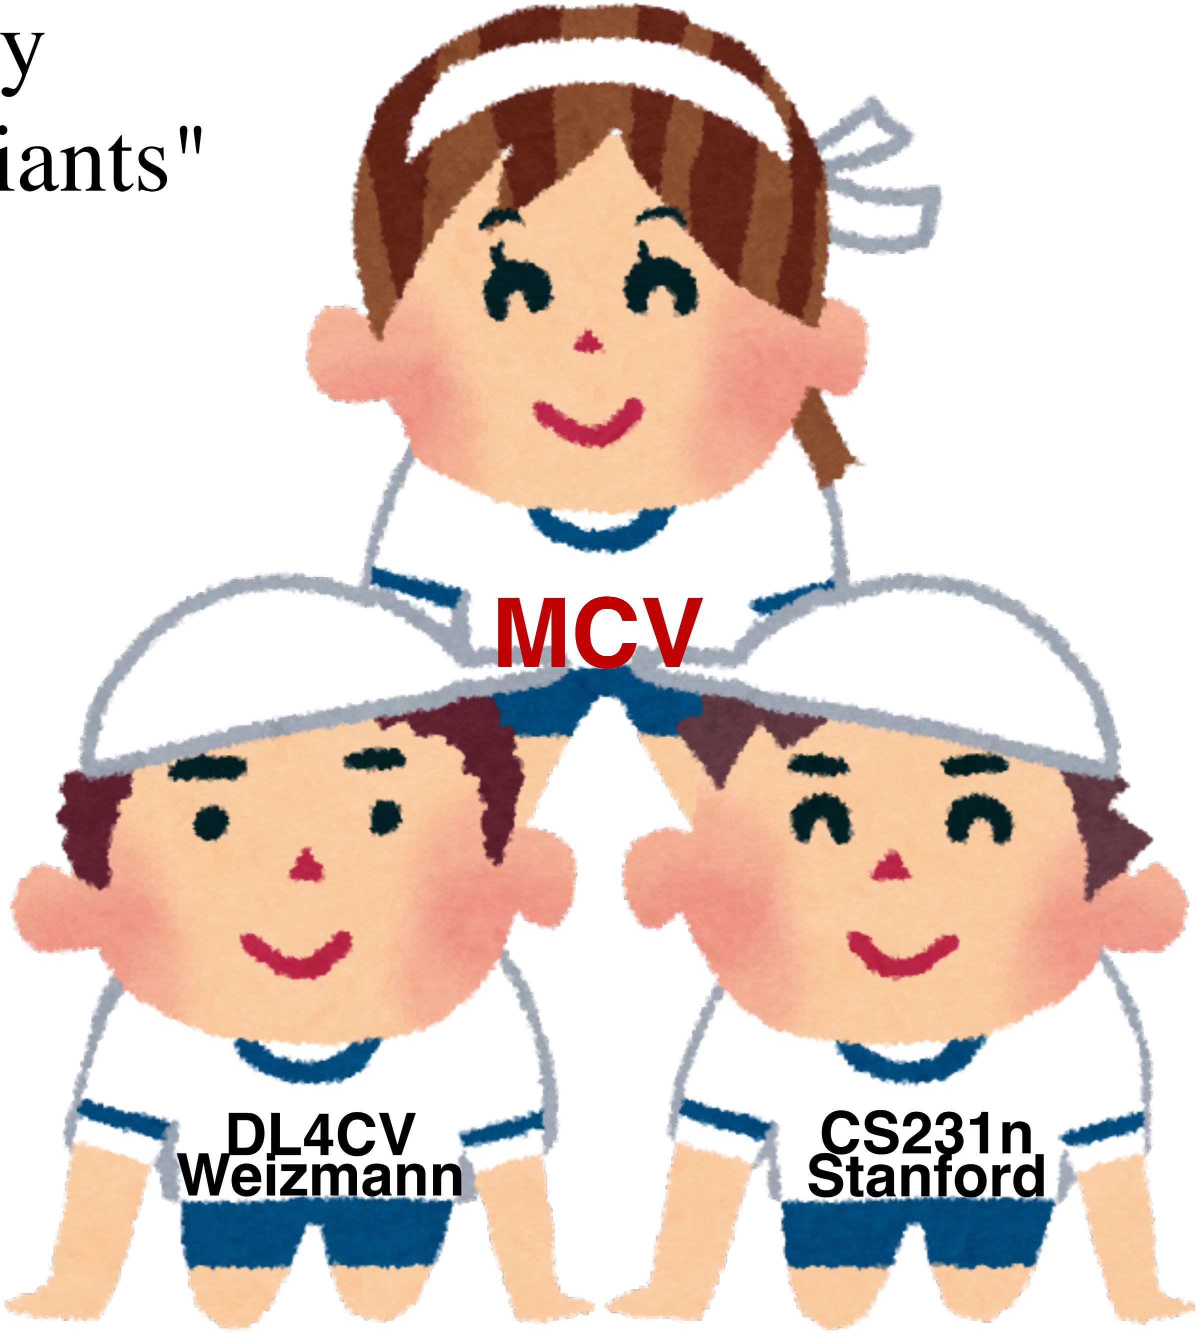

We Assume you…

We Assume you…

- Know Basic Calculus (e.g. know what a Gradient is).

We Assume you…

- Know Basic Calculus (e.g. know what a Gradient is).
- Know Basic Algebra (e.g. Vector spaces, Matrix multiplication, Eigen decomposition).

We Assume you…

- Know Basic Calculus (e.g. know what a Gradient is).
- Know Basic Algebra (e.g. Vector spaces, Matrix multiplication, Eigen decomposition).
- Know Basic probability (e.g. Cov matrix, Conditional expectation)

We Assume you…

- Know Basic Calculus (e.g. know what a Gradient is).
- Know Basic Algebra (e.g. Vector spaces, Matrix multiplication, Eigen decomposition).
- Know Basic probability (e.g. Cov matrix, Conditional expectation)
- Written code before (preferably Python).

We Assume you…

- Know Basic Calculus (e.g. know what a Gradient is).
- Know Basic Algebra (e.g. Vector spaces, Matrix multiplication, Eigen decomposition).
- Know Basic probability (e.g. Cov matrix, Conditional expectation)
- Written code before (preferably Python).
- Bumped into Machine Learning (e.g. heard the term "Overfitting").

We Assume you…

- Know Basic Calculus (e.g. know what a Gradient is).
- Know Basic Algebra (e.g. Vector spaces, Matrix multiplication, Eigen decomposition).
- Know Basic probability (e.g. Cov matrix, Conditional expectation)
- Written code before (preferably Python).
- Bumped into Machine Learning (e.g. heard the term “Overfitting”).

Homework

We Assume you…

- <u>Know</u> Basic Calculus (e.g. know what a Gradient is).
- <u>Know</u> Basic Algebra (e.g. Vector spaces, Matrix multiplication, Eigen decomposition).
- <u>Know</u> Basic probability (e.g. Cov matrix, Conditional expectation)
- <u>Written</u> code before (preferably Python).
- <u>Bumped into</u> Machine Learning (e.g. heard the term "Overfitting").

Homework

Theory

We Assume you…

- Know Basic Calculus (e.g. know what a Gradient is).
- Know Basic Algebra (e.g. Vector spaces, Matrix multiplication, Eigen decomposition).
- Know Basic probability (e.g. Cov matrix, Conditional expectation)
- Written code before (preferably Python).
- Bumped into Machine Learning (e.g. heard the term “Overfitting”).

Homework

Theory	From Scratch

We Assume you…

- Know Basic Calculus (e.g. know what a Gradient is).
- Know Basic Algebra (e.g. Vector spaces, Matrix multiplication, Eigen decomposition).
- Know Basic probability (e.g. Cov matrix, Conditional expectation)
- Written code before (preferably Python).
- Bumped into Machine Learning (e.g. heard the term "Overfitting").

Homework

Theory	From Scratch	Applied

We Assume you…

- Know Basic Calculus (e.g. know what a Gradient is).
- Know Basic Algebra (e.g. Vector spaces, Matrix multiplication, Eigen decomposition).
- Know Basic probability (e.g. Cov matrix, Conditional expectation)
- Written code before (preferably Python).
- Bumped into Machine Learning (e.g. heard the term "Overfitting").

Homework

Theory	From Scratch	Applied	HW1 is online!

Road map

Slide idea: Justin Johnson

Road map

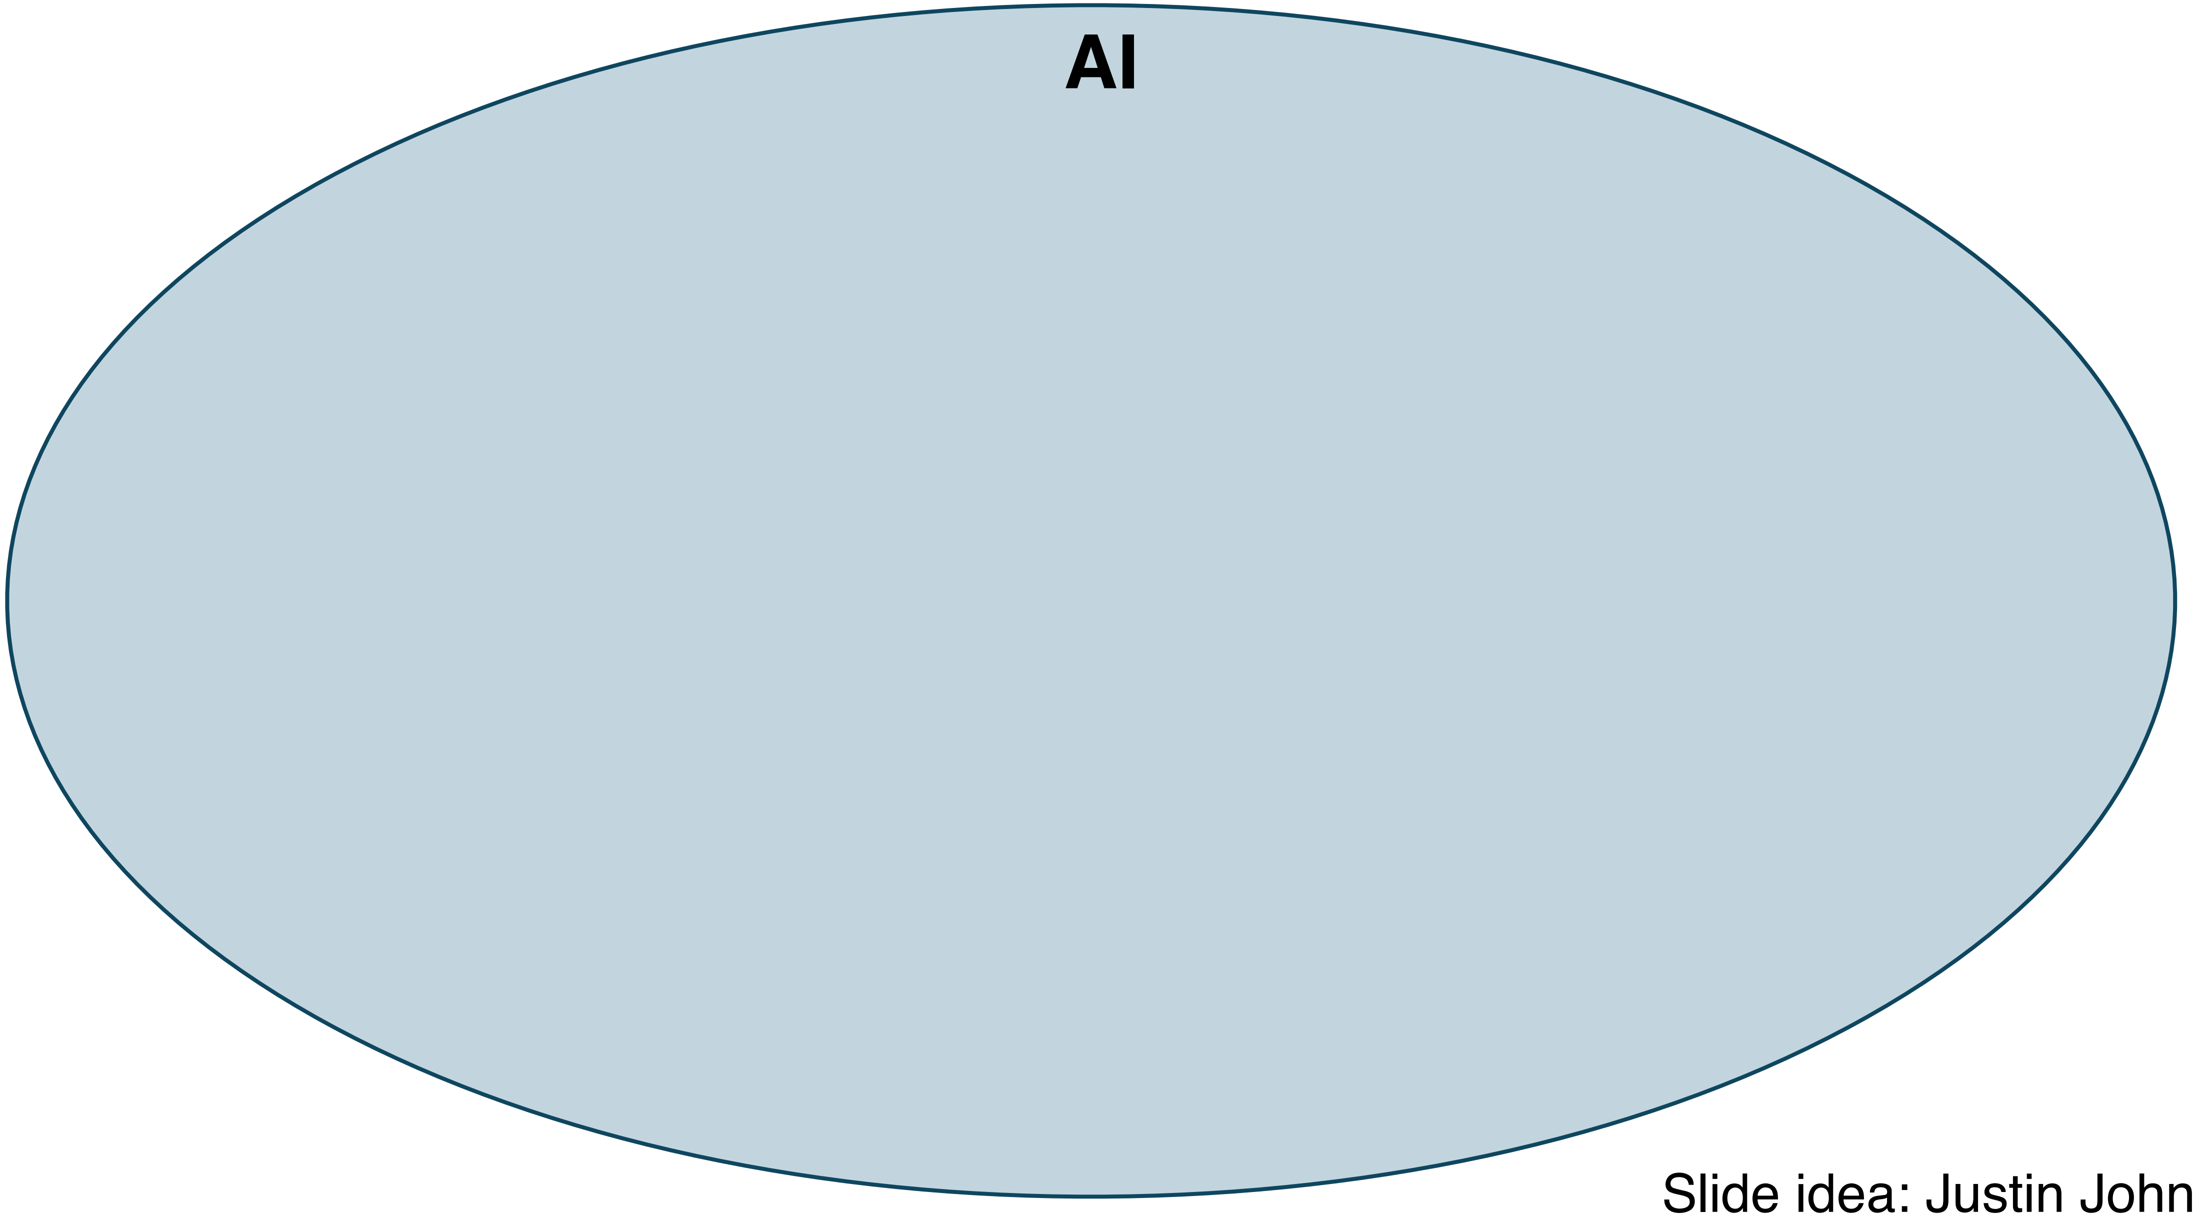

Slide idea: Justin Johnson

Road map

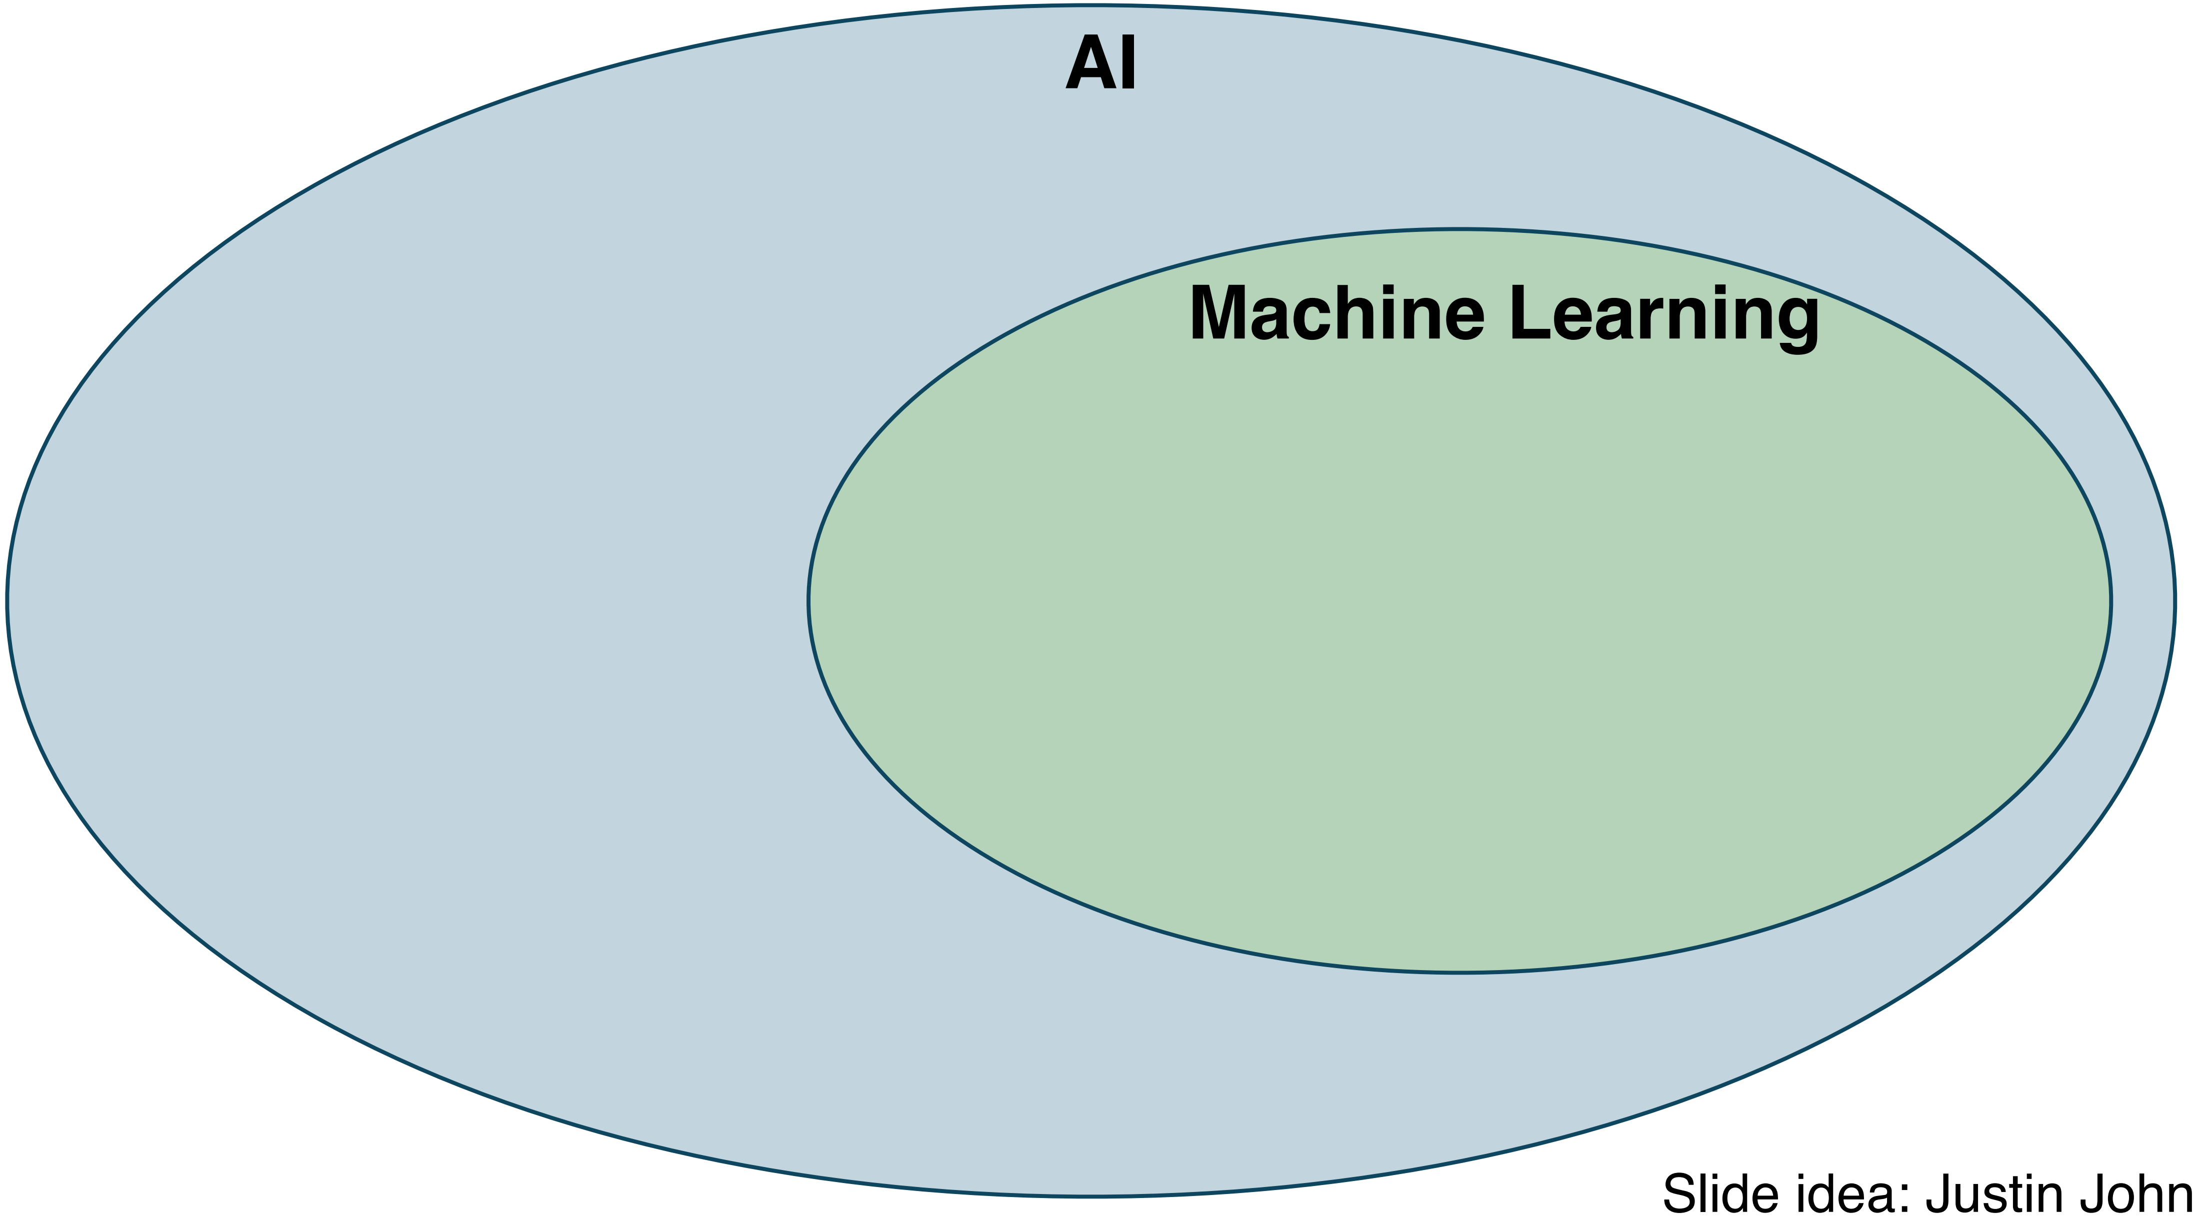

Slide idea: Justin Johnson

Road map

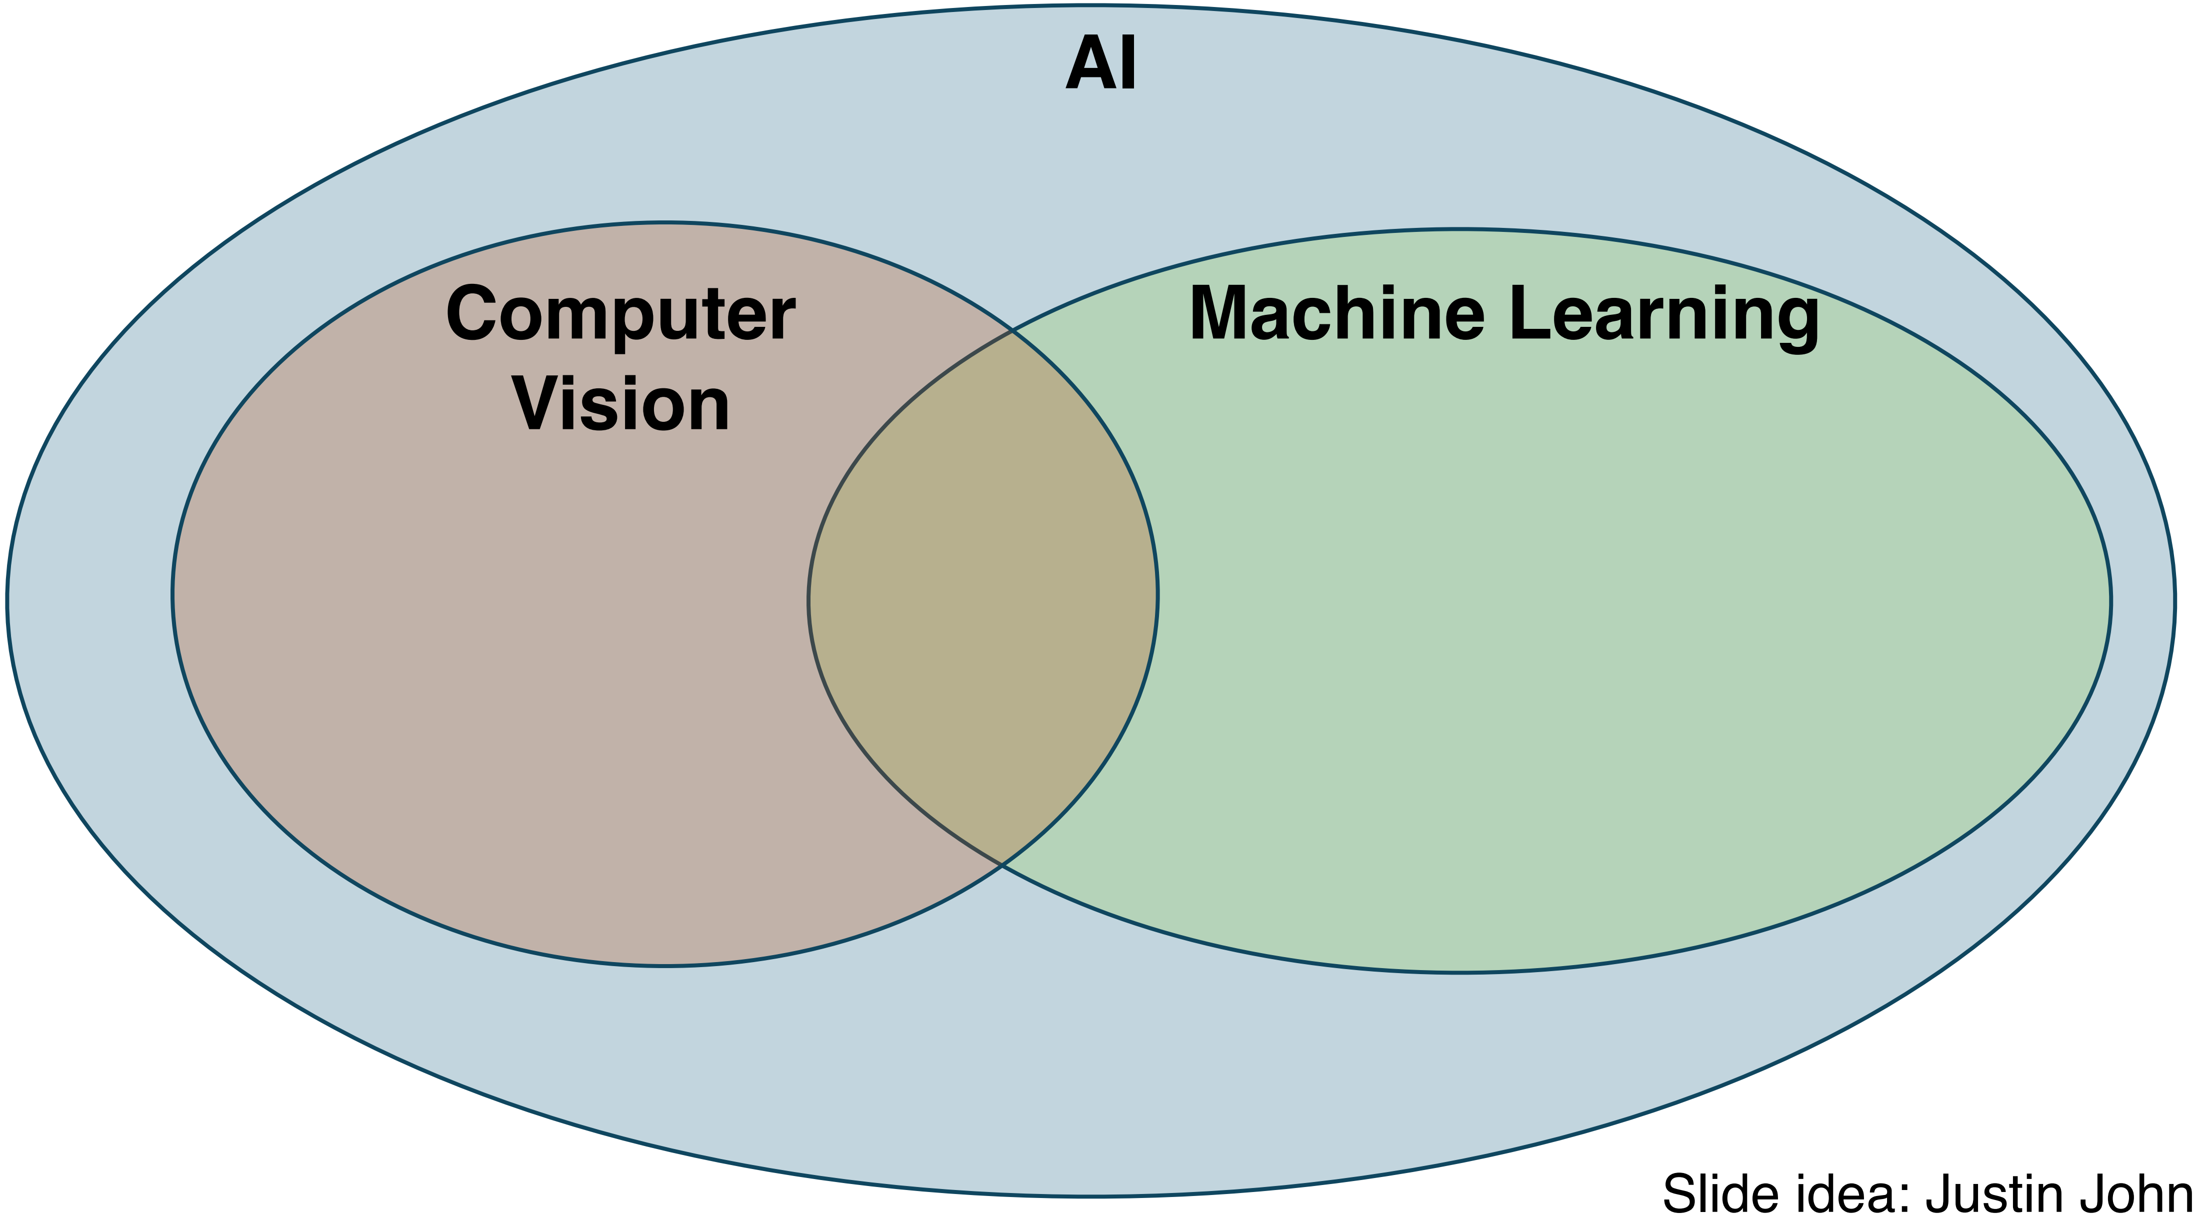

Slide idea: Justin Johnson

Road map

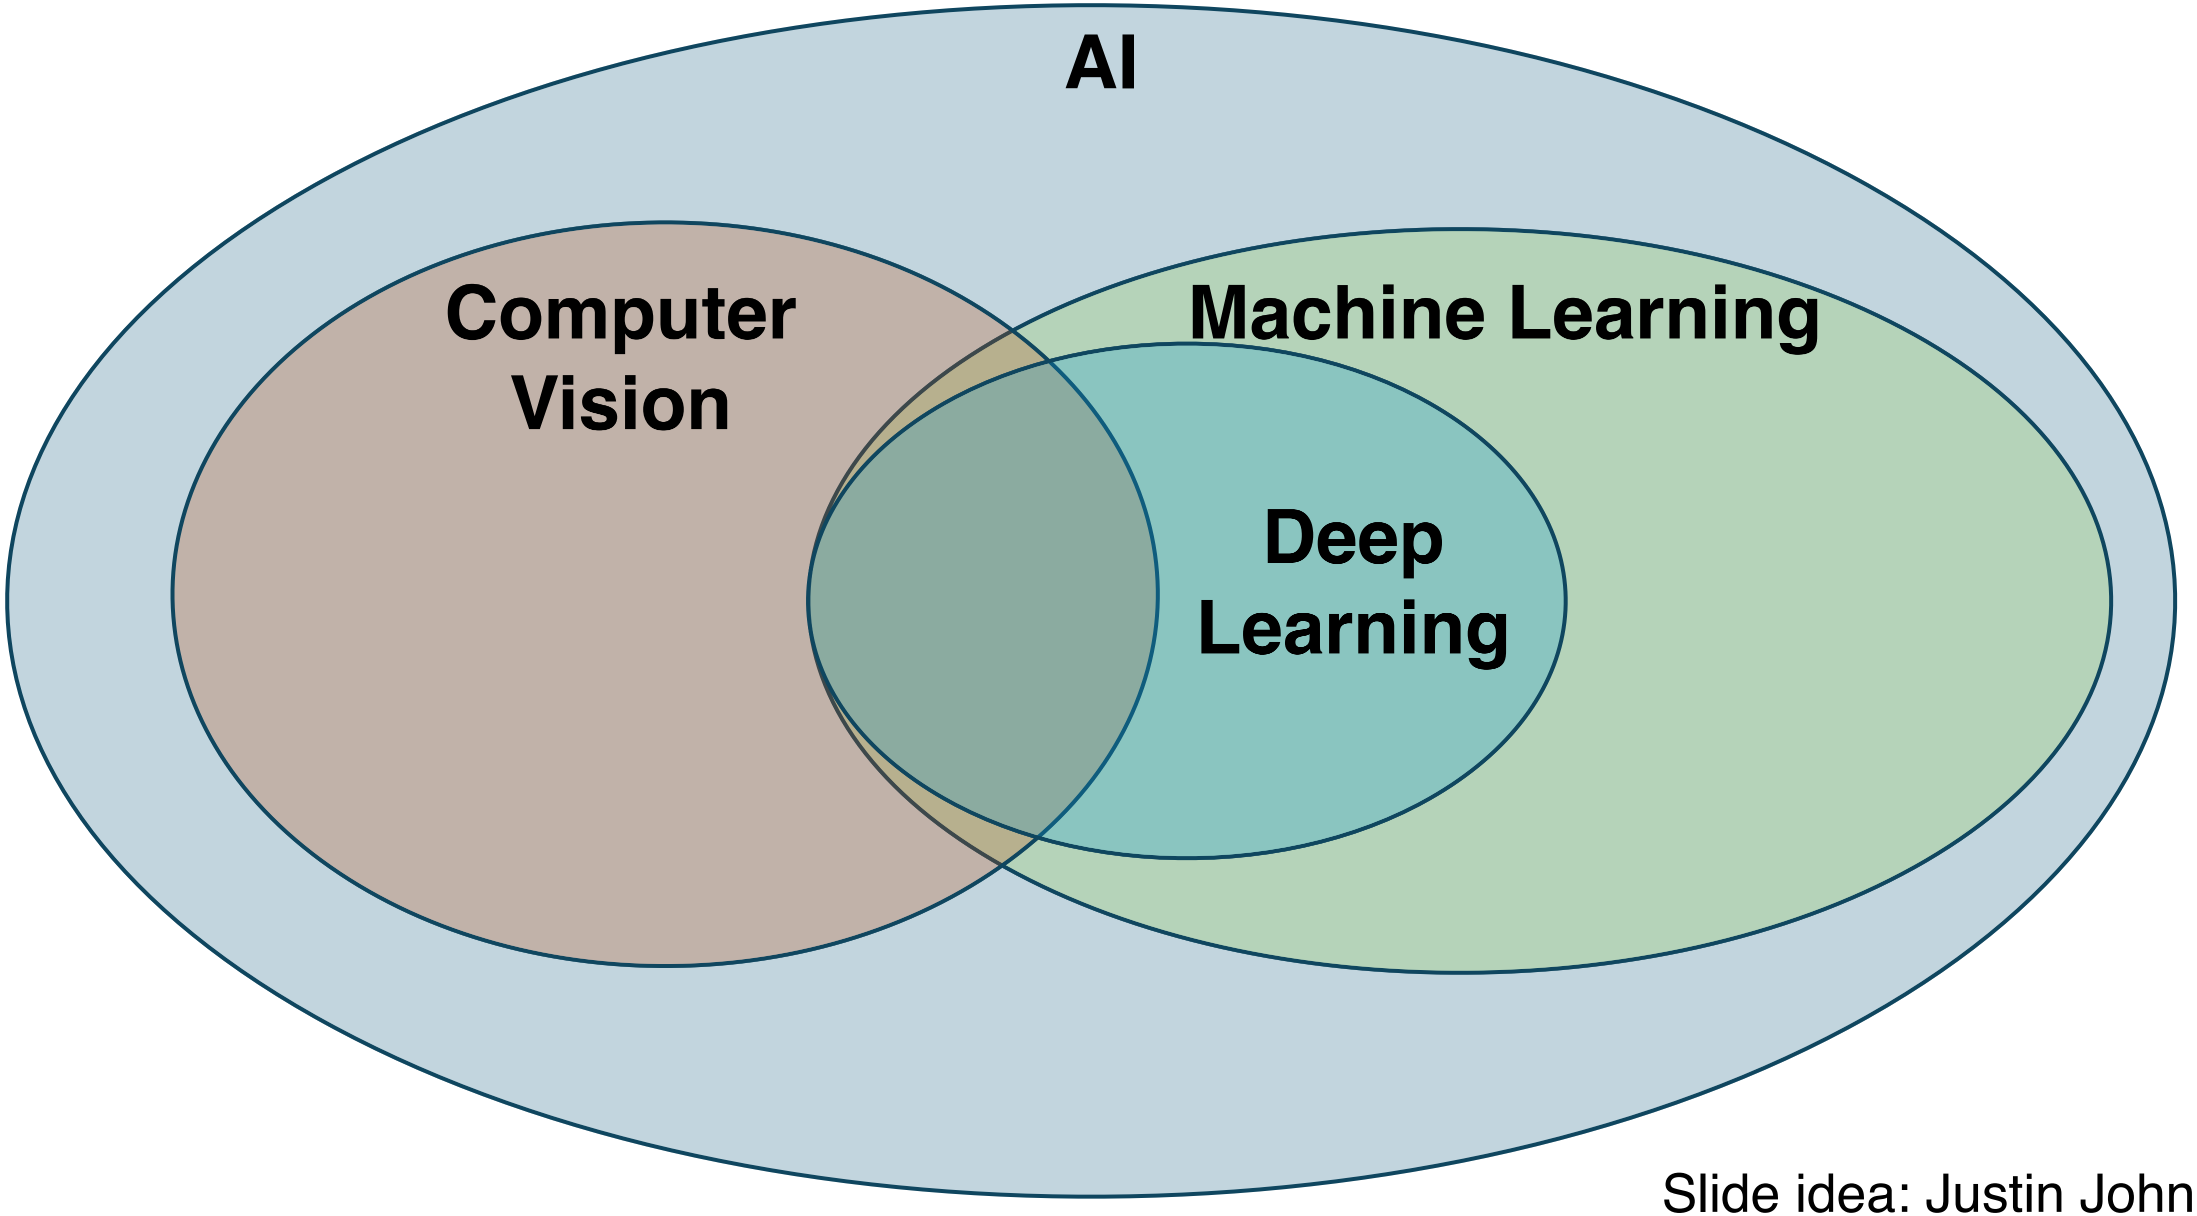

Slide idea: Justin Johnson

Road map

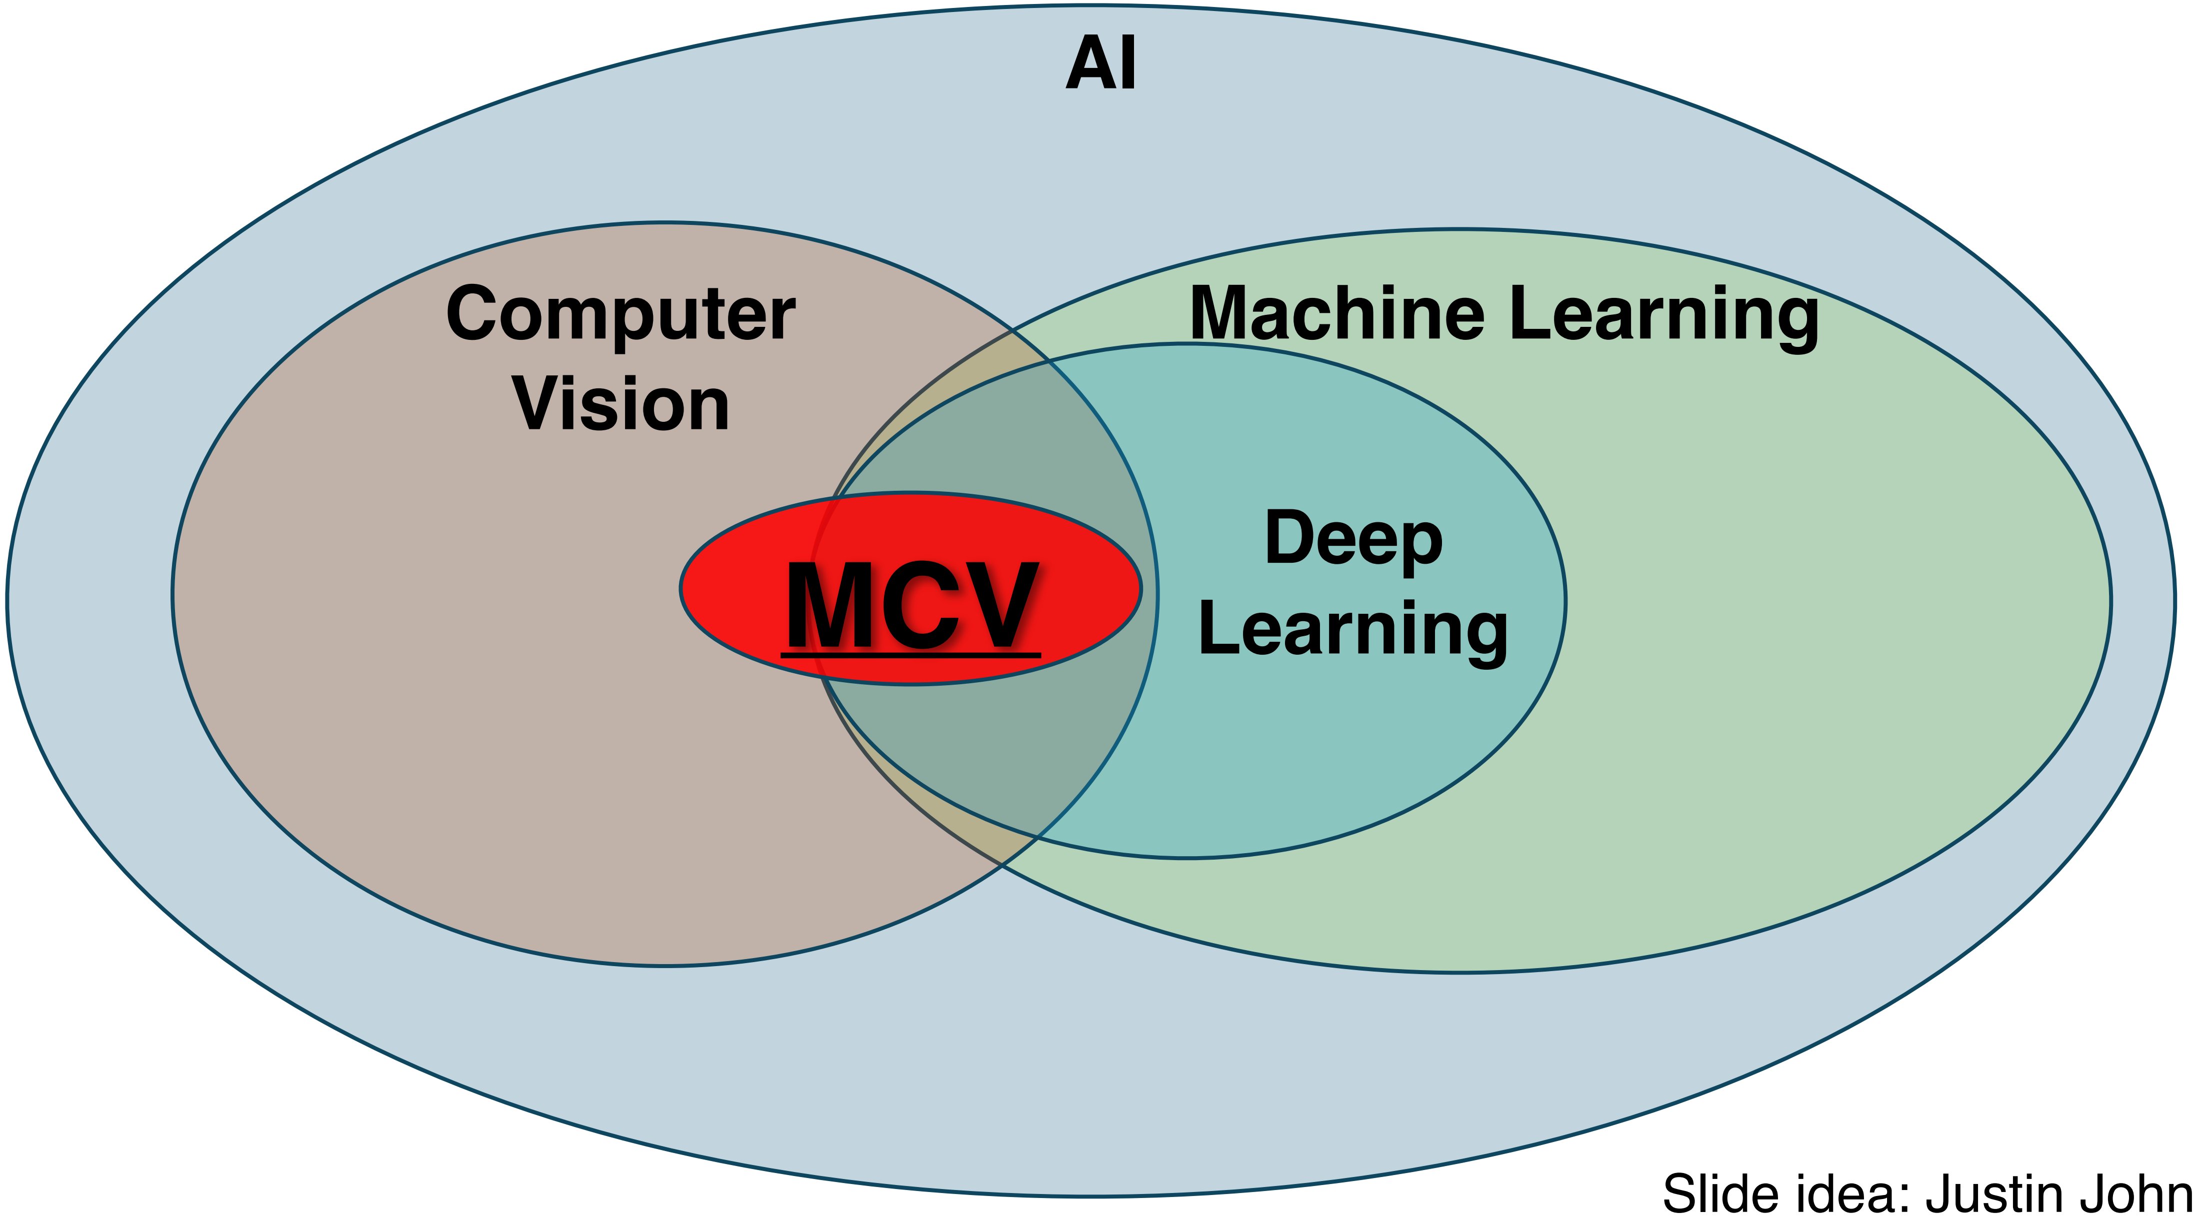

Slide idea: Justin Johnson

Road map

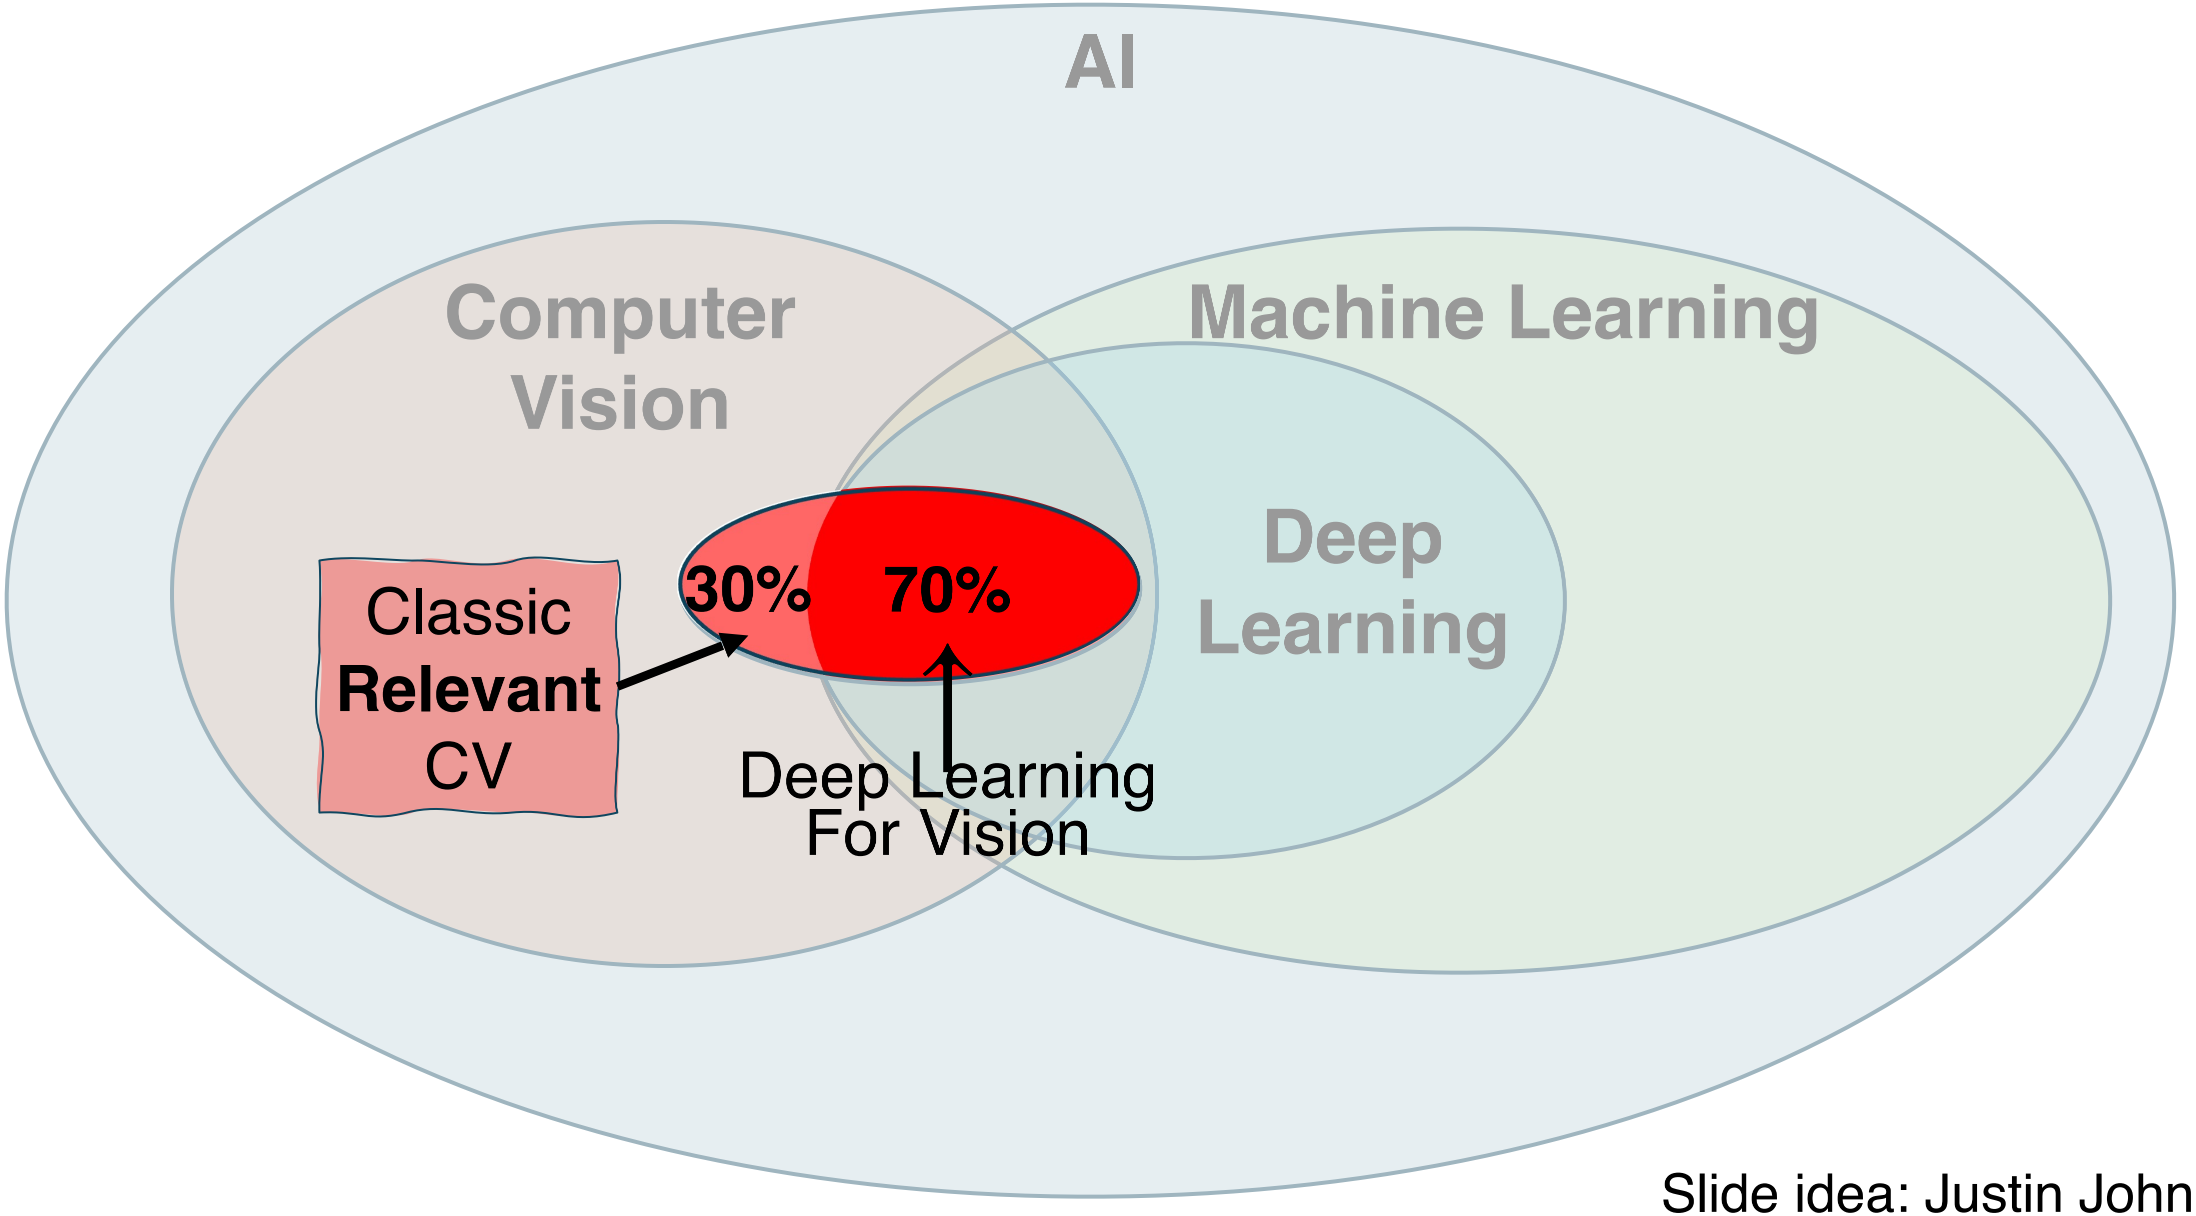

Slide idea: Justin Johnson

Syllabus

Classic Vision

Deep Learning

- Week 1: Welcome & The Basics of Imaging
- Week 2: ML Basics from Representation Learning View
- Week 3: Neural Networks
- Week 4: Frequencies and Sampling
- Week 5: Convolutional Neural Networks
- Week 6: Practical Training Tools & Architectures
- Week 7: Epipolar Geometry
- Week 8: Motion + Localization
- Week 9: Generative Models 1
- Week 10: Generative Models 2
- Week 11: Self-Supervision + Interpretability
- Week 12: Neural Rendering & Implicit Representations

McCulloch Pitts
Non learned
Neuron

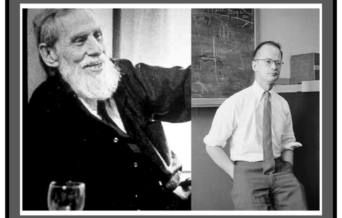

1943

McCulloch Pitts
Non learned
Neuron

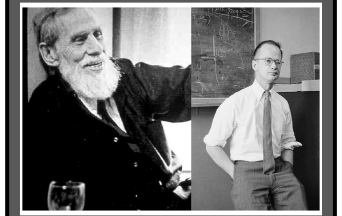

1957

1943

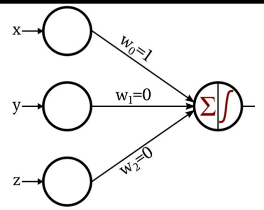

Perceptron
(Rosenblatt)
learned weights

McCulloch Pitts
Non learned
Neuron

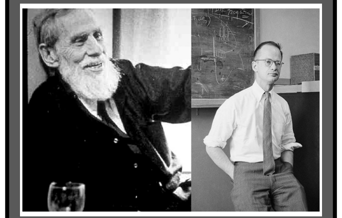

Hubel & Wiesel

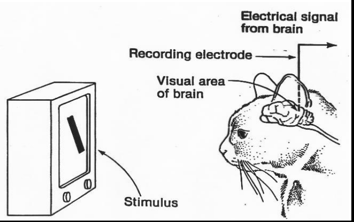

1943

1957

1959

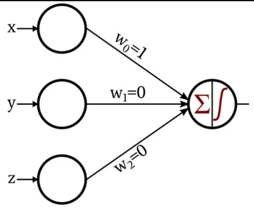

Perceptron
(Rosenblatt)
learned weights

McCulloch Pitts
Non learned
Neuron

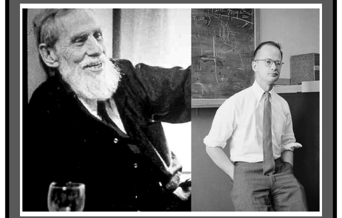

XOR kills
perceptron

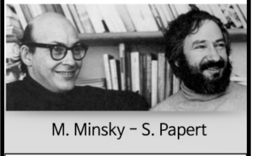

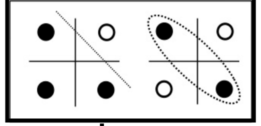

Hubel & Wiesel

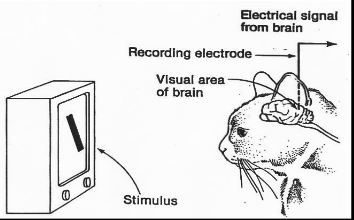

1957

1943

1959

1969

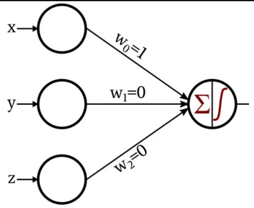

Perceptron
(Rosenblatt)
learned weights

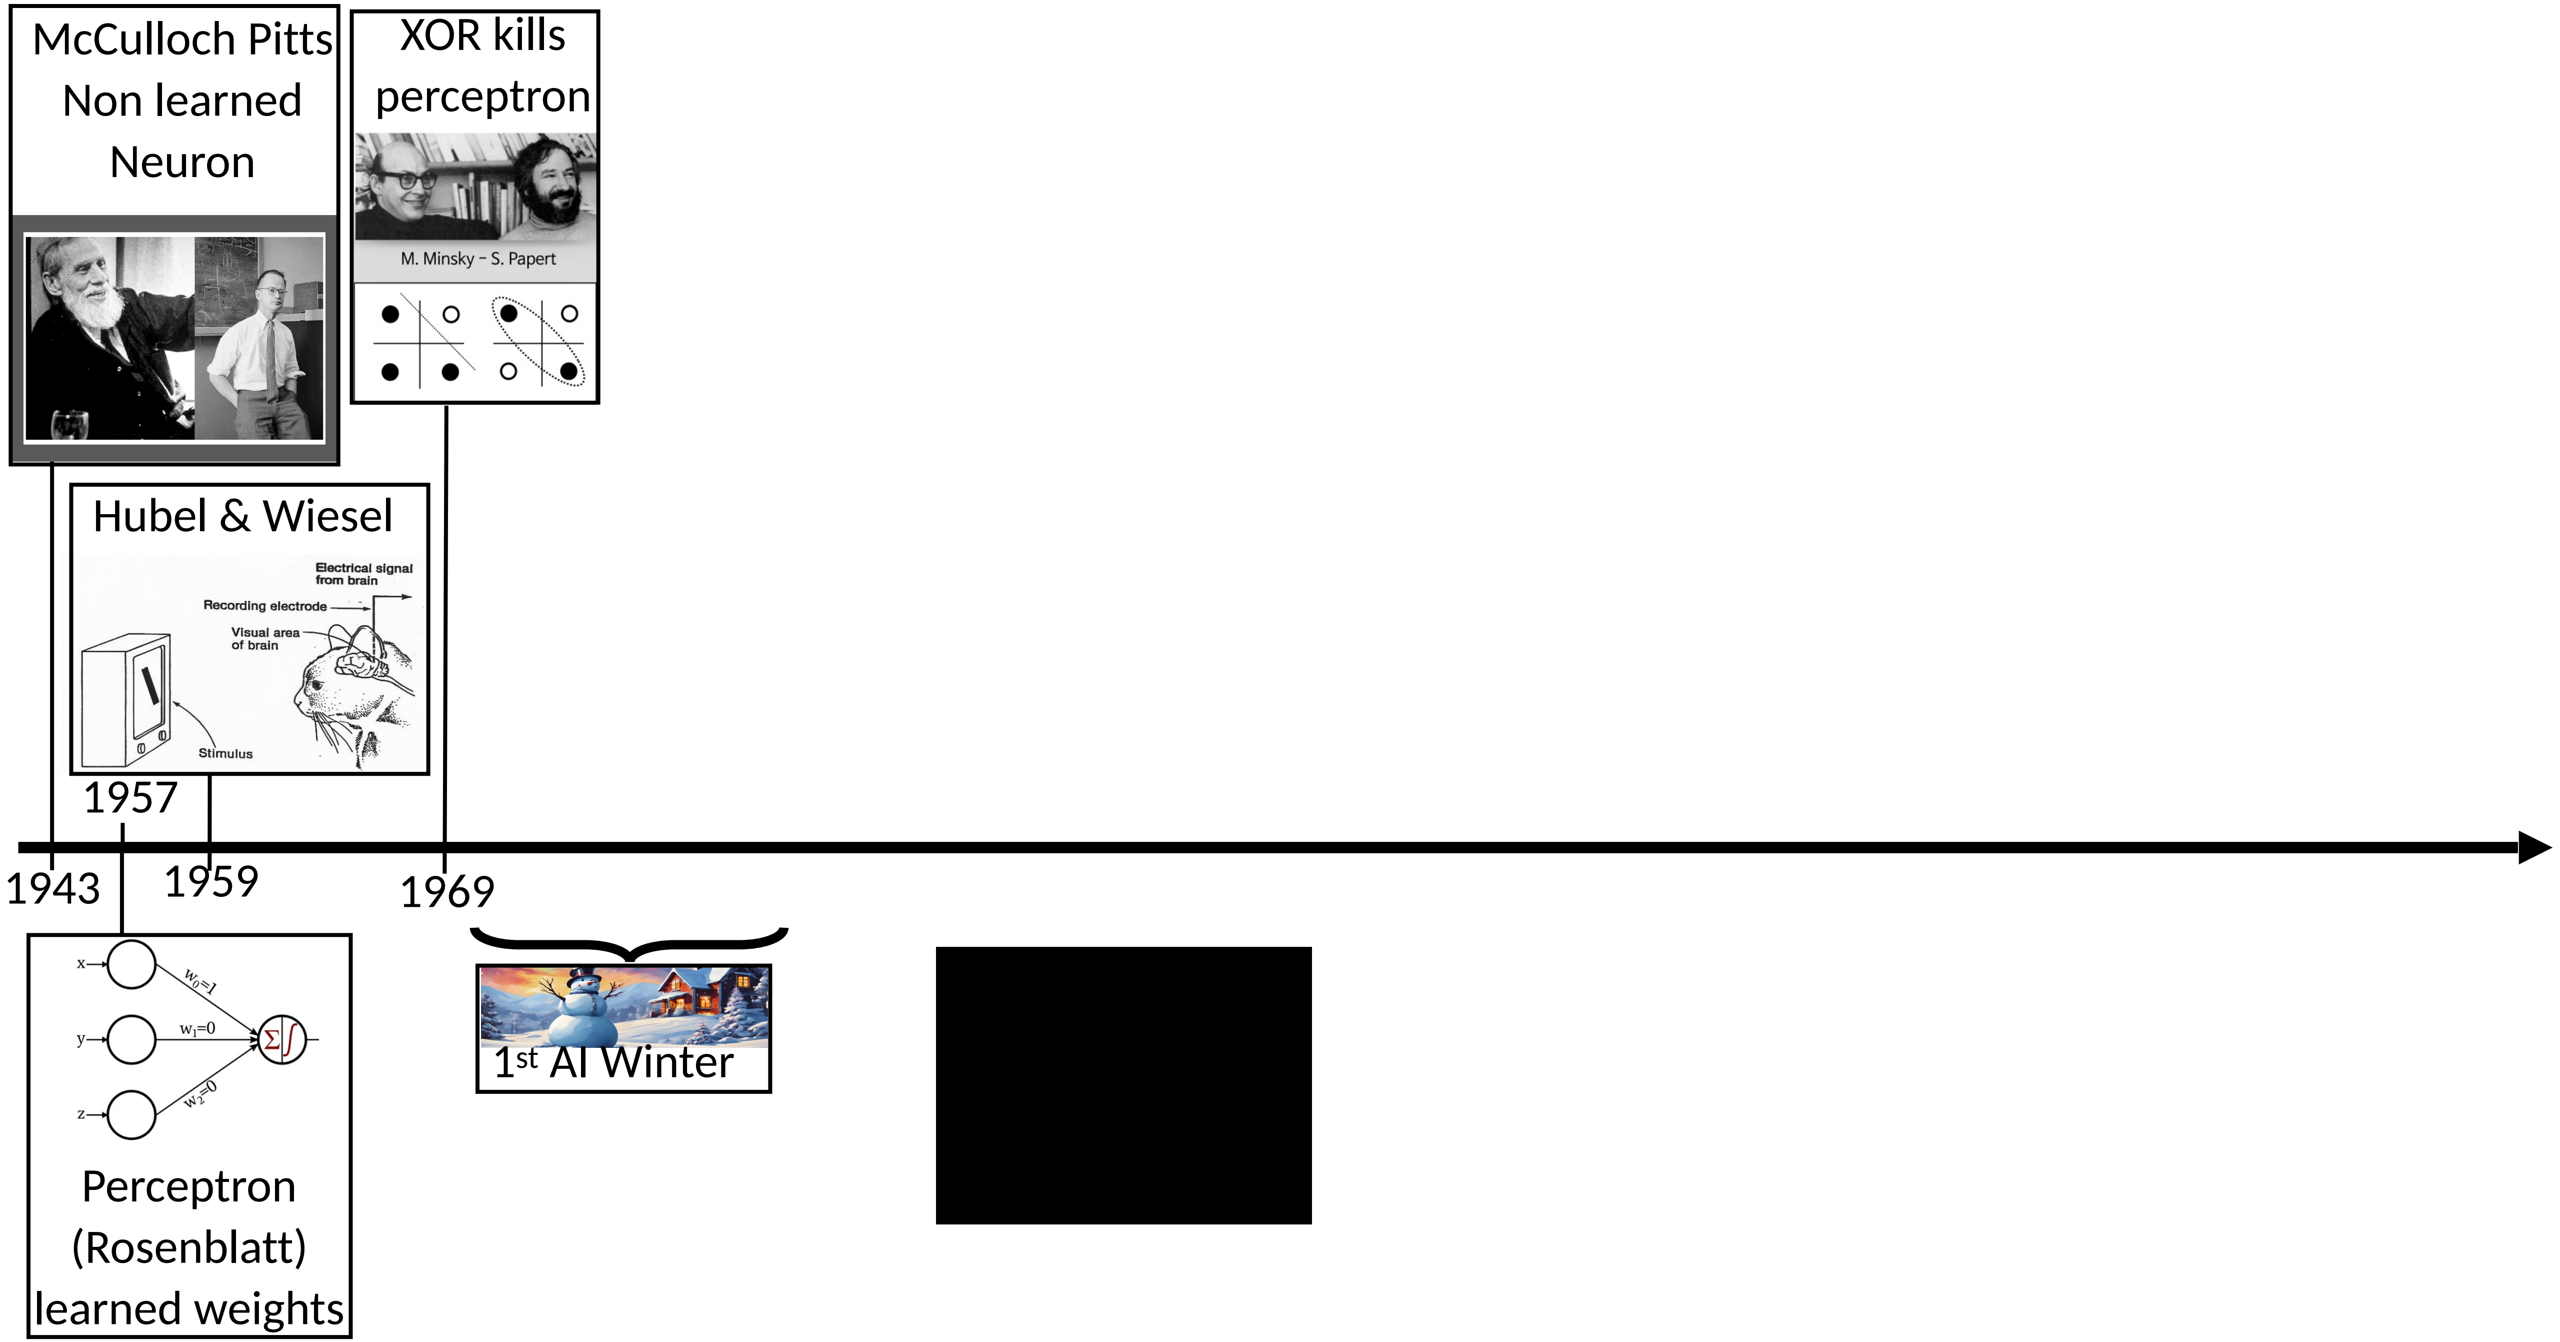
McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Hubel & Wiesel
Electrical signal from brain
Recording electrode
Visual area of brain
Stimulus
1957
1943
1959
1969
x
y
z
w_0=1
w_1=0
w_2=0
Perceptron
(Rosenblatt)
learned weights
1st AI Winter
[redacted]

McCulloch Pitts
Non learned
Neuron

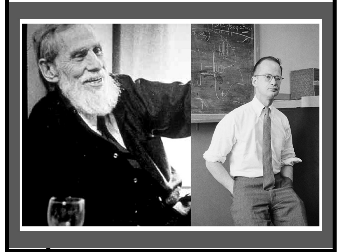

XOR kills
perceptron

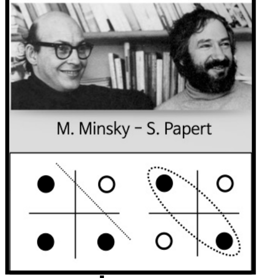

Hubel & Wiesel

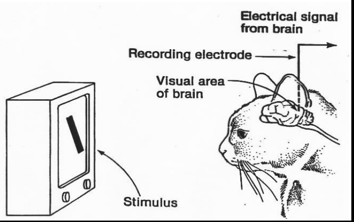

1943

1957

1959

1969

1970

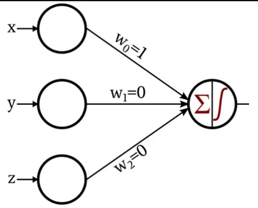

Perceptron
(Rosenblatt)
learned weights

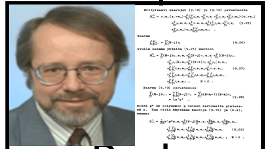

Back
Propagation
(Linnainmaa)

1st AI Winter

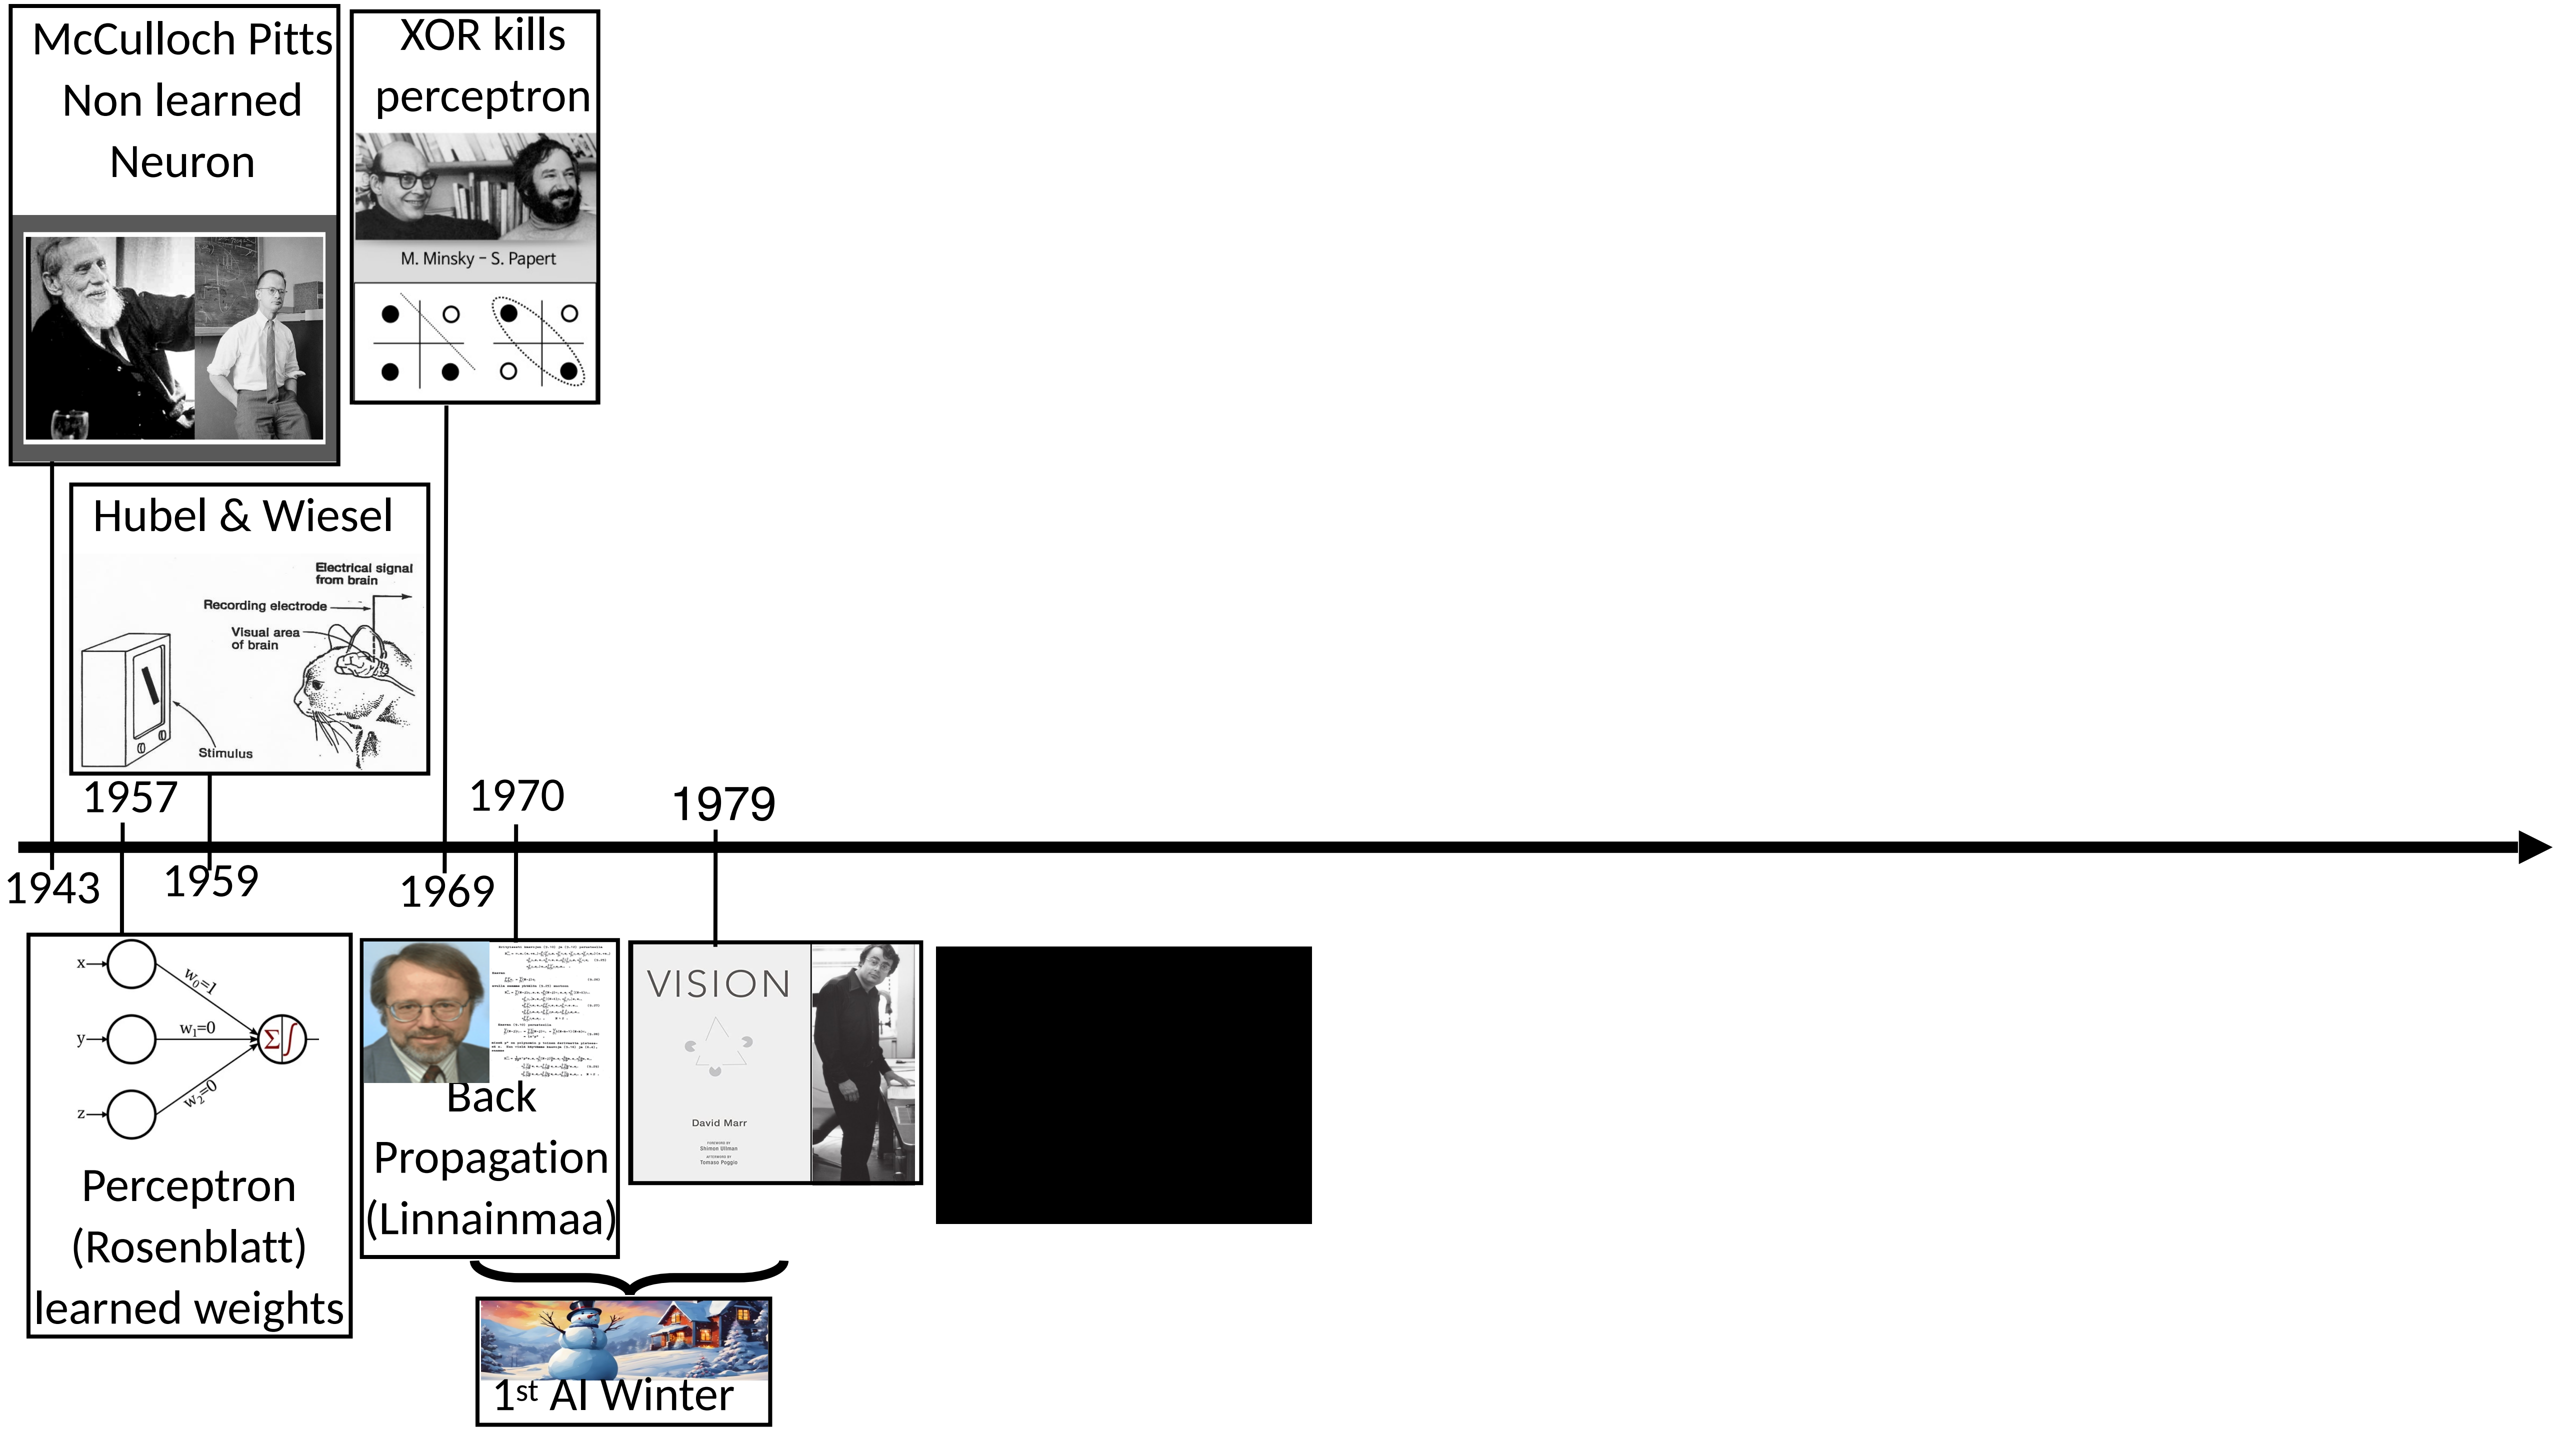

McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Hubel & Wiesel
Electrical signal
from brain
Recording electrode
Visual area
of brain
Stimulus
1957
1970
1979
1943
1959
1969
x
y
z
w0=1
w1=0
w2=0
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
VISION
David Marr
1st AI Winter

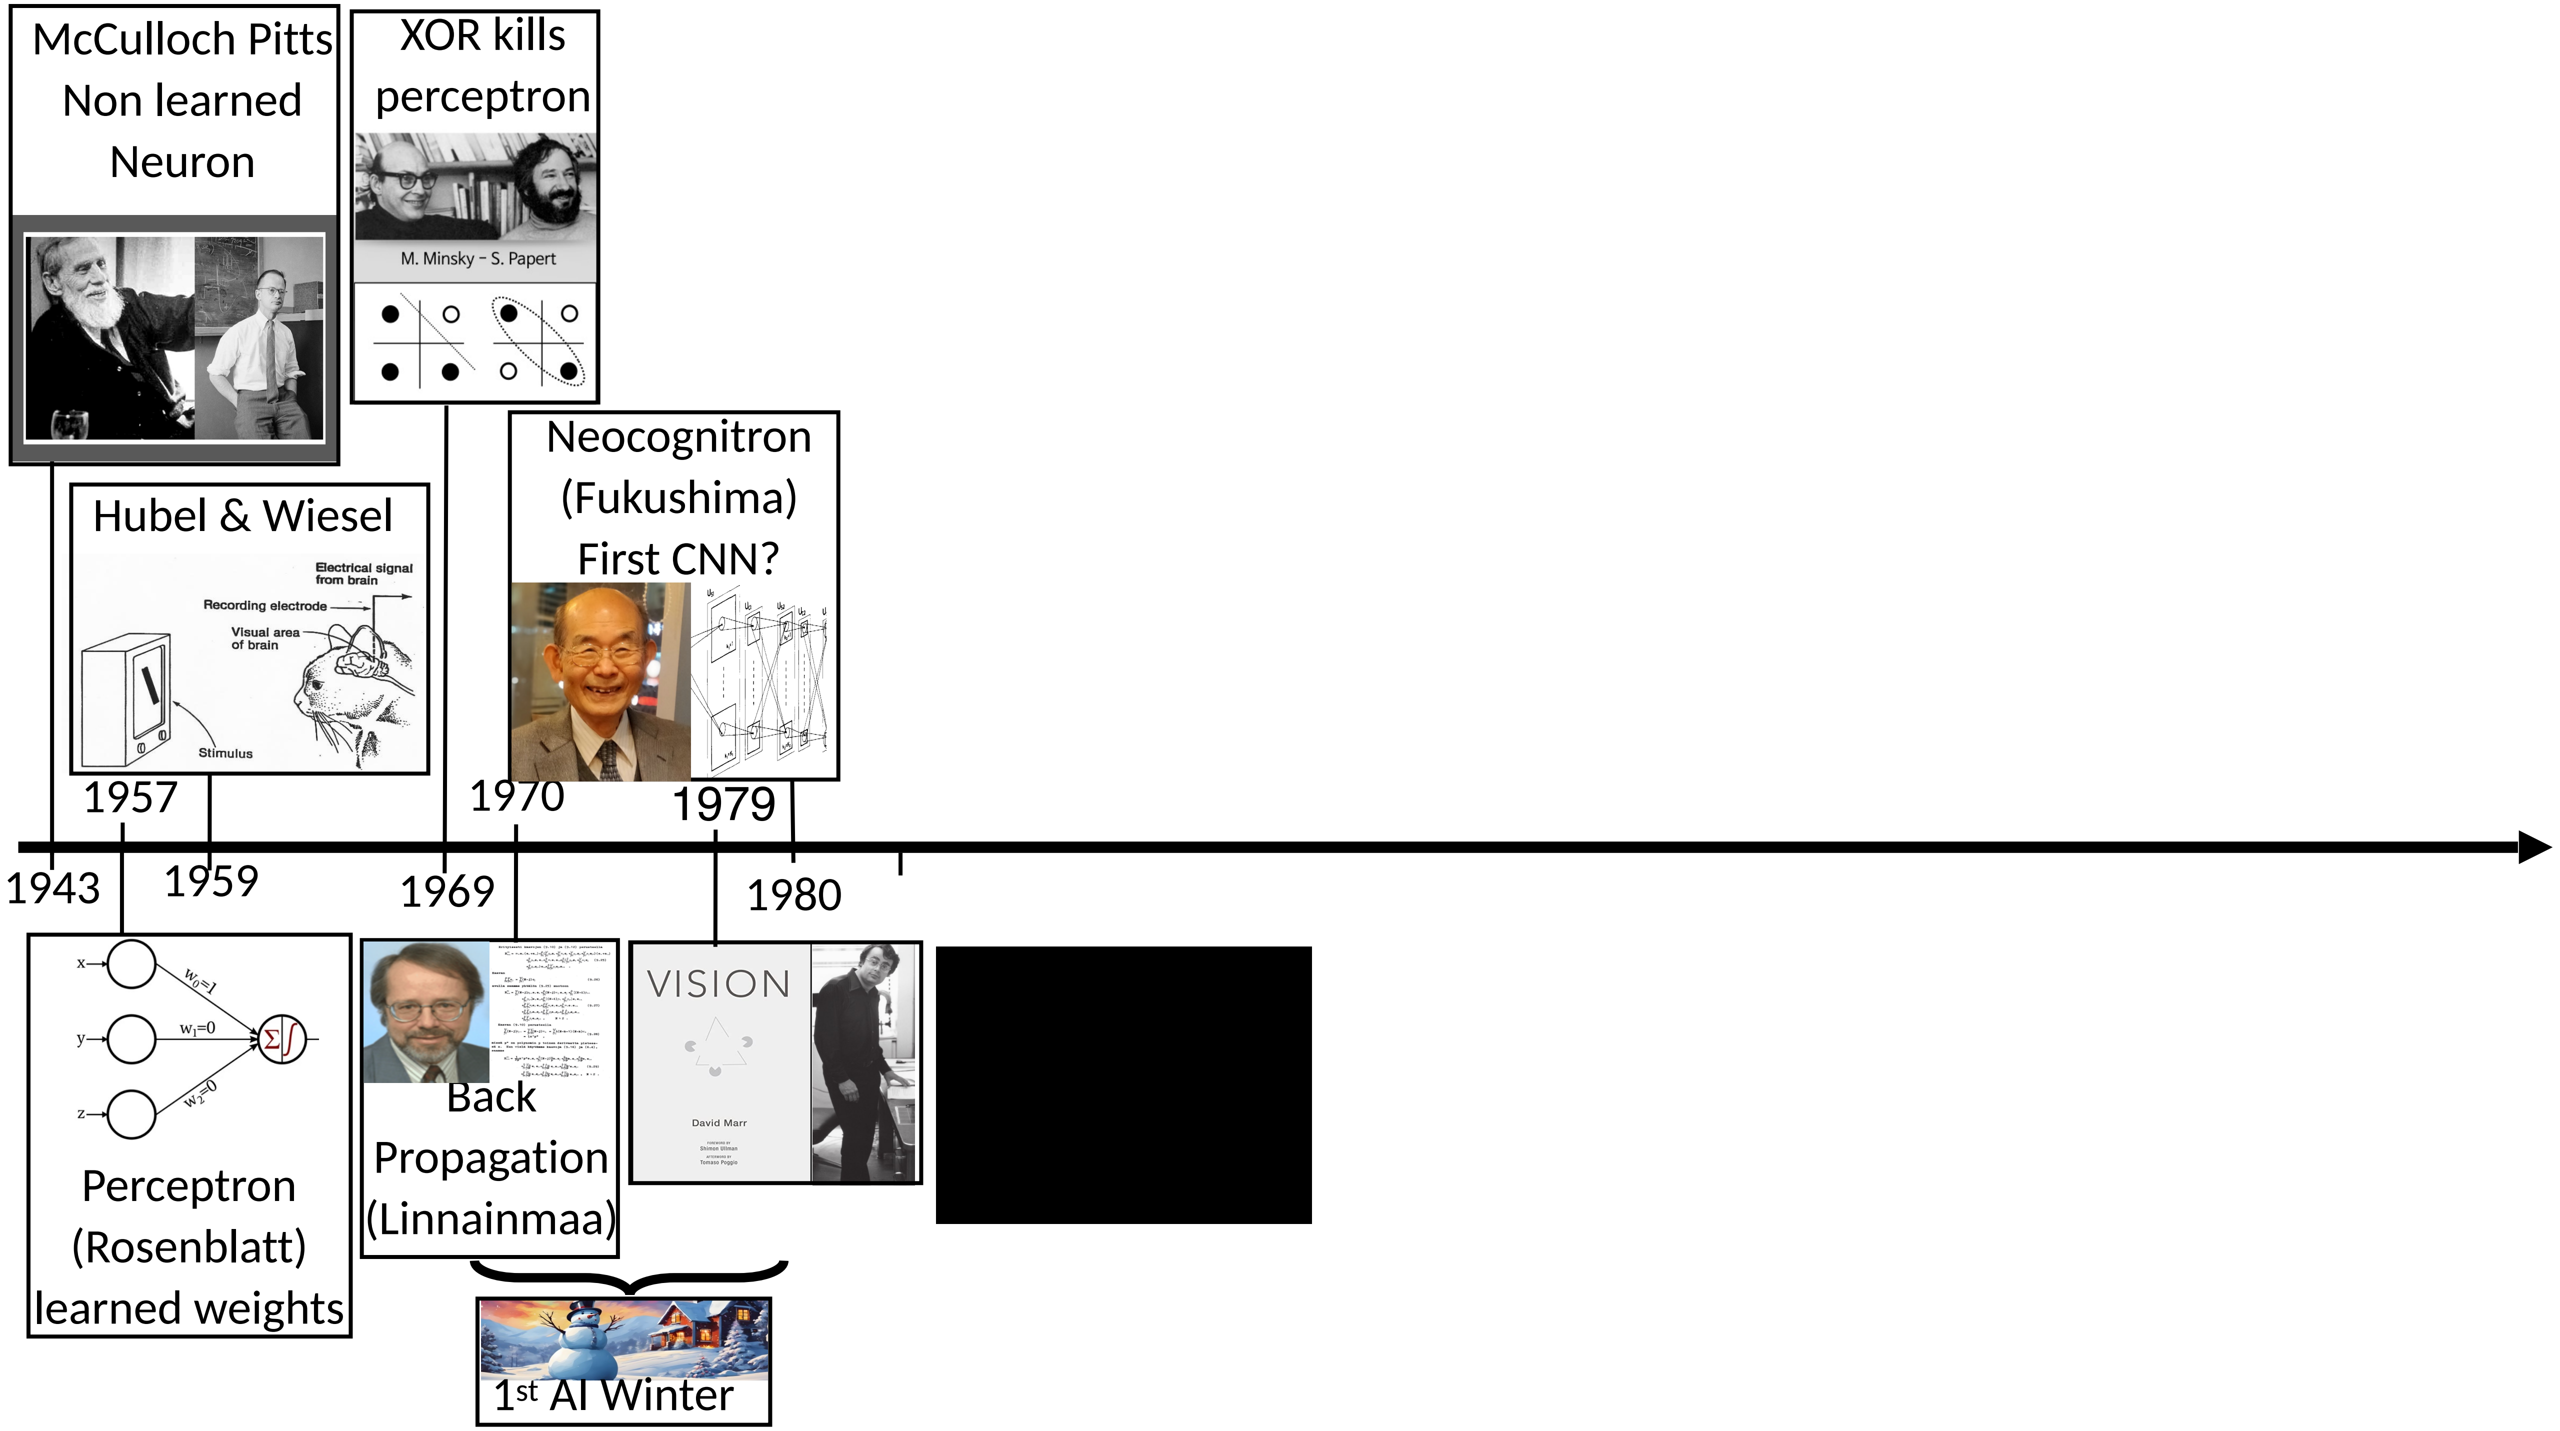
McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Neocognitron
(Fukushima)
First CNN?
Hubel & Wiesel
Electrical signal
from brain
Recording electrode
Visual area
of brain
Stimulus
1957
1970
1979
1943
1959
1969
1980
x
y
z
w0=1
w1=0
w2=0
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
VISION
David Marr
[redacted]
1st AI Winter

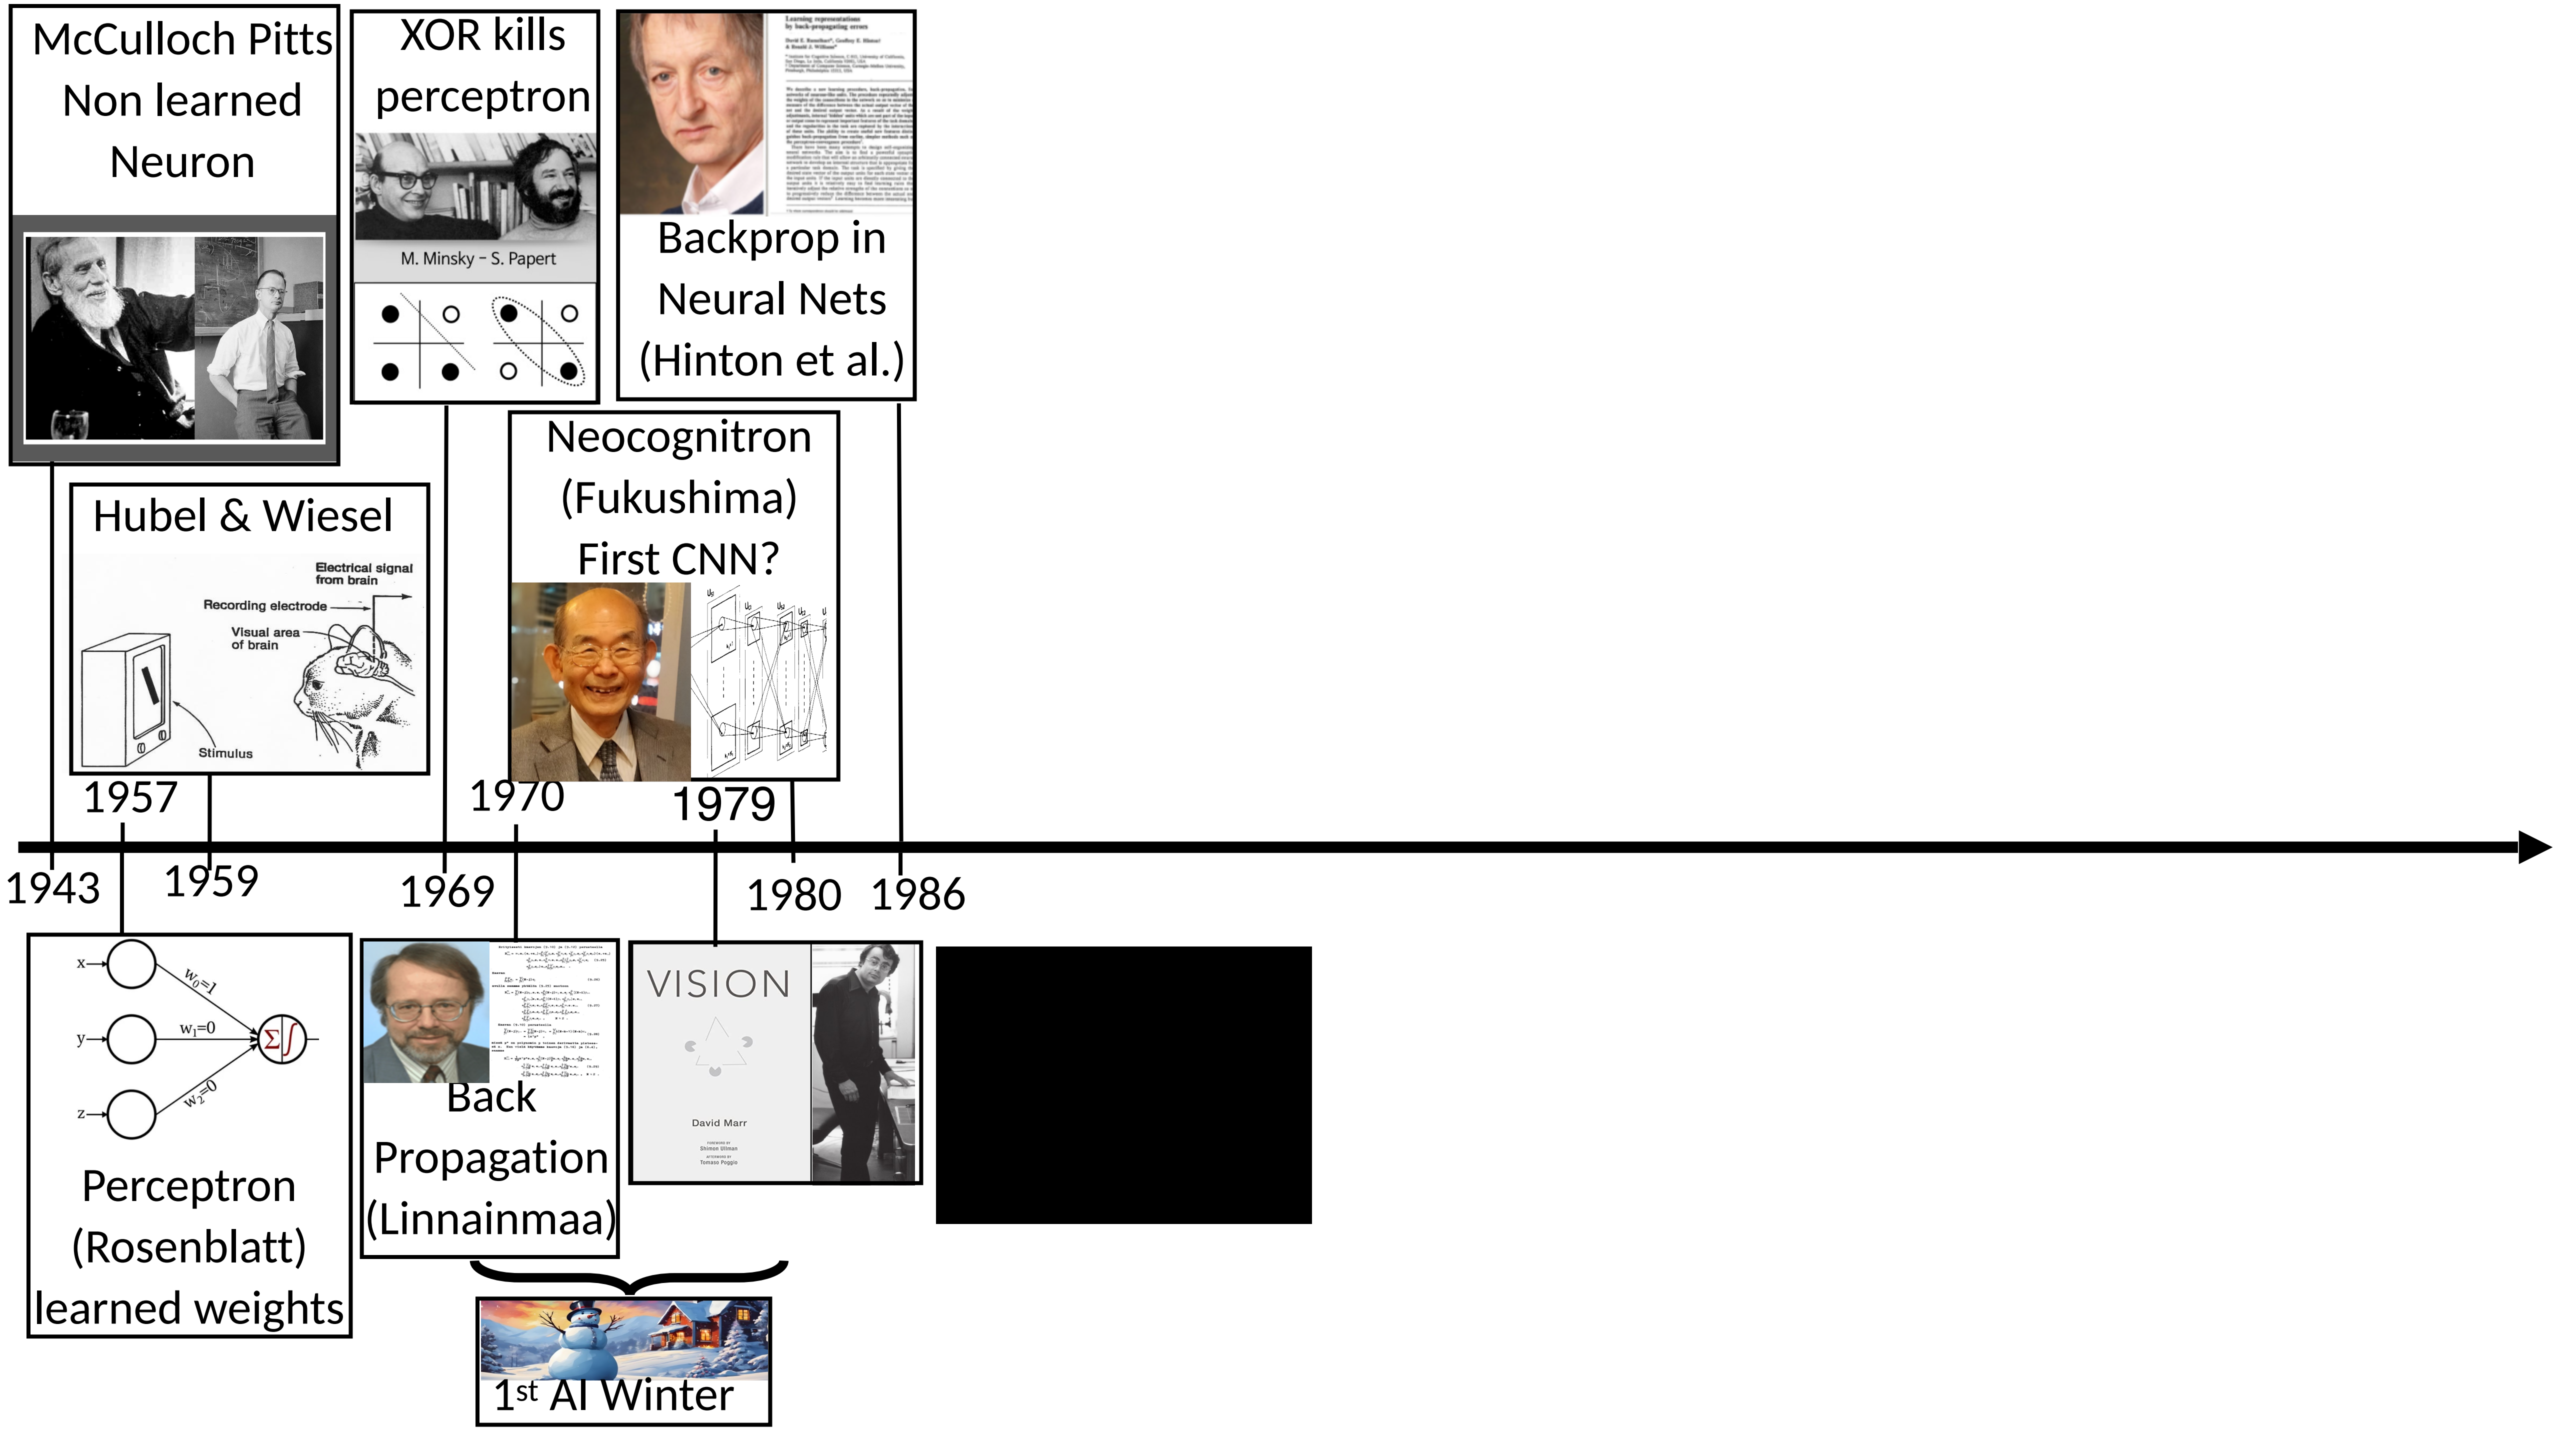
McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Backprop in
Neural Nets
(Hinton et al.)
Learning representations
by back-propagating errors
Hubel & Wiesel
Electrical signal
from brain
Recording electrode
Visual area
of brain
Stimulus
Neocognitron
(Fukushima)
First CNN?
1957
1970
1979
1943
1959
1969
1980
1986
Perceptron
(Rosenblatt)
learned weights
$w_0=1$
$w_1=0$
$w_2=0$
Back
Propagation
(Linnainmaa)
VISION
David Marr
1st AI Winter

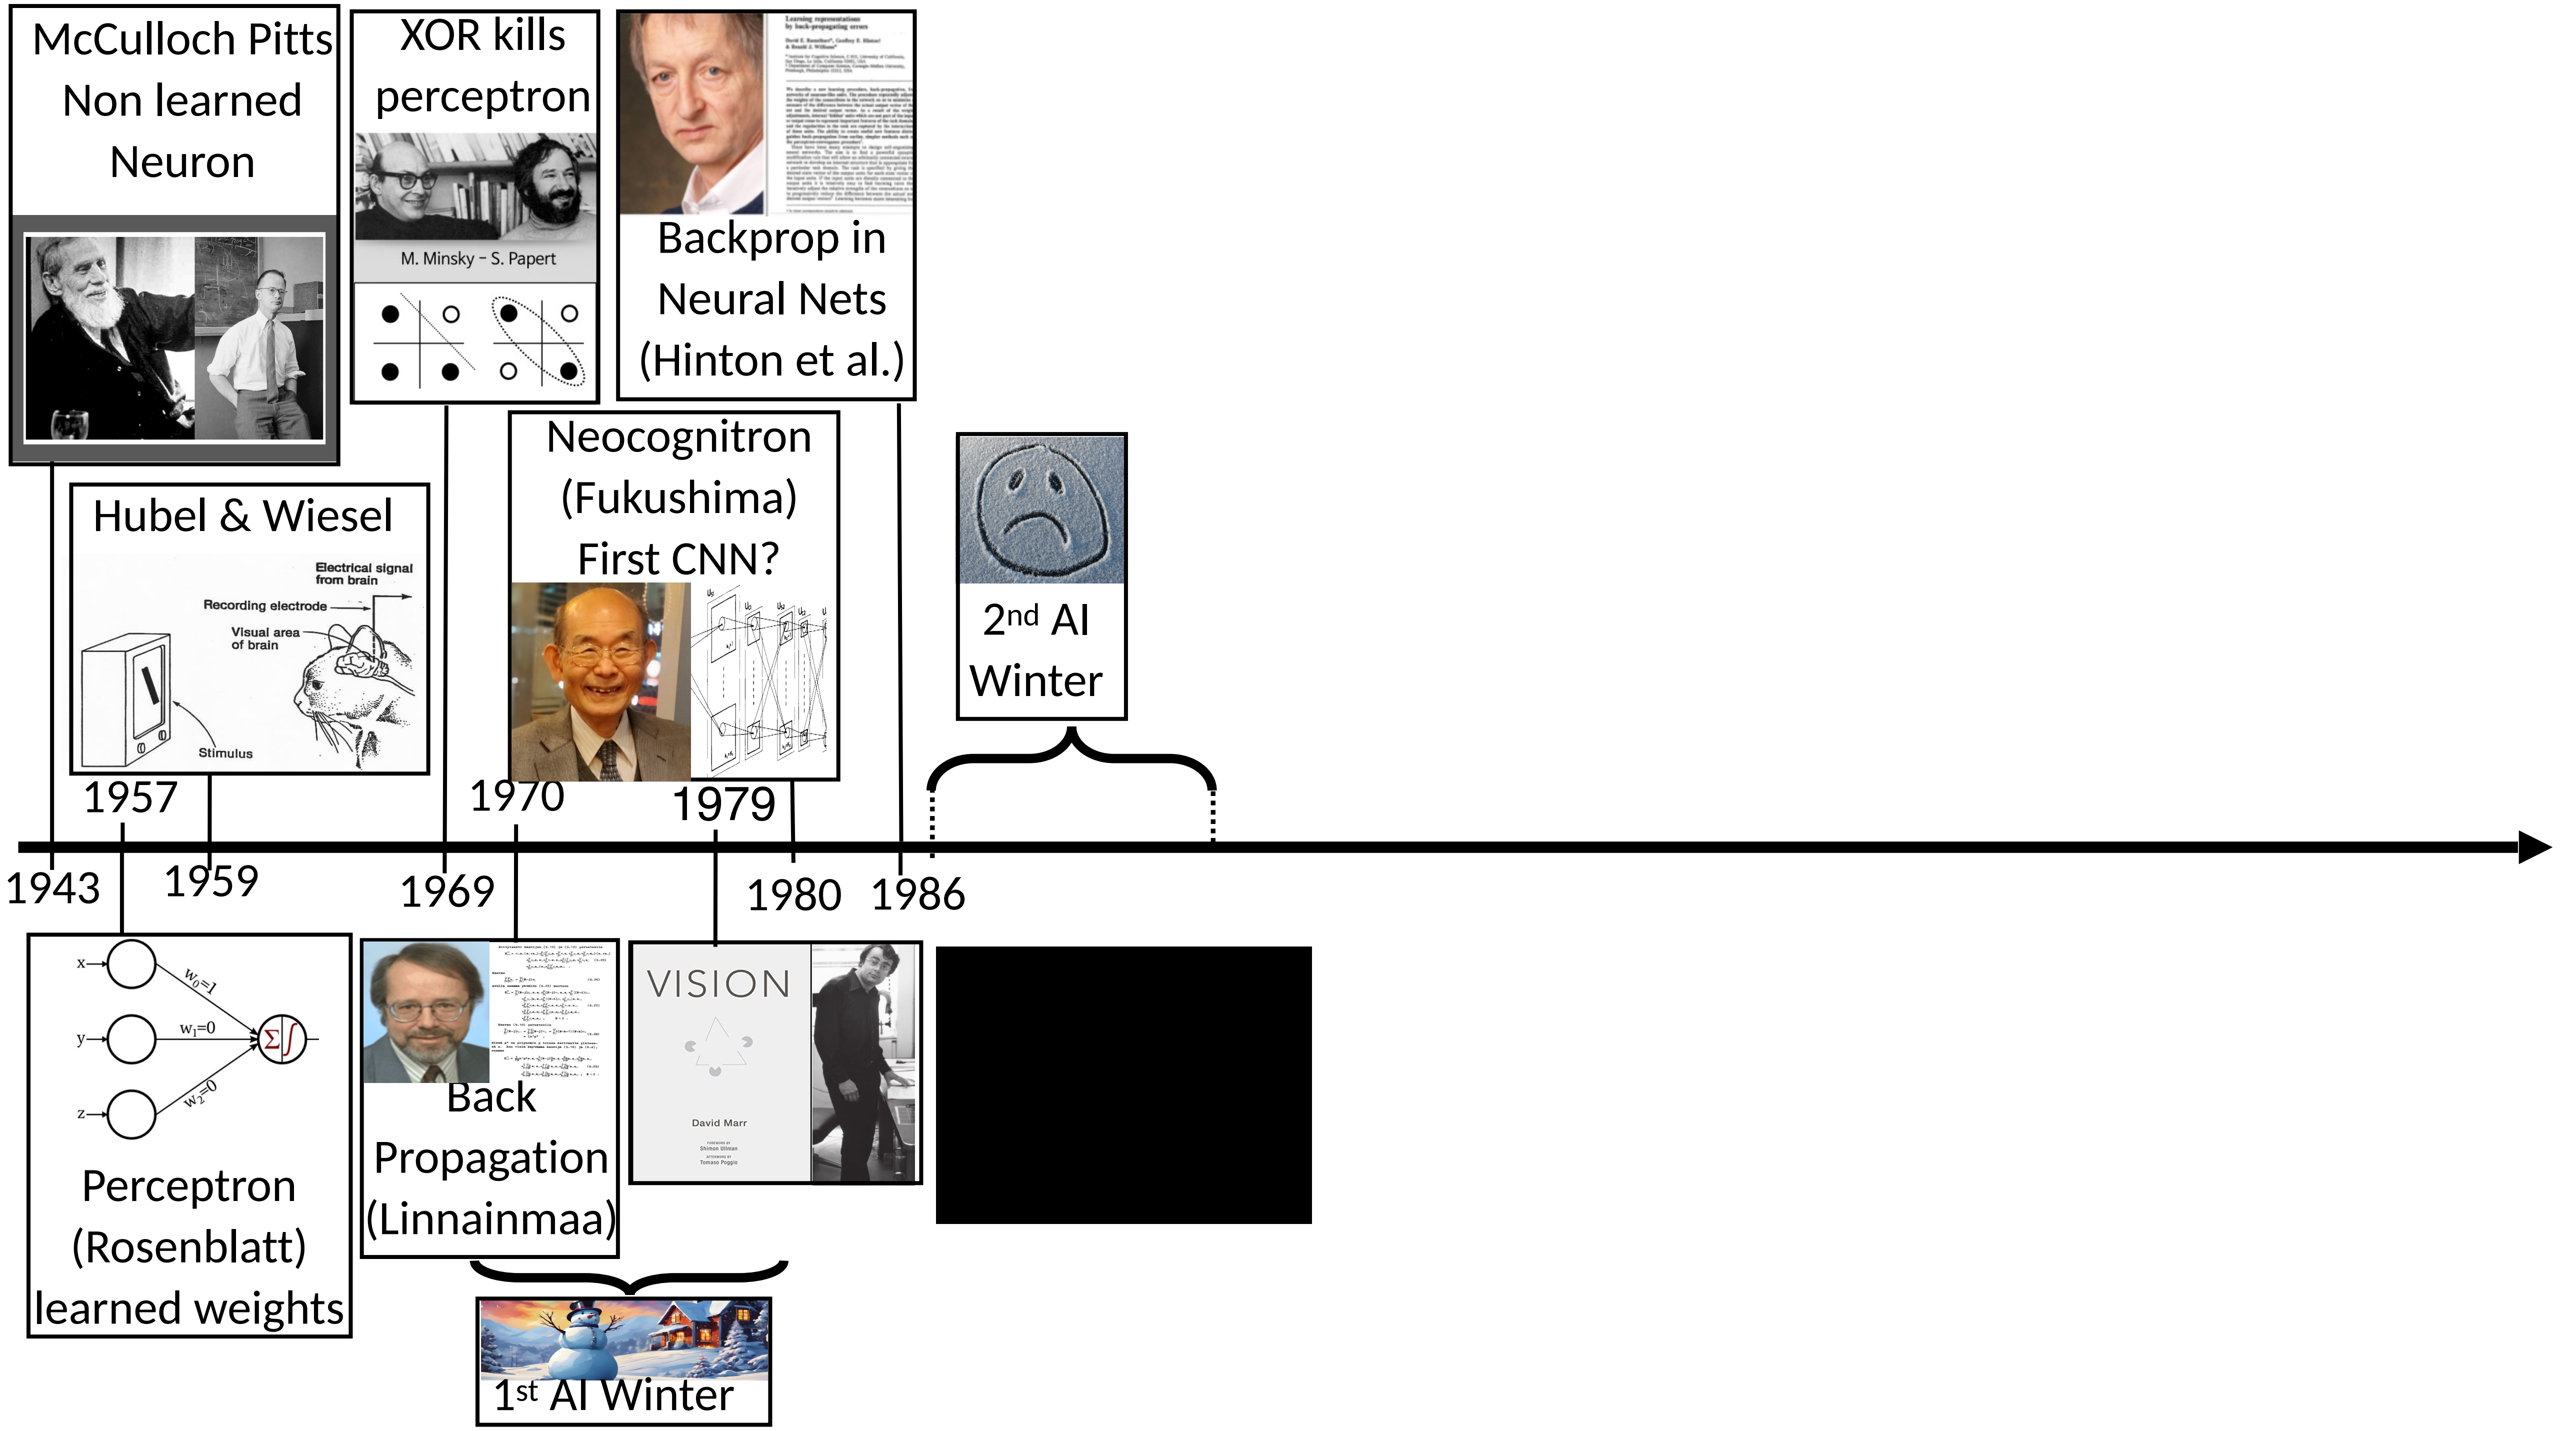

McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Backprop in
Neural Nets
(Hinton et al.)
Hubel & Wiesel
Electrical signal
from brain
Recording electrode
Visual area
of brain
Stimulus
Neocognitron
(Fukushima)
First CNN?
2nd AI
Winter
1957
1970
1979
1943
1959
1969
1980
1986
x
y
z
w0=1
w1=0
w2=0
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
VISION
David Marr
1st AI Winter

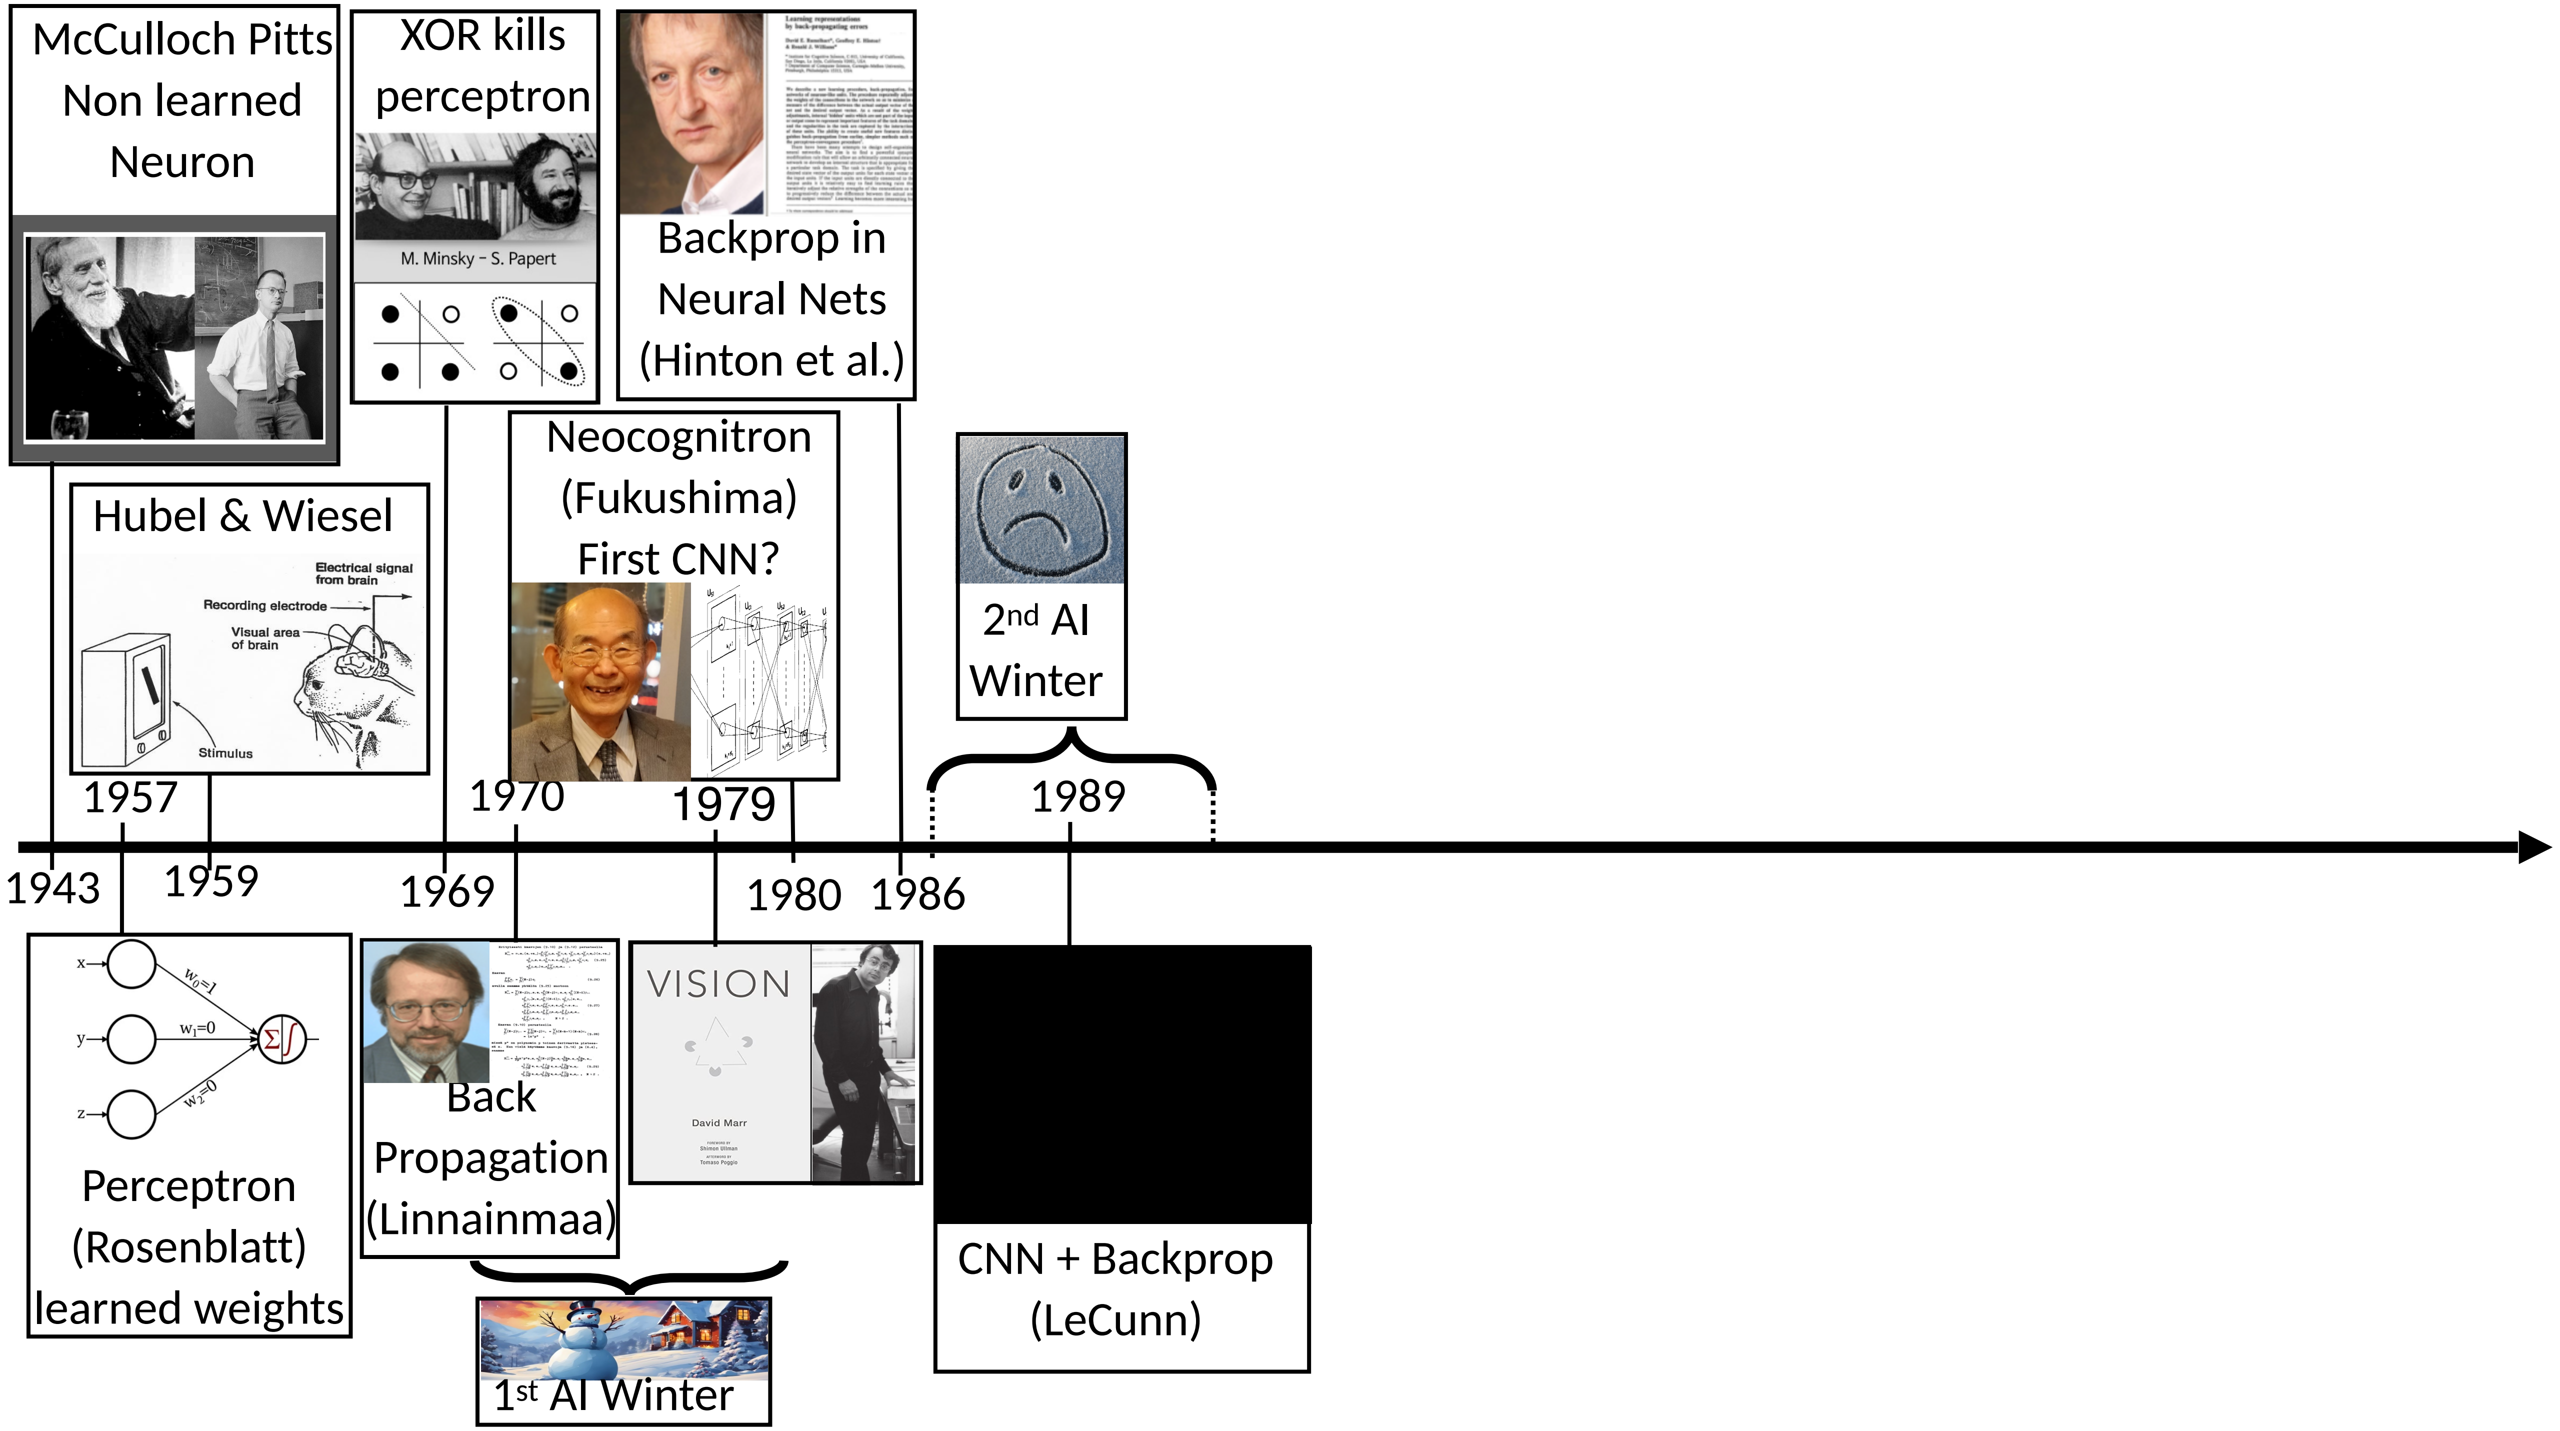
McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Backprop in
Neural Nets
(Hinton et al.)
Hubel & Wiesel
Electrical signal from brain
Recording electrode
Visual area of brain
Stimulus
Neocognitron
(Fukushima)
First CNN?
2nd AI
Winter
1957
1970
1979
1989
1943
1959
1969
1980
1986
x
y
z
w0=1
w1=0
w2=0
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
VISION
David Marr
CNN + Backprop
(LeCunn)
1st AI Winter

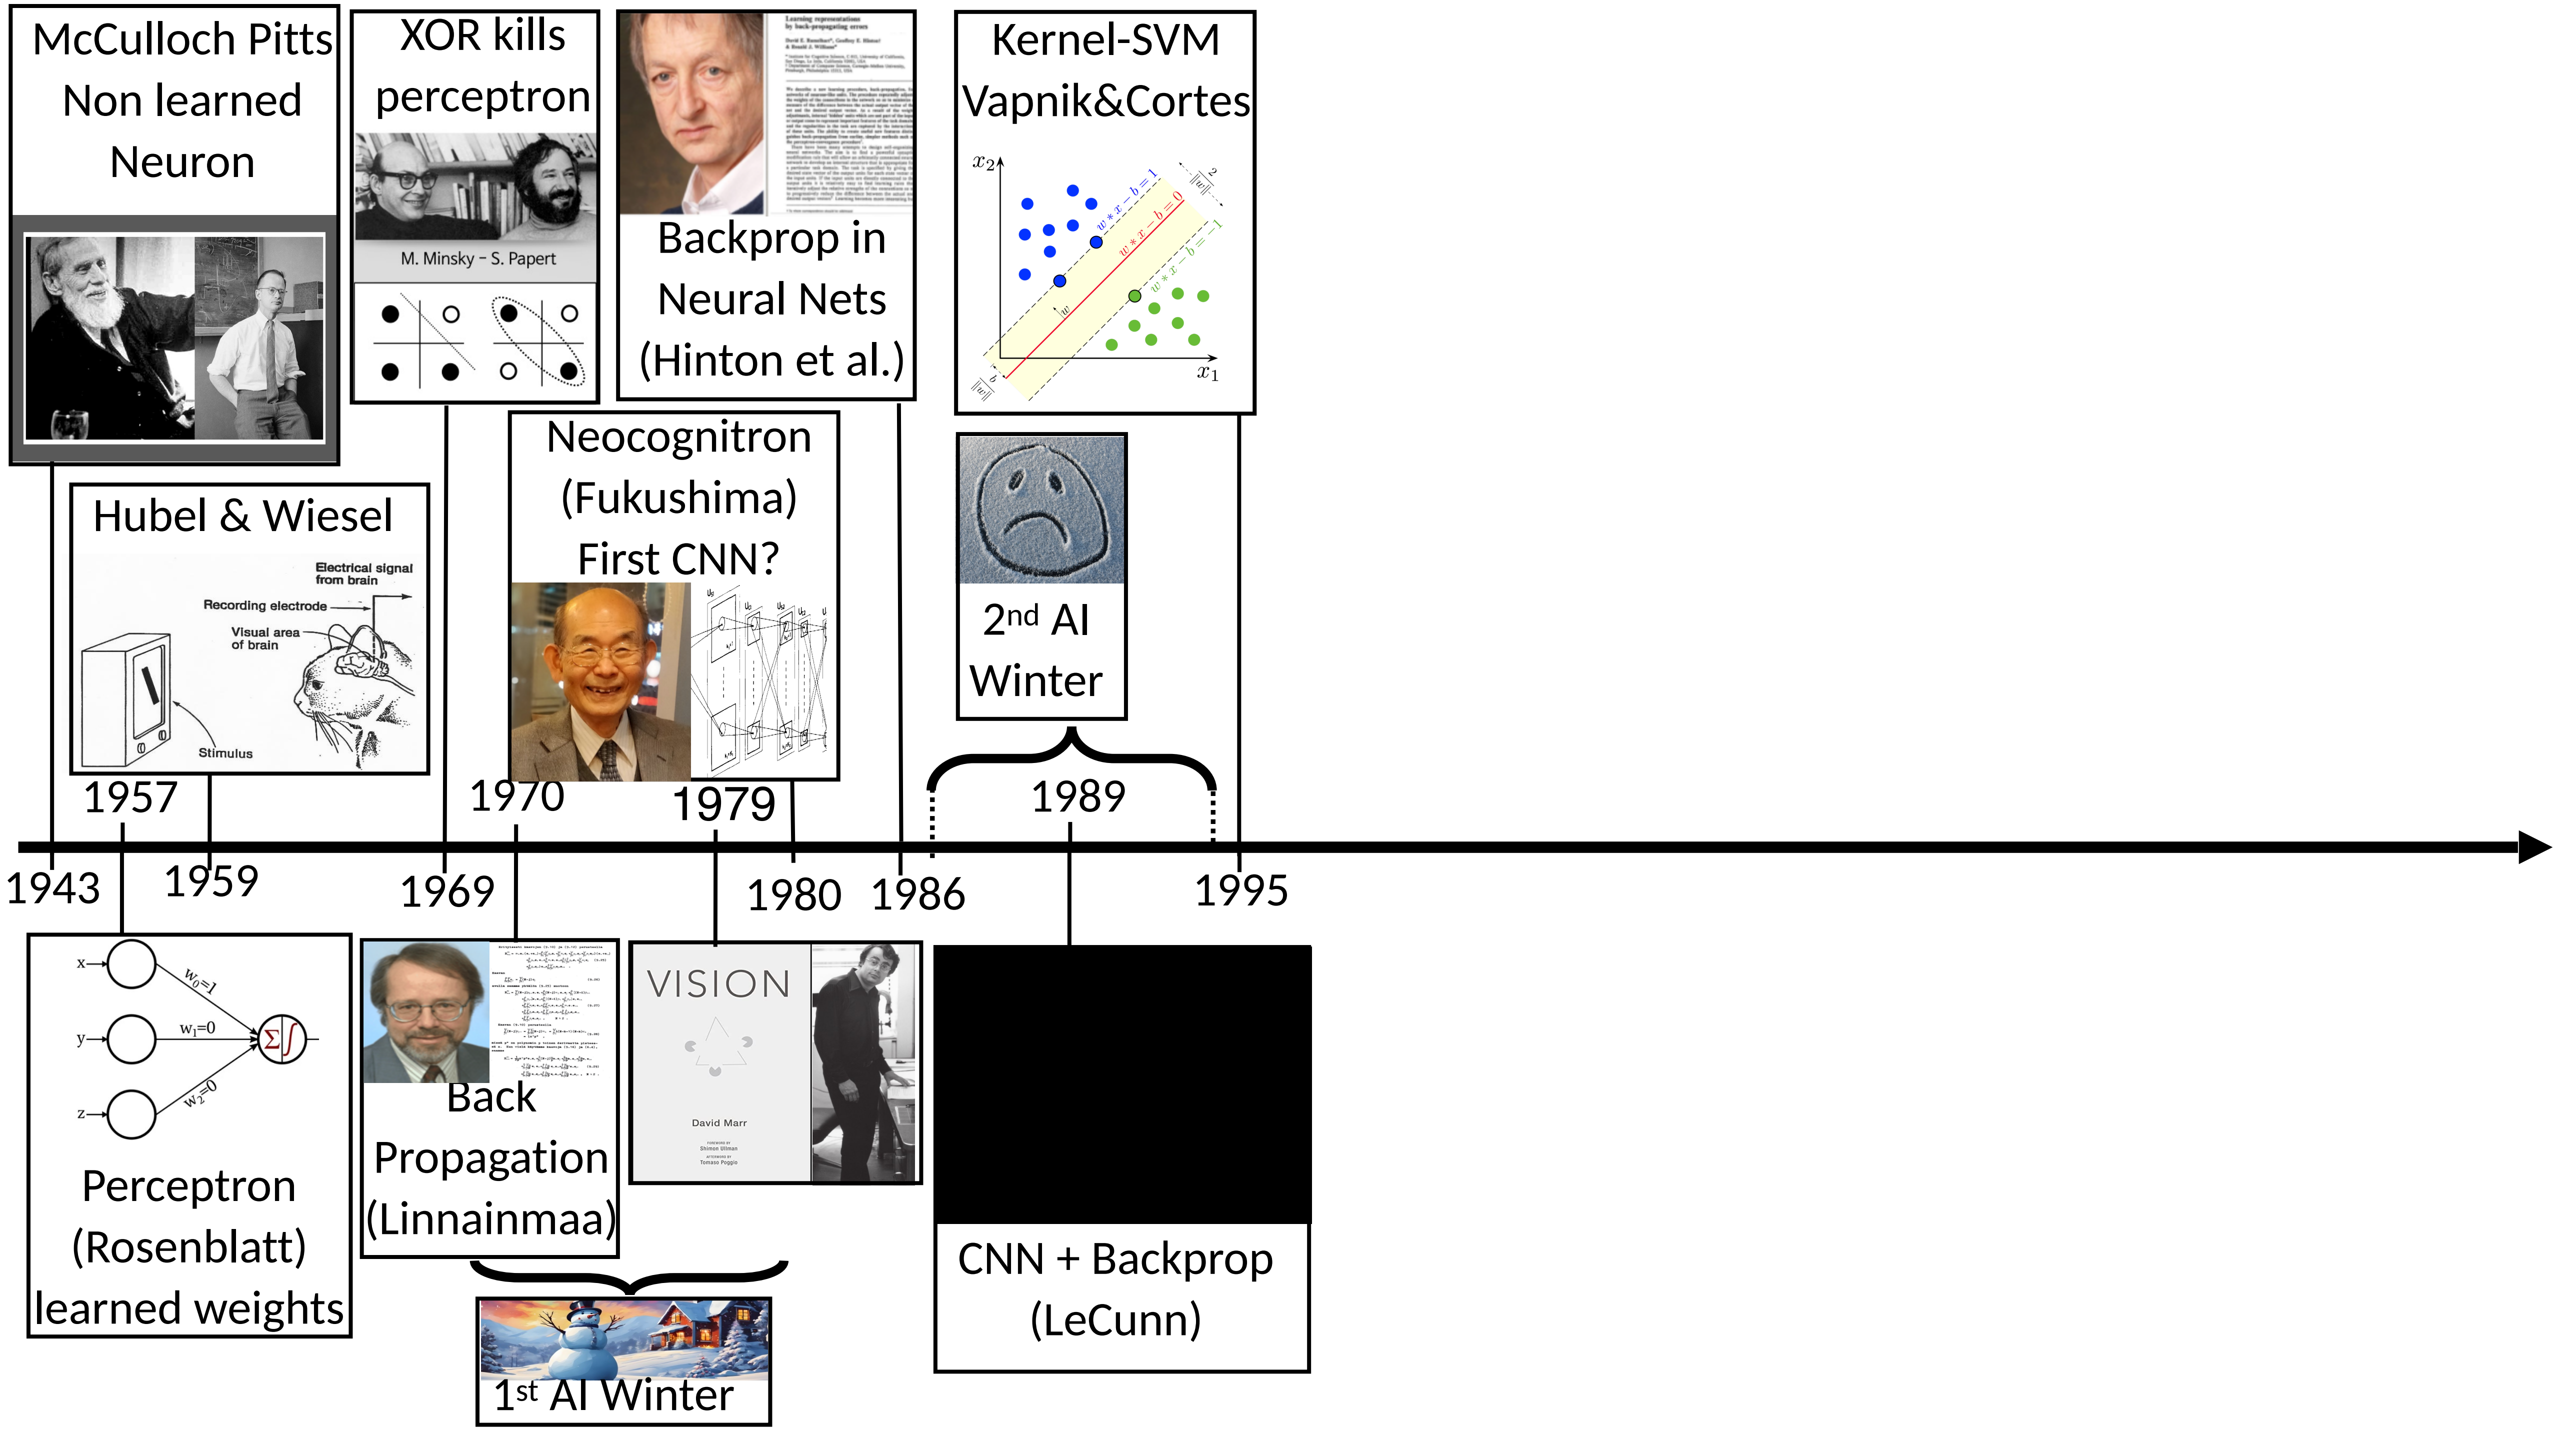

McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Backprop in
Neural Nets
(Hinton et al.)
Kernel-SVM
Vapnik&Cortes
Hubel & Wiesel
Electrical signal from brain
Recording electrode
Visual area of brain
Stimulus
Neocognitron
(Fukushima)
First CNN?
2nd AI
Winter
1957
1970
1979
1989
1943
1959
1969
1980
1986
1995
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
VISION
David Marr
CNN + Backprop
(LeCunn)
1st AI Winter

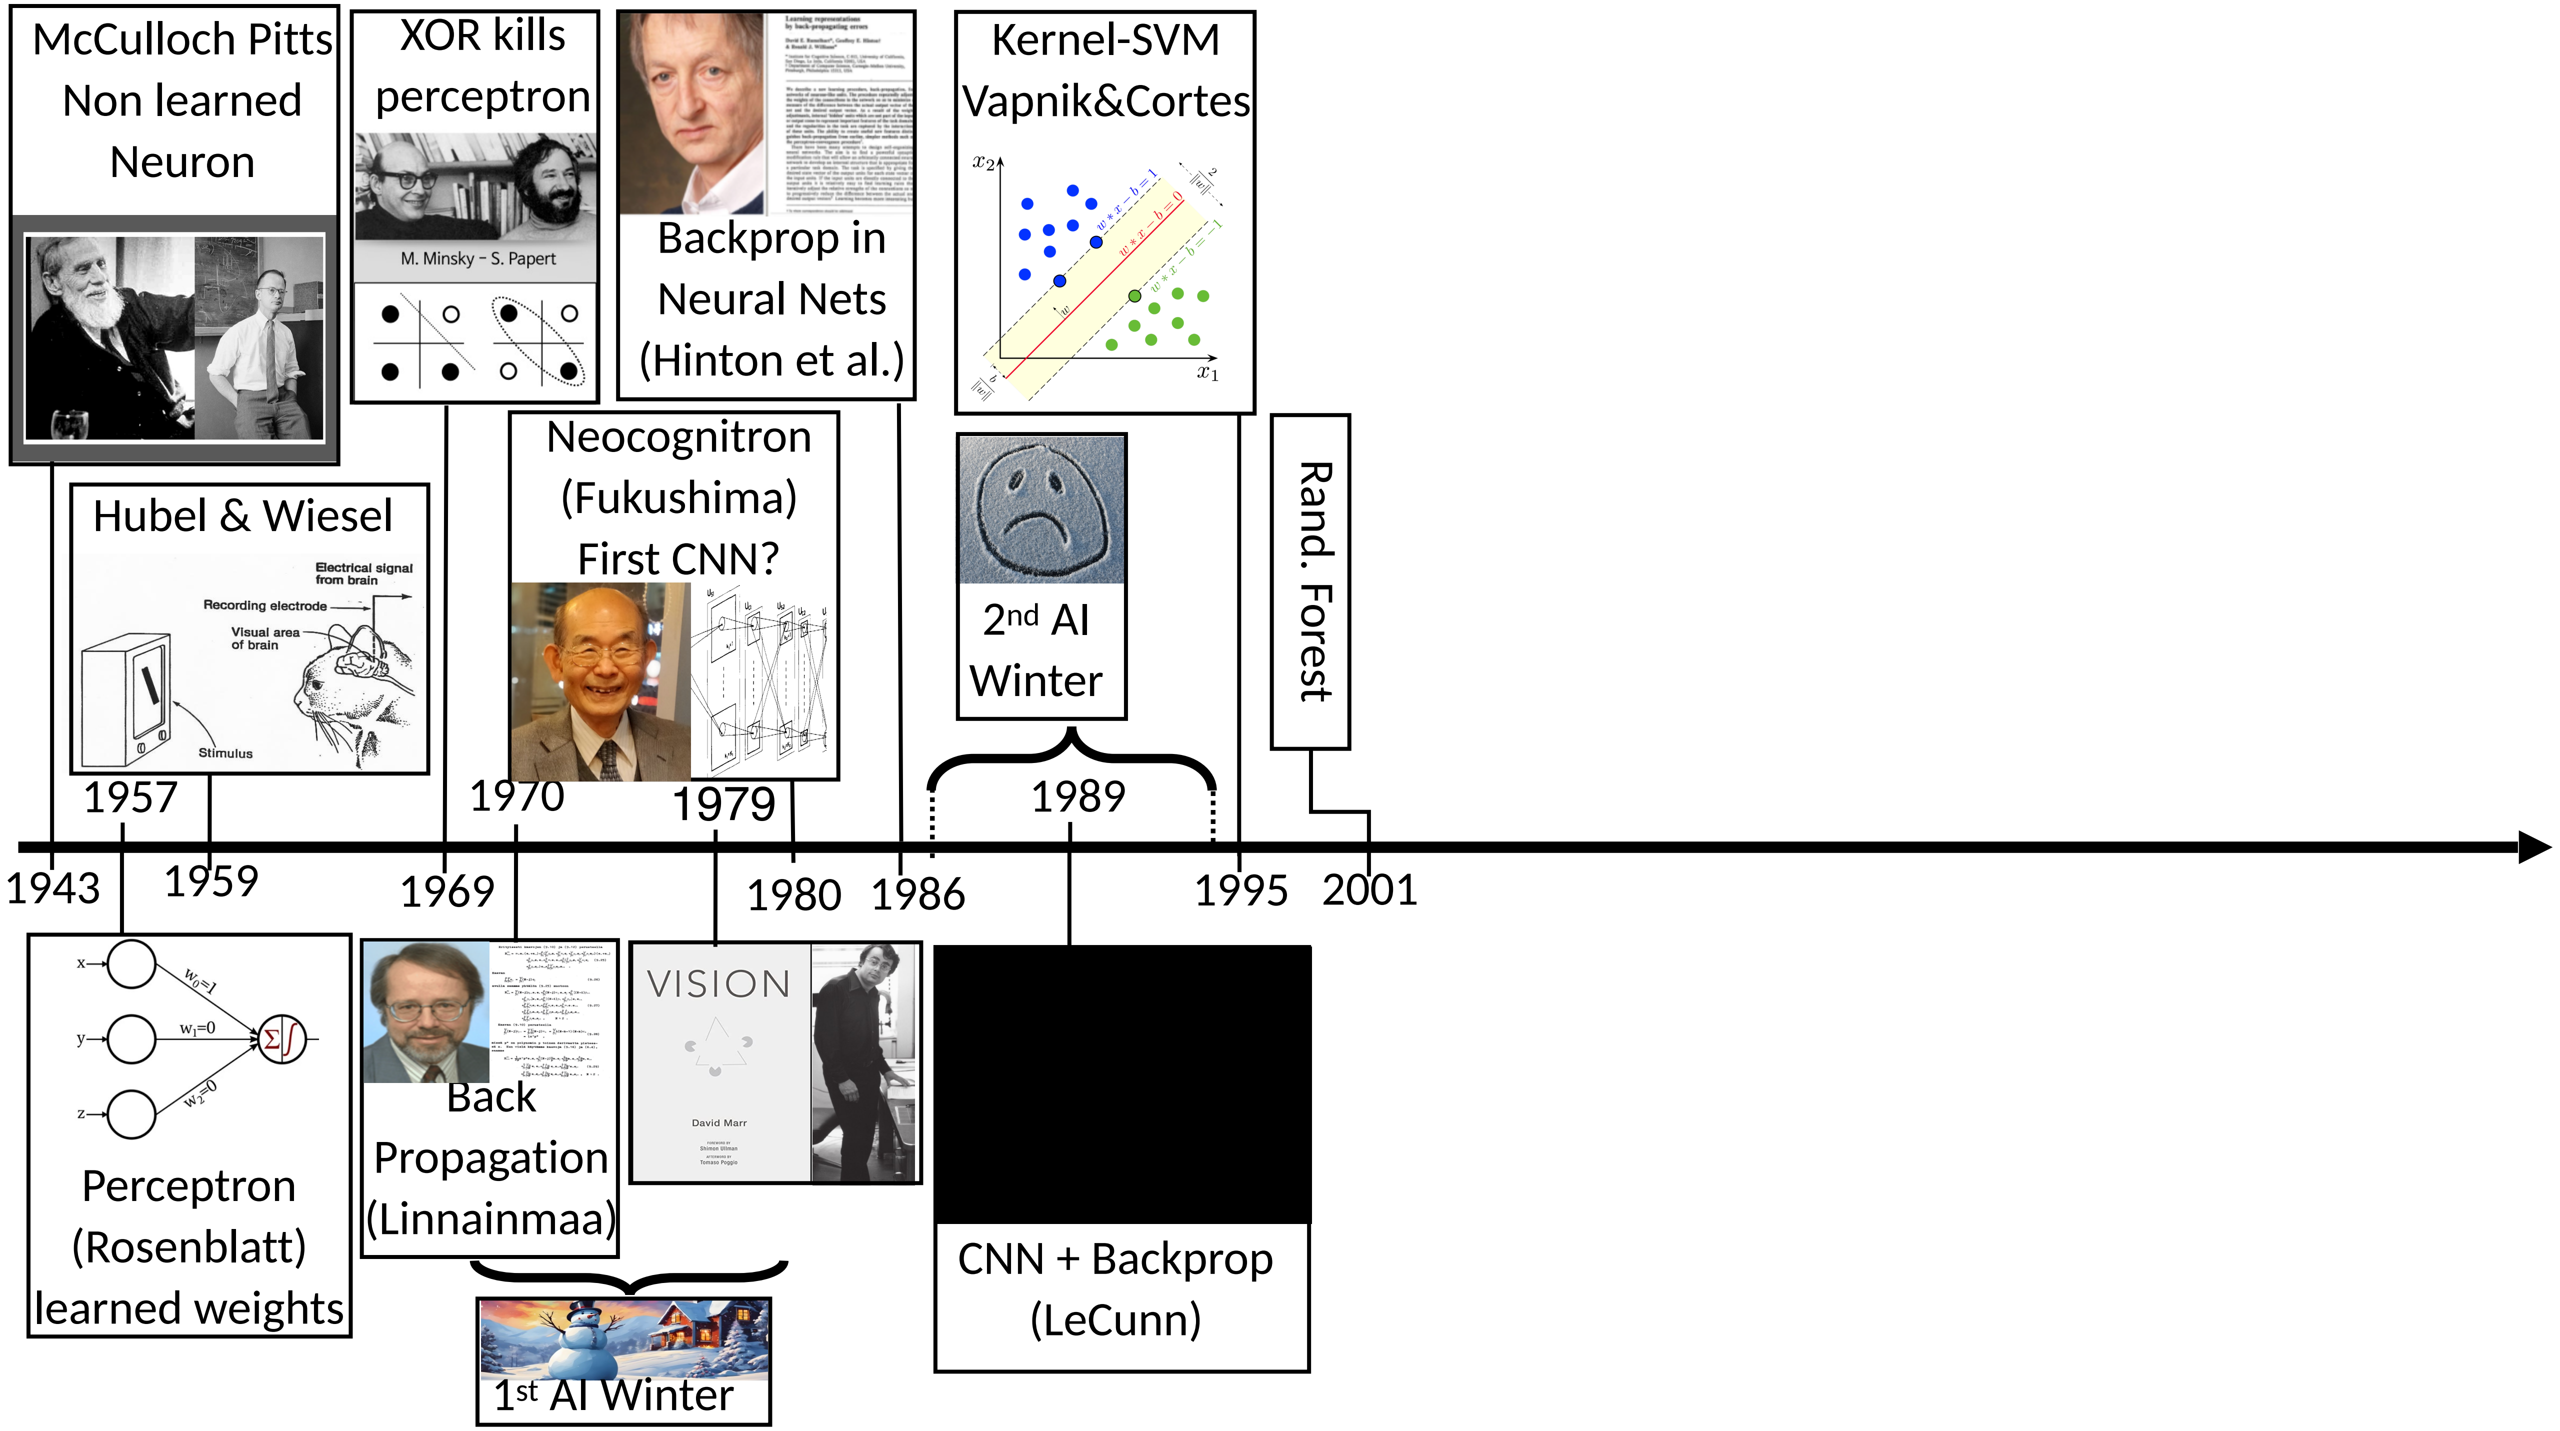
McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Backprop in
Neural Nets
(Hinton et al.)
Kernel-SVM
Vapnik&Cortes
Hubel & Wiesel
Electrical signal from brain
Recording electrode
Visual area of brain
Stimulus
Neocognitron
(Fukushima)
First CNN?
2nd AI
Winter
Rand. Forest
1957
1970
1979
1989
1943
1959
1969
1980
1986
1995
2001
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
VISION
David Marr
CNN + Backprop
(LeCunn)
1st AI Winter

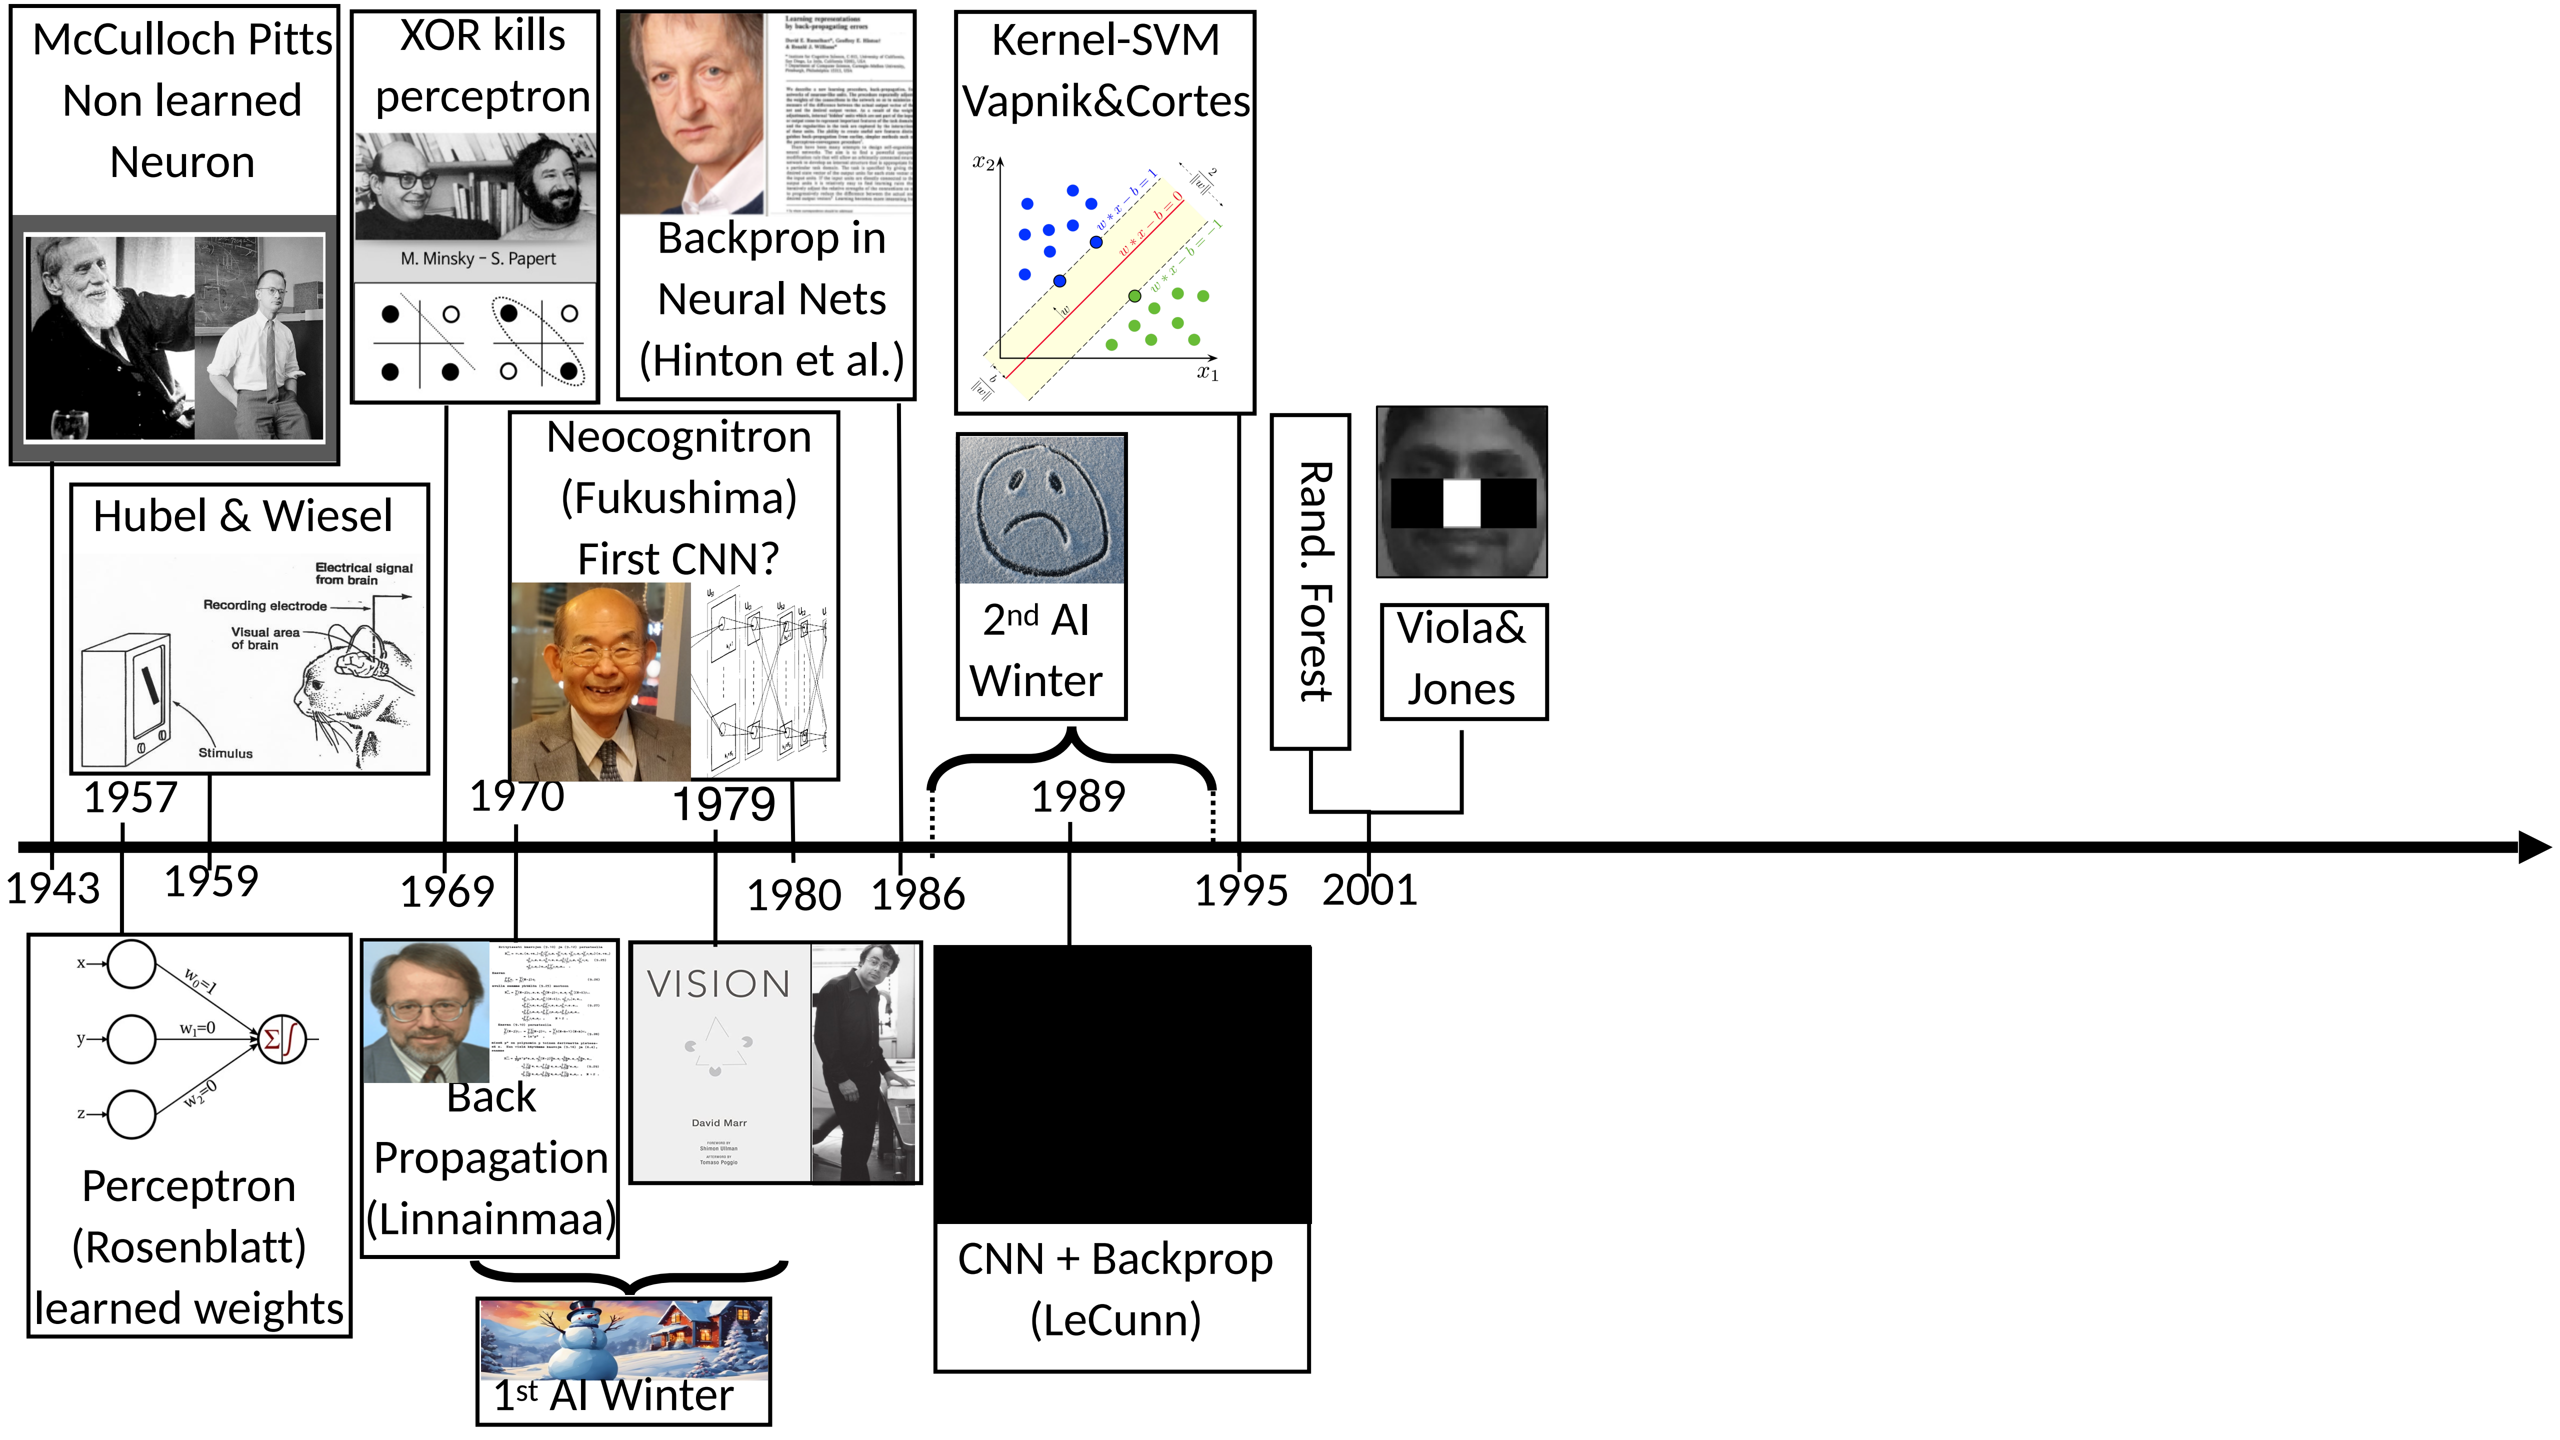
McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Backprop in
Neural Nets
(Hinton et al.)
Kernel-SVM
Vapnik&Cortes
Hubel & Wiesel
Electrical signal
from brain
Recording electrode
Visual area
of brain
Stimulus
Neocognitron
(Fukushima)
First CNN?
2nd AI
Winter
Rand. Forest
Viola&
Jones
1957
1970
1979
1989
1943
1959
1969
1980
1986
1995
2001
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
VISION
David Marr
CNN + Backprop
(LeCunn)
1st AI Winter

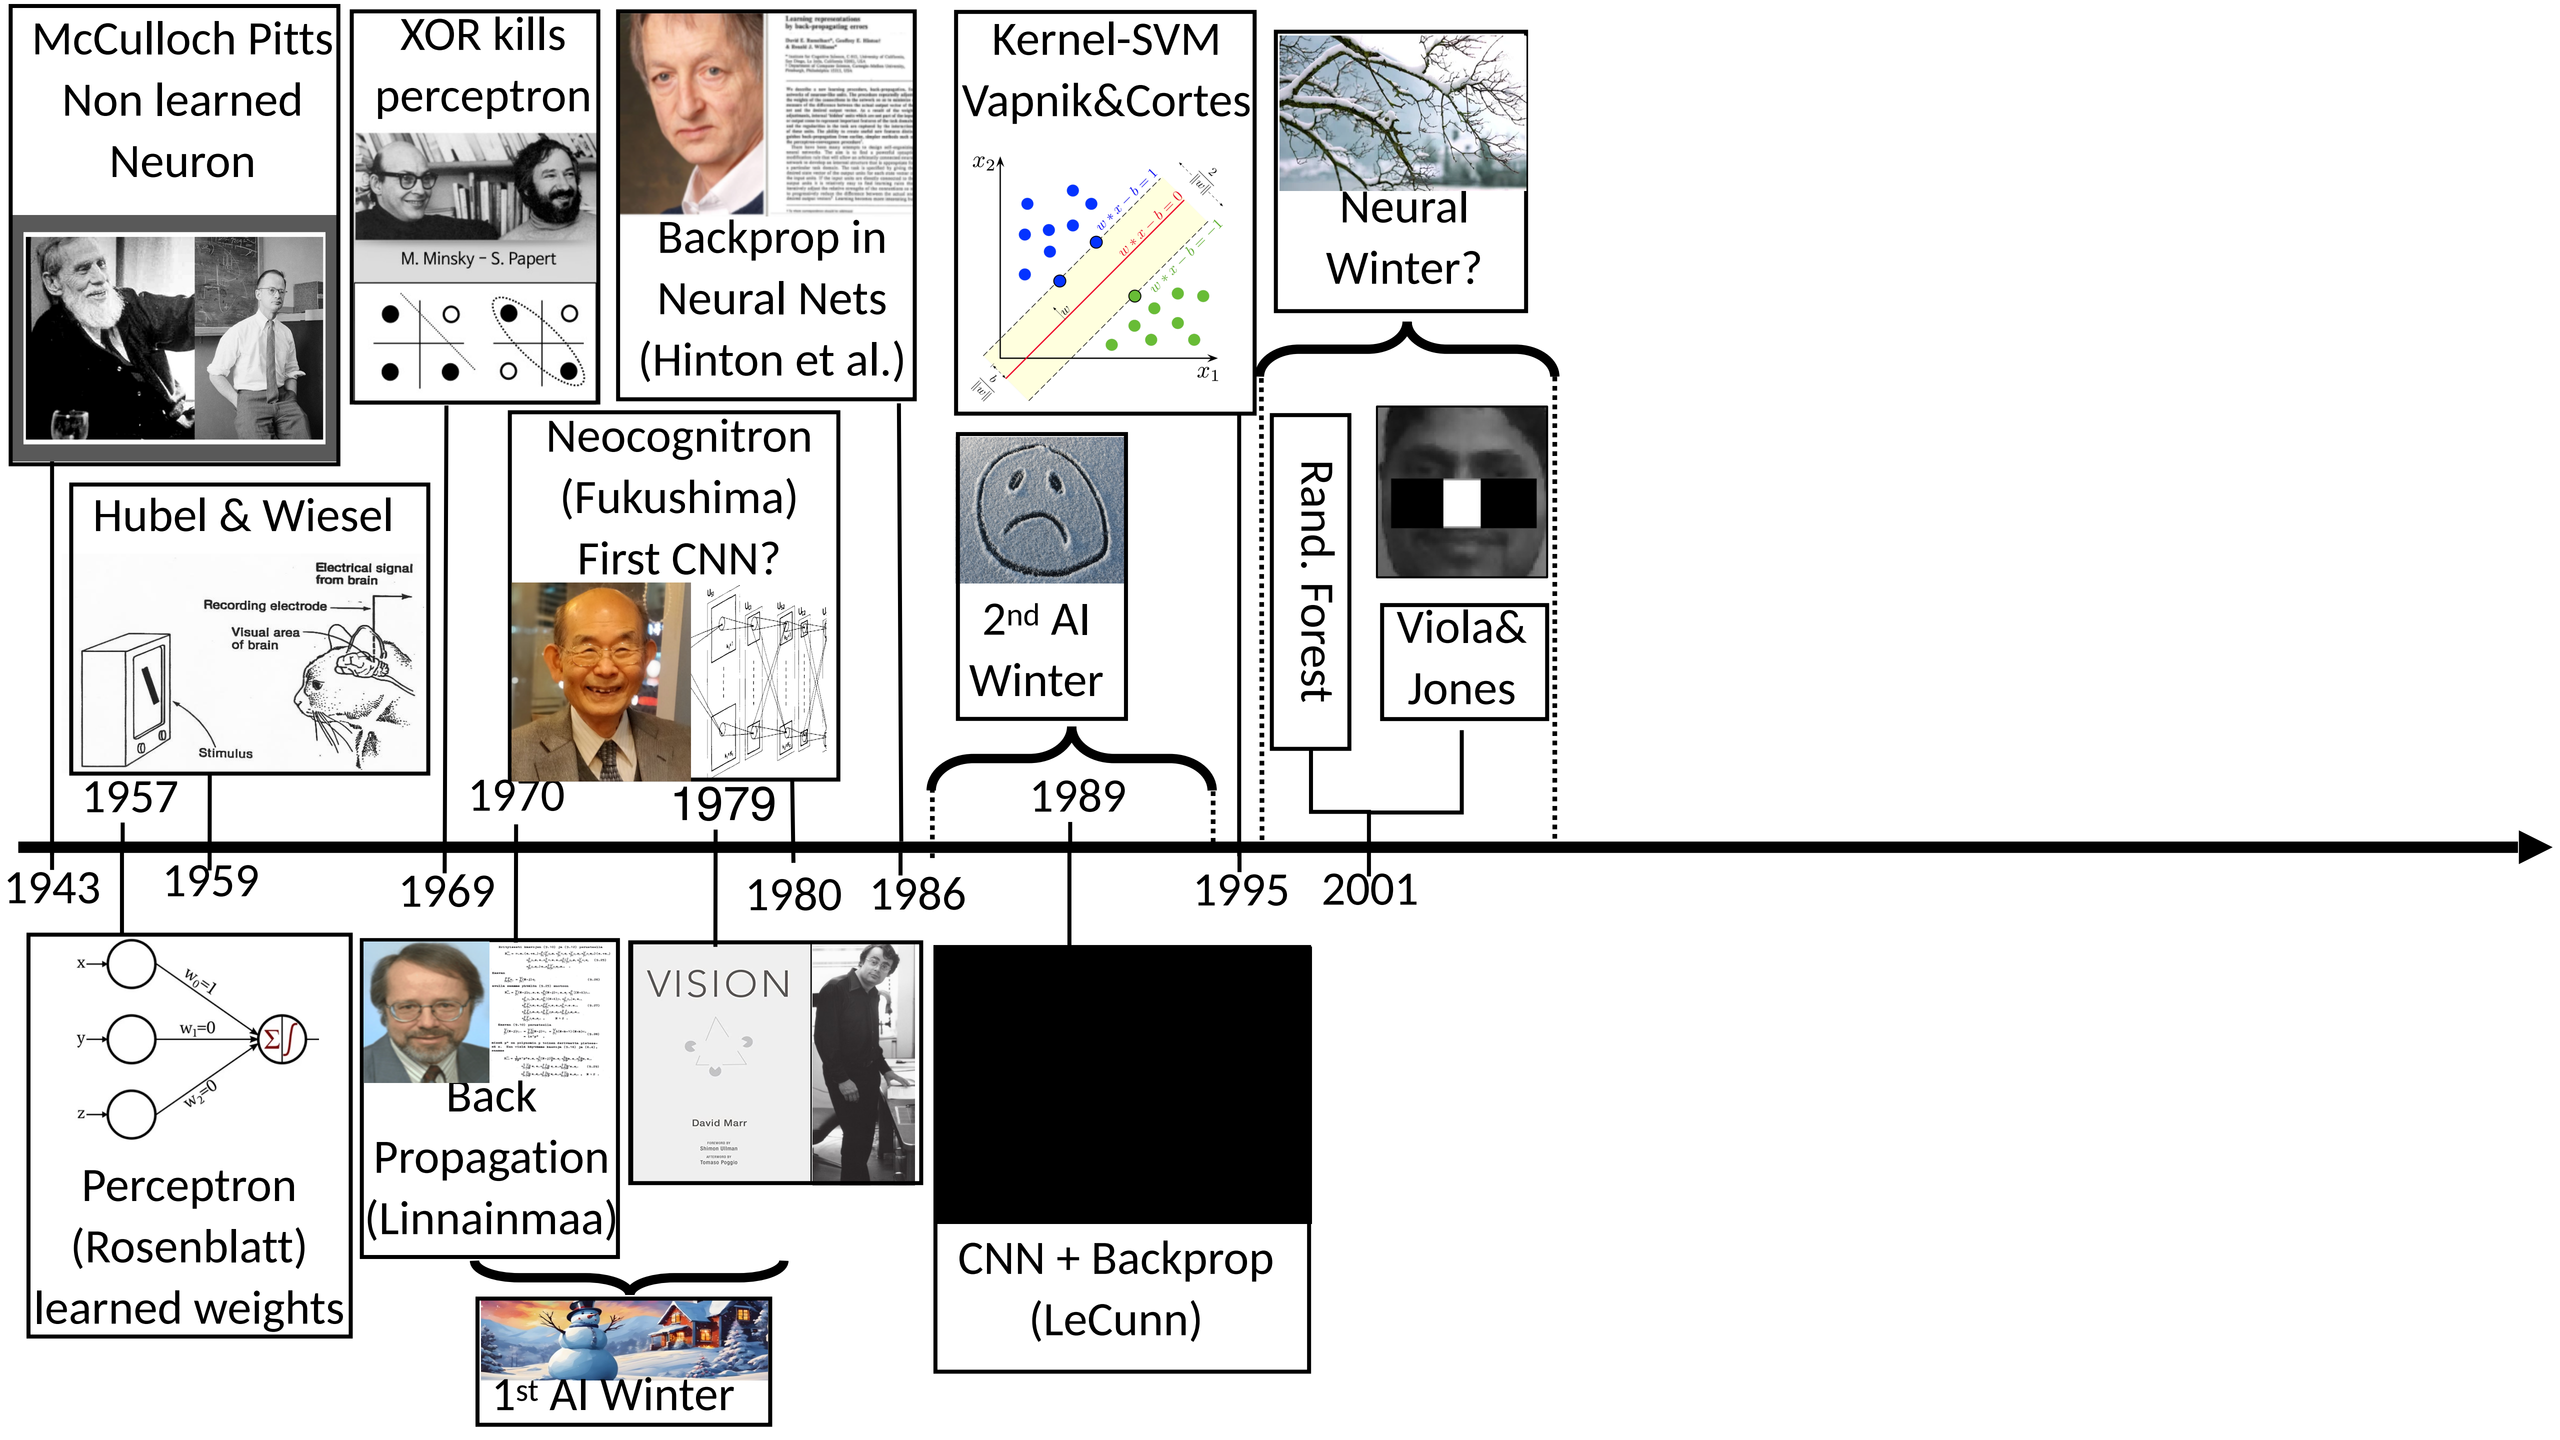
McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Backprop in
Neural Nets
(Hinton et al.)
Kernel-SVM
Vapnik&Cortes
Neural
Winter?
Hubel & Wiesel
Electrical signal from brain
Recording electrode
Visual area of brain
Stimulus
Neocognitron
(Fukushima)
First CNN?
2nd AI
Winter
Rand. Forest
Viola&
Jones
1957
1970
1979
1989
1943
1959
1969
1980
1986
1995
2001
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
VISION
David Marr
CNN + Backprop
(LeCunn)
1st AI Winter

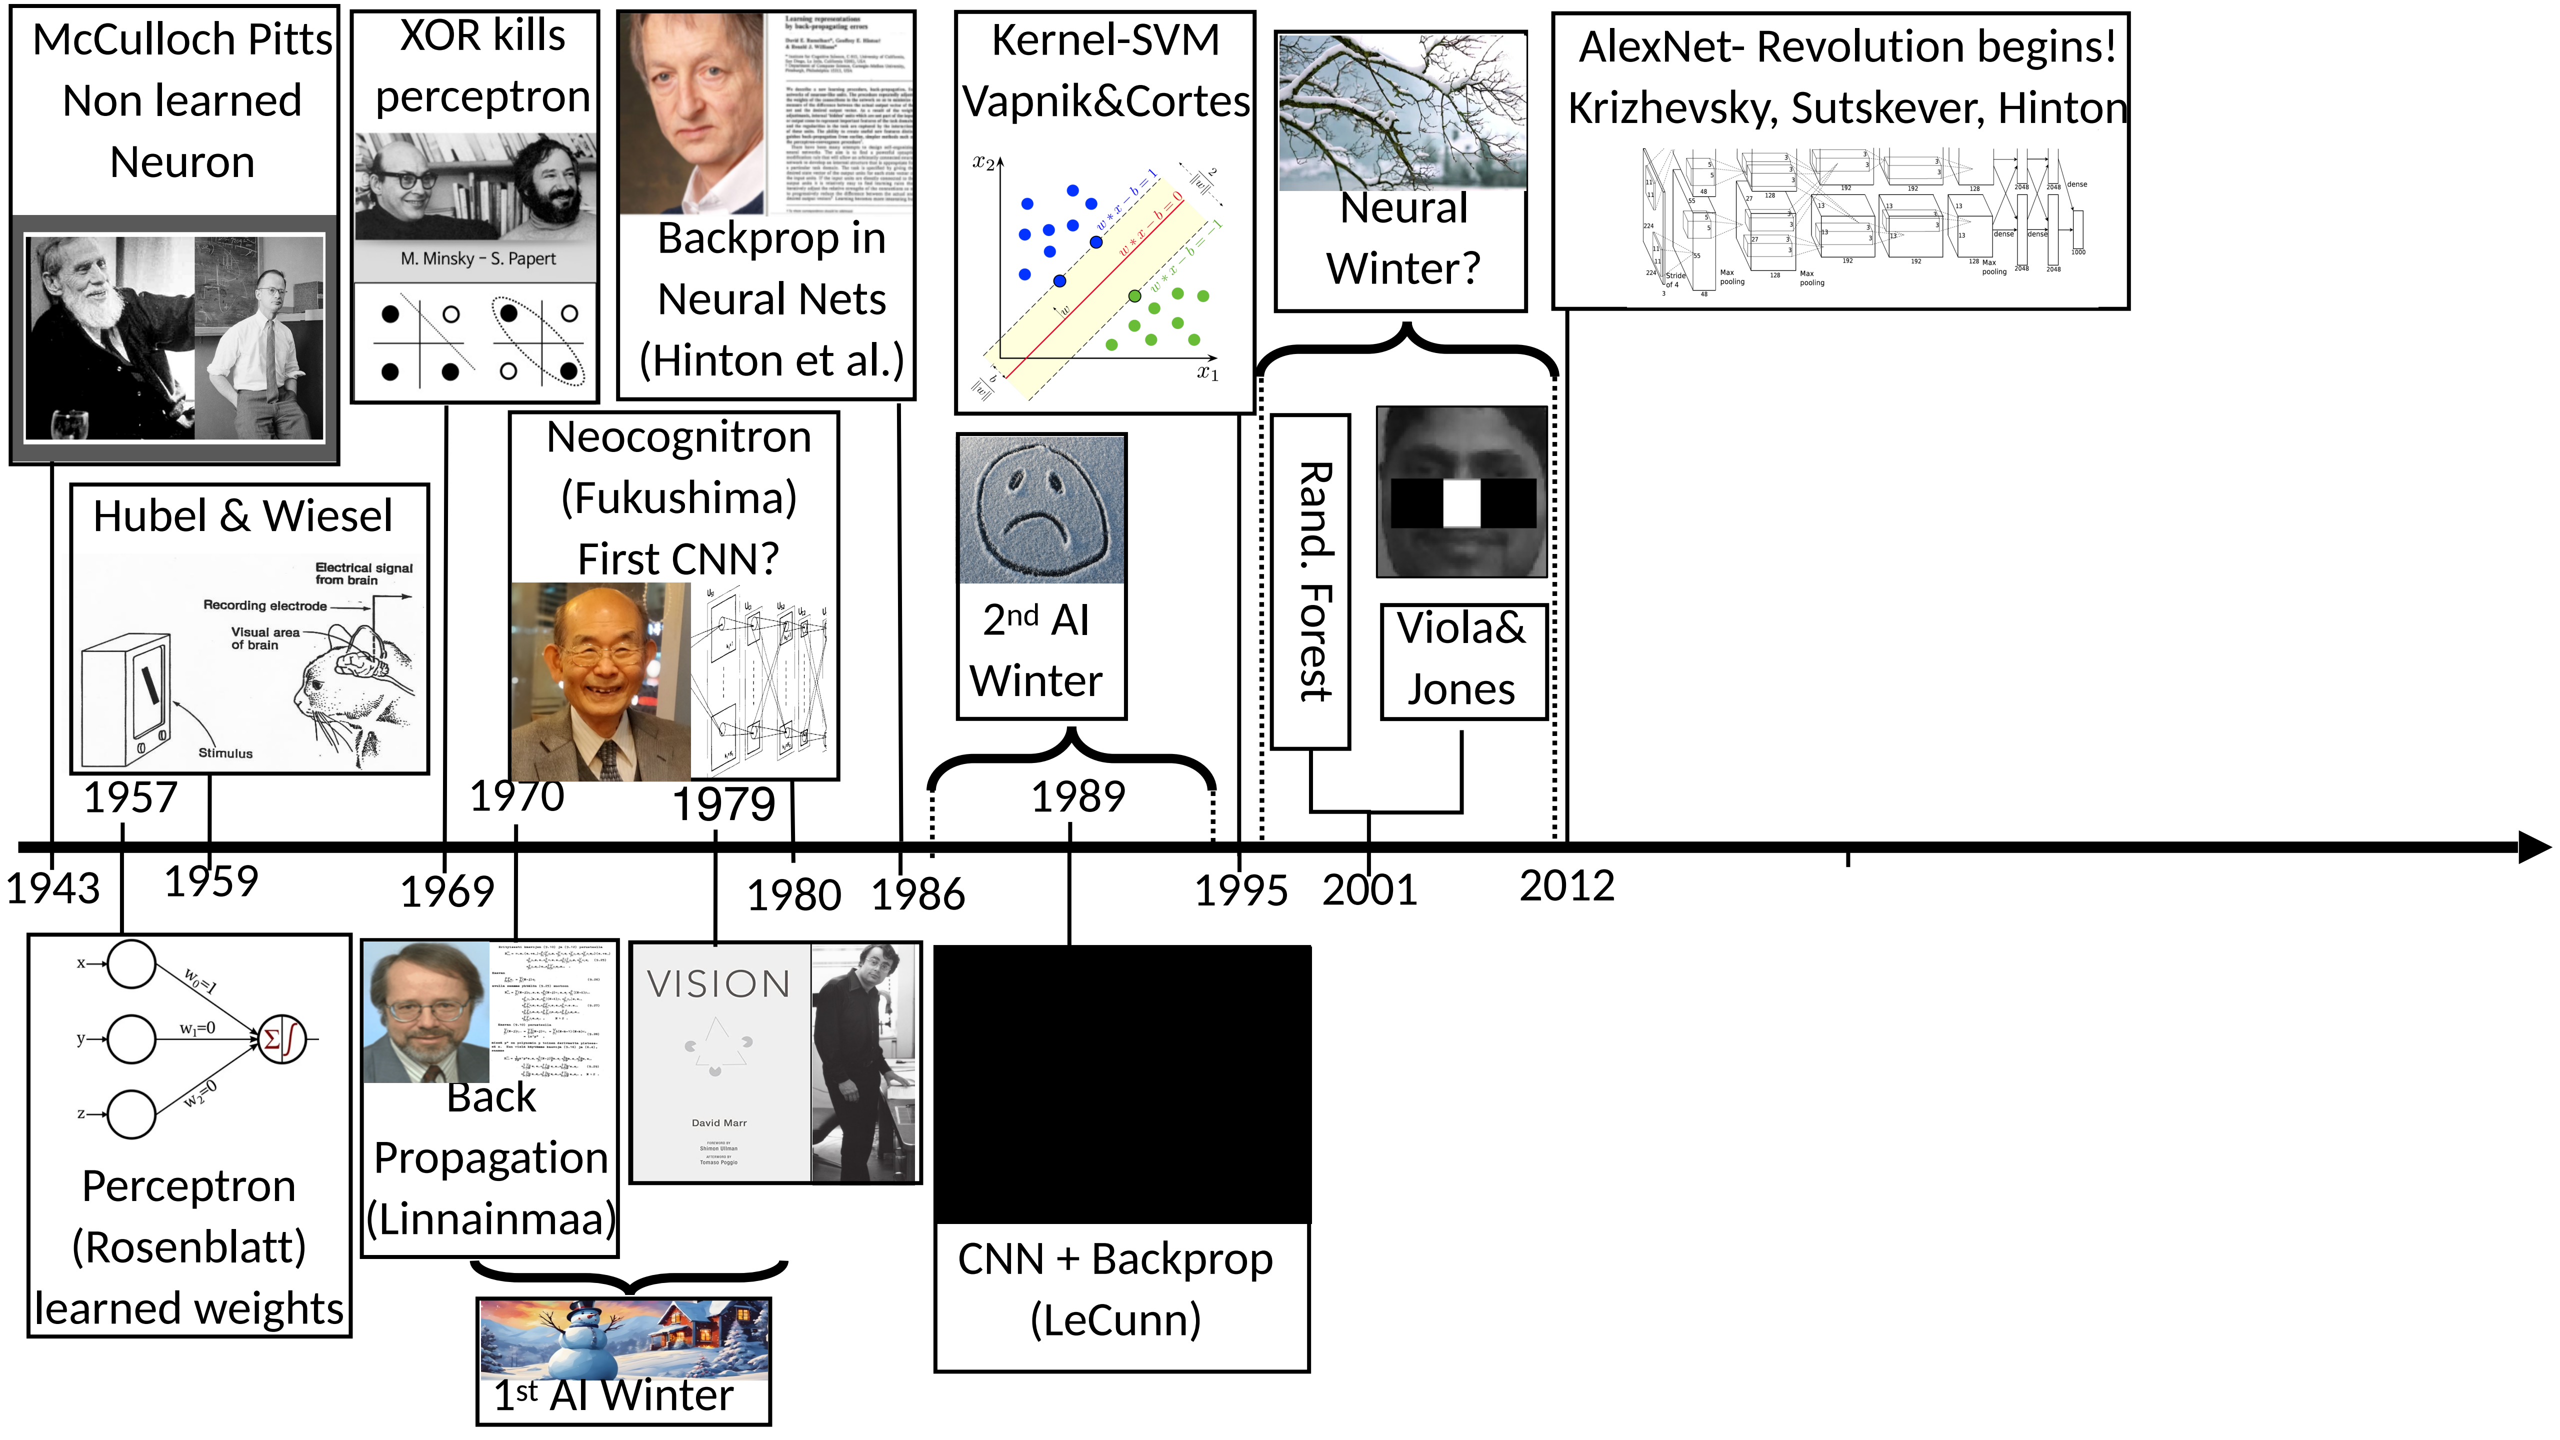
McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Backprop in
Neural Nets
(Hinton et al.)
Kernel-SVM
Vapnik&Cortes
Neural
Winter?
AlexNet- Revolution begins!
Krizhevsky, Sutskever, Hinton
Hubel & Wiesel
Electrical signal
from brain
Recording electrode
Visual area
of brain
Stimulus
Neocognitron
(Fukushima)
First CNN?
2nd AI
Winter
Rand. Forest
Viola&
Jones
1957
1970
1979
1989
1943
1959
1969
1980
1986
1995
2001
2012
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
VISION
David Marr
CNN + Backprop
(LeCunn)
1st AI Winter

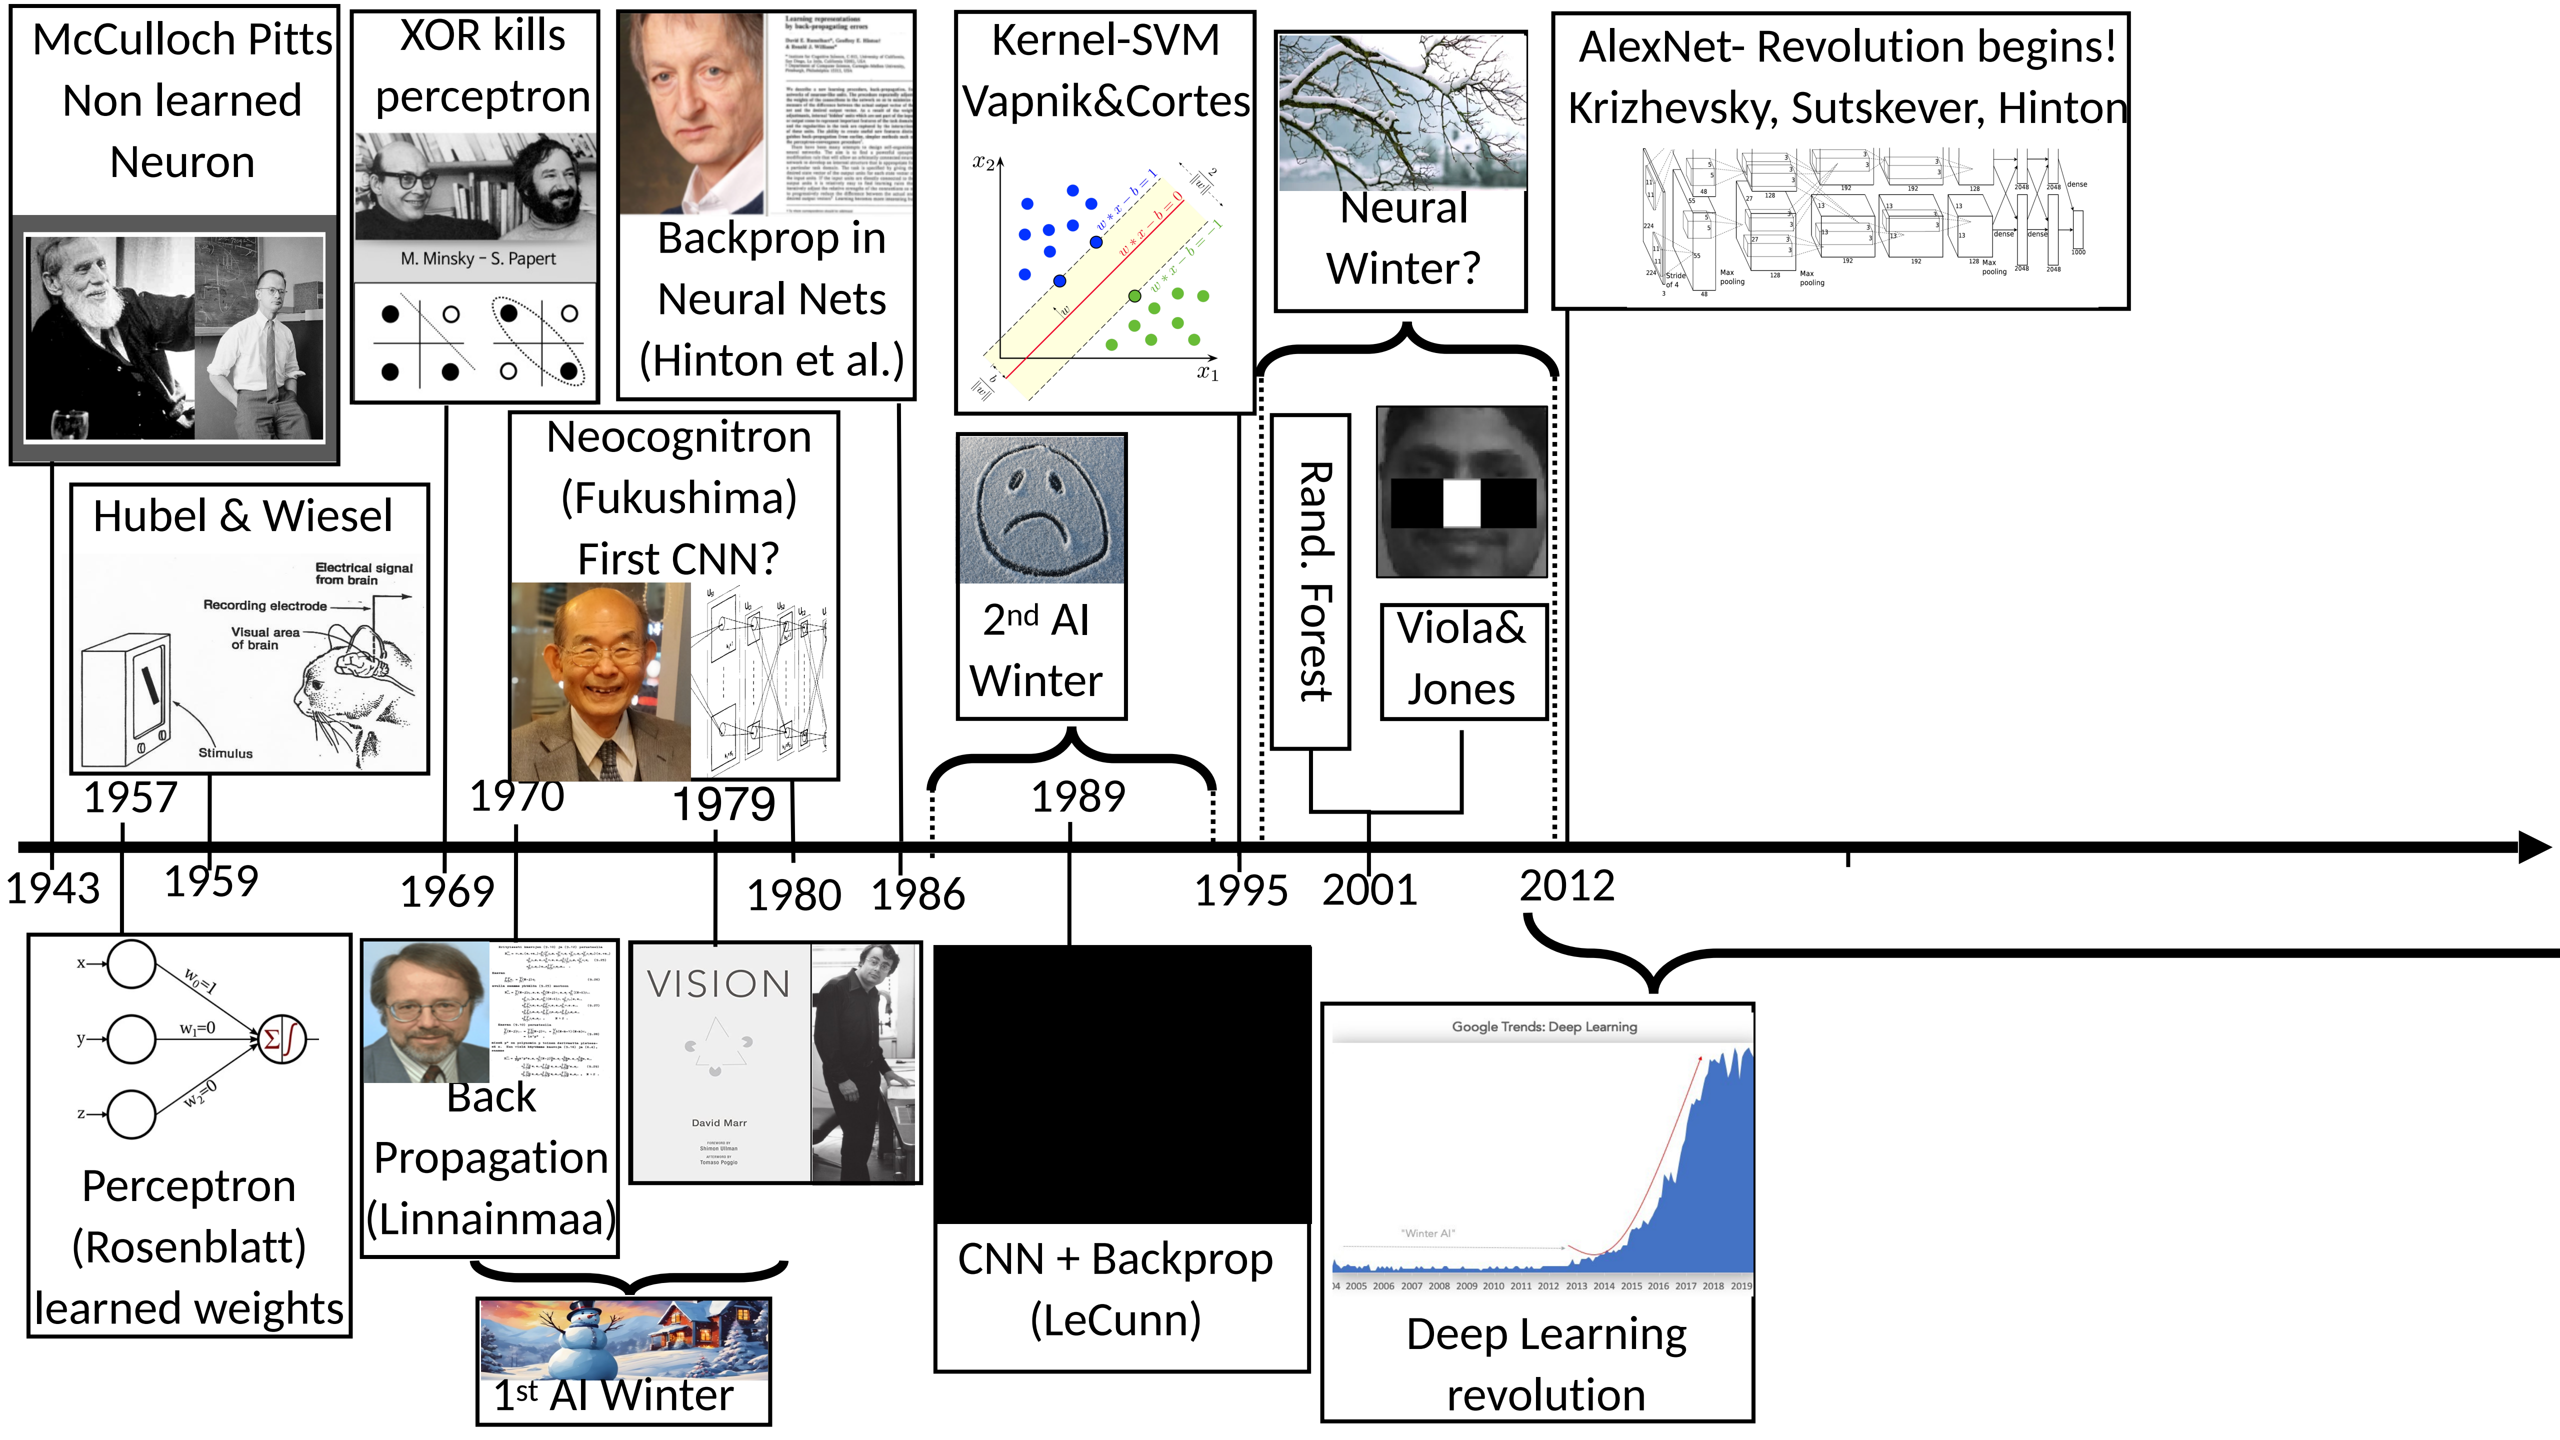
McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Backprop in
Neural Nets
(Hinton et al.)
Learning representations by back-propagating errors
Kernel-SVM
Vapnik&Cortes
Neural
Winter?
AlexNet- Revolution begins!
Krizhevsky, Sutskever, Hinton
Hubel & Wiesel
Electrical signal from brain
Recording electrode
Visual area of brain
Stimulus
Neocognitron
(Fukushima)
First CNN?
2nd AI
Winter
Rand. Forest
Viola&
Jones
1957
1970
1979
1989
1943
1959
1969
1980
1986
1995
2001
2012
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
1st AI Winter
VISION
David Marr
CNN + Backprop
(LeCunn)
Google Trends: Deep Learning
"Winter AI"
2005 2006 2007 2008 2009 2010 2011 2012 2013 2014 2015 2016 2017 2018 2019
Deep Learning
revolution

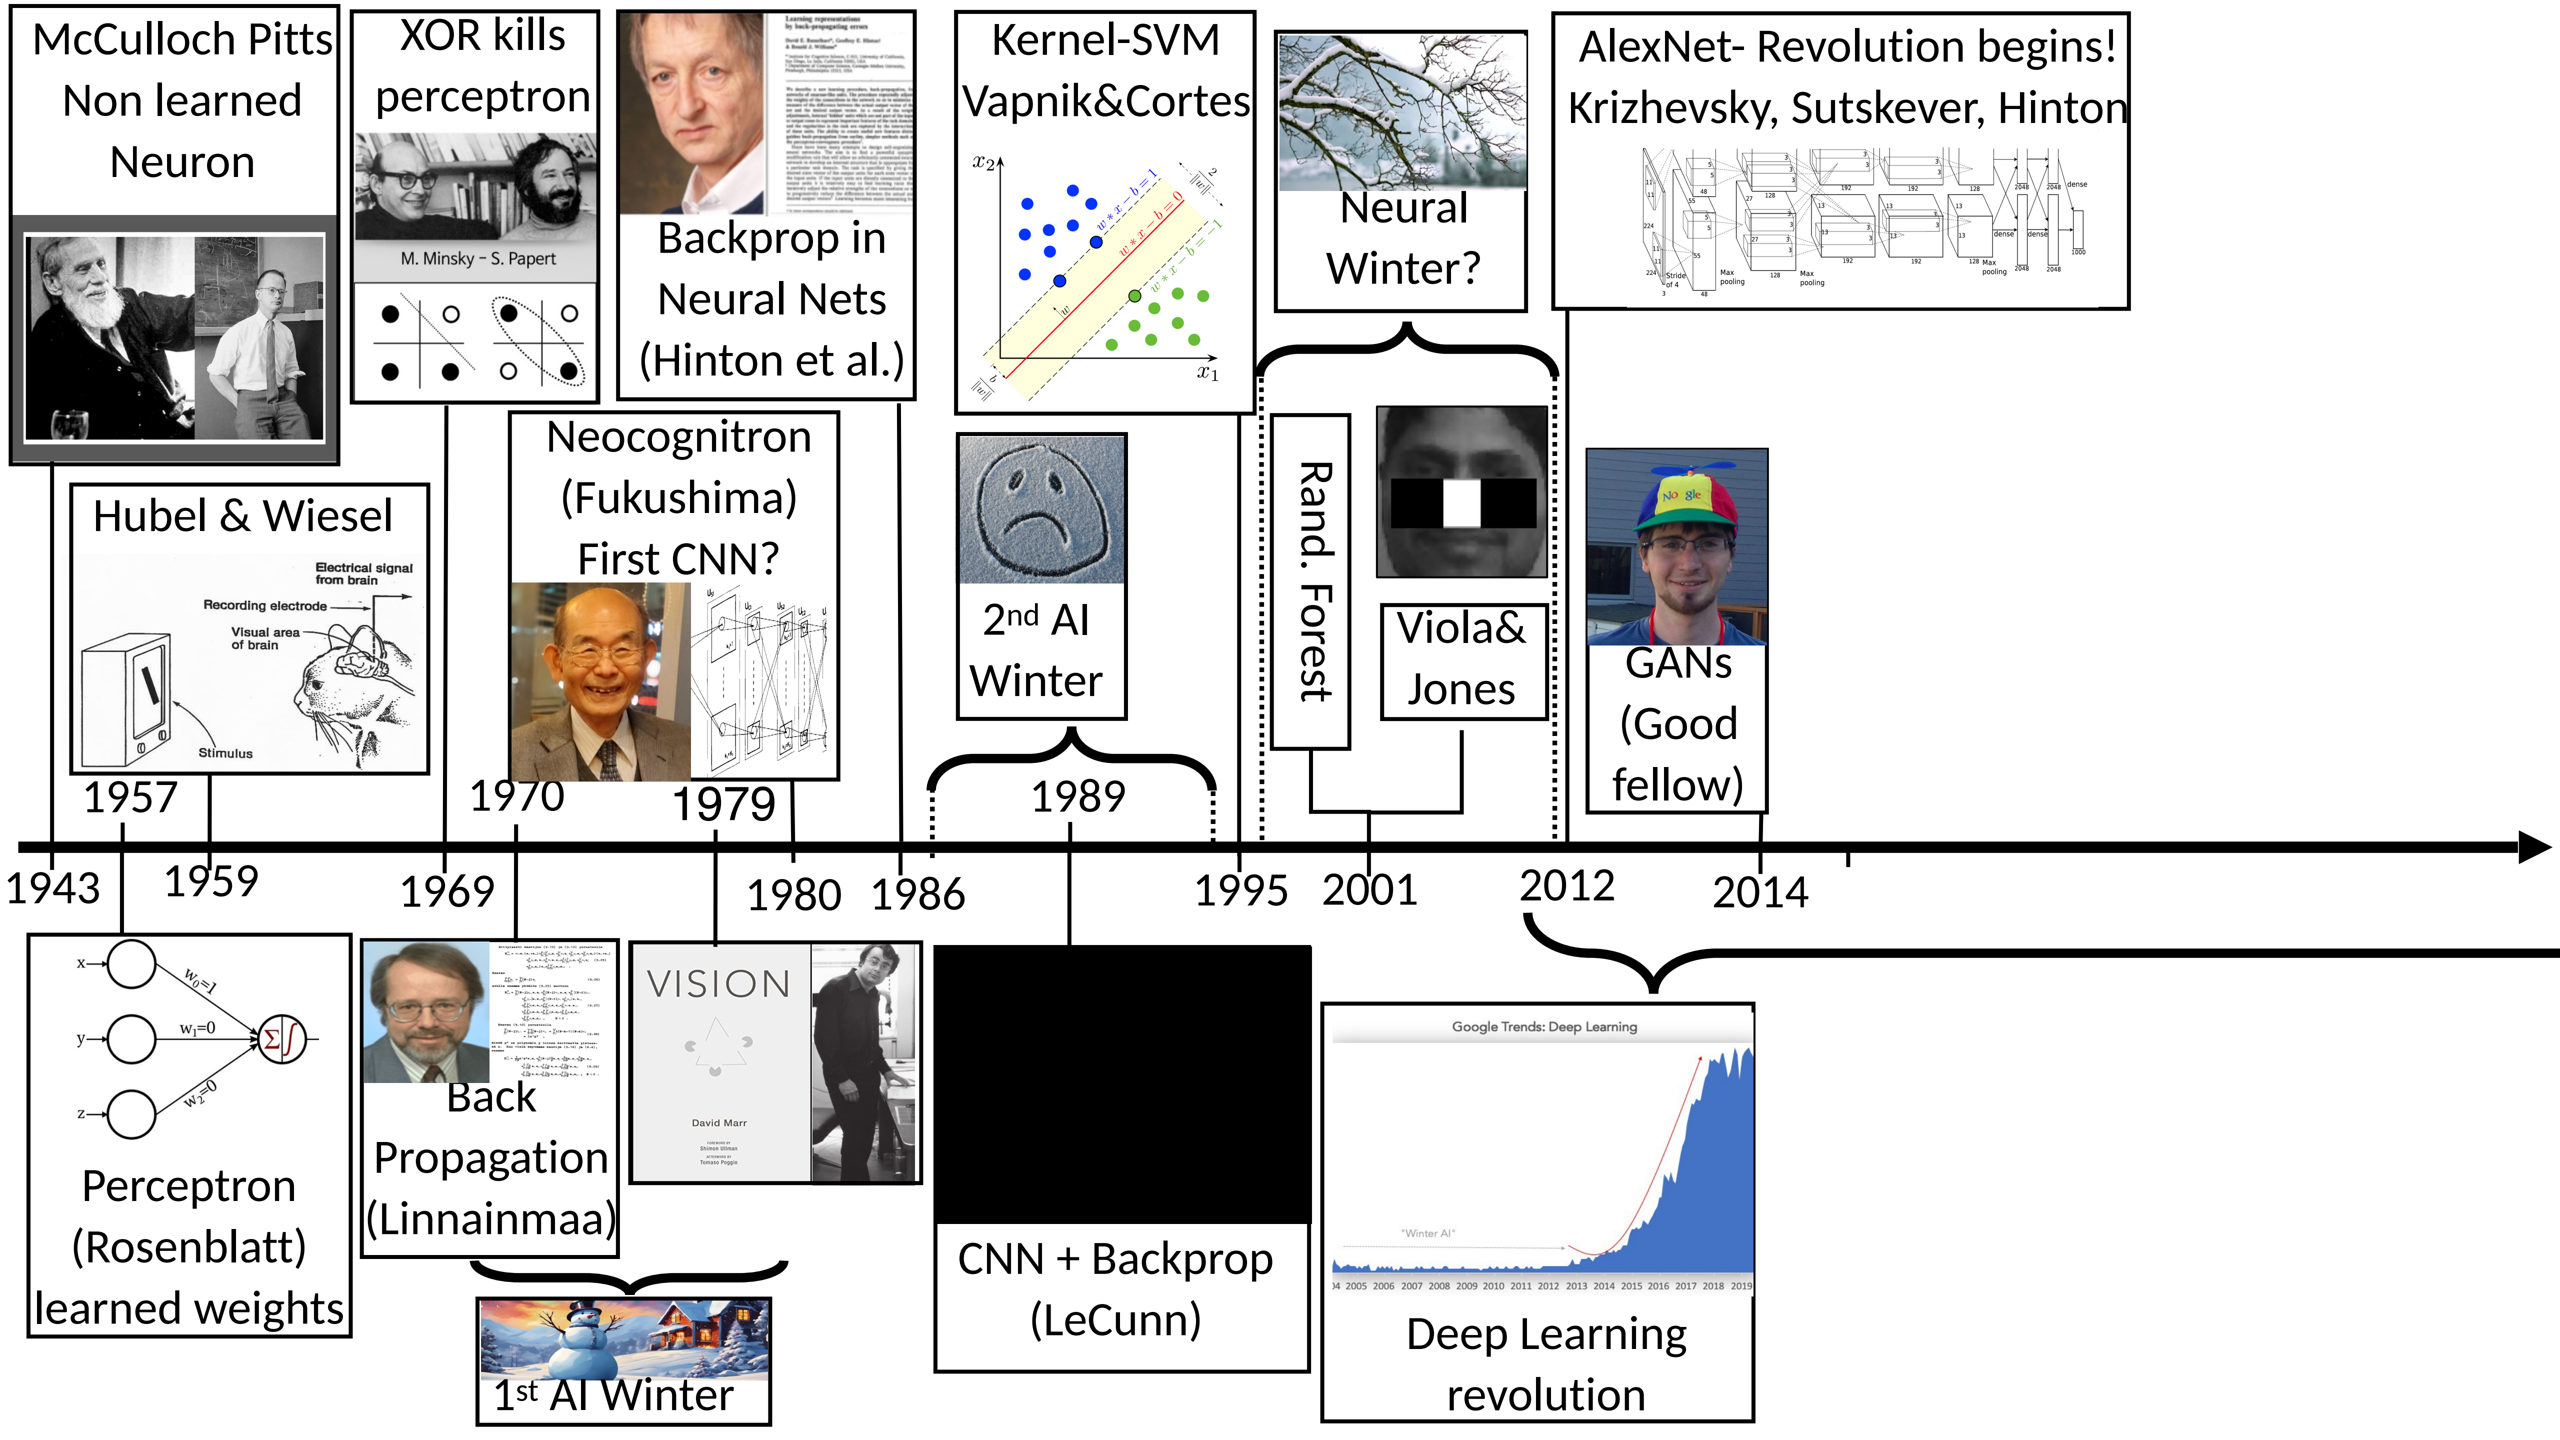

McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Backprop in
Neural Nets
(Hinton et al.)
Kernel-SVM
Vapnik&Cortes
Neural
Winter?
AlexNet- Revolution begins!
Krizhevsky, Sutskever, Hinton
Hubel & Wiesel
Electrical signal
from brain
Recording electrode
Visual area
of brain
Stimulus
Neocognitron
(Fukushima)
First CNN?
2nd AI
Winter
Rand. Forest
Viola&
Jones
GANs
(Good
fellow)
1957
1970
1979
1989
1943
1959
1969
1980
1986
1995
2001
2012
2014
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
VISION
David Marr
1st AI Winter
CNN + Backprop
(LeCunn)
Google Trends: Deep Learning
"Winter AI"
2005 2006 2007 2008 2009 2010 2011 2012 2013 2014 2015 2016 2017 2018 2019
Deep Learning
revolution

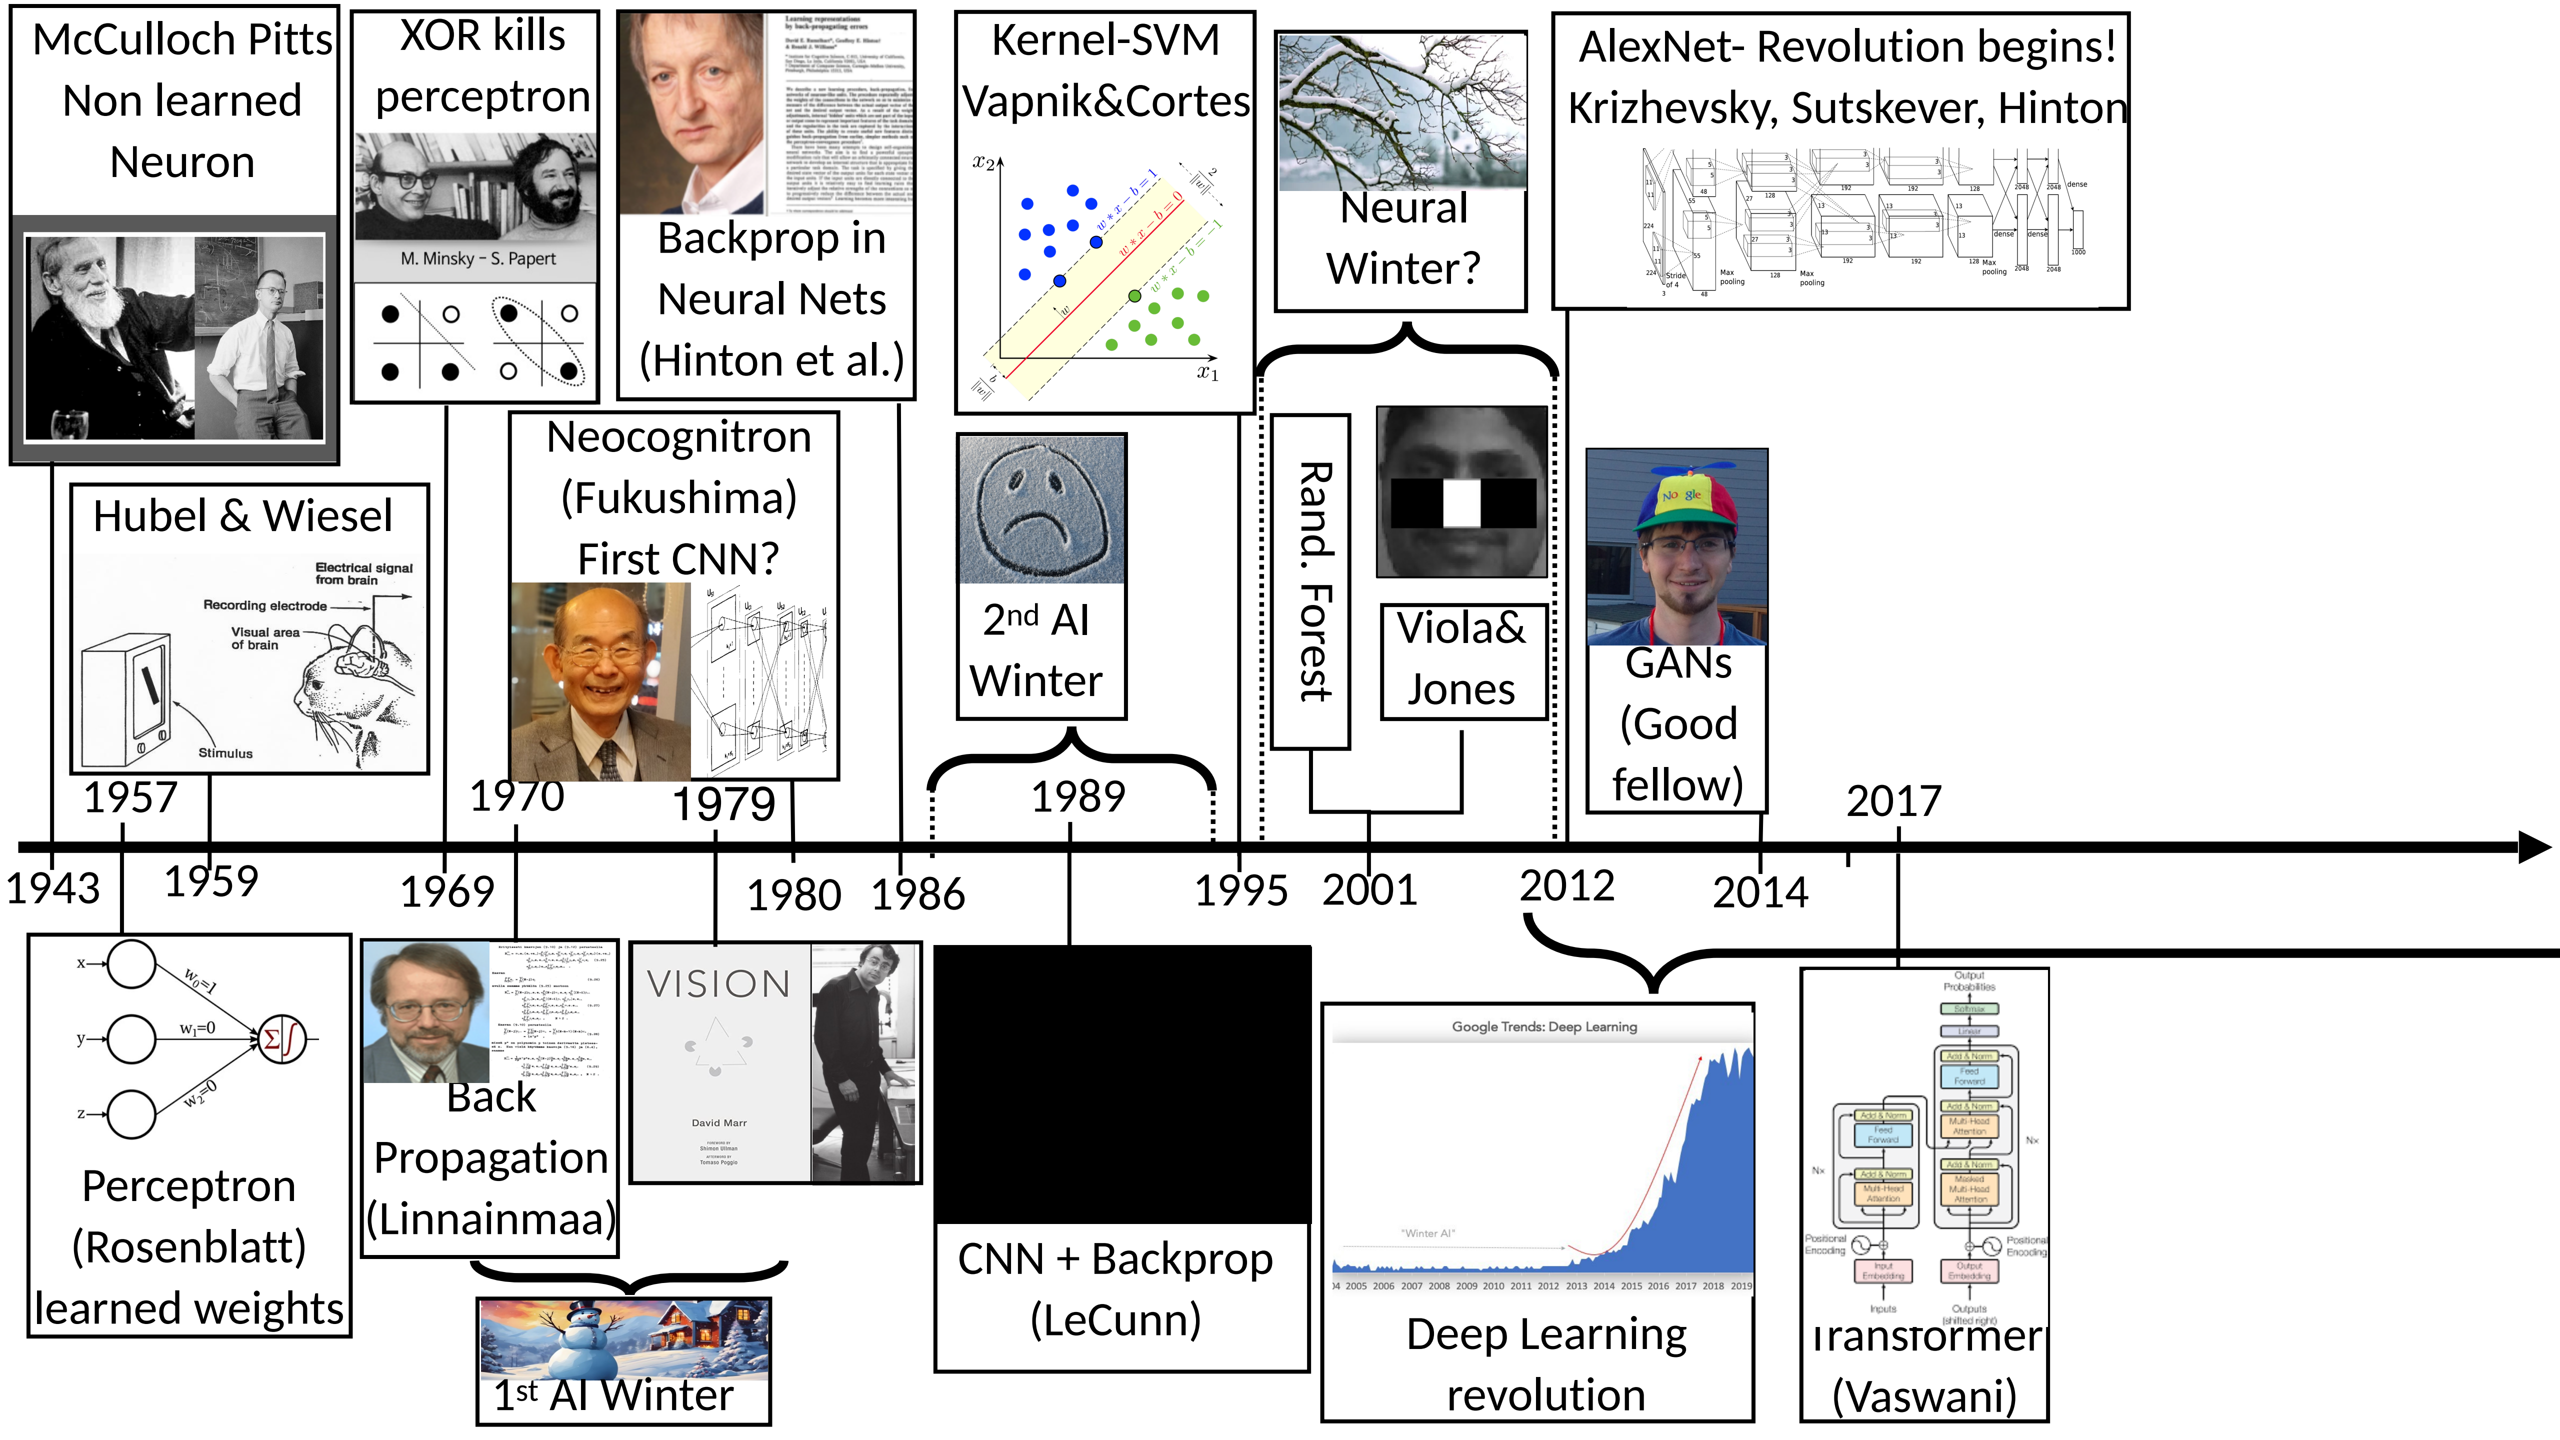
McCulloch Pitts Non learned Neuron
1943
XOR kills perceptron
M. Minsky – S. Papert
1969
Backprop in Neural Nets (Hinton et al.)
1986
Kernel-SVM Vapnik&Cortes
1995
Neural Winter?
AlexNet- Revolution begins! Krizhevsky, Sutskever, Hinton
2012
Hubel & Wiesel
Electrical signal from brain
Recording electrode
Visual area of brain
Stimulus
1957
1959
Neocognitron (Fukushima) First CNN?
1979
1980
1970
2nd AI Winter
1989
Rand. Forest
Viola& Jones
2001
GANs (Good fellow)
2014
2017
Perceptron (Rosenblatt) learned weights
Back Propagation (Linnainmaa)
1st AI Winter
VISION
David Marr
CNN + Backprop (LeCunn)
Google Trends: Deep Learning
"Winter AI"
Deep Learning revolution
Transformer (Vaswani)

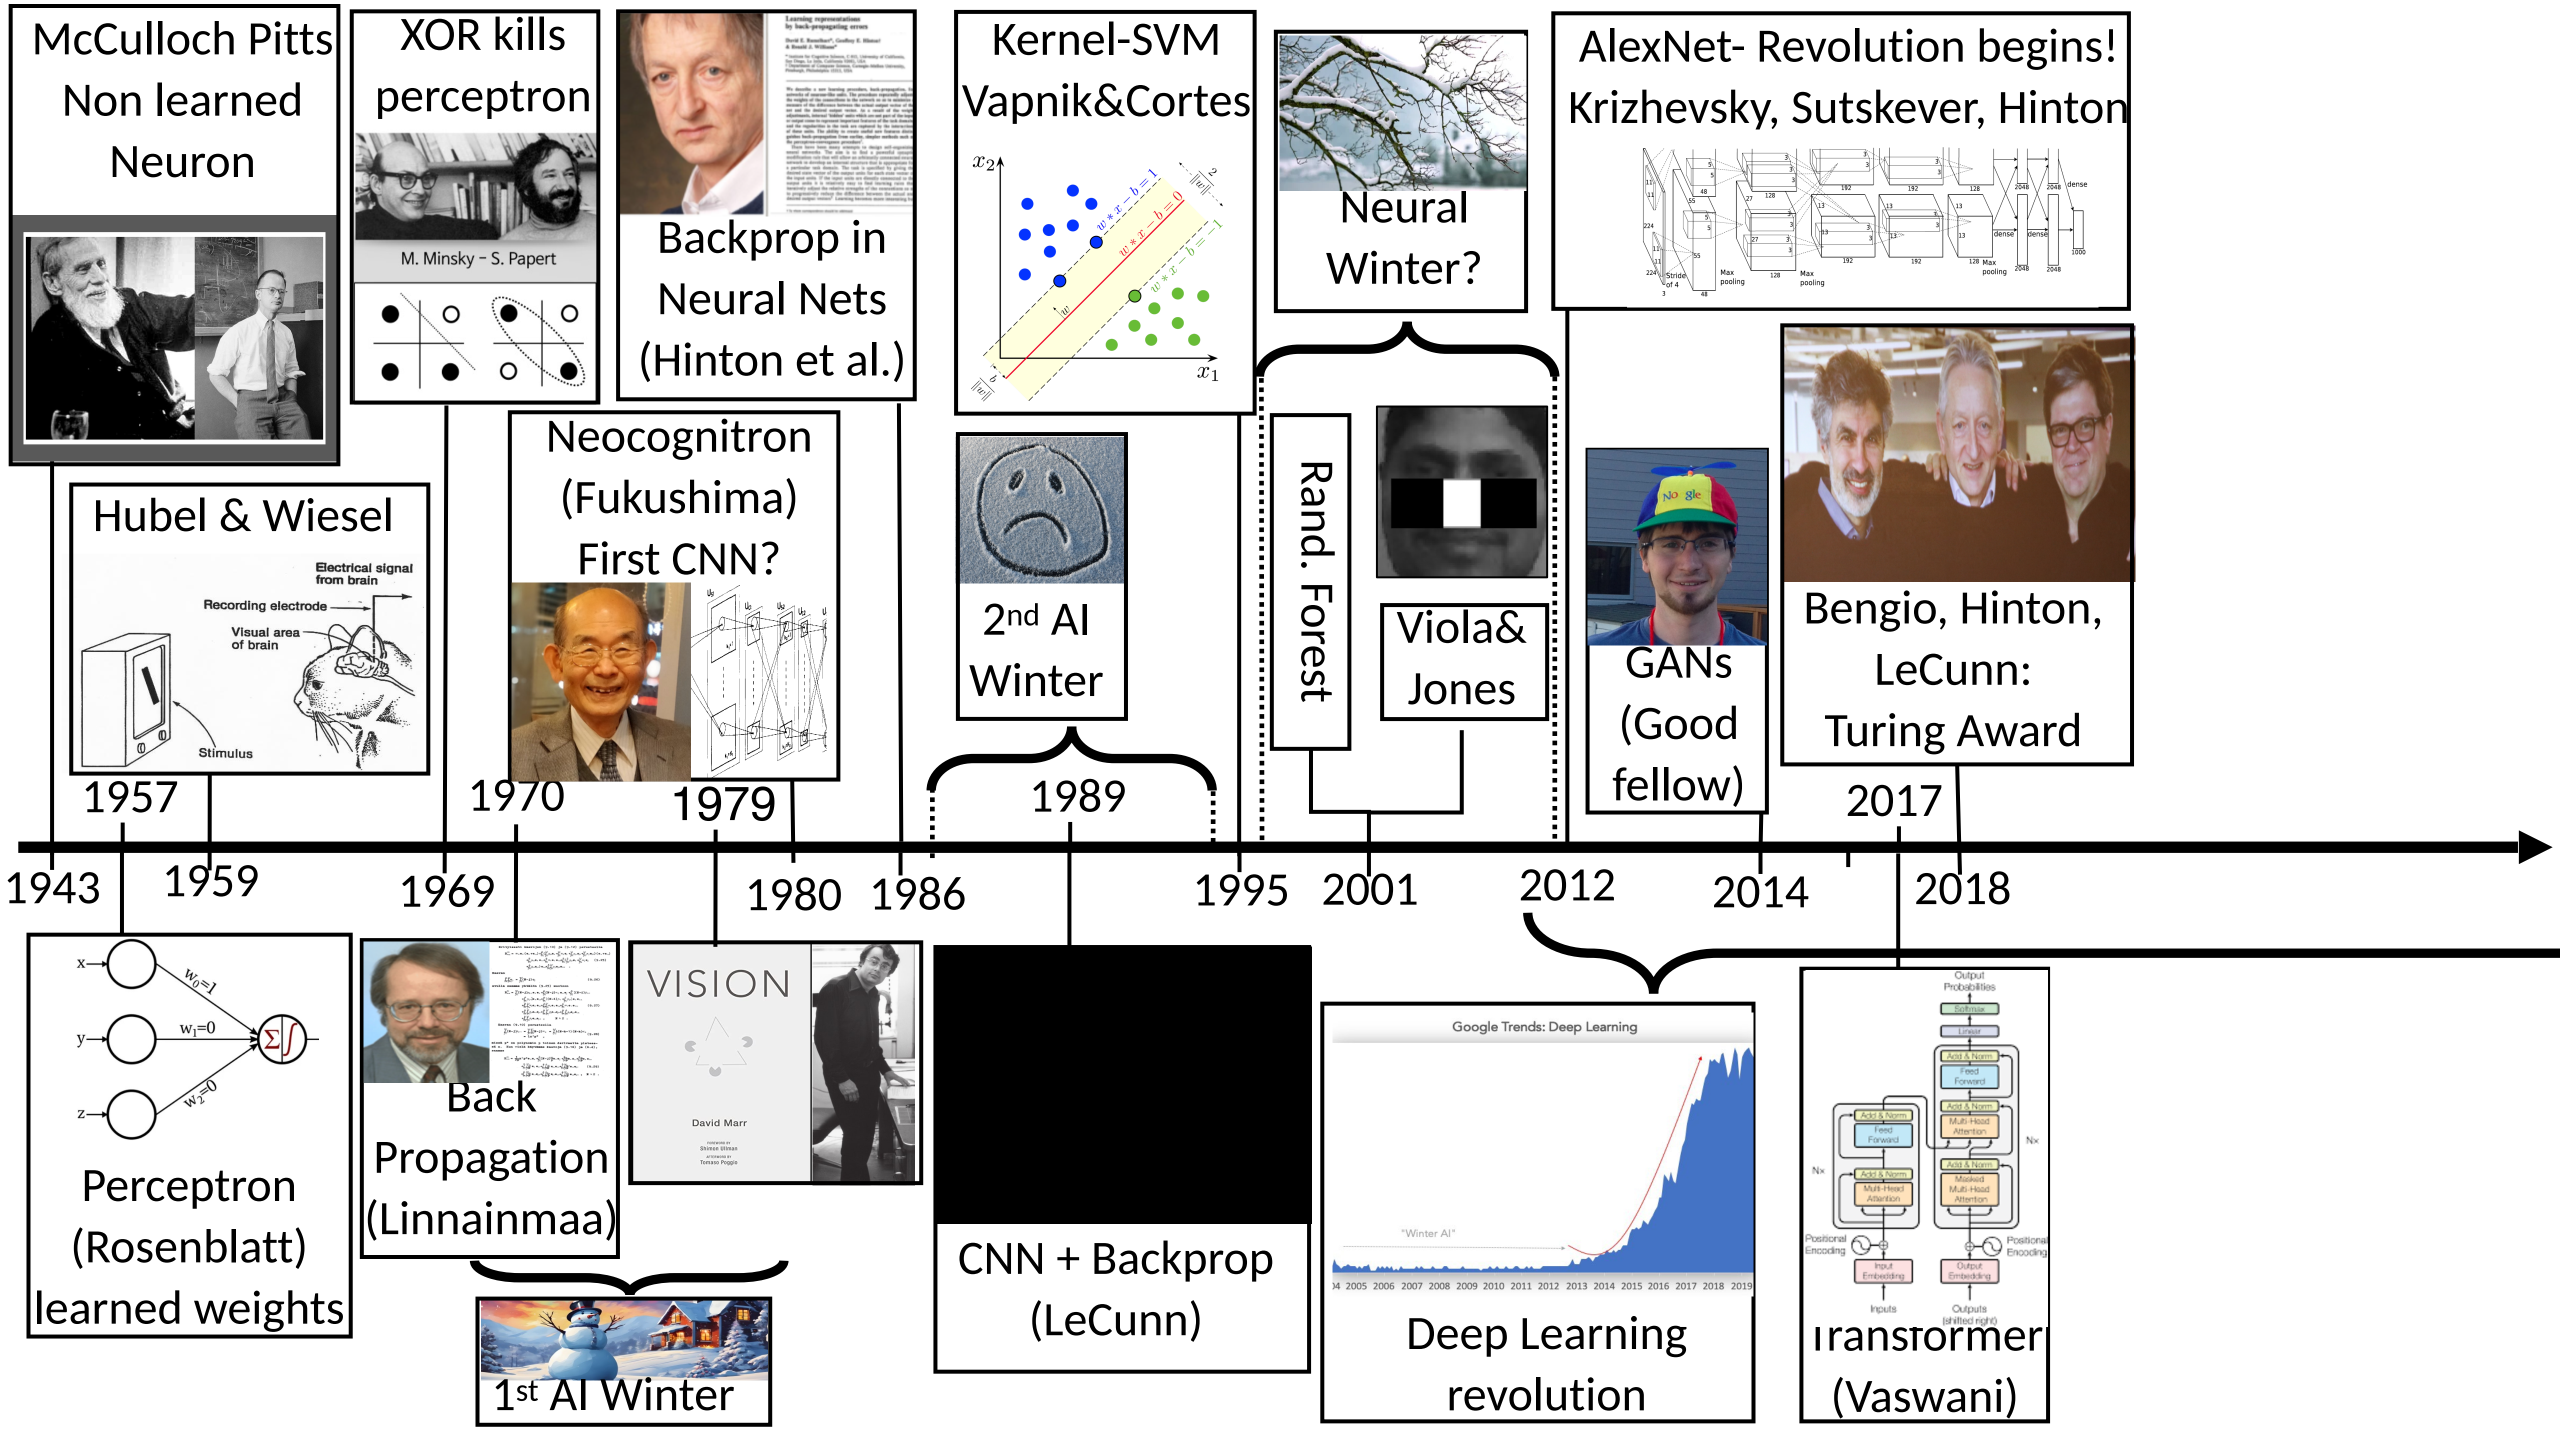
McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Backprop in
Neural Nets
(Hinton et al.)
Kernel-SVM
Vapnik&Cortes
Neural
Winter?
AlexNet- Revolution begins!
Krizhevsky, Sutskever, Hinton
Hubel & Wiesel
Electrical signal
from brain
Recording electrode
Visual area
of brain
Stimulus
Neocognitron
(Fukushima)
First CNN?
2nd AI
Winter
Rand. Forest
Viola&
Jones
GANs
(Good
fellow)
Bengio, Hinton,
LeCunn:
Turing Award
1957
1970
1979
1989
2017
1943
1959
1969
1980
1986
1995
2001
2012
2014
2018
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
1st AI Winter
VISION
David Marr
CNN + Backprop
(LeCunn)
Google Trends: Deep Learning
"Winter AI"
Deep Learning
revolution
Transformer
(Vaswani)

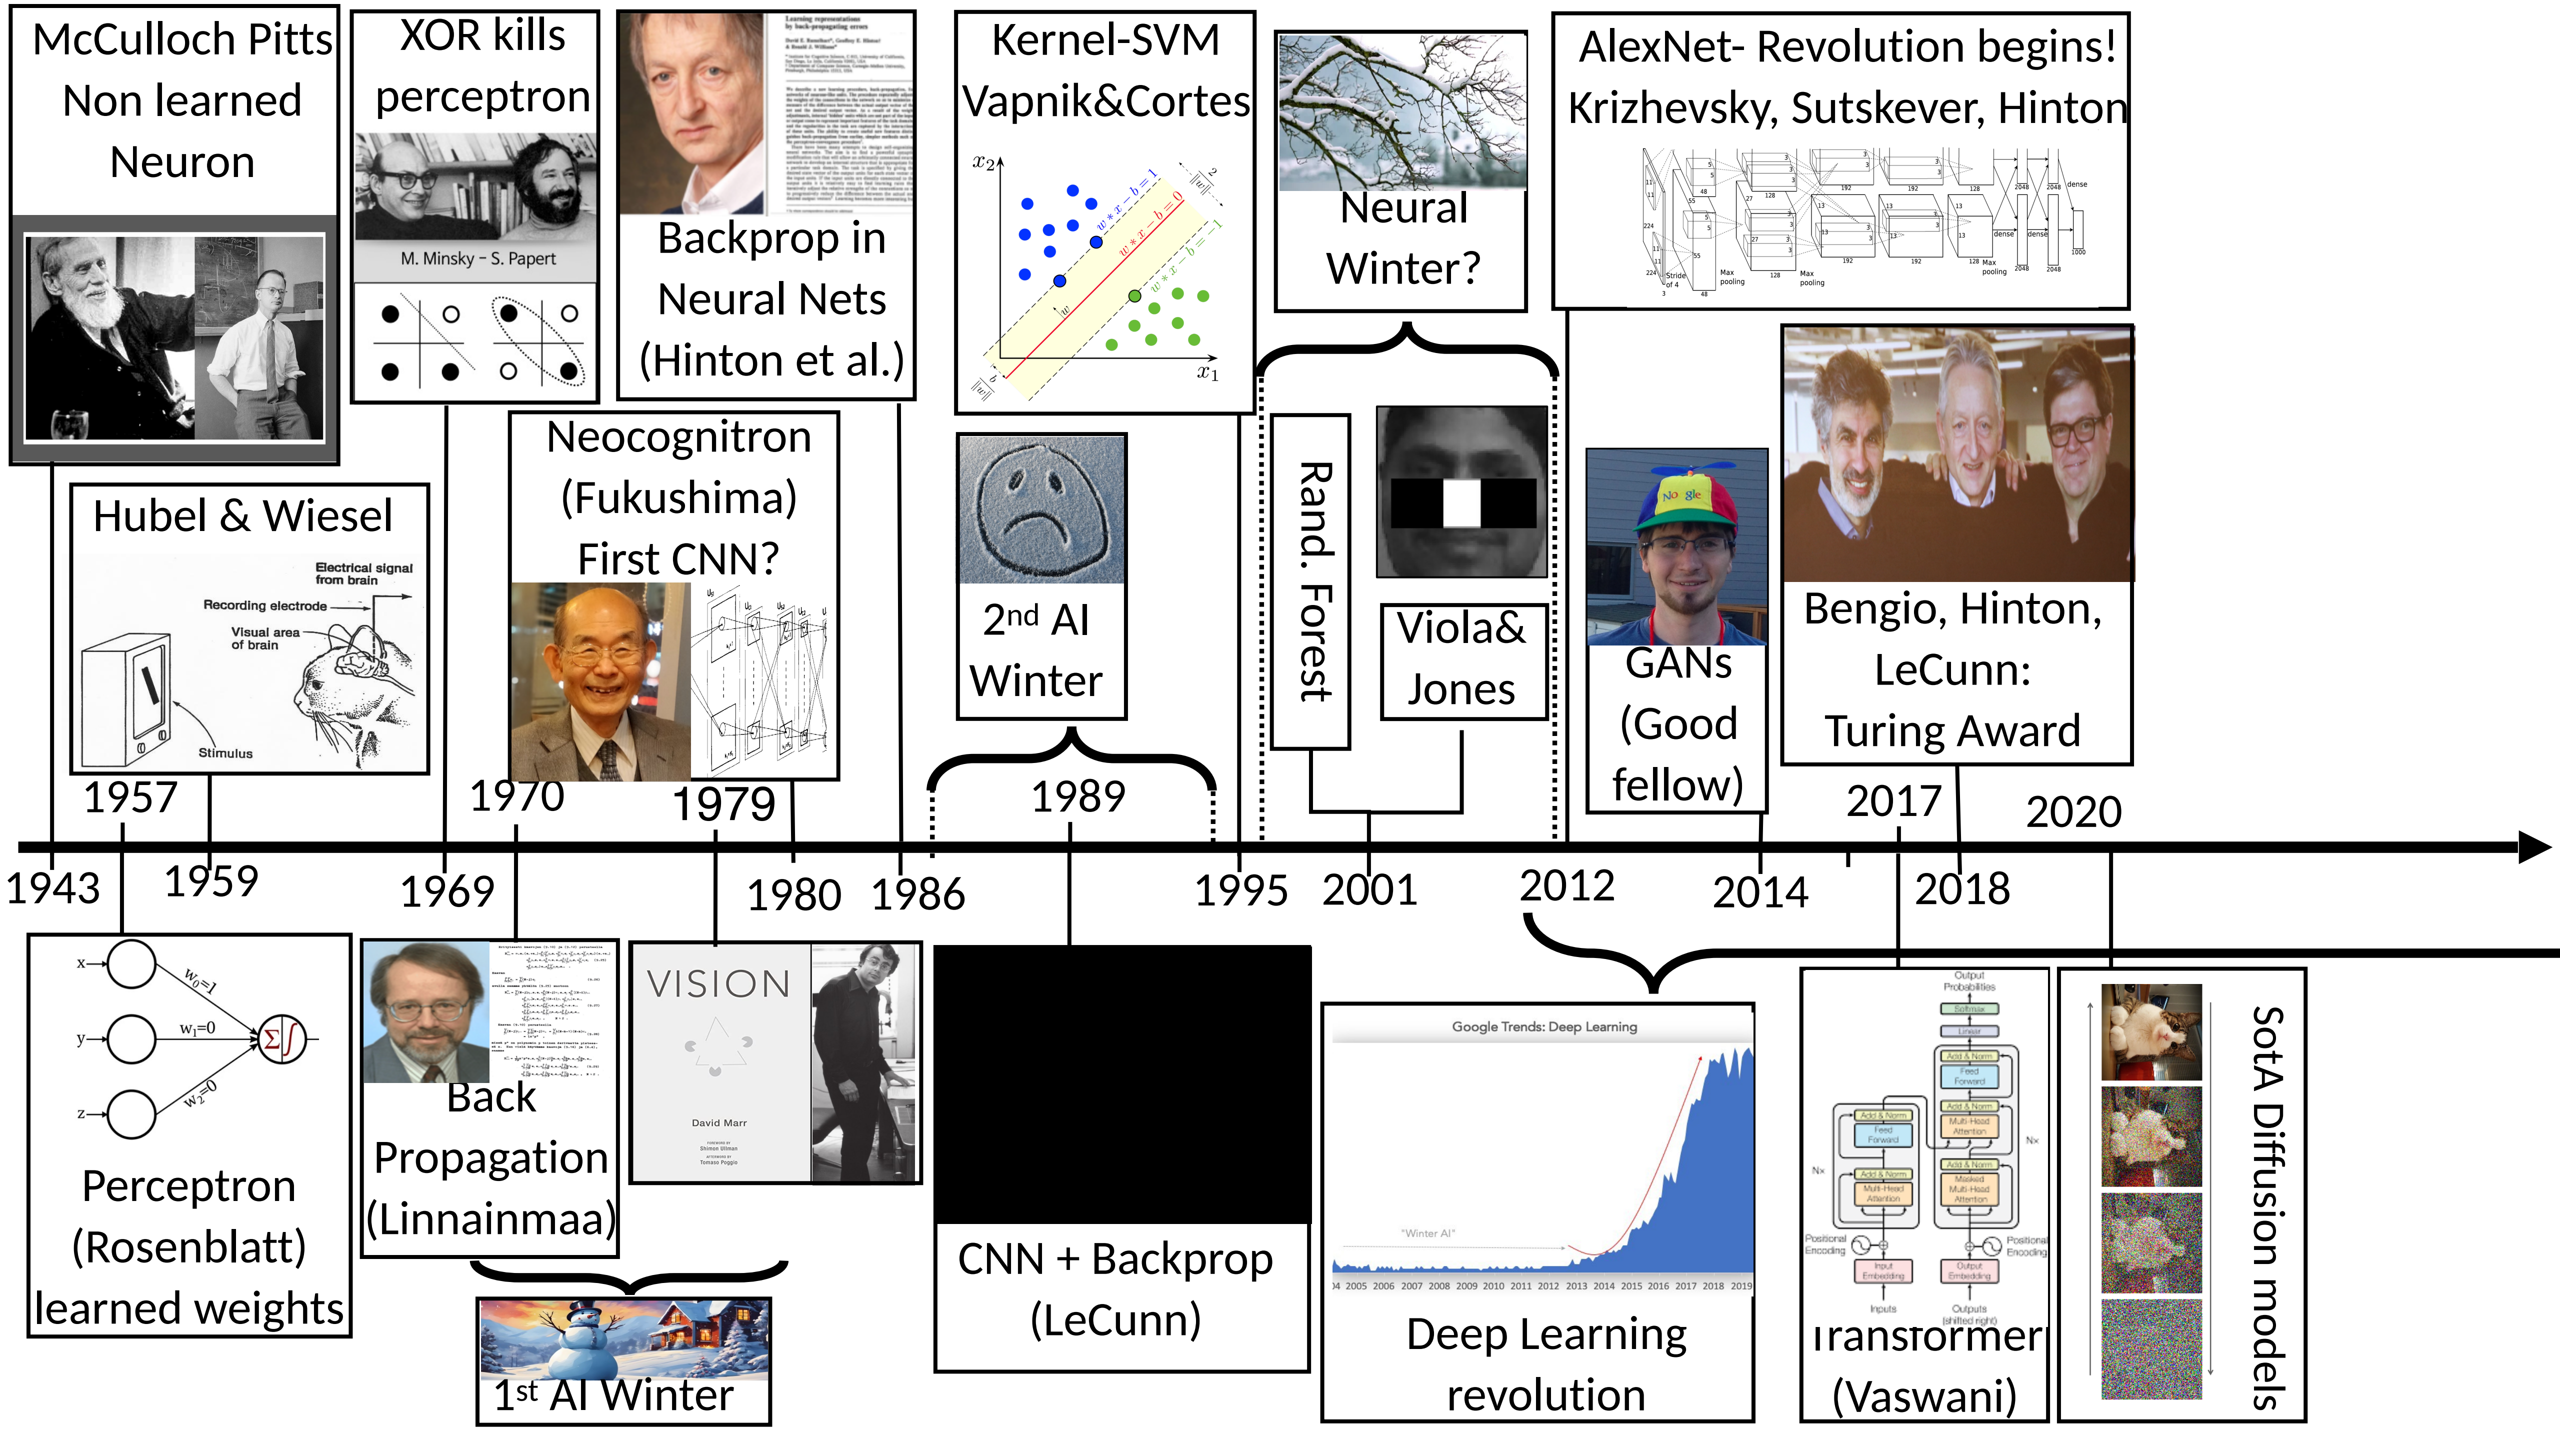
McCulloch Pitts
Non learned
Neuron
XOR kills
perceptron
M. Minsky – S. Papert
Backprop in
Neural Nets
(Hinton et al.)
Kernel-SVM
Vapnik&Cortes
Neural
Winter?
AlexNet- Revolution begins!
Krizhevsky, Sutskever, Hinton
Hubel & Wiesel
Electrical signal
from brain
Recording electrode
Visual area
of brain
Stimulus
Neocognitron
(Fukushima)
First CNN?
2nd AI
Winter
Rand. Forest
Viola&
Jones
GANs
(Good
fellow)
Bengio, Hinton,
LeCunn:
Turing Award
1957
1970
1979
1989
2017
2020
1943
1959
1969
1980
1986
1995
2001
2012
2014
2018
Perceptron
(Rosenblatt)
learned weights
Back
Propagation
(Linnainmaa)
1st AI Winter
VISION
David Marr
CNN + Backprop
(LeCunn)
Google Trends: Deep Learning
"Winter AI"
Deep Learning
revolution
Output
Probabilities
Softmax
Linear
Add & Norm
Feed
Forward
Multi-Head
Attention
Masked
Multi-Head
Attention
Positional
Encoding
Input
Embedding
Output
Embedding
Inputs
Outputs
(shifted right)
Transformer
(Vaswani)
SotA Diffusion models

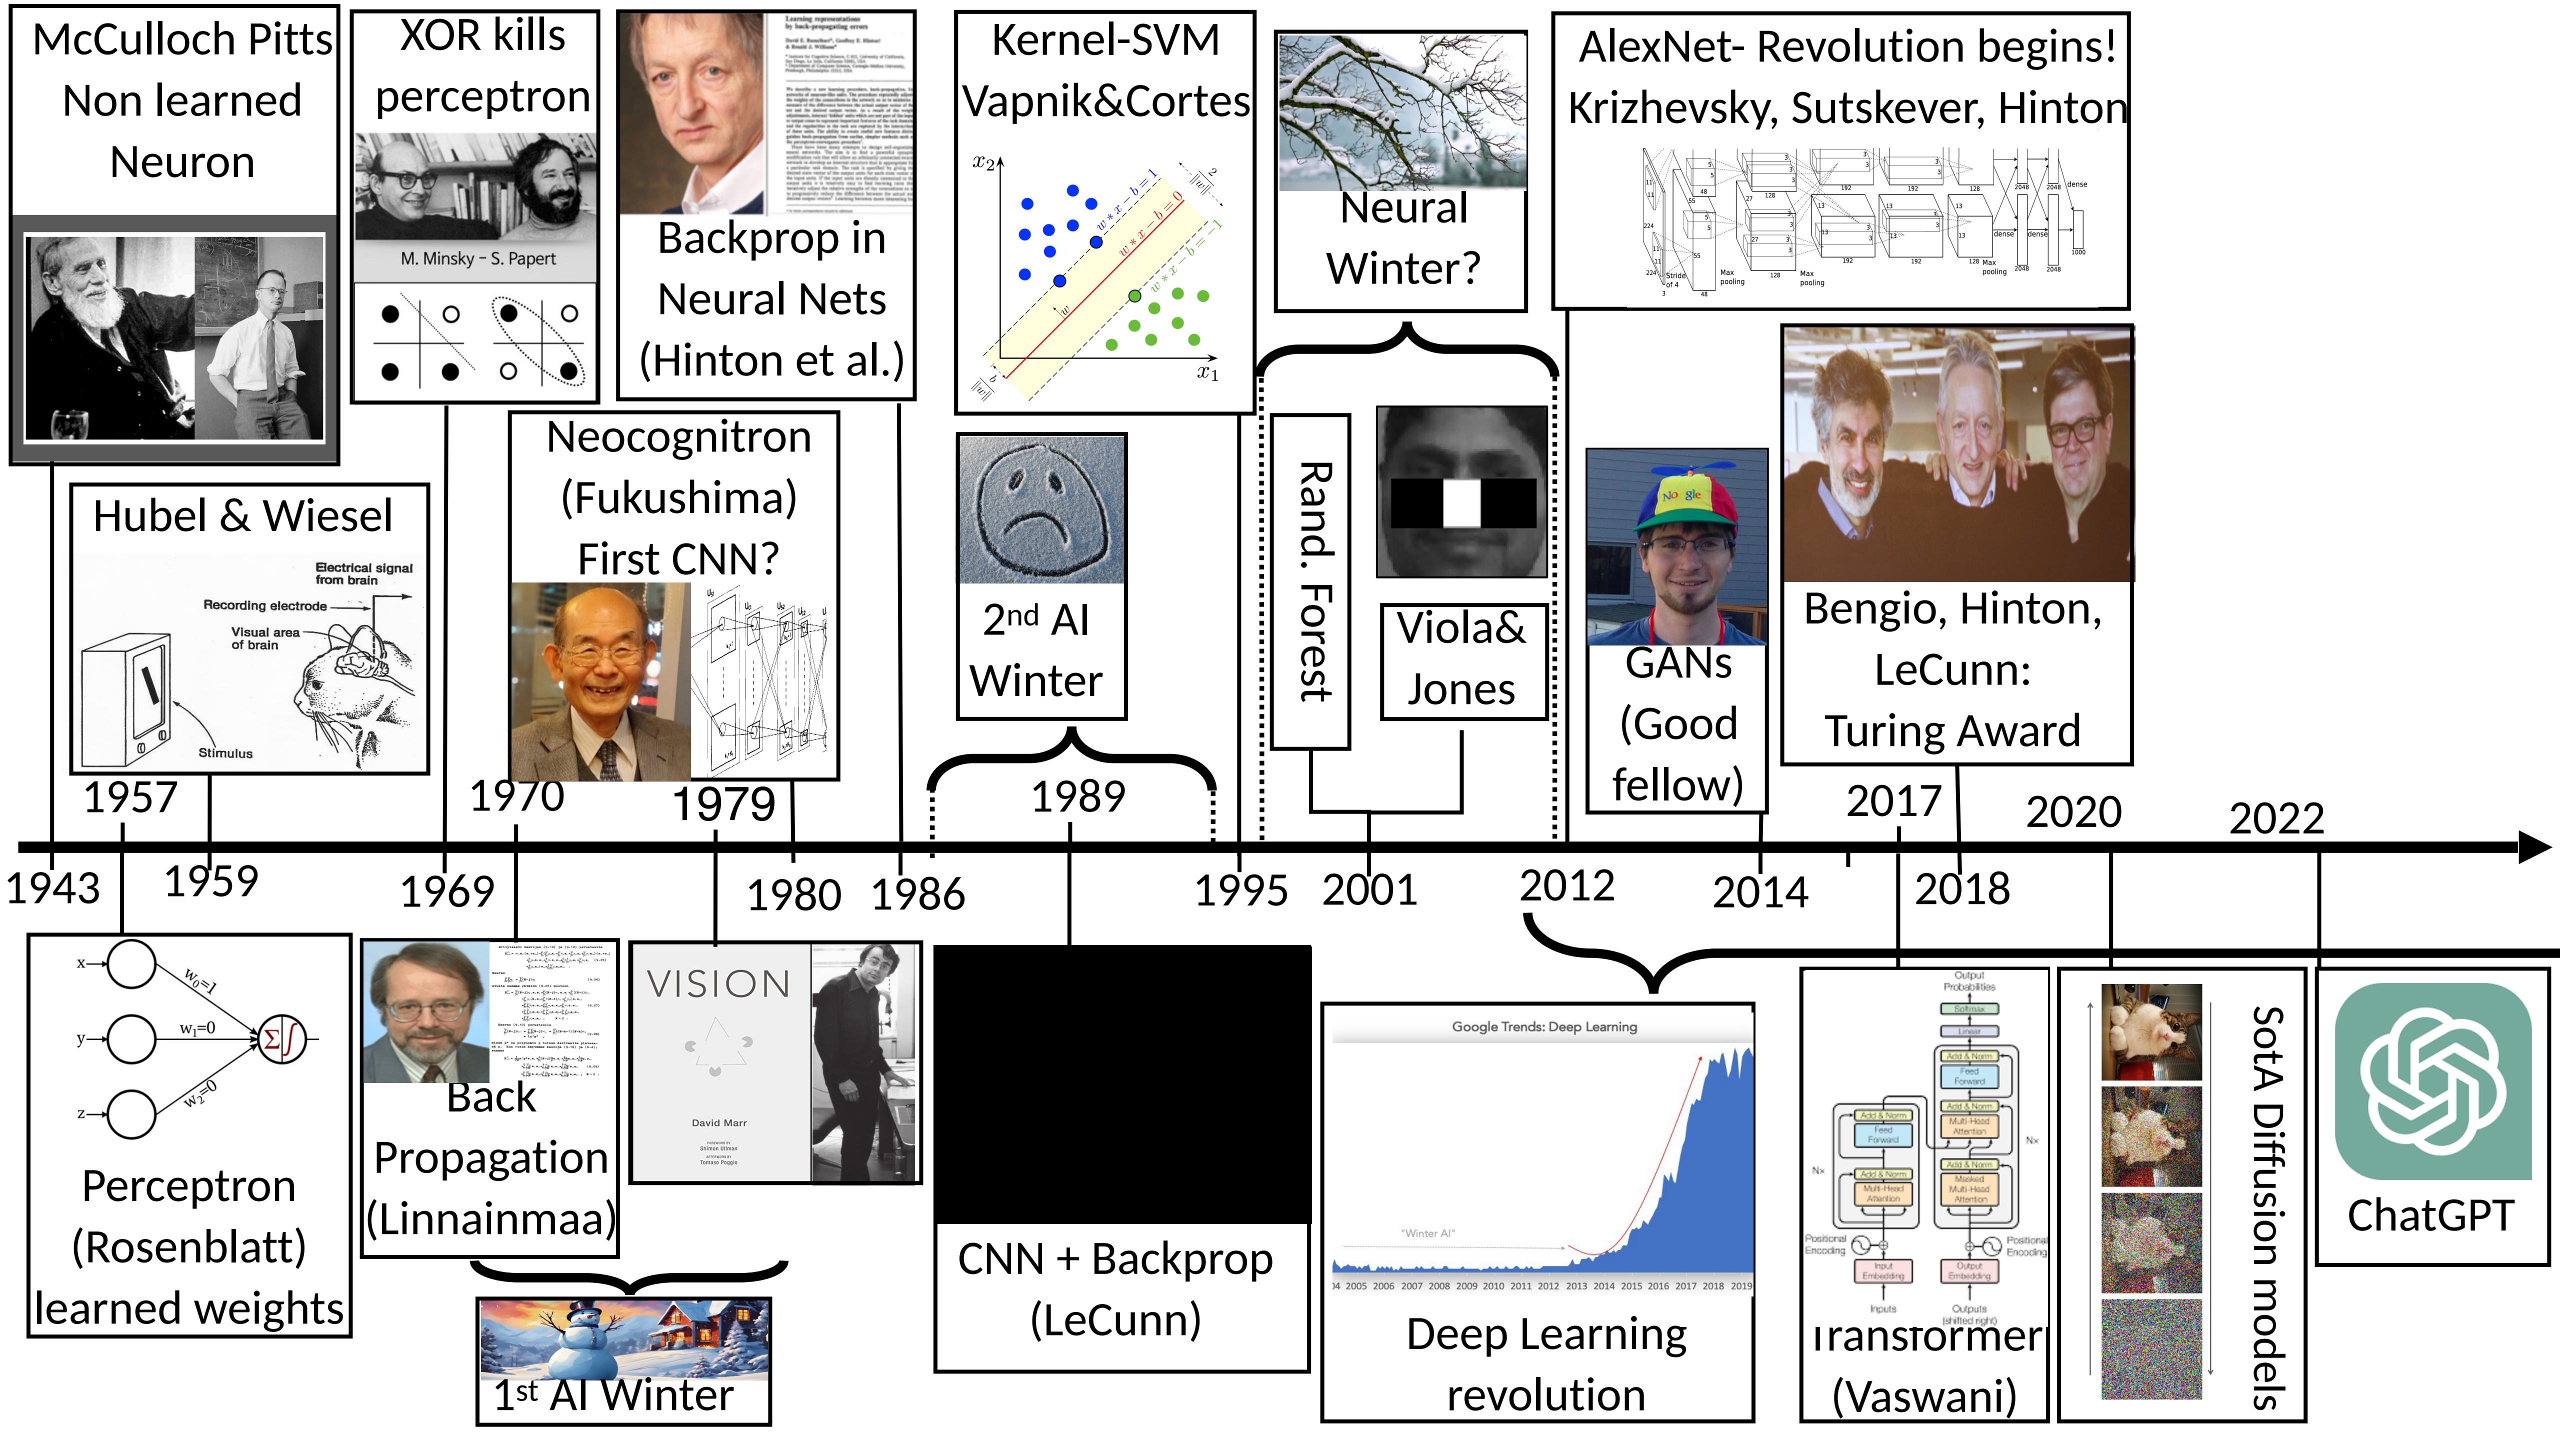

McCulloch Pitts Non learned Neuron
XOR kills perceptron
M. Minsky – S. Papert
Backprop in Neural Nets (Hinton et al.)
Kernel-SVM Vapnik&Cortes
Neural Winter?
AlexNet- Revolution begins! Krizhevsky, Sutskever, Hinton
Hubel & Wiesel
Electrical signal from brain
Recording electrode
Visual area of brain
Stimulus
Neocognitron (Fukushima) First CNN?
2nd AI Winter
Rand. Forest
Viola& Jones
GANs (Good fellow)
Bengio, Hinton, LeCunn: Turing Award
1957
1970
1979
1989
2017
2020
2022
1943
1959
1969
1980
1986
1995
2001
2012
2014
2018
Perceptron (Rosenblatt) learned weights
Back Propagation (Linnainmaa)
1st AI Winter
VISION
David Marr
CNN + Backprop (LeCunn)
Google Trends: Deep Learning
"Winter AI"
Deep Learning revolution
Transformer (Vaswani)
SotA Diffusion models
ChatGPT

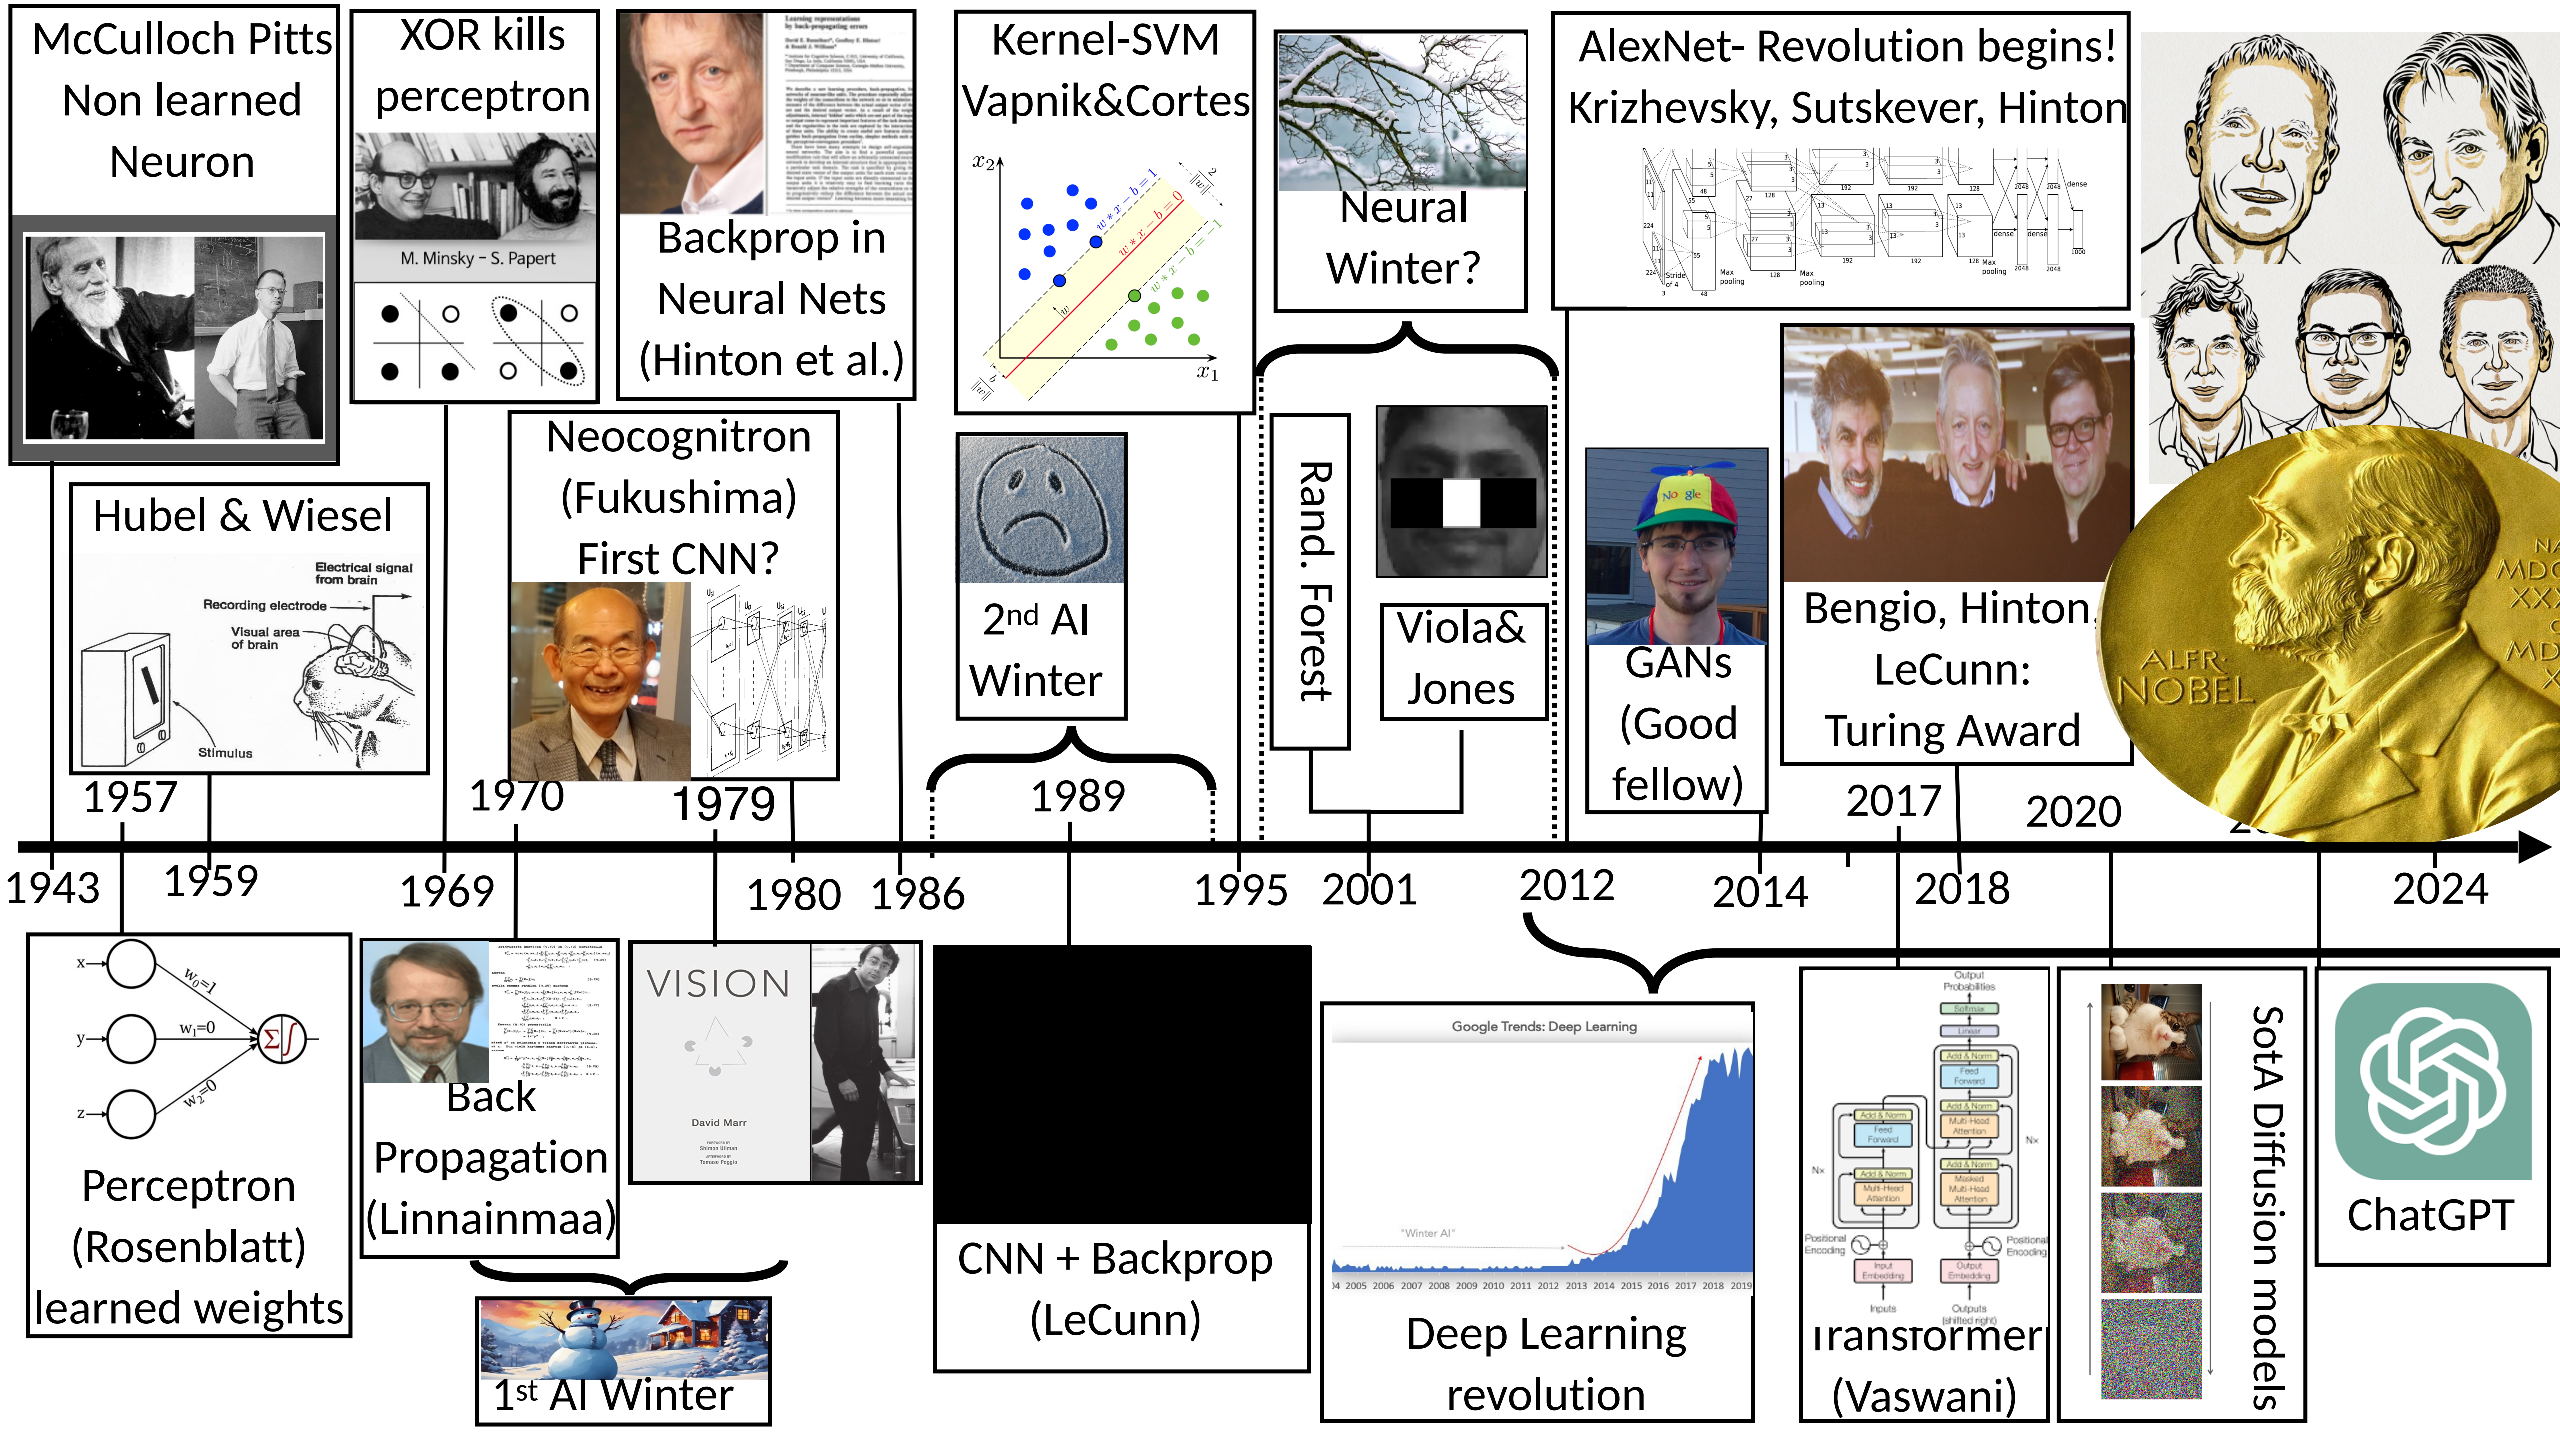

McCulloch Pitts Non learned Neuron
XOR kills perceptron
M. Minsky – S. Papert
Backprop in Neural Nets (Hinton et al.)
Kernel-SVM Vapnik&Cortes
Neural Winter?
AlexNet- Revolution begins! Krizhevsky, Sutskever, Hinton
Hubel & Wiesel
Electrical signal from brain
Recording electrode
Visual area of brain
Stimulus
Neocognitron (Fukushima) First CNN?
2nd AI Winter
Rand. Forest
Viola& Jones
GANs (Good fellow)
Bengio, Hinton, LeCunn: Turing Award
1957
1970
1979
1989
2017
2020
1943
1959
1969
1980
1986
1995
2001
2012
2014
2018
2024
Perceptron (Rosenblatt) learned weights
Back Propagation (Linnainmaa)
1st AI Winter
VISION
David Marr
CNN + Backprop (LeCunn)
Google Trends: Deep Learning
"Winter AI"
Deep Learning revolution
Transformer (Vaswani)
SotA Diffusion models
ChatGPT

From light to insight

From light to insight

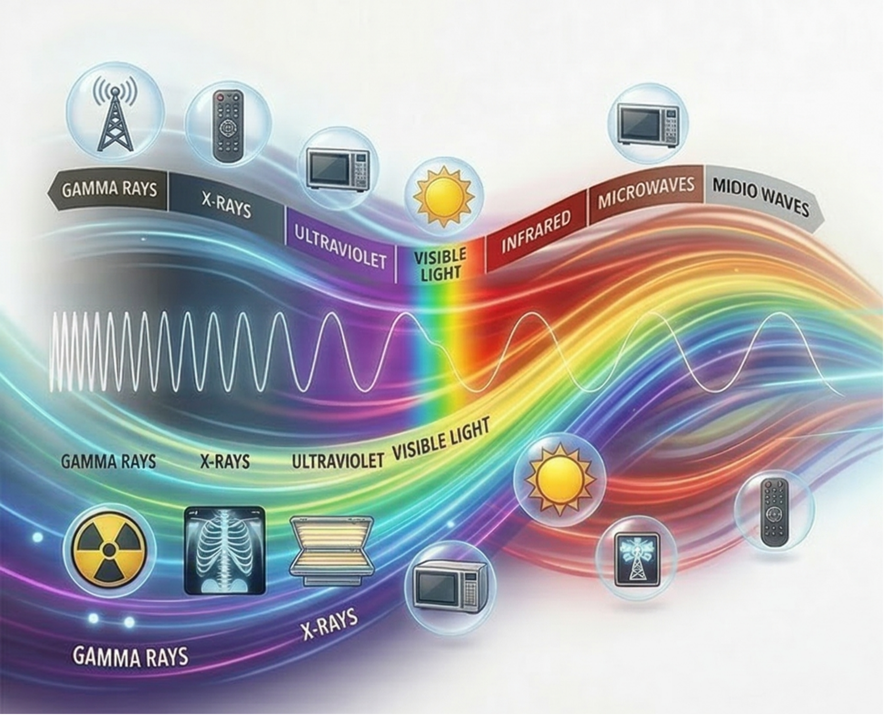

Physics of light

From light to insight

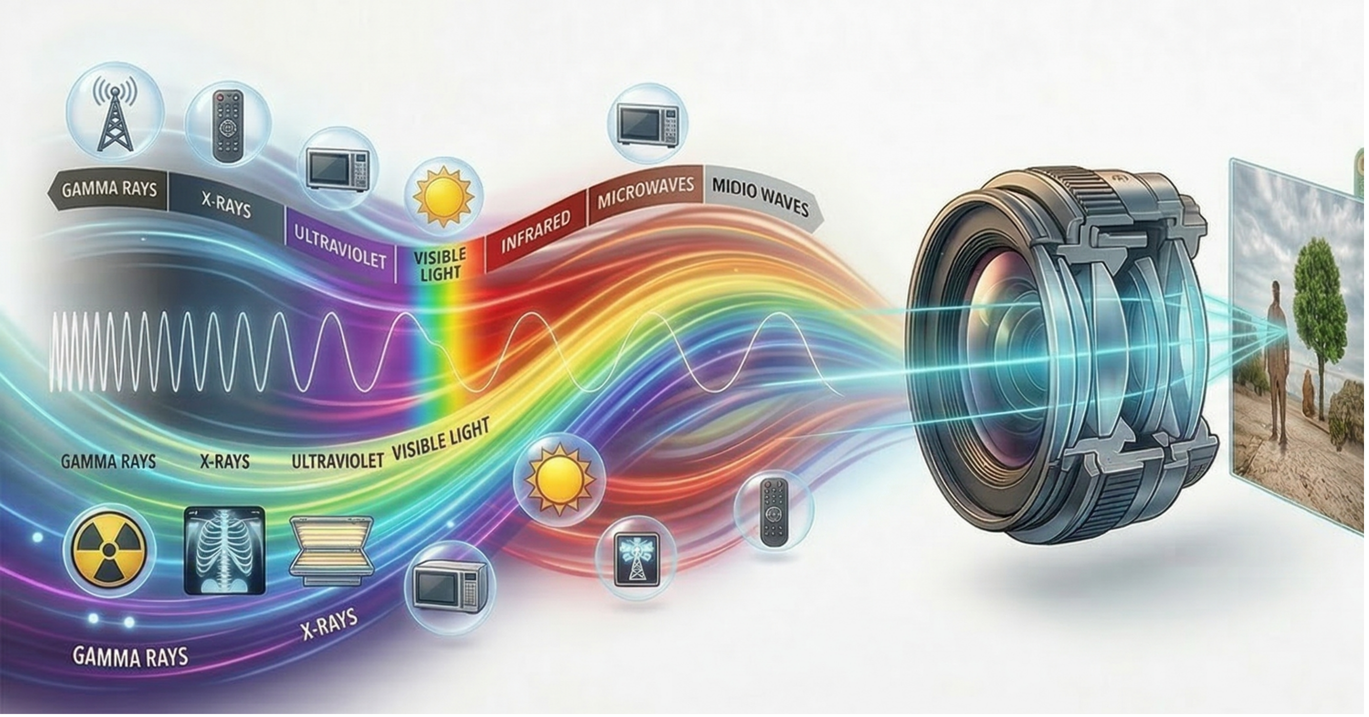

Physics of light

Optics of image formation

From light to insight

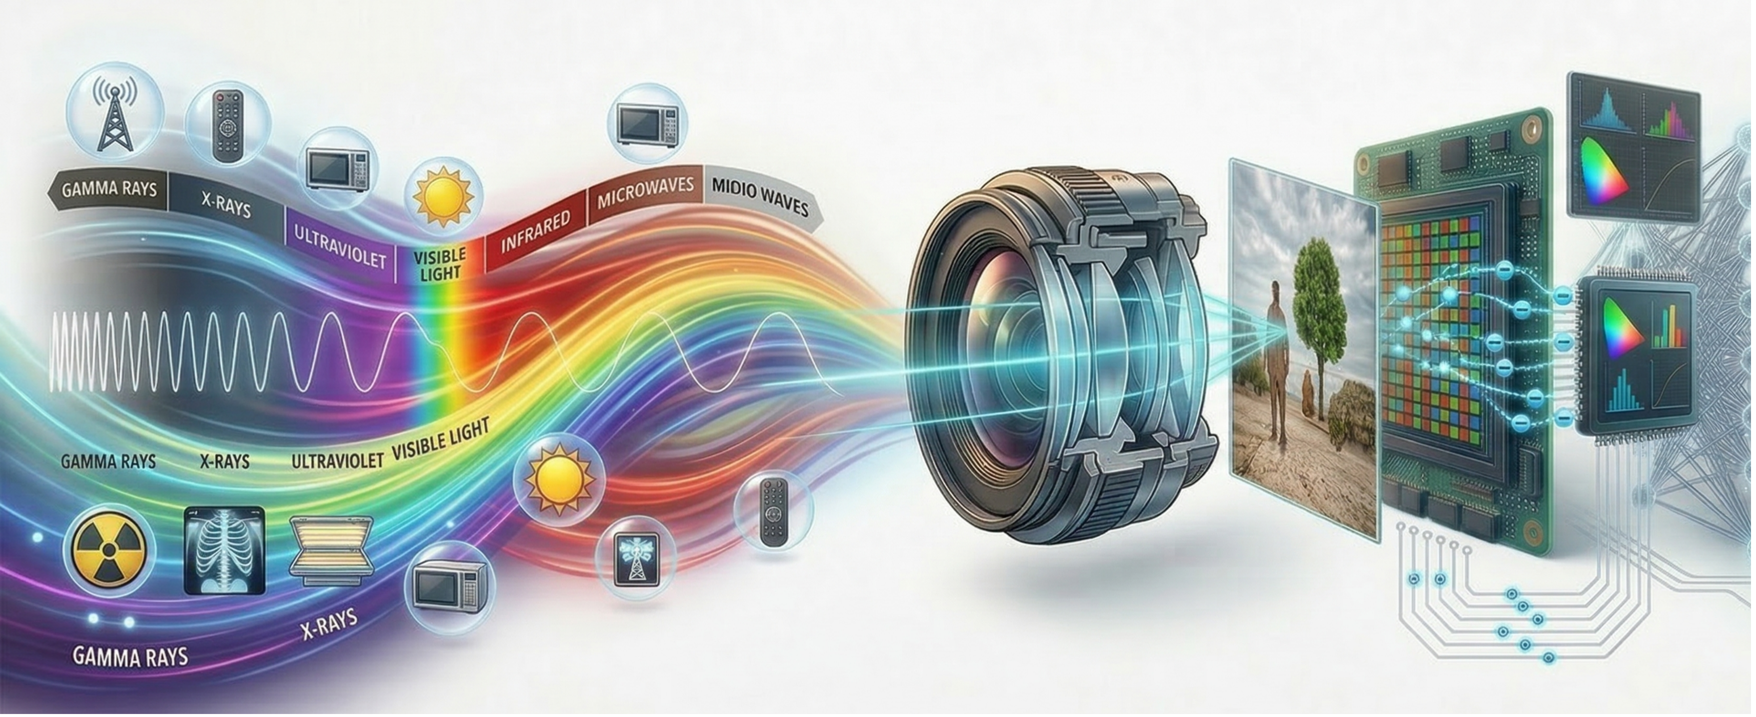

Physics of light

Optics of image formation

Digital image pipeline

From light to insight

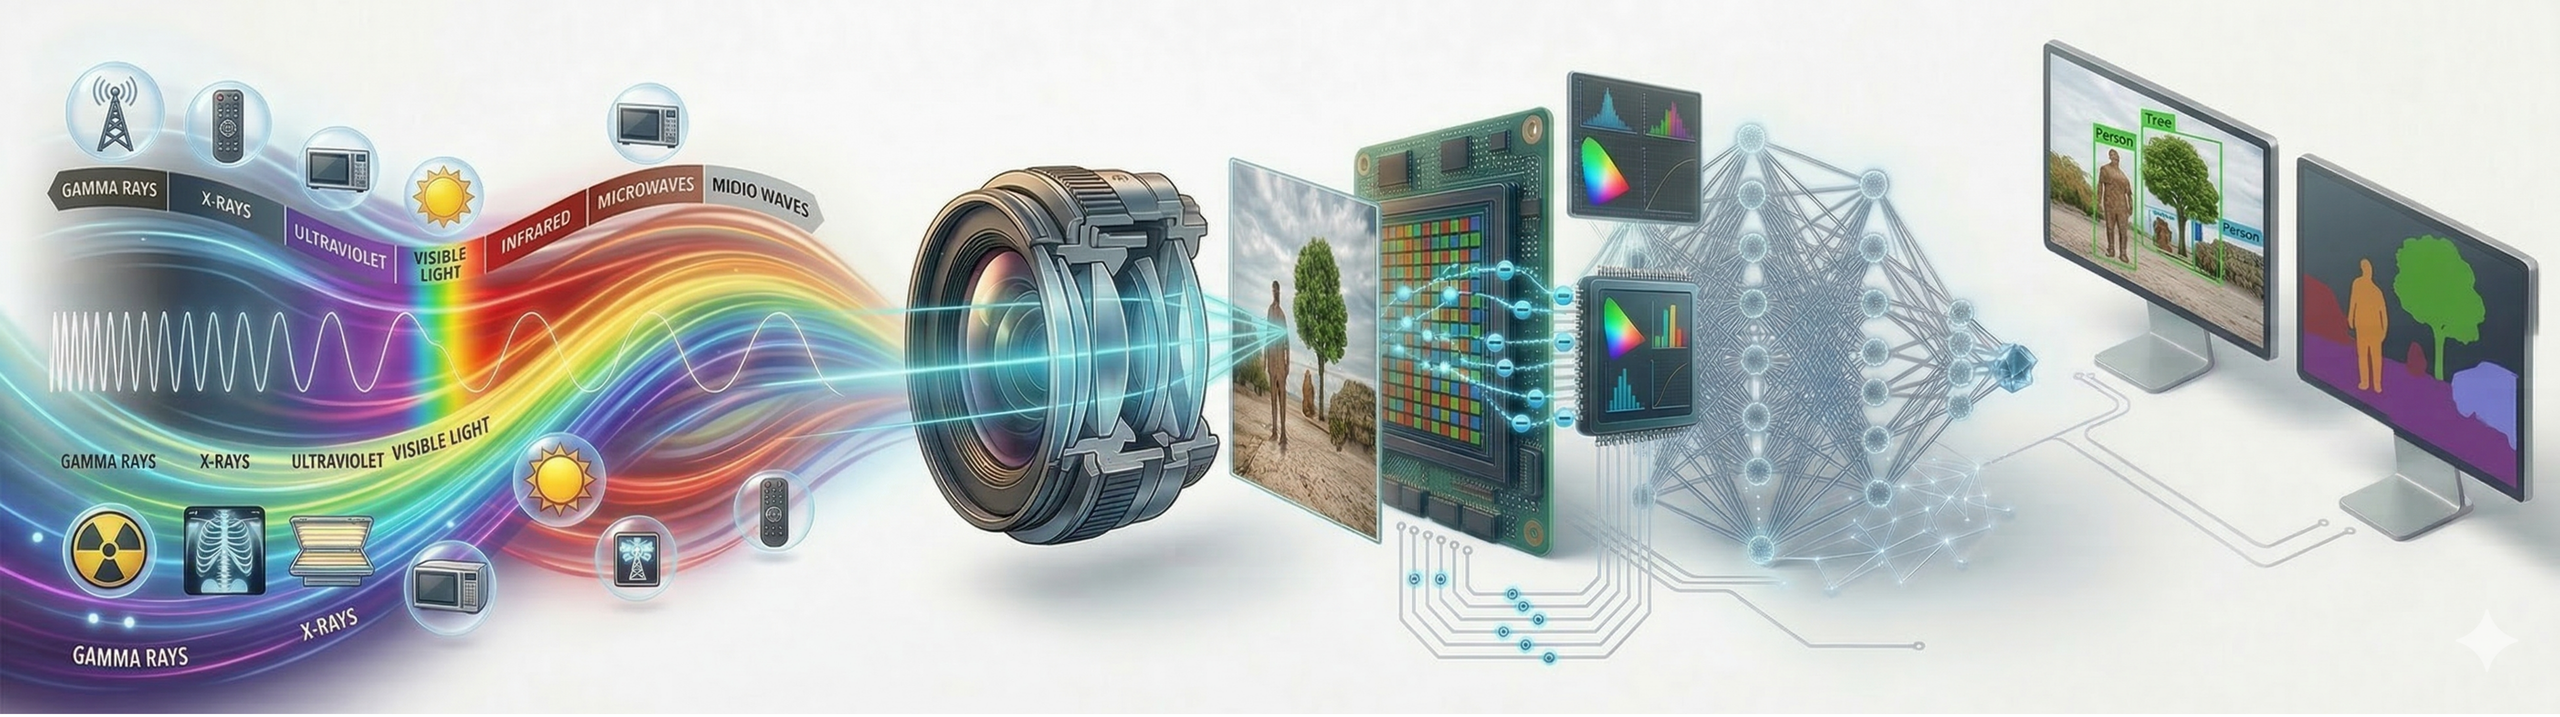

Physics of light

Optics of image formation

Digital image pipeline

Algorithms of computer vision

From light to insight

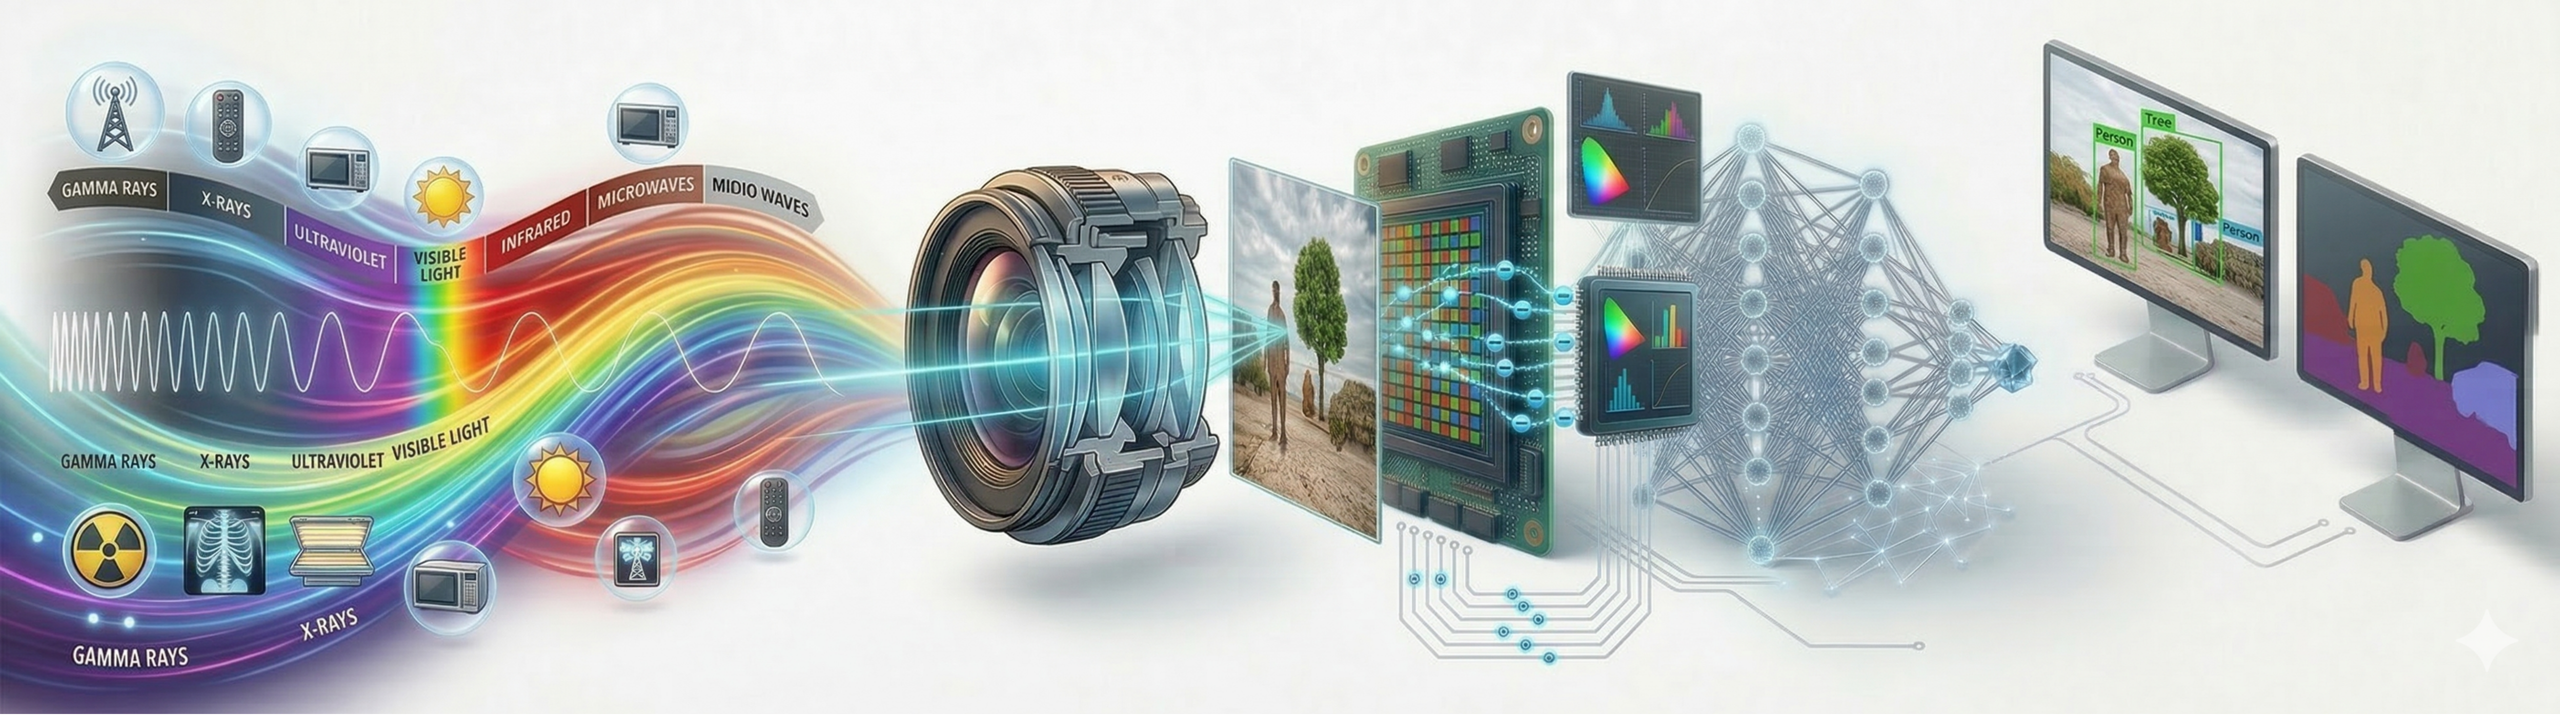

Physics of light

Optics of image formation

Digital image pipeline

Algorithms of computer vision

Today

From light to insight

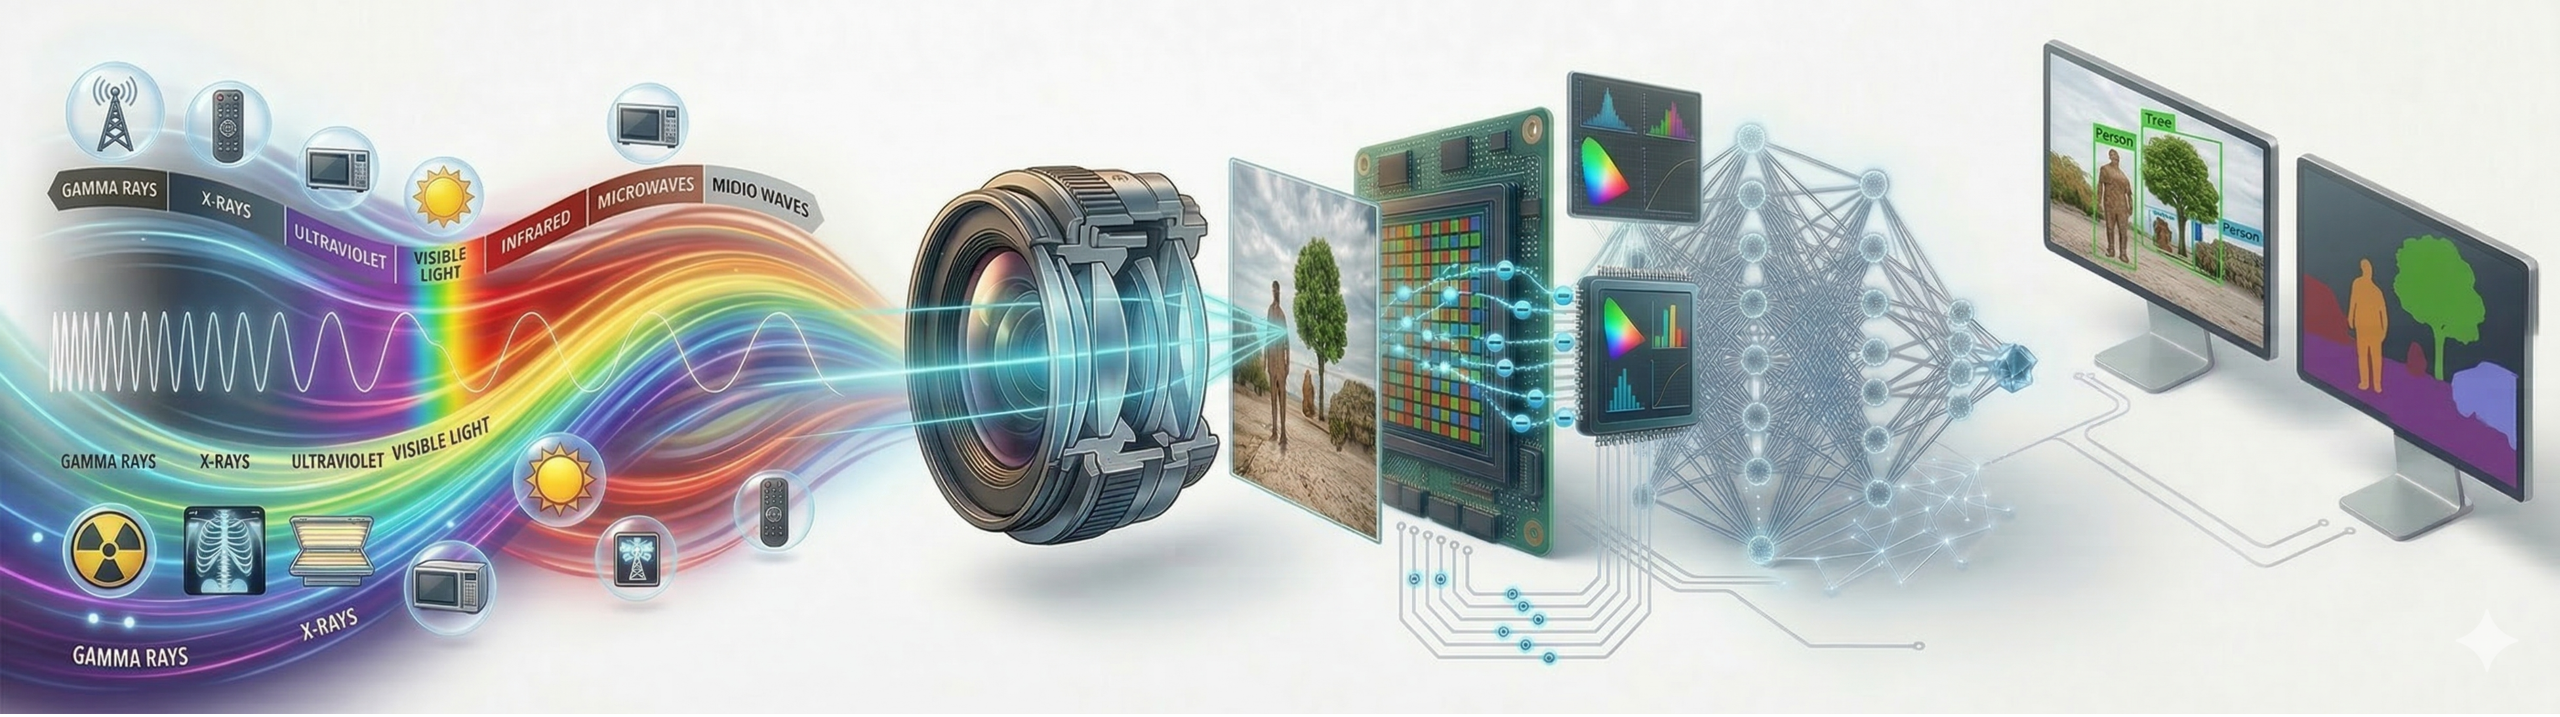

Physics of light | **Optics of image formation** | **Digital image pipeline** | **Algorithms of computer vision**

Today

Rest of the

Human optical system

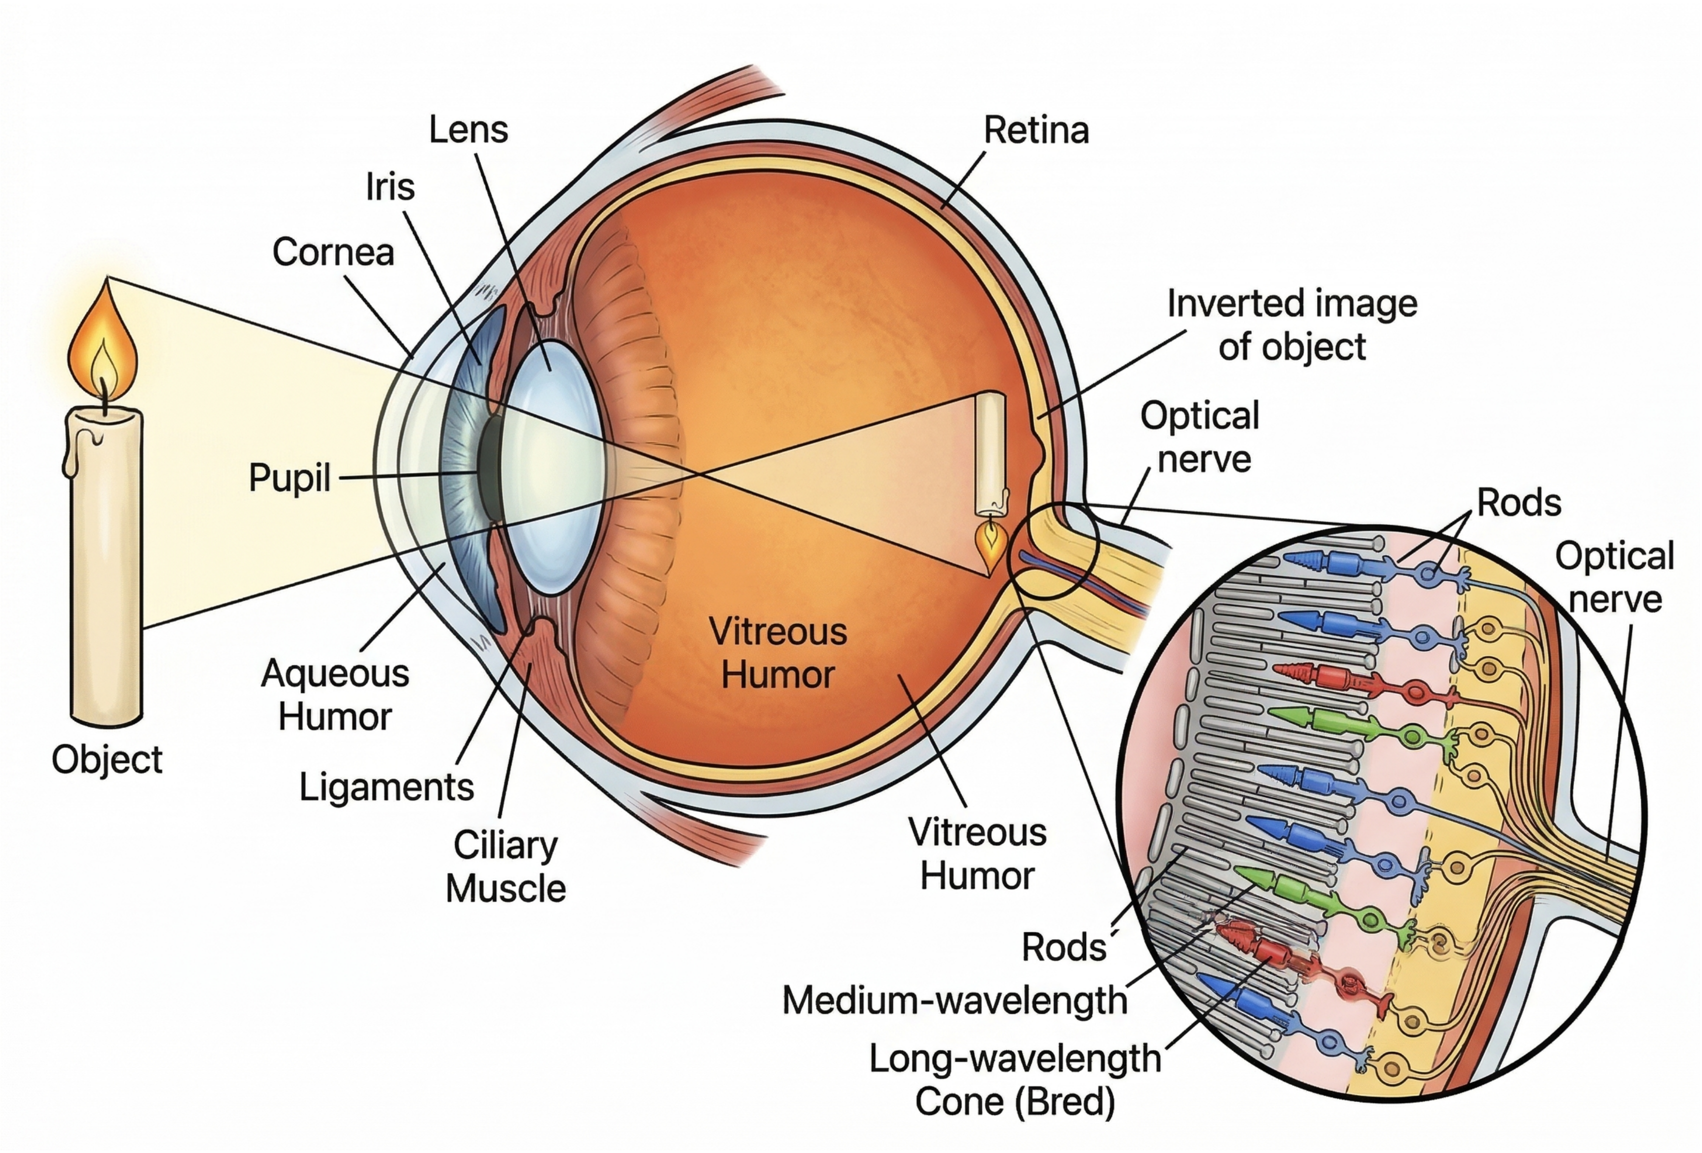

What is Light?

What is Light?

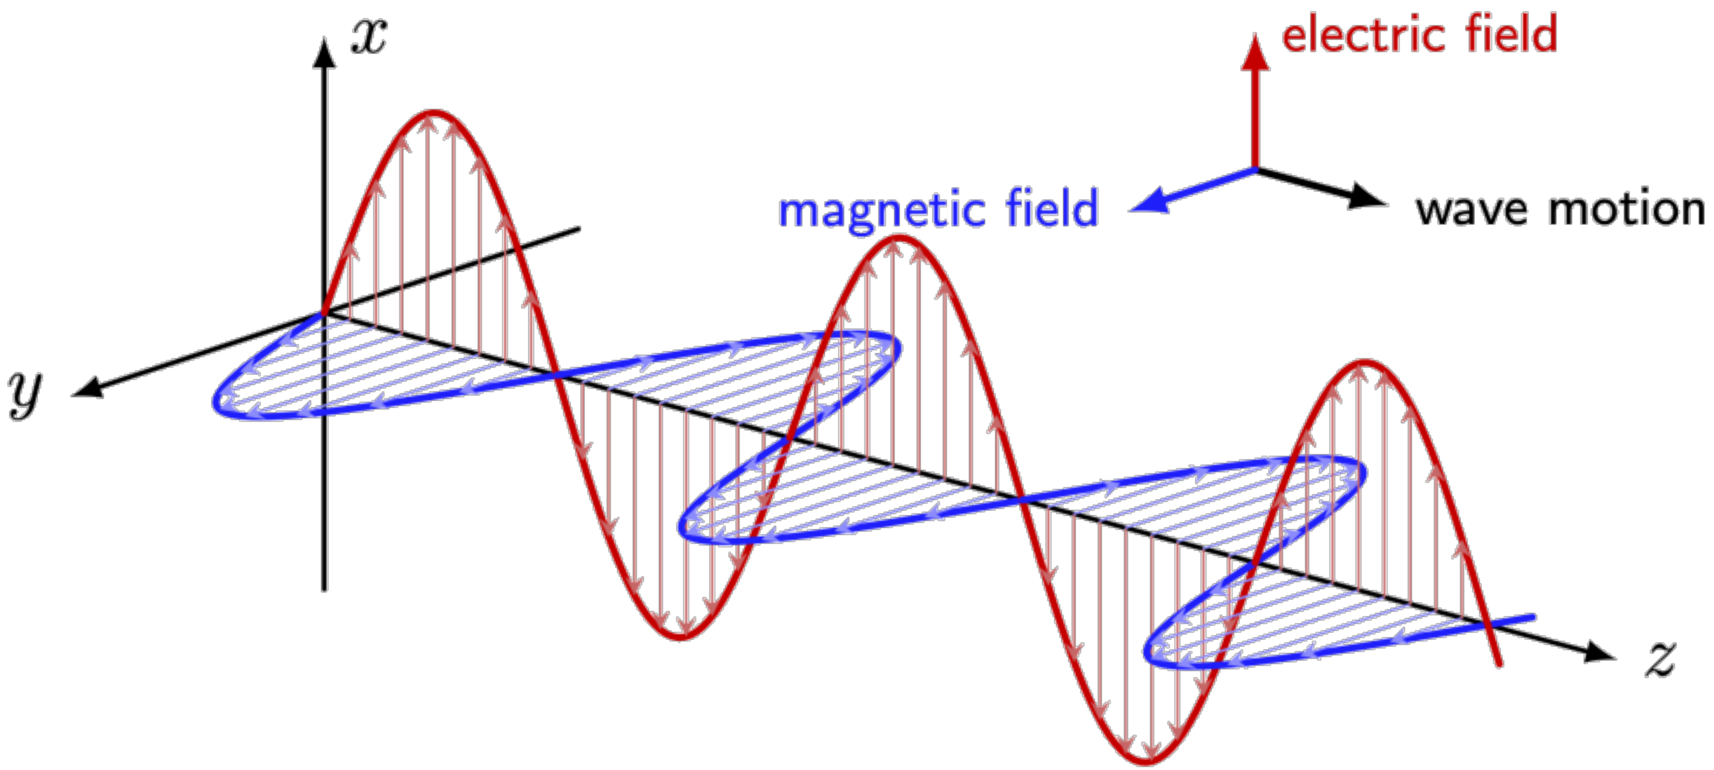

What is Light?

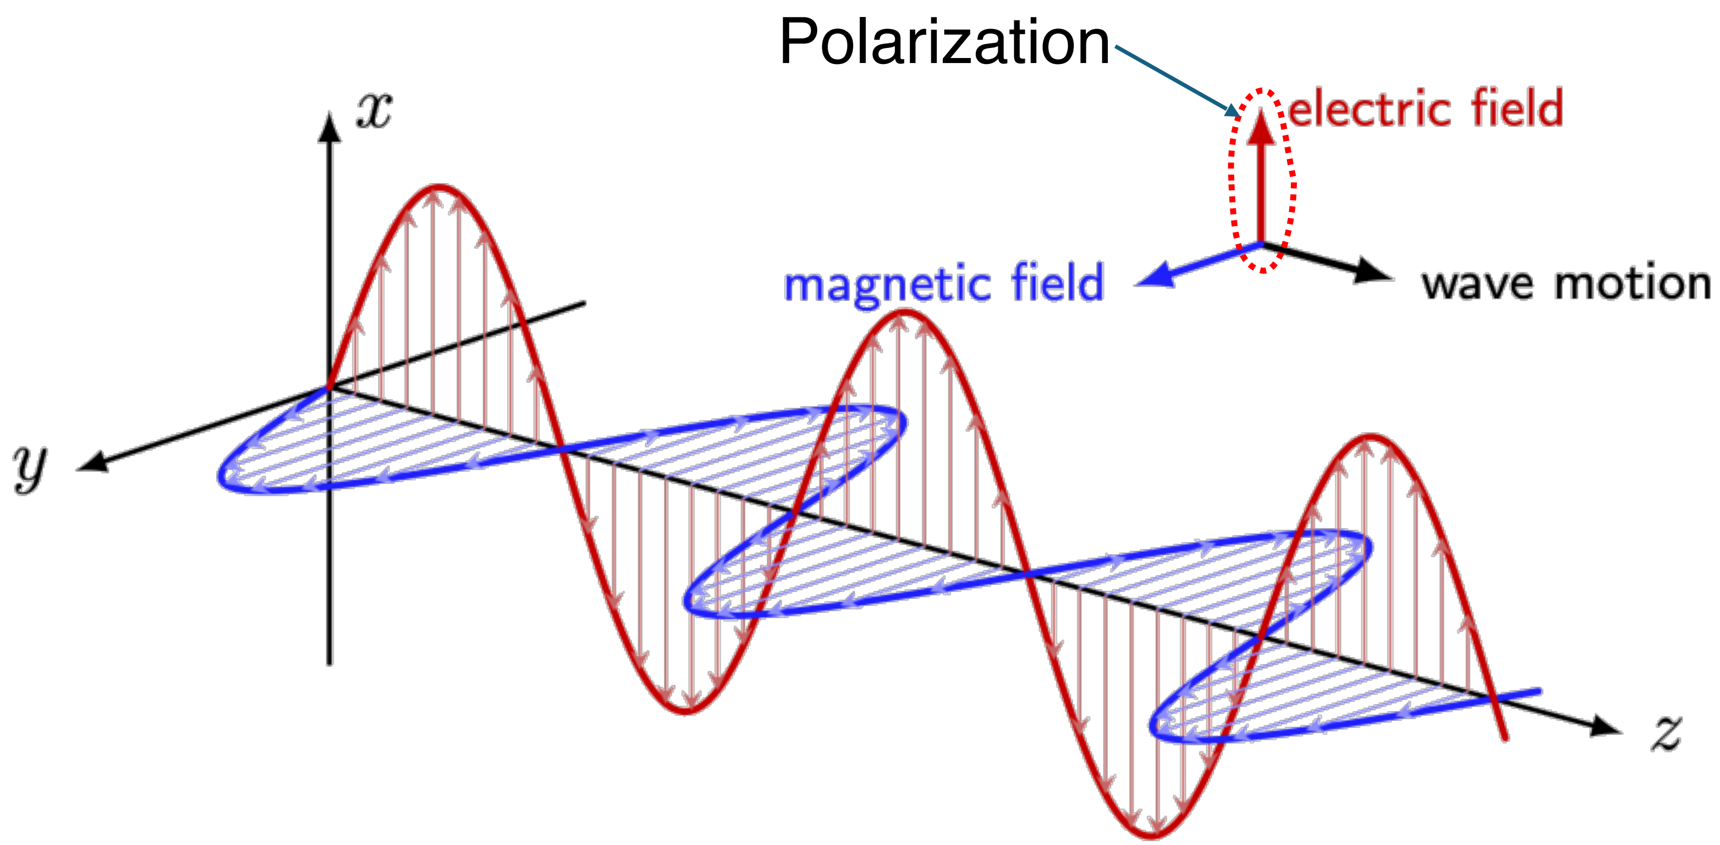

What is Light?

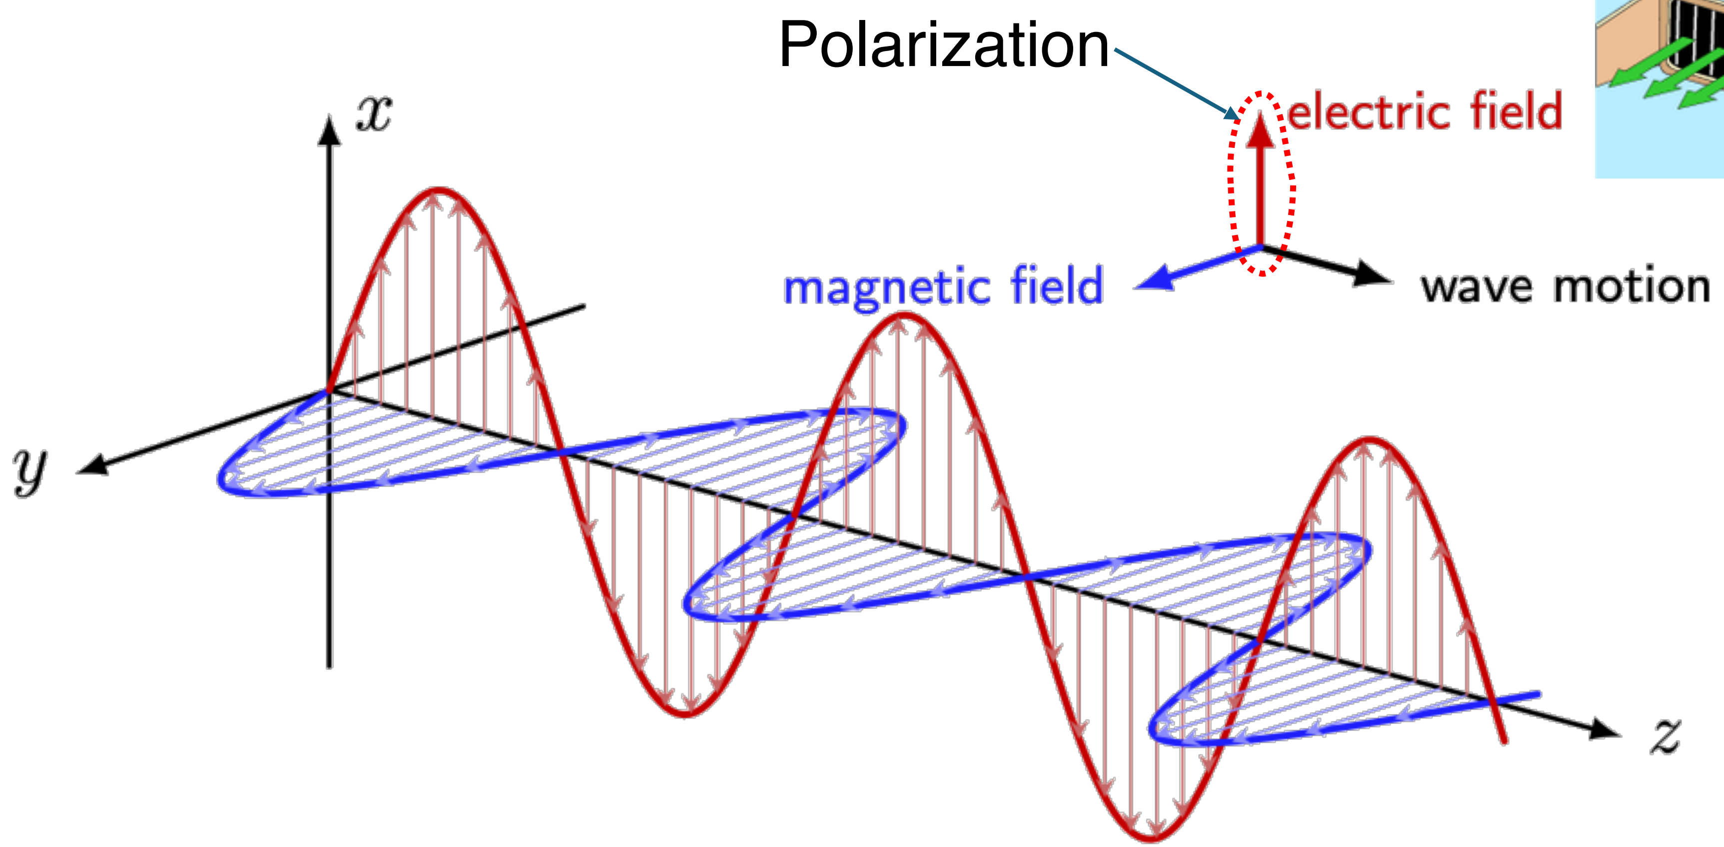

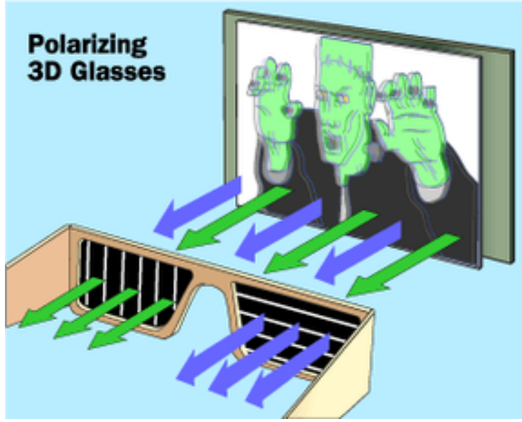

What is Light?

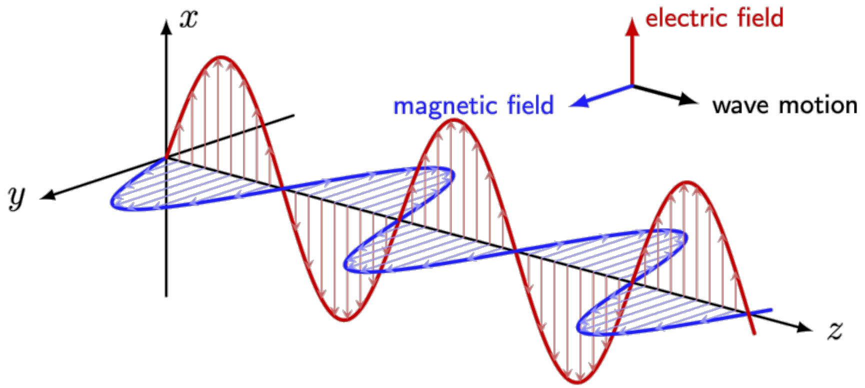

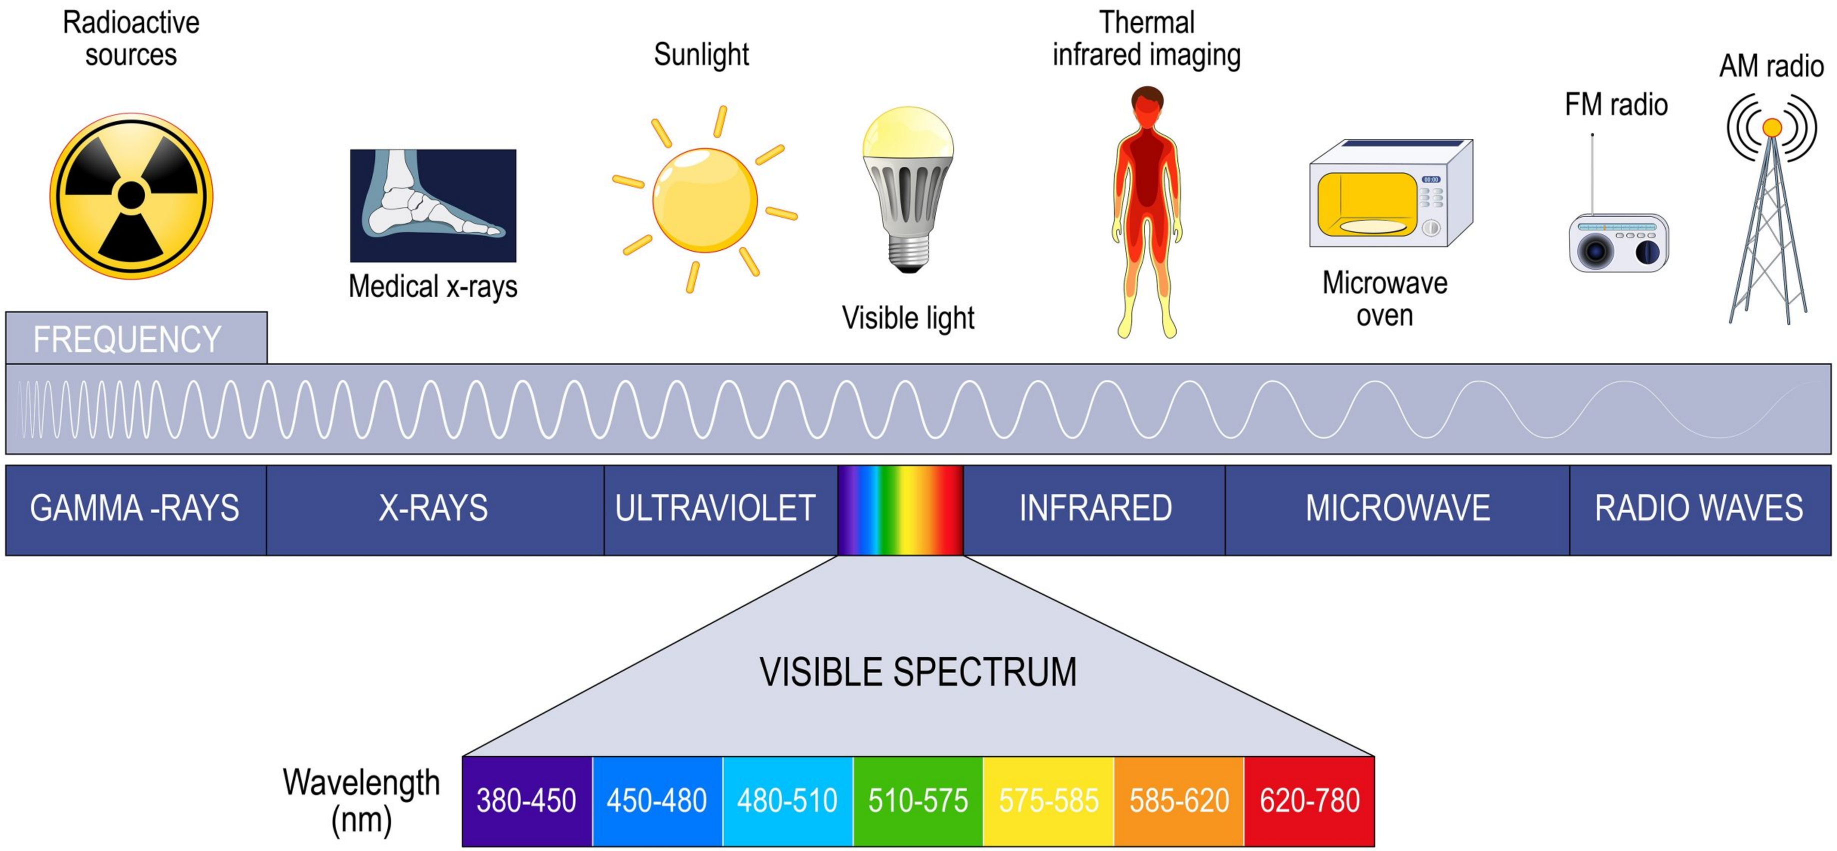

So light is a wave, Einstein?

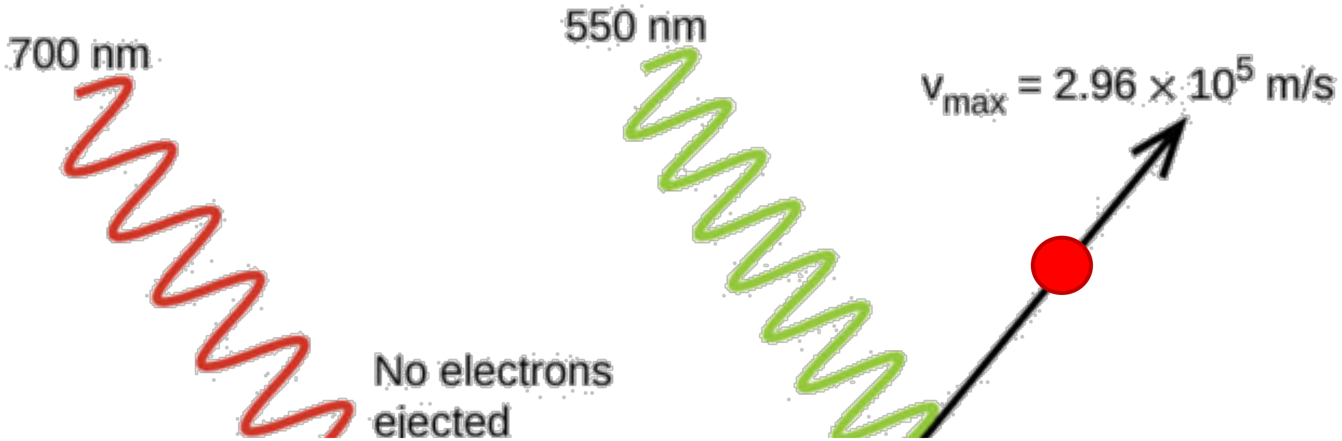

So light is a wave, Einstein?

So light is particles?

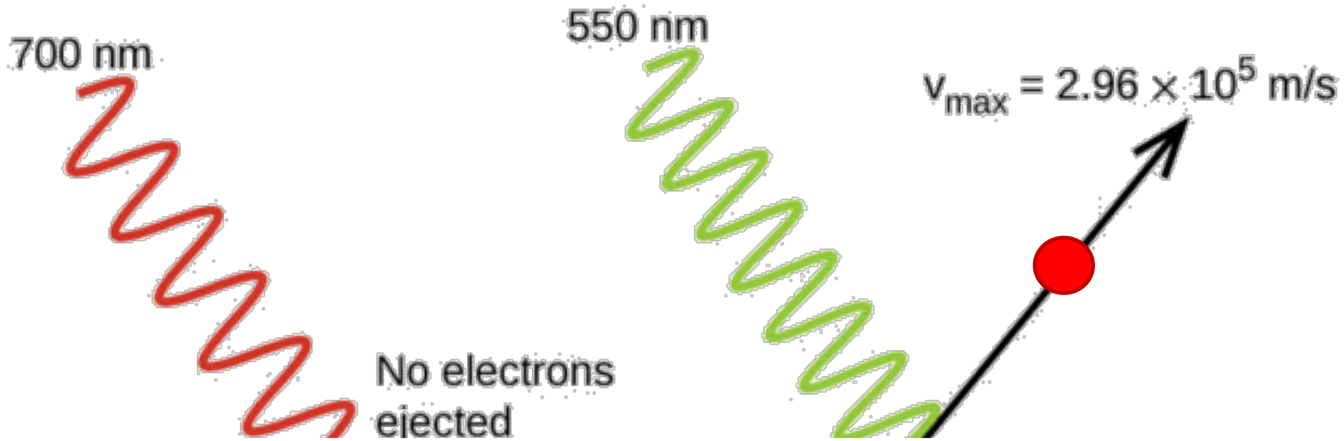

So light is a wave, Einstein?

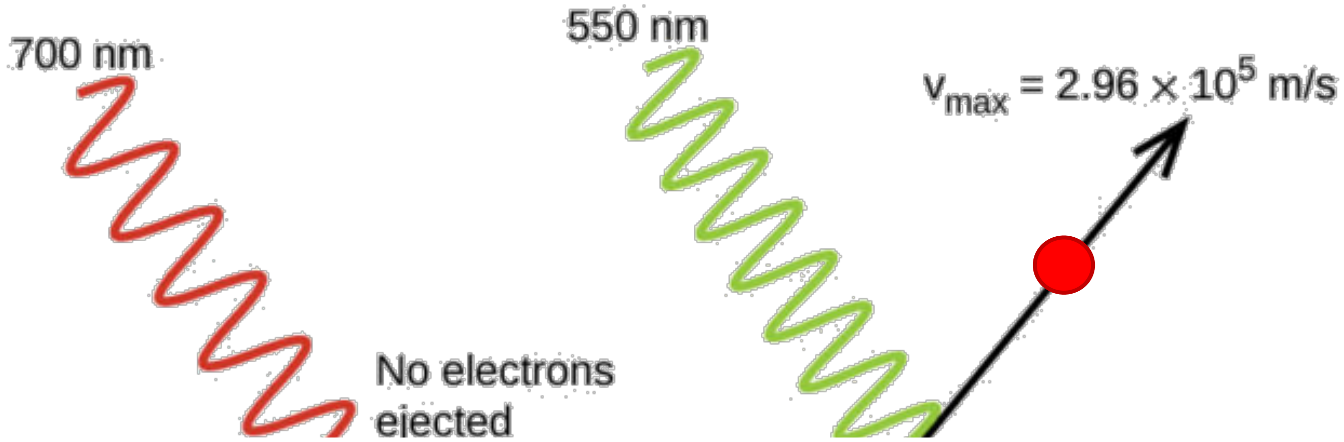

So light is particles?

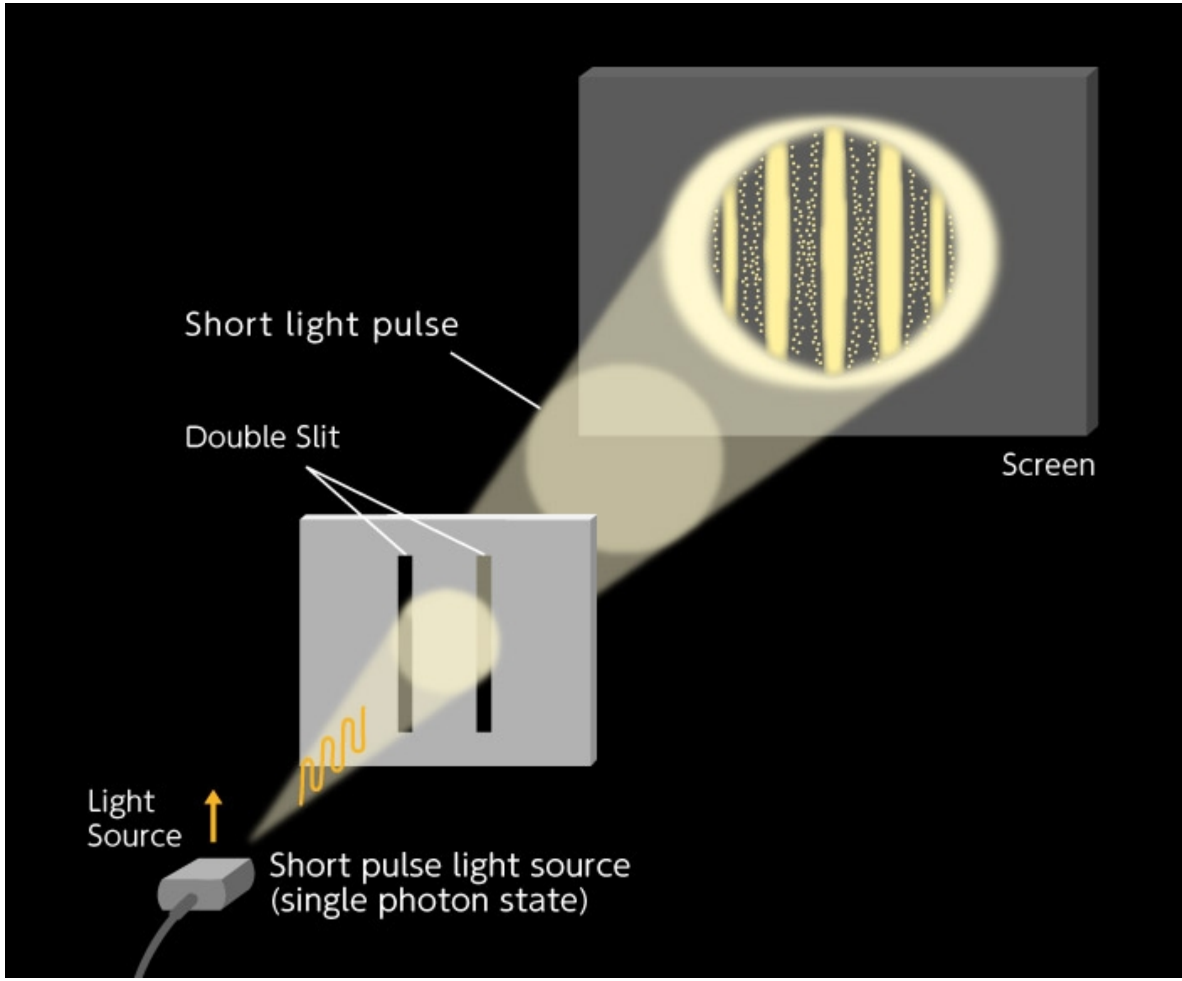

So light is a wave, Einstein?

So light is particles?

Reflection	✓	✓
Refraction	✓	✓
Interference	✗	✓
Diffraction	✗	✓
Polarization	✗	✓
Photoelectric effect	✓	✗

So light is a wave, Einstein?

So light is particles?

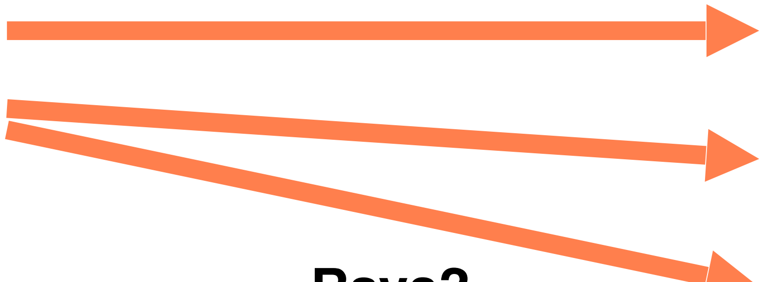

	●→	∿∿∿
Reflection	✓	✓
Refraction	✓	✓
Interference	✗	✓
Diffraction	✗	✓
Polarization	✗	✓
Photoelectric effect	✓	✗

So light is a wave, Einstein?

So light is particles?

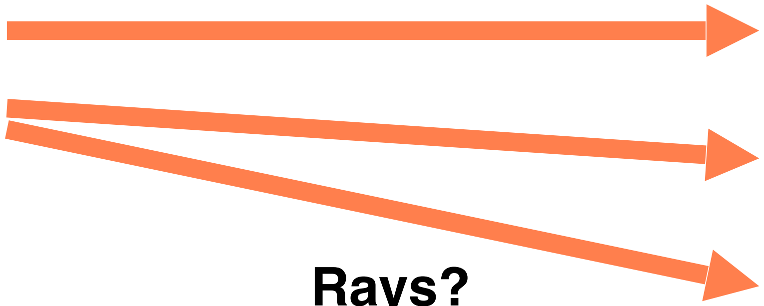

Rays?

stationary phase approximation

Reflection	✓	✓
Refraction	✓	✓
Interference	✗	✓
Diffraction	✗	✓
Polarization	✗	✓
Photoelectric effect	✓	✗

“View from the Window at Le Gras”; 1826

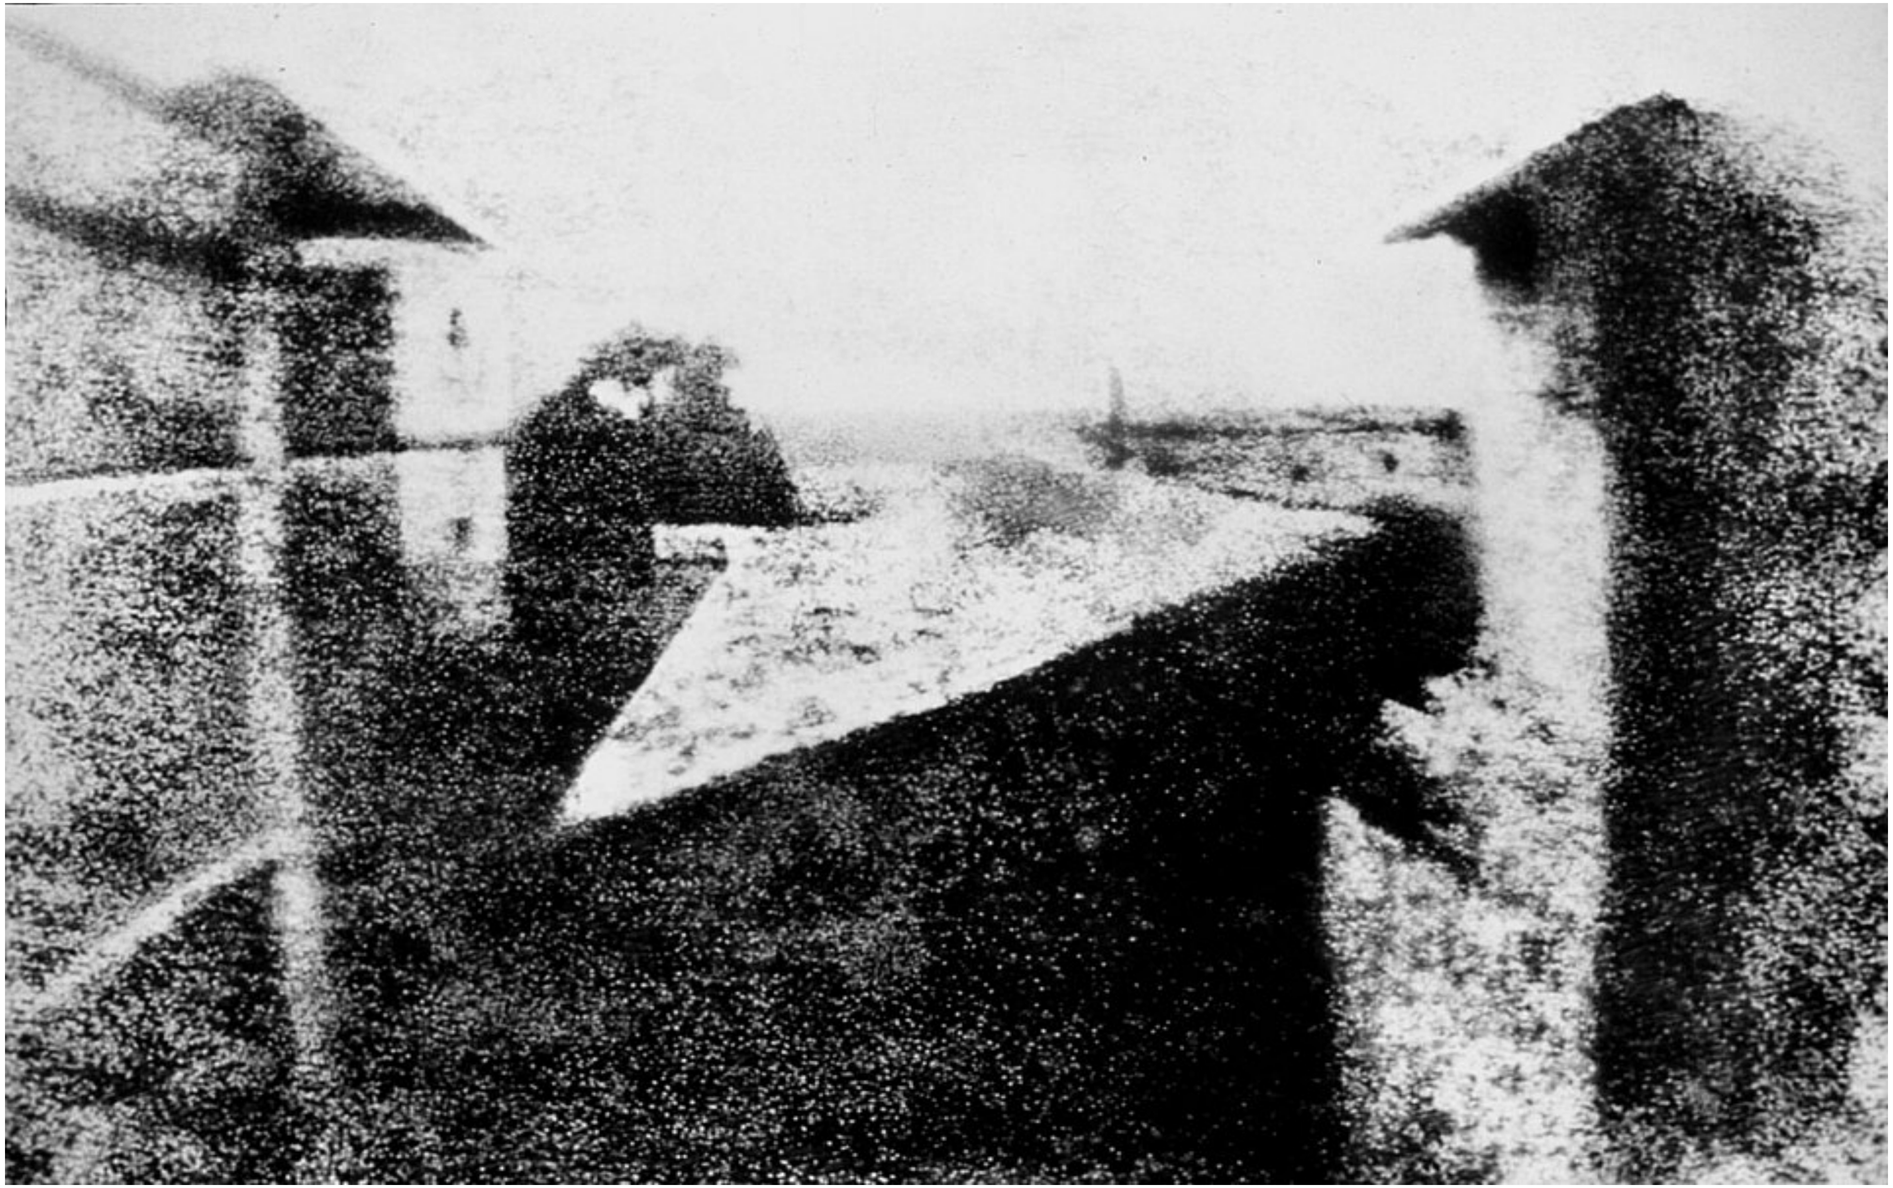

“View from the Window at Le Gras”; 1826

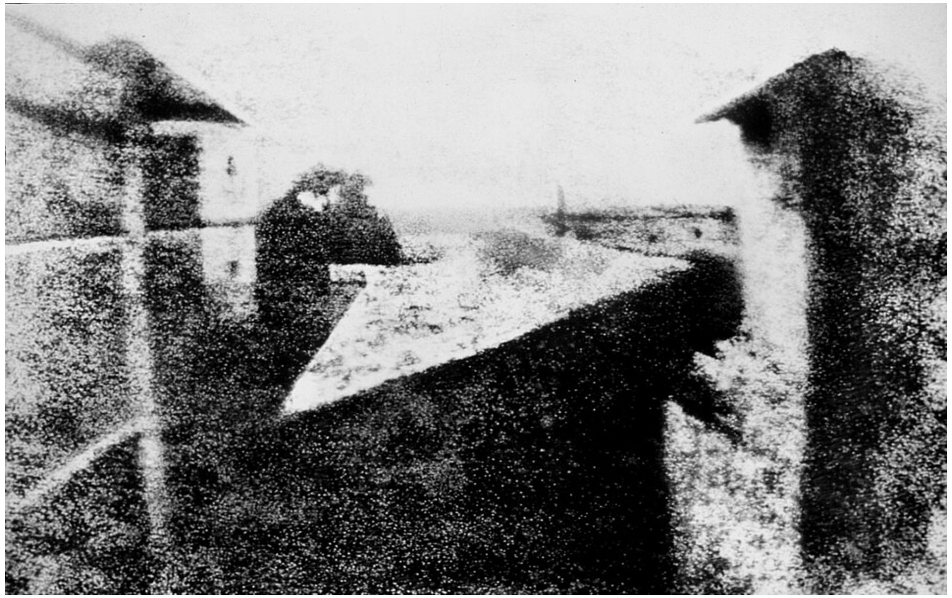

Pinhole Camera Model

“View from the Window at Le Gras”; 1826

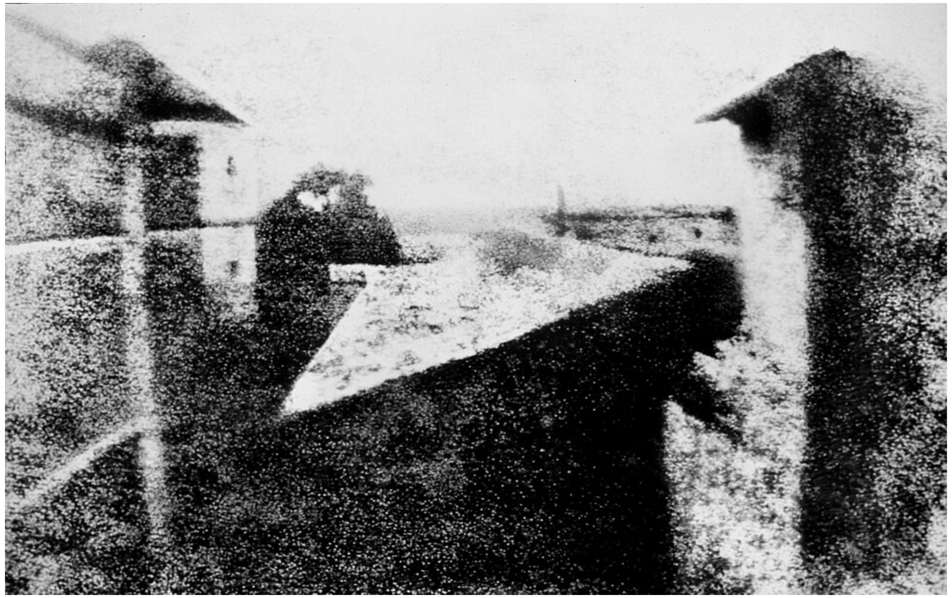

Trade-off

Time ←→ Intensity ←→ Resolution

Pinhole Camera Model

“View from the Window at Le Gras”; 1826

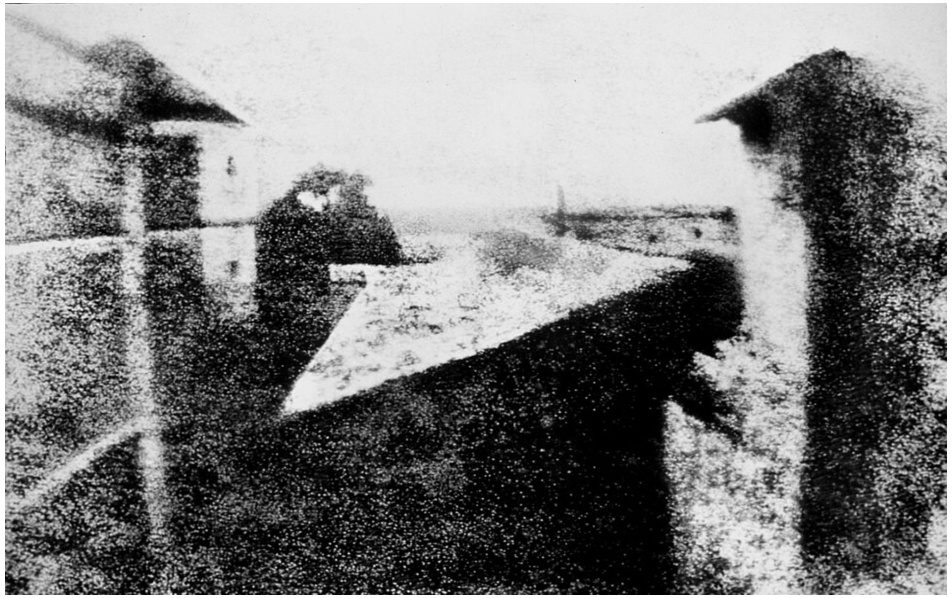

Trade-off

Time ←→ Intensity ←→ Resolution

Pinhole Camera Model

“You press the button, we do the rest”; Kodak 1888

Refraction

Refraction

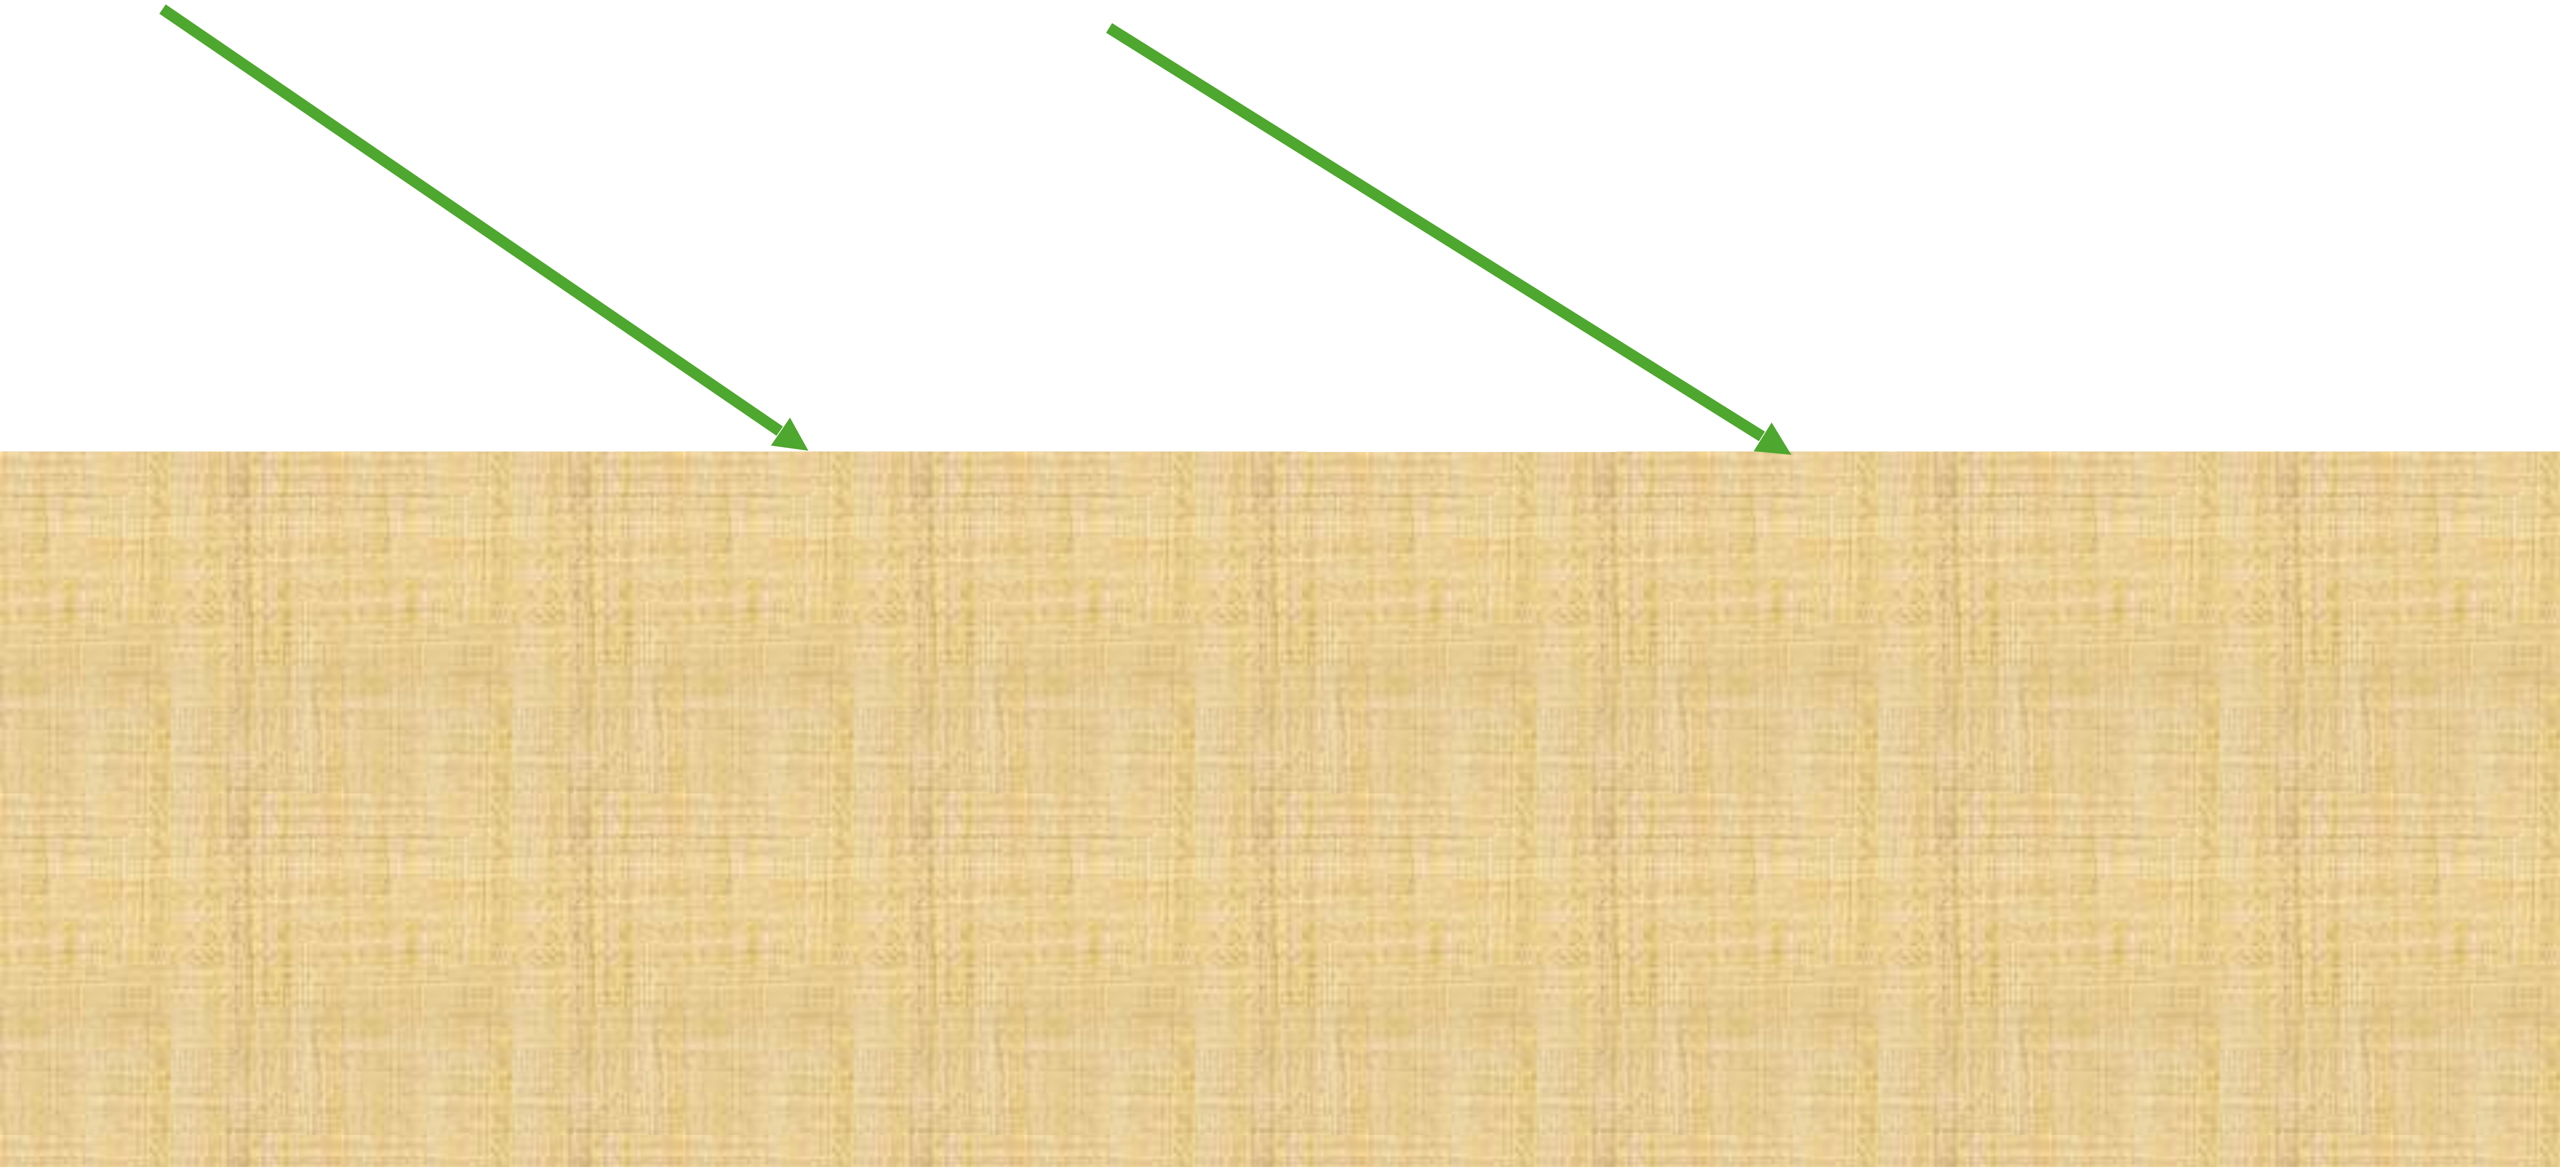

Refraction

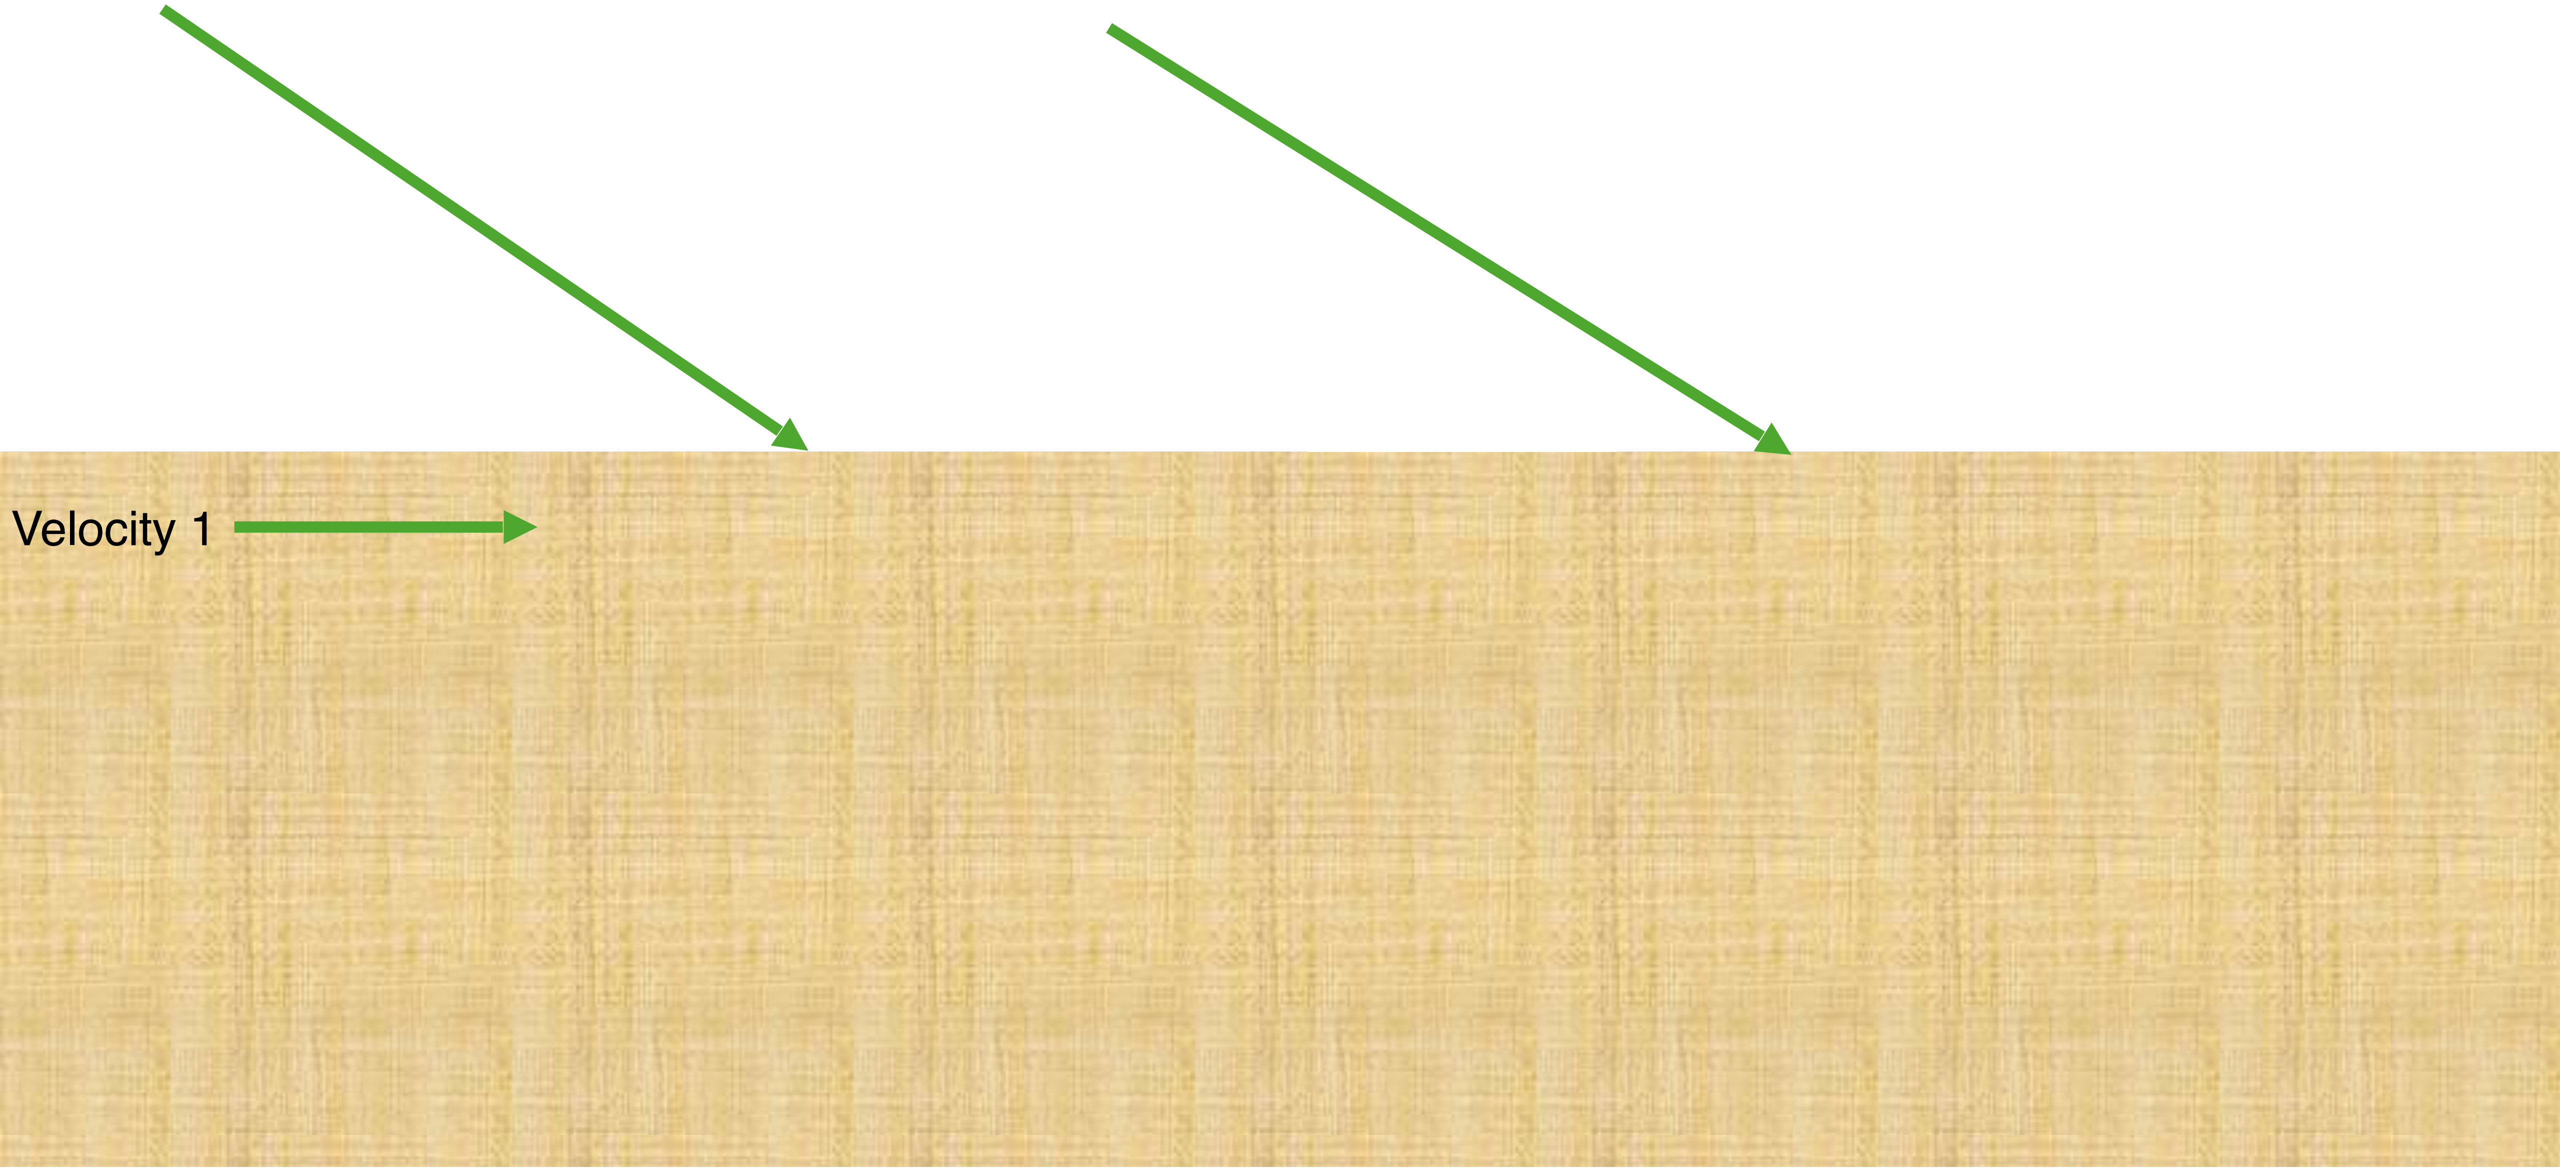

Refraction

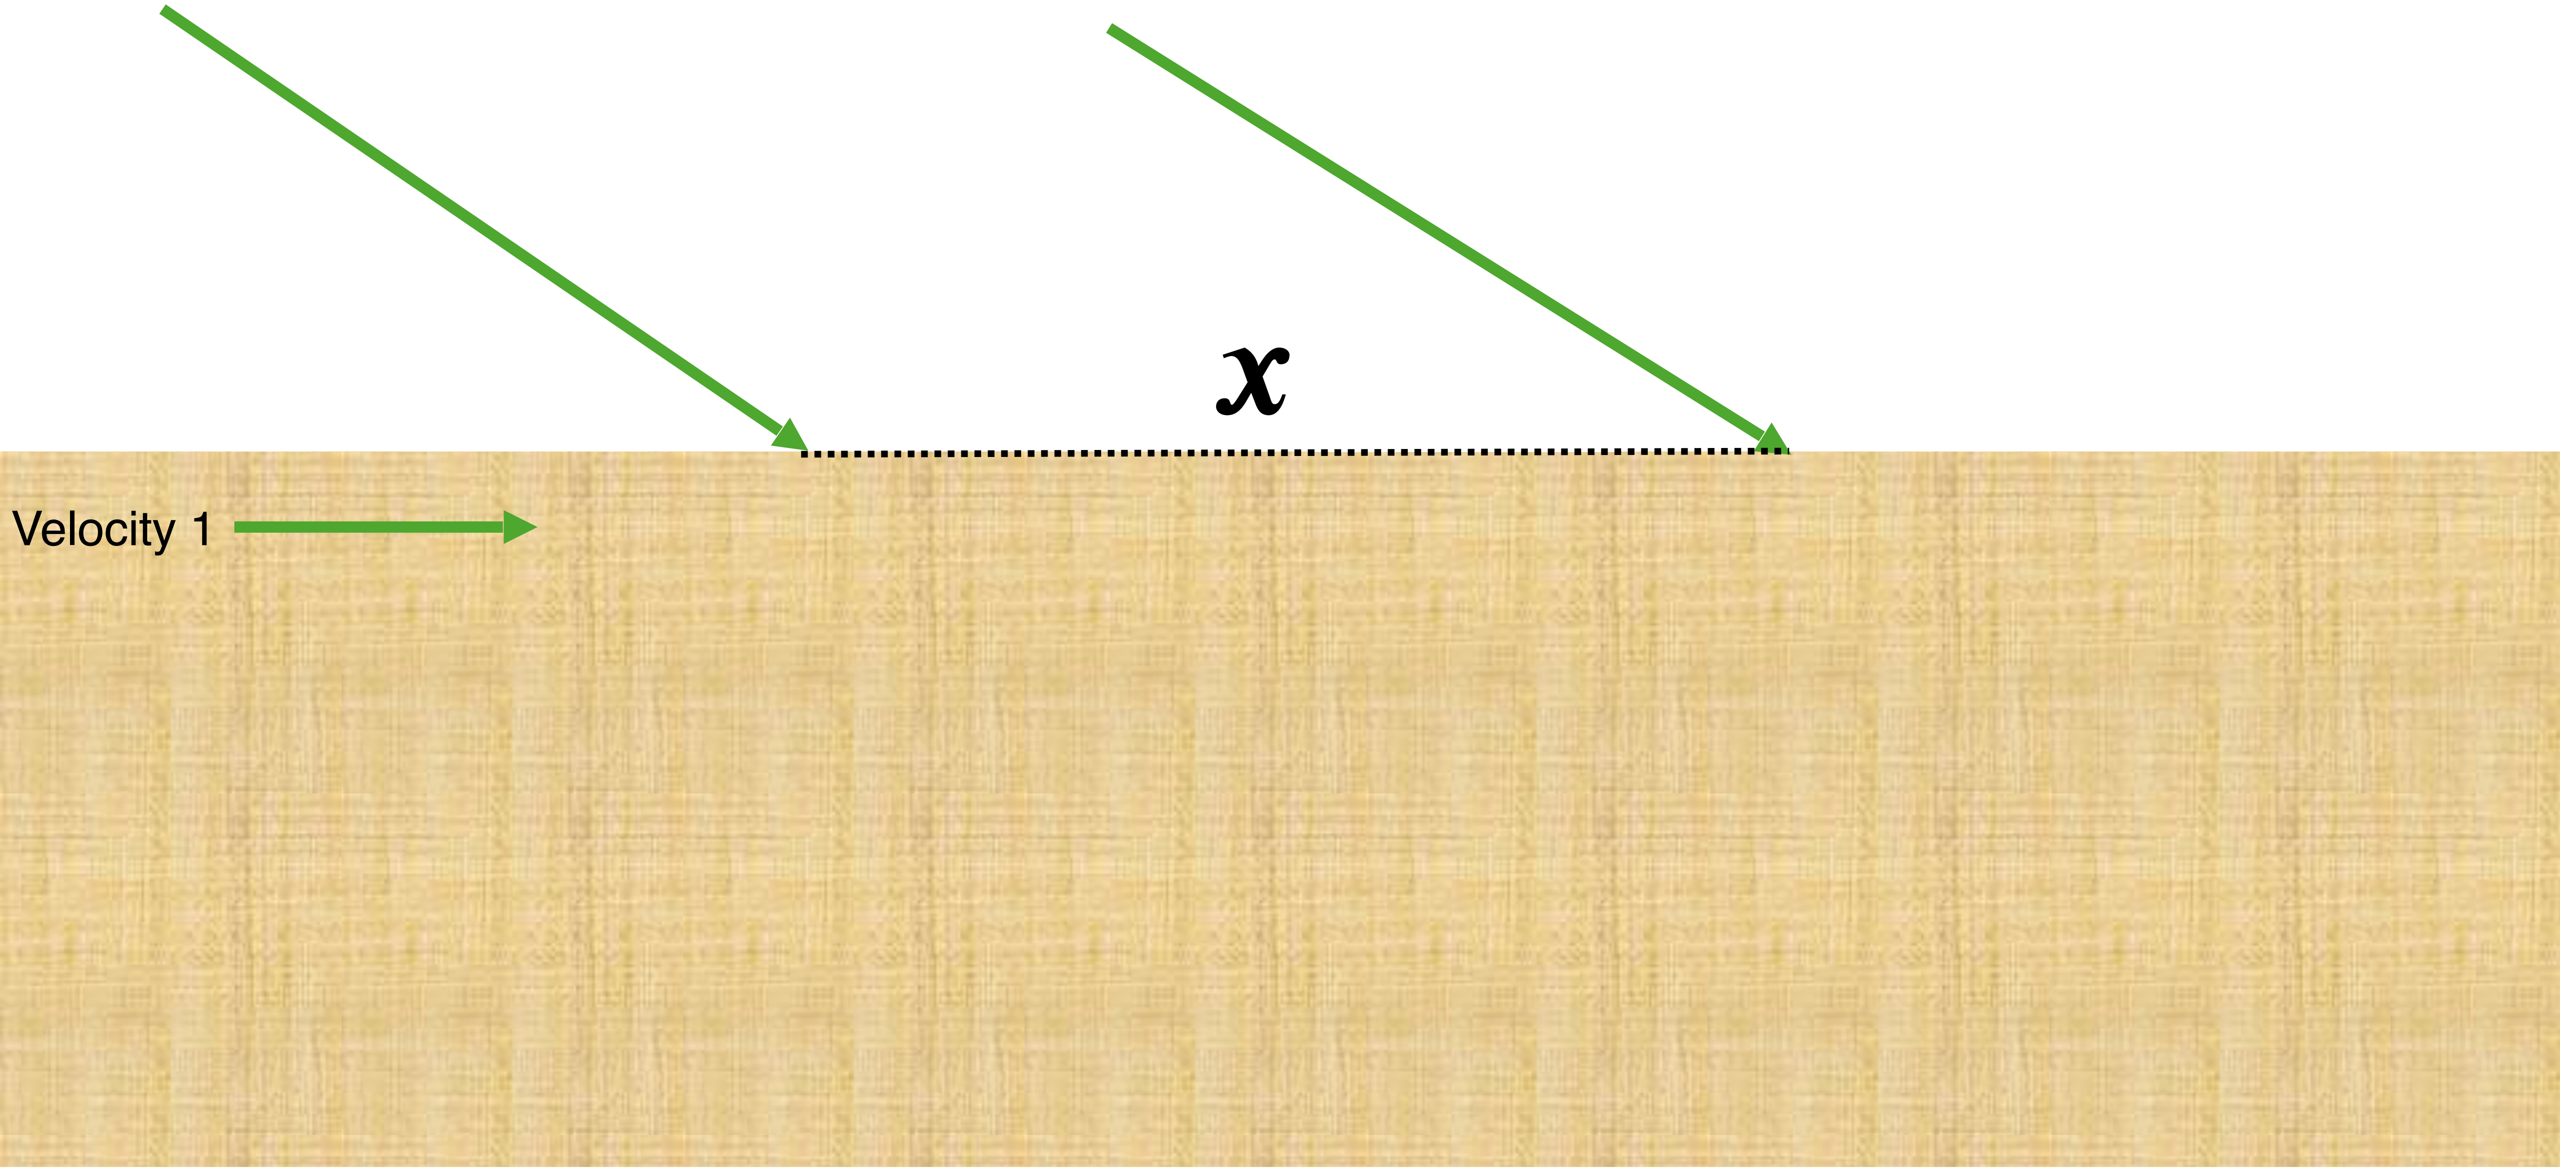

Refraction

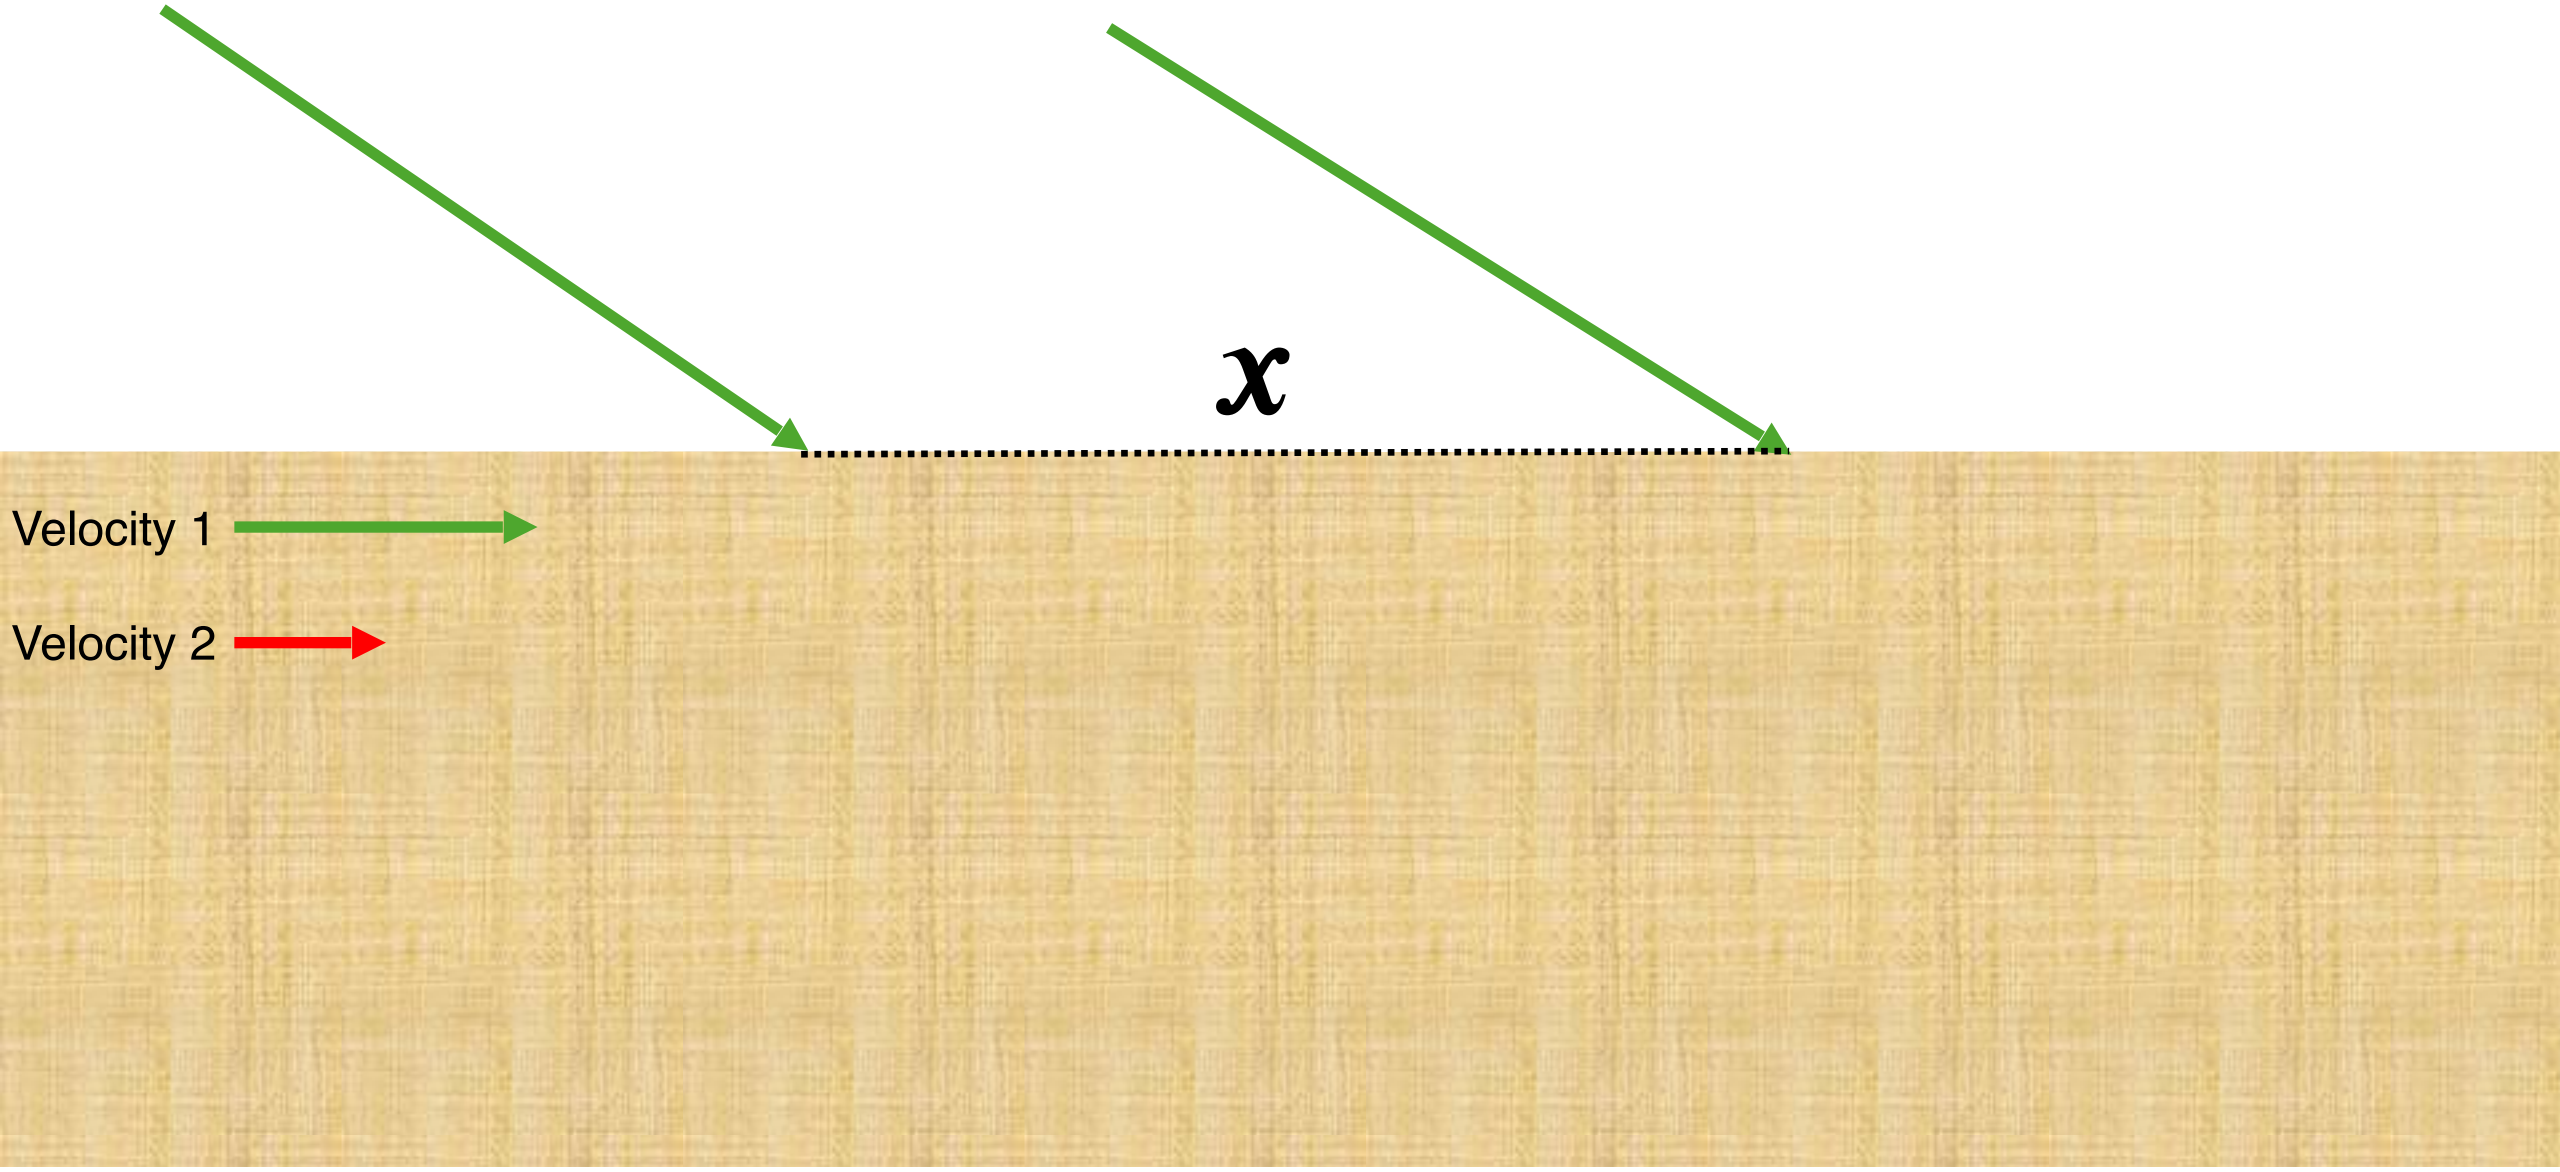

Refraction

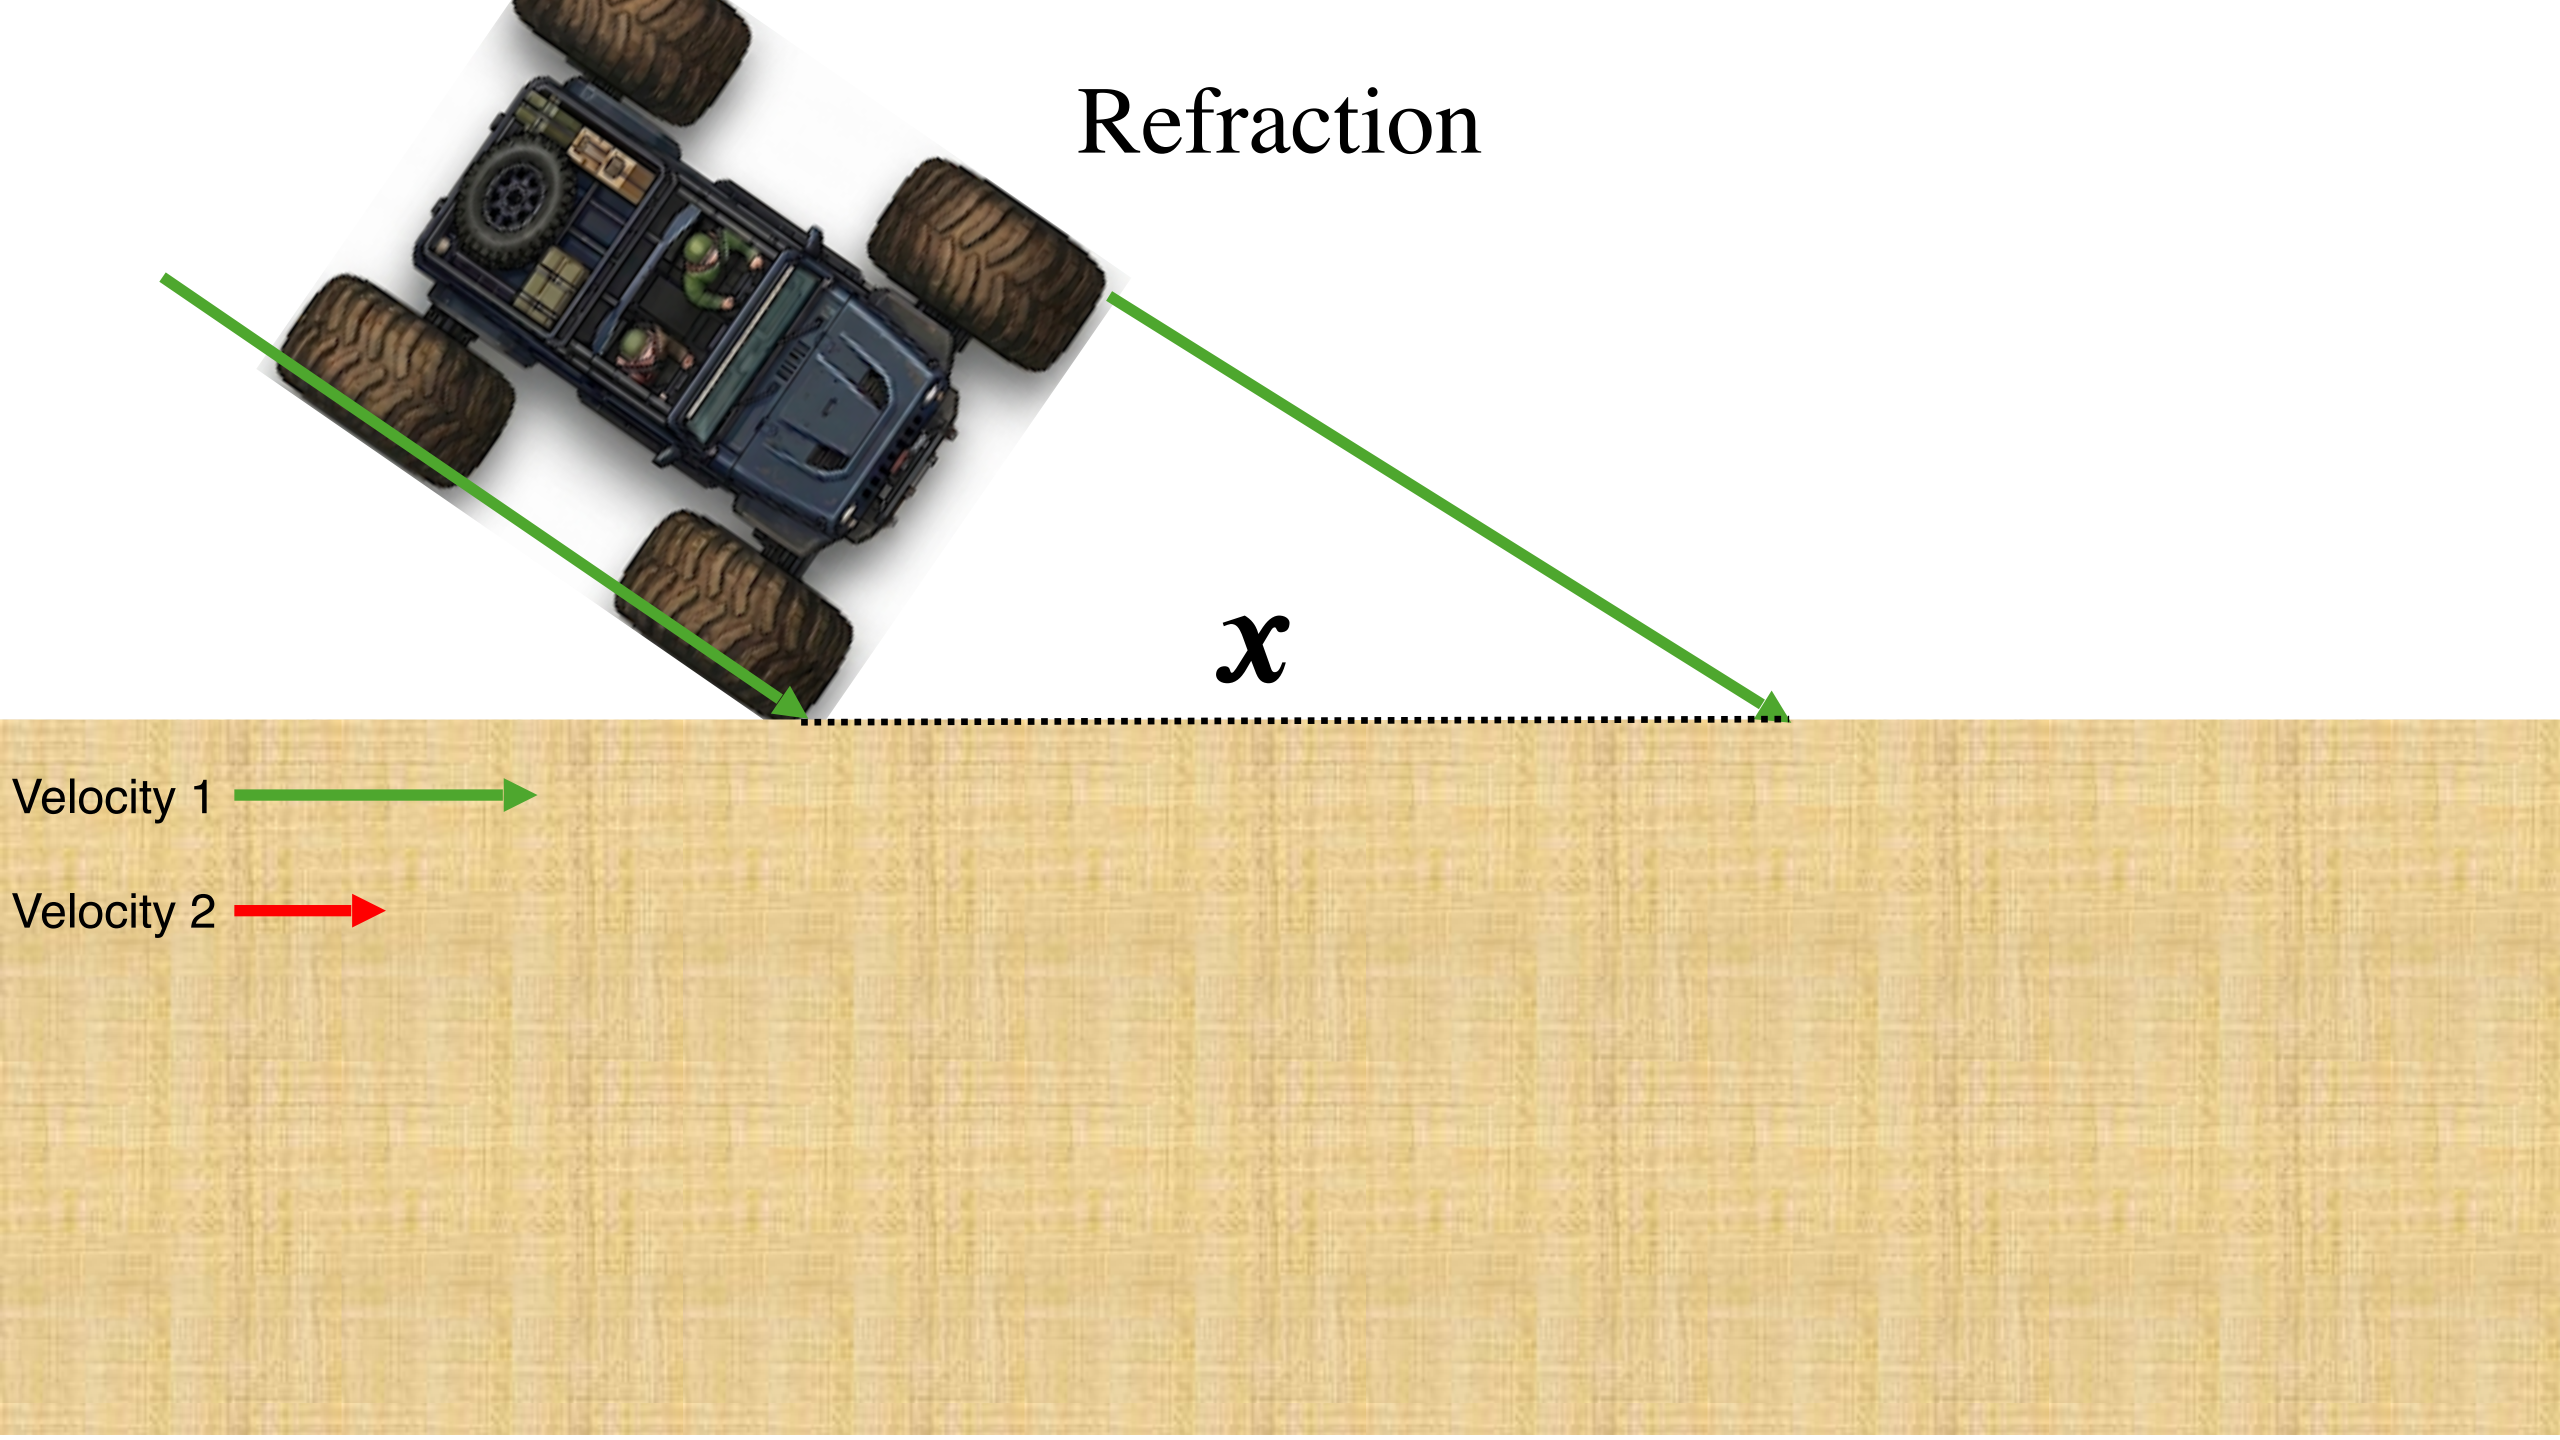

Refraction

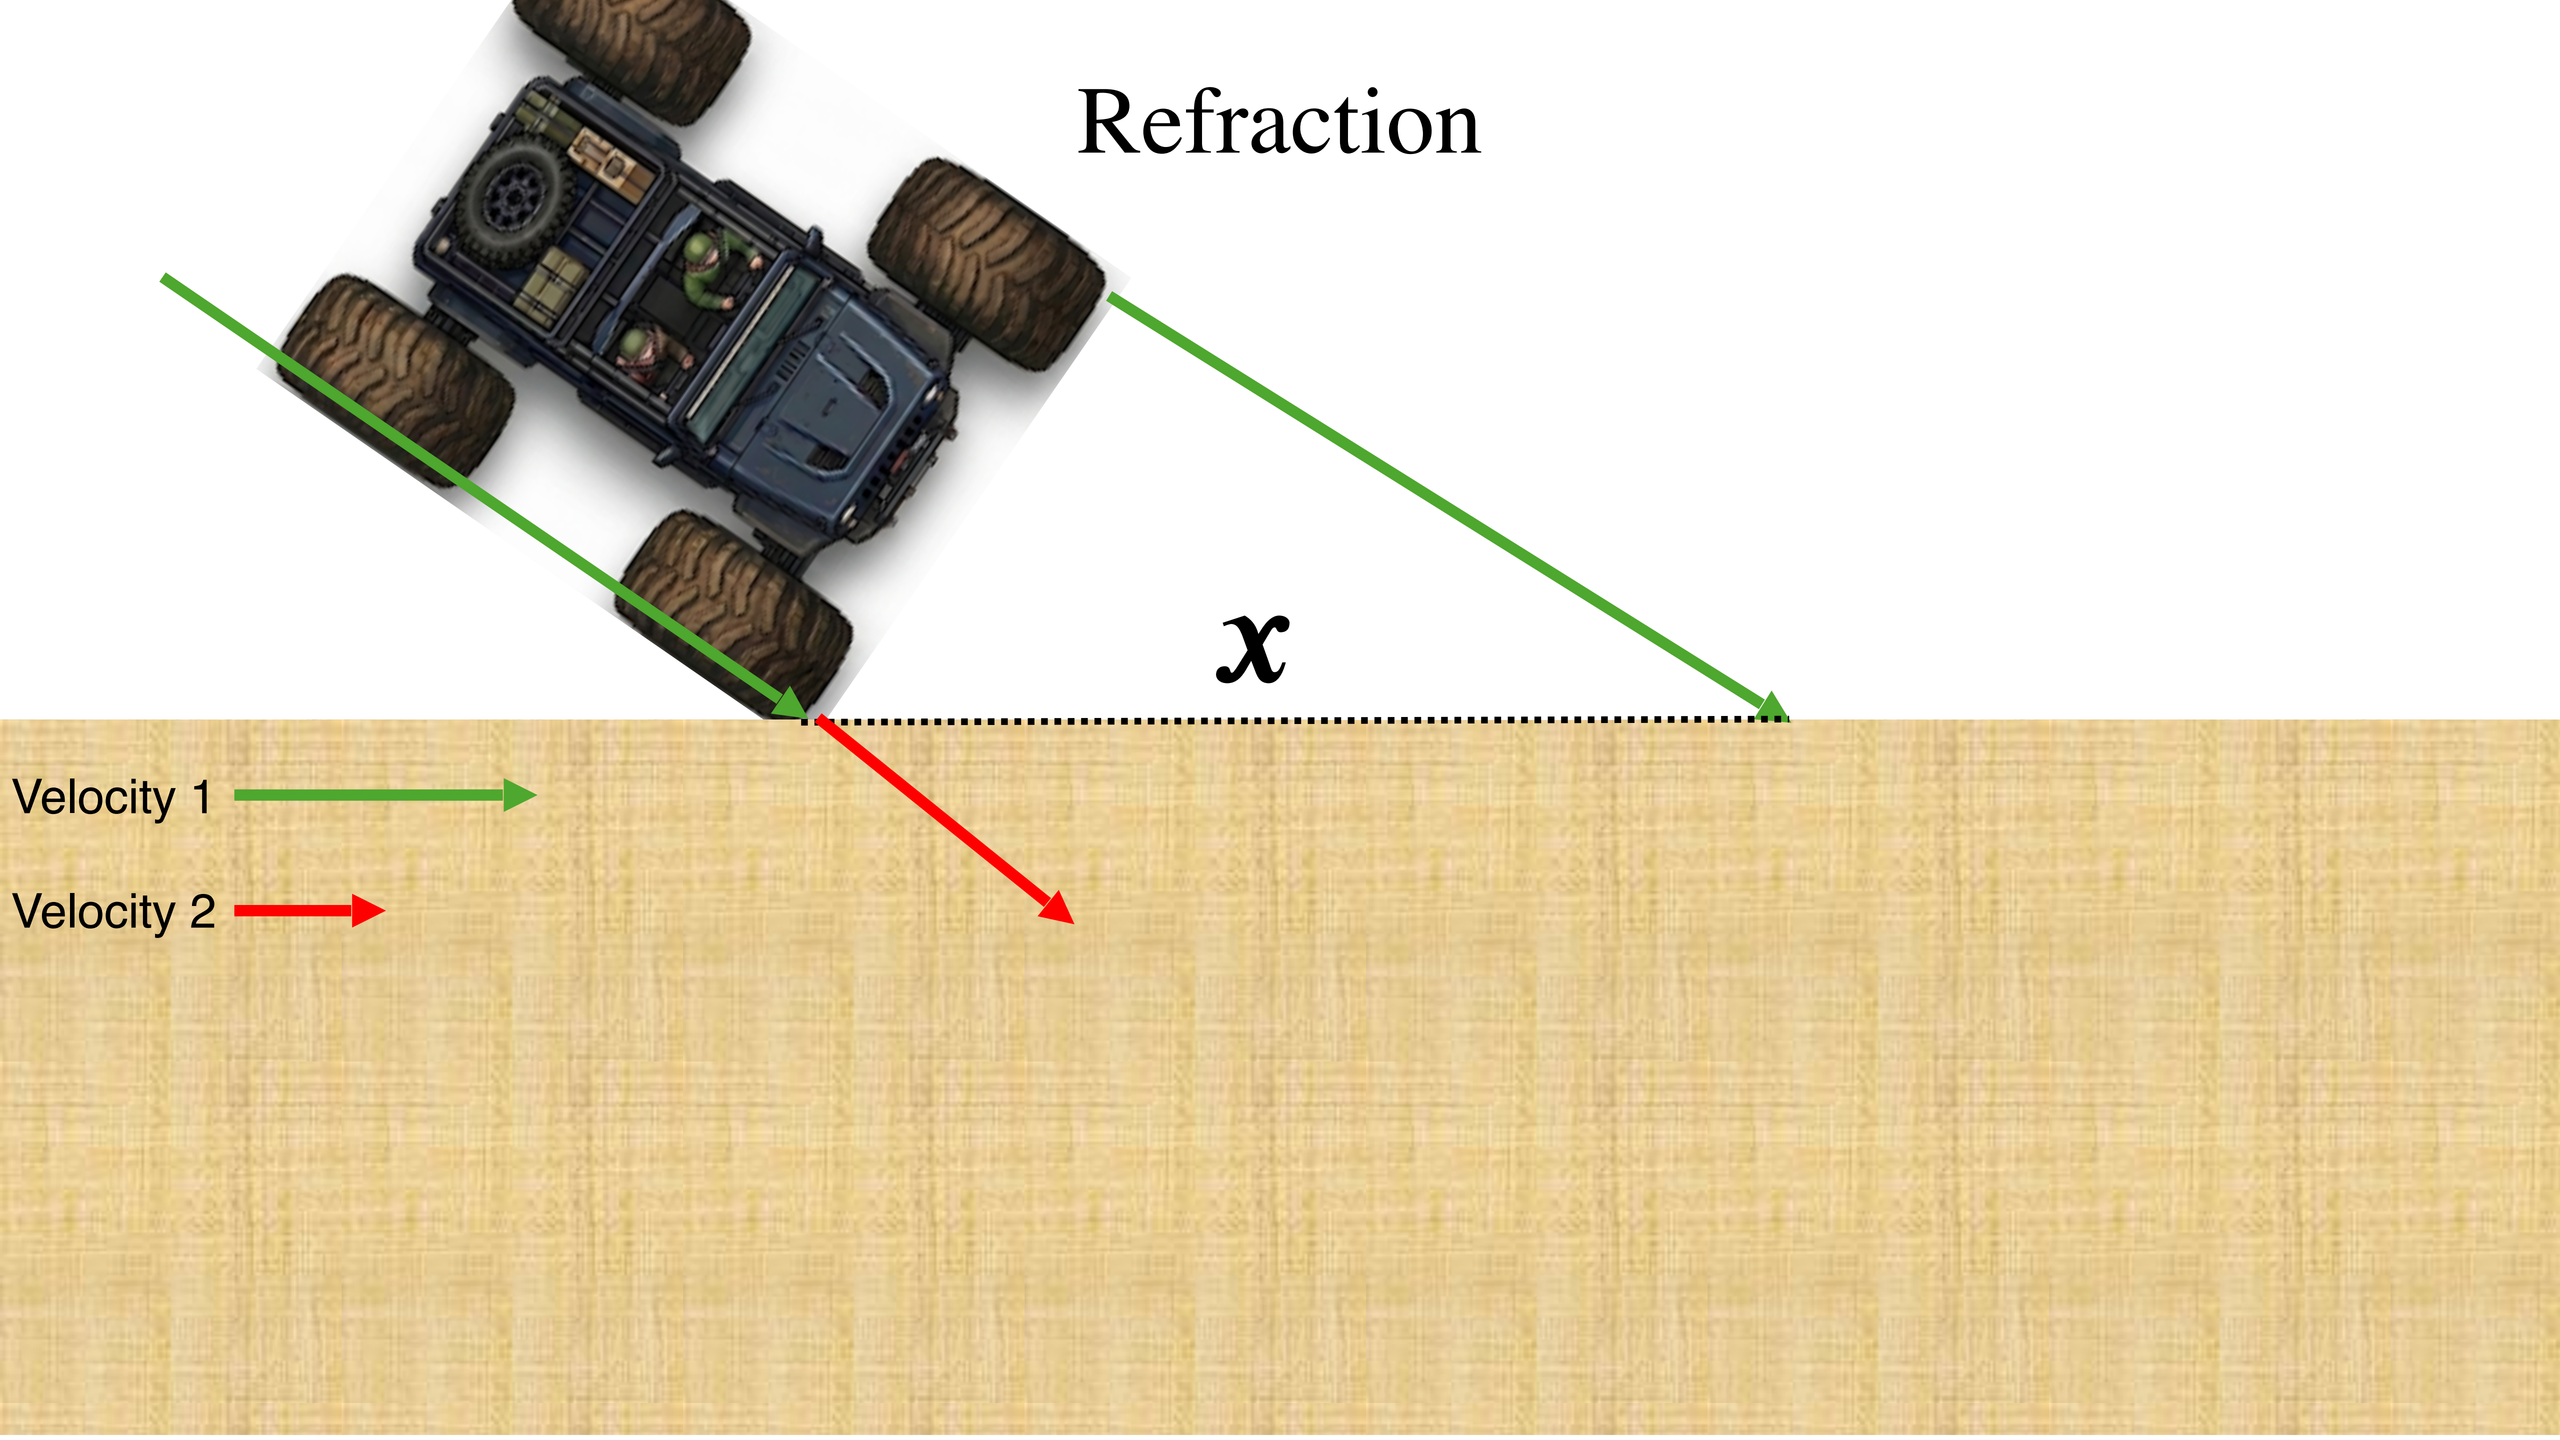

Refraction

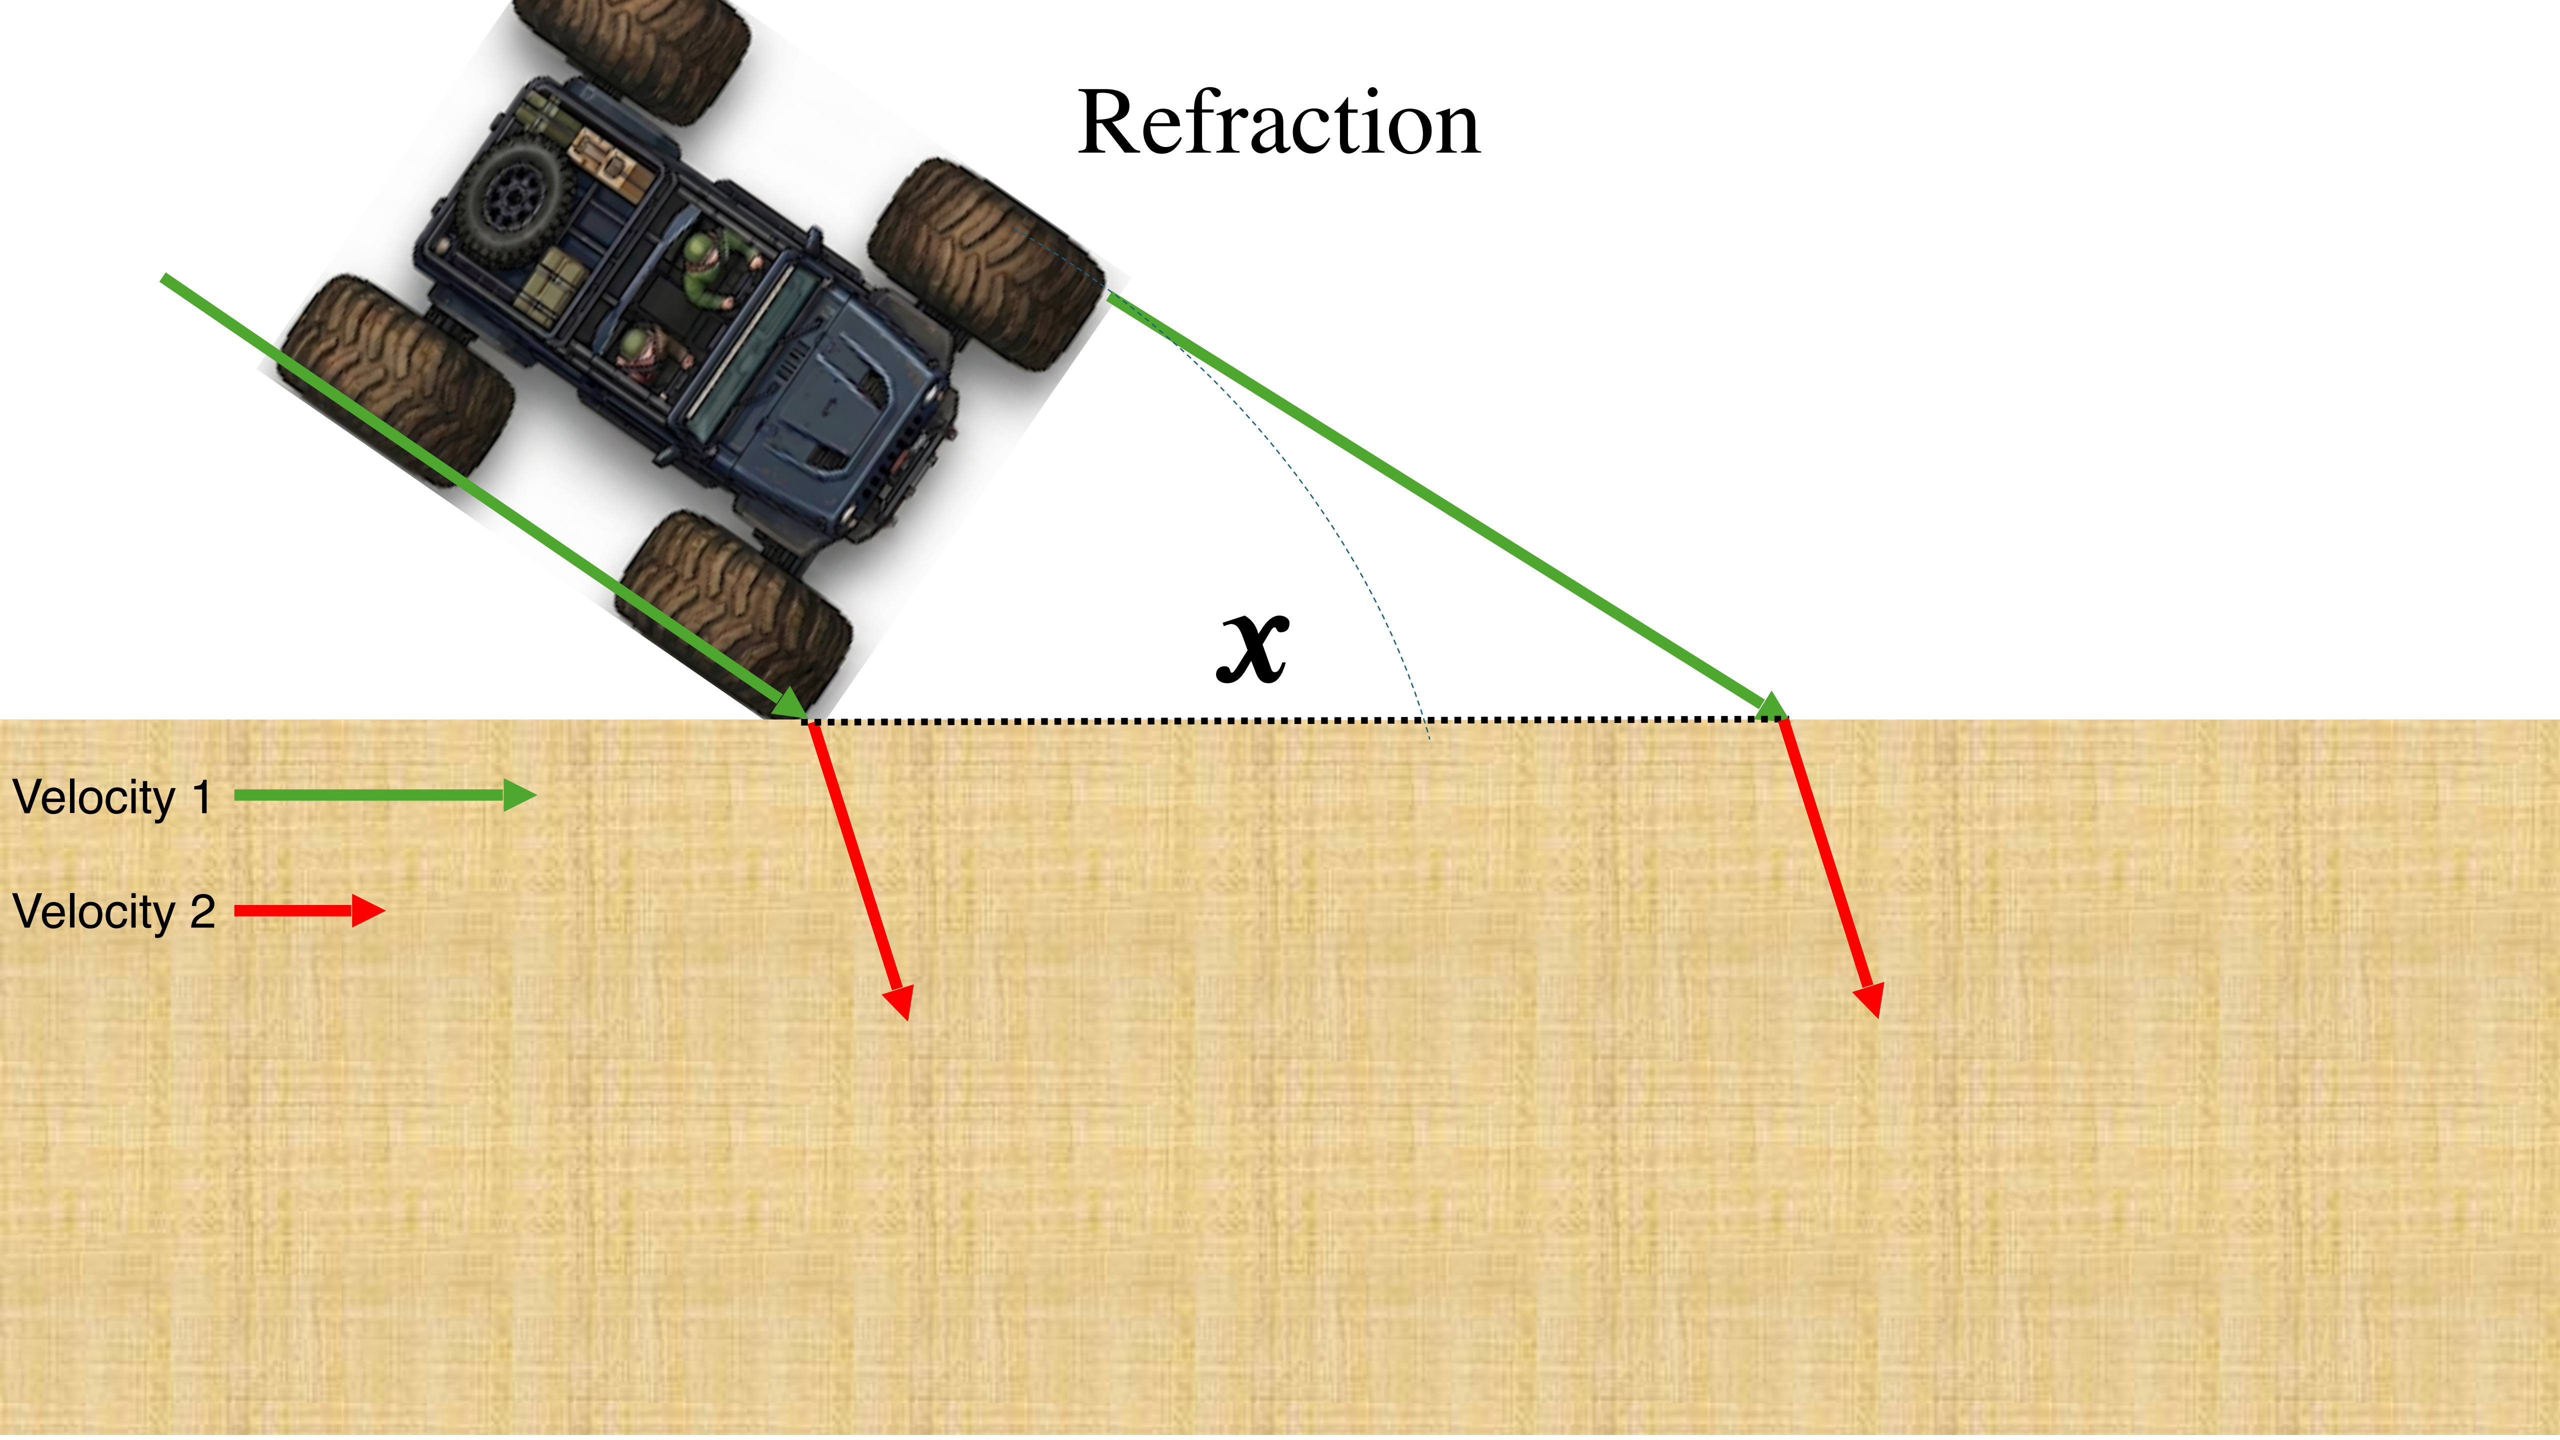

Refraction

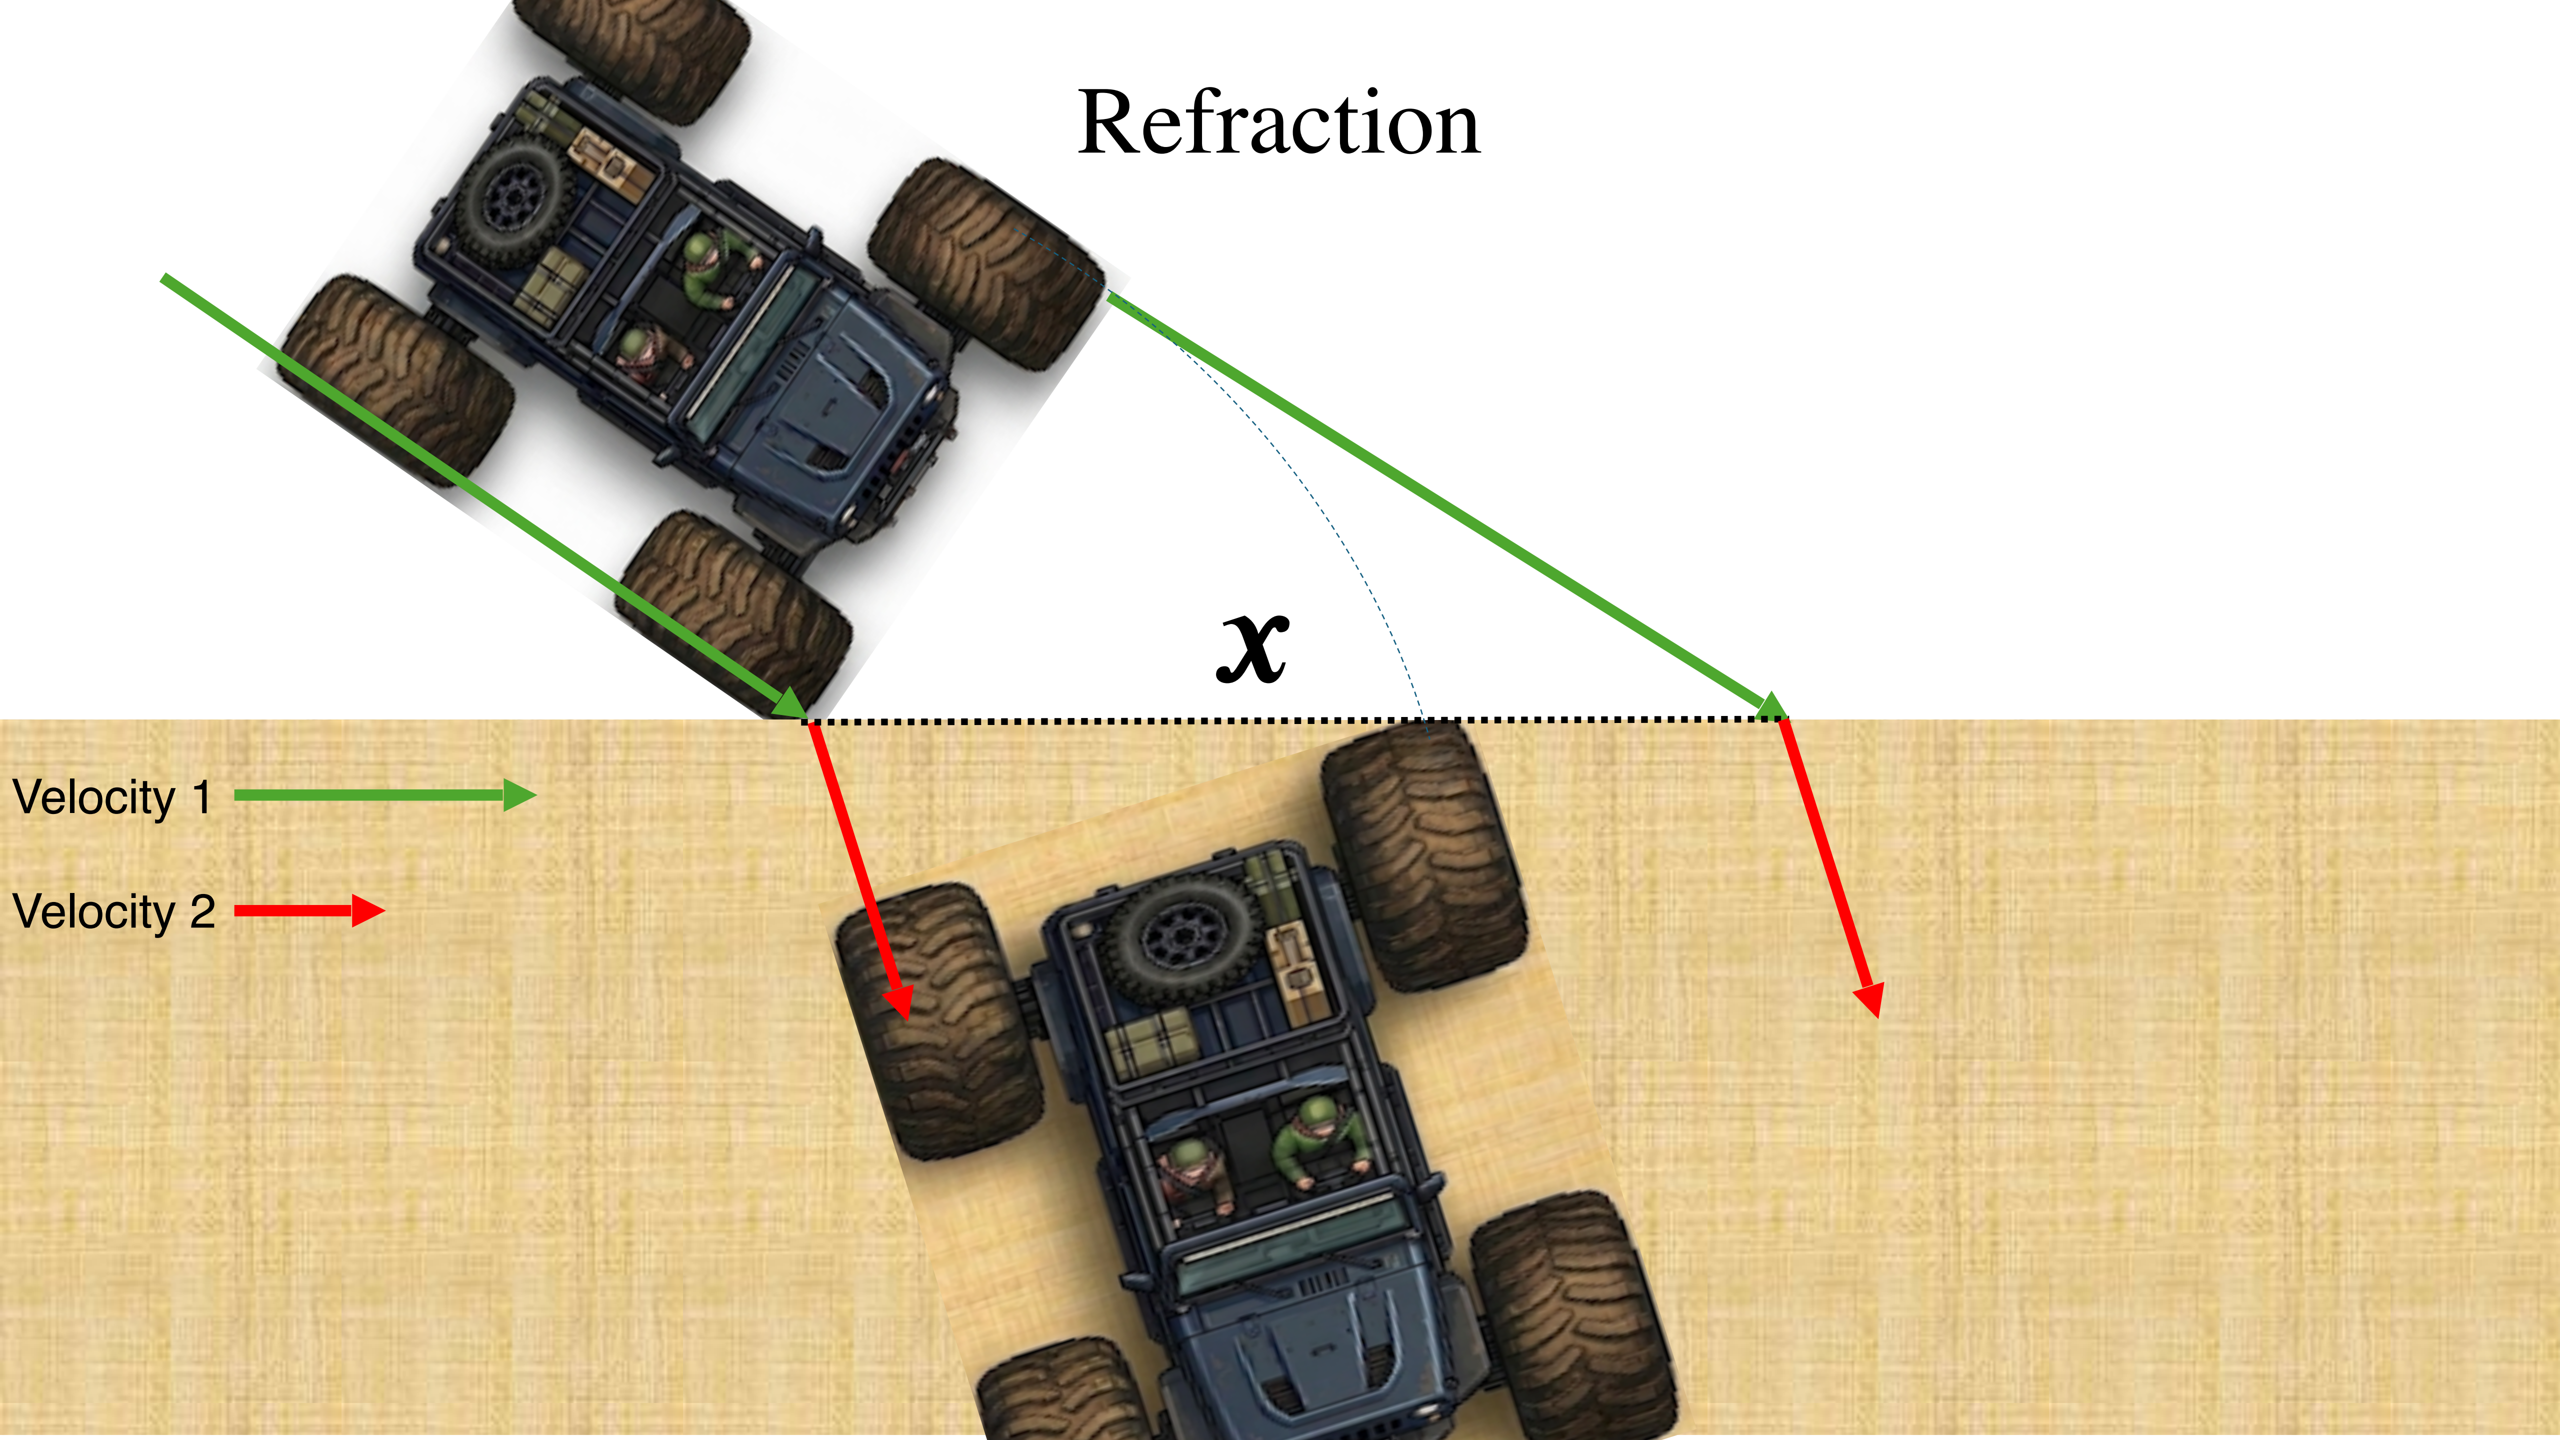

Refraction

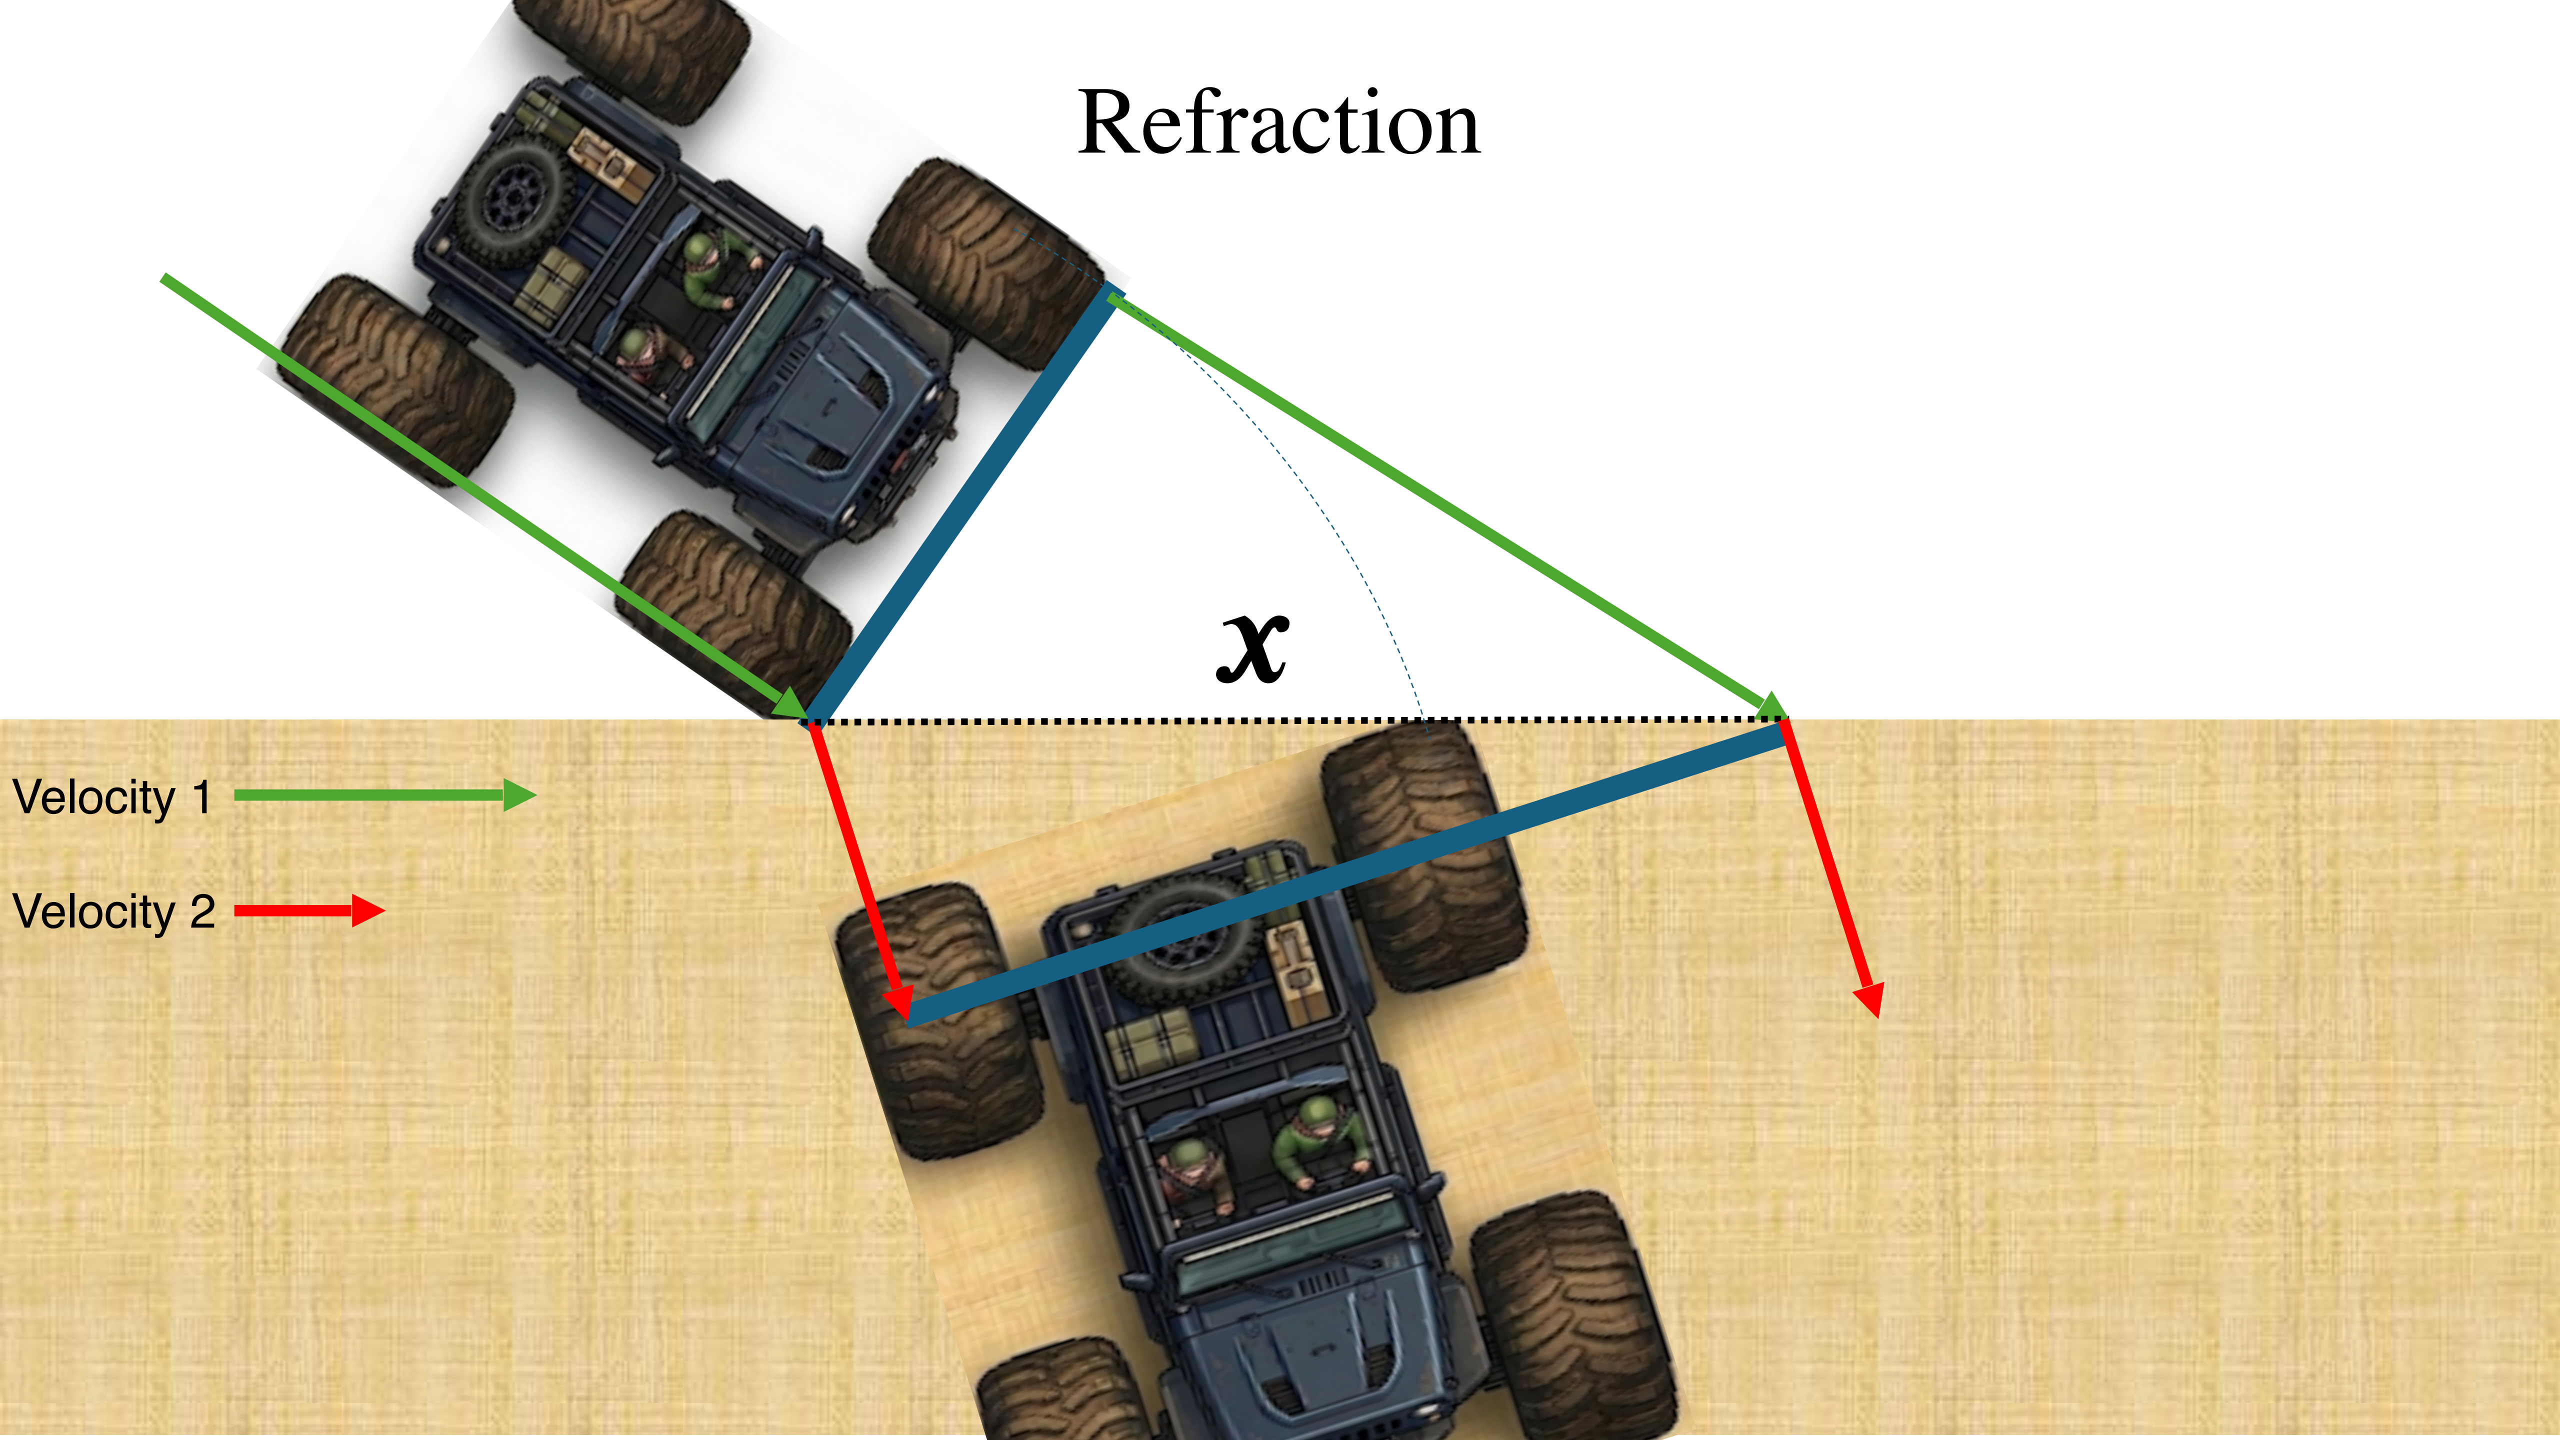

Refraction

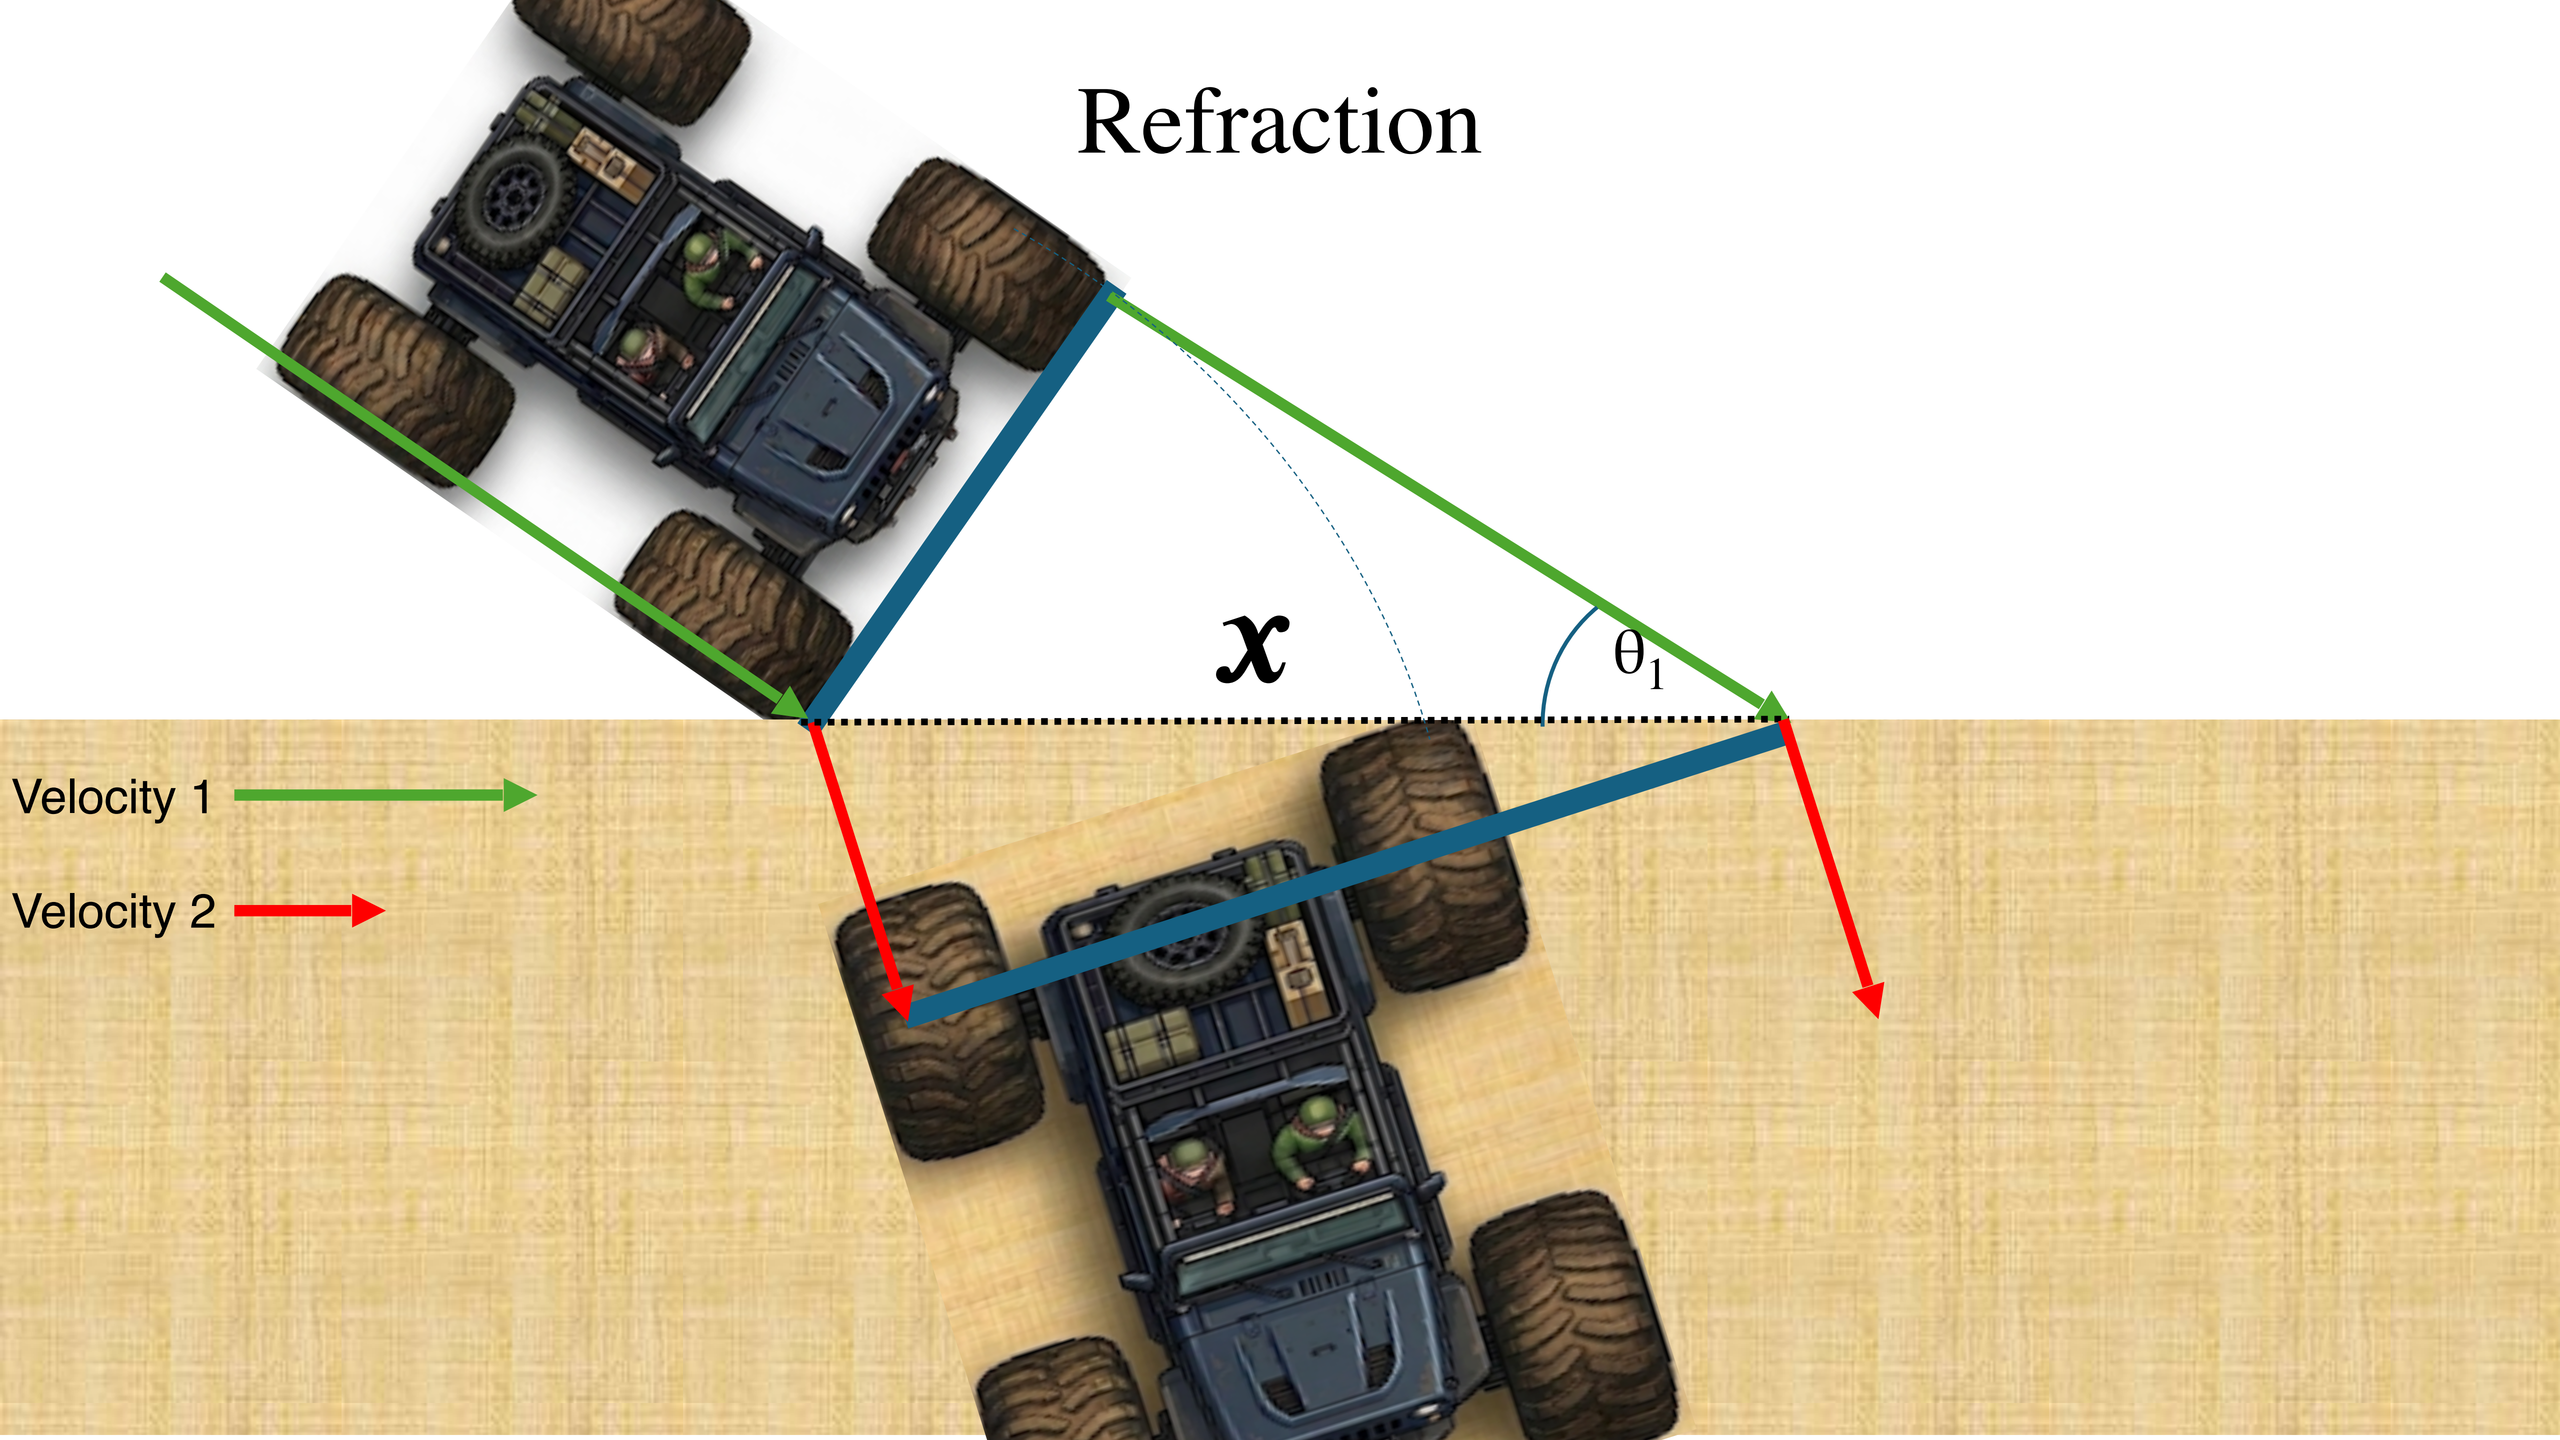

Refraction

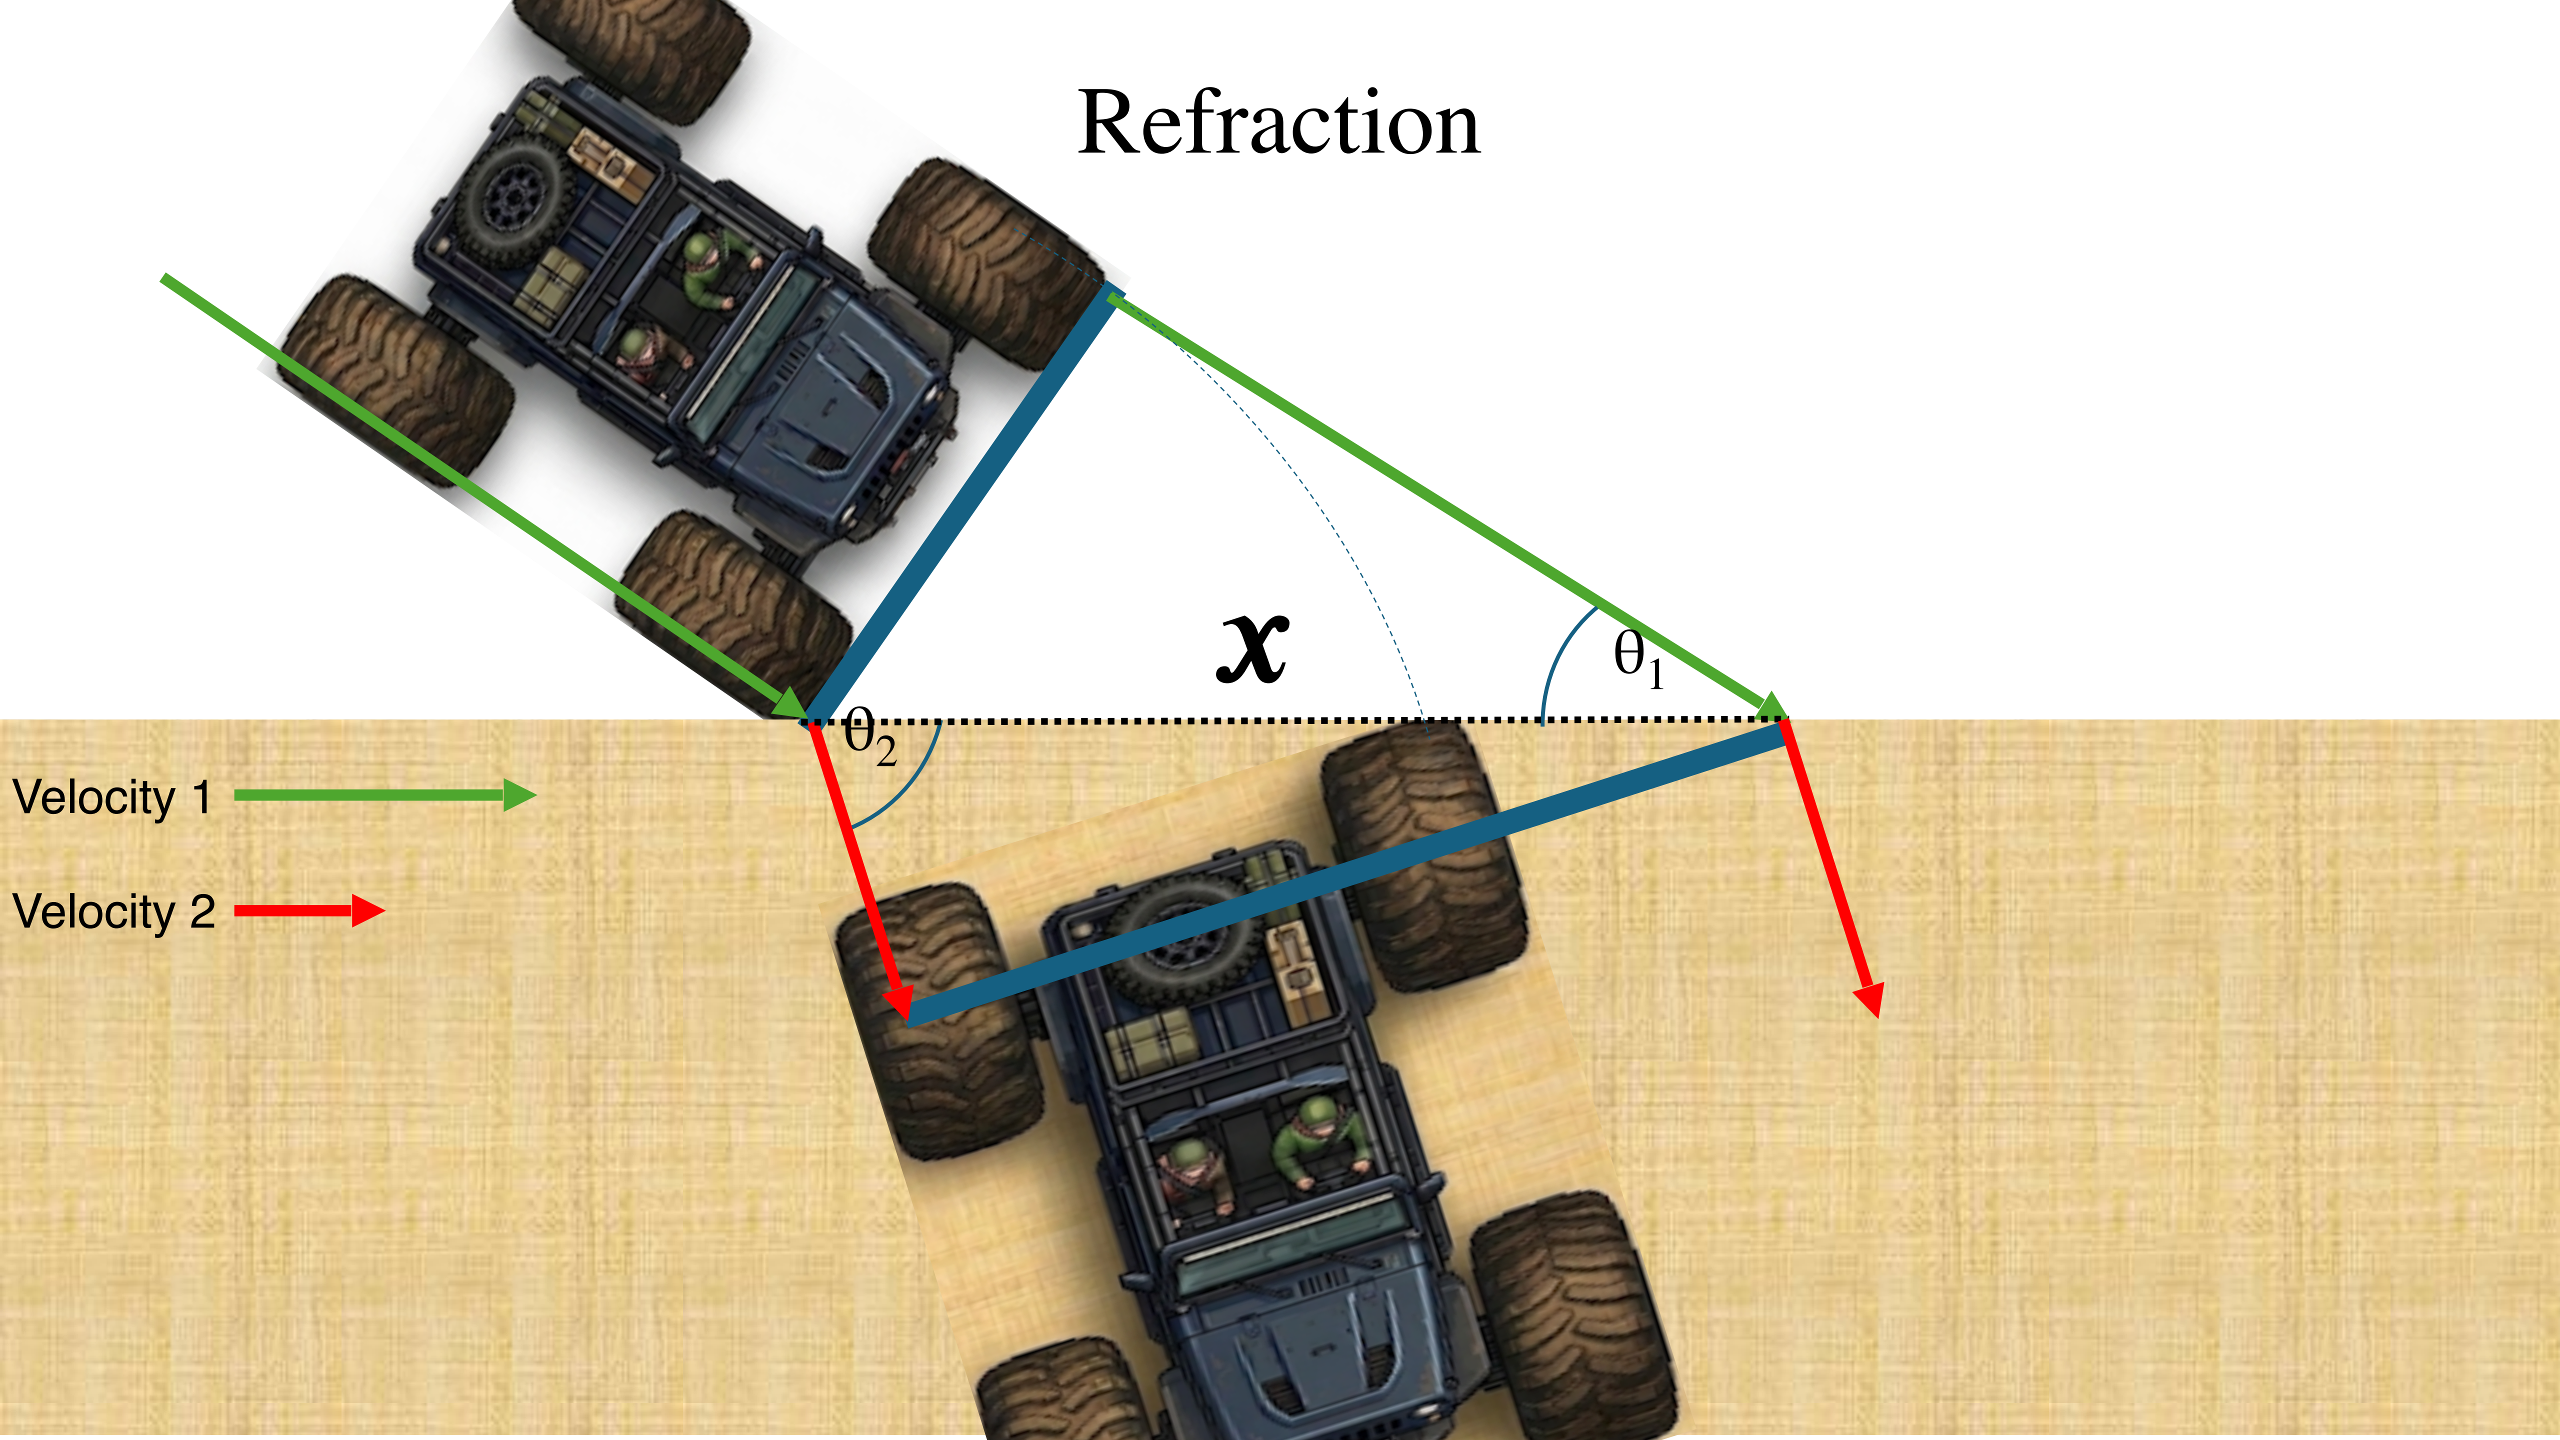

Refraction

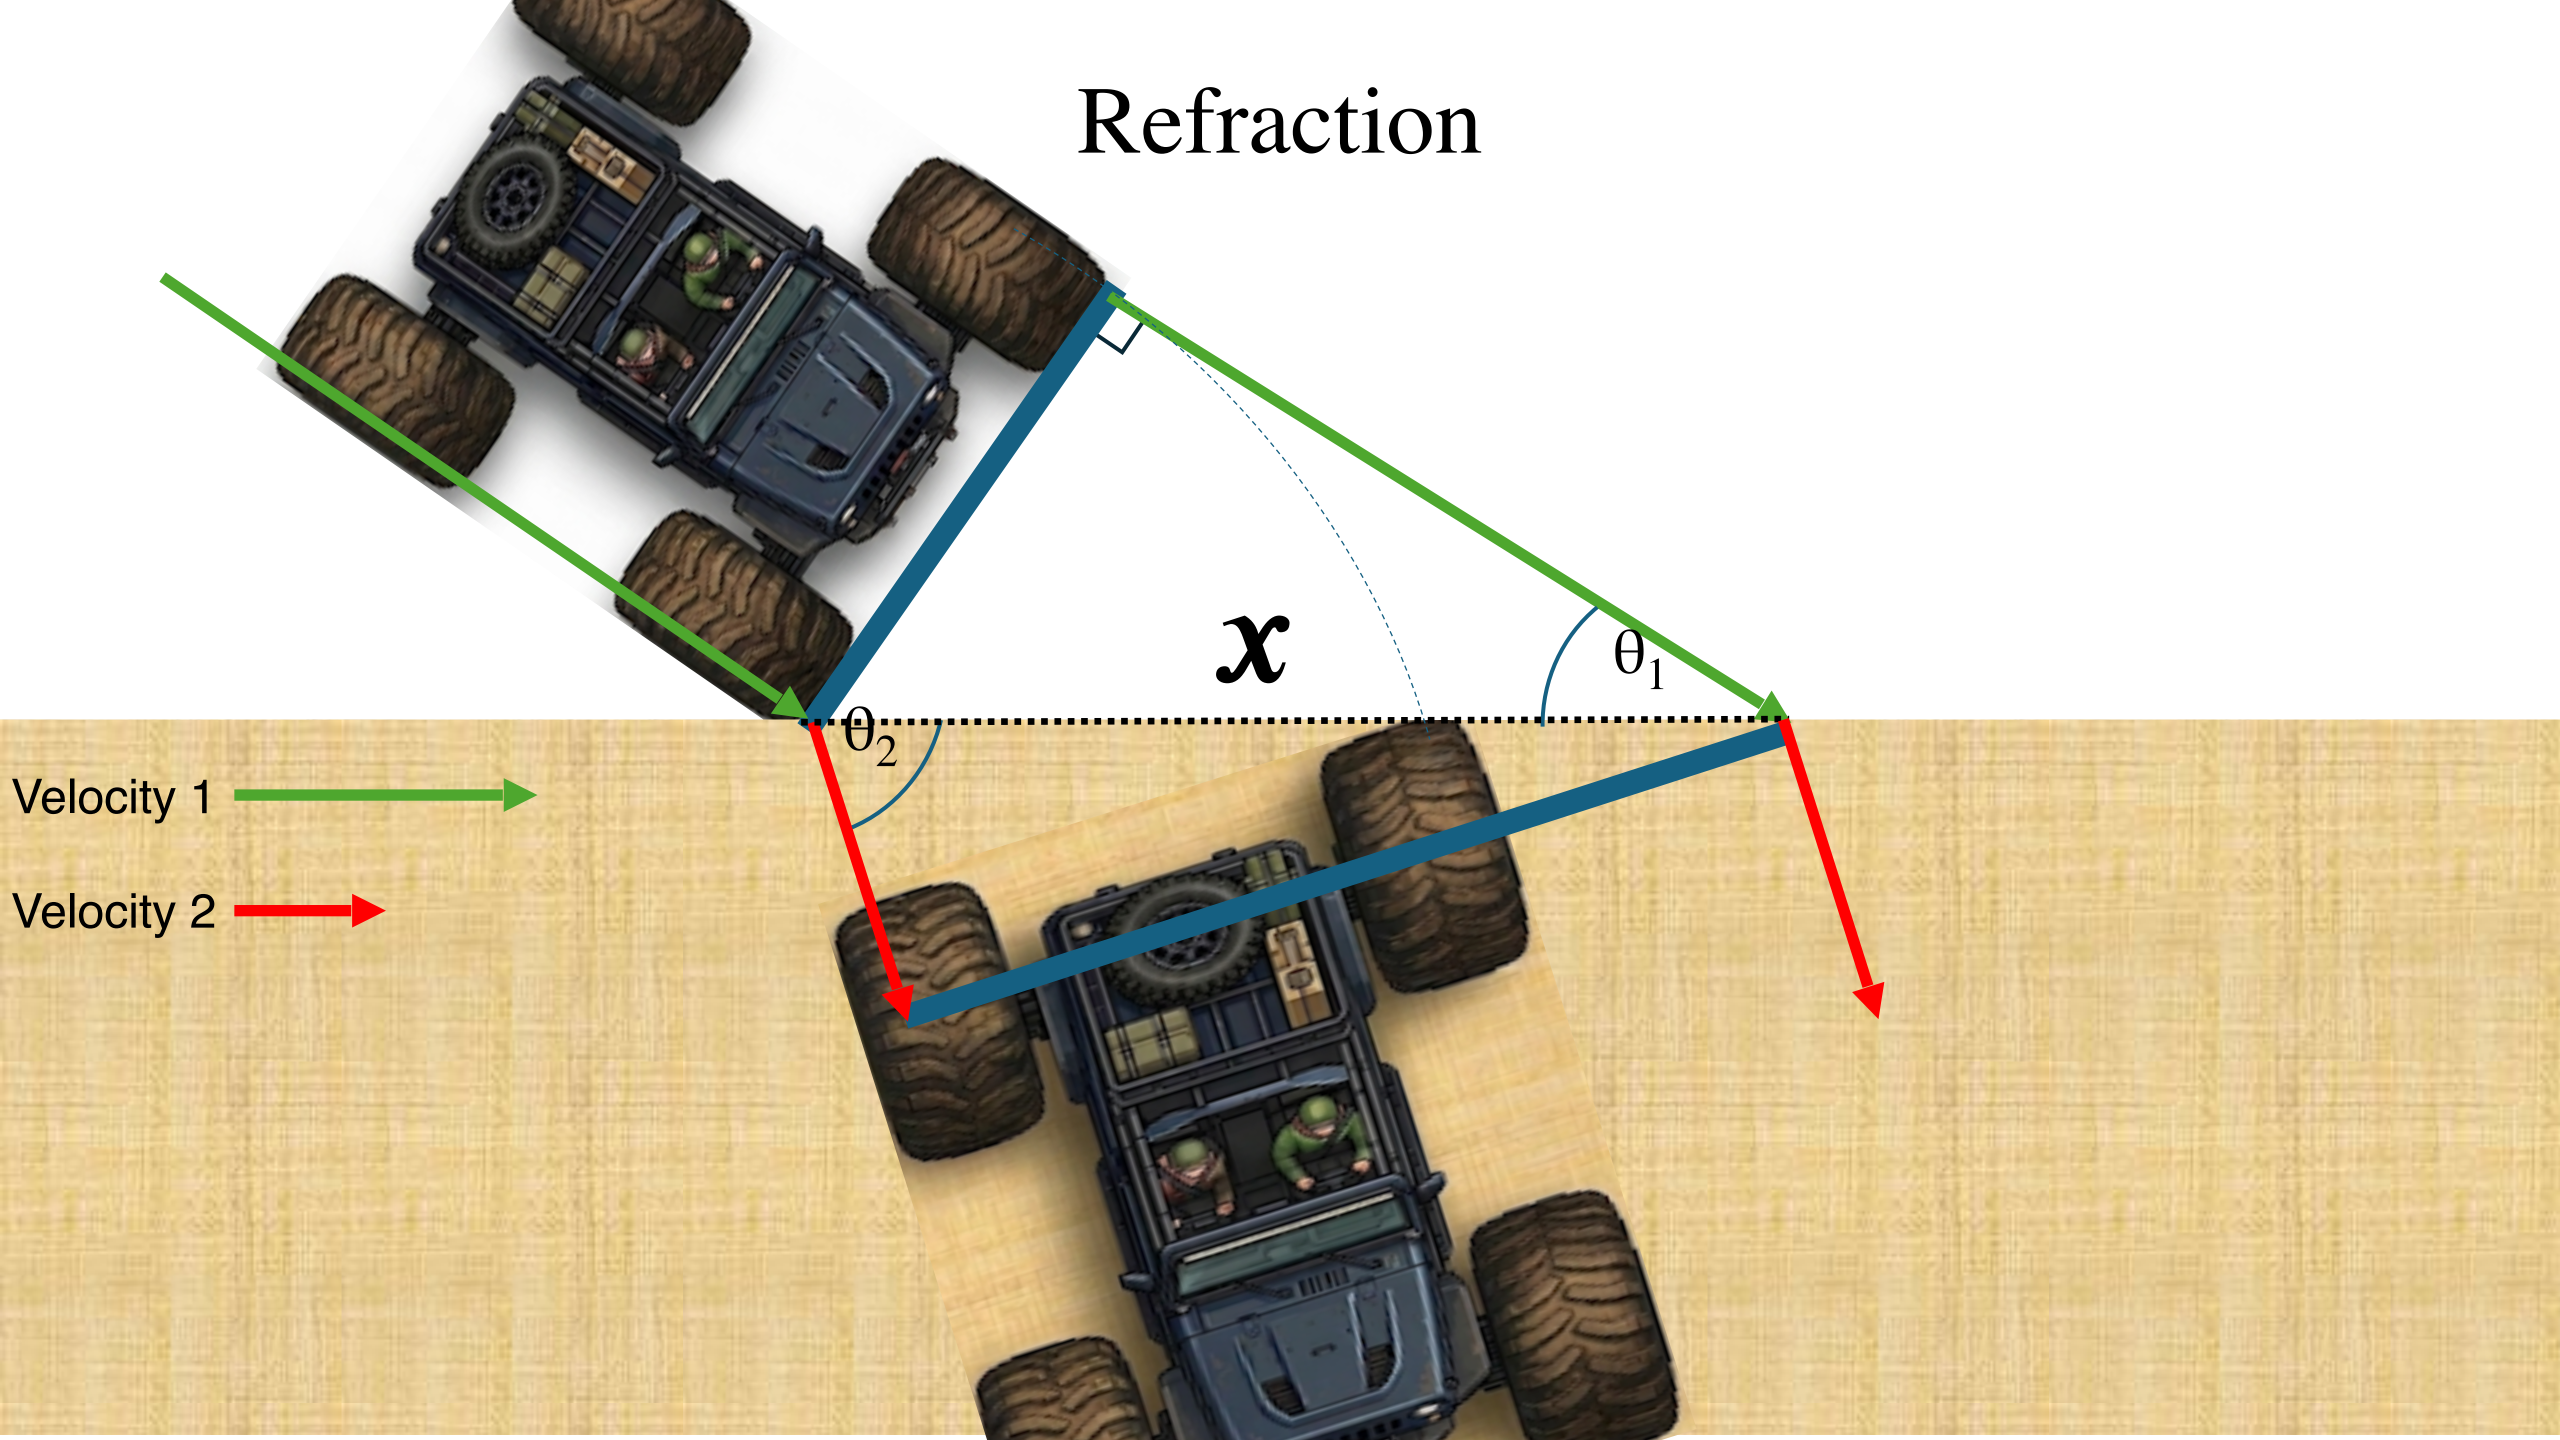

Refraction

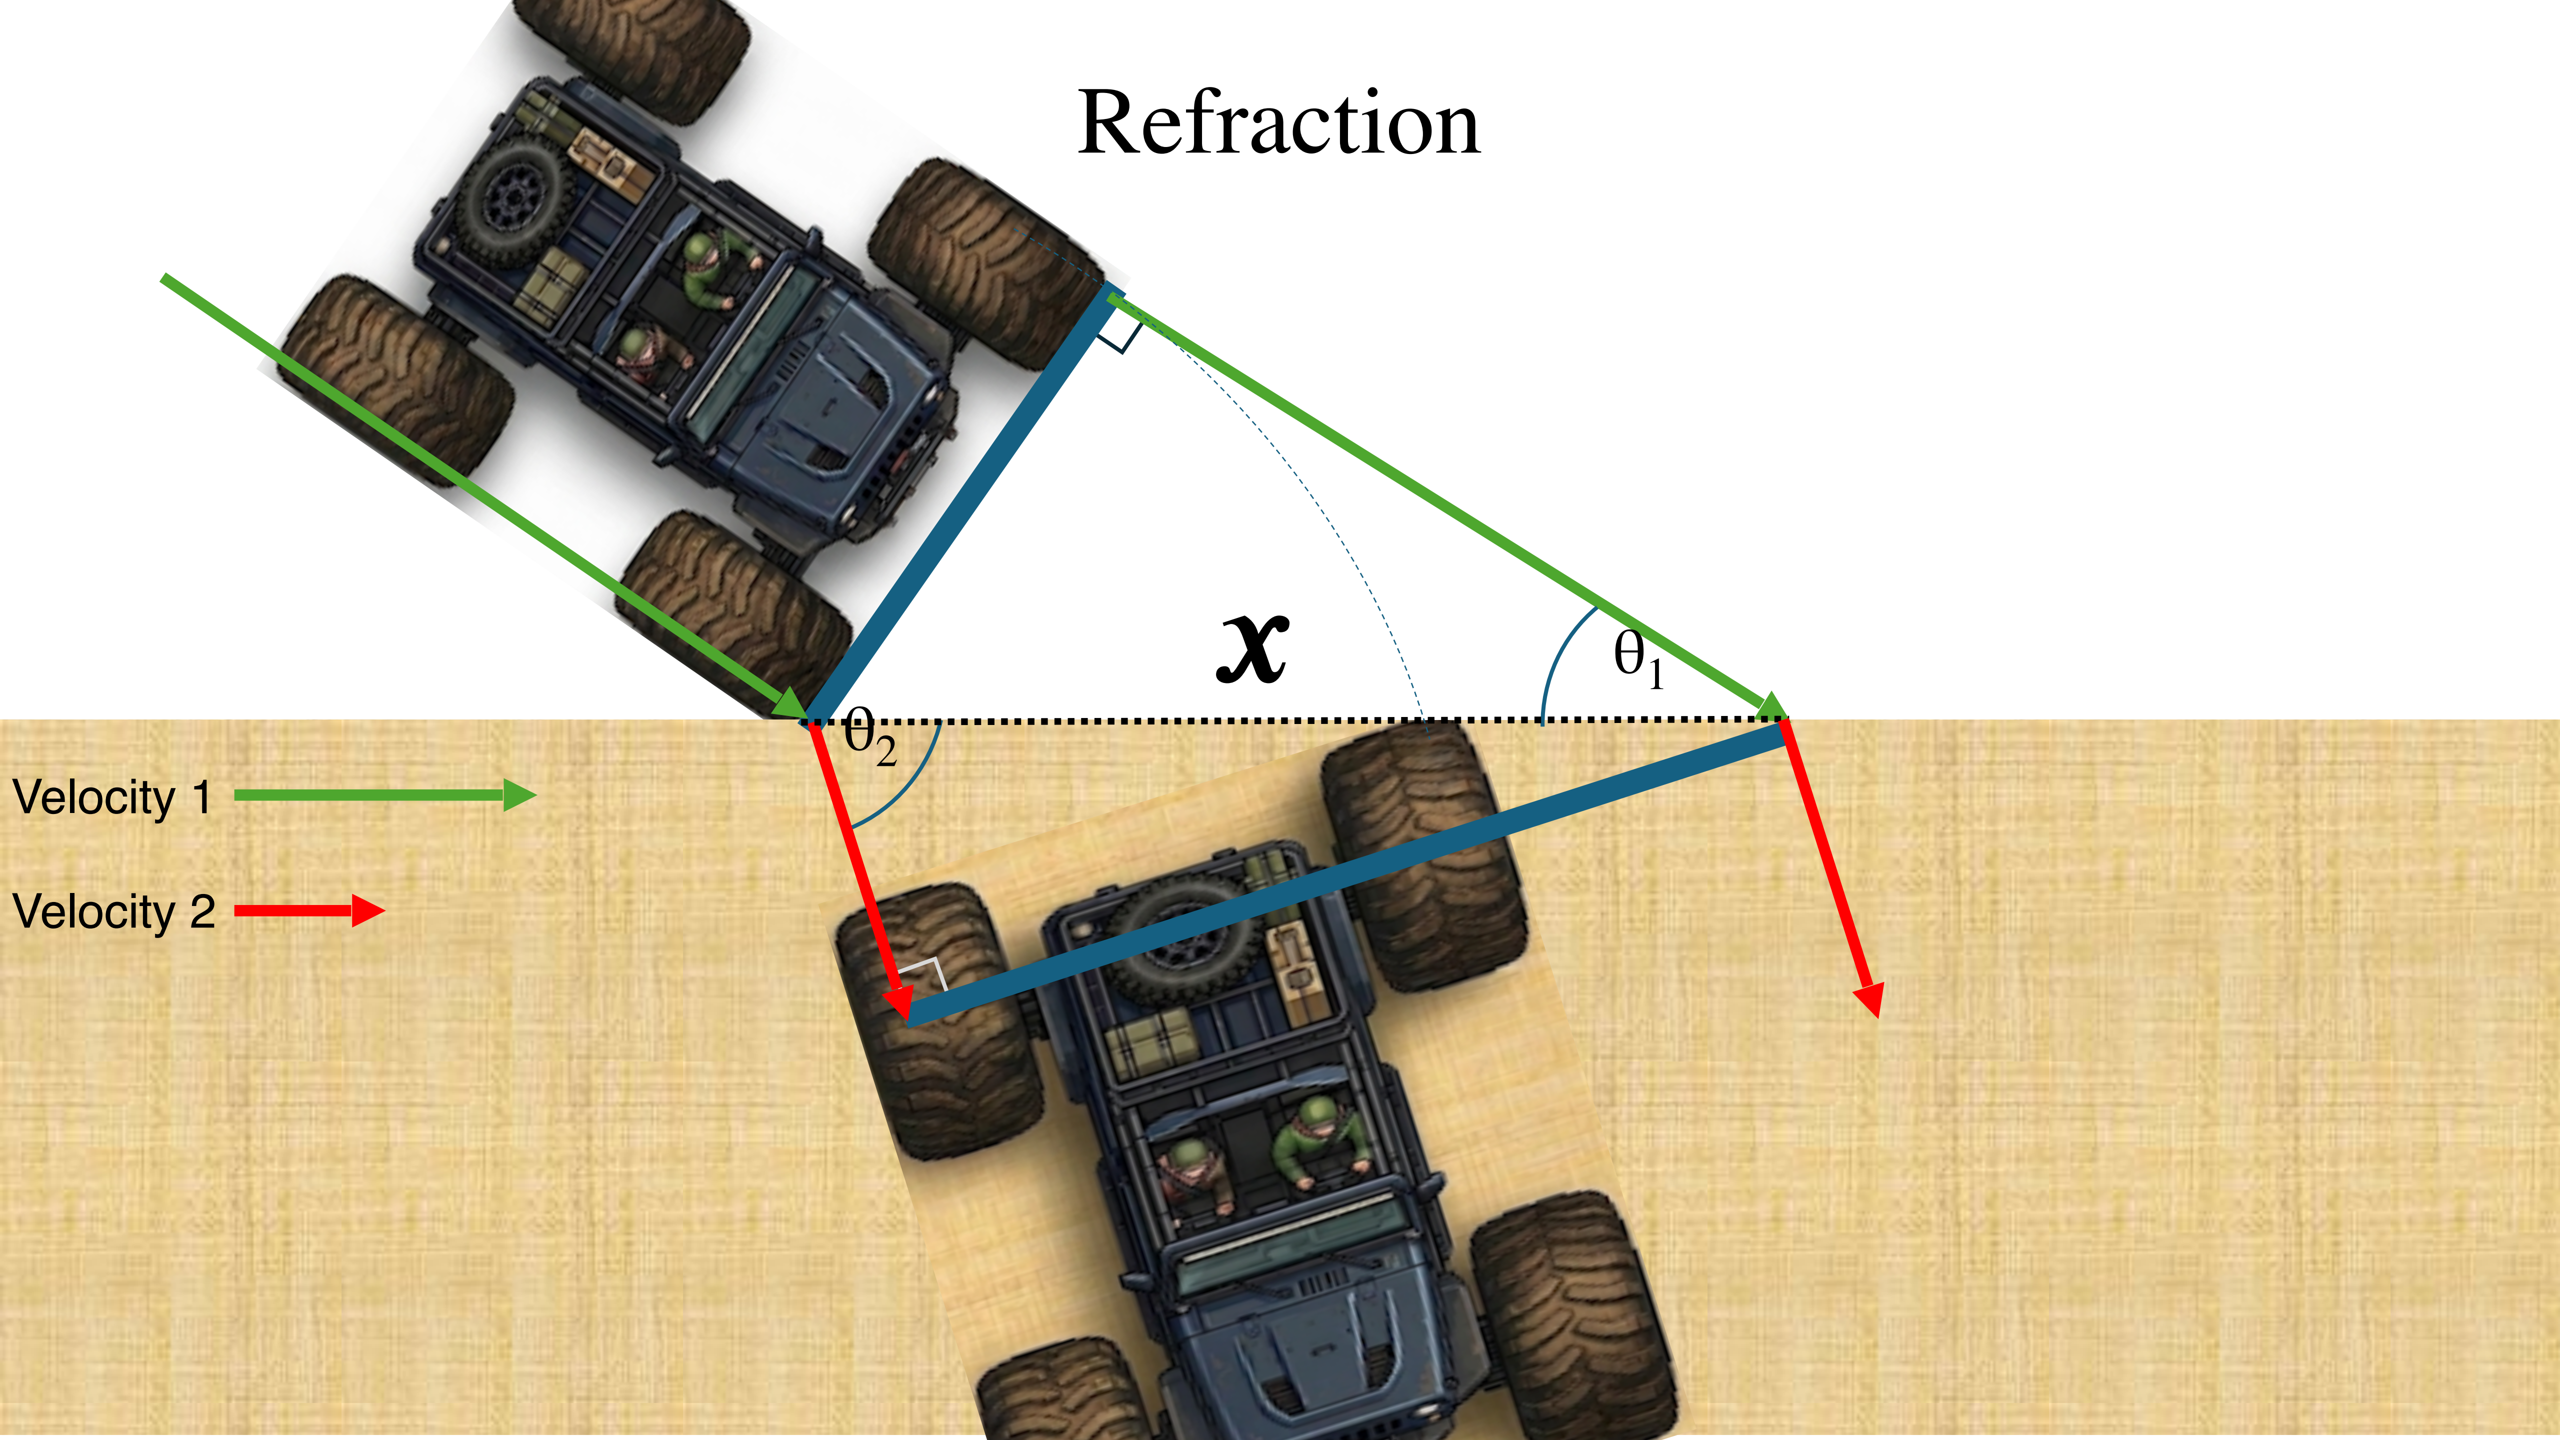

Refraction

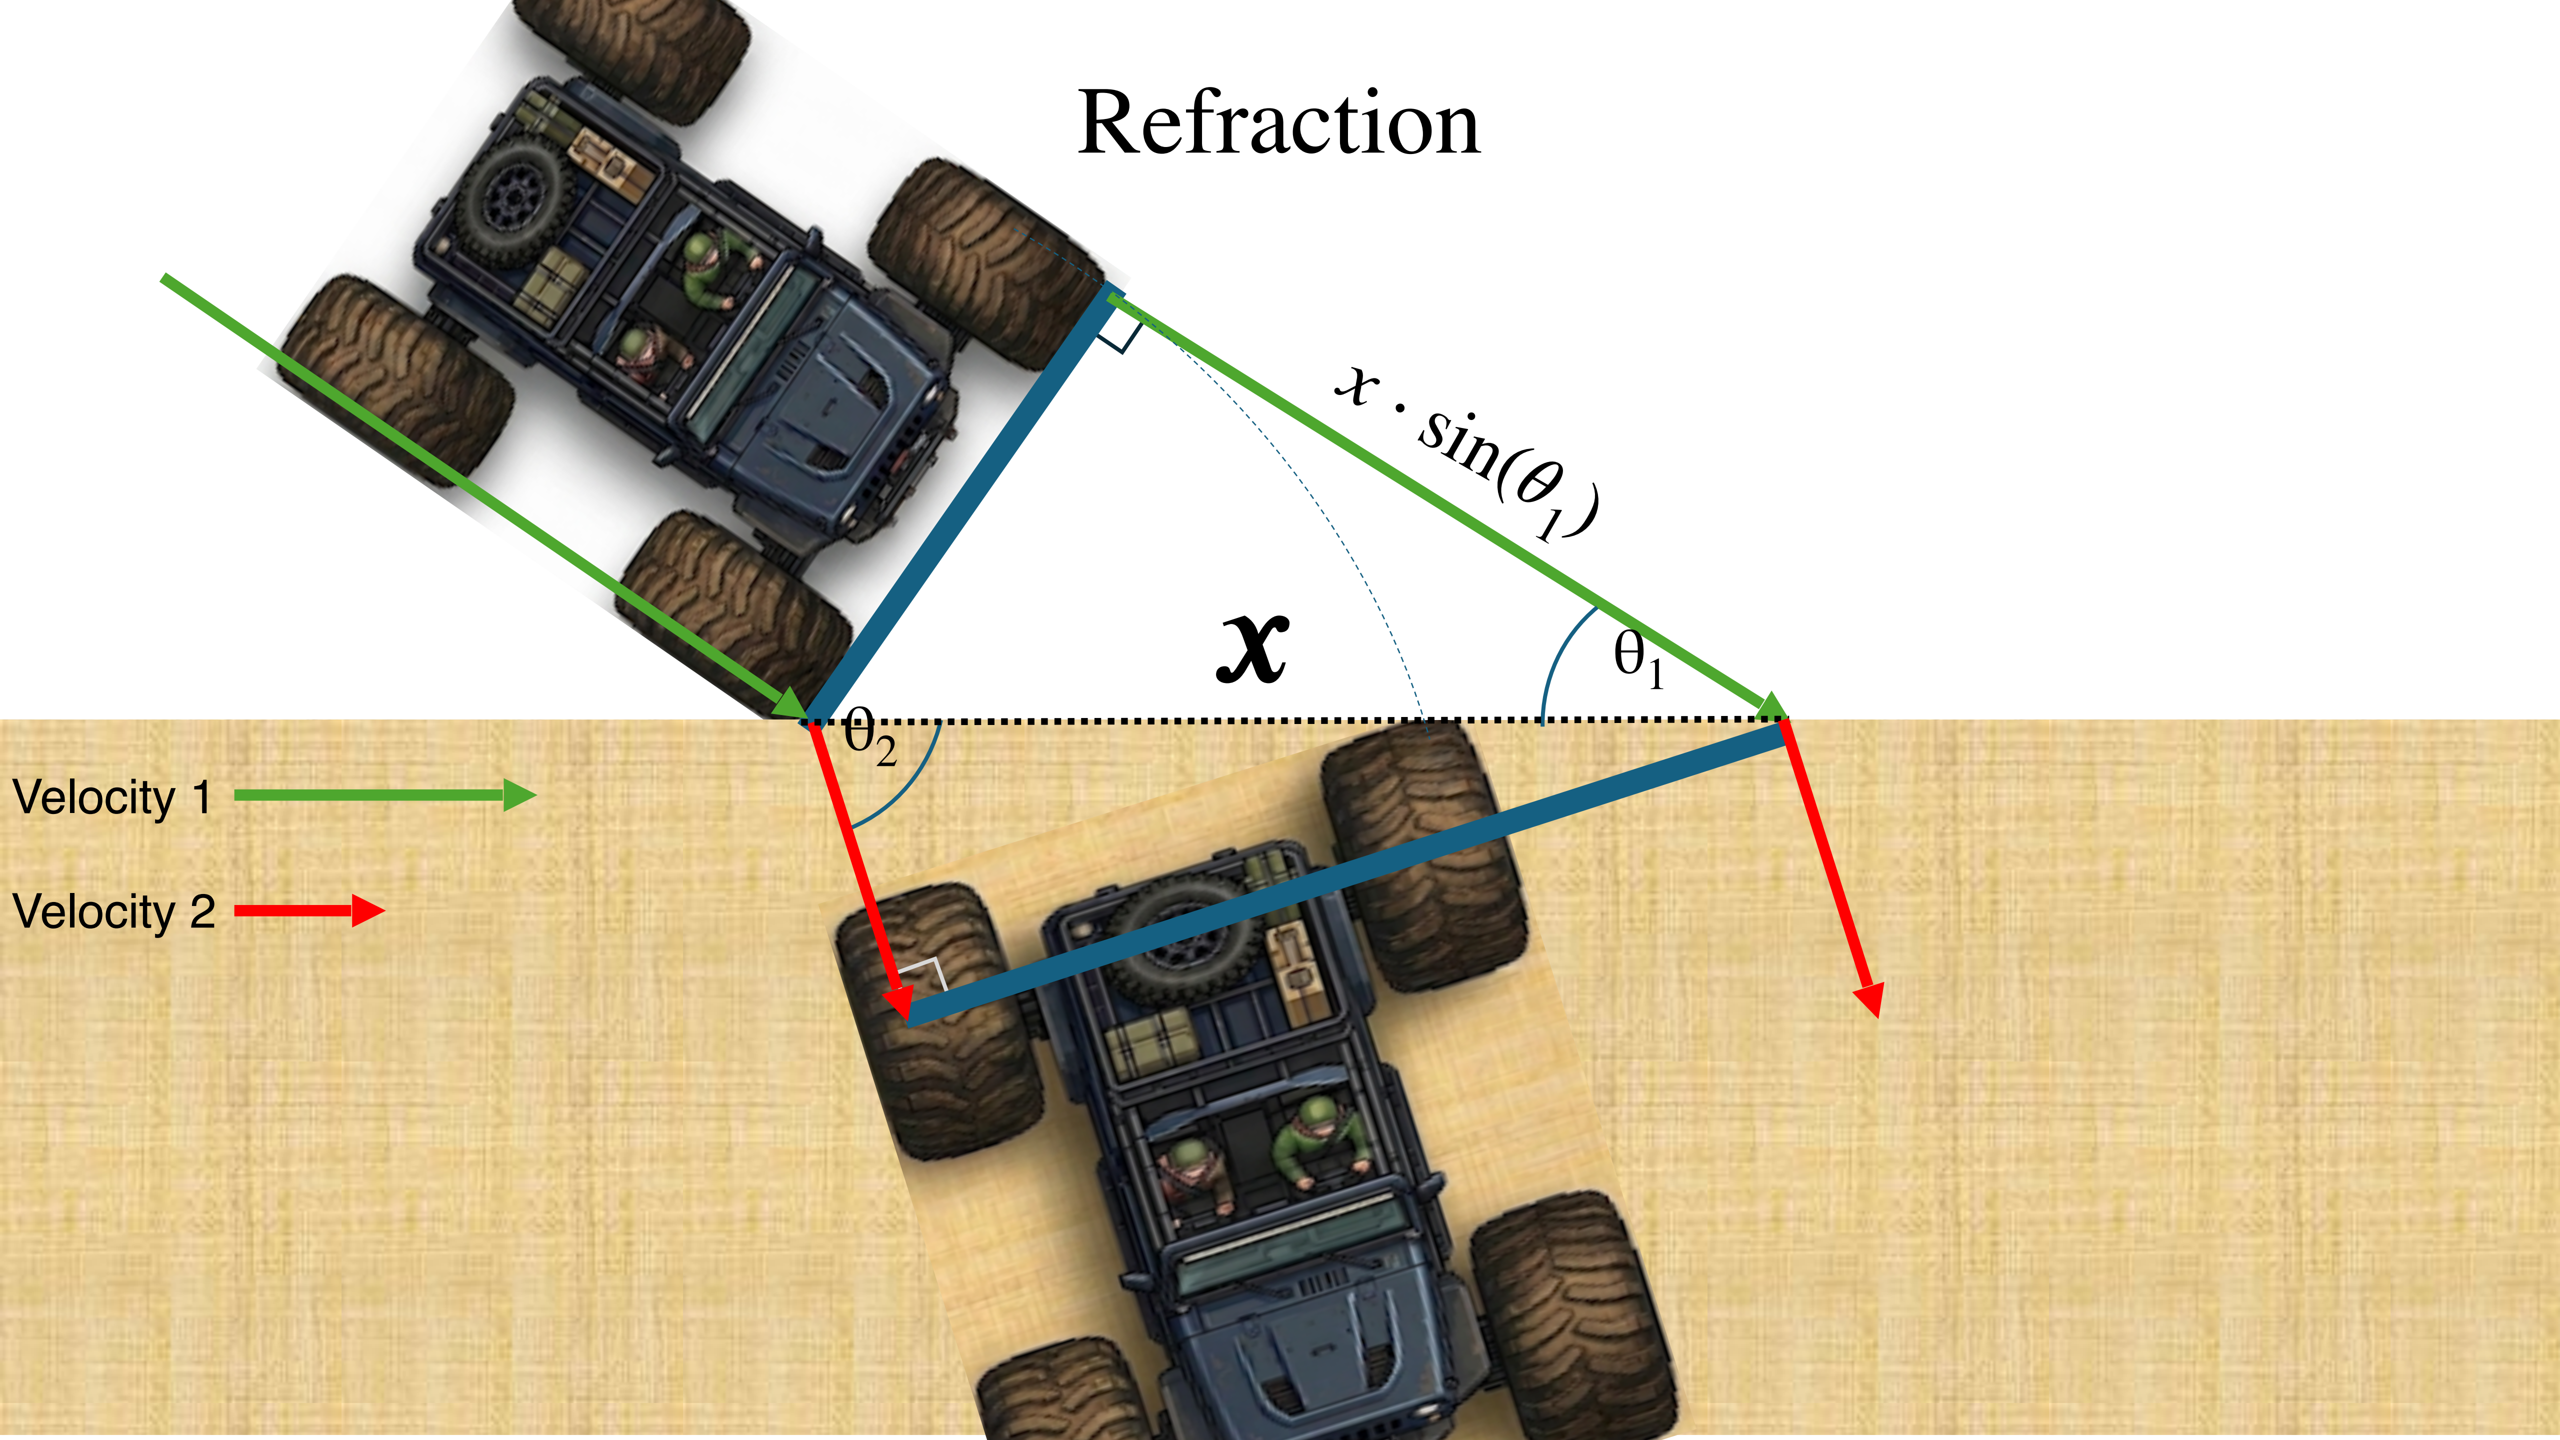

$x \cdot \sin(\theta_1)$
$\boldsymbol{x}$
θ_1
θ_2
Velocity 1
Velocity 2

Refraction

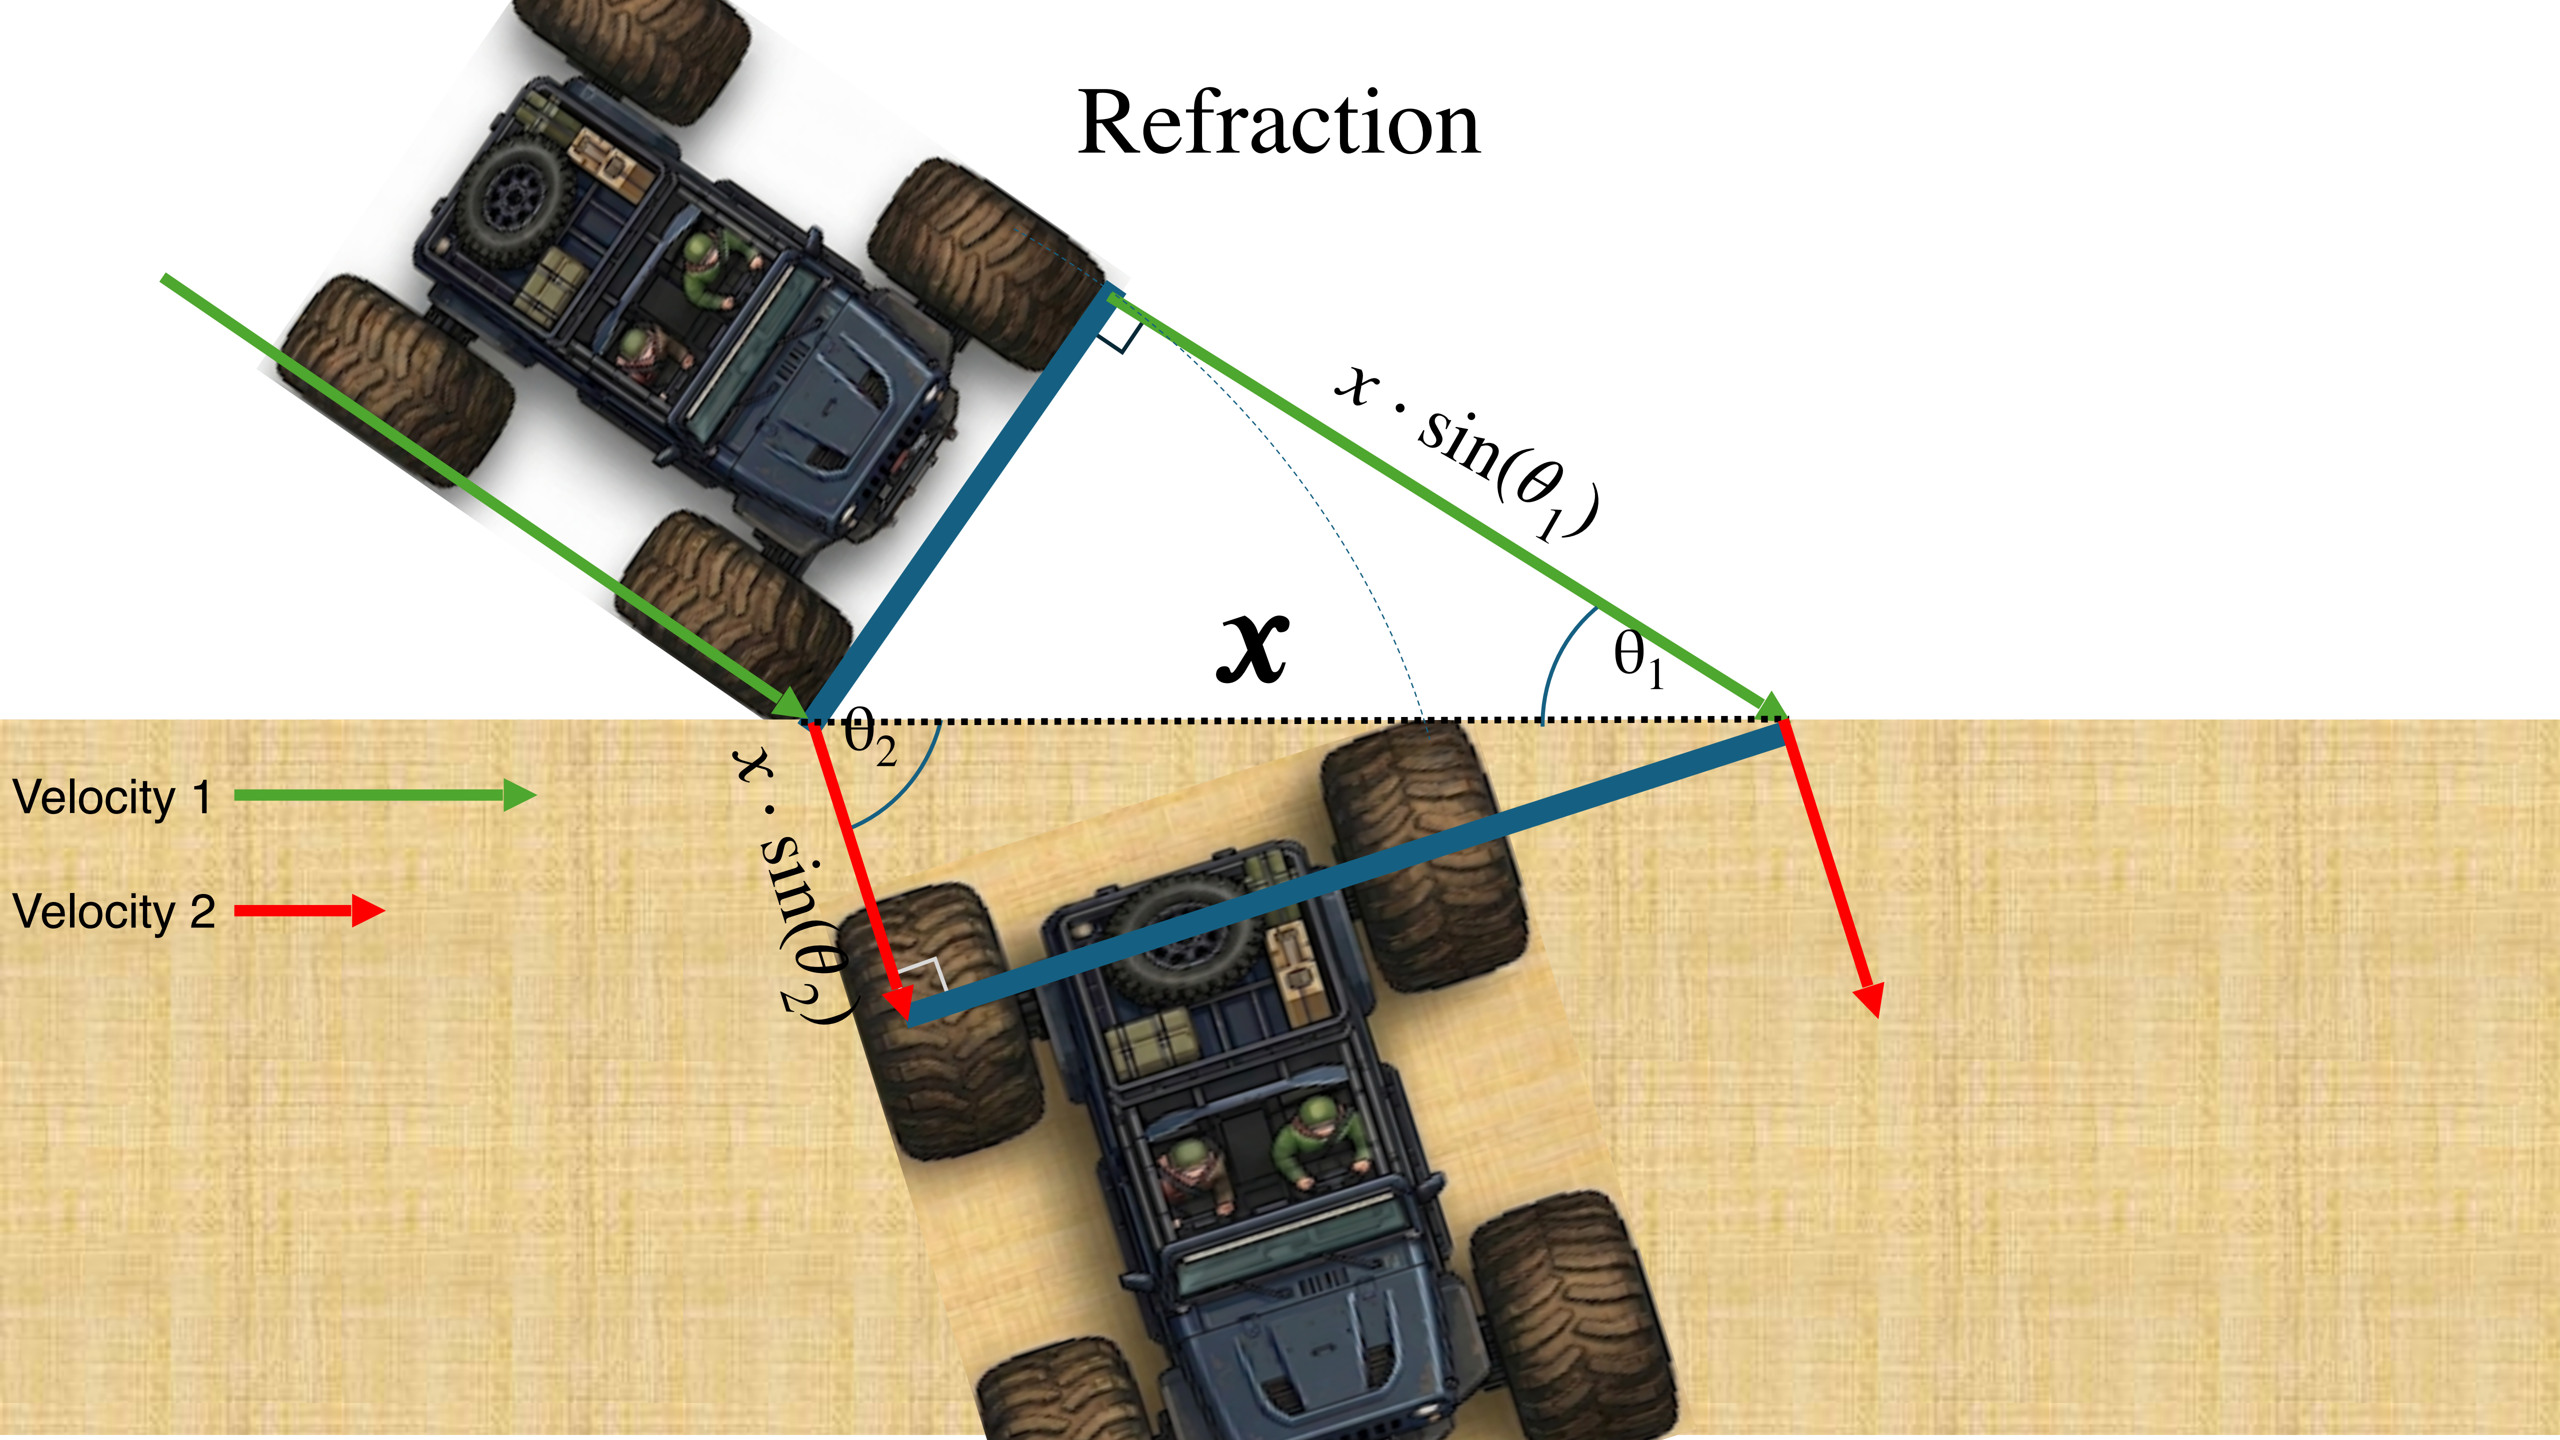

$x \cdot \sin(\theta_1)$
θ_1
x
θ_2
$x \cdot \sin(\theta_2)$
Velocity 1
Velocity 2

Refraction

$$\frac{x \cdot \sin(\theta_1)}{x \cdot \sin(\theta_2)} = \frac{v_1}{v_2} = \frac{n_2}{n_1}$$

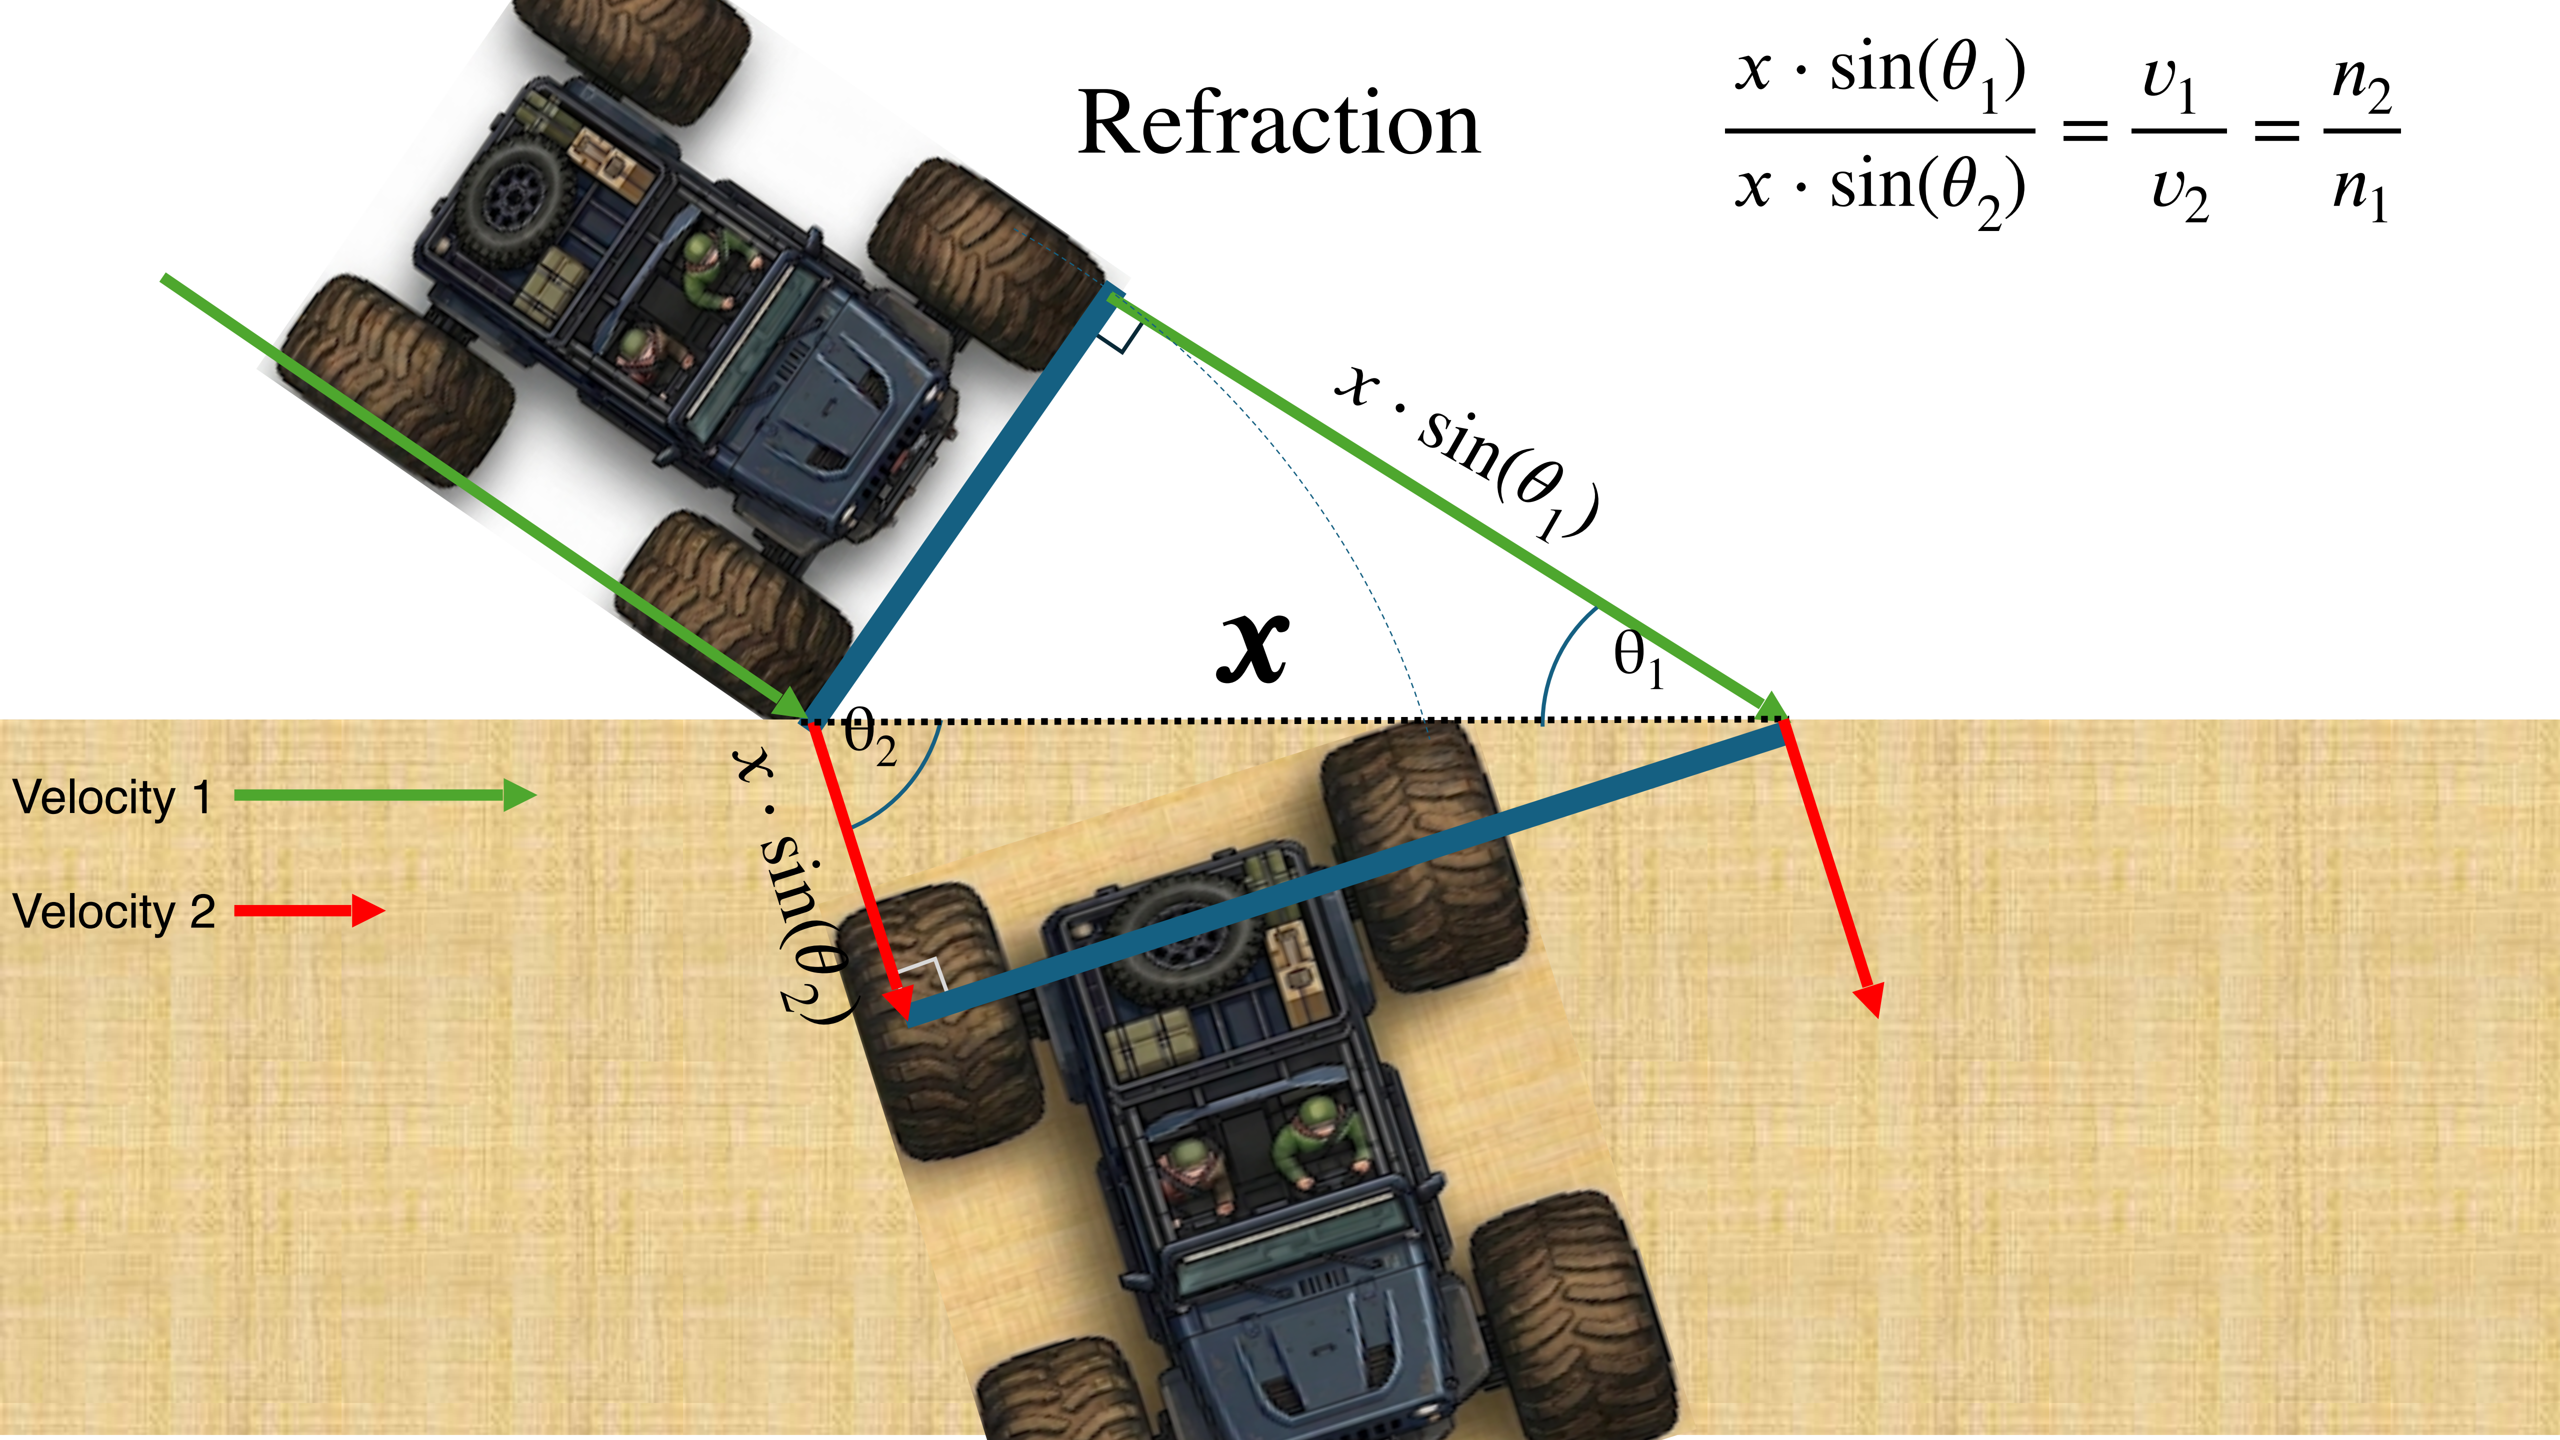

Refraction

$$\frac{x \cdot \sin(\theta_1)}{x \cdot \sin(\theta_2)} = \frac{v_1}{v_2} = \frac{n_2}{n_1}$$

$$n_1 \sin(\theta_1) = n_2 \sin(\theta_2)$$

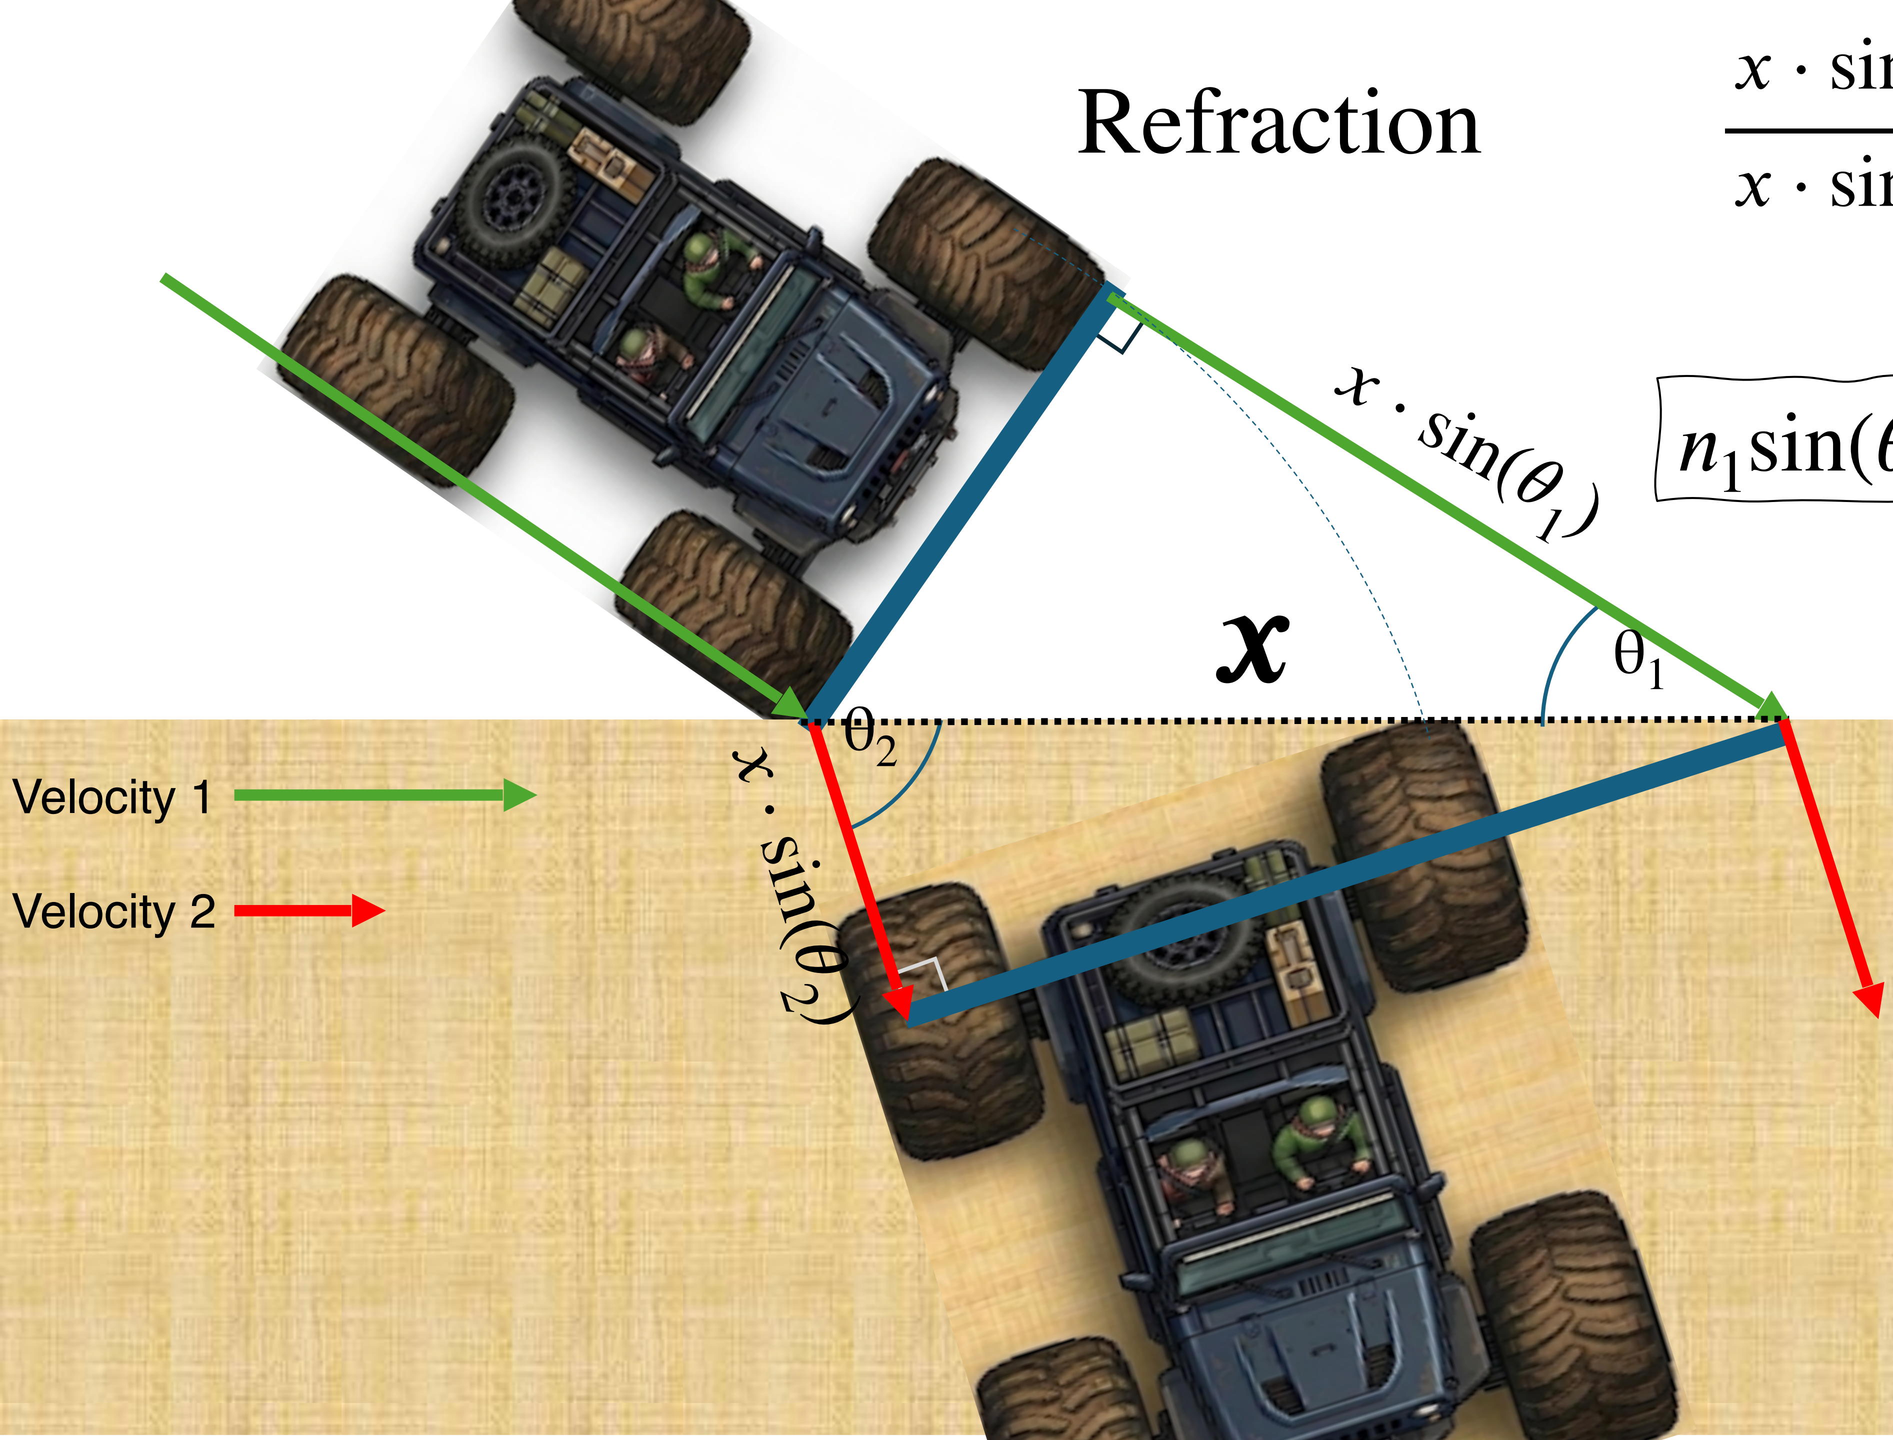

Refraction

$$\frac{x \cdot \sin(\theta_1)}{x \cdot \sin(\theta_2)} = \frac{v_1}{v_2} = \frac{n_2}{n_1}$$

$$n_1 \sin(\theta_1) = n_2 \sin(\theta_2)$$

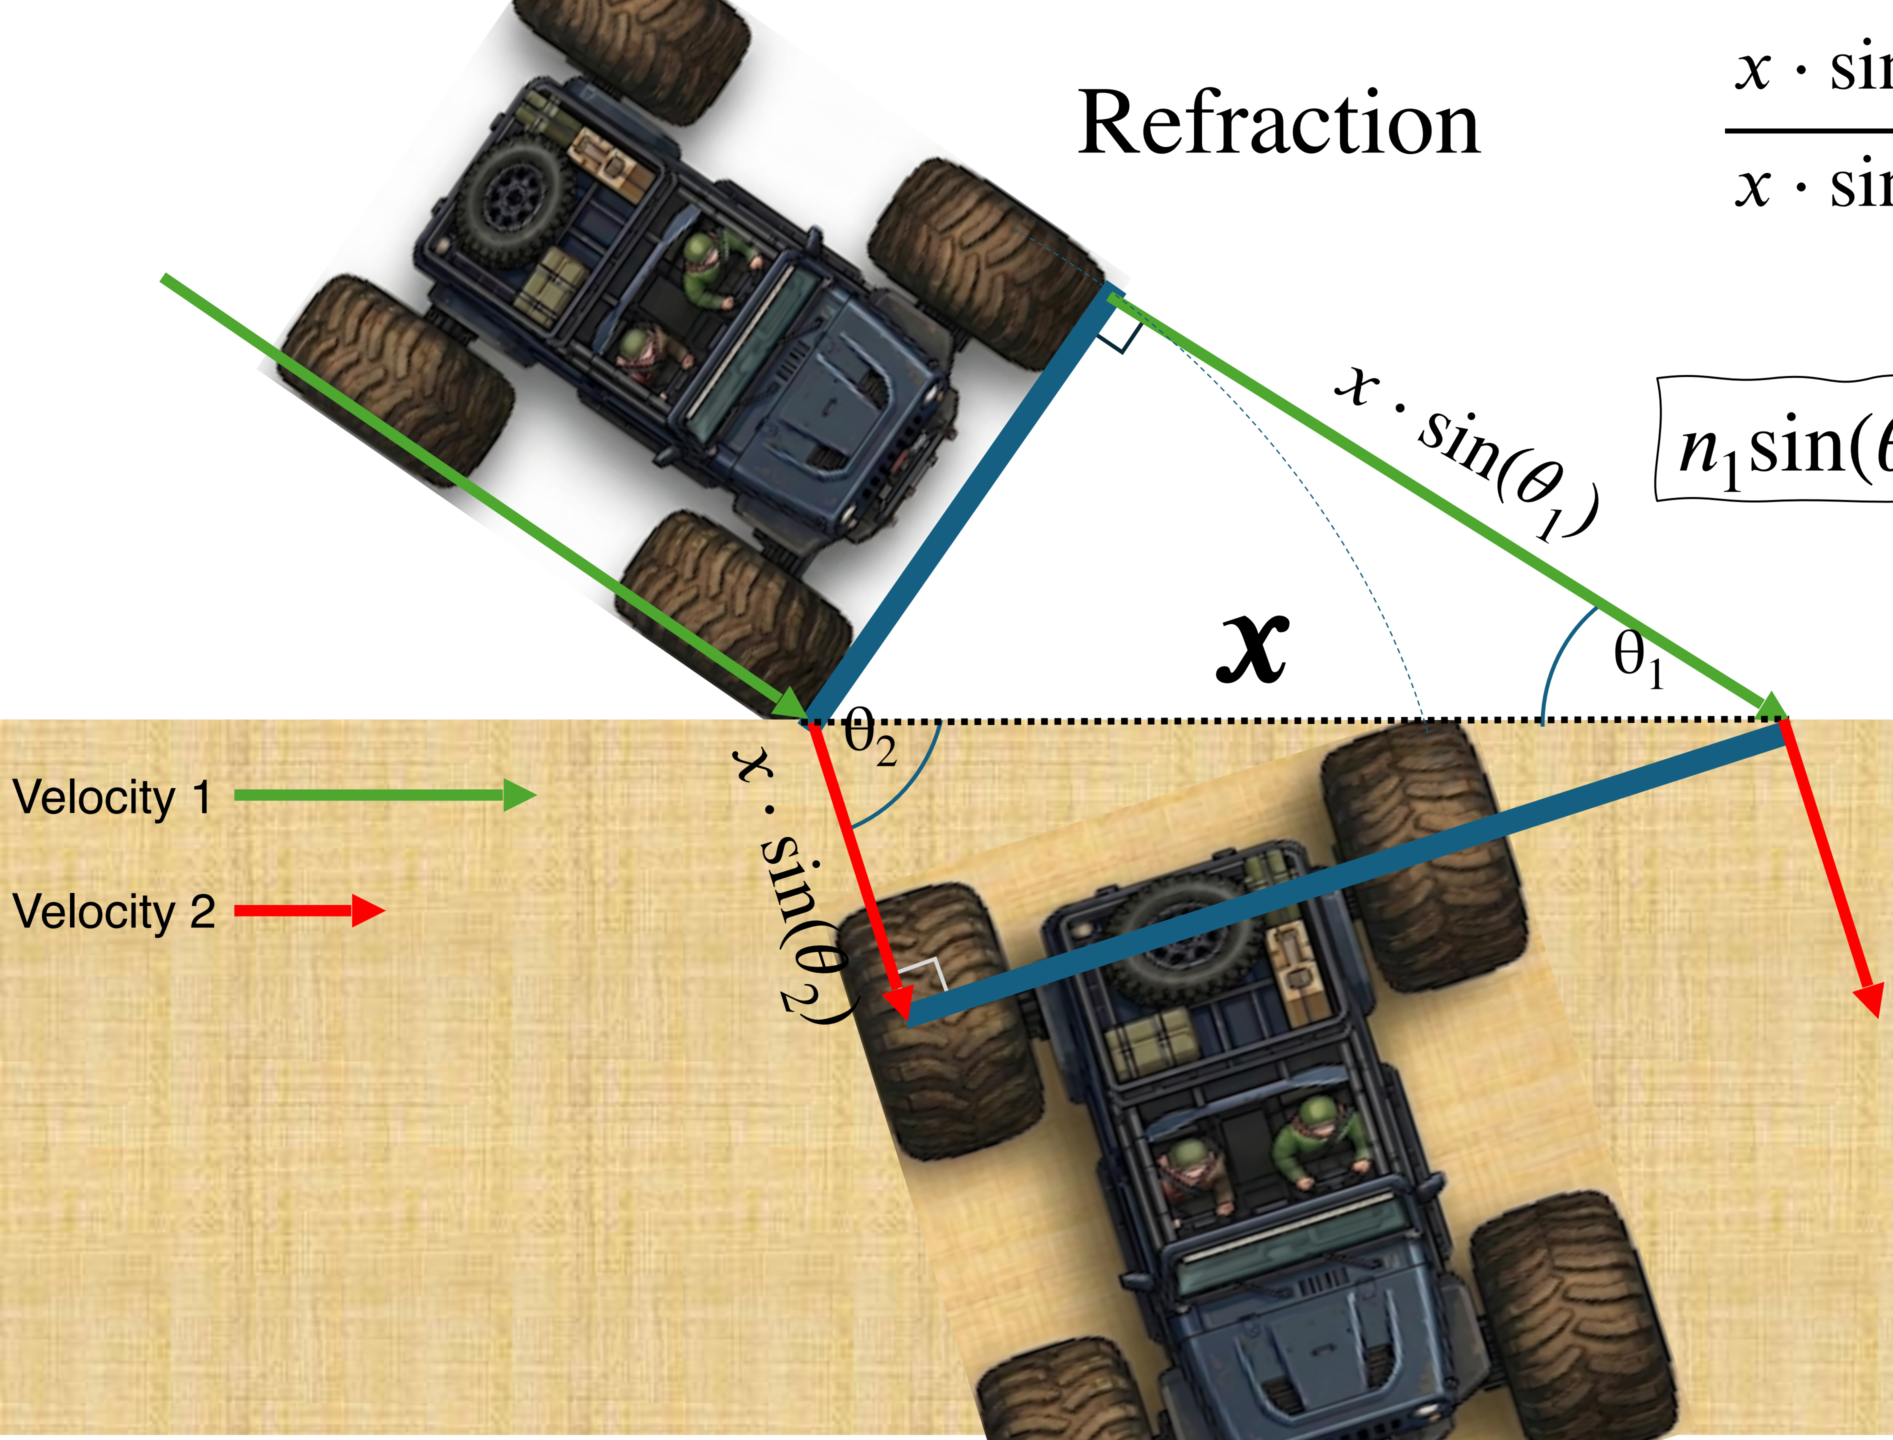

Willebrord Snellius

Thin lens

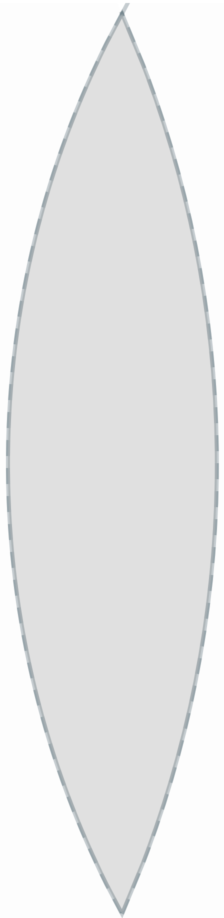

Thin lens

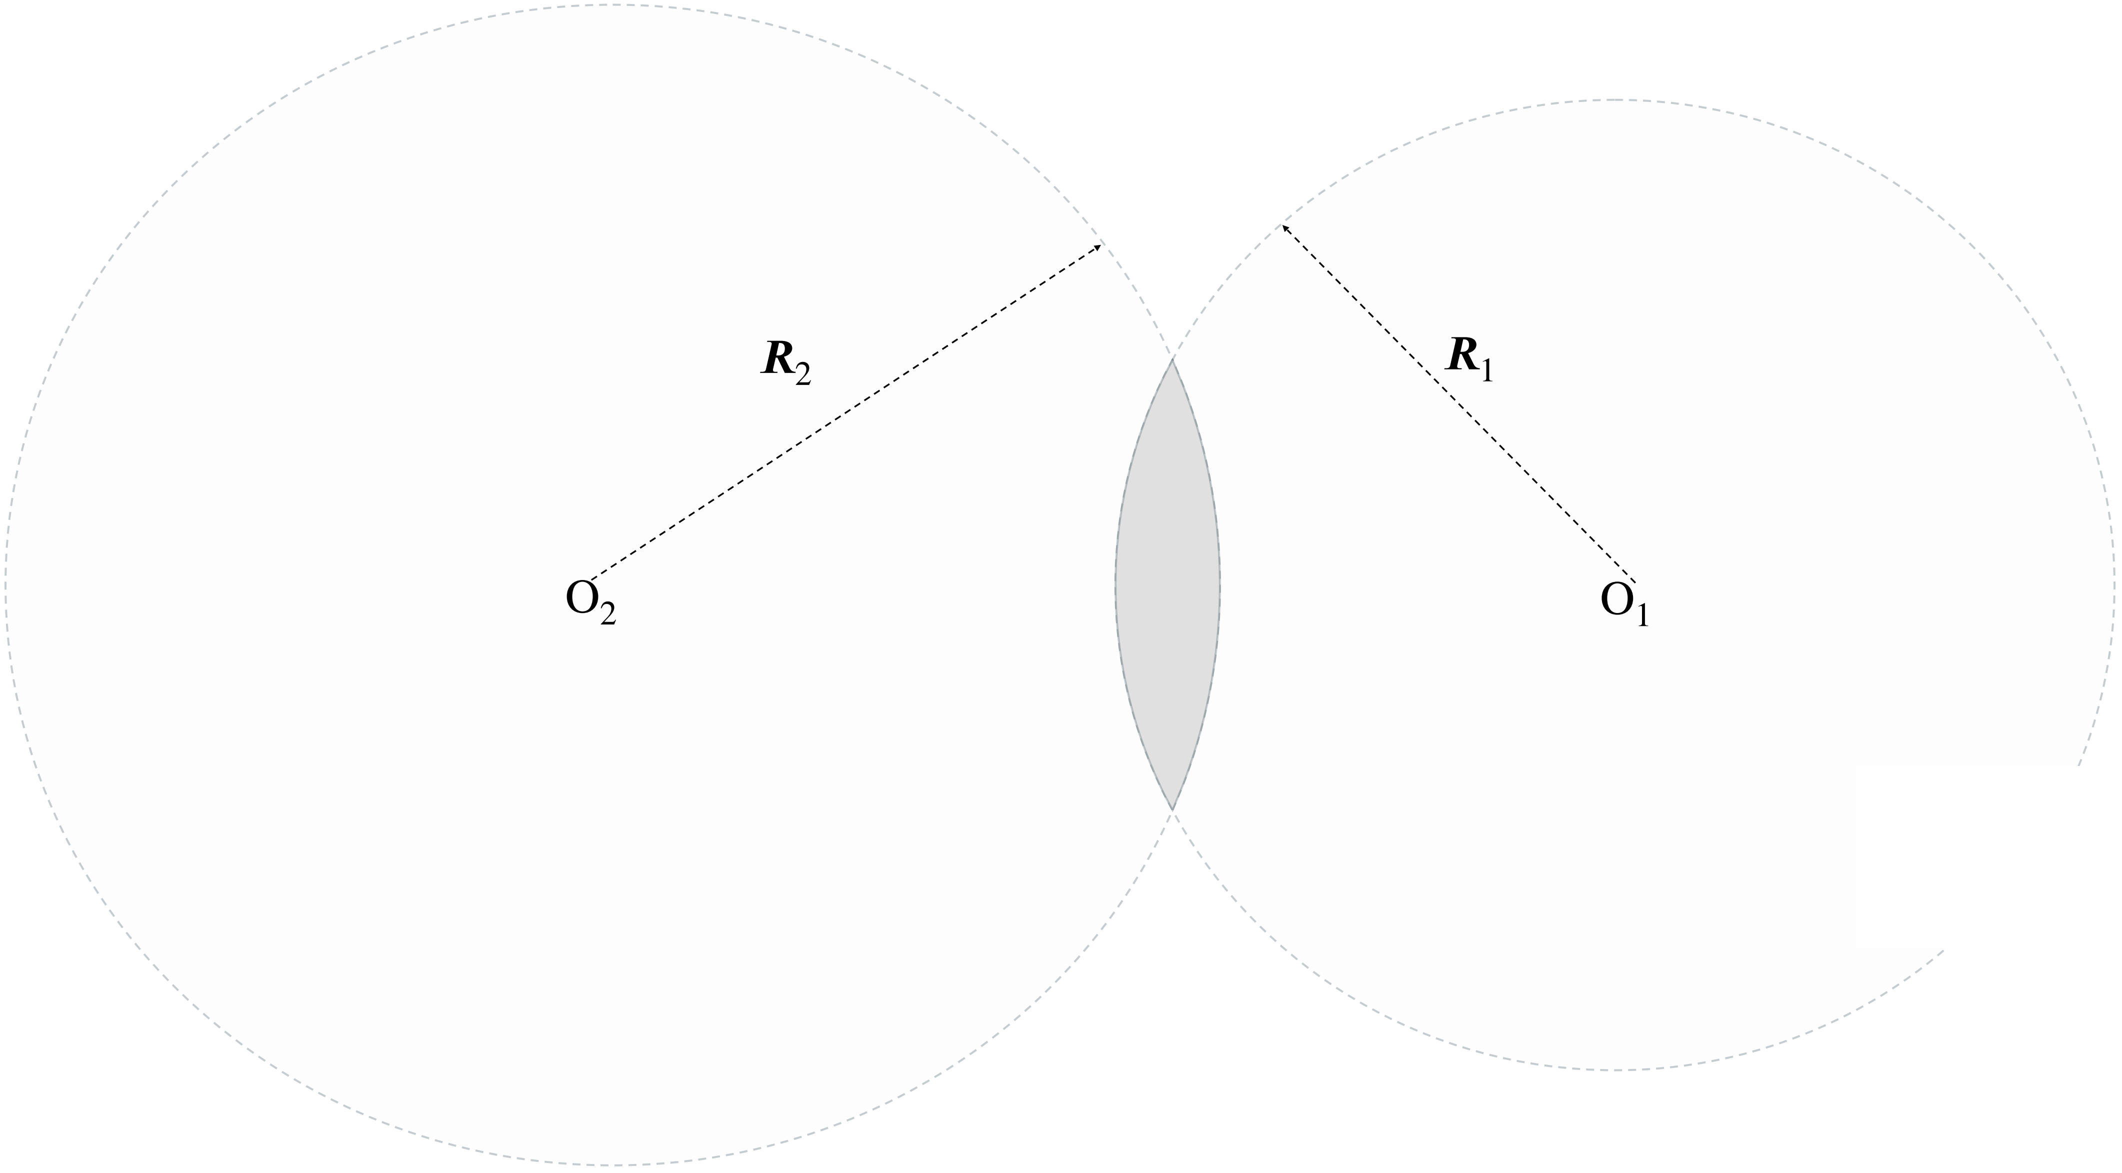

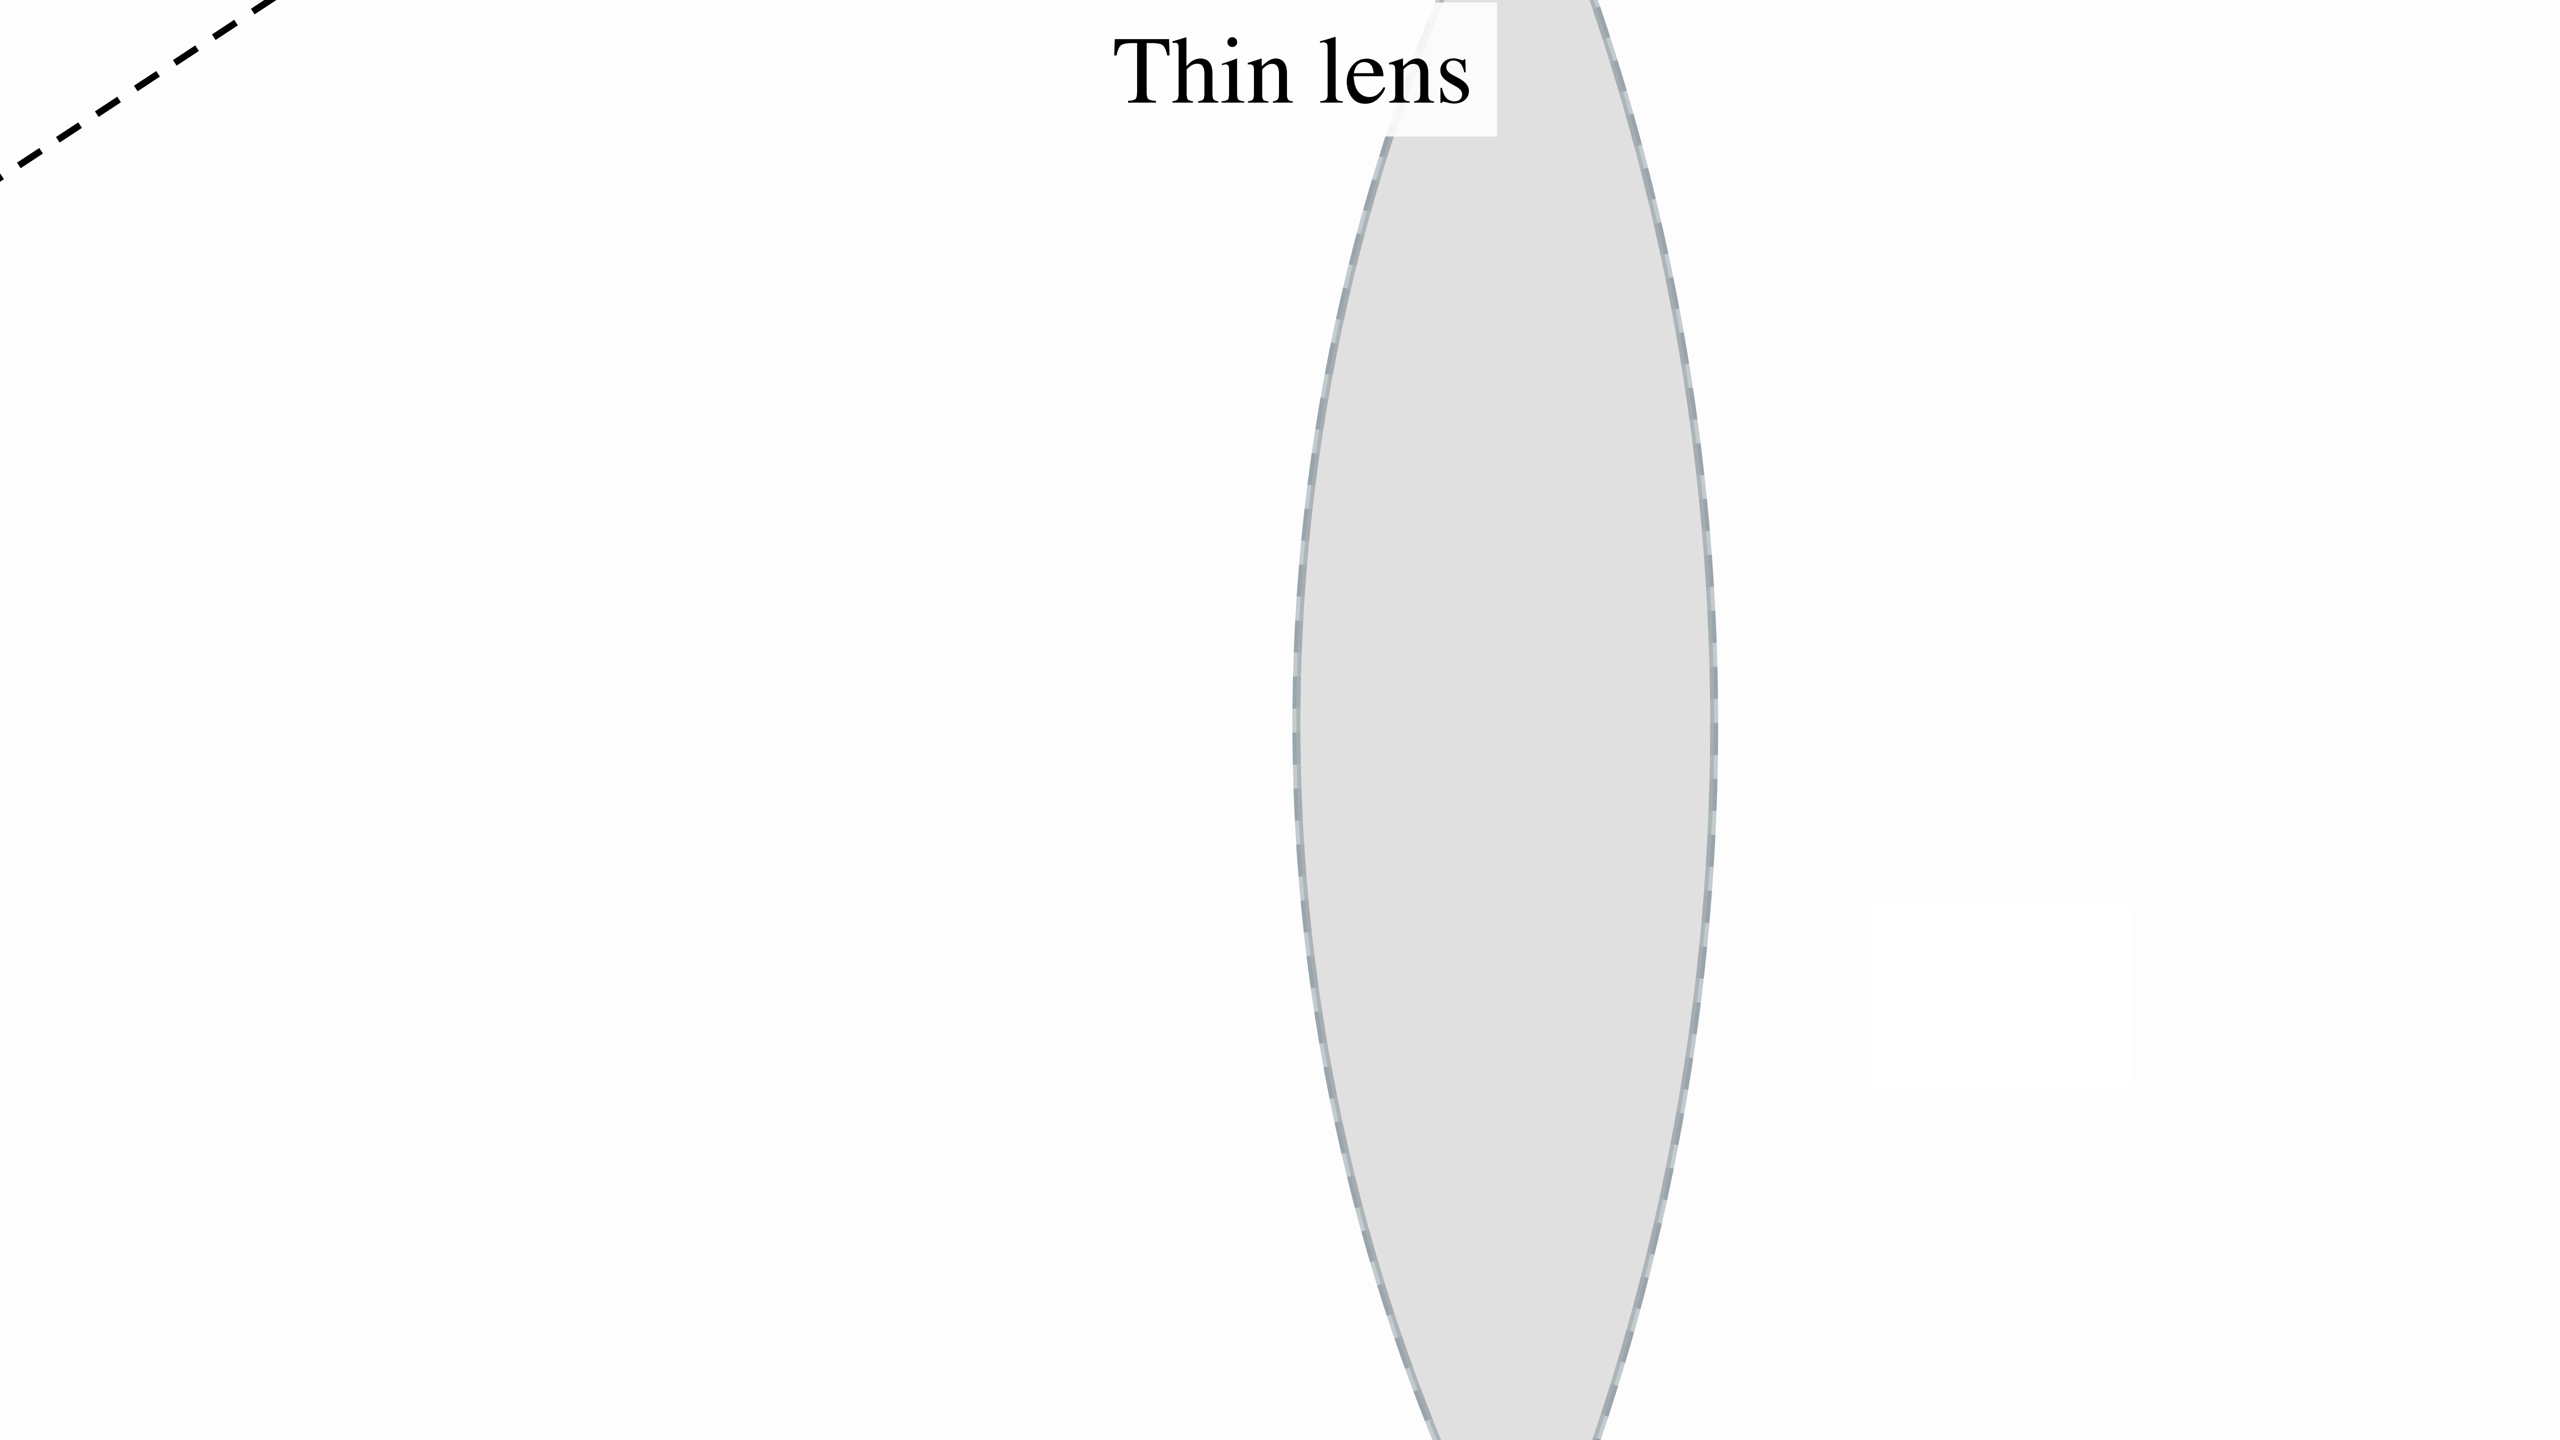
Thin lens

Thin lens

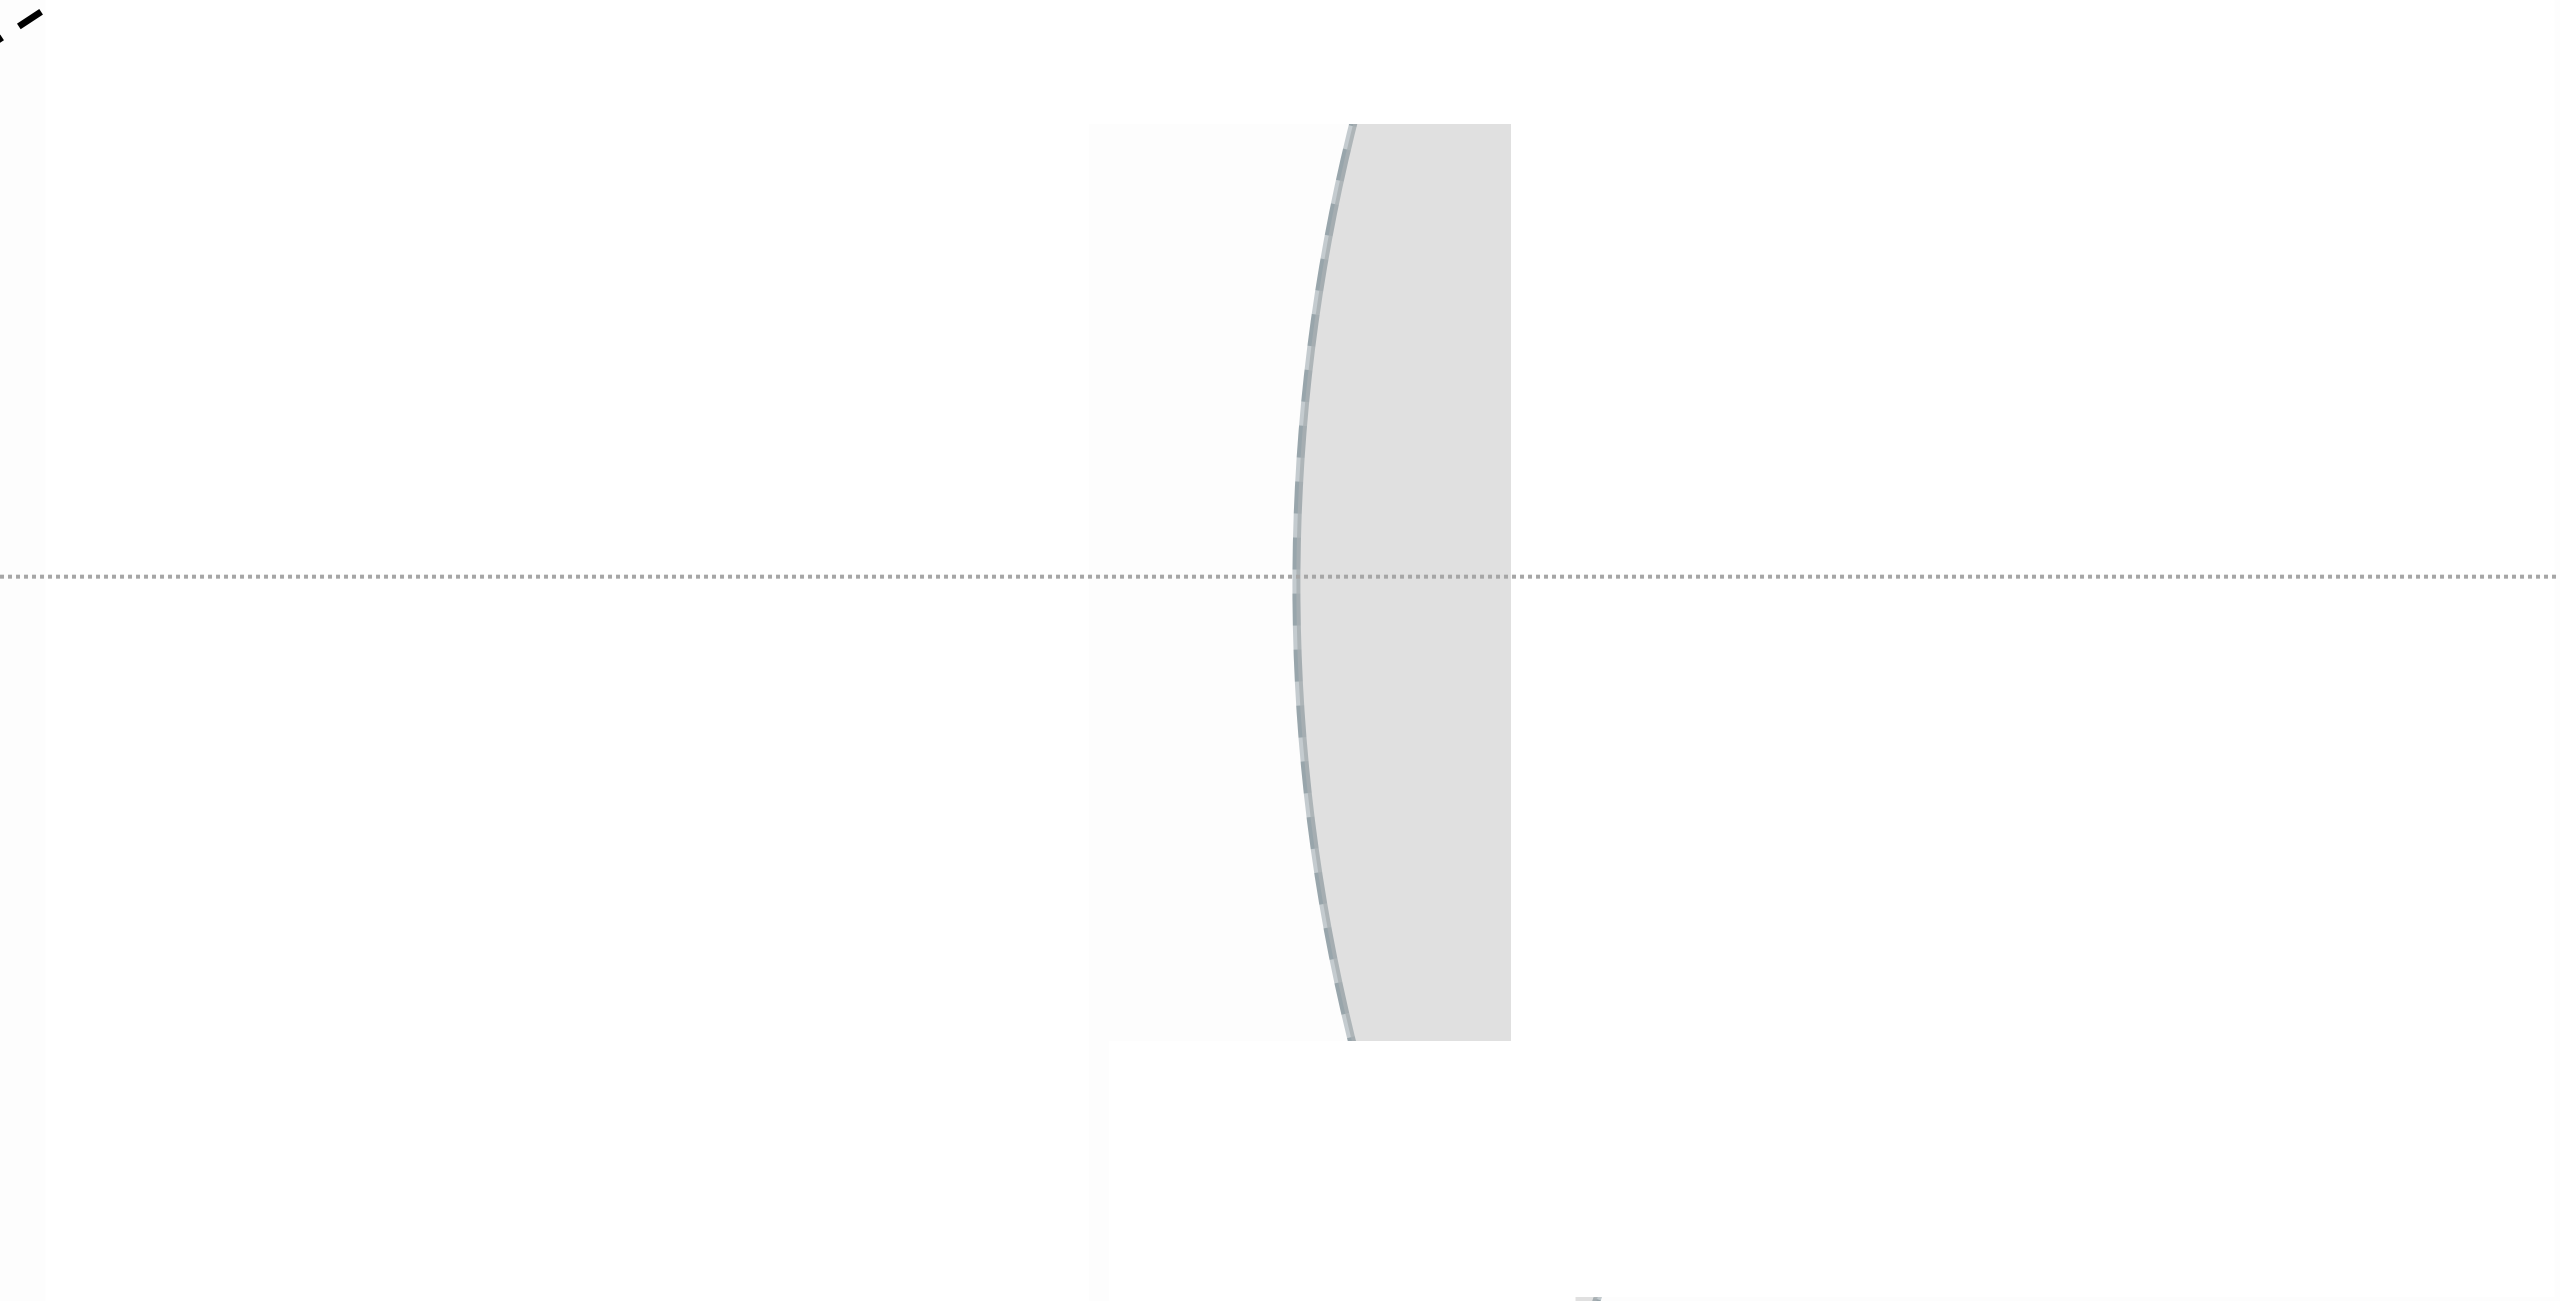

Thin lens

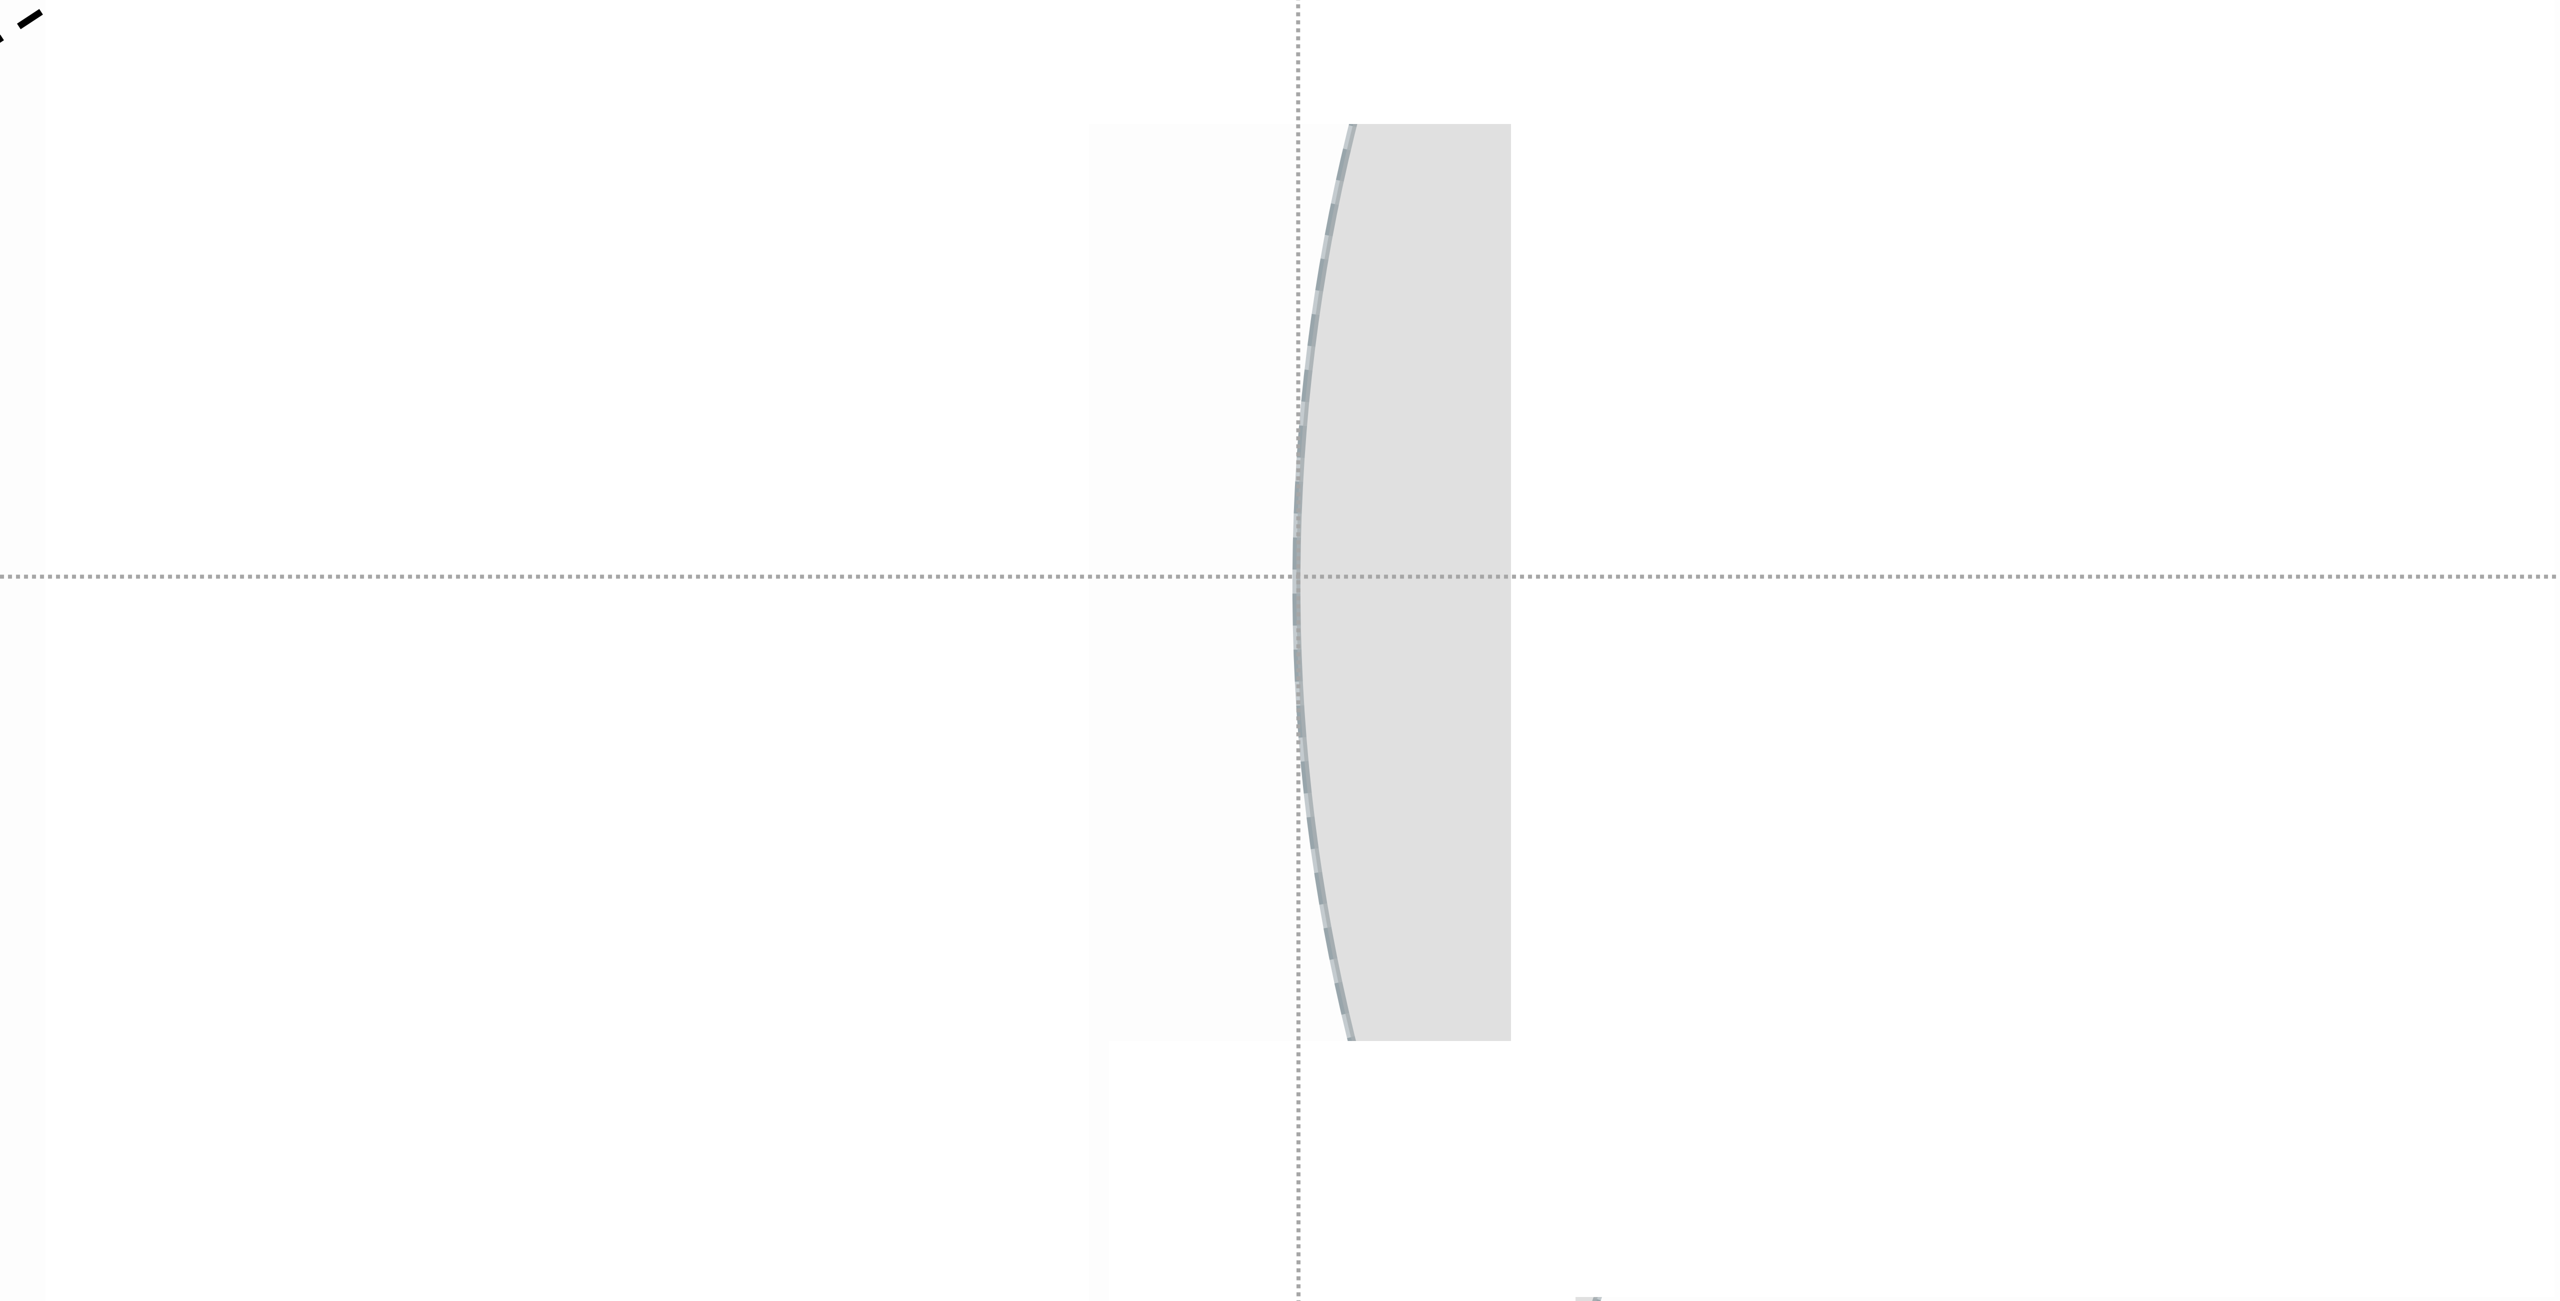

Thin lens

Thin lens

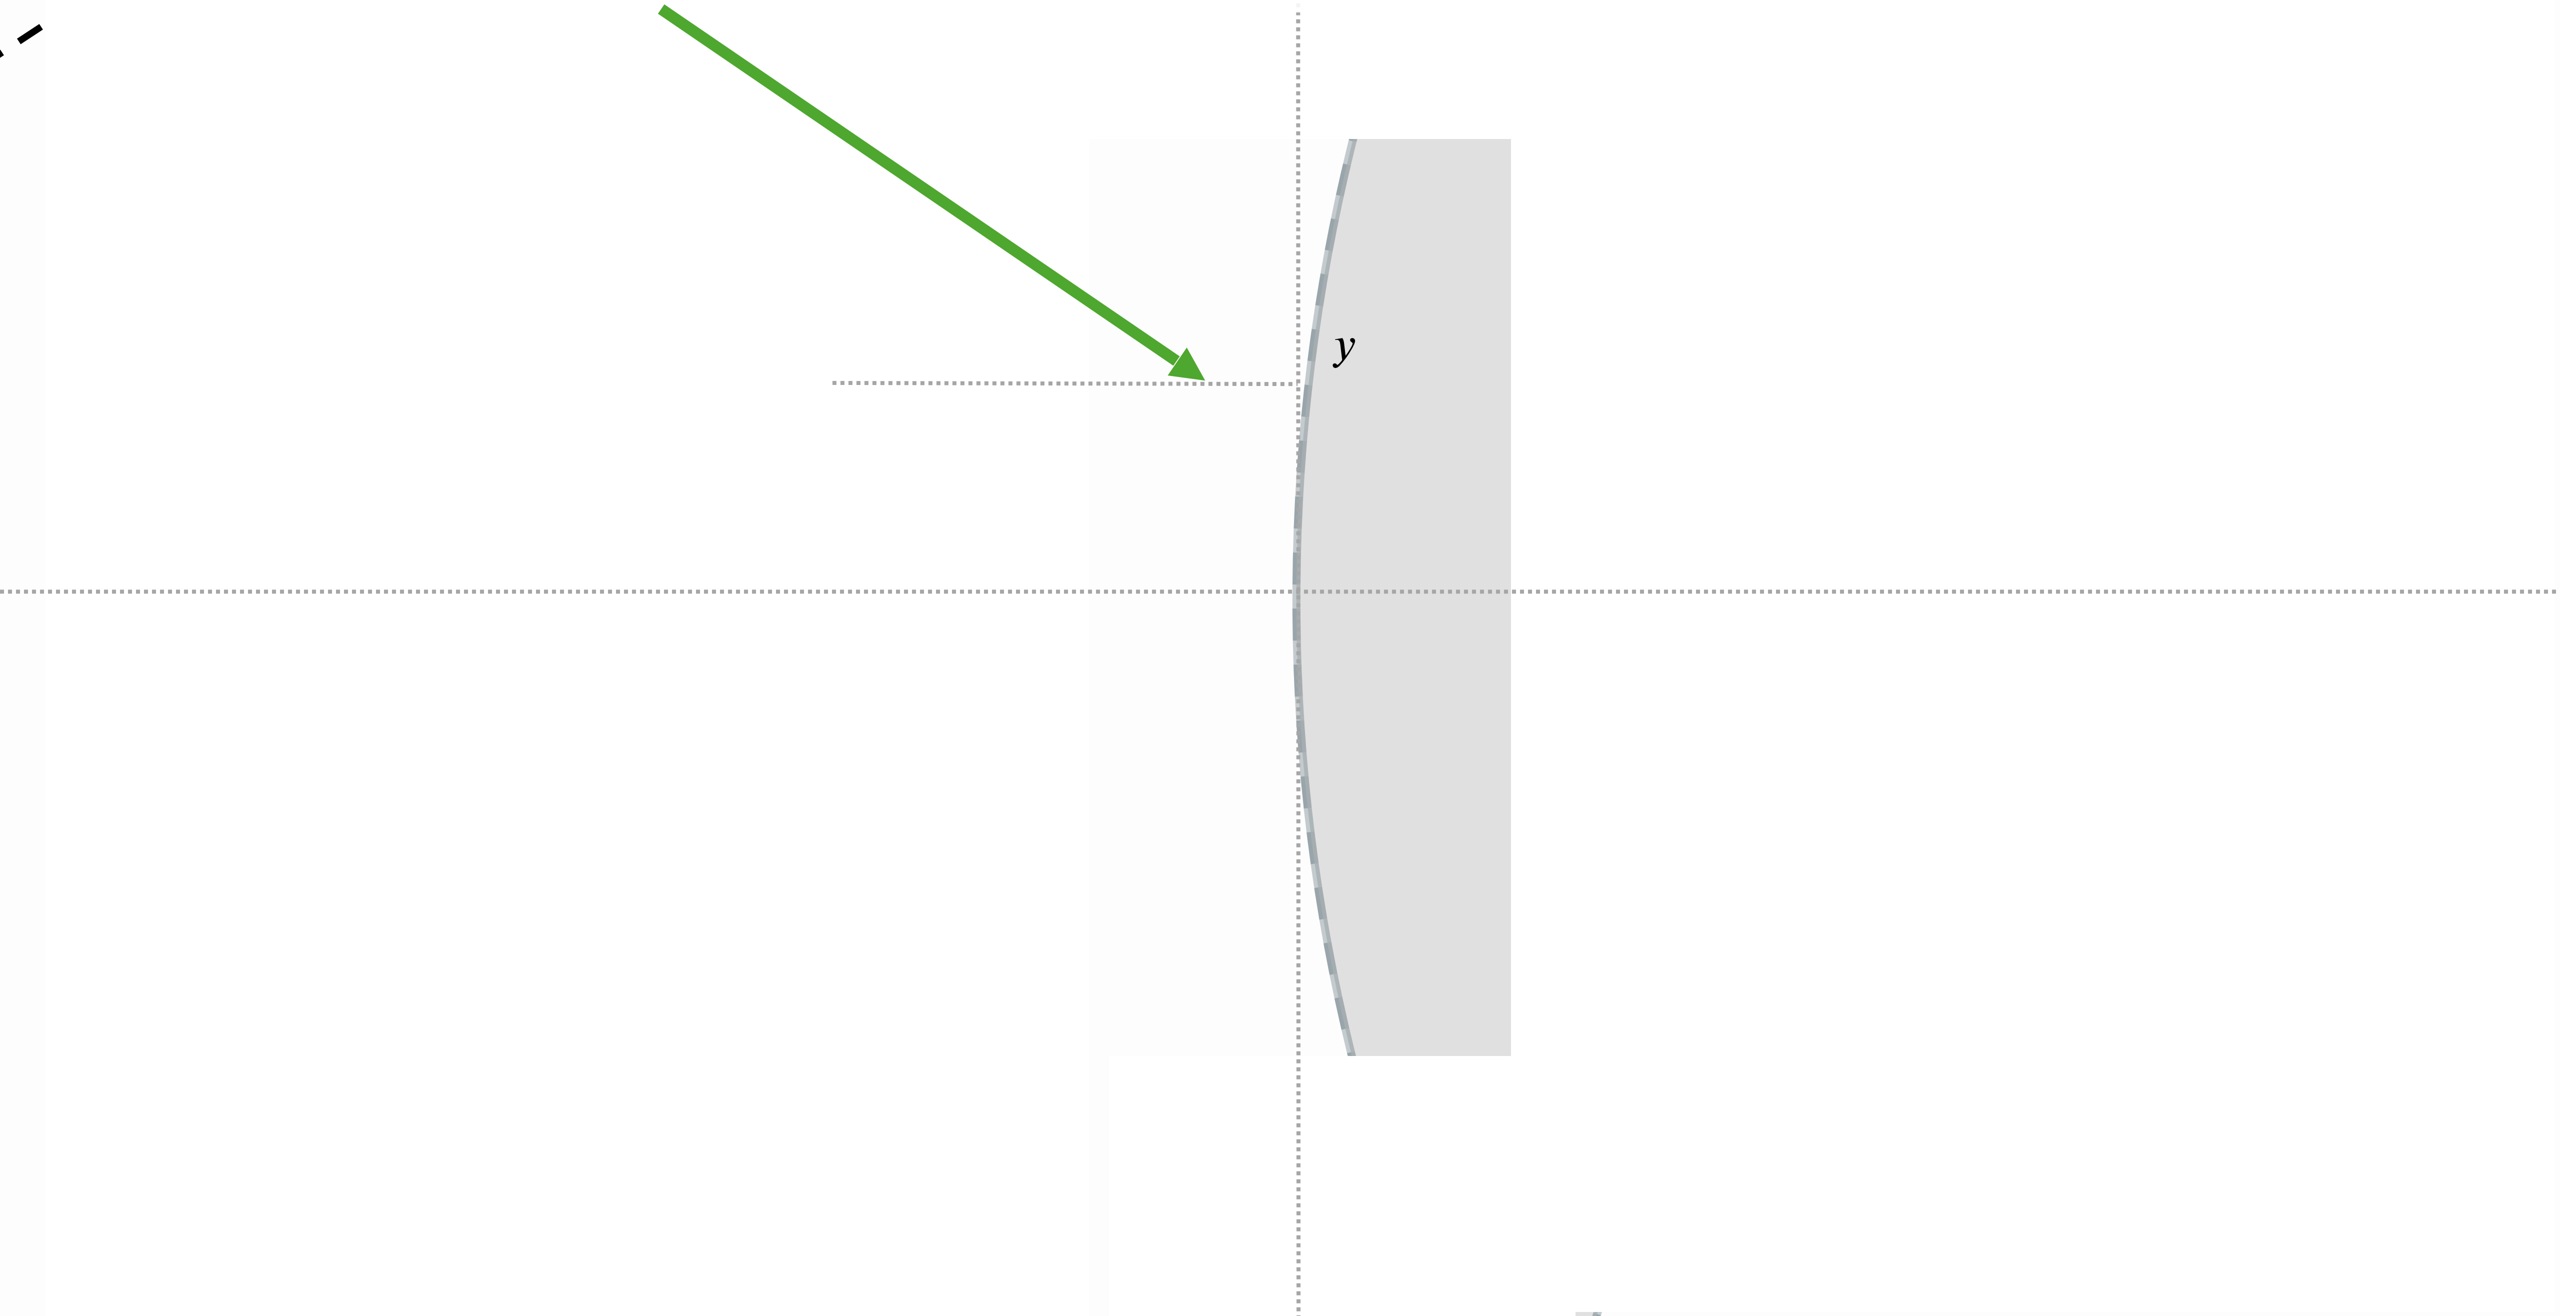

Thin lens

Thin lens

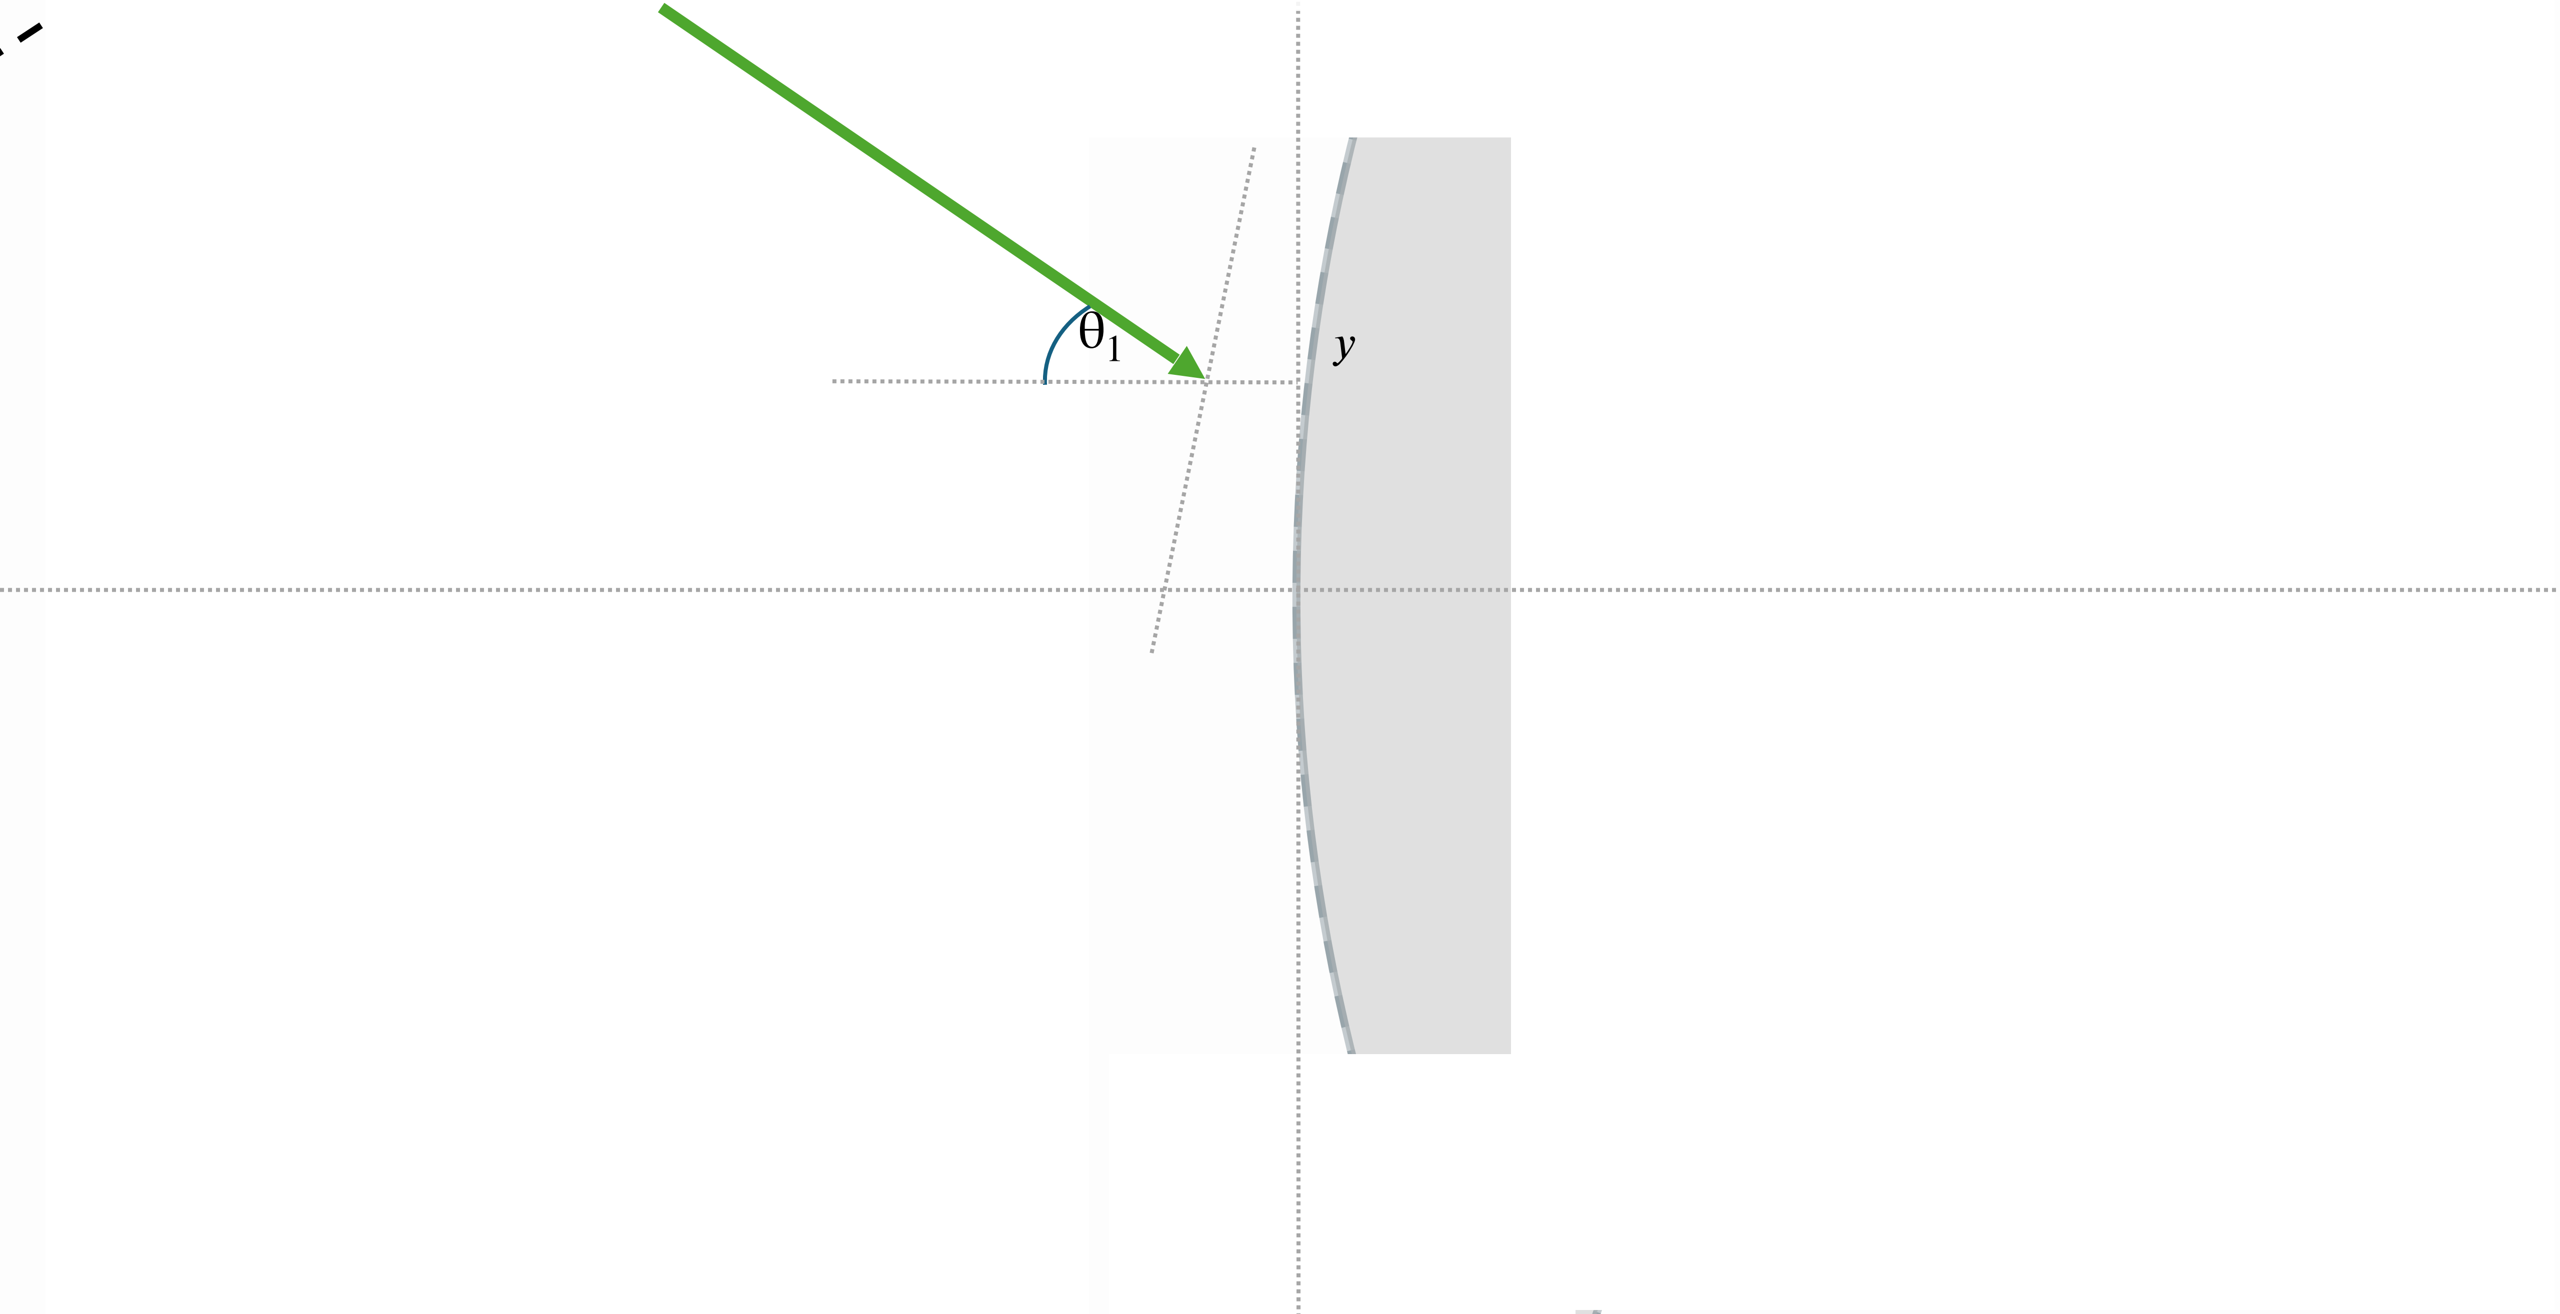

Thin lens

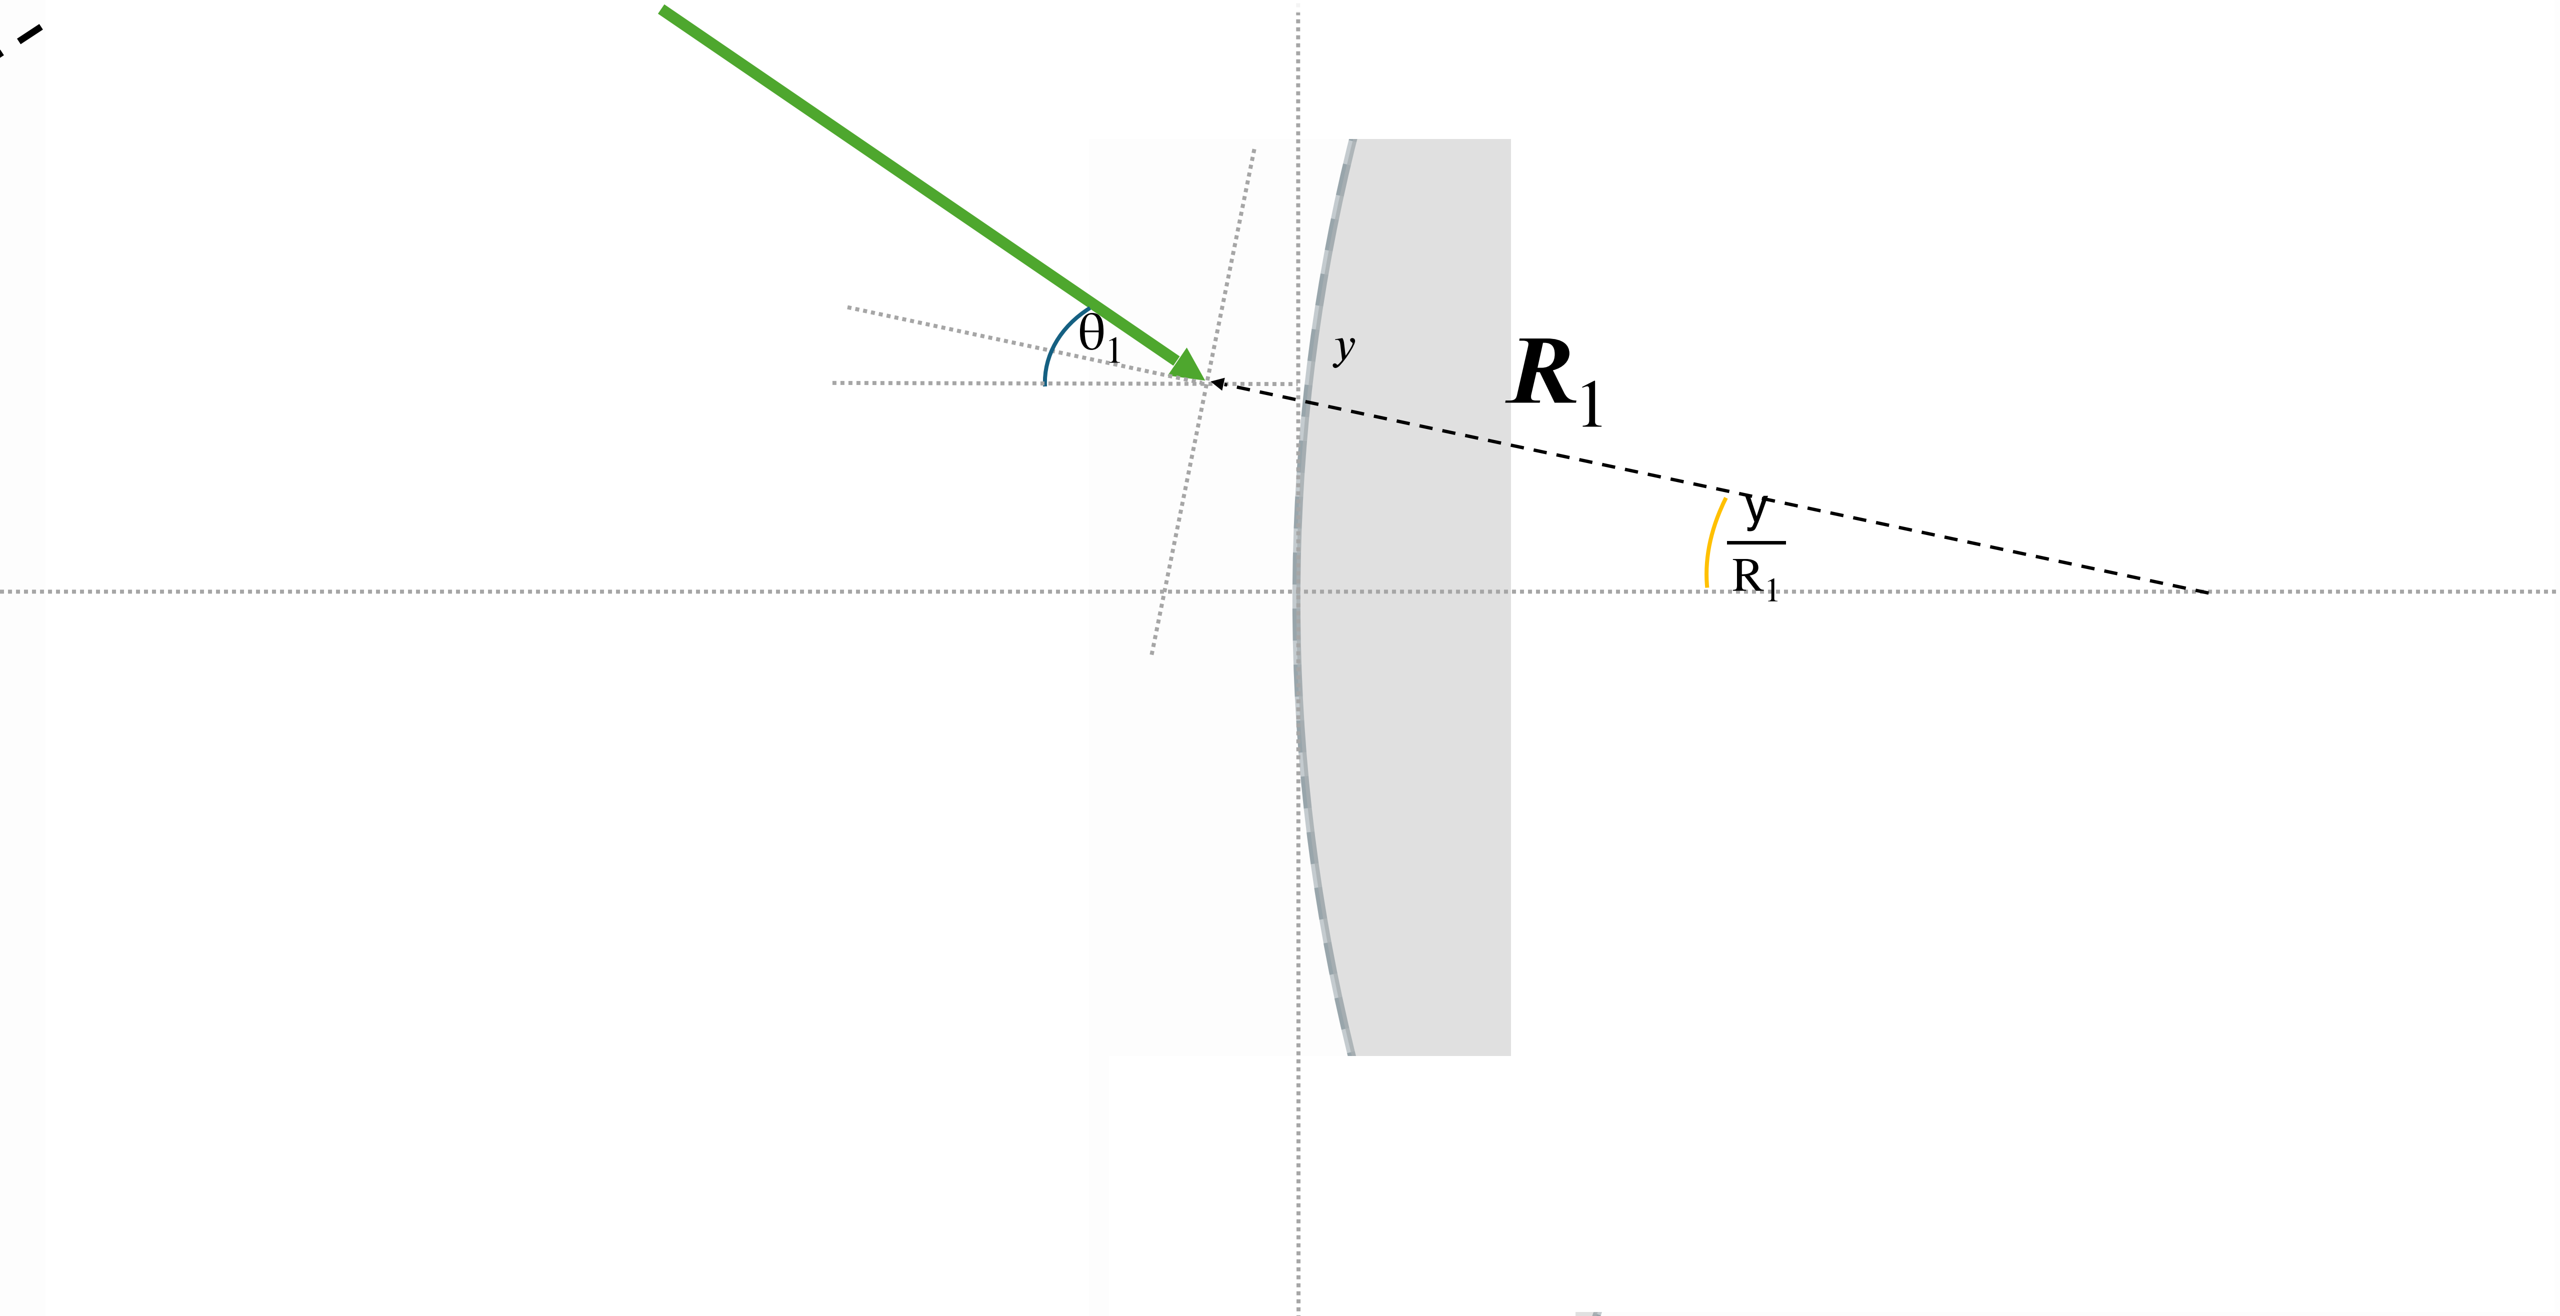

Thin lens

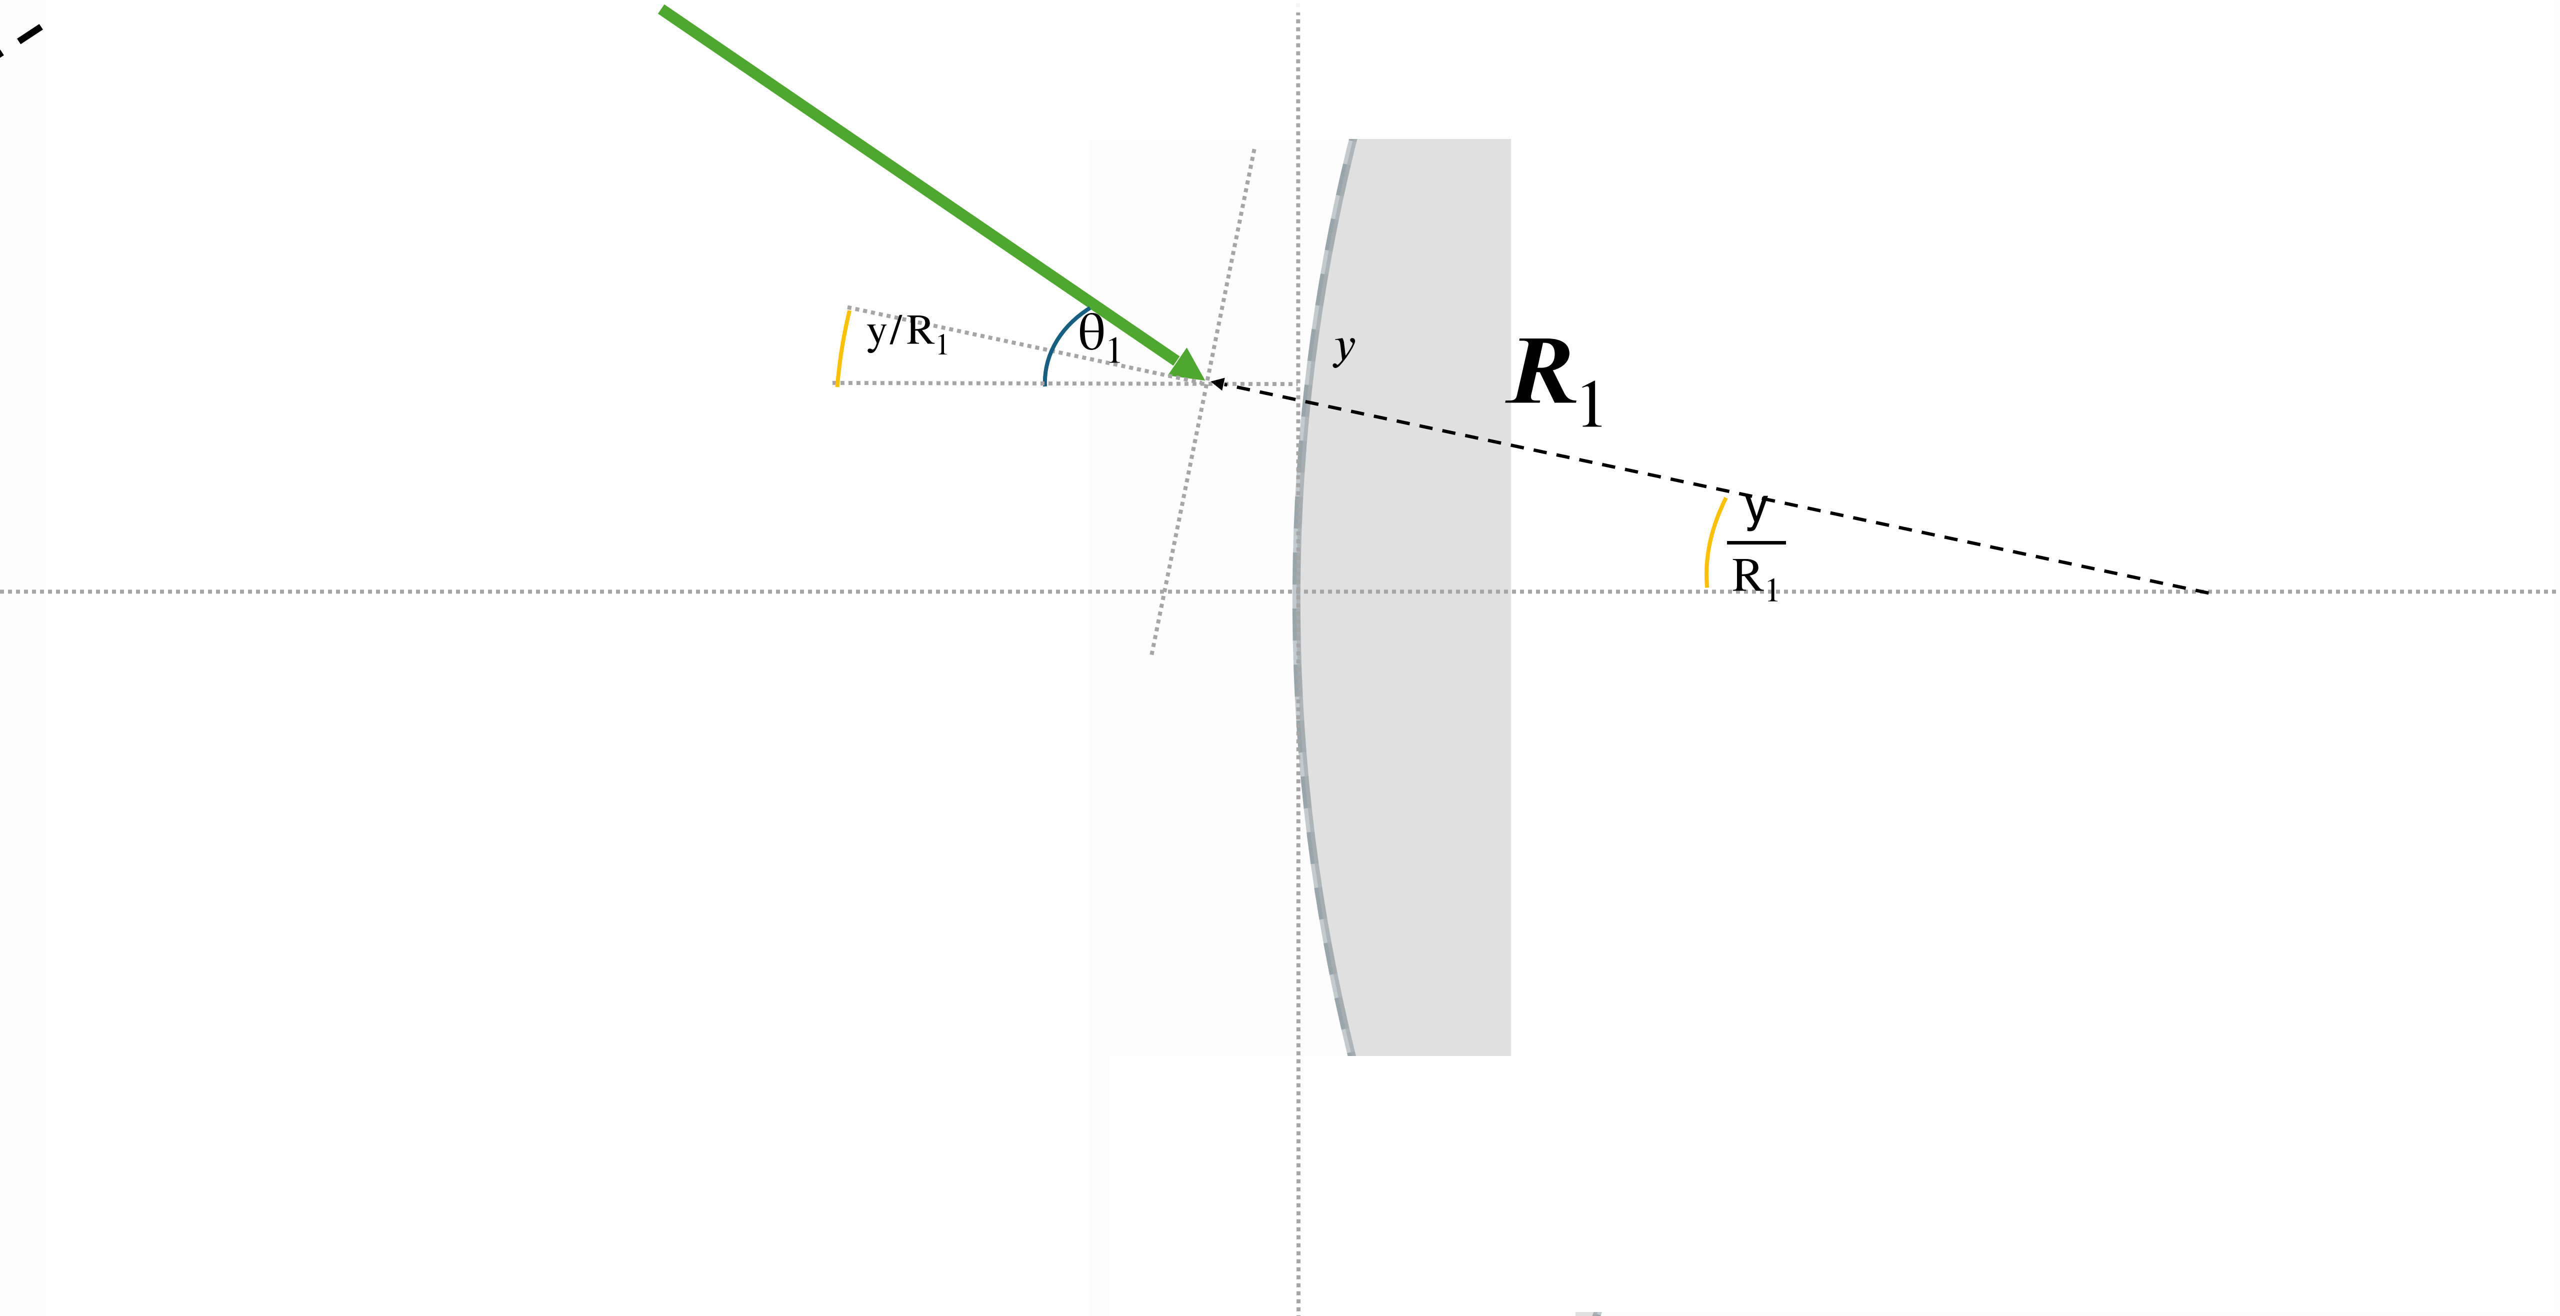

Thin lens

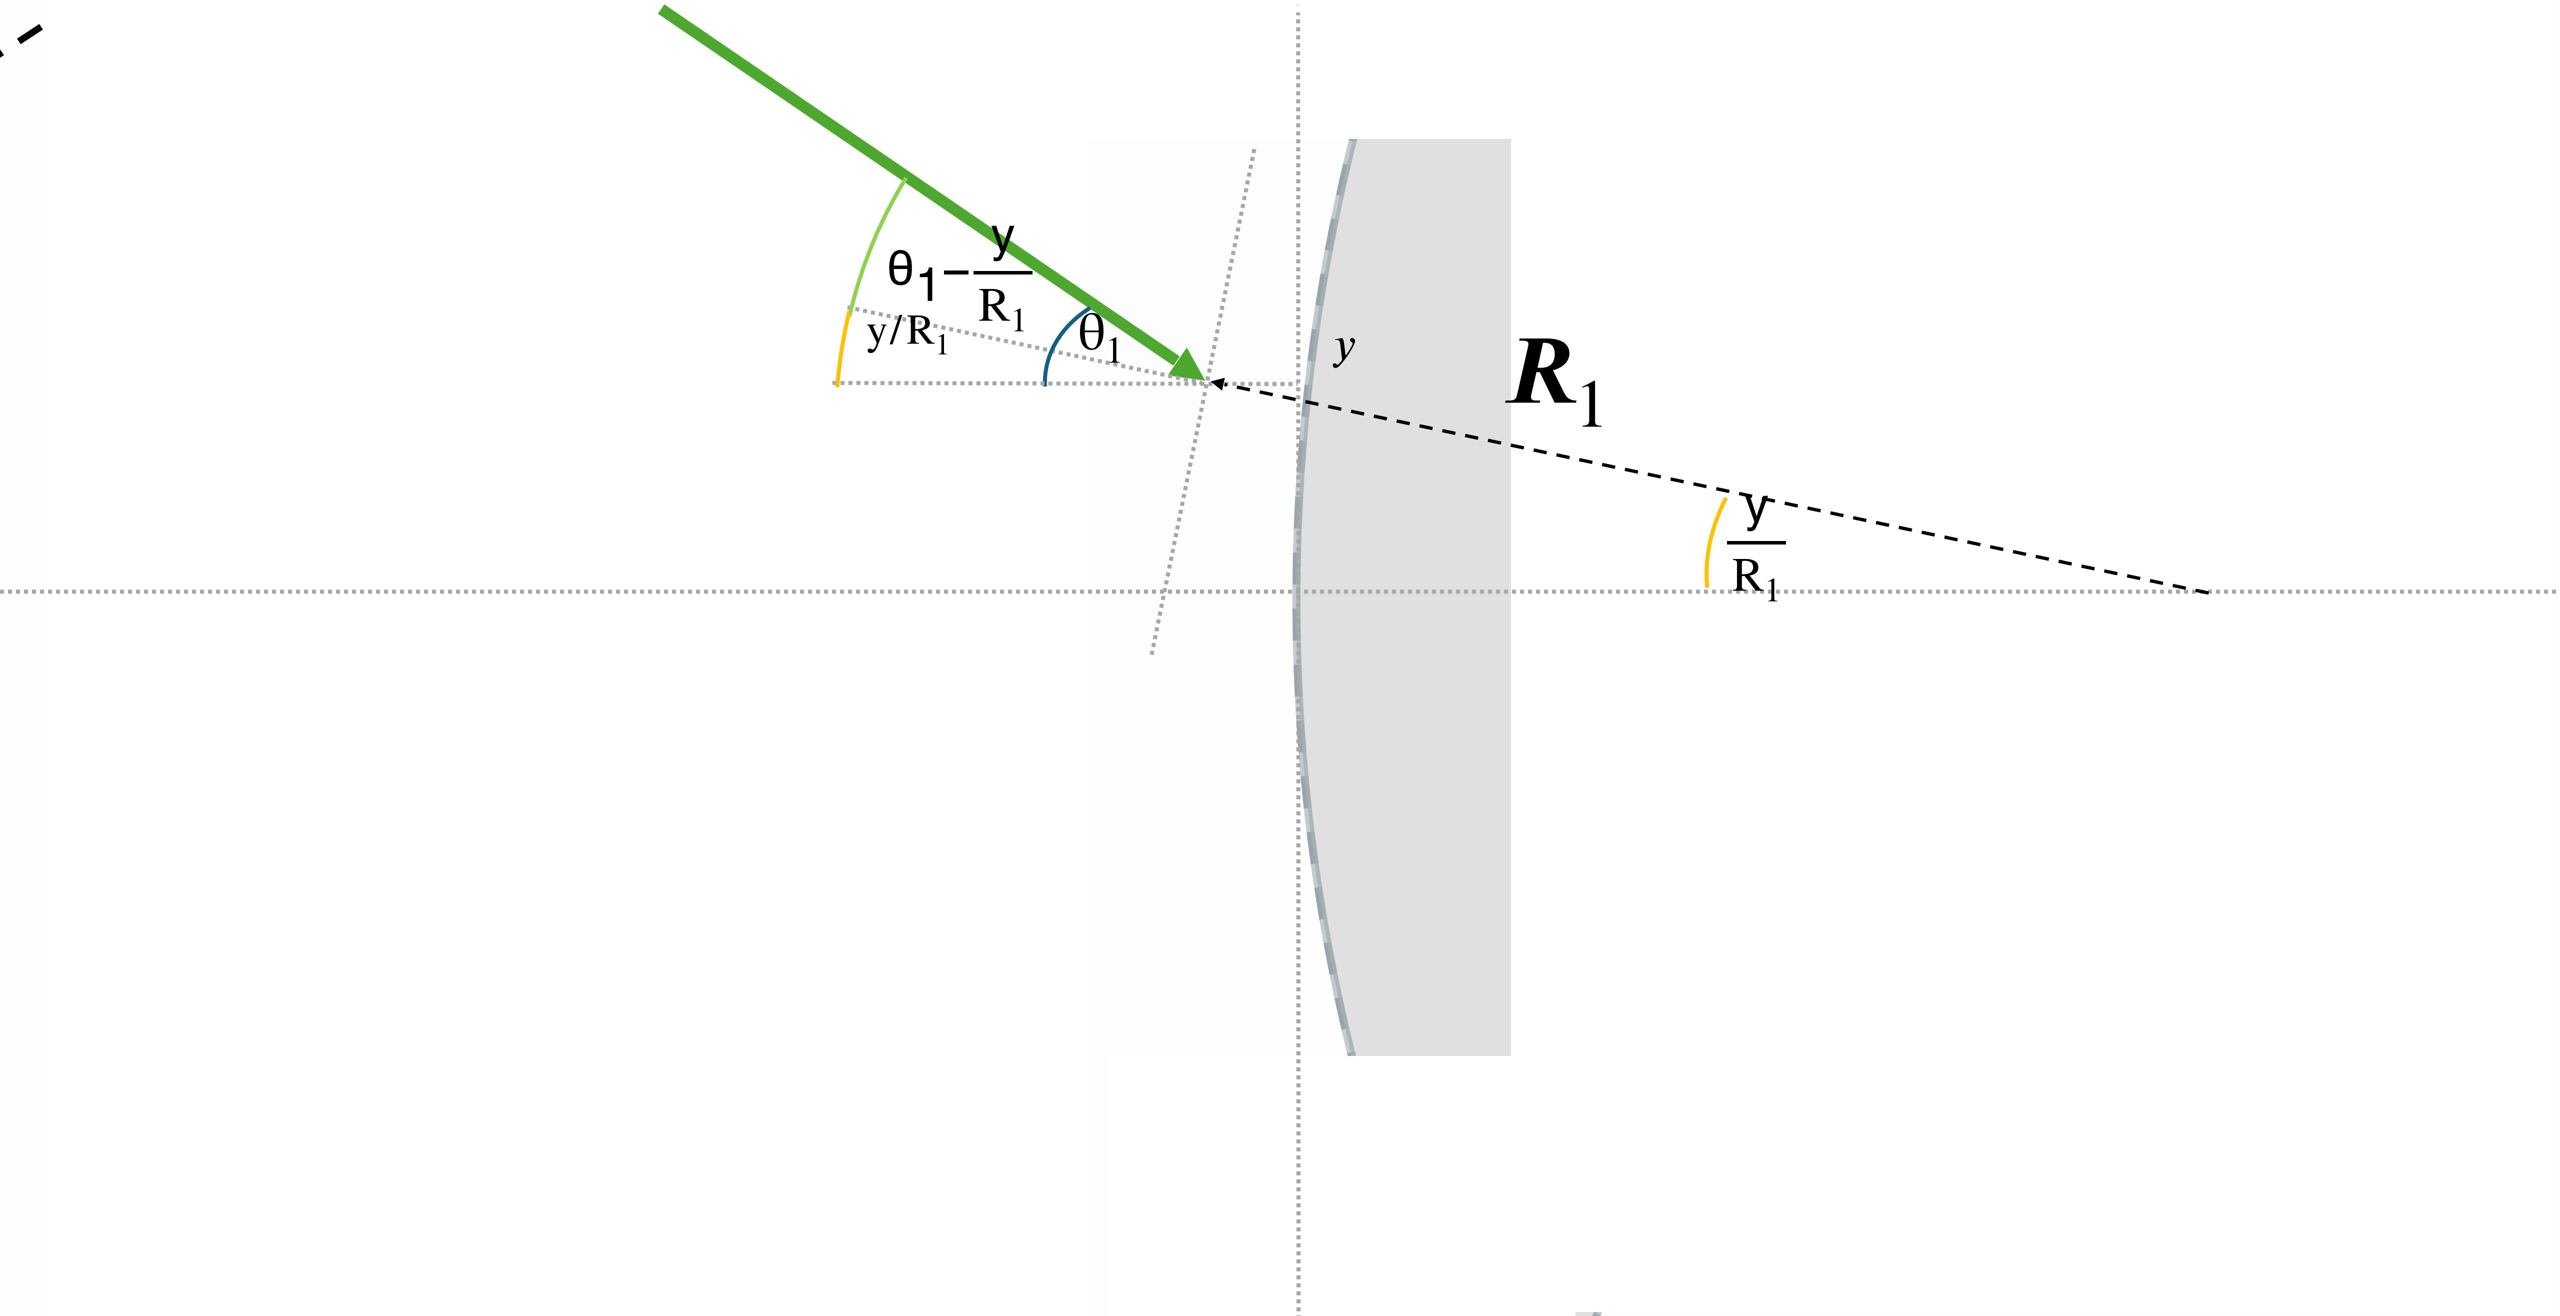

Thin lens

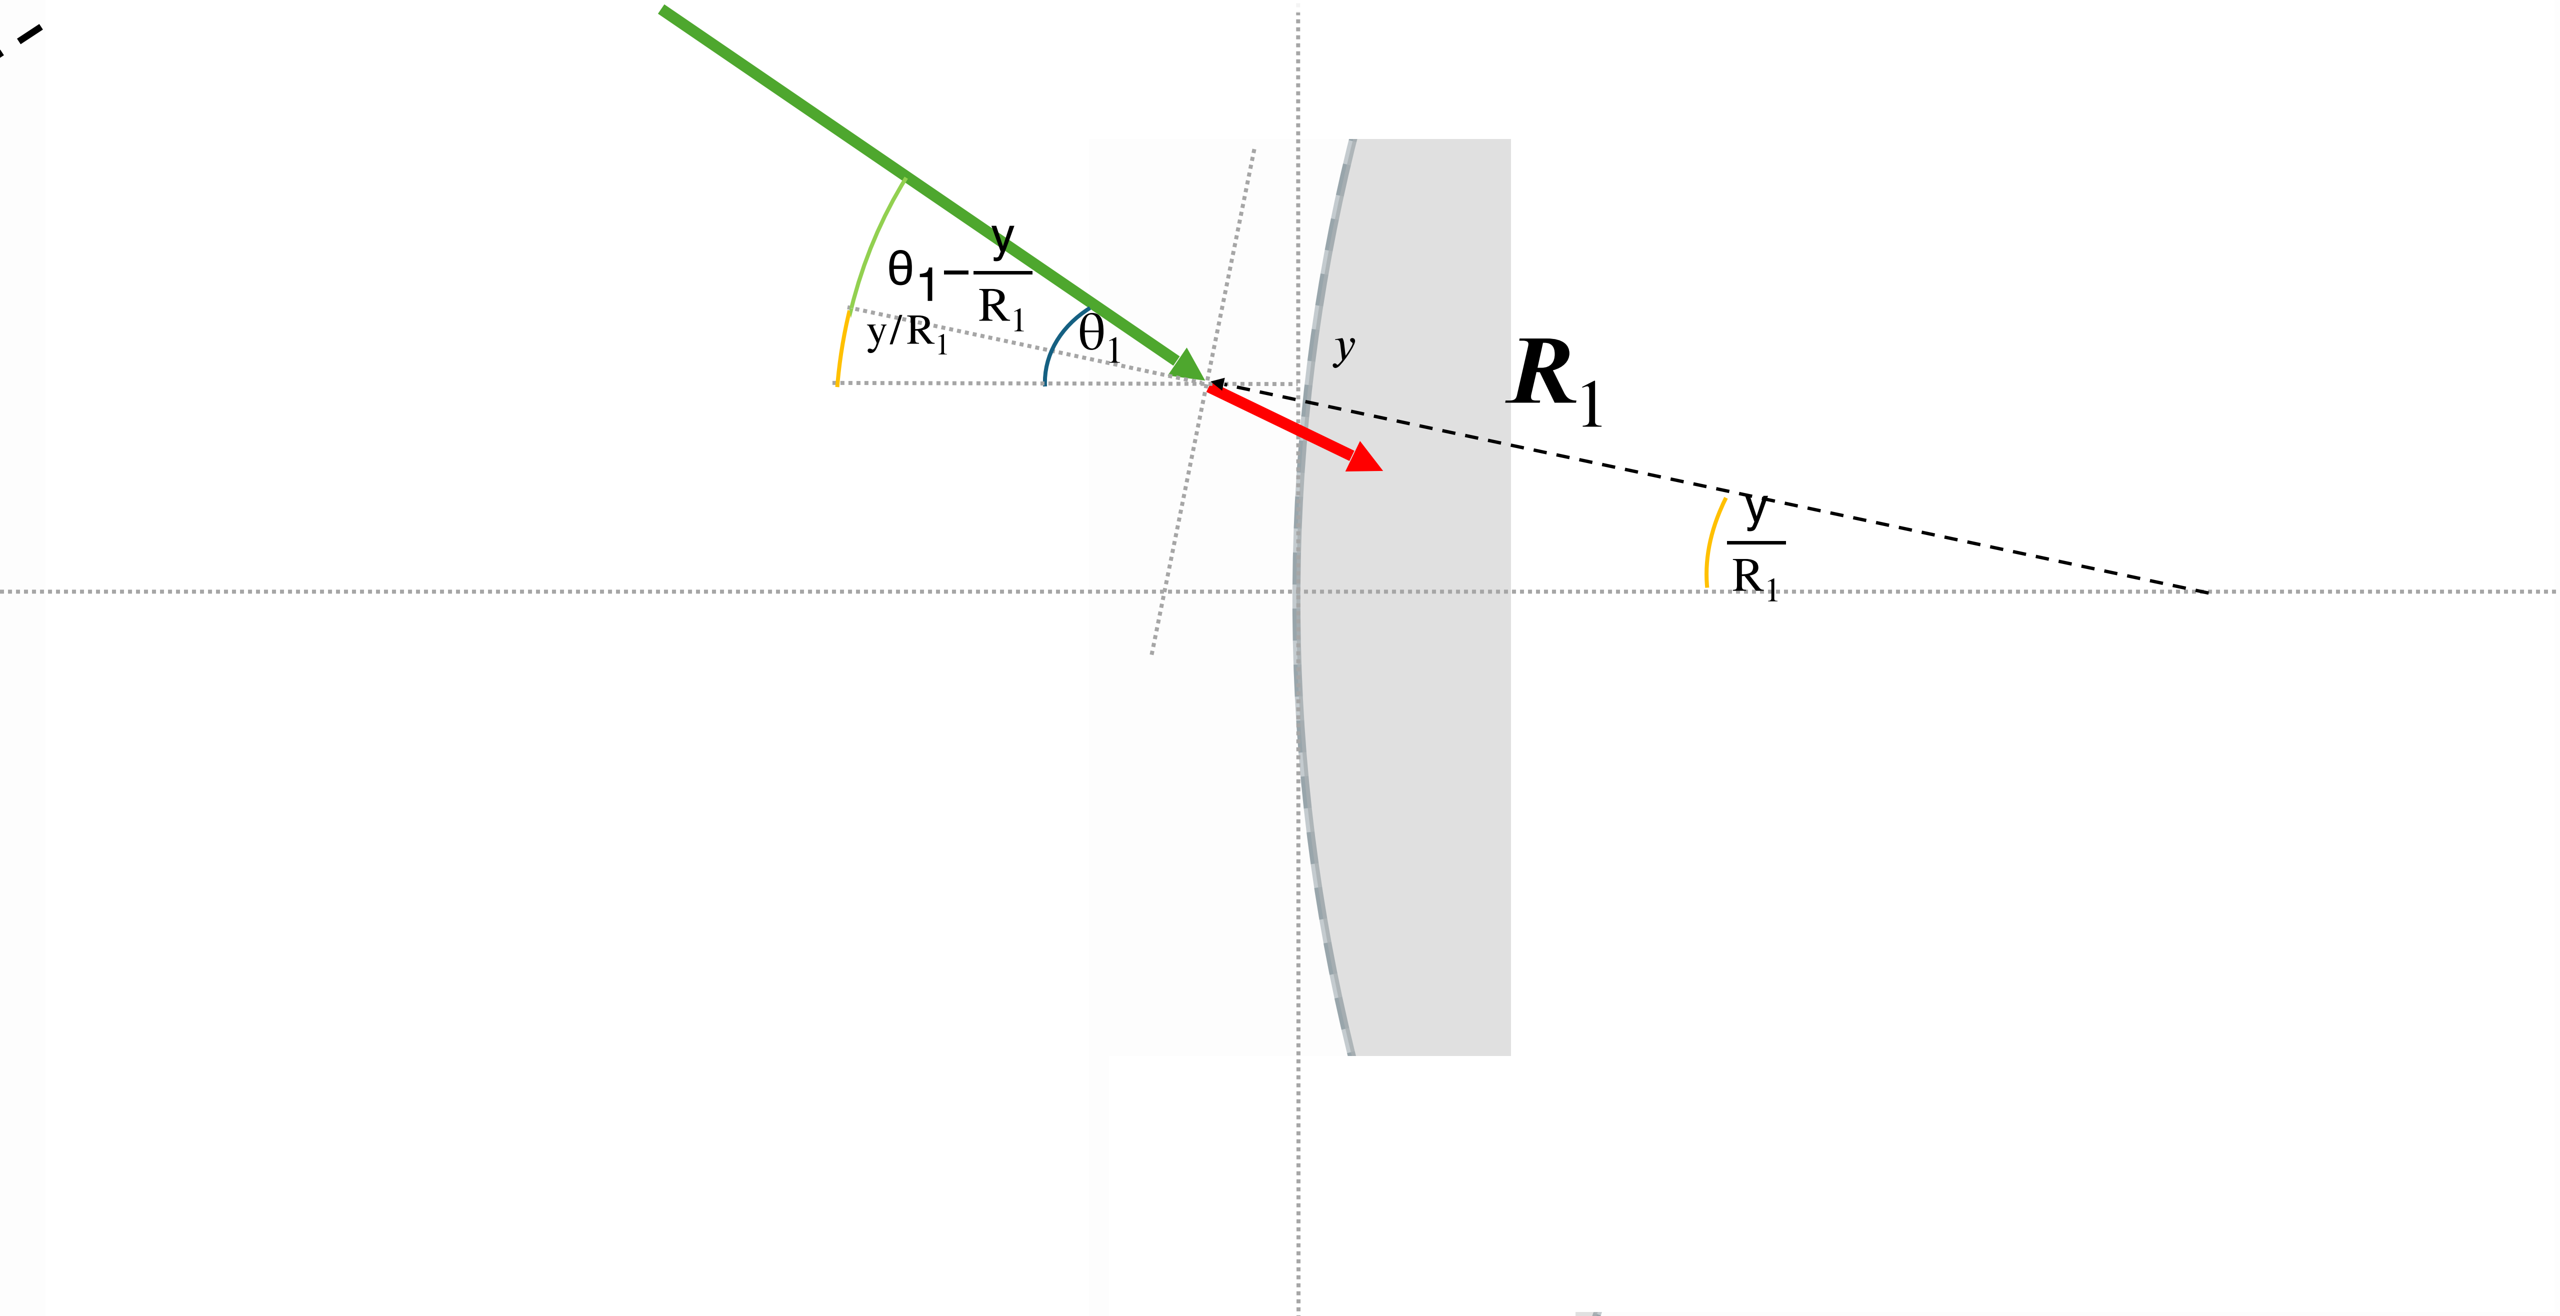

Thin lens

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$$

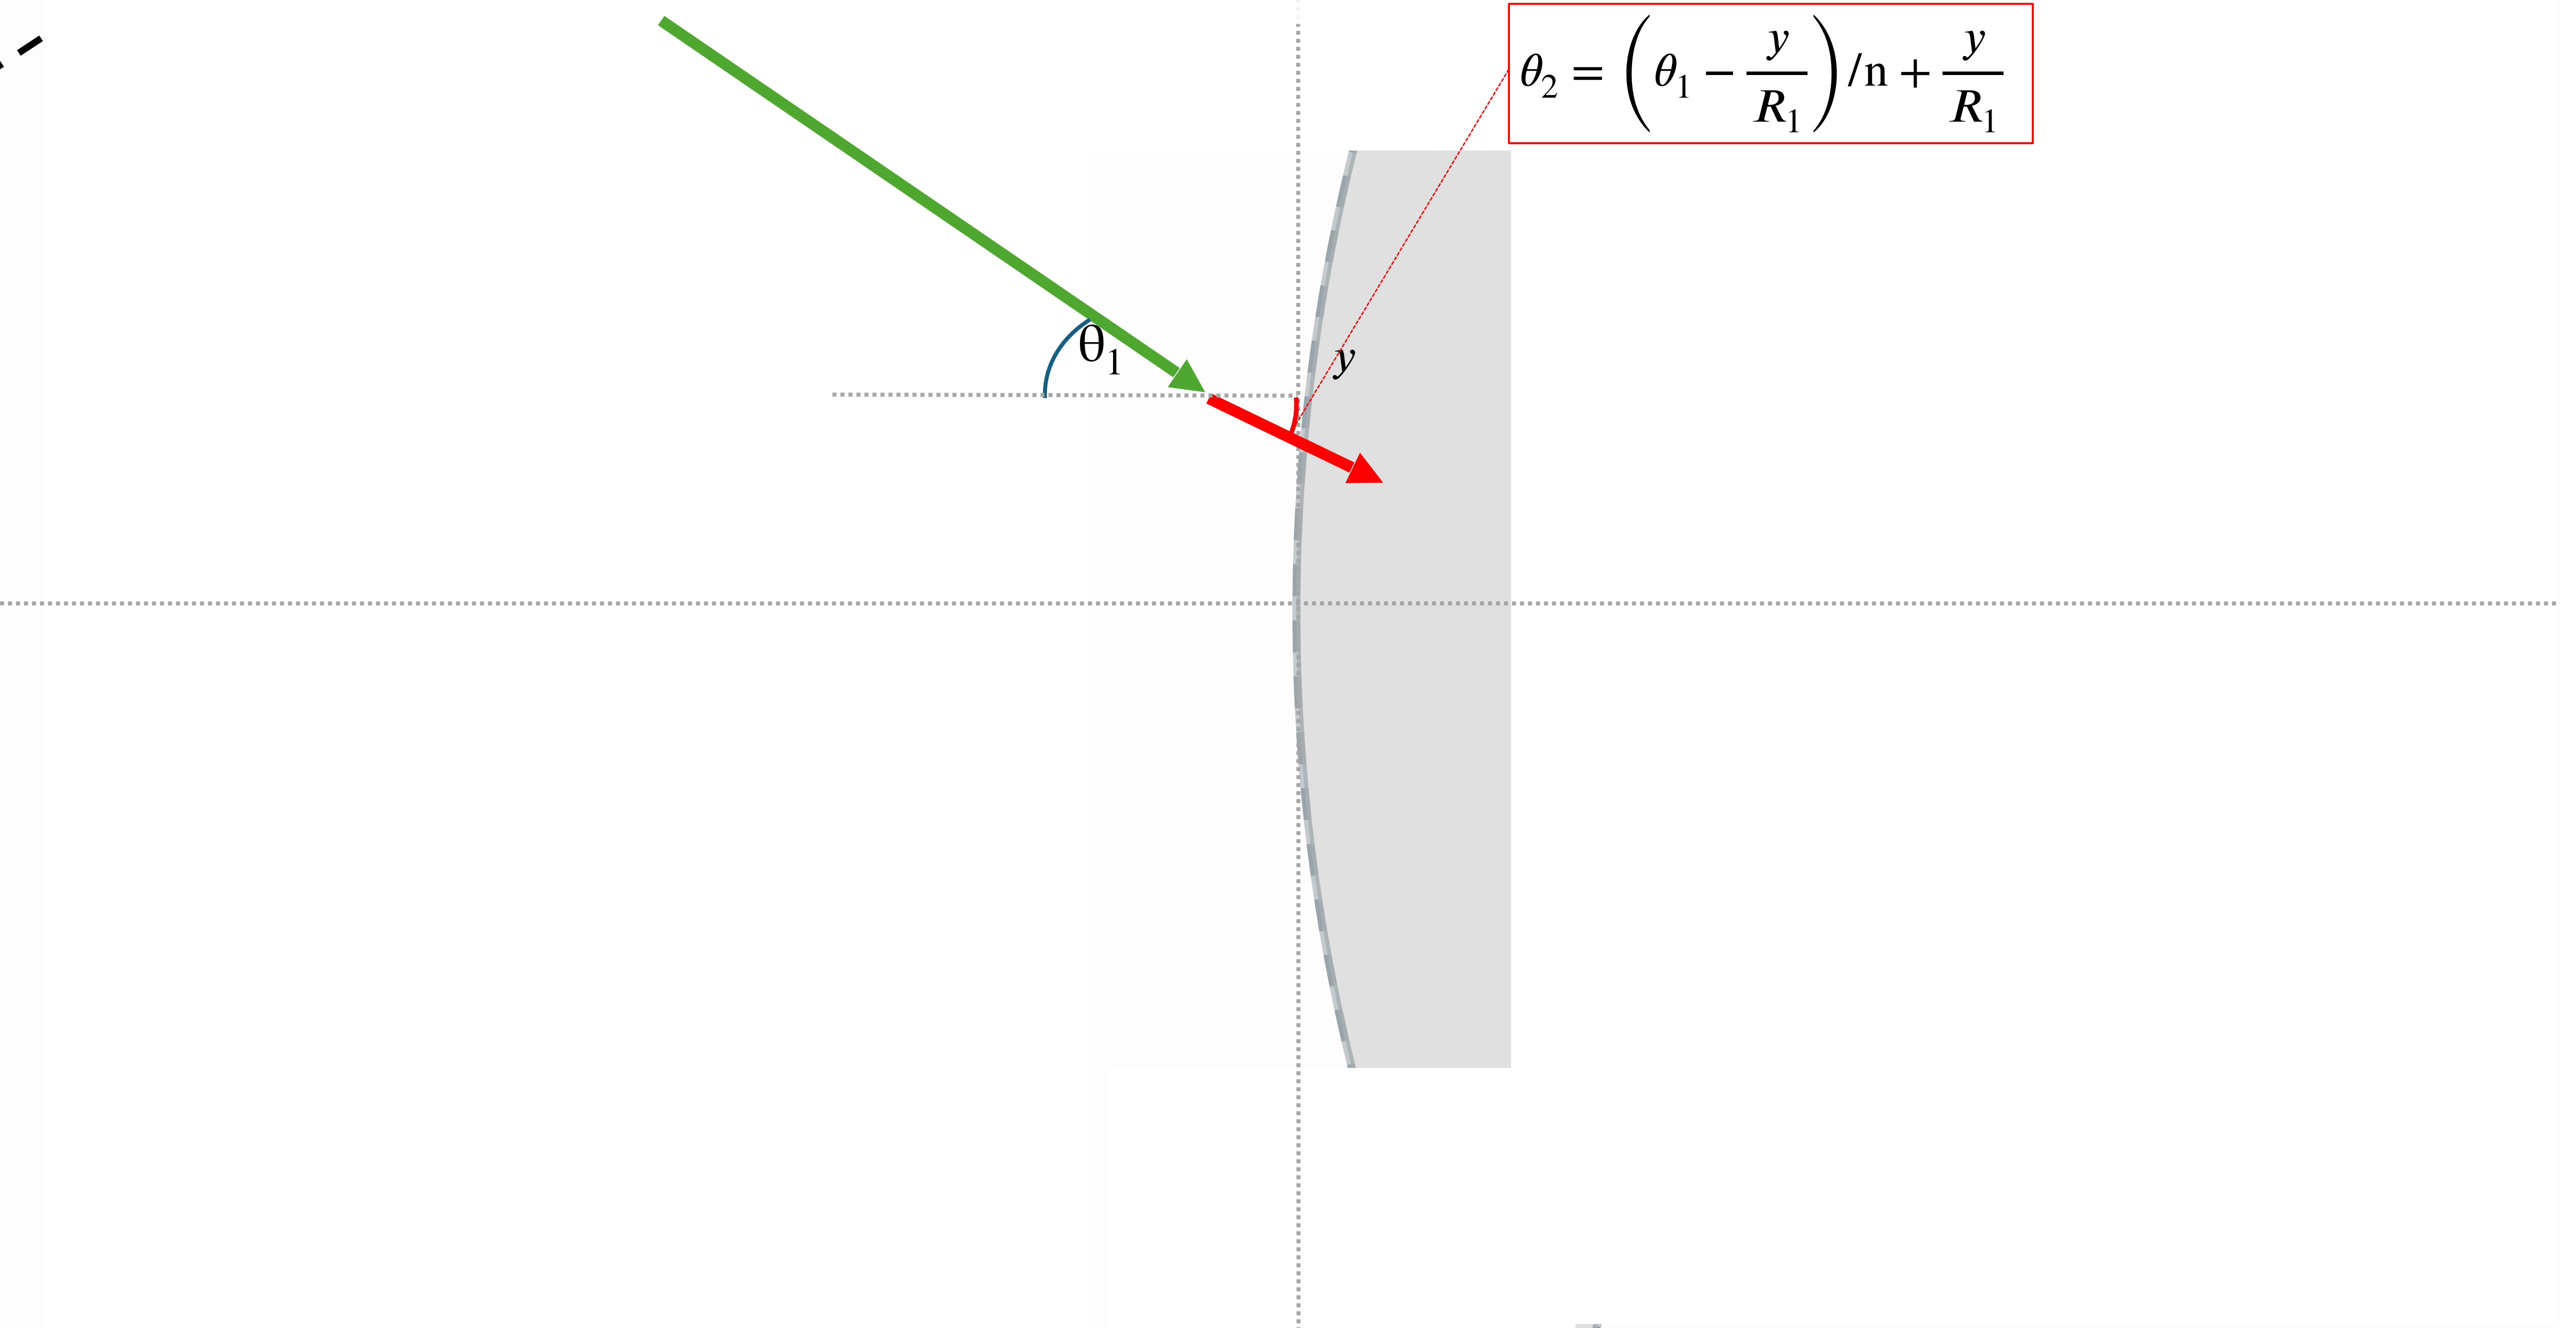

Thin lens

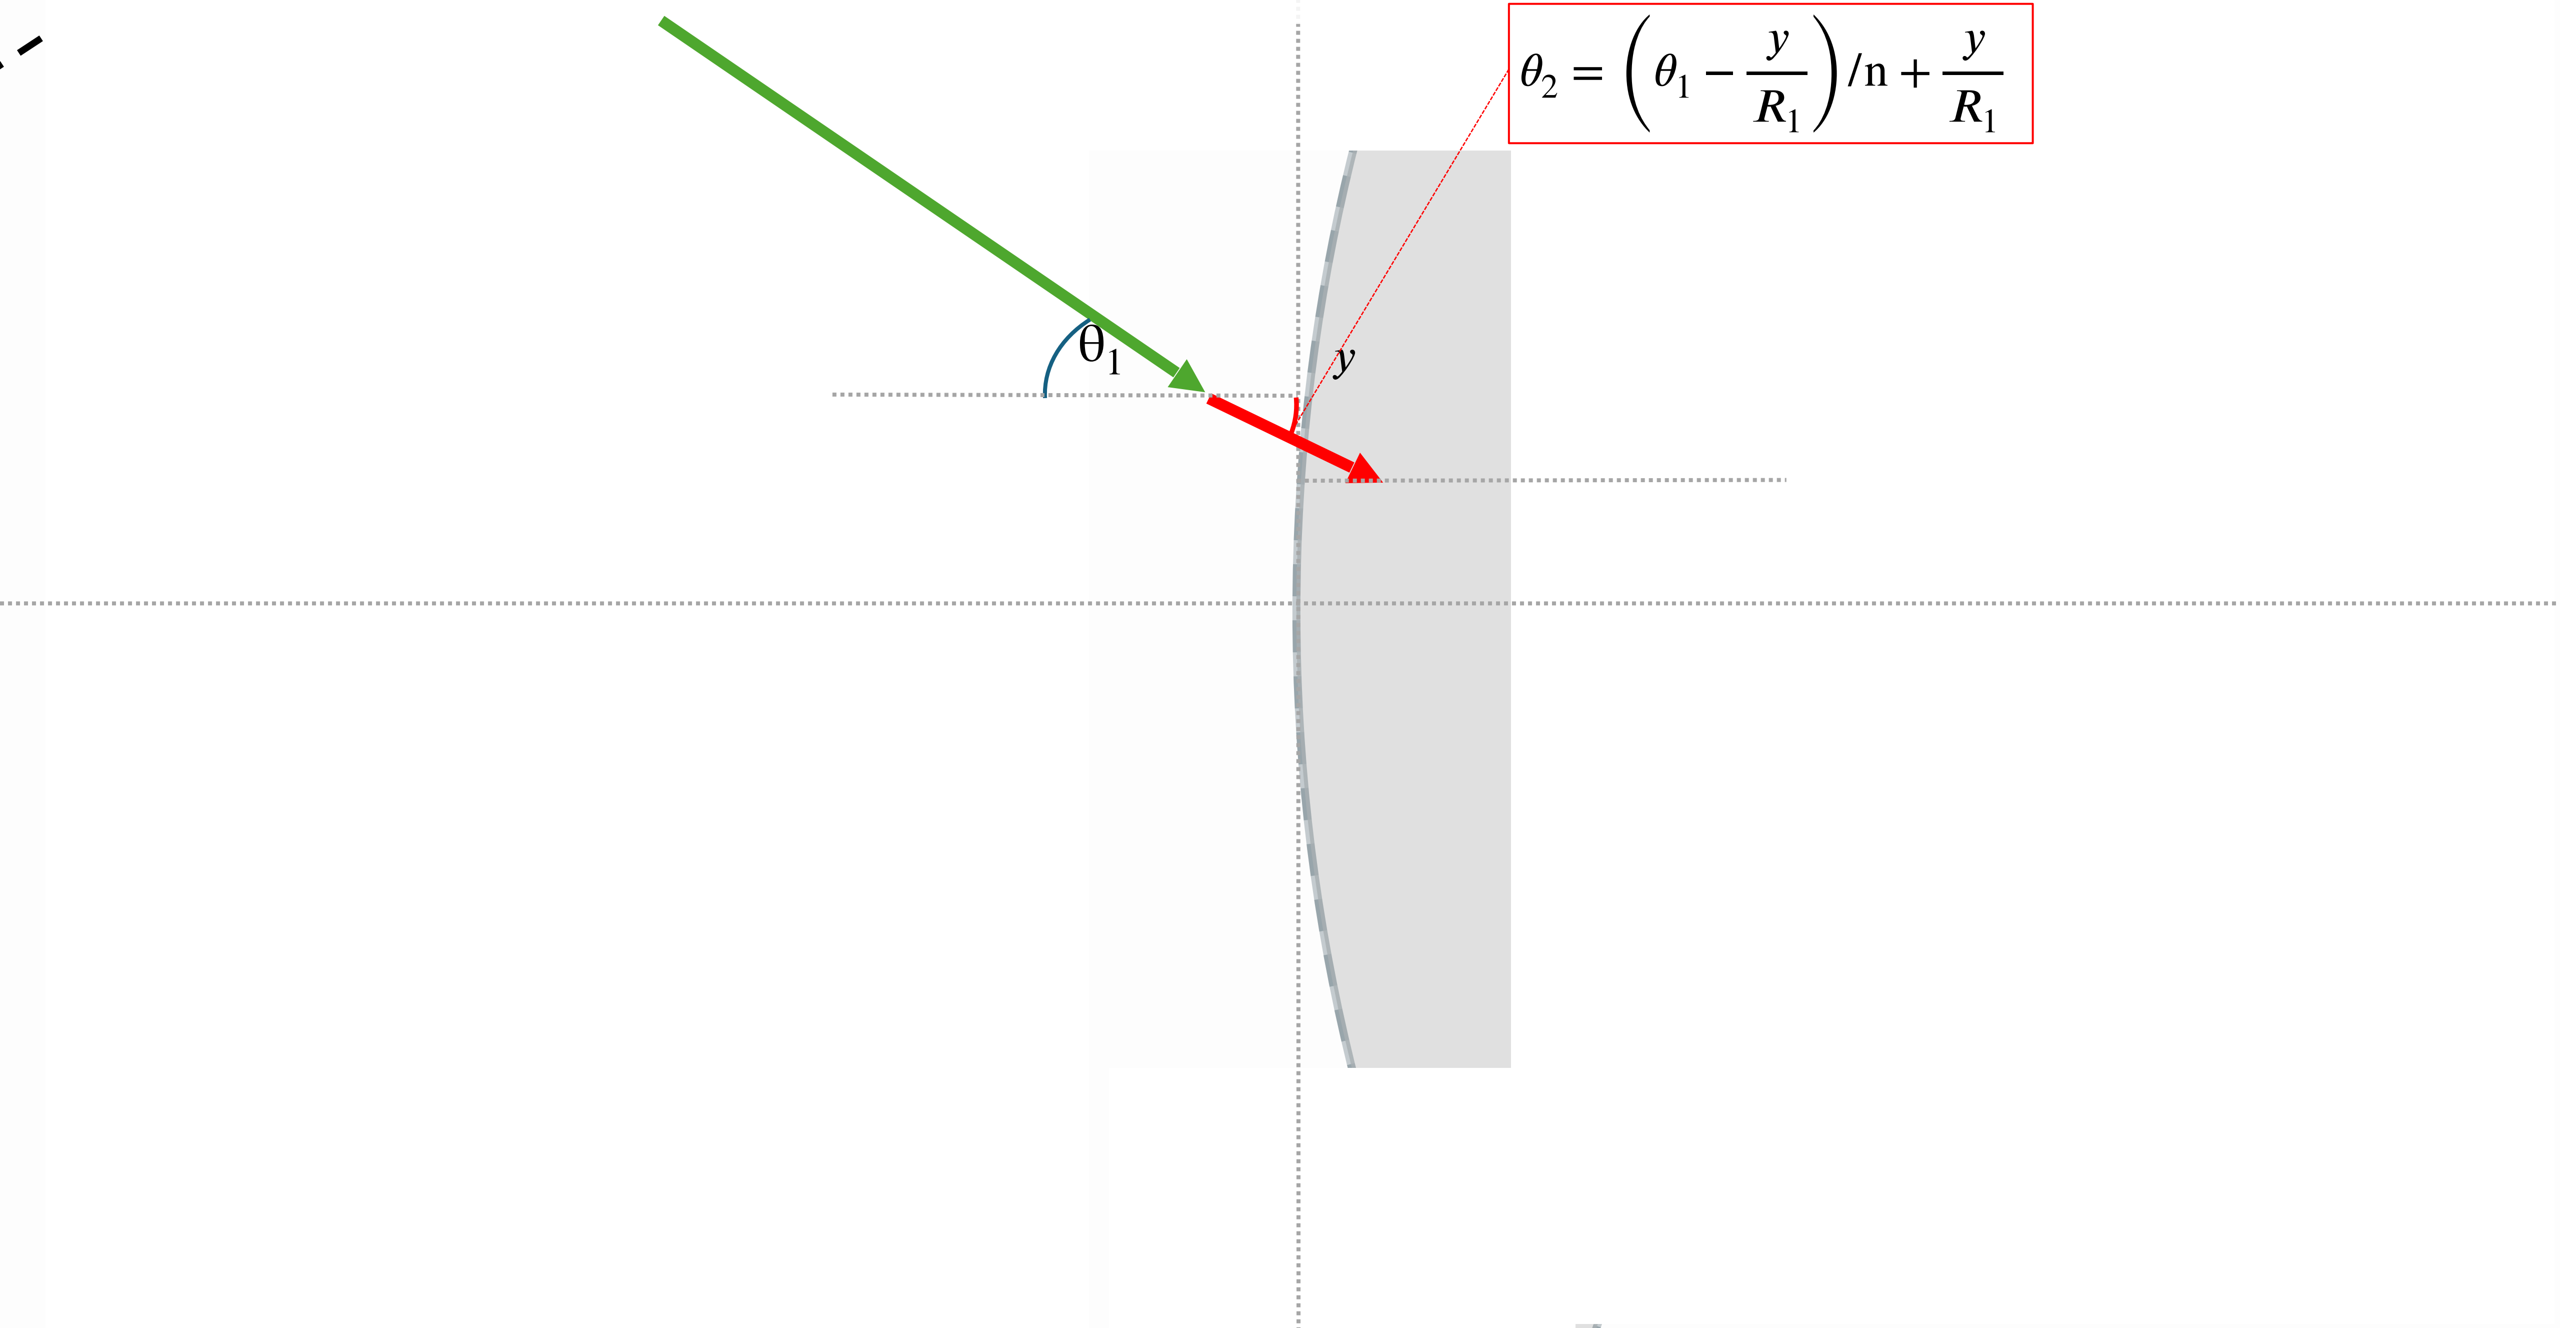

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$$

Thin lens

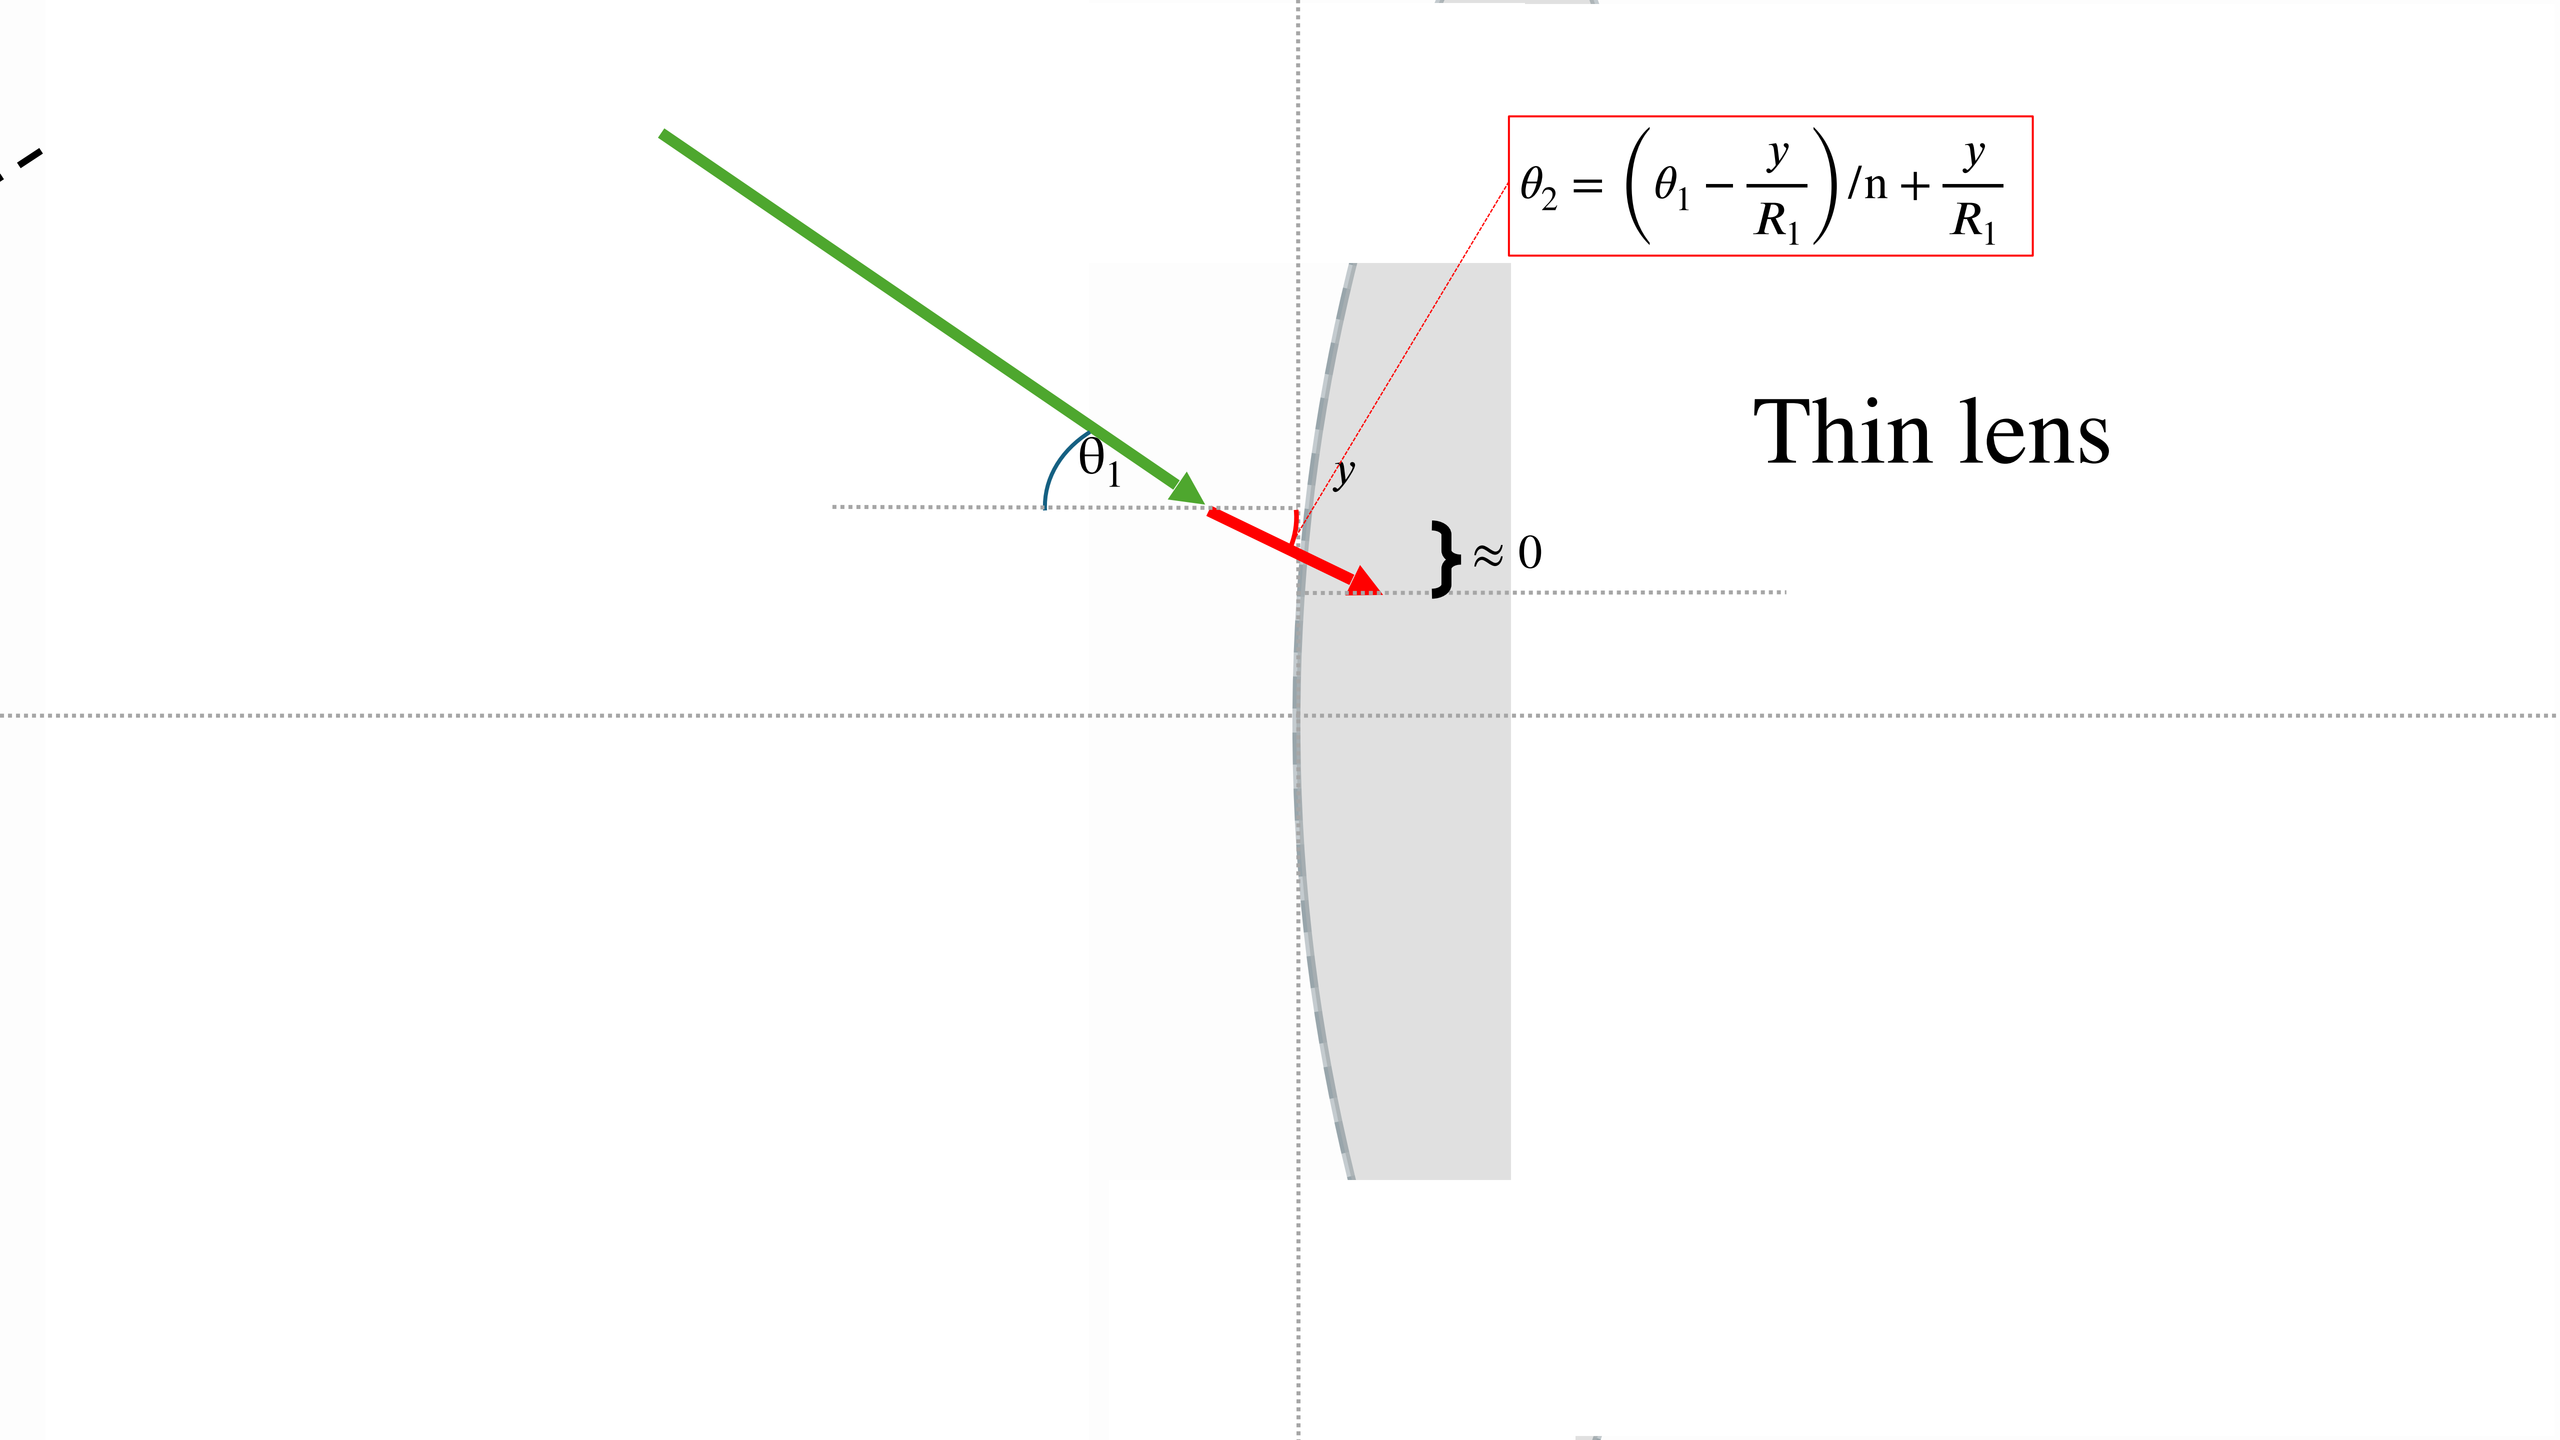
$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$
θ_1
y
≈ 0

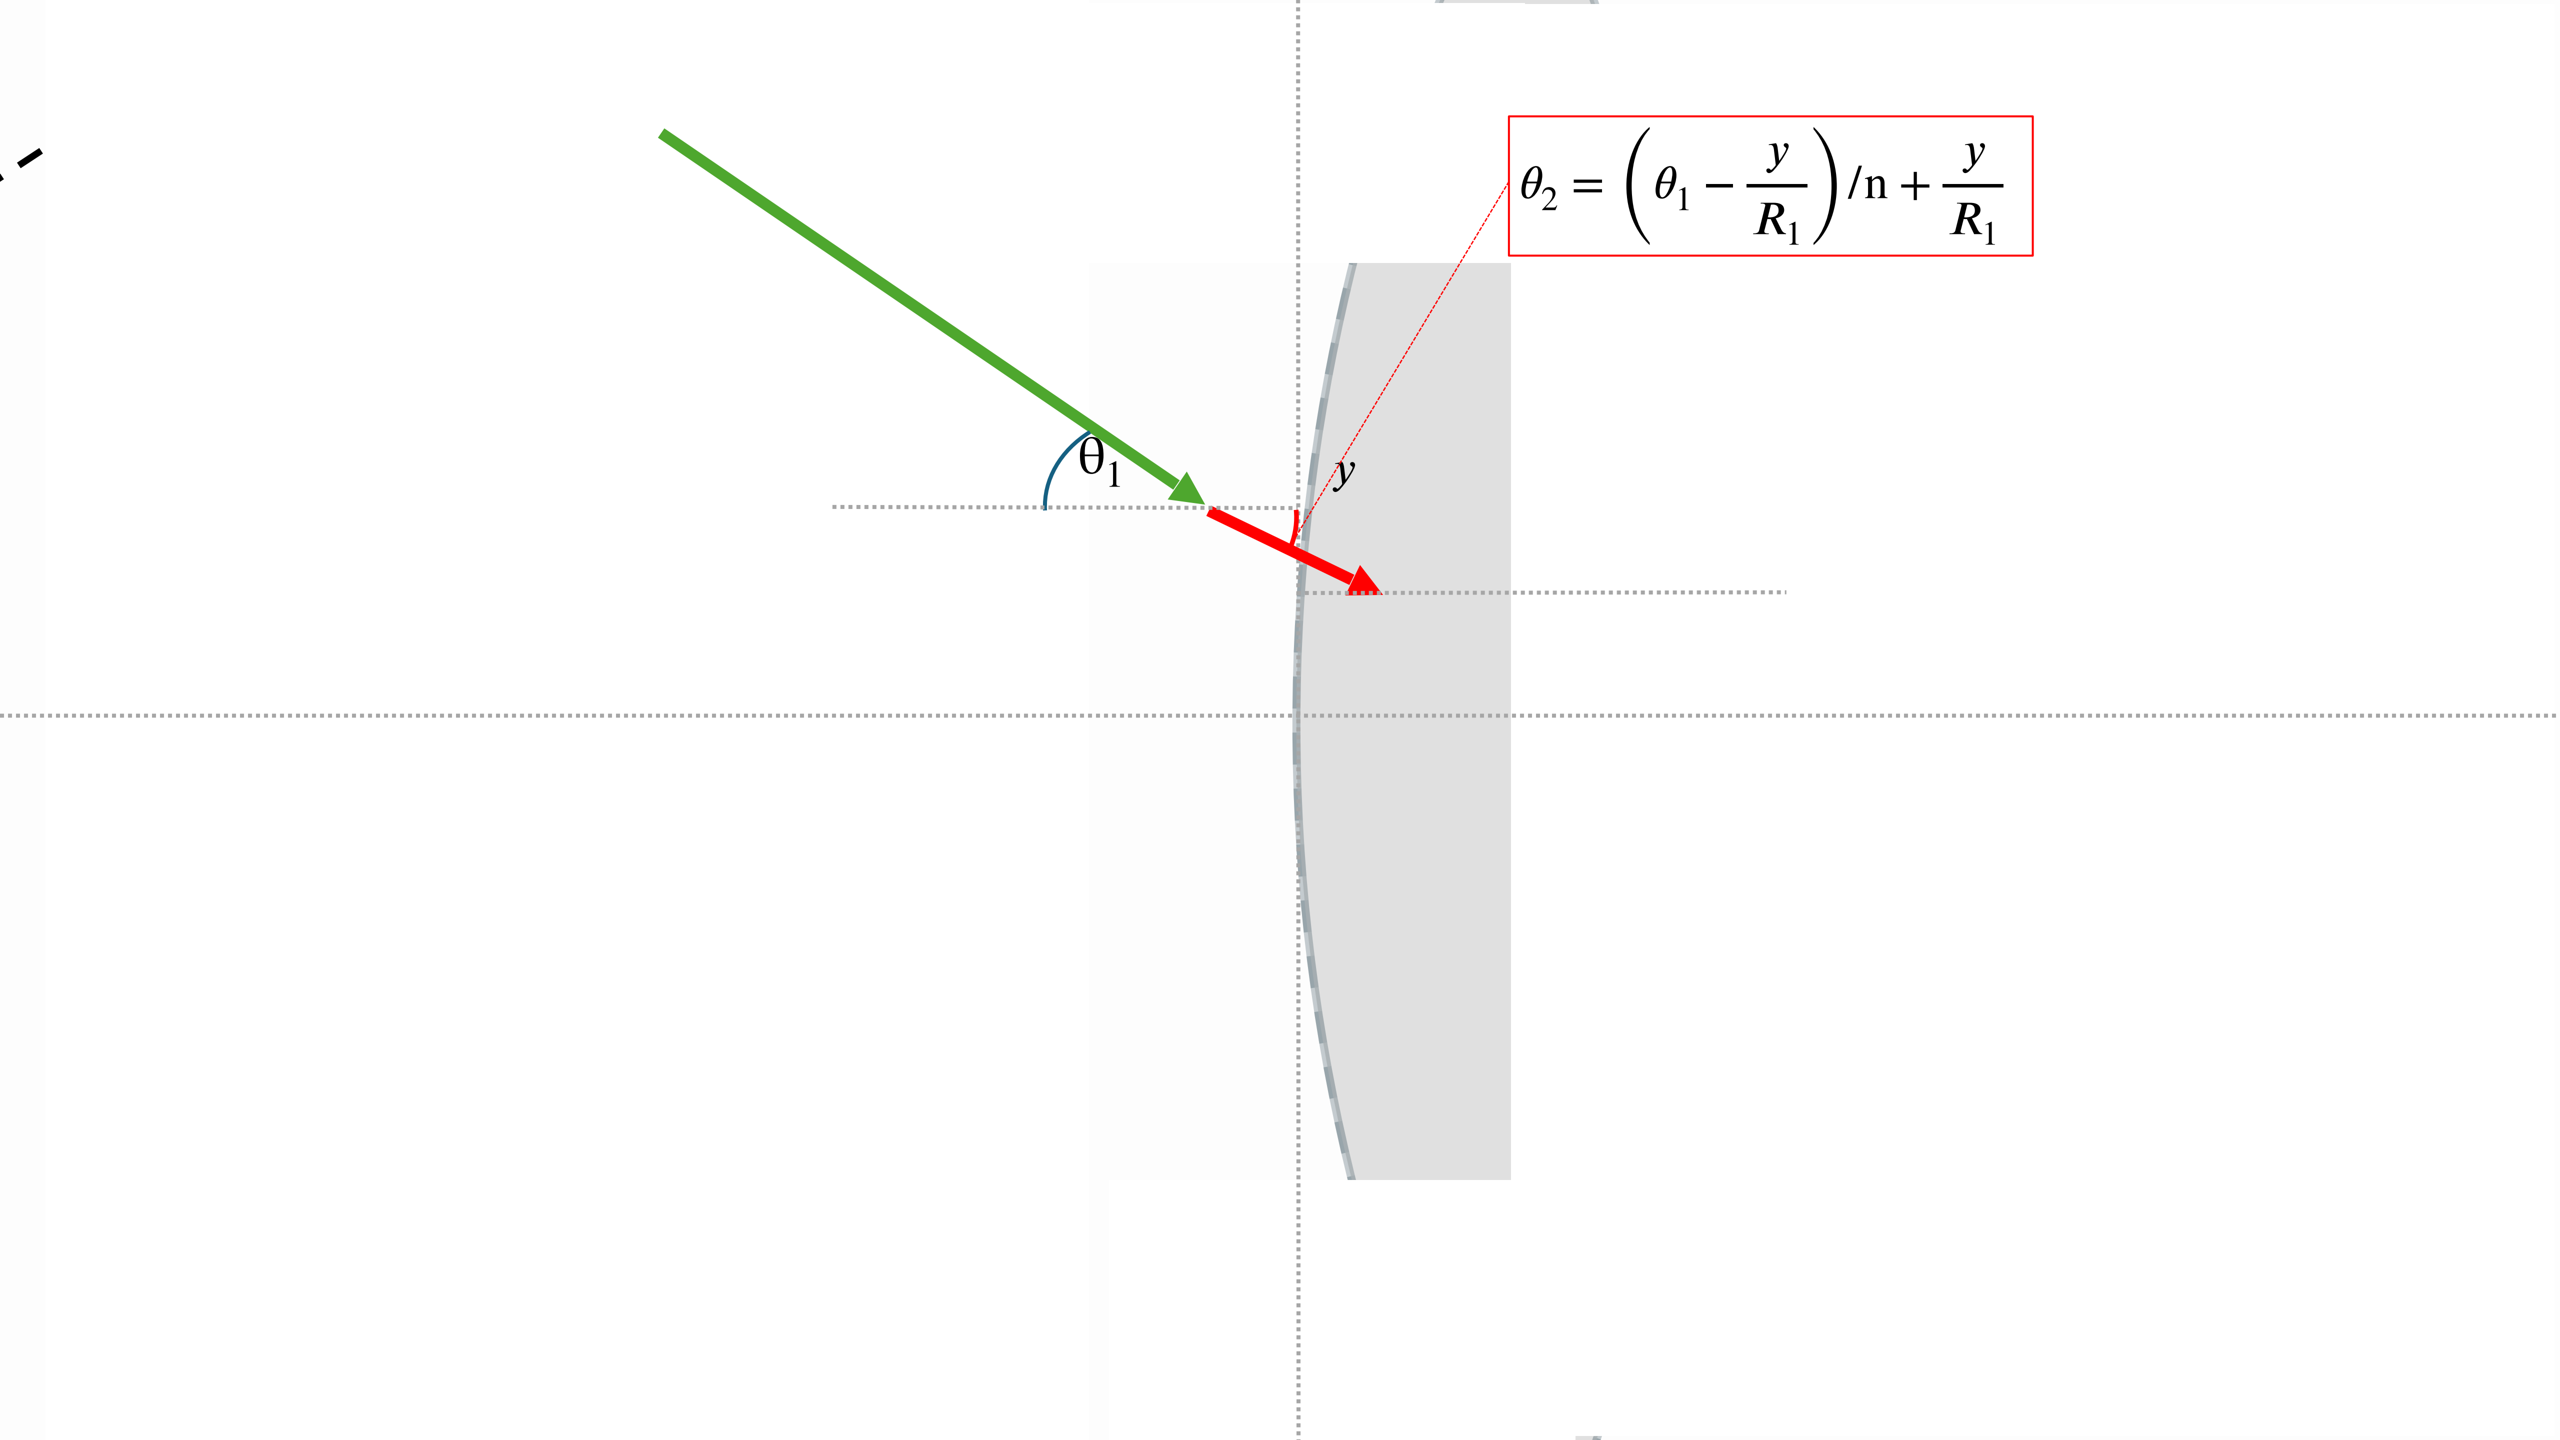

$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\text{n} + \frac{y}{R_1}$
θ_1
y

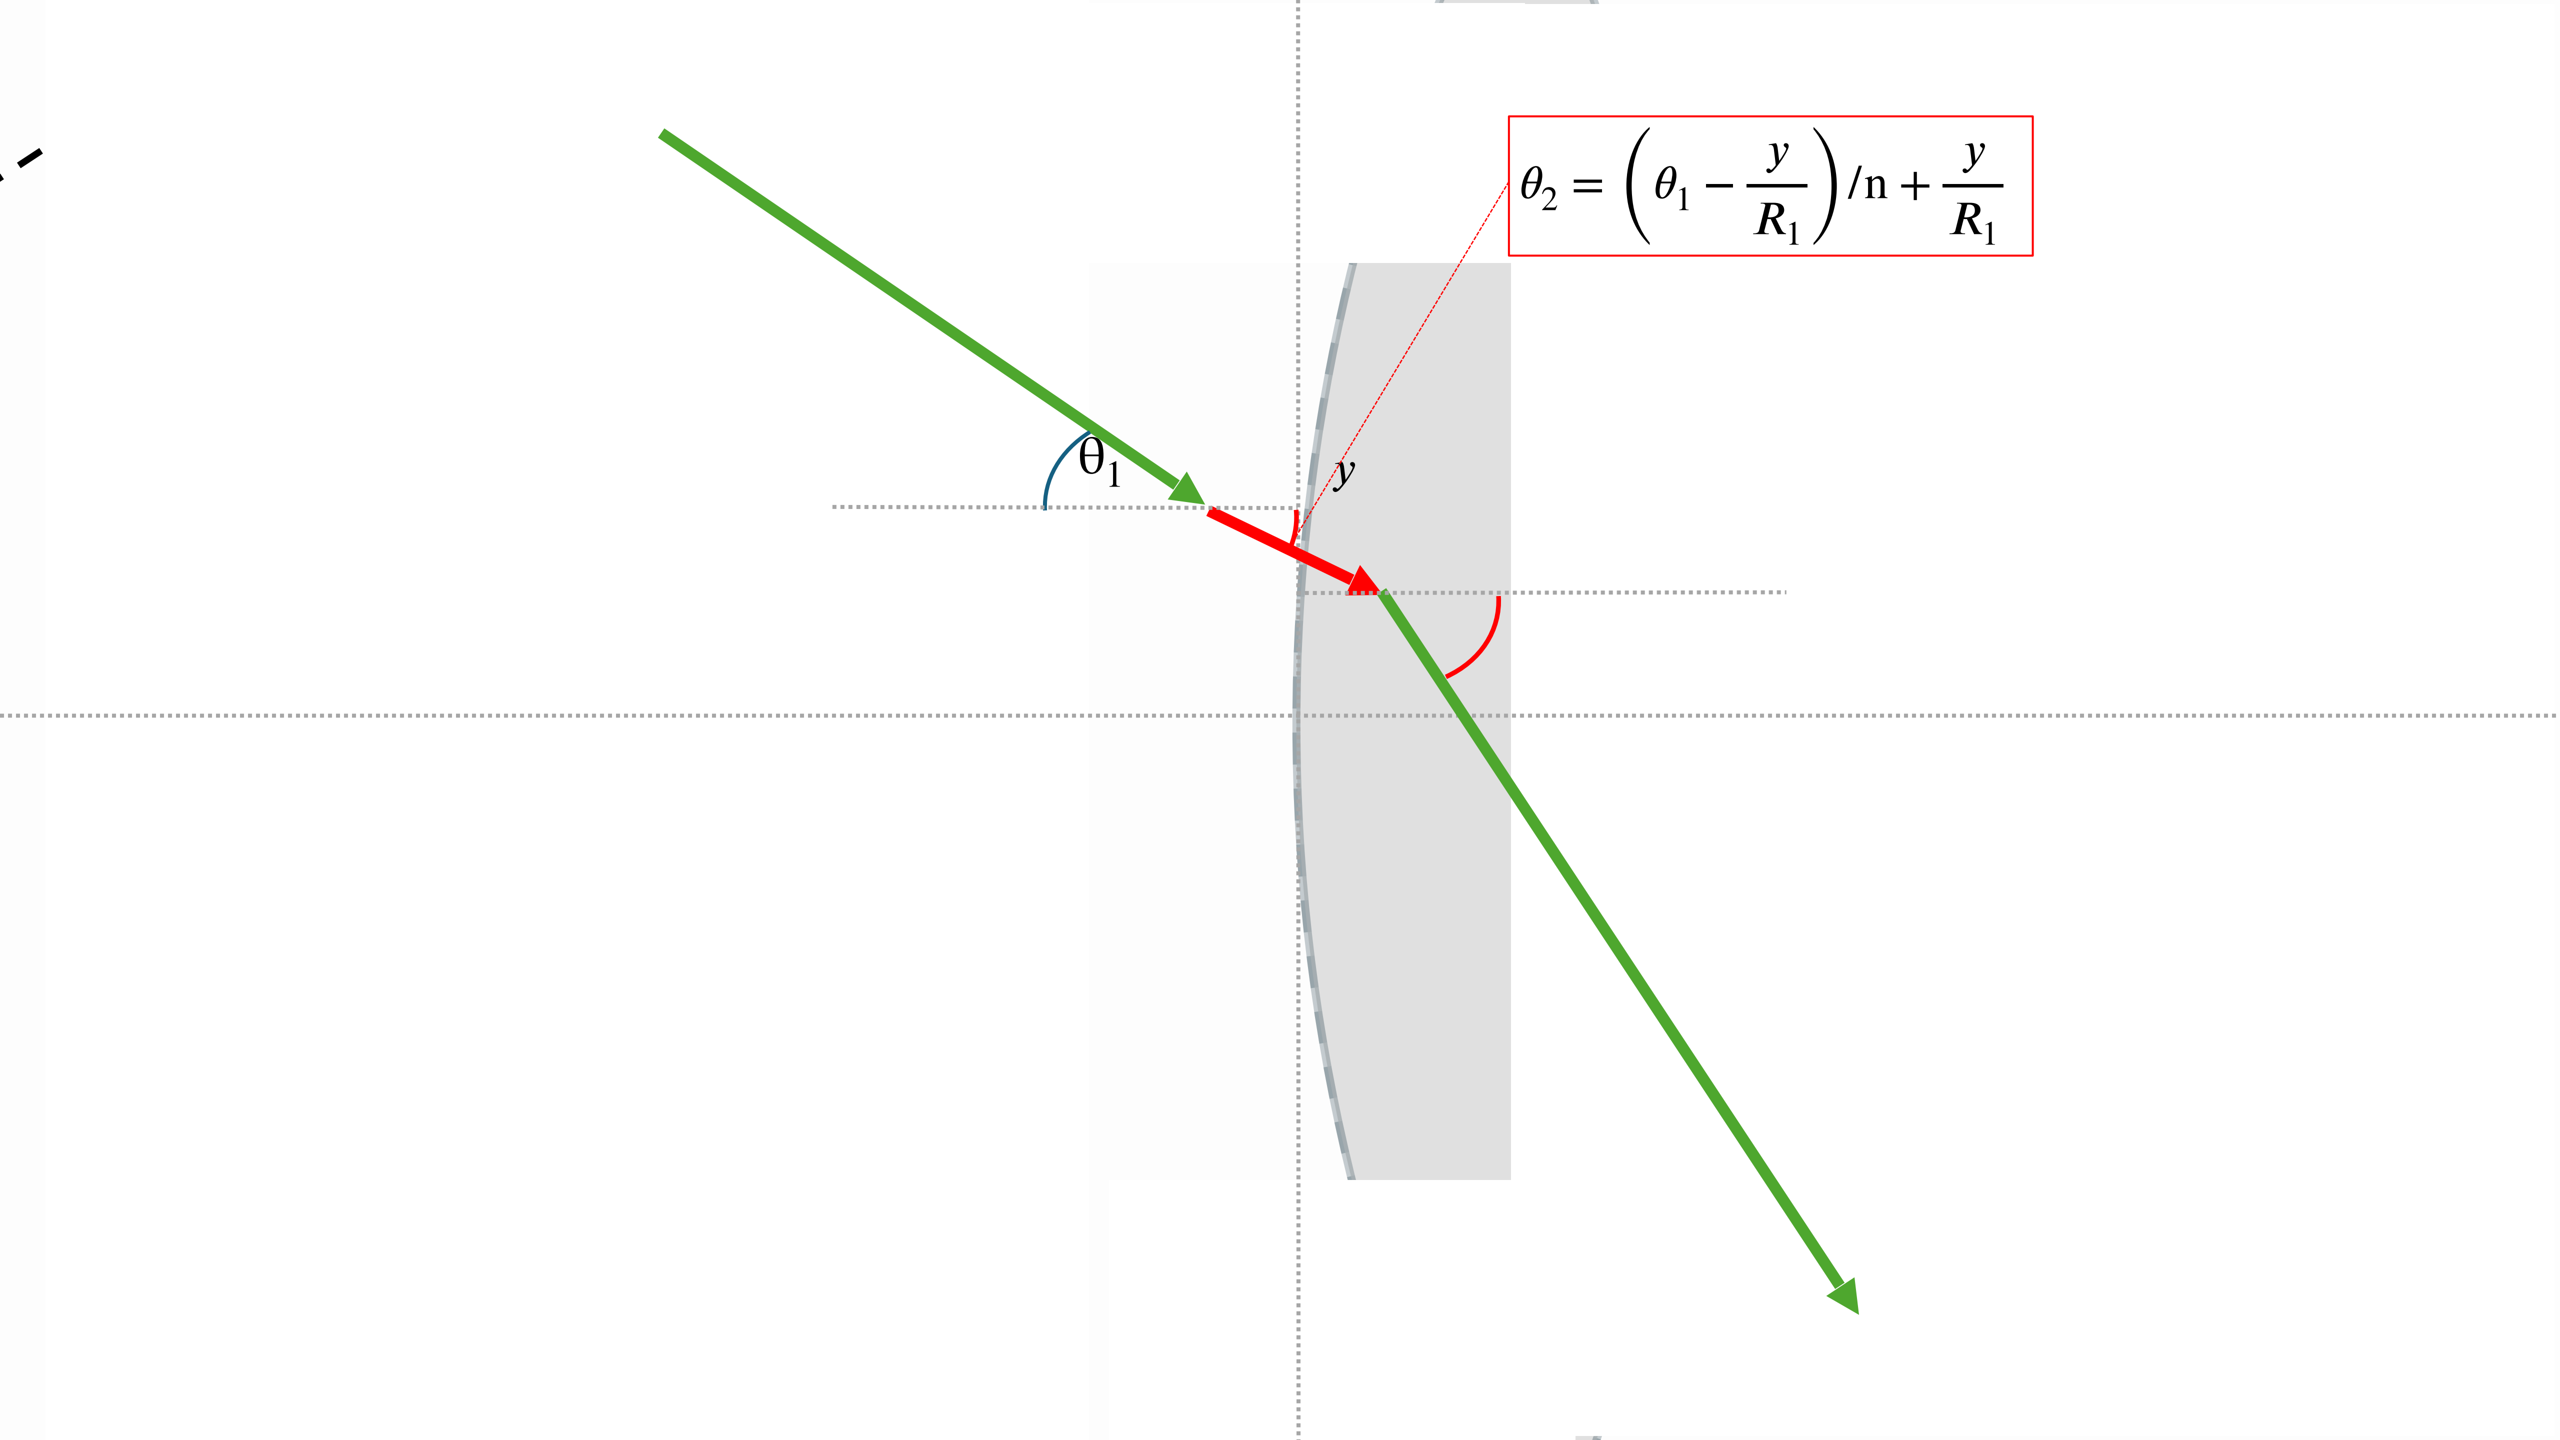

$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$
θ_1
y

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$$

$$\begin{aligned}\theta_3 &= \left(\theta_2 + \frac{y}{R_2}\right)/\mathrm{n} + \frac{y}{R_2} \\ &= \boldsymbol{\theta_1 + y(1-n)\left(\frac{1}{R_1} + \frac{1}{R_2}\right)}\end{aligned}$$

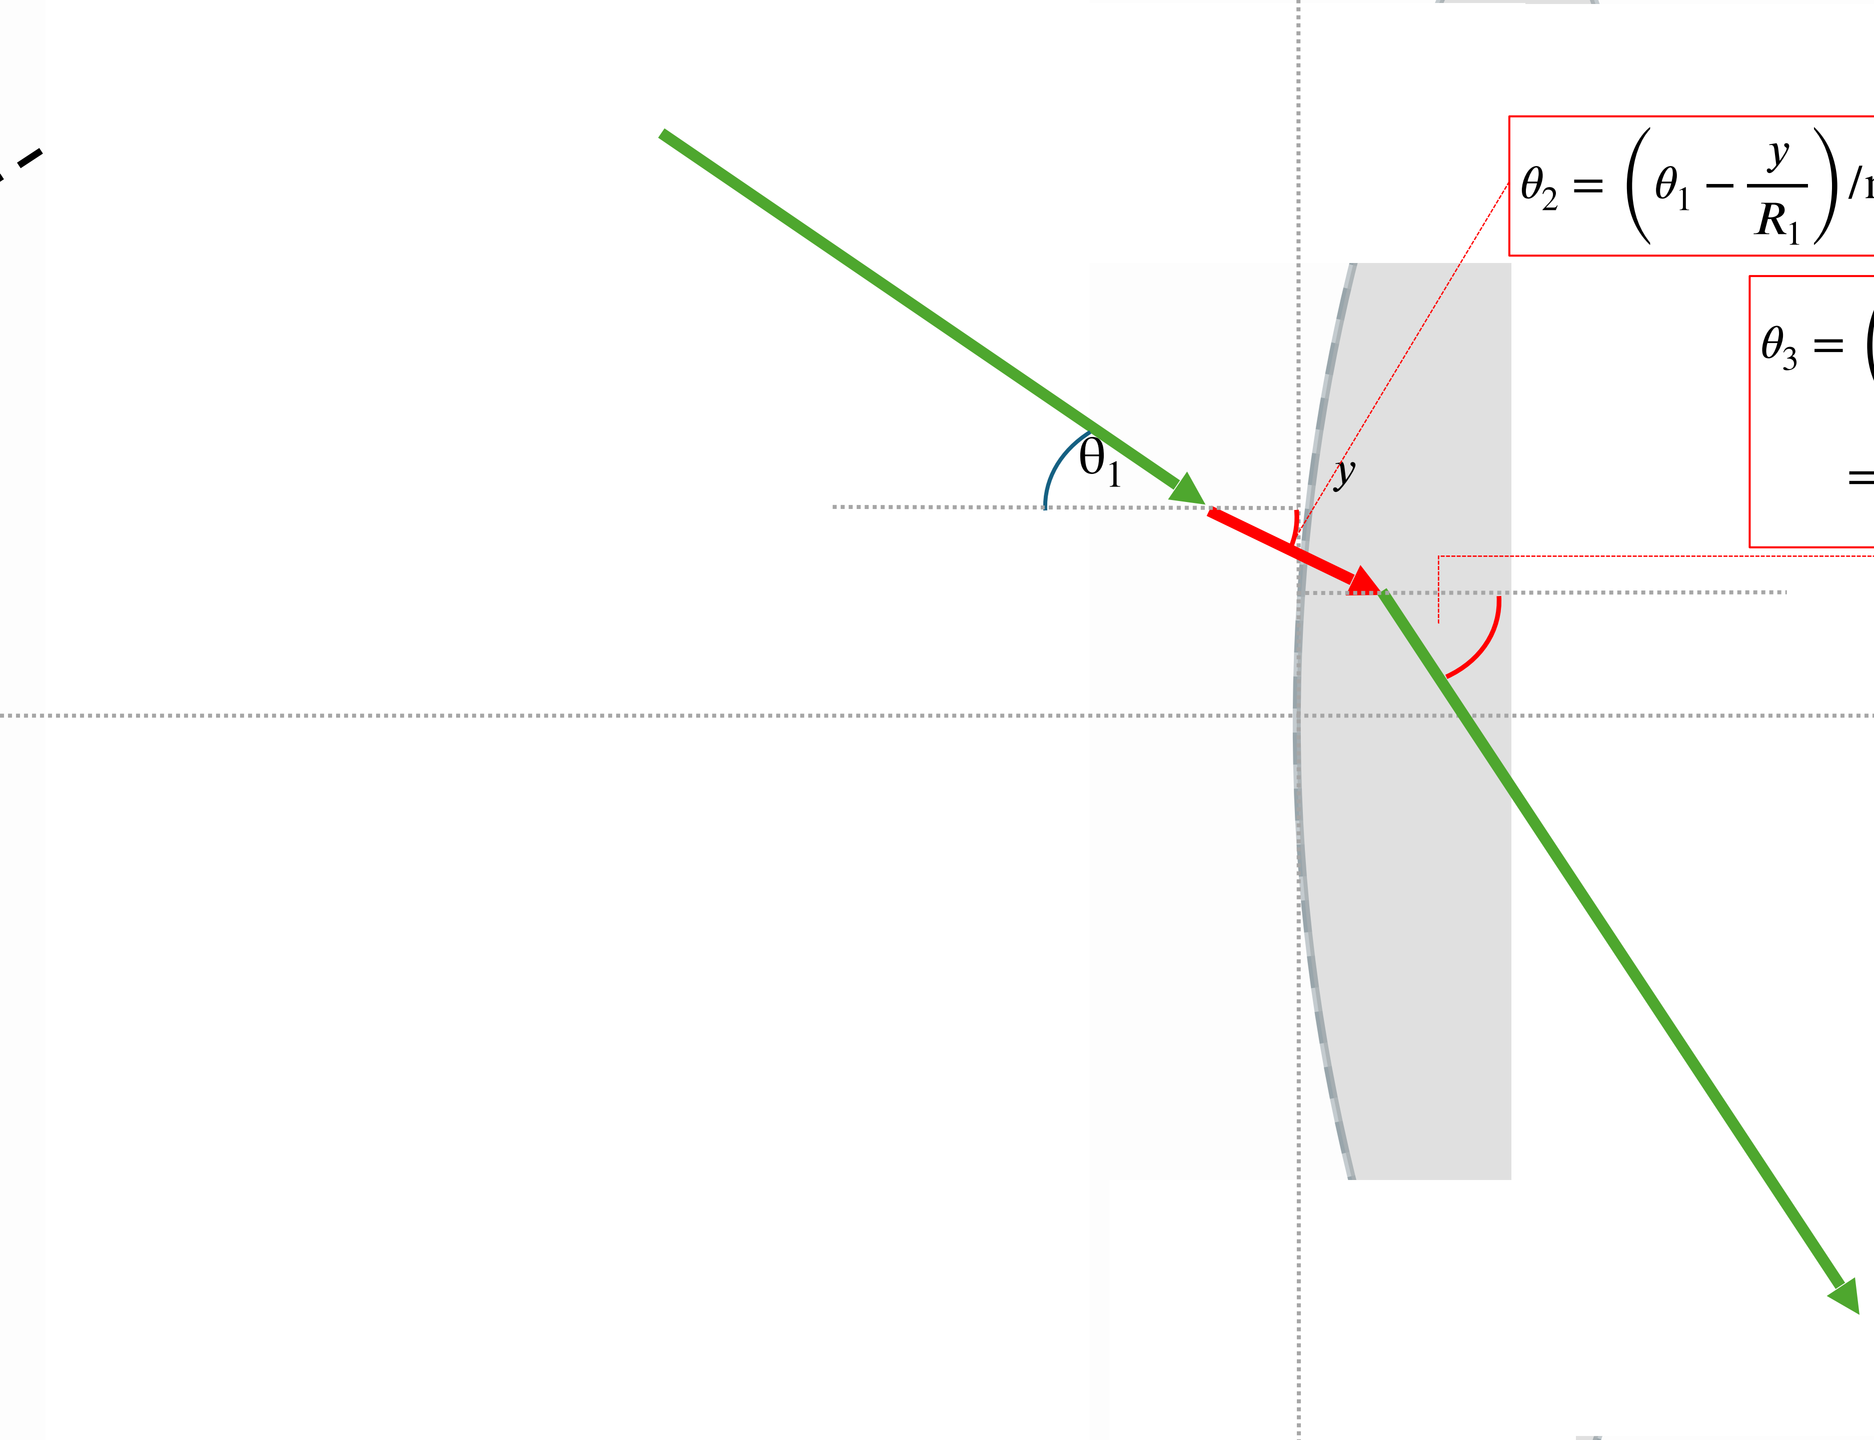

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\text{n} + \frac{y}{R_1}$$

$$\theta_3 = \left(\theta_2 + \frac{y}{R_2}\right)/\text{n} + \frac{y}{R_2}$$

$$= \boldsymbol{\theta_1} + \underbrace{\boldsymbol{y(1-n)\left(\frac{1}{R_1} + \frac{1}{R_2}\right)}}_{\boldsymbol{\frac{1}{f}}}$$

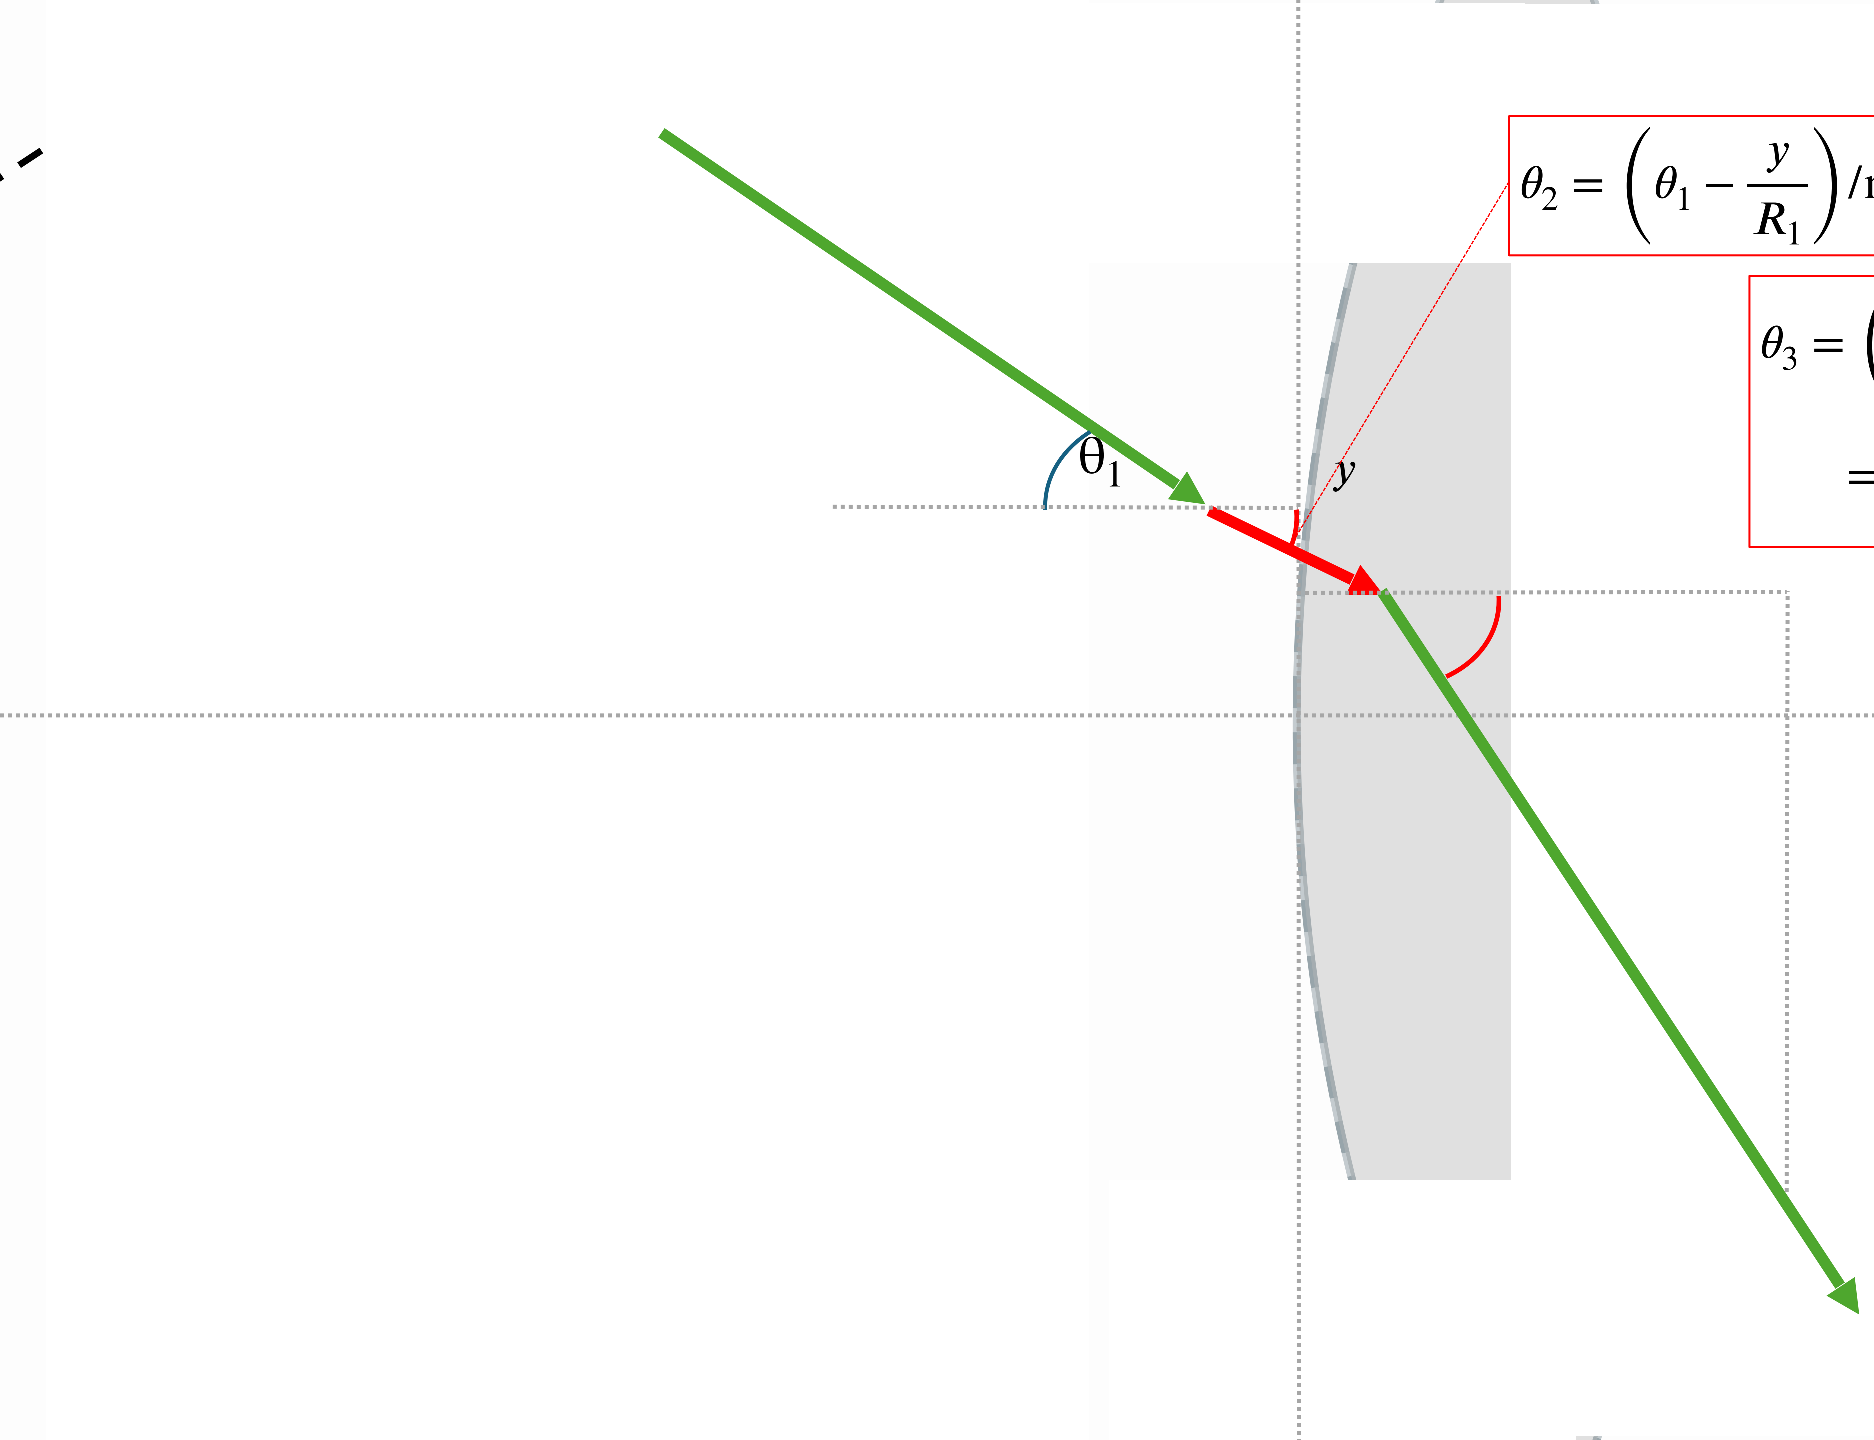

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$$

$$\theta_3 = \left(\theta_2 + \frac{y}{R_2}\right)/\mathrm{n} + \frac{y}{R_2}$$
$$= \boldsymbol{\theta_1} + \underbrace{\boldsymbol{y(1-n)\left(\frac{1}{R_1} + \frac{1}{R_2}\right)}}_{\boldsymbol{\frac{1}{f}}}$$

$$: \theta_3 = \theta_1 + y/f$$

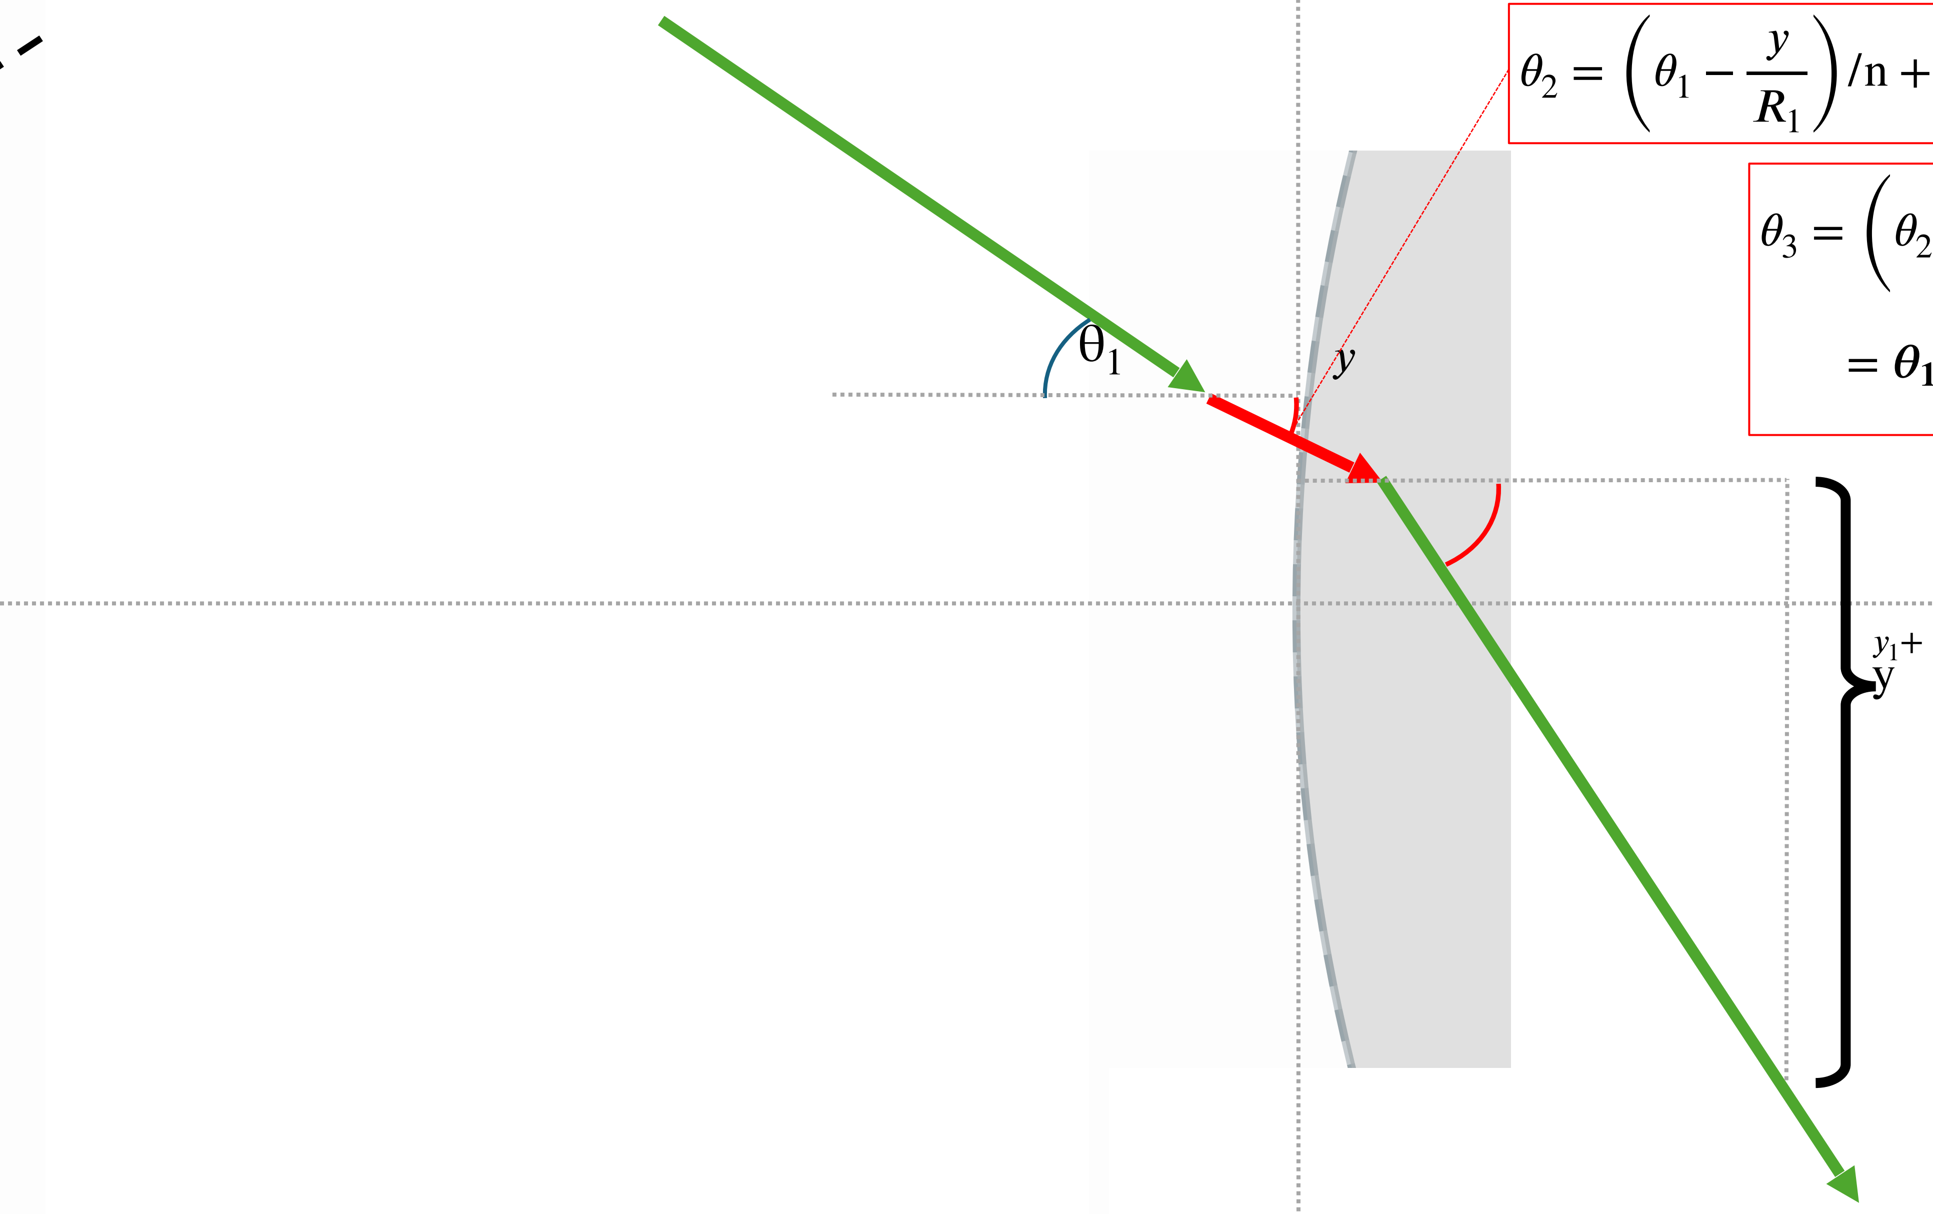

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\text{n} + \frac{y}{R_1}$$

$$\theta_3 = \left(\theta_2 + \frac{y}{R_2}\right)/\text{n} + \frac{y}{R_2}$$

$$= \boldsymbol{\theta_1} + \underbrace{\boldsymbol{y(1-n)\left(\frac{1}{R_1} + \frac{1}{R_2}\right)}}_{\boldsymbol{\frac{1}{f}}}$$

$$: \theta_3 = \theta_1 + y/f$$

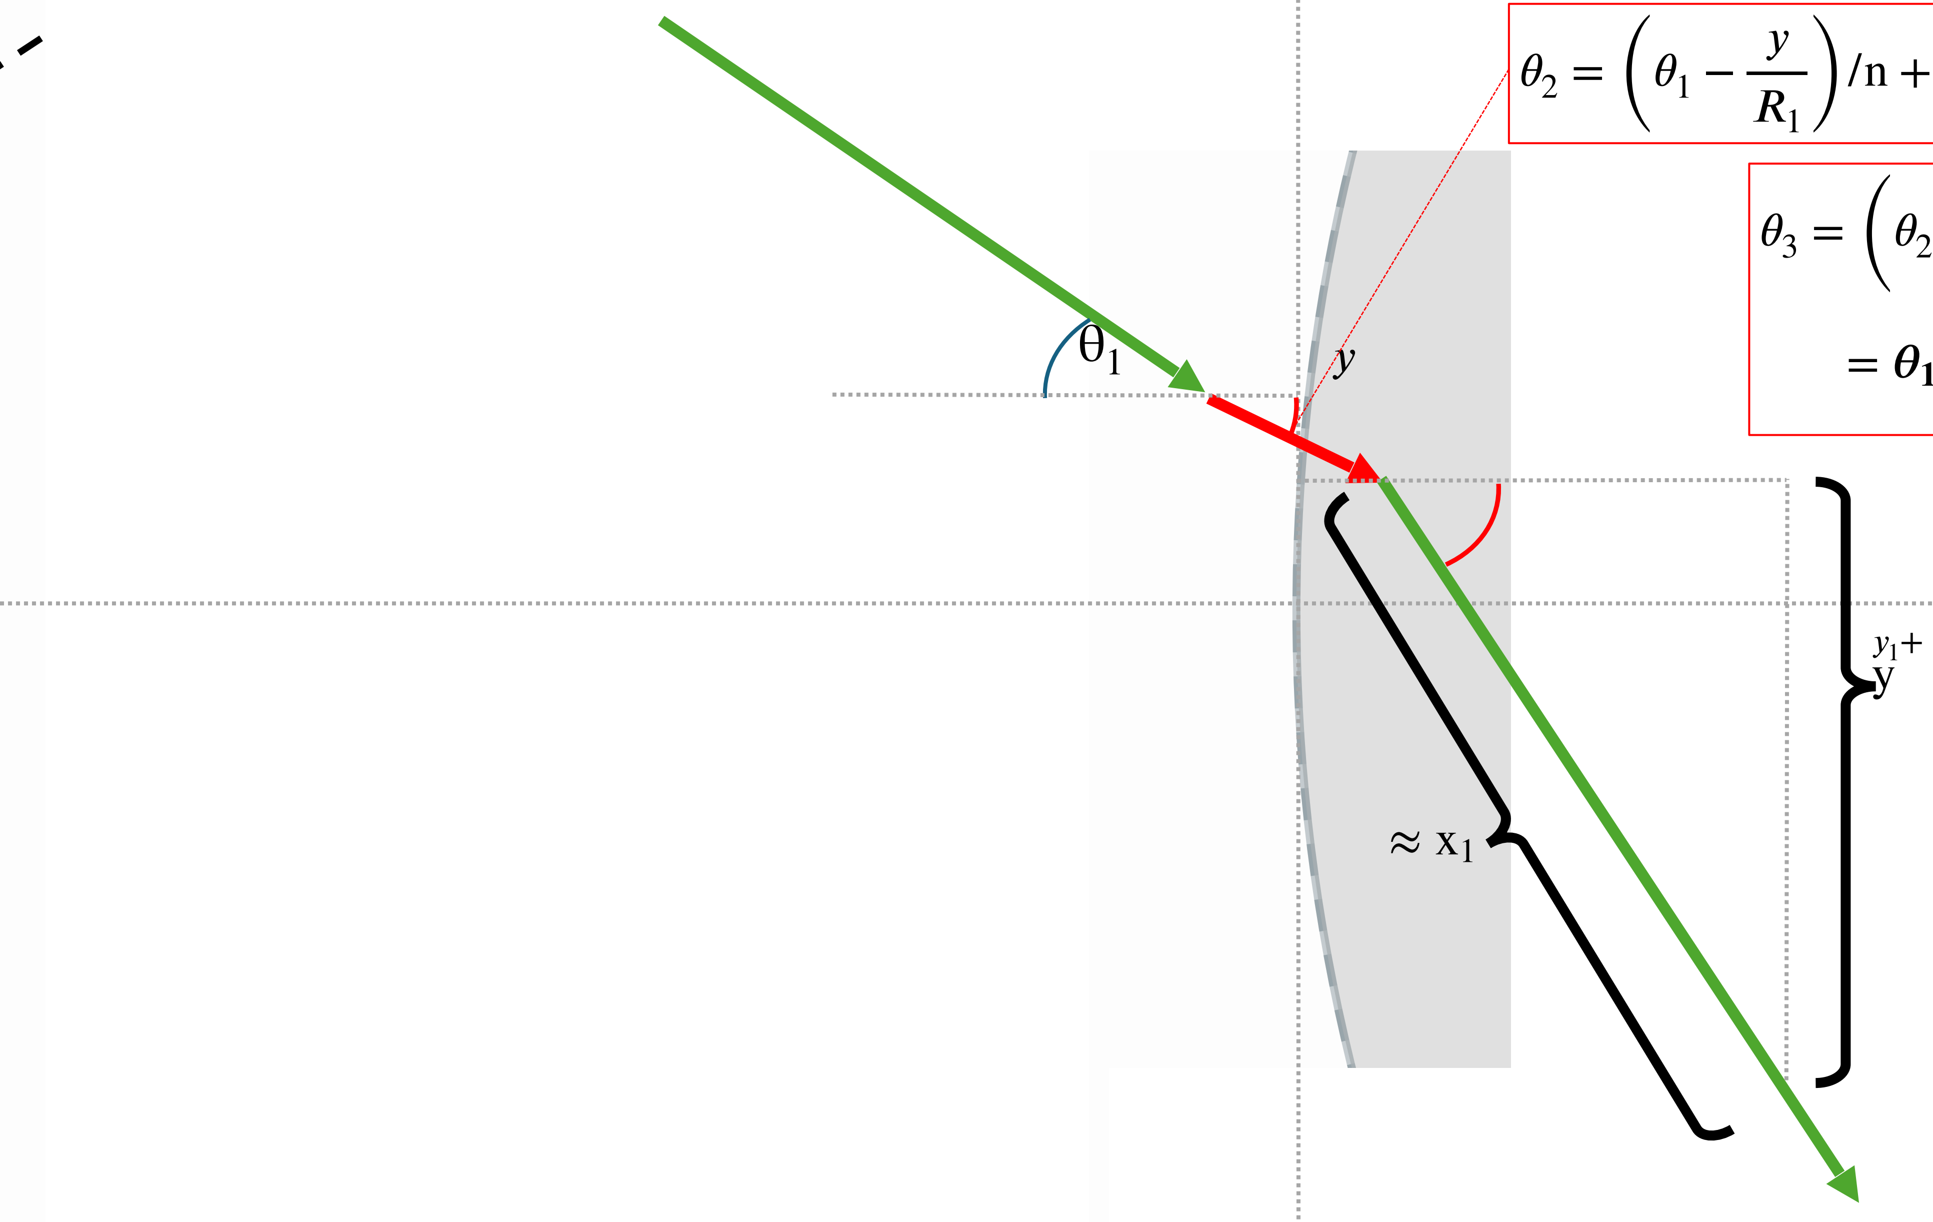

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$$

$$\theta_3 = \left(\theta_2 + \frac{y}{R_2}\right)/\mathrm{n} + \frac{y}{R_2}$$

$$= \boldsymbol{\theta_1 + y\underbrace{(1-n)\left(\frac{1}{R_1} + \frac{1}{R_2}\right)}_{\frac{1}{f}}}$$

$$: \theta_3 = \theta_1 + y/f$$

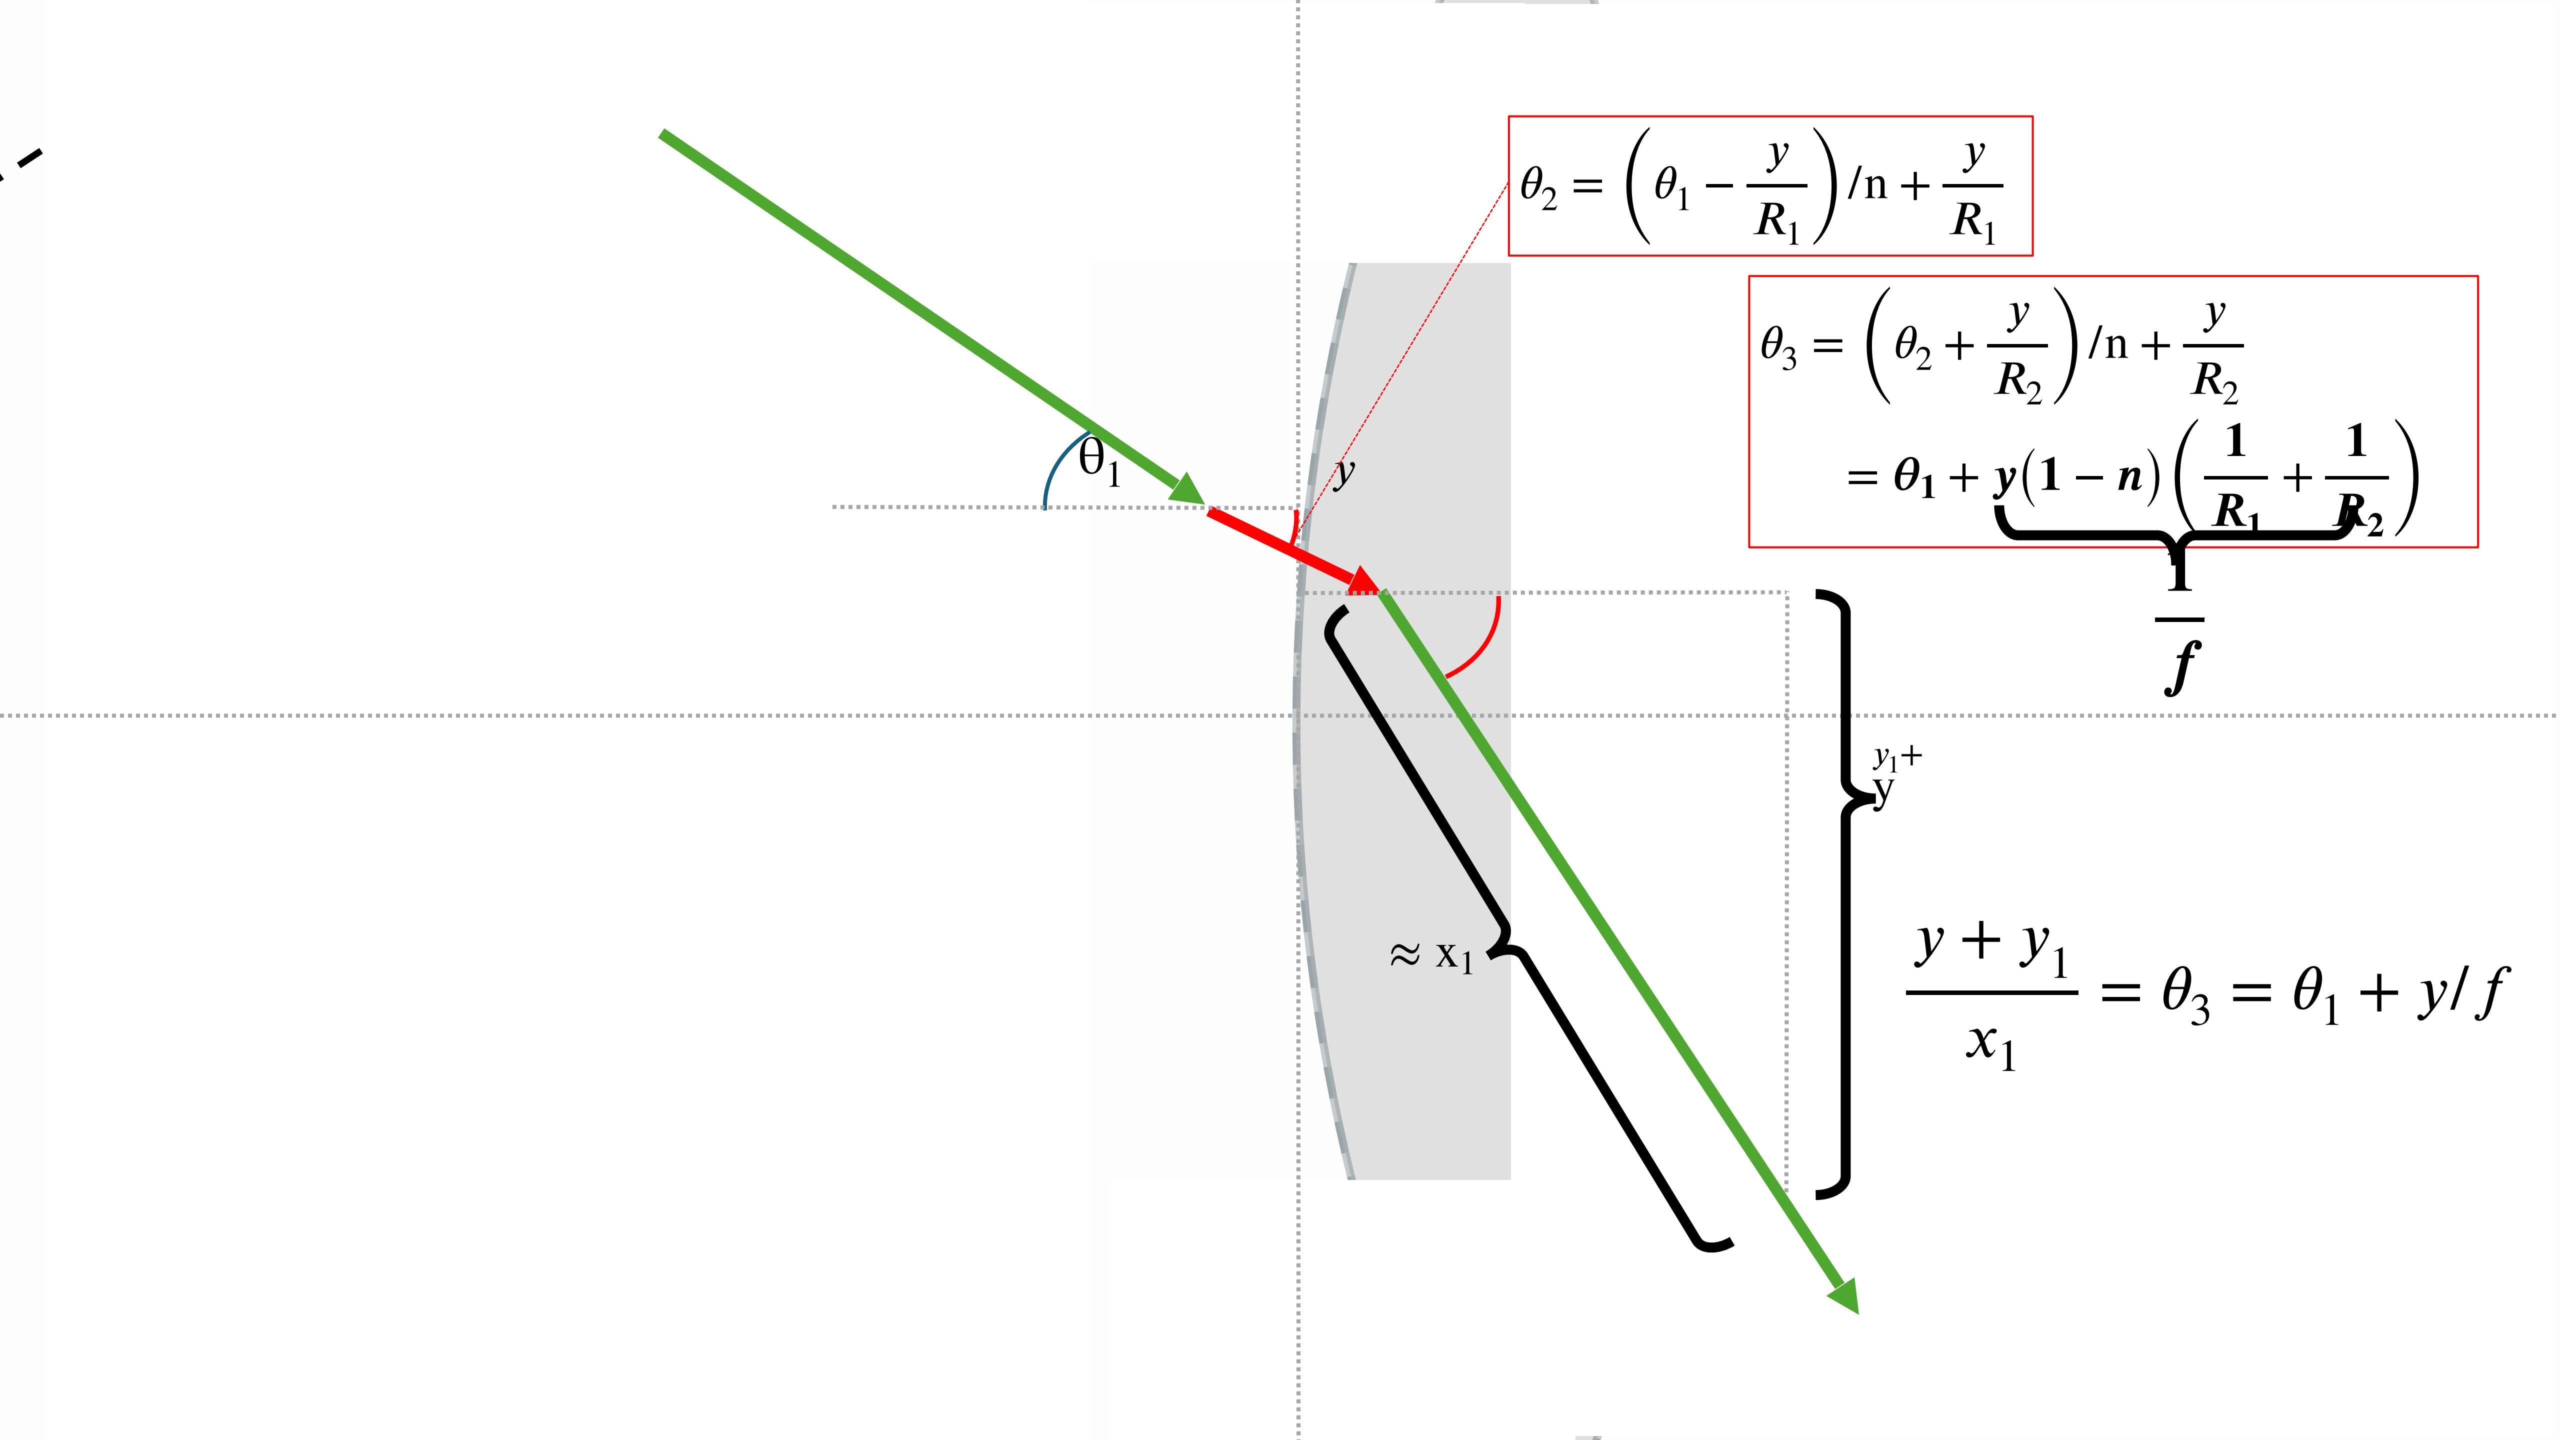

θ_1
y
$\approx x_1$
$y_1 +$ y
$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$
$\theta_3 = \left(\theta_2 + \frac{y}{R_2}\right)/\mathrm{n} + \frac{y}{R_2}$
$= \boldsymbol{\theta_1 + y(1-n)\left(\frac{1}{R_1} + \frac{1}{R_2}\right)}$
$\frac{1}{f}$
$\frac{y + y_1}{x_1} = \theta_3 = \theta_1 + y/f$

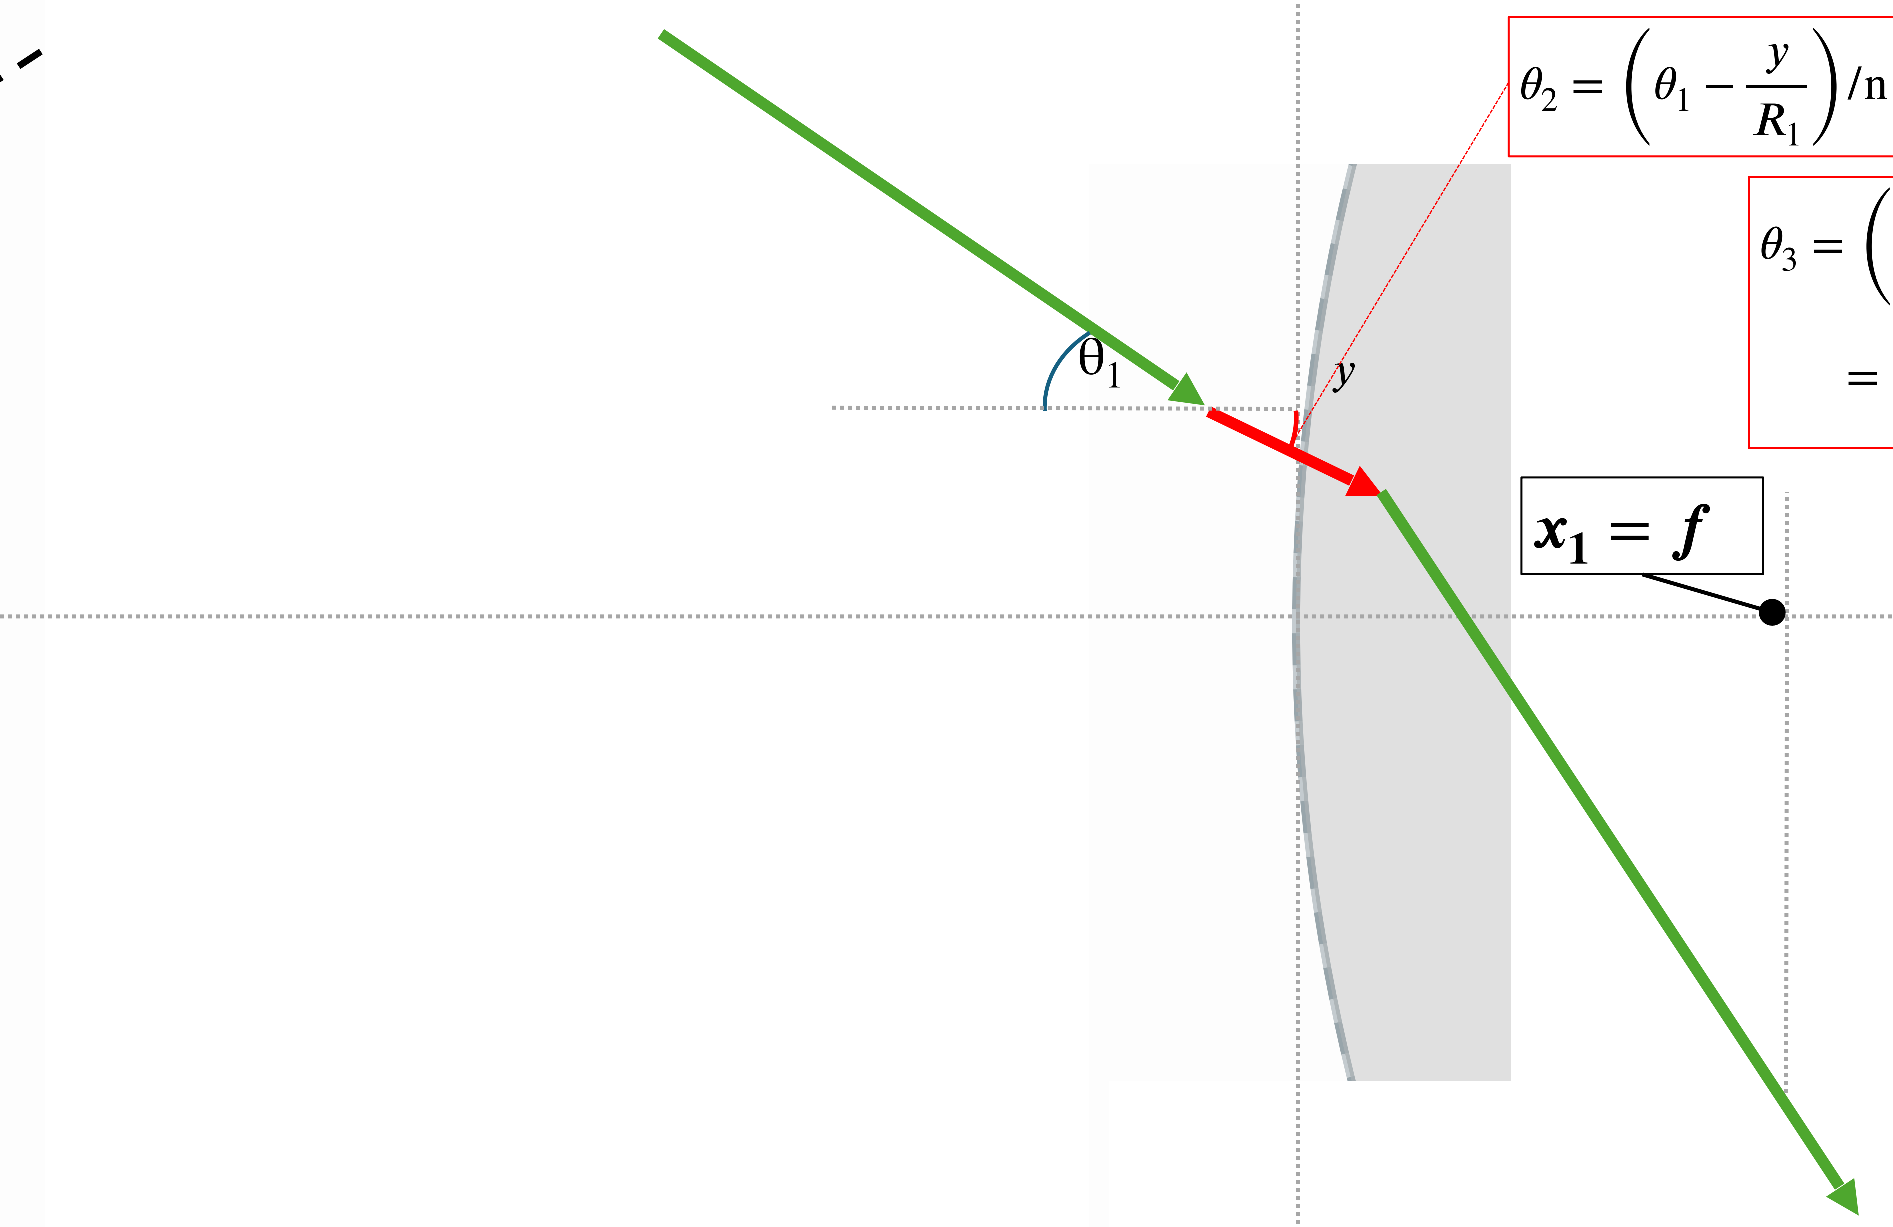

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$$

$$\theta_3 = \left(\theta_2 + \frac{y}{R_2}\right)/\mathrm{n} + \frac{y}{R_2}$$
$$= \boldsymbol{\theta_1} + \underbrace{\boldsymbol{y(1-n)\left(\frac{1}{R_1} + \frac{1}{R_2}\right)}}_{\frac{1}{f}}$$

$$\frac{y + y_1}{x_1} = \theta_3 = \theta_1 + y/f$$

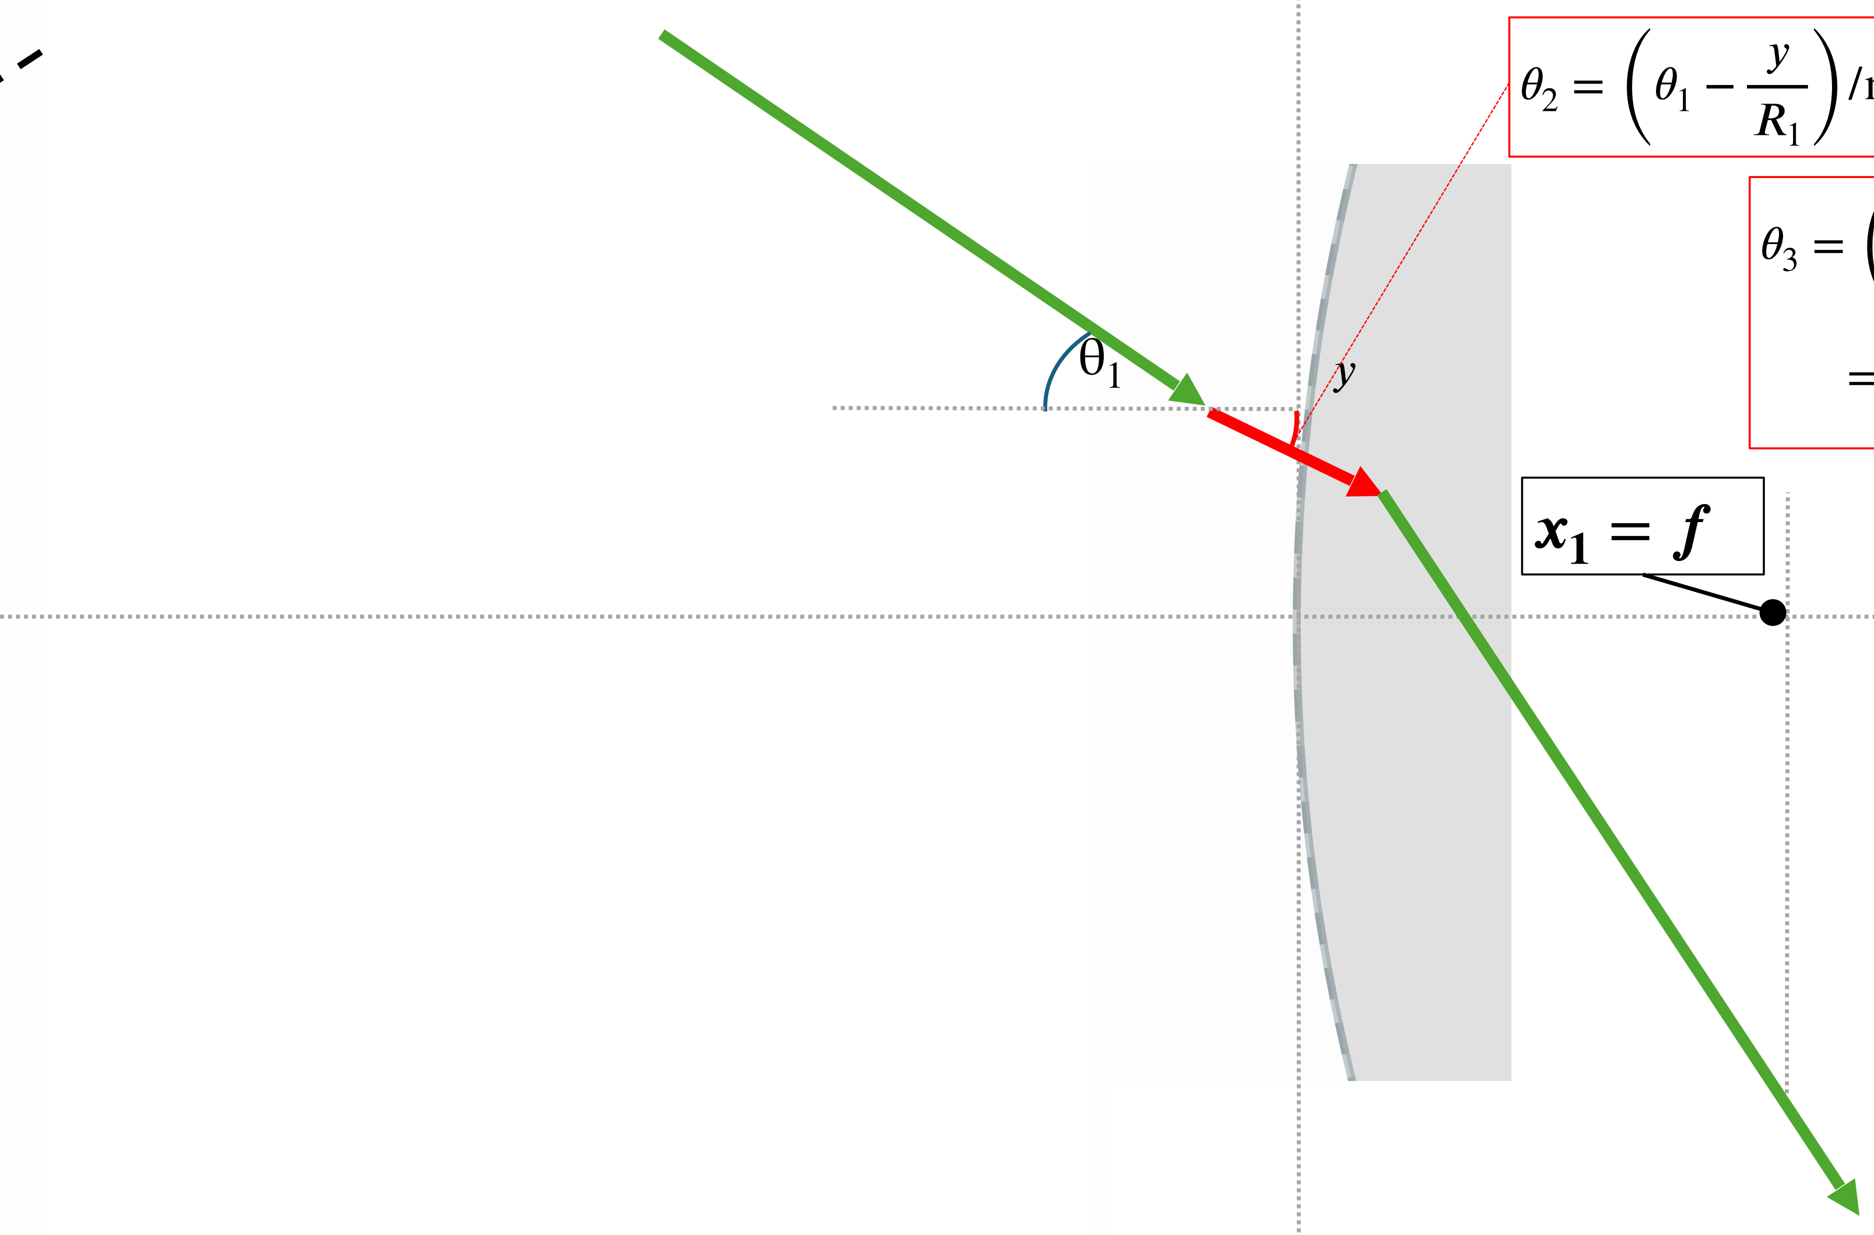

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$$

$$\theta_3 = \left(\theta_2 + \frac{y}{R_2}\right)/\mathrm{n} + \frac{y}{R_2}$$

$$= \theta_1 + \underbrace{y(1-n)\left(\frac{1}{R_1} + \frac{1}{R_2}\right)}_{\frac{1}{f}}$$

$$\frac{y + y_1}{\cancel{x_1}\, f} = \theta_3 = \theta_1 + y/f$$

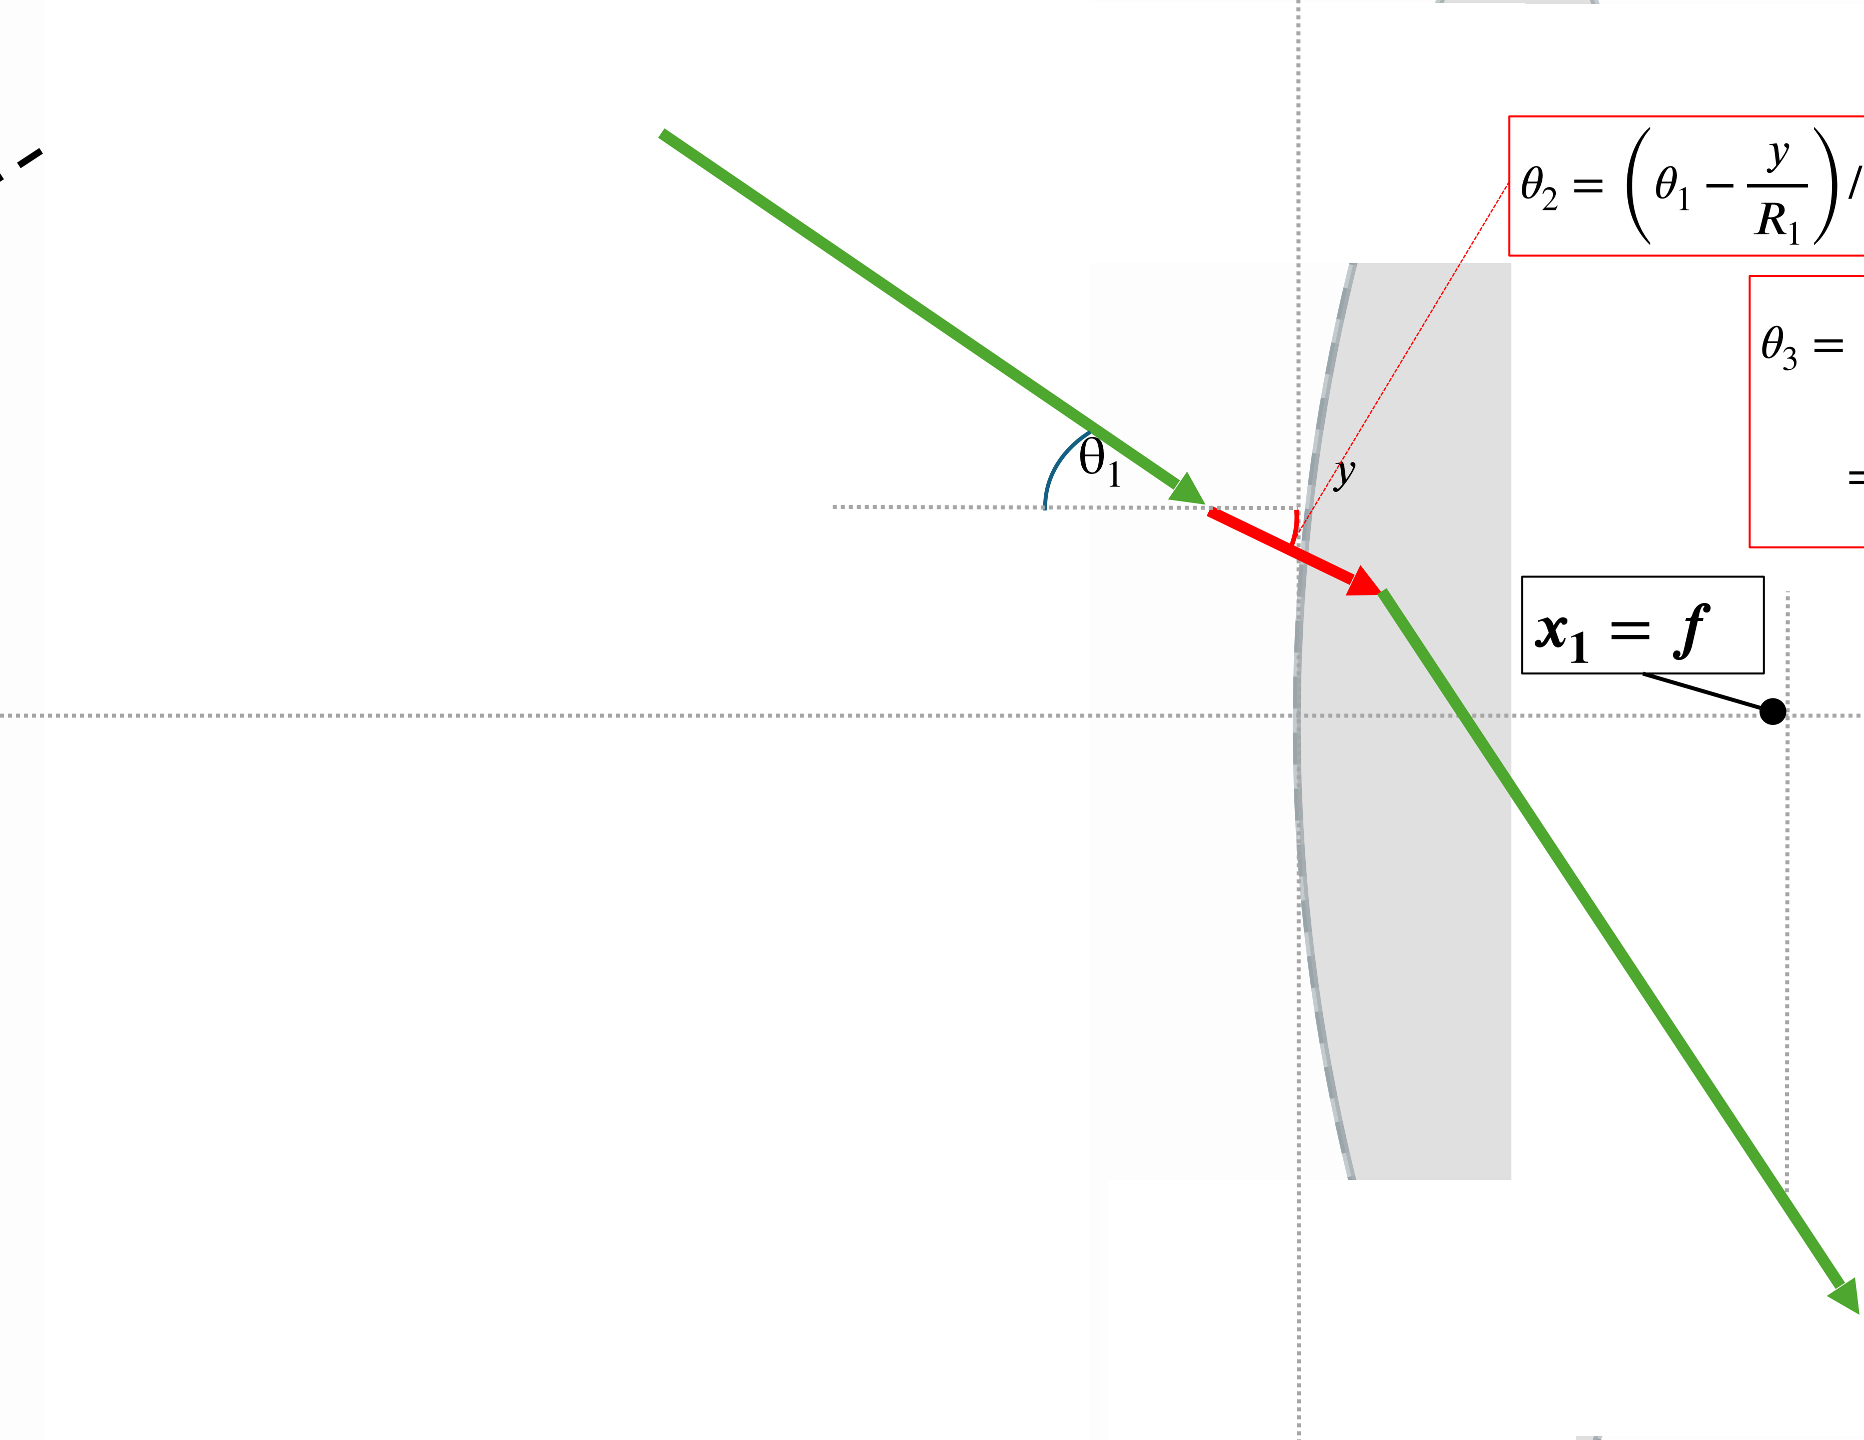

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$$

$$\theta_3 = \left(\theta_2 + \frac{y}{R_2}\right)/\mathrm{n} + \frac{y}{R_2}$$

$$= \theta_1 + \underbrace{y(1-n)\left(\frac{1}{R_1} + \frac{1}{R_2}\right)}_{\frac{1}{f}}$$

$$\frac{y + y_1}{\cancel{x_1}\, f} = \theta_3 = \theta_1 + y/f$$

$$y_1 + y = f\theta_1 + y$$

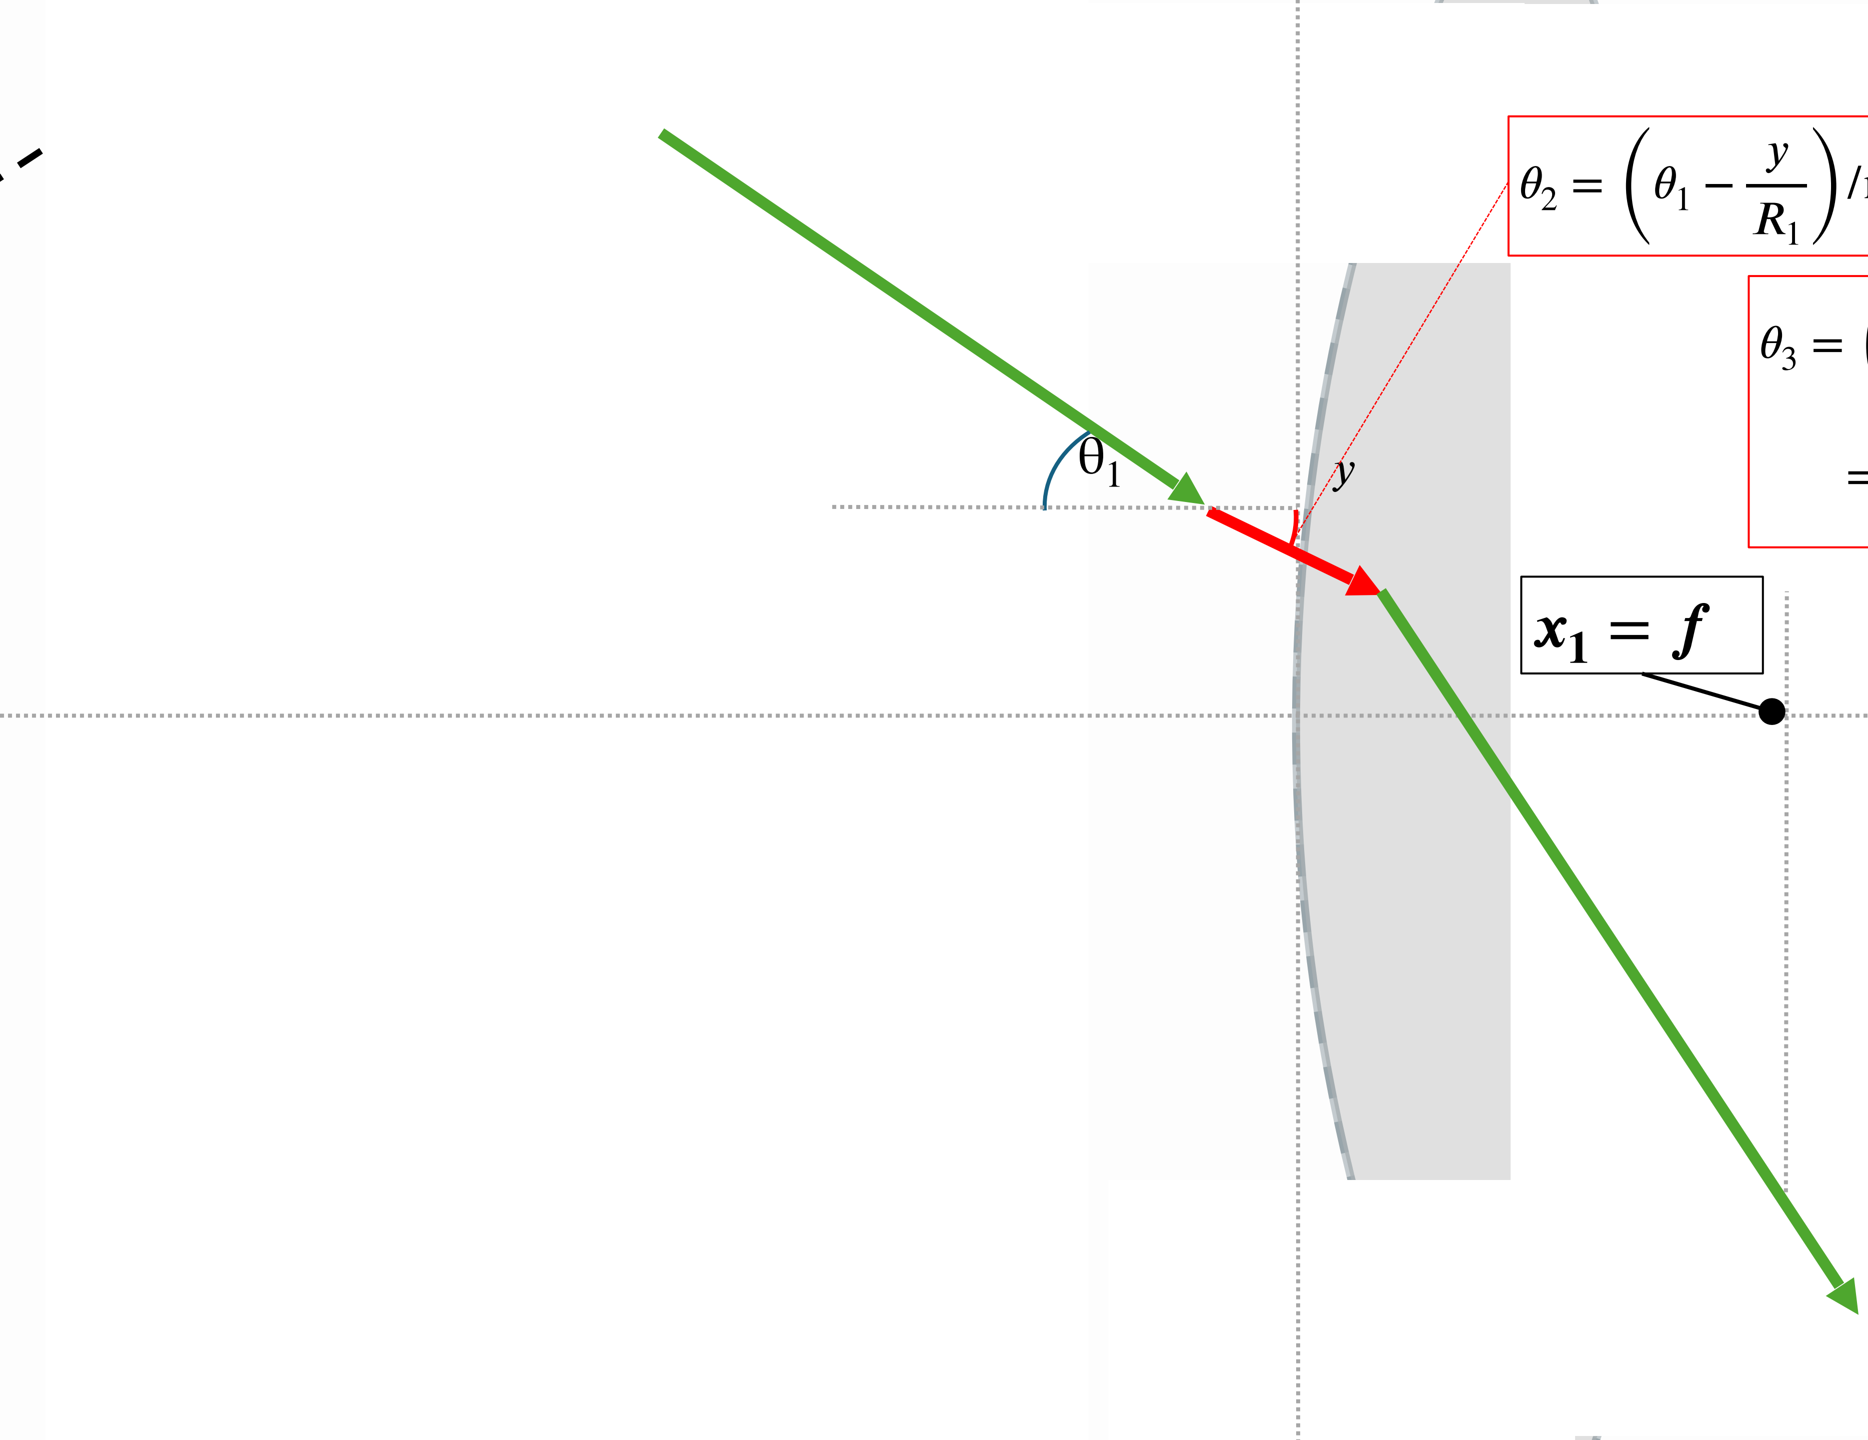

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$$

$$\theta_3 = \left(\theta_2 + \frac{y}{R_2}\right)/\mathrm{n} + \frac{y}{R_2}$$
$$= \theta_1 + \underbrace{y(1-n)\left(\frac{1}{R_1} + \frac{1}{R_2}\right)}_{\frac{1}{f}}$$

$$\frac{y + y_1}{\cancel{x_1}\,f} = \theta_3 = \theta_1 + y/f$$

$$y_1 + y = f\theta_1 + y$$

$$\boxed{y_1 = f\theta_1}$$

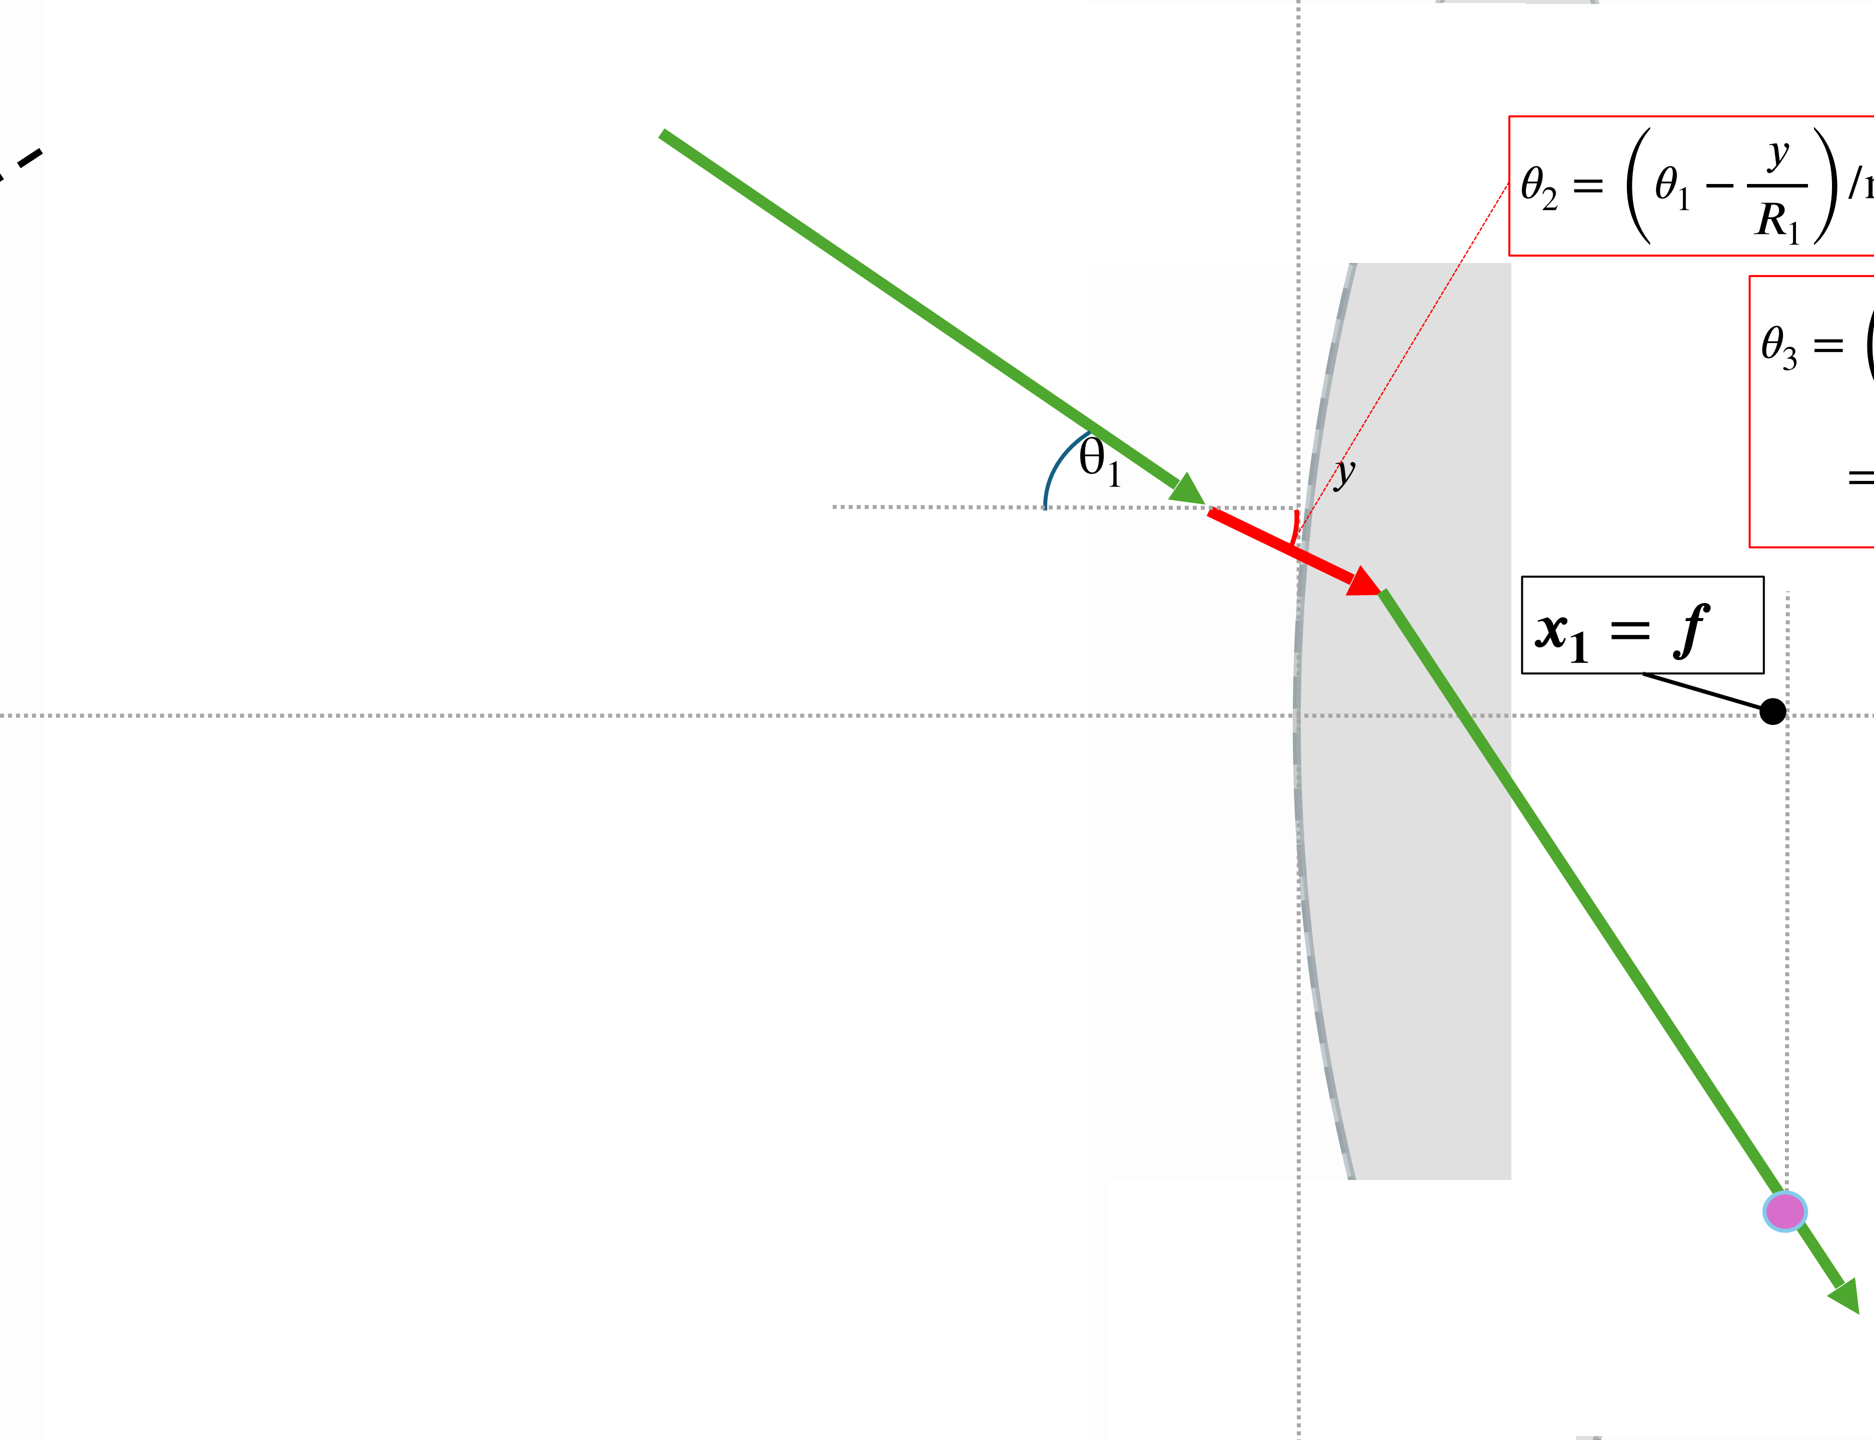

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$$

$$\theta_3 = \left(\theta_2 + \frac{y}{R_2}\right)/\mathrm{n} + \frac{y}{R_2}$$

$$= \boldsymbol{\theta_1} + \boldsymbol{y}\underbrace{(\mathbf{1} - \boldsymbol{n})\left(\frac{\mathbf{1}}{\boldsymbol{R_1}} + \frac{\mathbf{1}}{\boldsymbol{R_2}}\right)}_{\frac{1}{f}}$$

$$\frac{y + y_1}{\cancel{x_1}\, f} = \theta_3 = \theta_1 + y/f$$

$$y_1 + y = f\theta_1 + y$$

$$\boxed{\boldsymbol{y_1 = f\theta_1}}$$

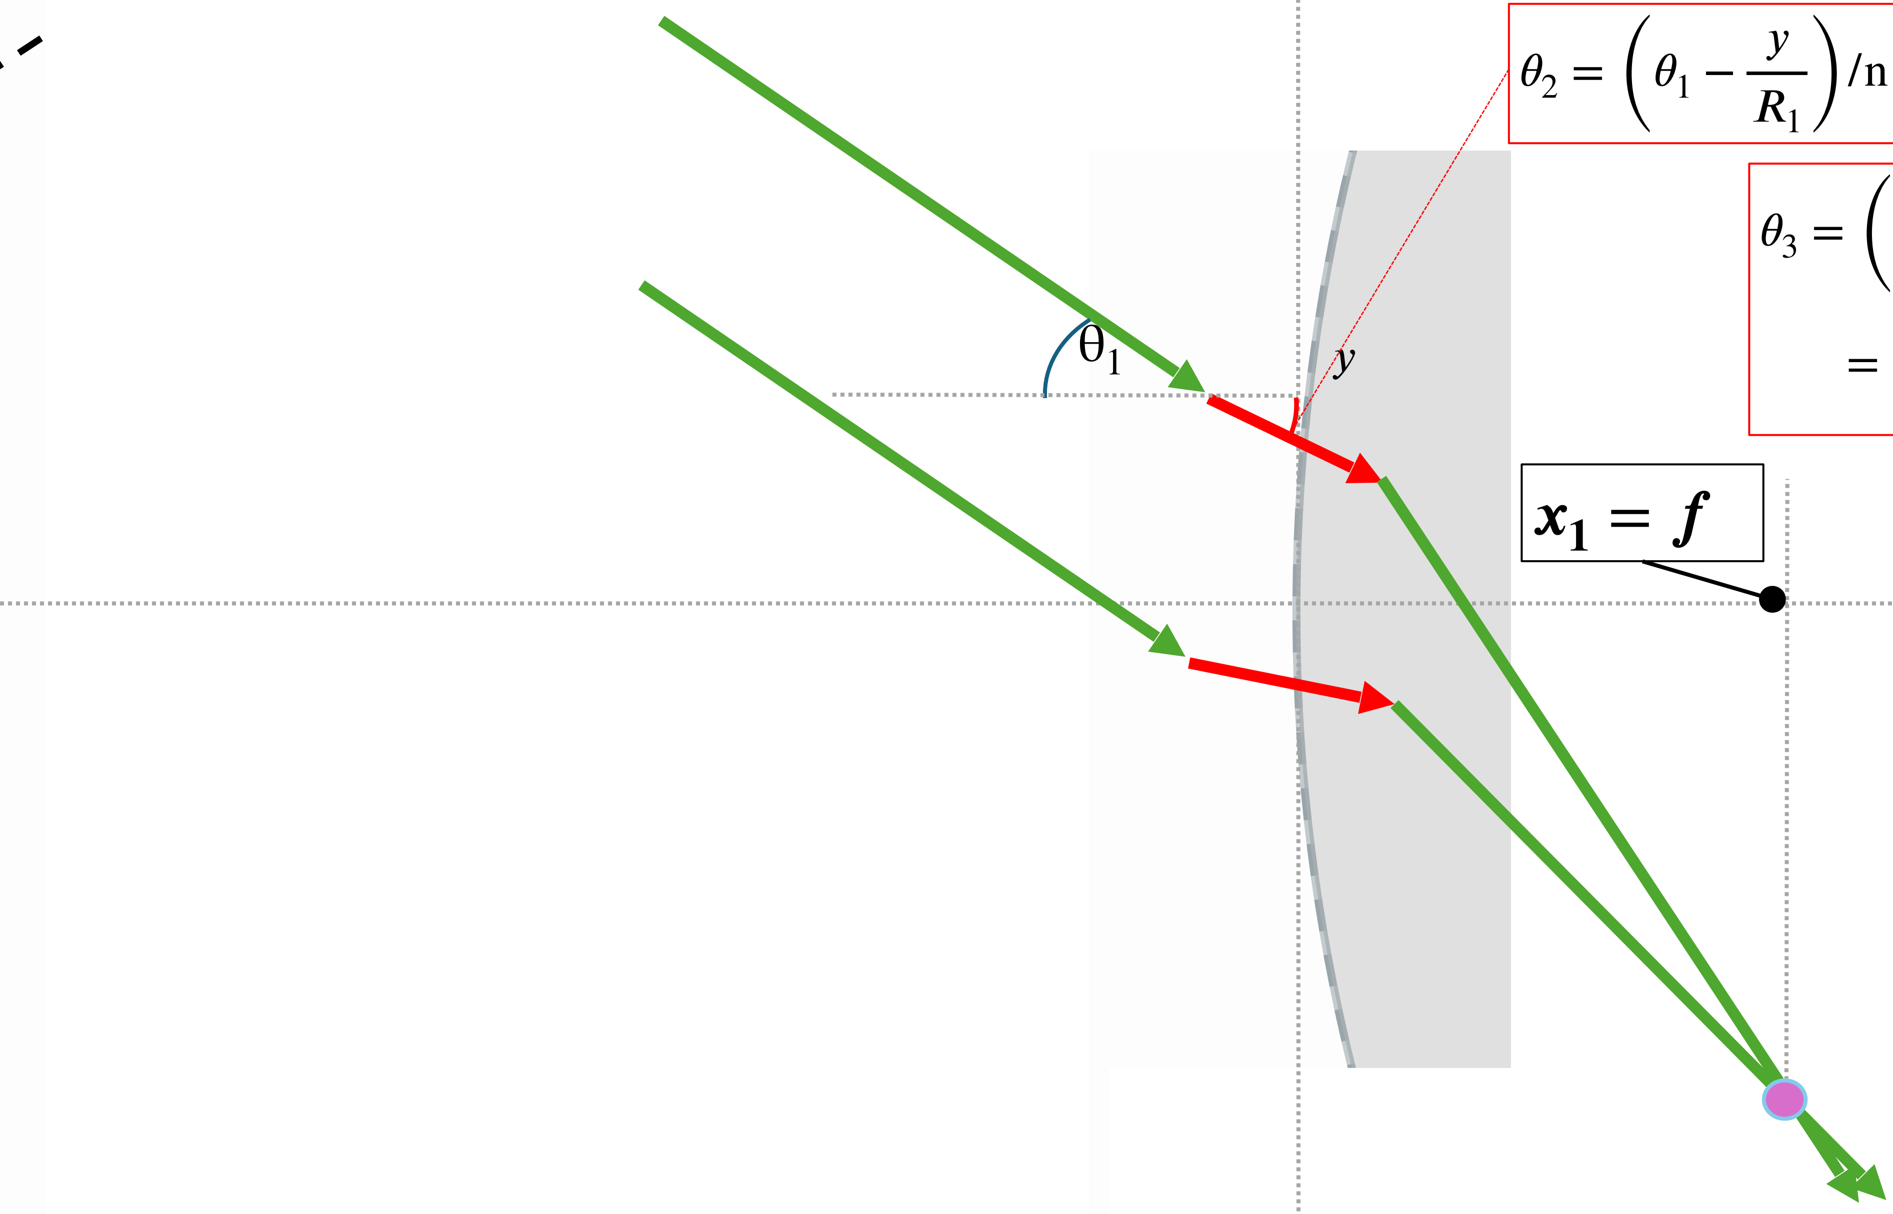

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\text{n} + \frac{y}{R_1}$$

$$\theta_3 = \left(\theta_2 + \frac{y}{R_2}\right)/\text{n} + \frac{y}{R_2}$$

$$= \theta_1 + \underbrace{y(1-n)\left(\frac{1}{R_1} + \frac{1}{R_2}\right)}_{\frac{1}{f}}$$

$$\frac{y + y_1}{f} = \theta_3 = \theta_1 + y/f$$

$$y_1 + y = f\theta_1 + y$$

$$\boxed{y_1 = f\theta_1}$$

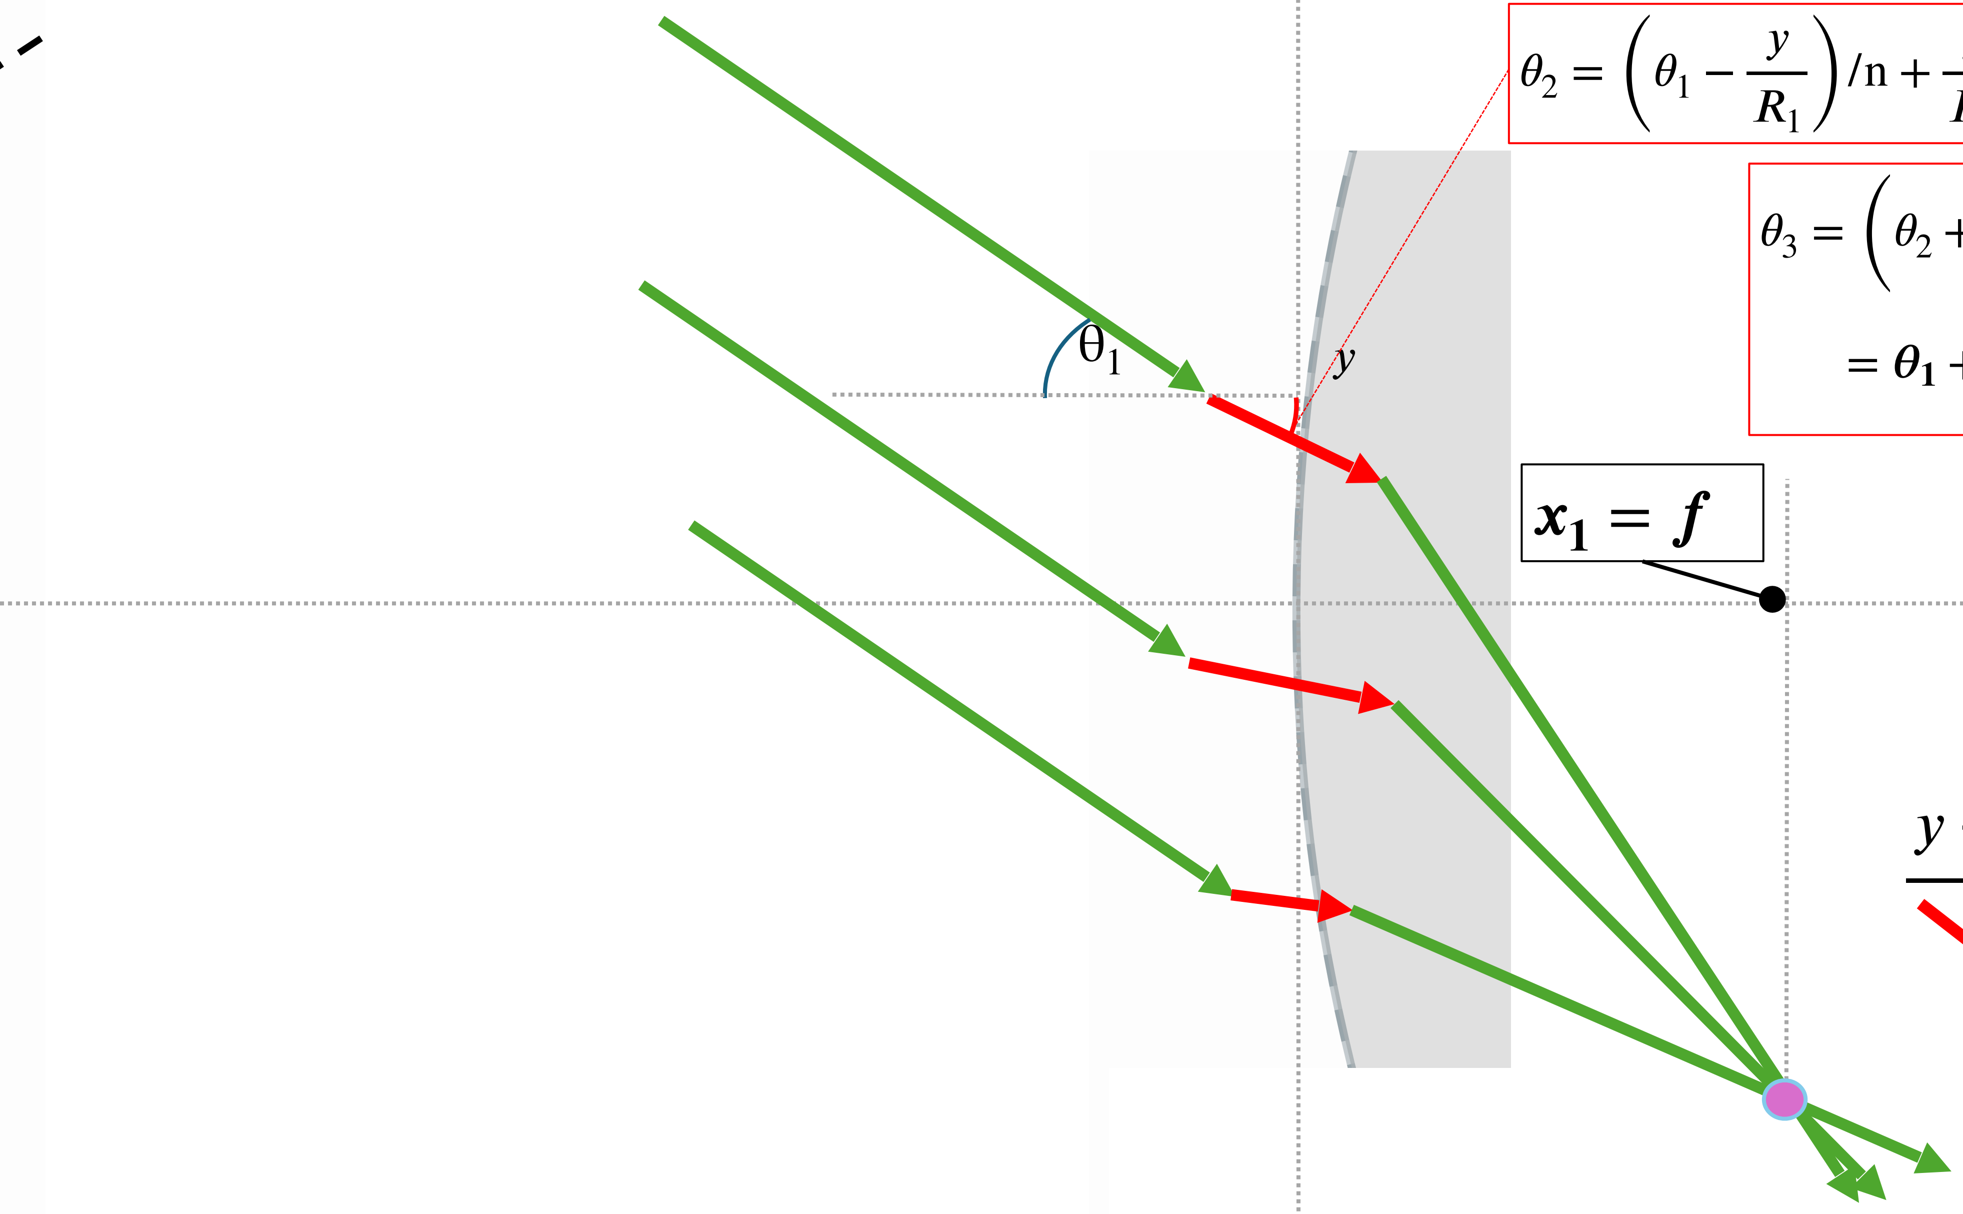

$$\theta_2 = \left(\theta_1 - \frac{y}{R_1}\right)/\mathrm{n} + \frac{y}{R_1}$$

$$\theta_3 = \left(\theta_2 + \frac{y}{R_2}\right)/\mathrm{n} + \frac{y}{R_2}$$

$$= \boldsymbol{\theta_1} + \underbrace{\boldsymbol{y}(\mathbf{1} - \boldsymbol{n})\left(\frac{\mathbf{1}}{\boldsymbol{R_1}} + \frac{\mathbf{1}}{\boldsymbol{R_2}}\right)}_{\frac{\mathbf{1}}{\boldsymbol{f}}}$$

$$\frac{y + y_1}{x_1 \; f} = \theta_3 = \theta_1 + y/f$$

$$y_1 + y = f\theta_1 + y$$

$$\boldsymbol{y_1 = f\theta_1}$$

Demo:

1_lens_demo.html

Convolution (starting with discrete 🤯)

Convolution (starting with discrete 🤯)

What if we want the same function to be applied
at all times/places?

Convolution (starting with discrete 🤯)

What if we want the same function to be applied at all times/places?

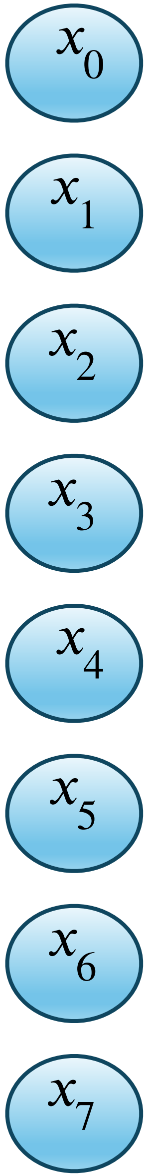

Convolution (starting with discrete 🤯)

What if we want the same function to be applied at all times/places?

$$\boldsymbol{Y}_t = \boldsymbol{T}(\boldsymbol{x})_t$$

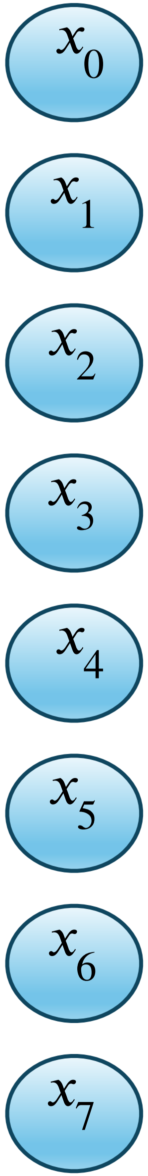

Convolution (starting with discrete 🤯)

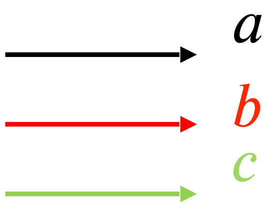

What if we want the same function to be applied at all times/places?

$$\boldsymbol{Y}_t = \boldsymbol{T}(\boldsymbol{x})_t$$

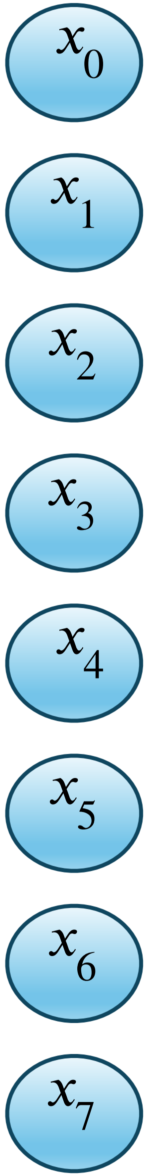

Convolution (starting with discrete 🤯)

What if we want the same function to be applied at all times/places?

$$Y_t = T(x)_t$$

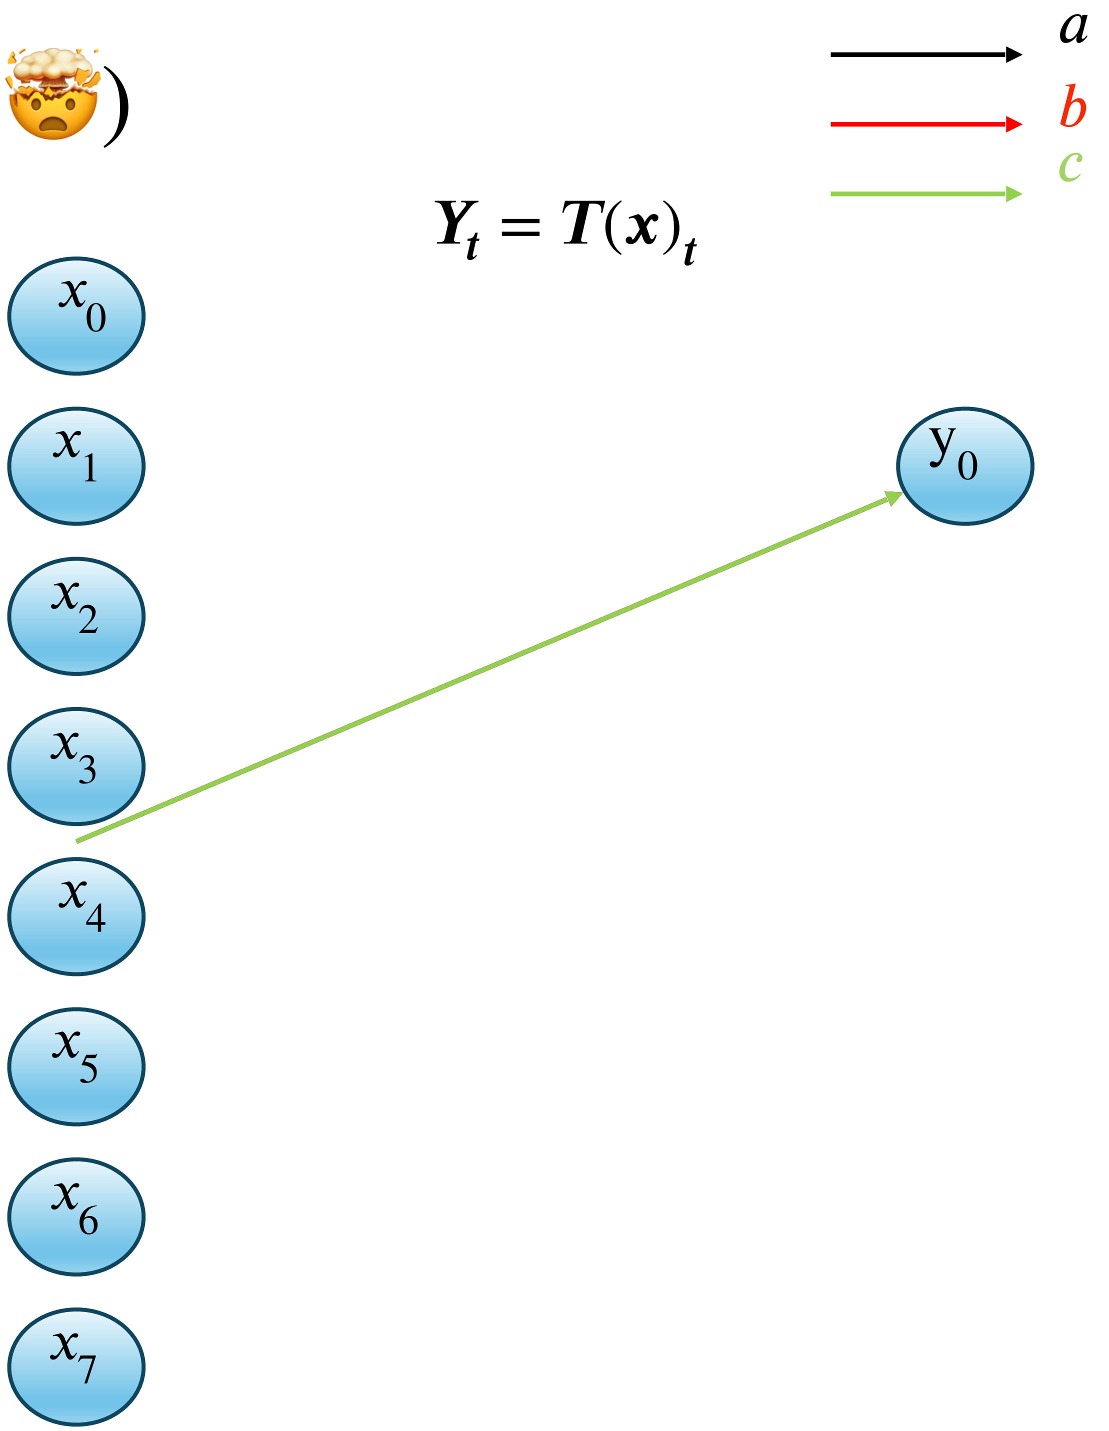

Convolution (starting with discrete 🤯)

What if we want the same function to be applied at all times/places?

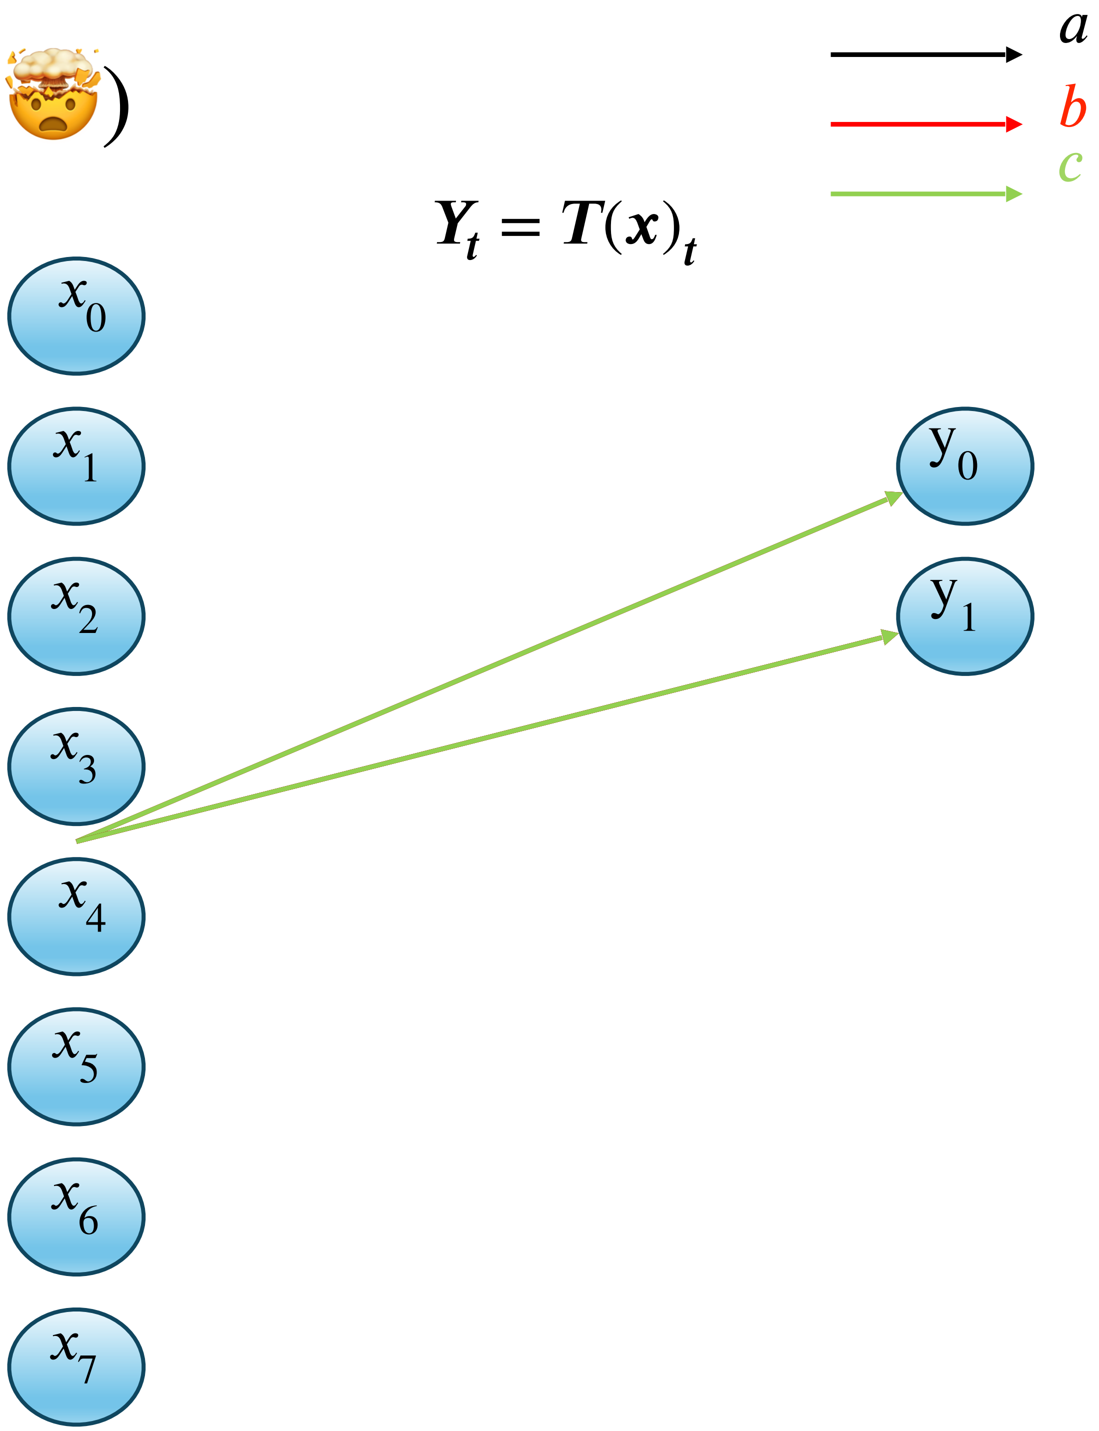

Convolution (starting with discrete 🤯)

What if we want the same function to be applied at all times/places?

$$T(x_{t-\tau})_t = T(x_t)_{t-\tau}$$

$$\boldsymbol{Y}_t = \boldsymbol{T}(\boldsymbol{x})_t$$

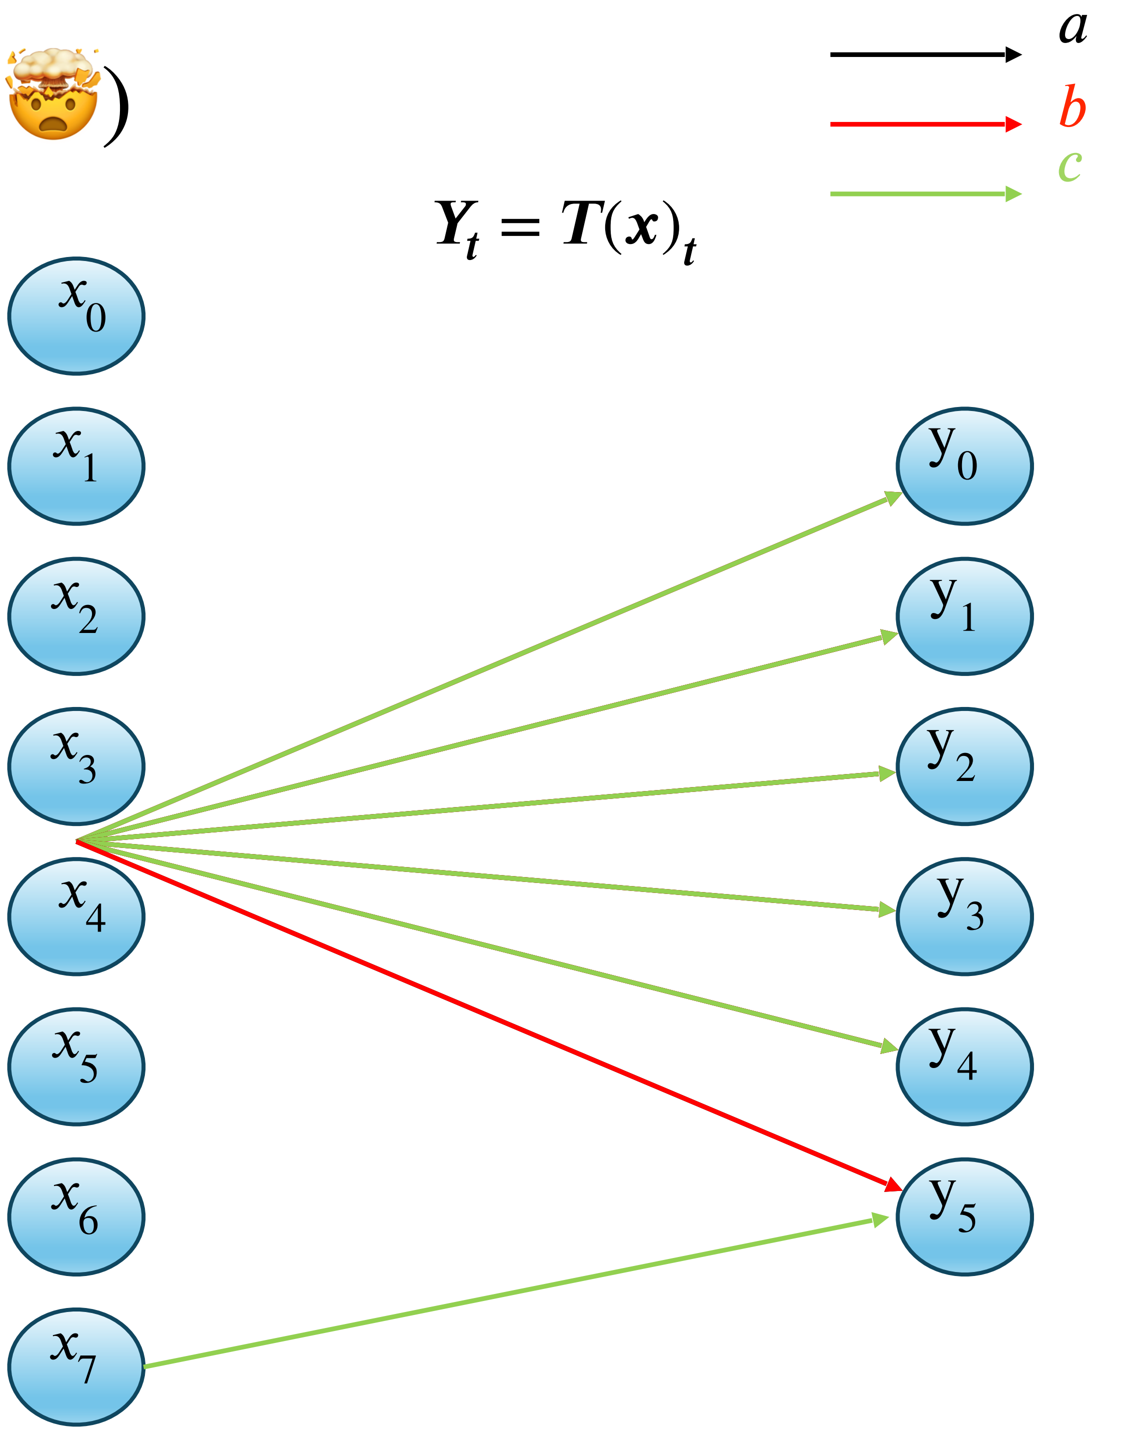

Convolution (starting with discrete 🤯)

What if we want the same function to be applied at all times/places?

$$T(x_{t-\tau})_t = T(x_t)_{t-\tau}$$

Convolution:

$$\{f * g\}[t] = \sum_{t=-\infty}^{\infty} f[t]g[t-\tau]$$

$$\boldsymbol{Y}_t = \boldsymbol{T}(\boldsymbol{x})_t$$

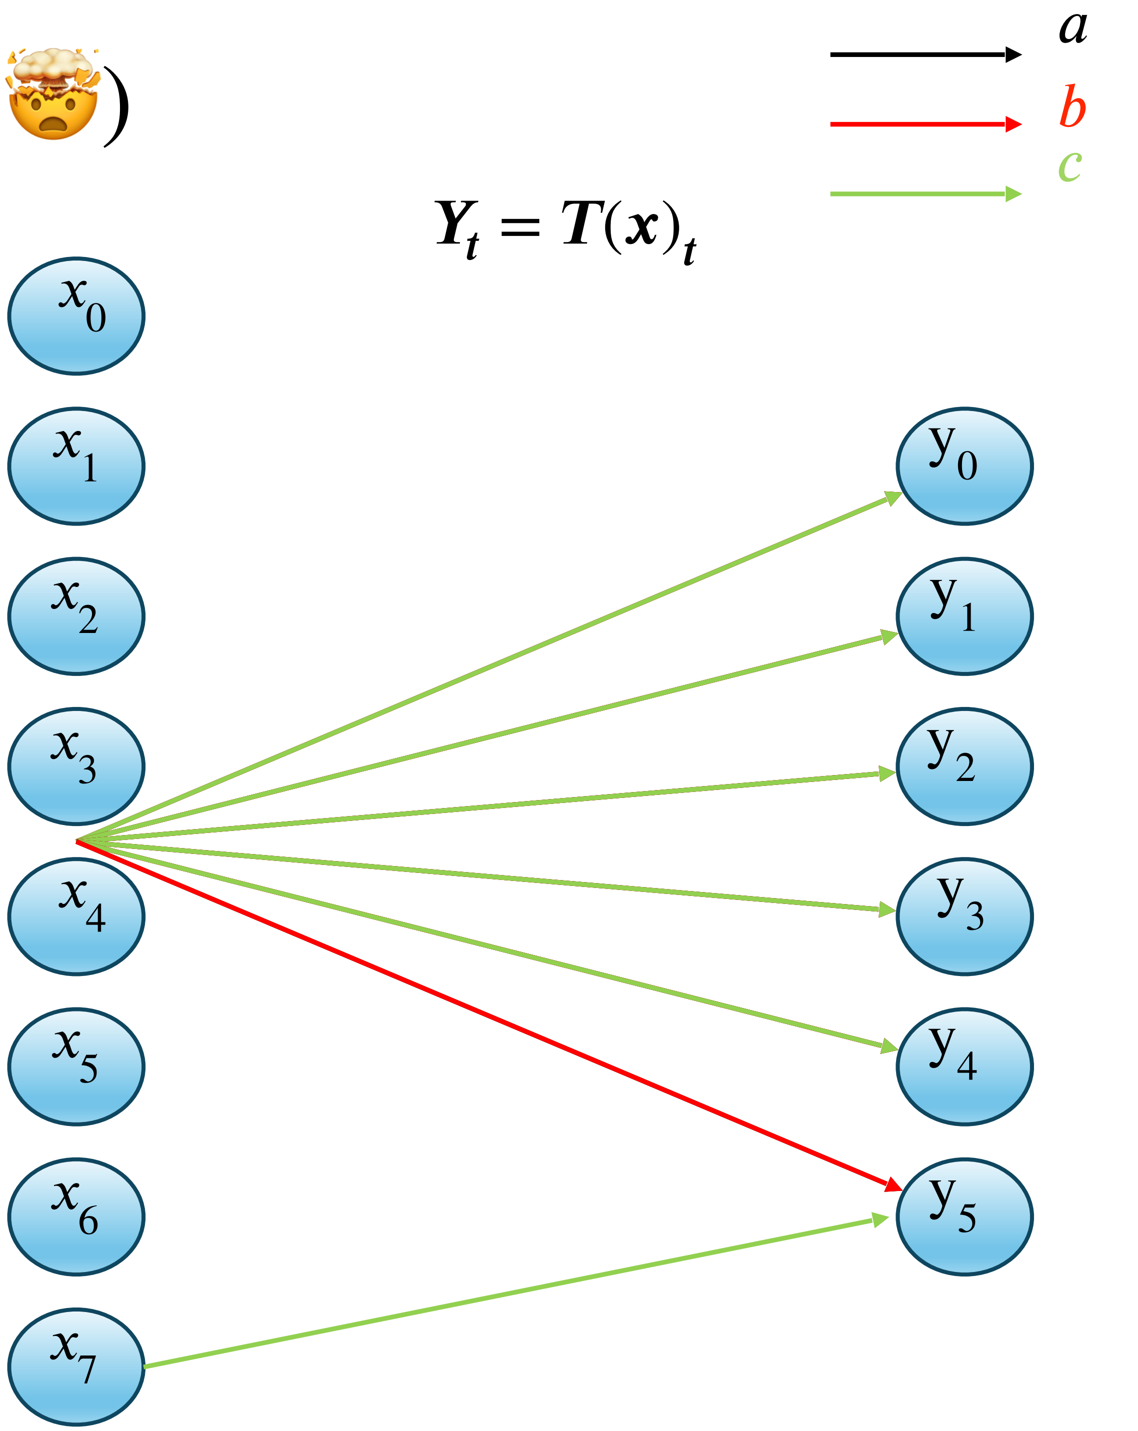

Convolution (starting with discrete 🤯)

What if we want the same function to be applied at all times/places?

$$T(x_{t-\tau})_t = T(x_t)_{t-\tau}$$

<u>Convolution</u>:

$$\{f * g\}[t] = \sum_{t=-\infty}^{\infty} f[t]g[t-\tau]$$

$$g = [c,\ b,\ a]$$

$$\boldsymbol{Y}_t = \boldsymbol{T}(\boldsymbol{x})_t$$

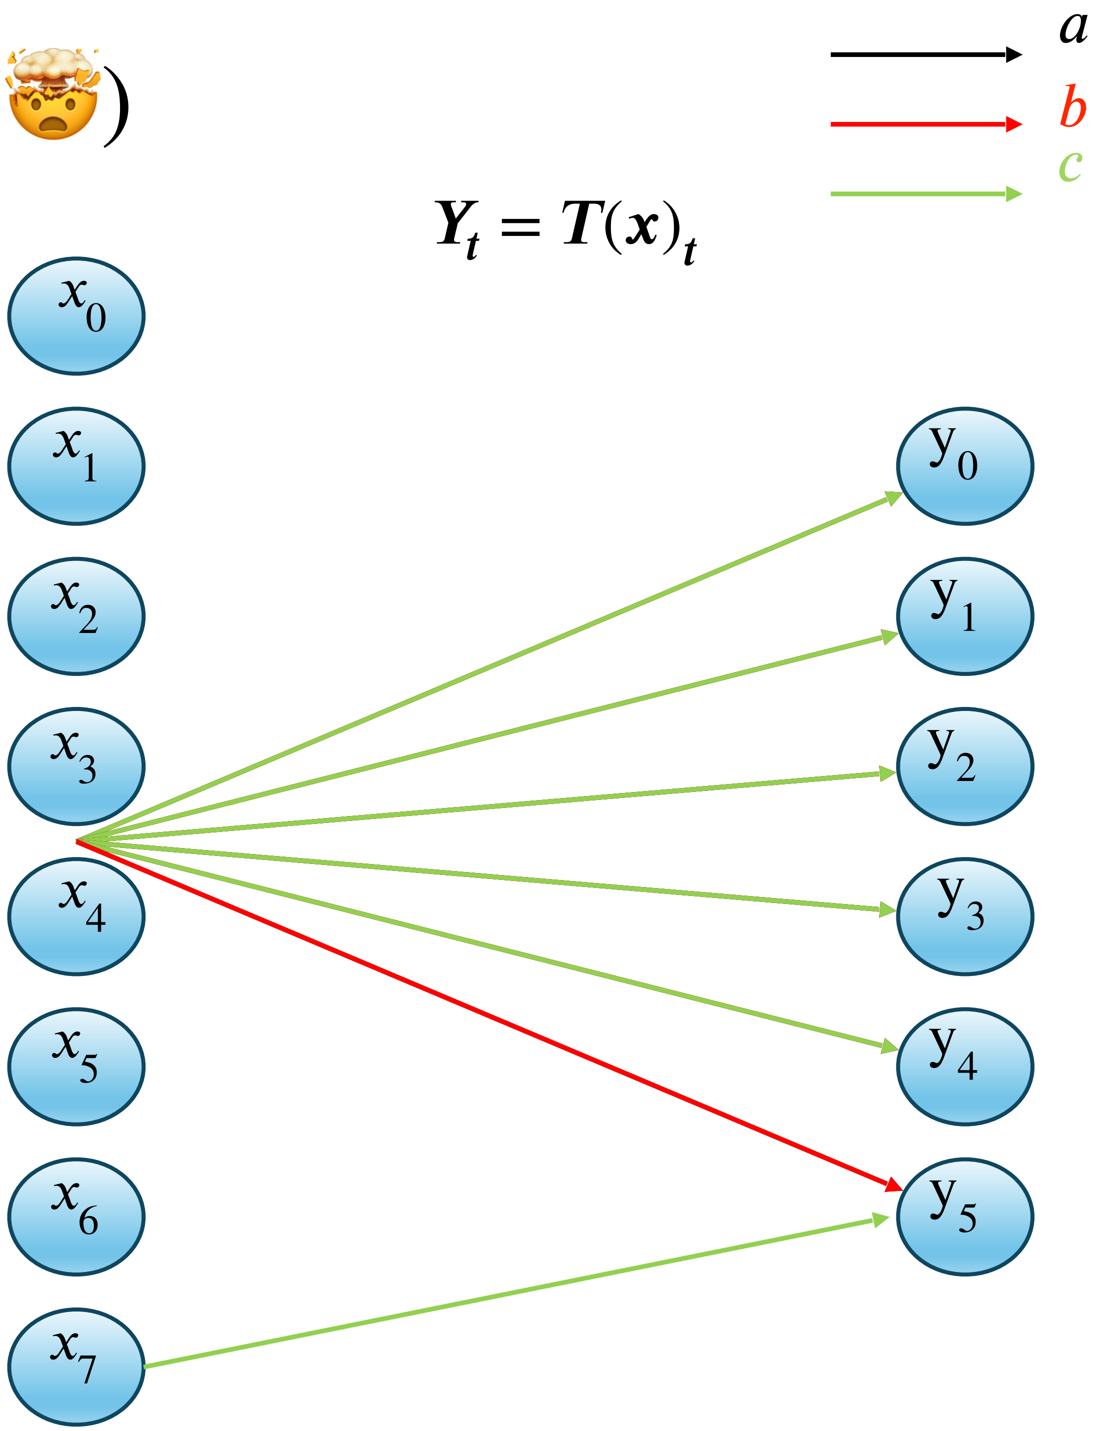

Convolution Filter

$g = [c,\ b,\ a]$

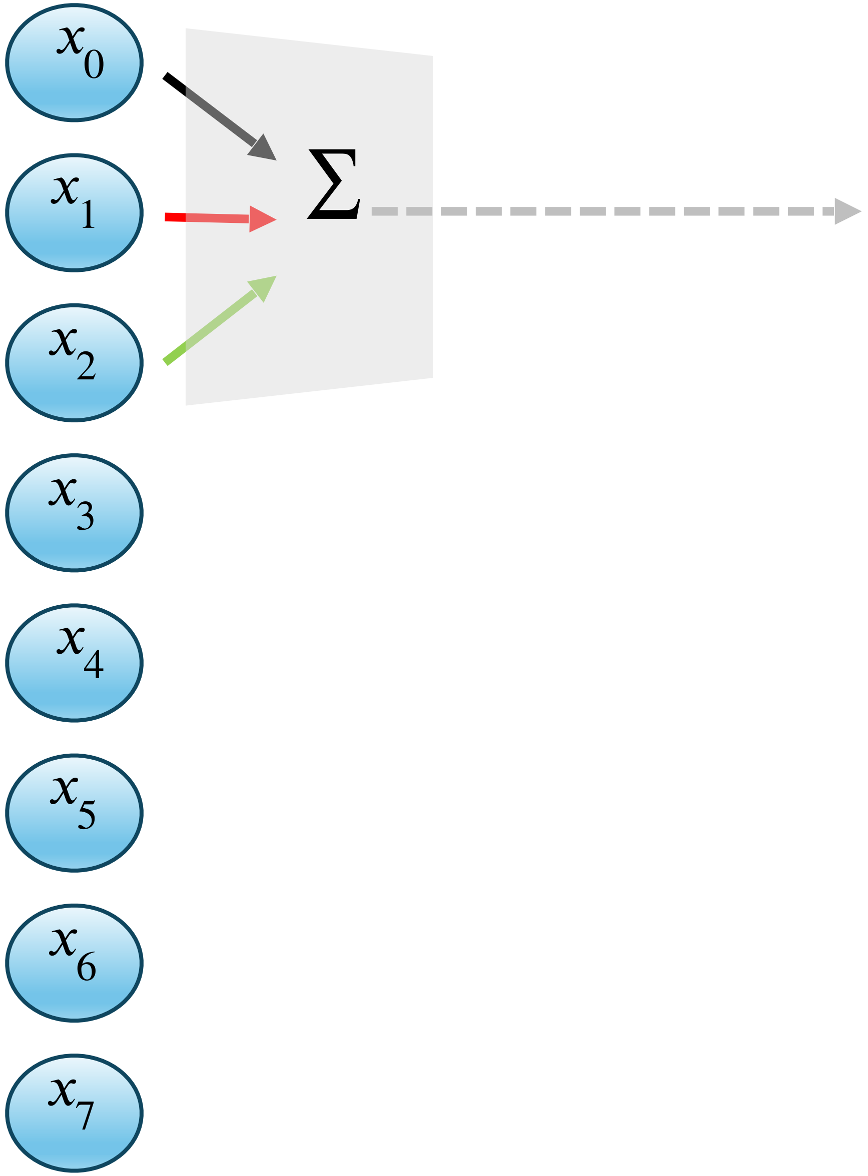

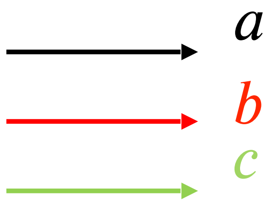

Convolution Filter

$g = [c,\ b,\ a]$

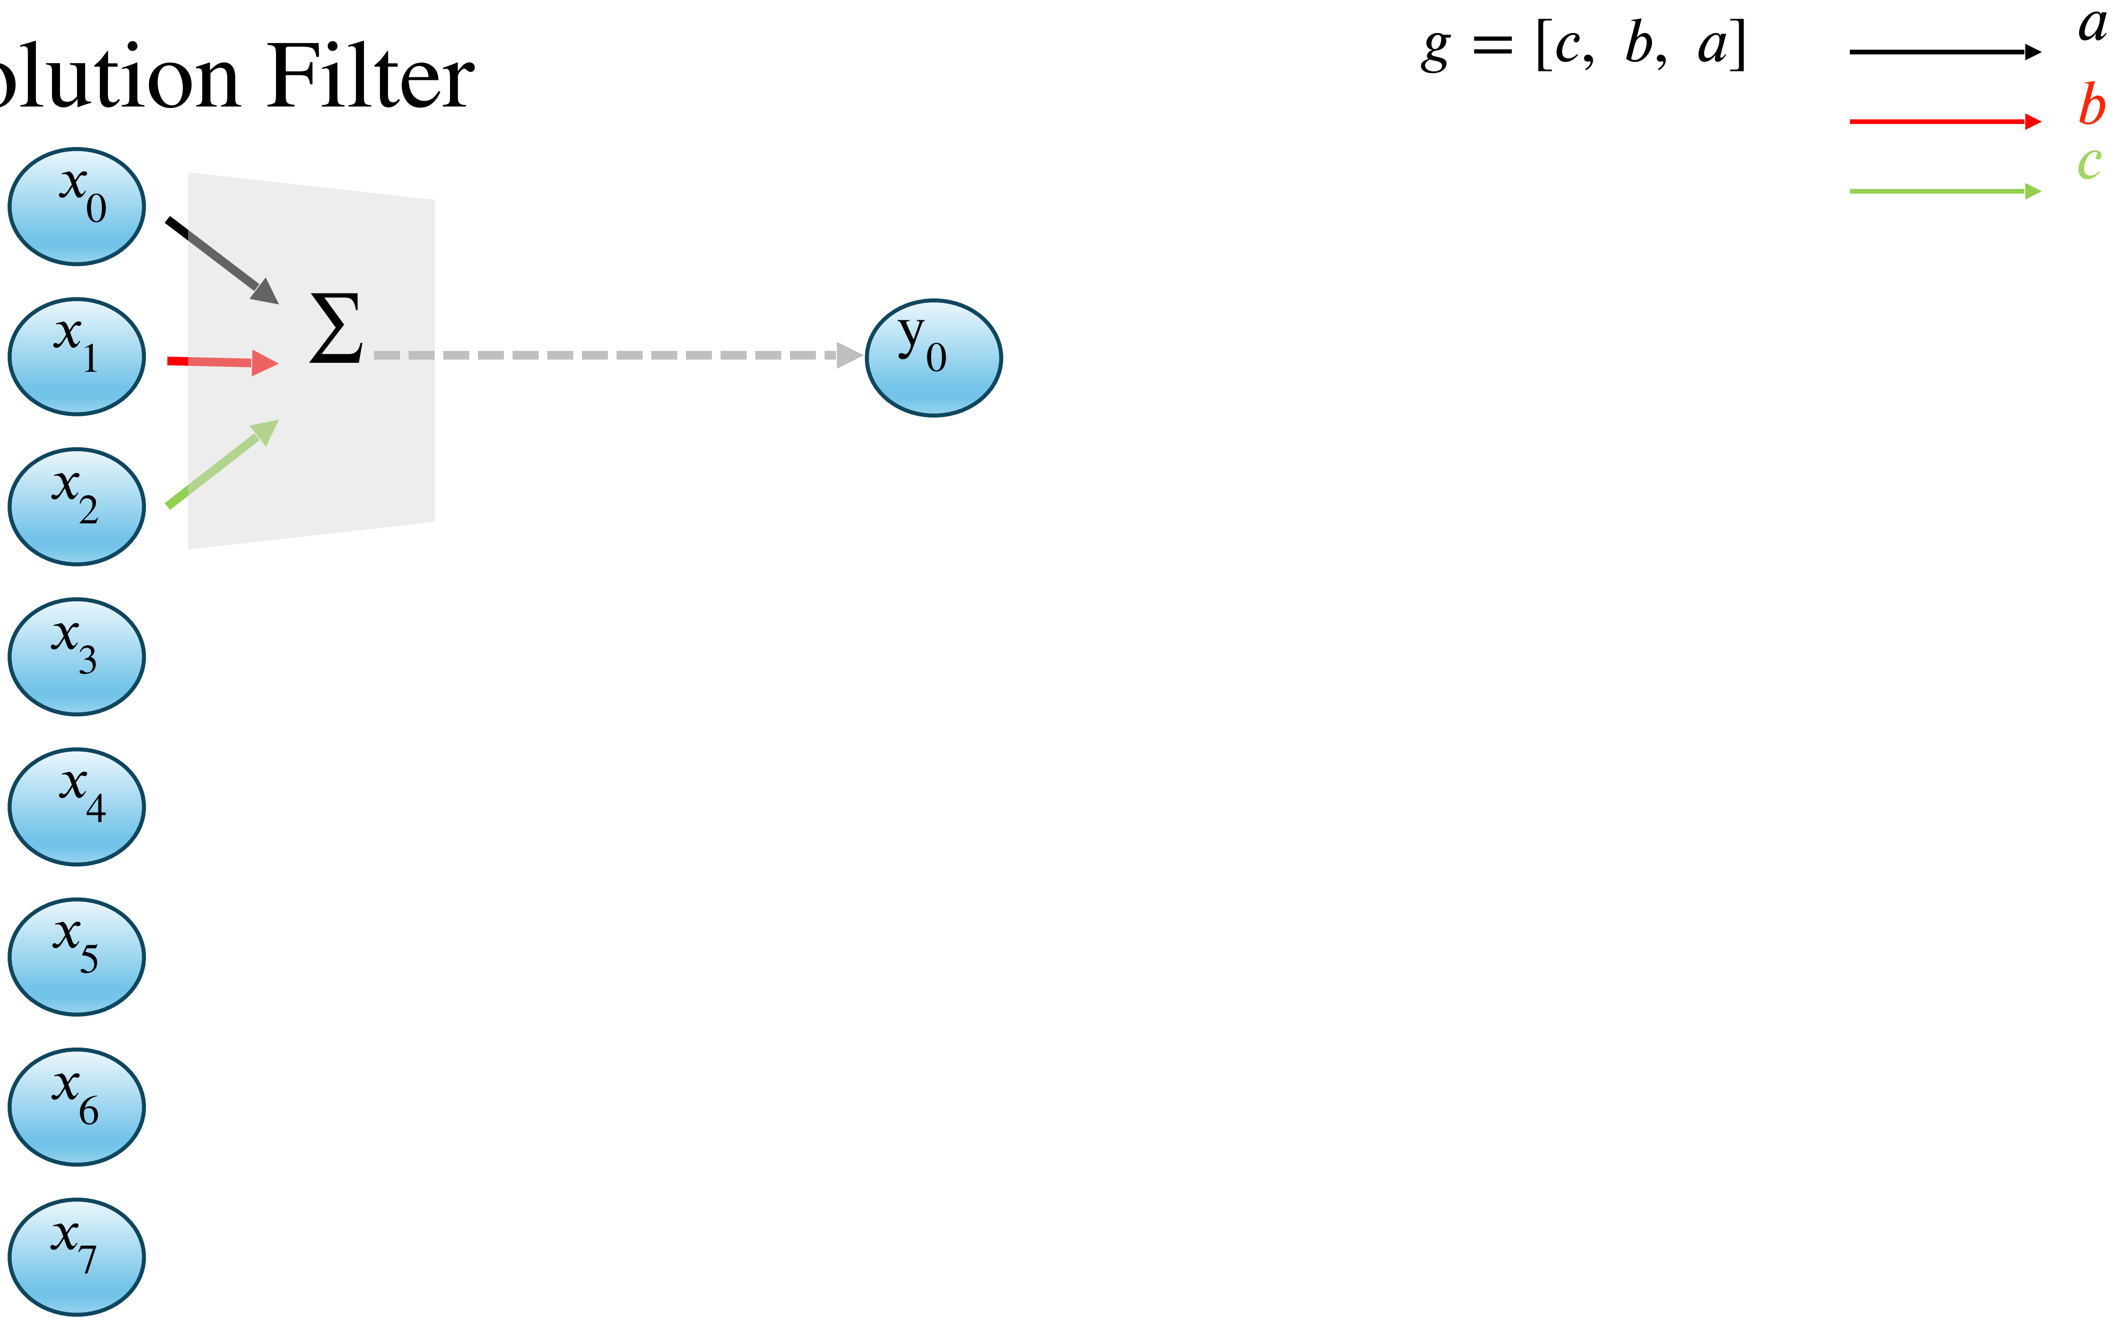

Convolution Filter

$$g = [c,\ b,\ a]$$

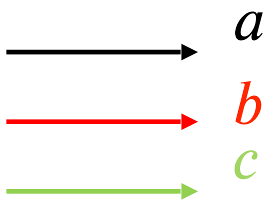

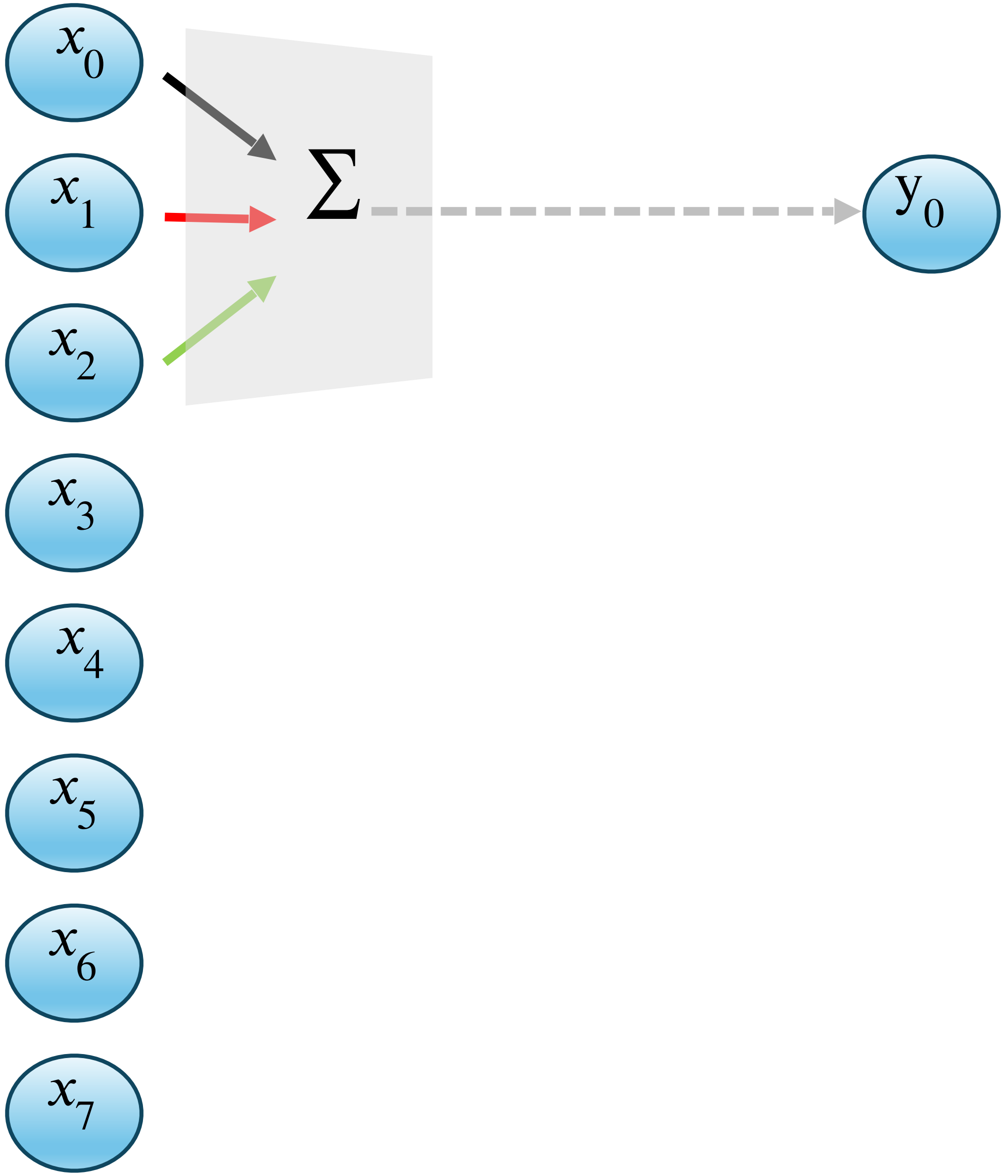

Standard inner product

$$y_0 = [x_0 \quad x_1 \quad x_2] \begin{bmatrix} c \\ b \\ a \end{bmatrix} = \langle x_{0:3}\ ,\ g \rangle$$

Convolution Filter

$$g = [c,\ b,\ a]$$

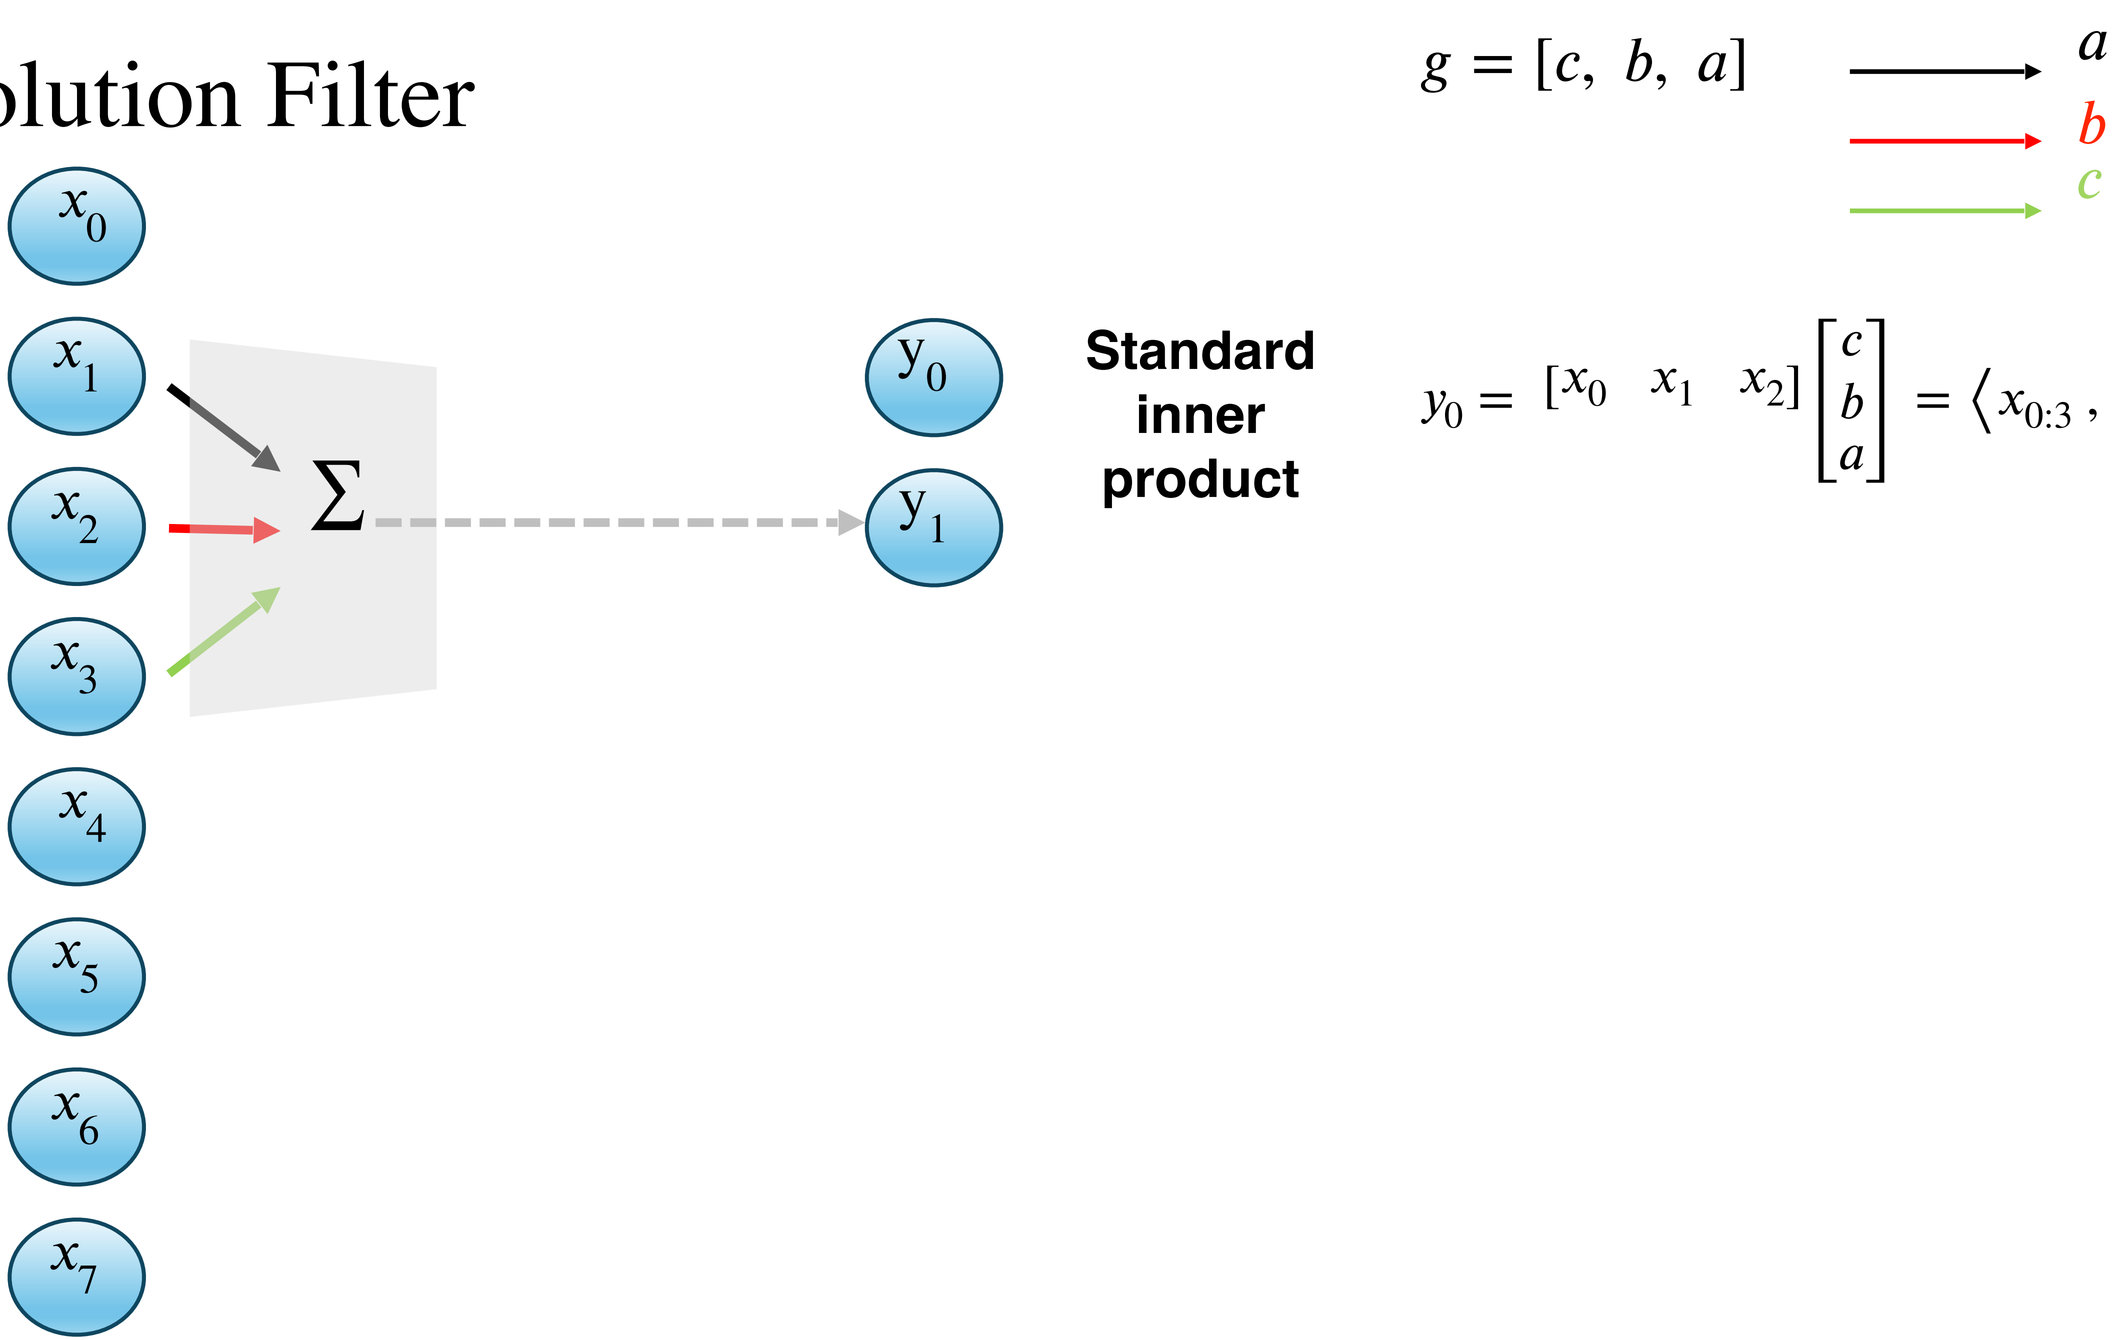

Standard inner product

$$y_0 = [x_0 \quad x_1 \quad x_2]\begin{bmatrix} c \\ b \\ a \end{bmatrix} = \langle x_{0:3}\,,\ g \rangle$$

Convolution Filter

$g = [c,\ b,\ a]$

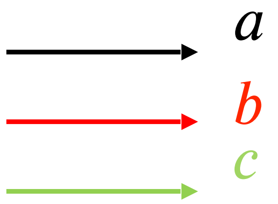

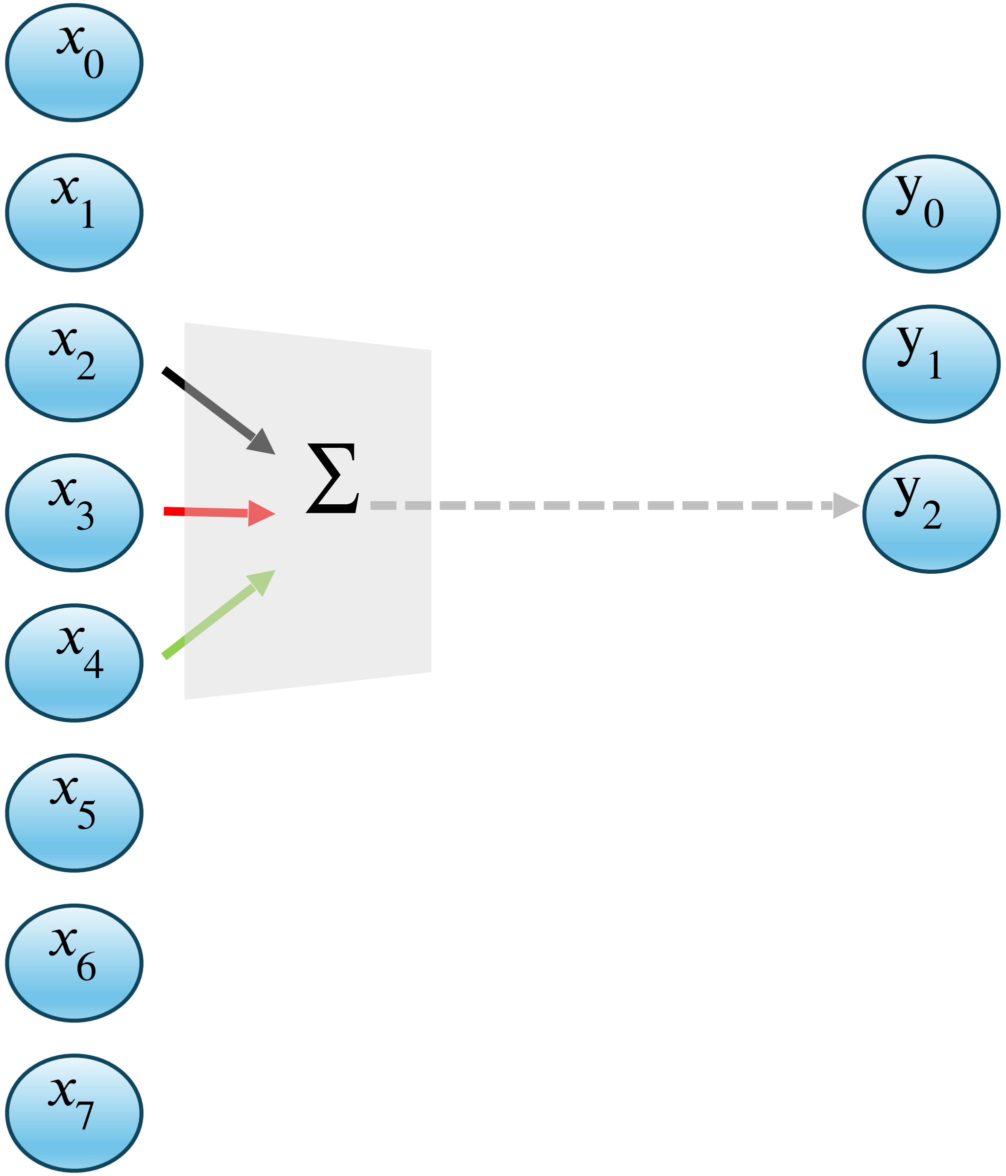

Standard inner product

$$y_0 = [x_0 \quad x_1 \quad x_2]\begin{bmatrix} c \\ b \\ a \end{bmatrix} = \langle x_{0:3}\,,\ g \rangle$$

Convolution Filter

$g = [c,\ b,\ a]$

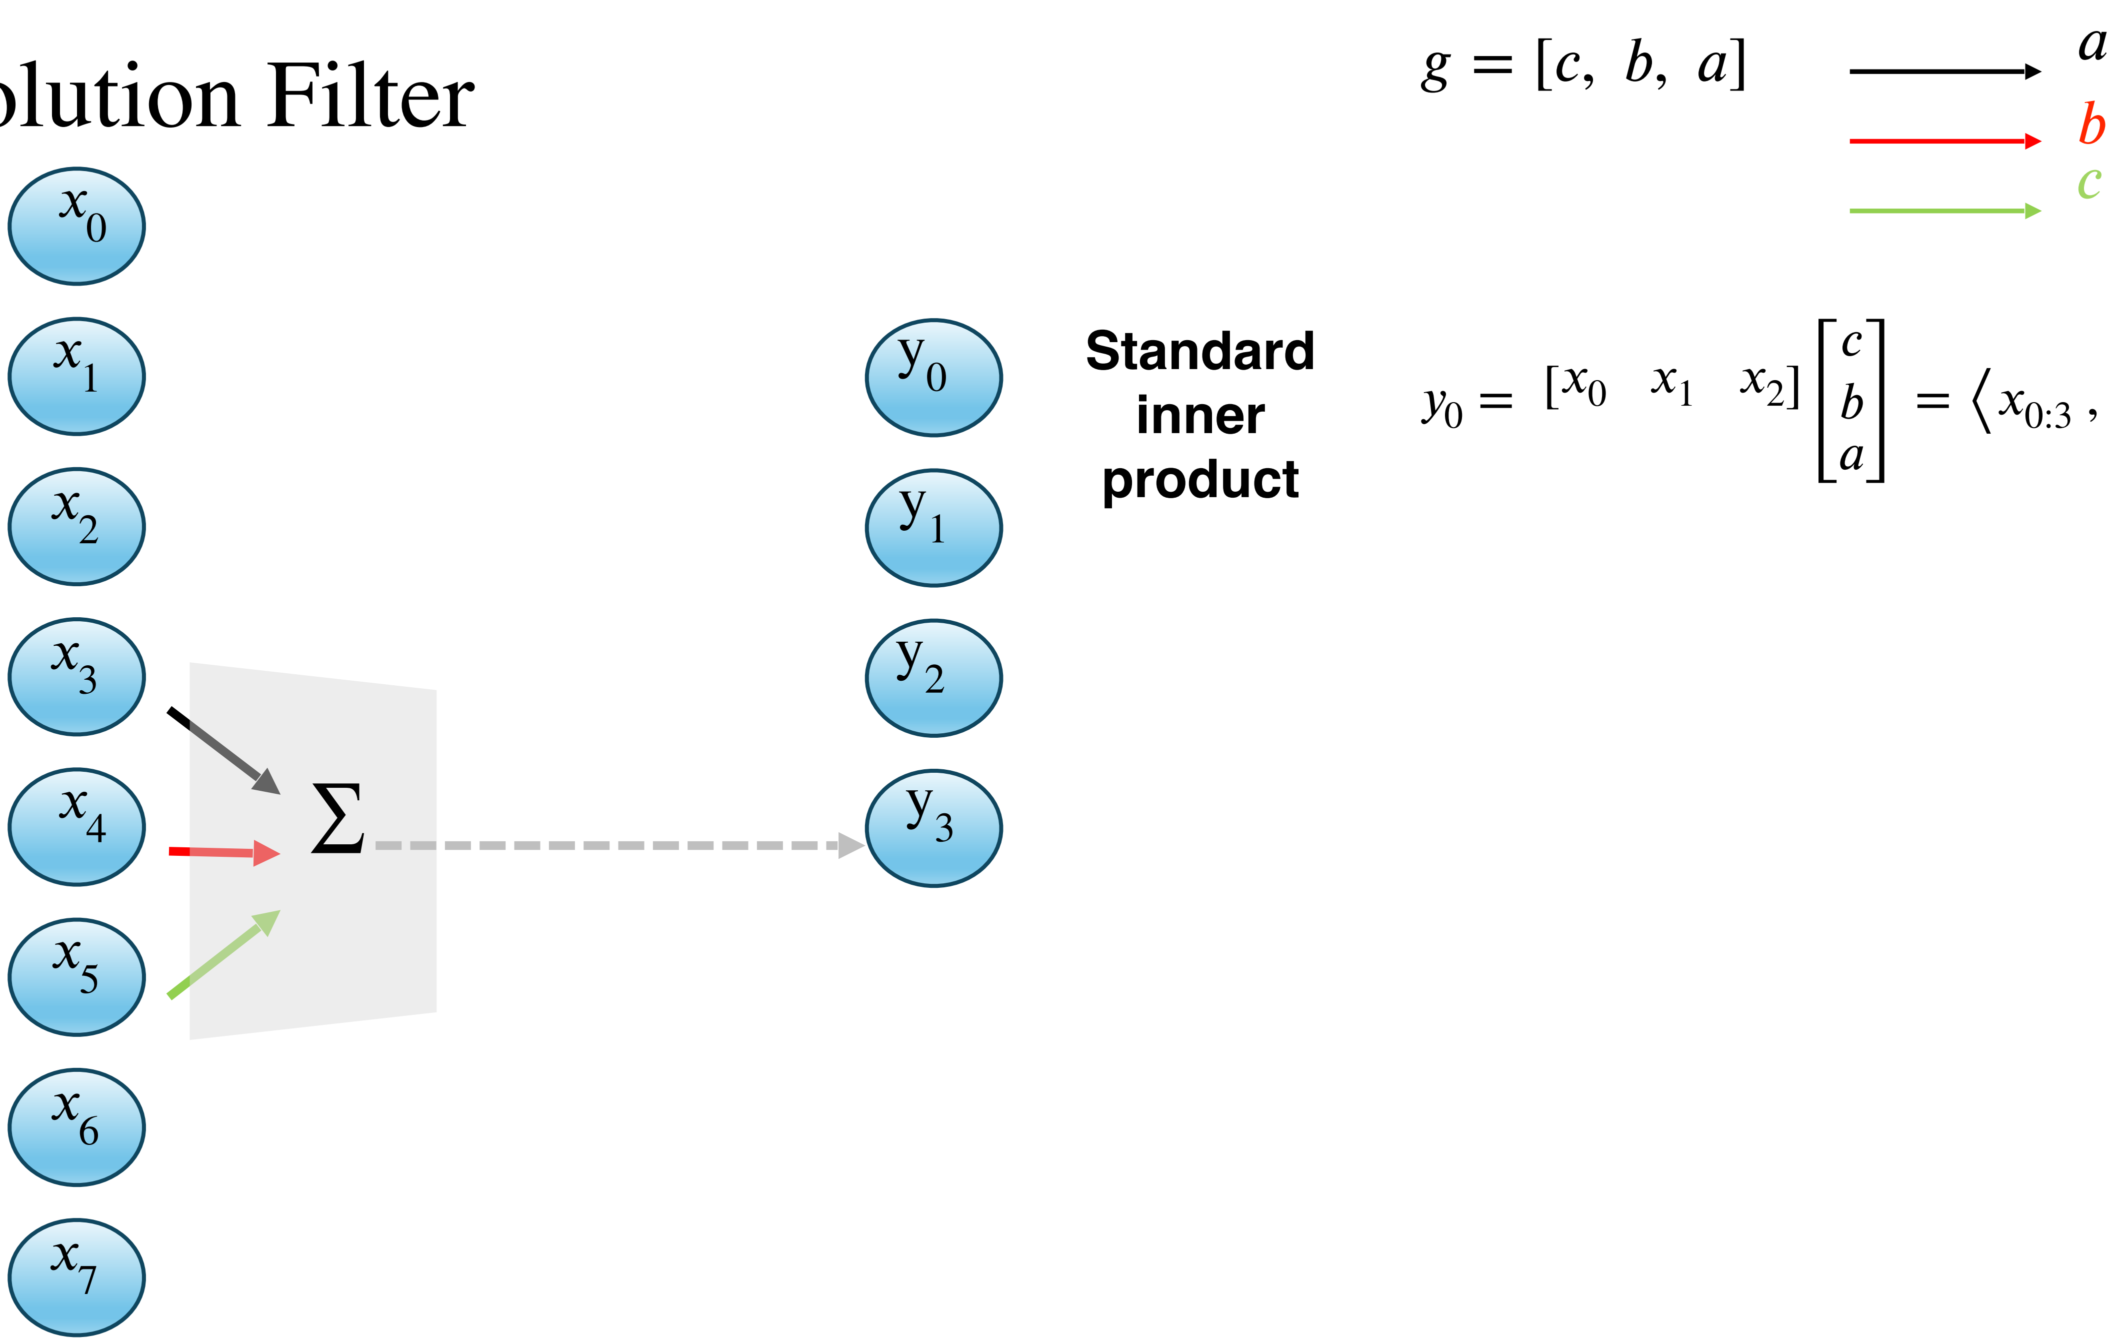

Standard inner product

$$y_0 = [x_0 \quad x_1 \quad x_2] \begin{bmatrix} c \\ b \\ a \end{bmatrix} = \langle x_{0:3}\,,\ g \rangle$$

Convolution Filter

$g = [c,\ b,\ a]$

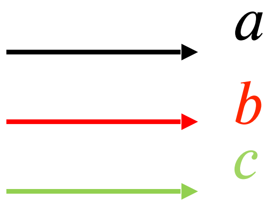

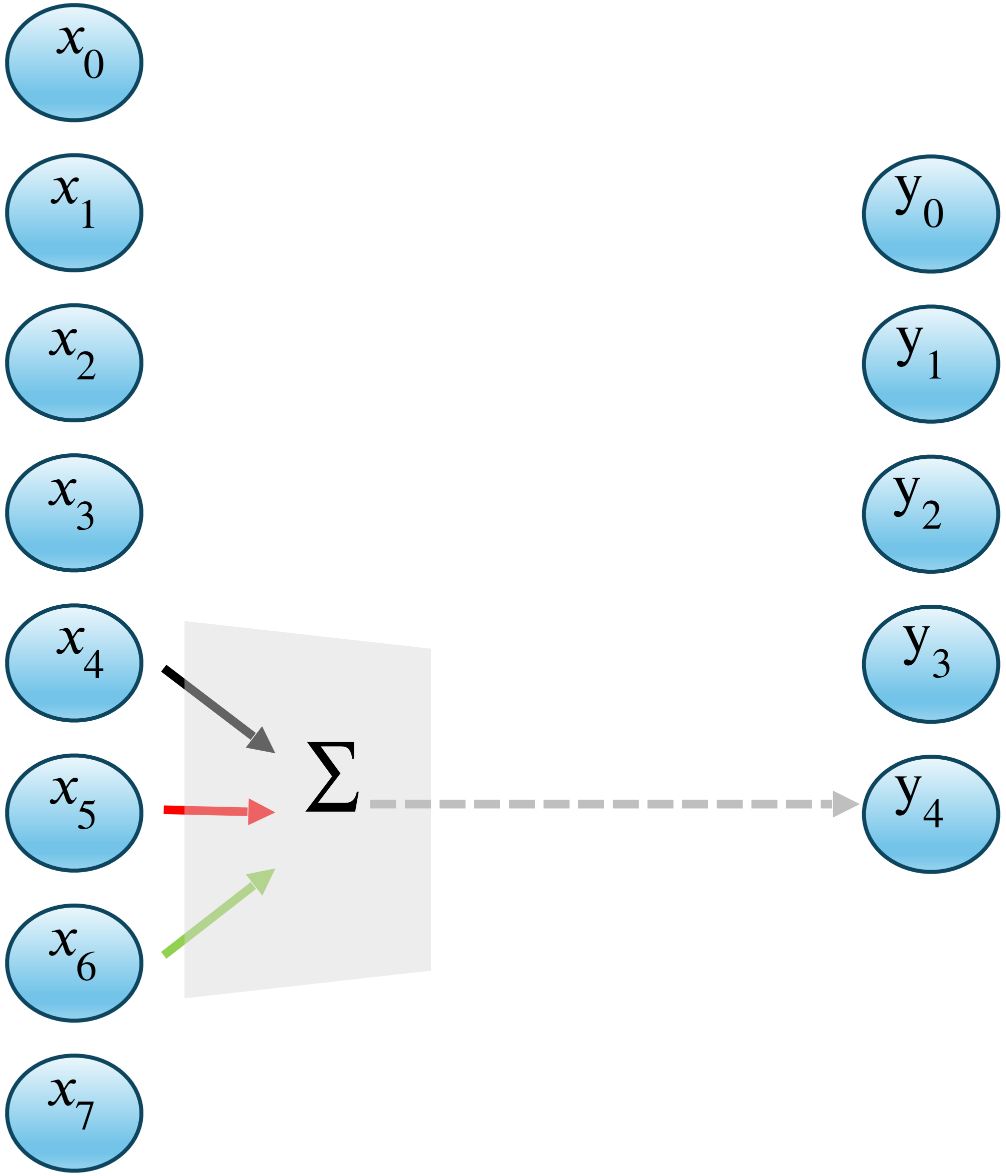

Standard inner product

$$y_0 = [x_0 \quad x_1 \quad x_2] \begin{bmatrix} c \\ b \\ a \end{bmatrix} = \langle x_{0:3}\,,\ g \rangle$$

Convolution Filter

$g = [c,\ b,\ a]$

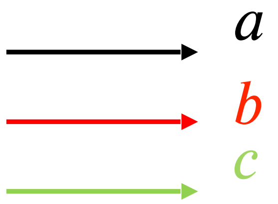

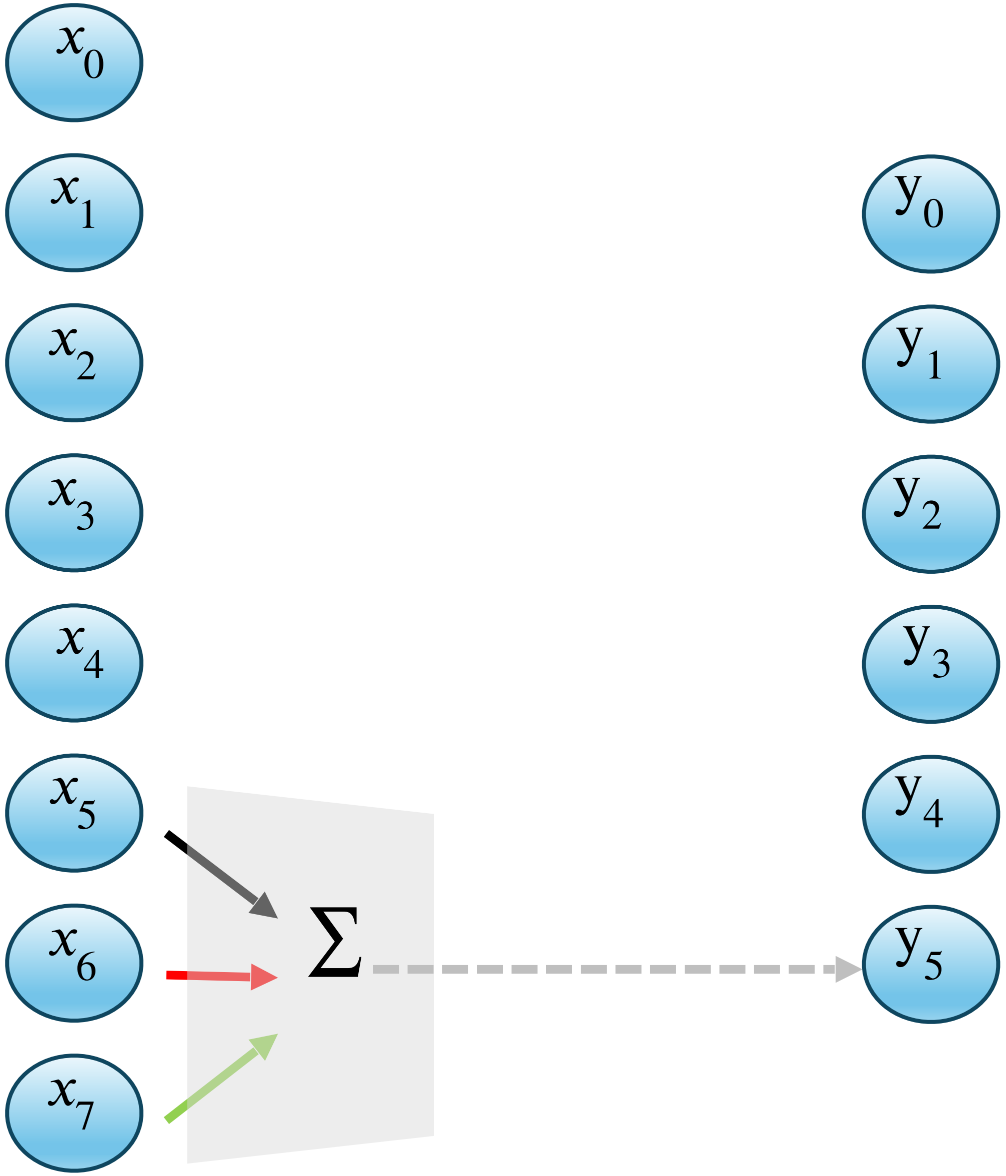

Standard inner product

$$y_0 = [x_0 \quad x_1 \quad x_2]\begin{bmatrix} c \\ b \\ a \end{bmatrix} = \langle x_{0:3}\,,\ g \rangle$$

Convolution Filter

$g = [c, \ b, \ a]$

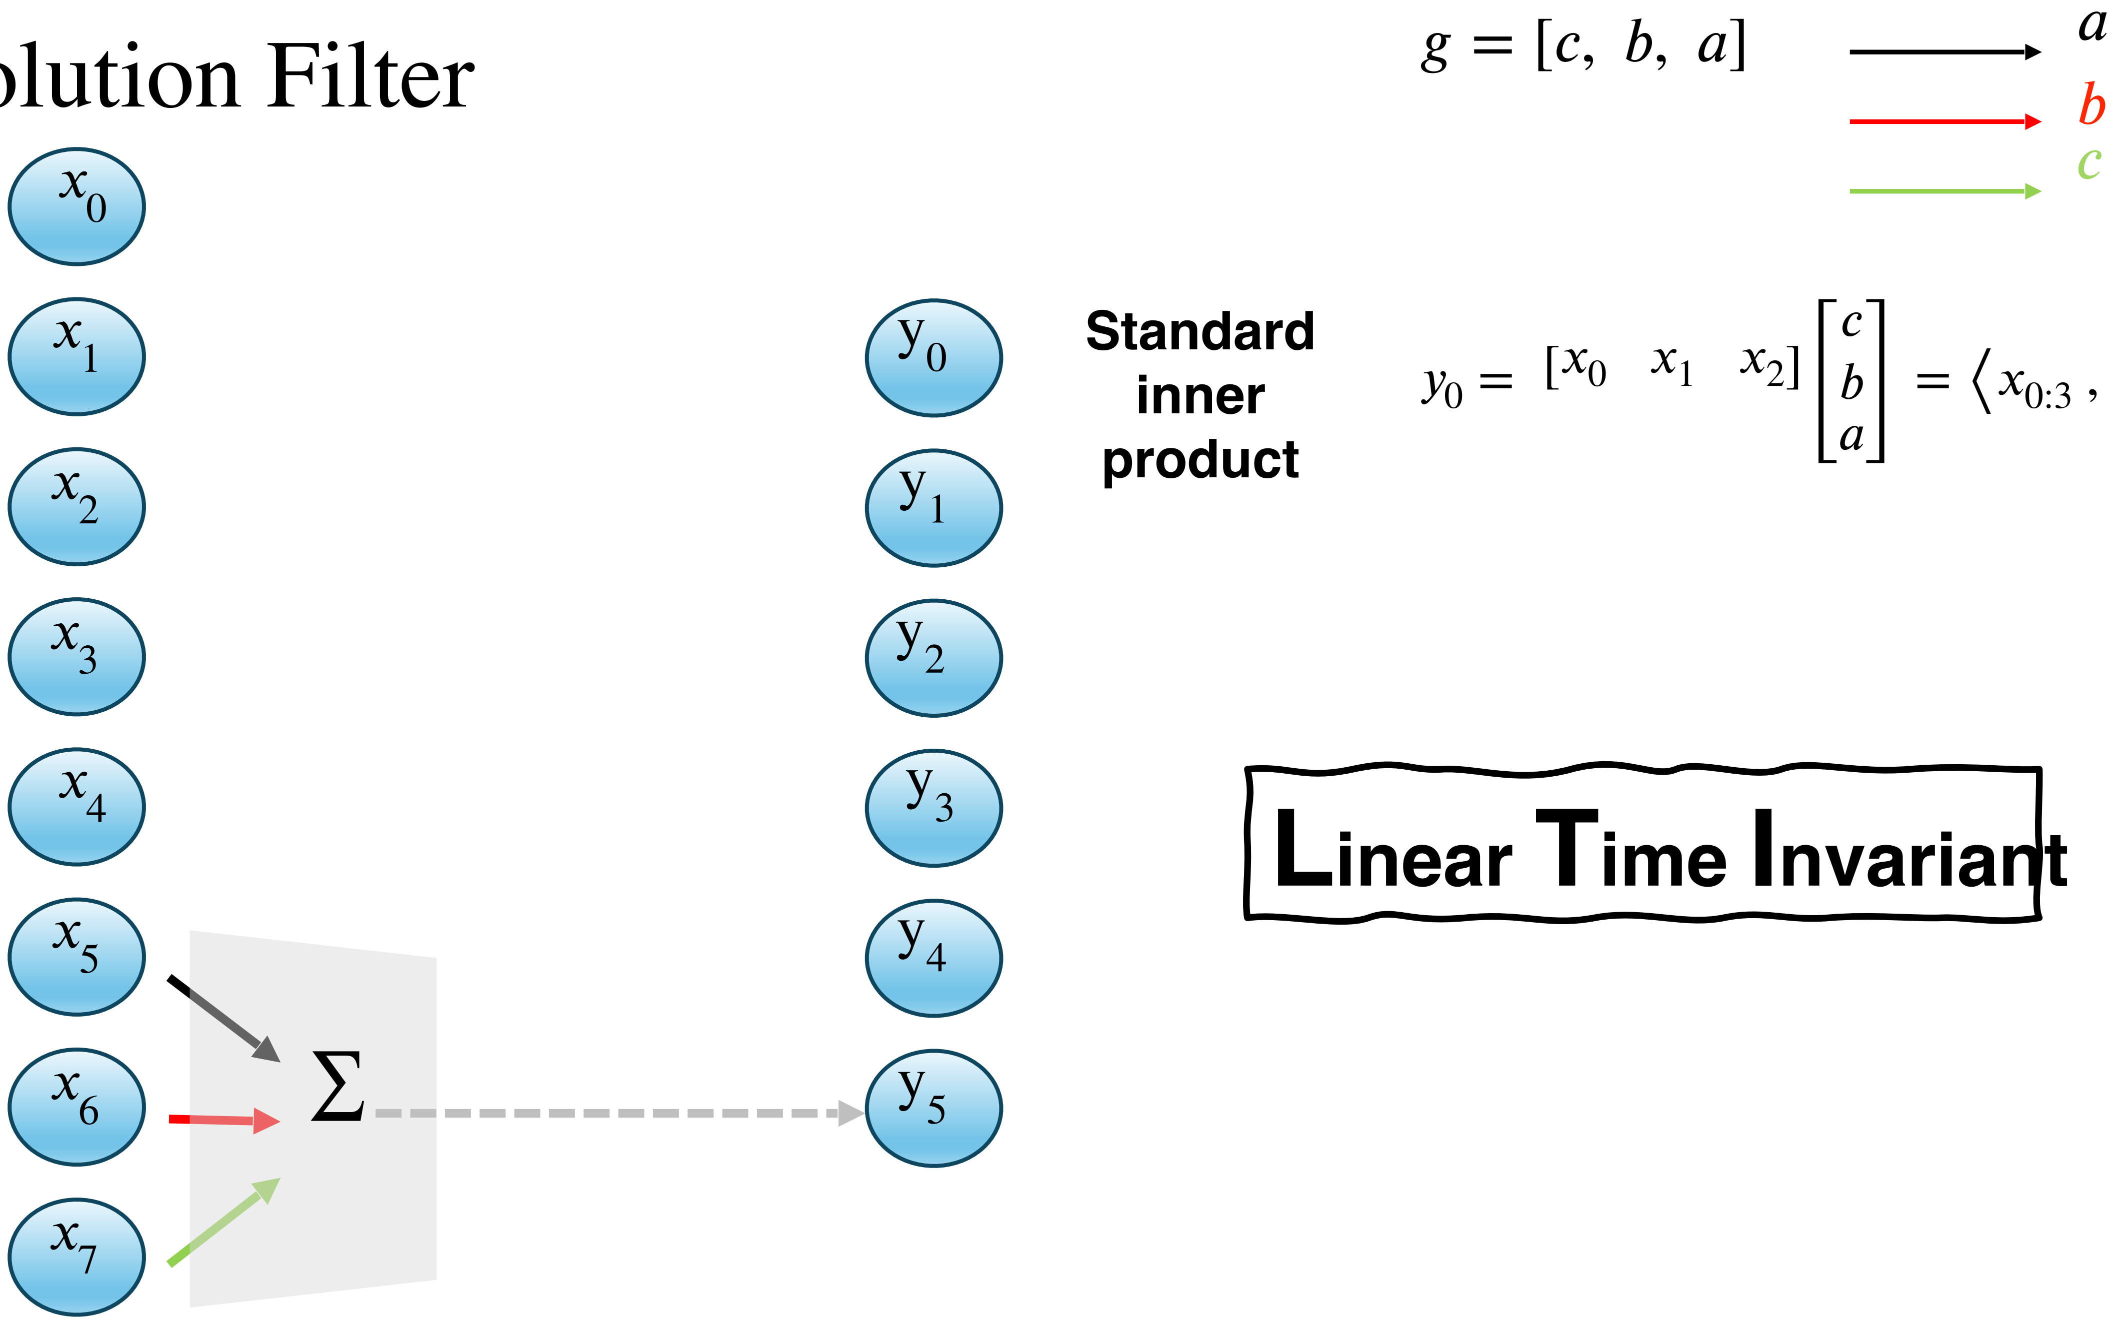

Standard inner product

$$y_0 = [x_0 \quad x_1 \quad x_2] \begin{bmatrix} c \\ b \\ a \end{bmatrix} = \langle x_{0:3} \, , \; g \rangle$$

Linear Time Invariant

Convolution Filter

$g = [c,\ b,\ a]$

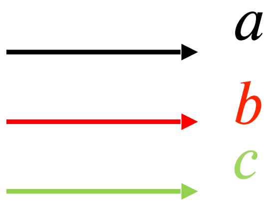

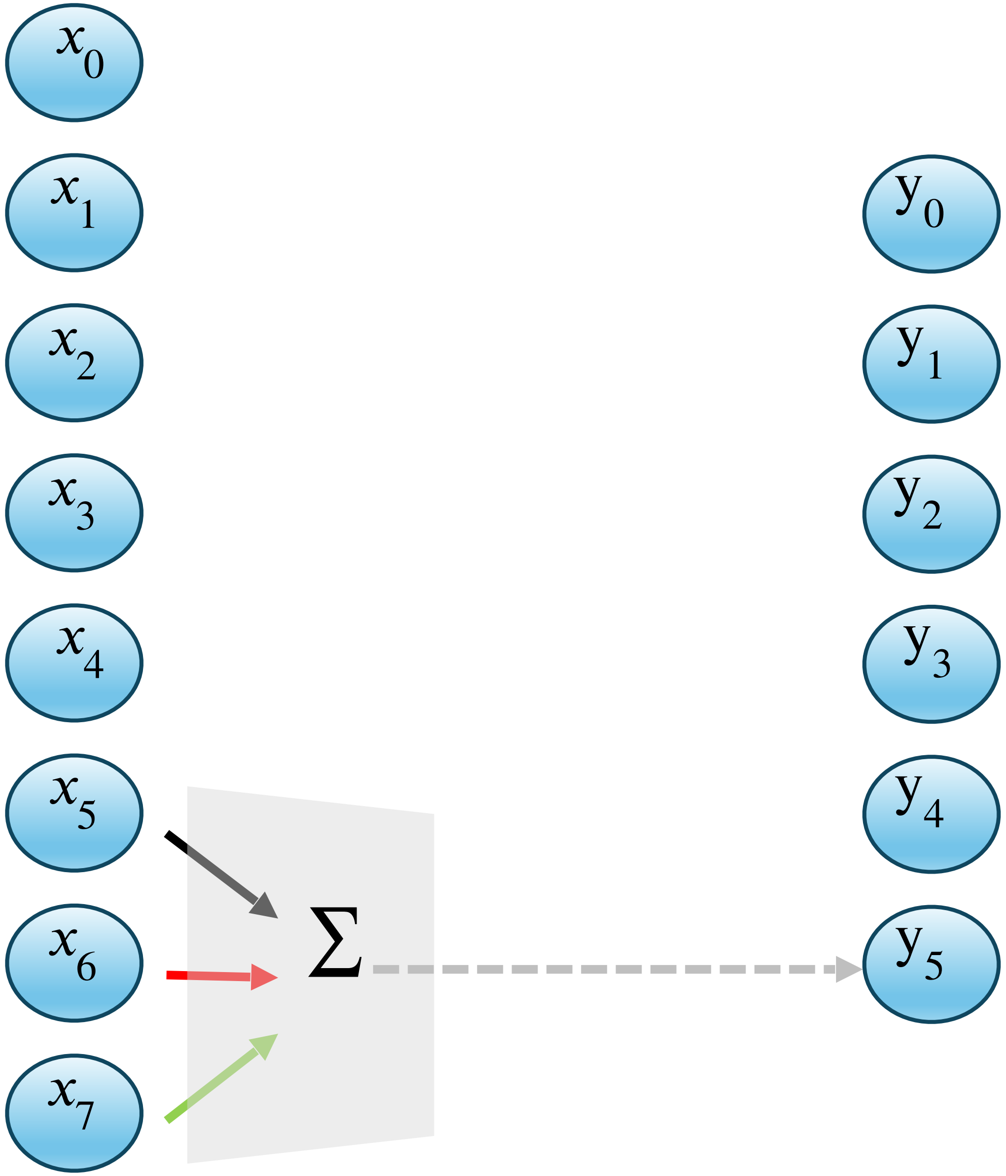

Standard inner product

$$y_0 = [x_0 \quad x_1 \quad x_2] \begin{bmatrix} c \\ b \\ a \end{bmatrix} = \langle x_{0:3}\, ,\ g \rangle$$

Linear Time Invariant

LTI is always a convolution!
Proof: HW

Is it even linear?

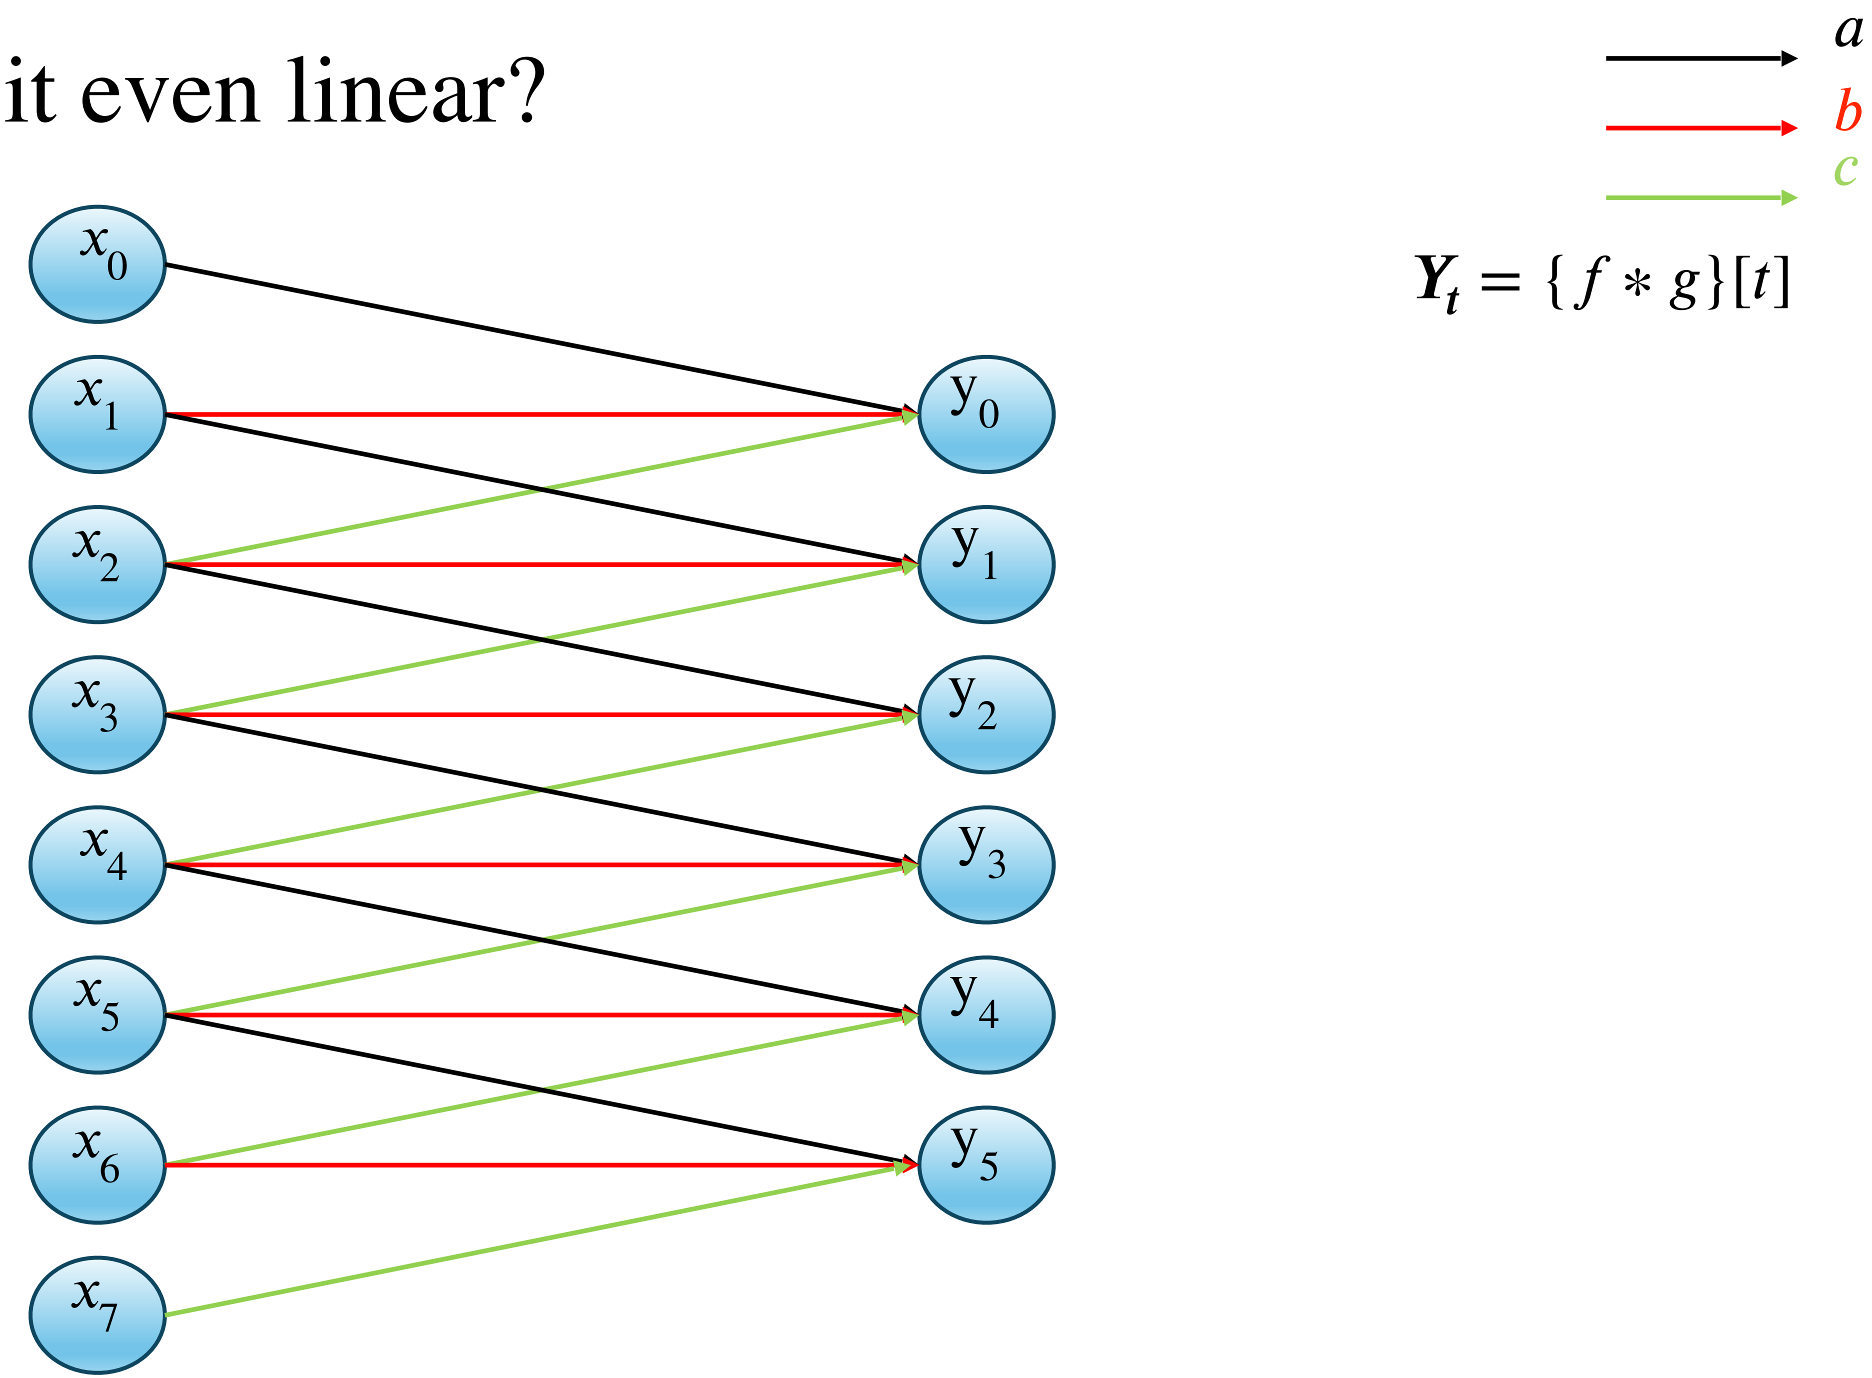

$$\boldsymbol{Y}_t = \{f * g\}[t] \qquad g = [c,\ b,\ a]$$

Is it even linear?

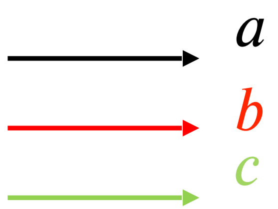

$$\boldsymbol{Y}_t = \{f * g\}[t] \qquad g = [c,\ b,\ a]$$

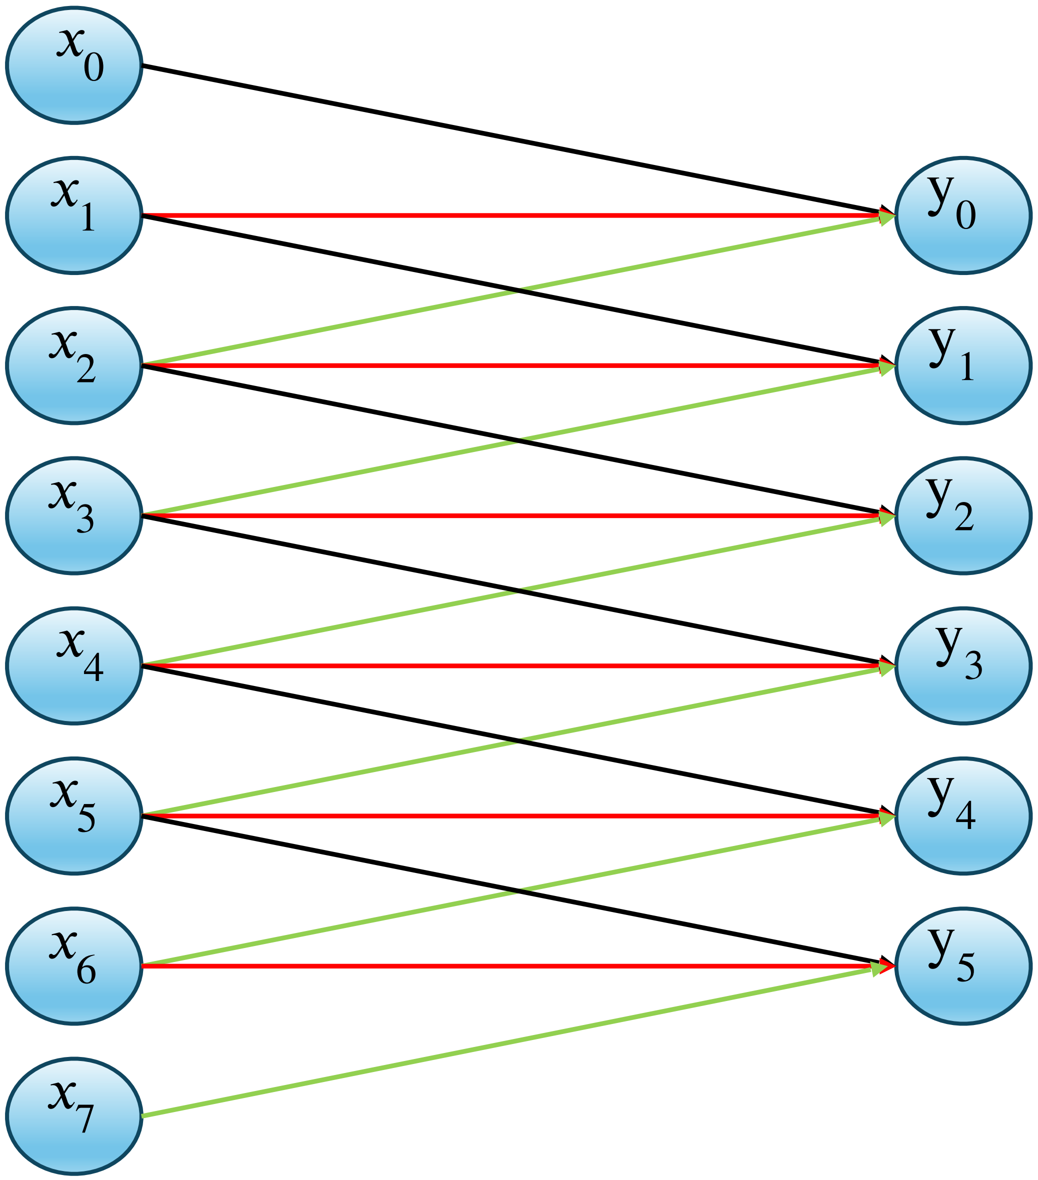

Is it even linear?

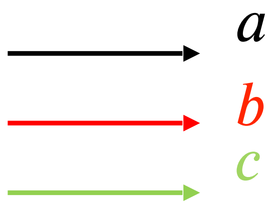

$$Y_t = \{f * g\}[t] \qquad g = [c,\ b,\ a]$$

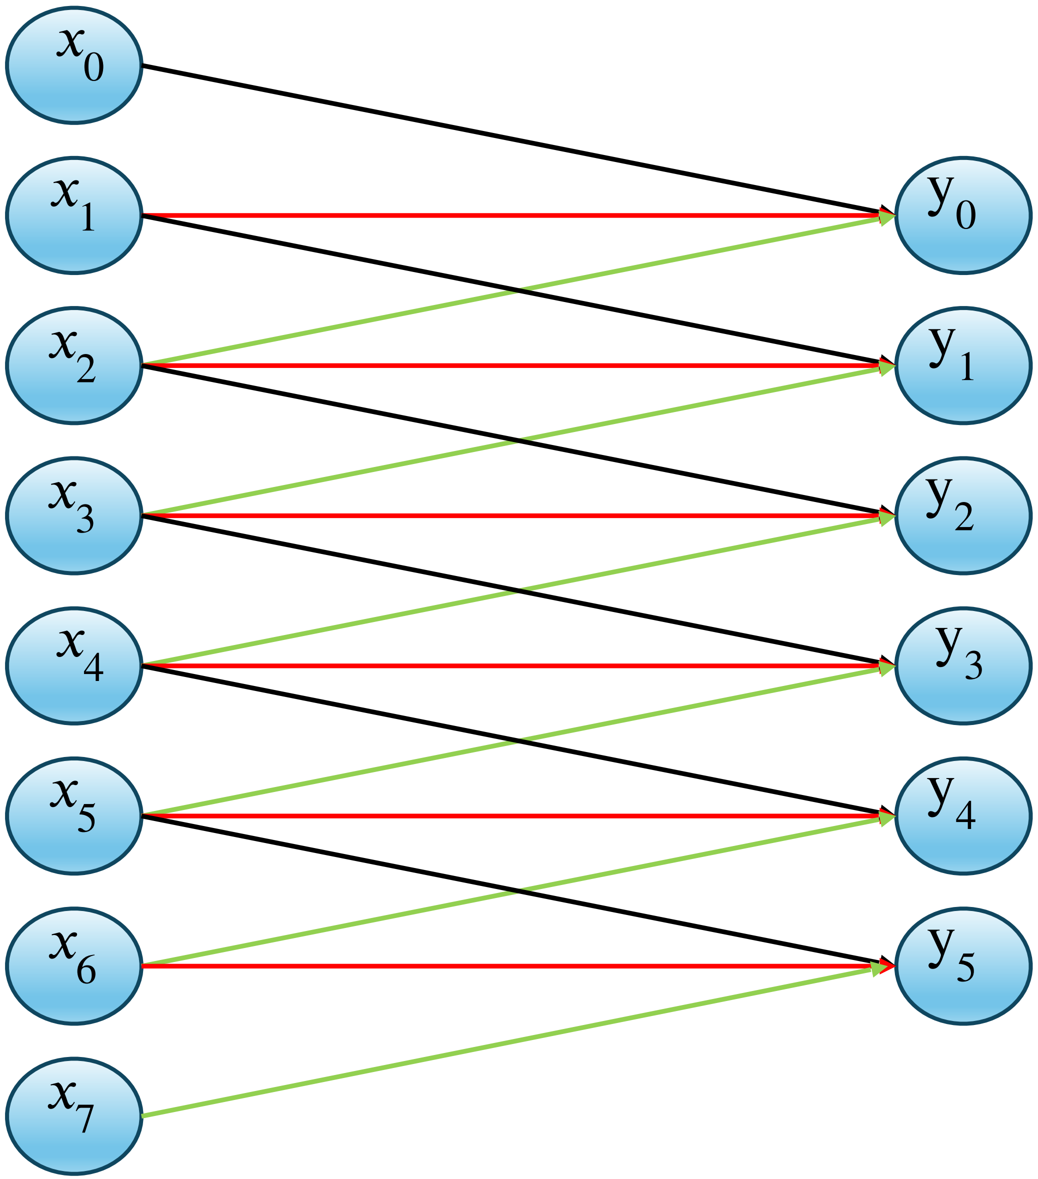

Is it even linear?

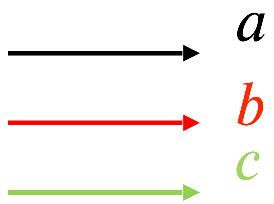

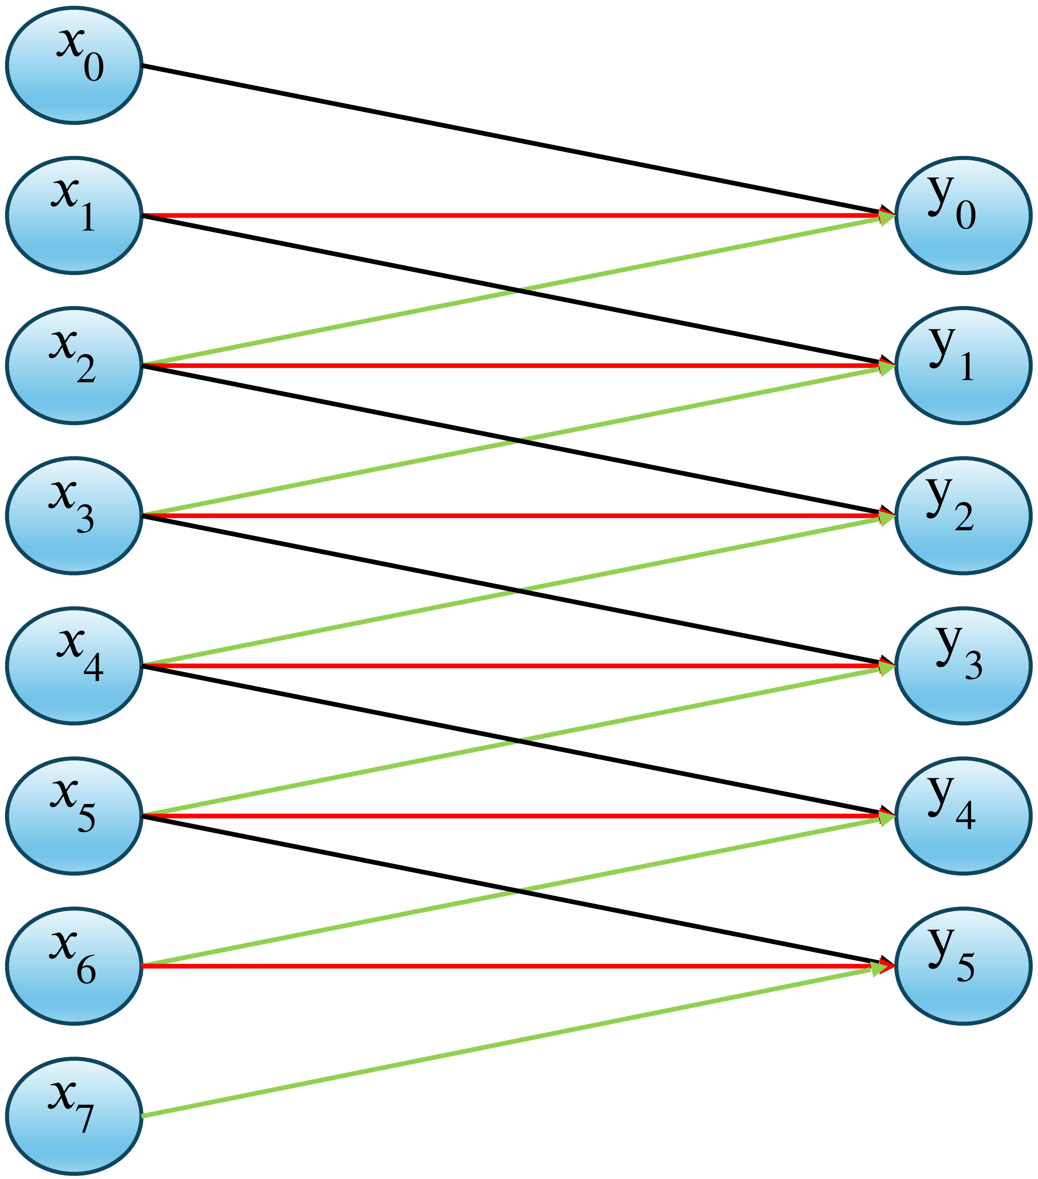

$$Y_t = \{f * g\}[t] \qquad g = [c,\ b,\ a]$$

$\boldsymbol{a}$ $\boldsymbol{b}$

Is it even linear?

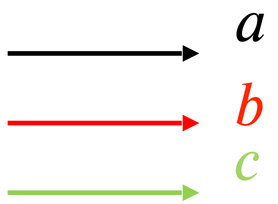

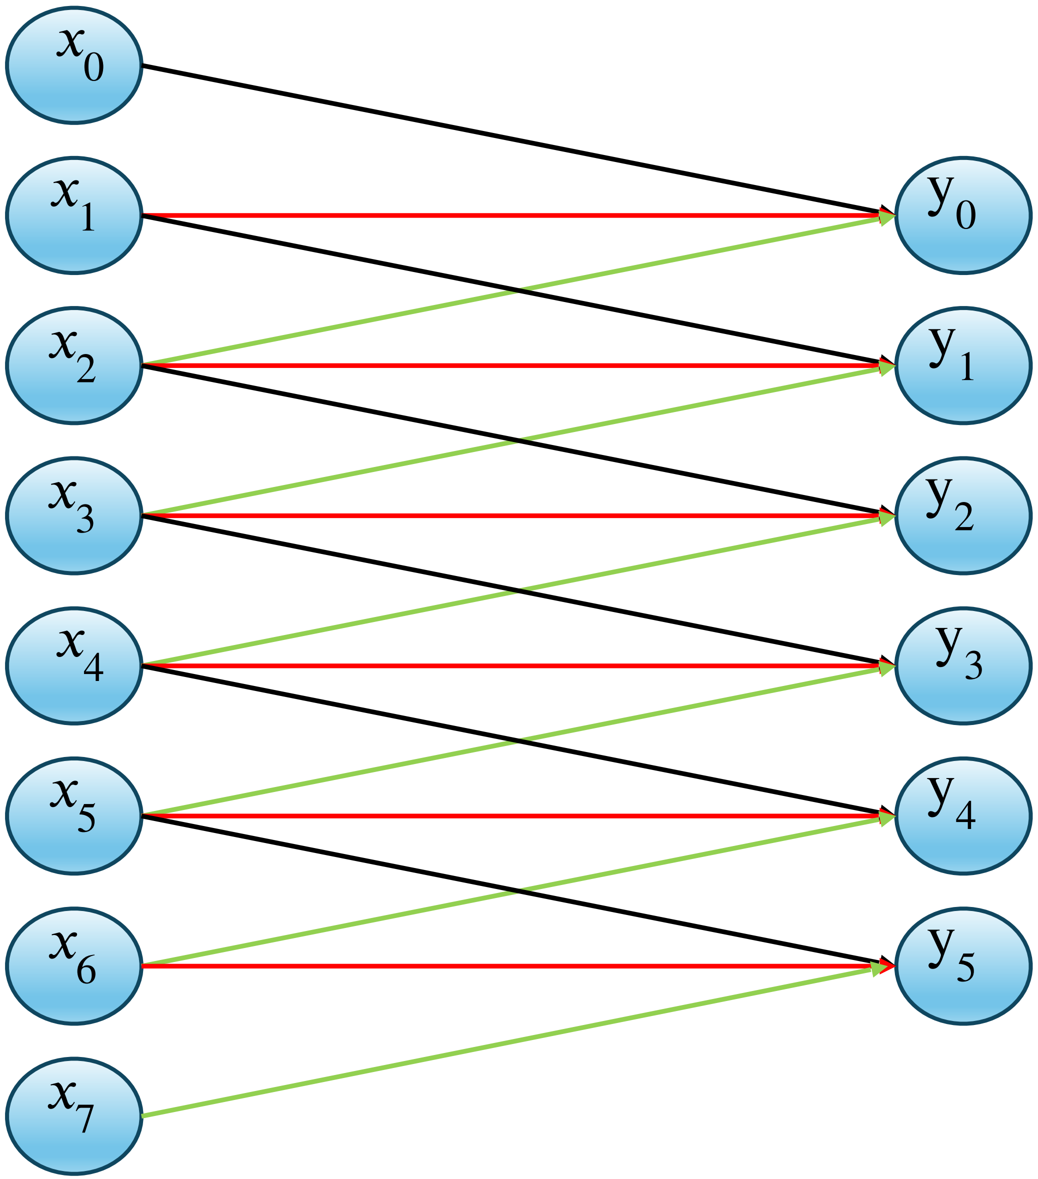

$$\boldsymbol{Y}_t = \{f * g\}[t] \qquad g = [c,\ b,\ a]$$

$$\begin{bmatrix} \boldsymbol{a} & \boldsymbol{b} & \boldsymbol{c} & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

Is it even linear?

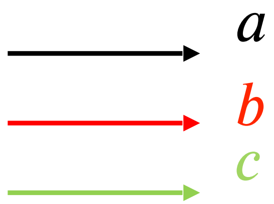

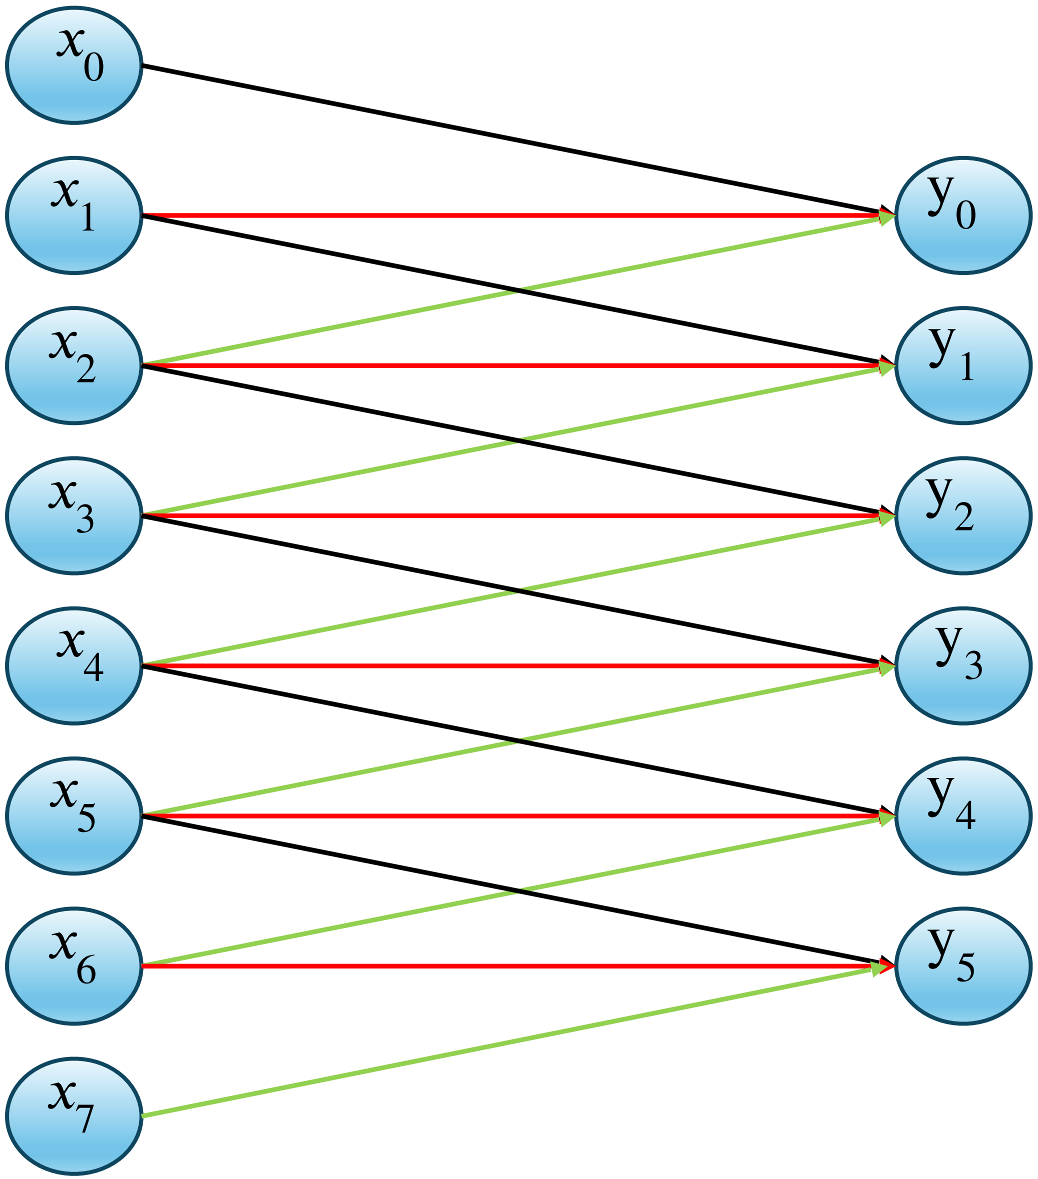

$$Y_t = \{f * g\}[t] \qquad g = [c,\ b,\ a]$$

$$\begin{bmatrix} a & b & c & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & a & b & c \end{bmatrix}$$

Is it even linear?

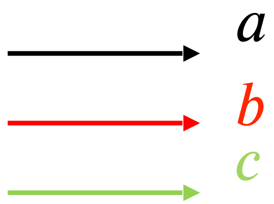

$$Y_t = \{f * g\}[t] \qquad g = [c,\ b,\ a]$$

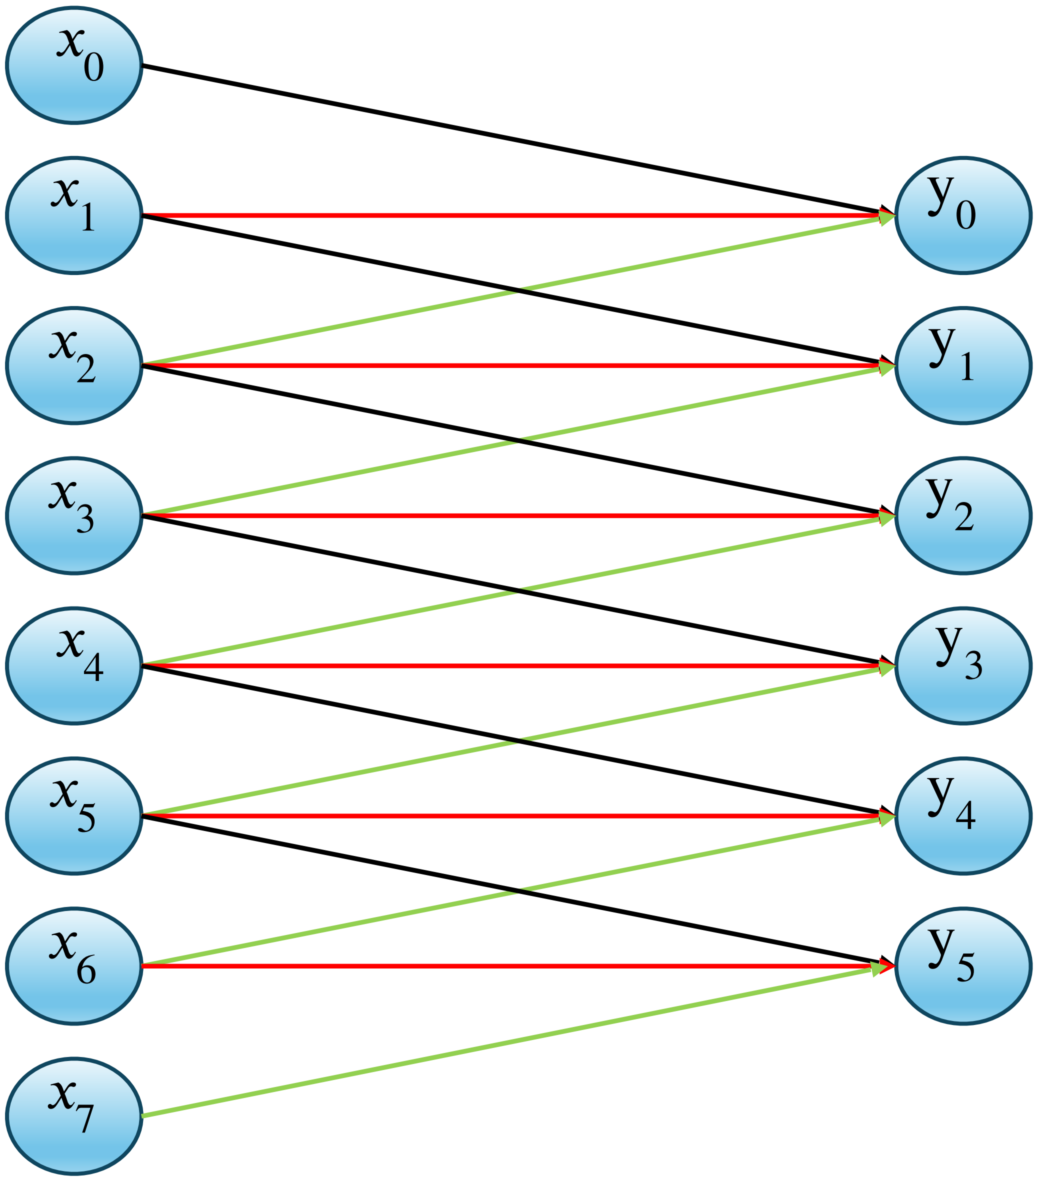

$$\begin{bmatrix} a & b & c & 0 & 0 & 0 & 0 & 0 \\ 0 & a & \cdots \\ \vdots \end{bmatrix}$$

Is it even linear?

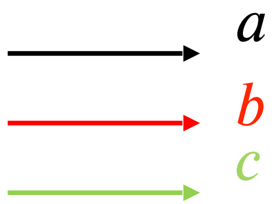

$$Y_t = \{f * g\}[t] \qquad g = [c,\ b,\ a]$$

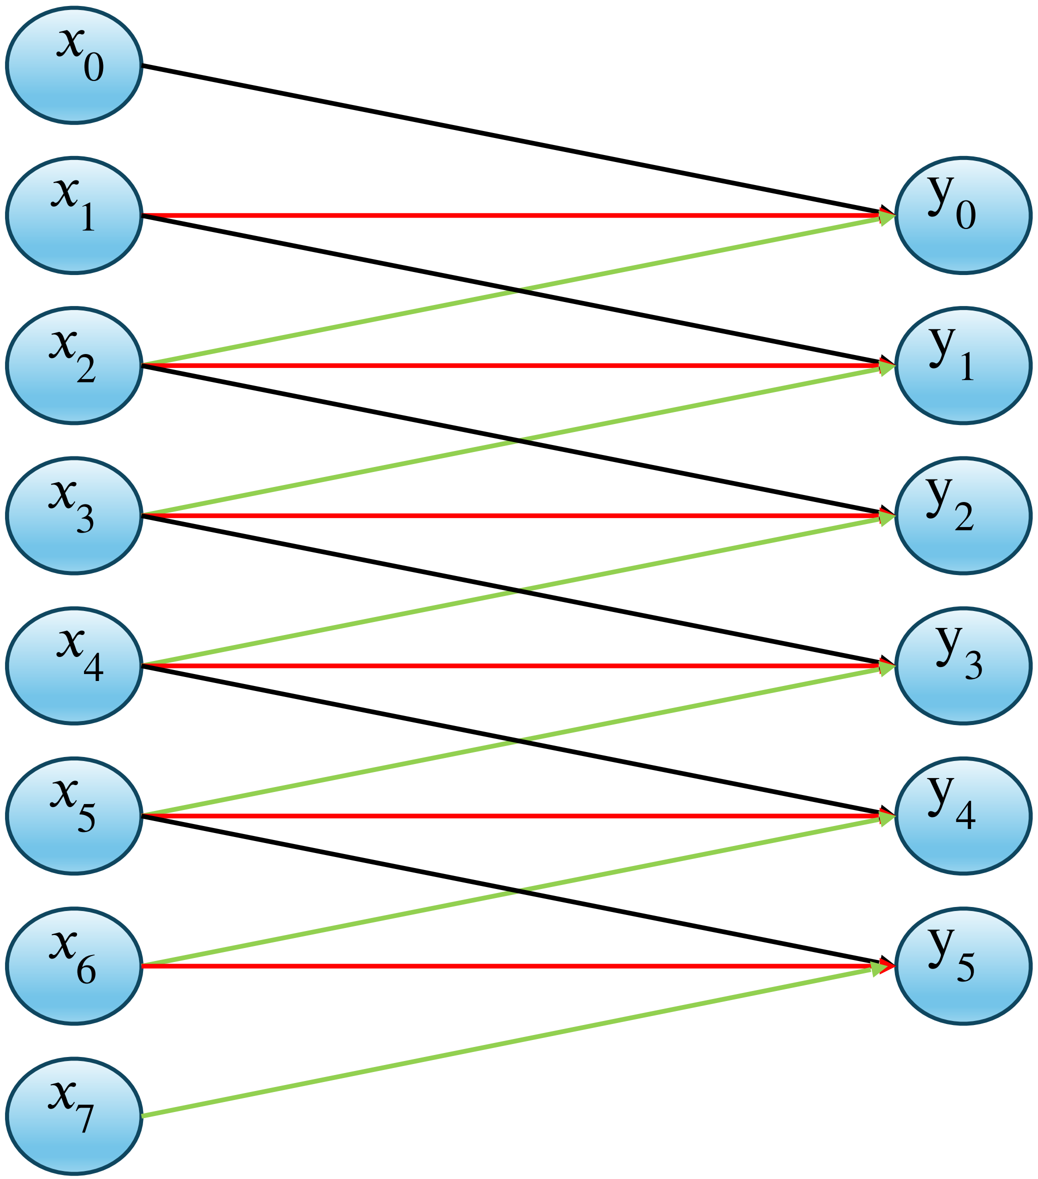

$$\begin{bmatrix} a & b & c & 0 & 0 & 0 & 0 & 0 \\ 0 & a & b & c & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & a & b & c \end{bmatrix}$$

Is it even linear?

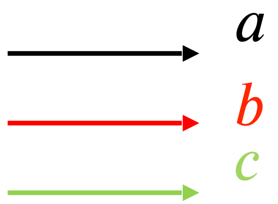

$$Y_t = \{f * g\}[t] \qquad g = [c,\ b,\ a]$$

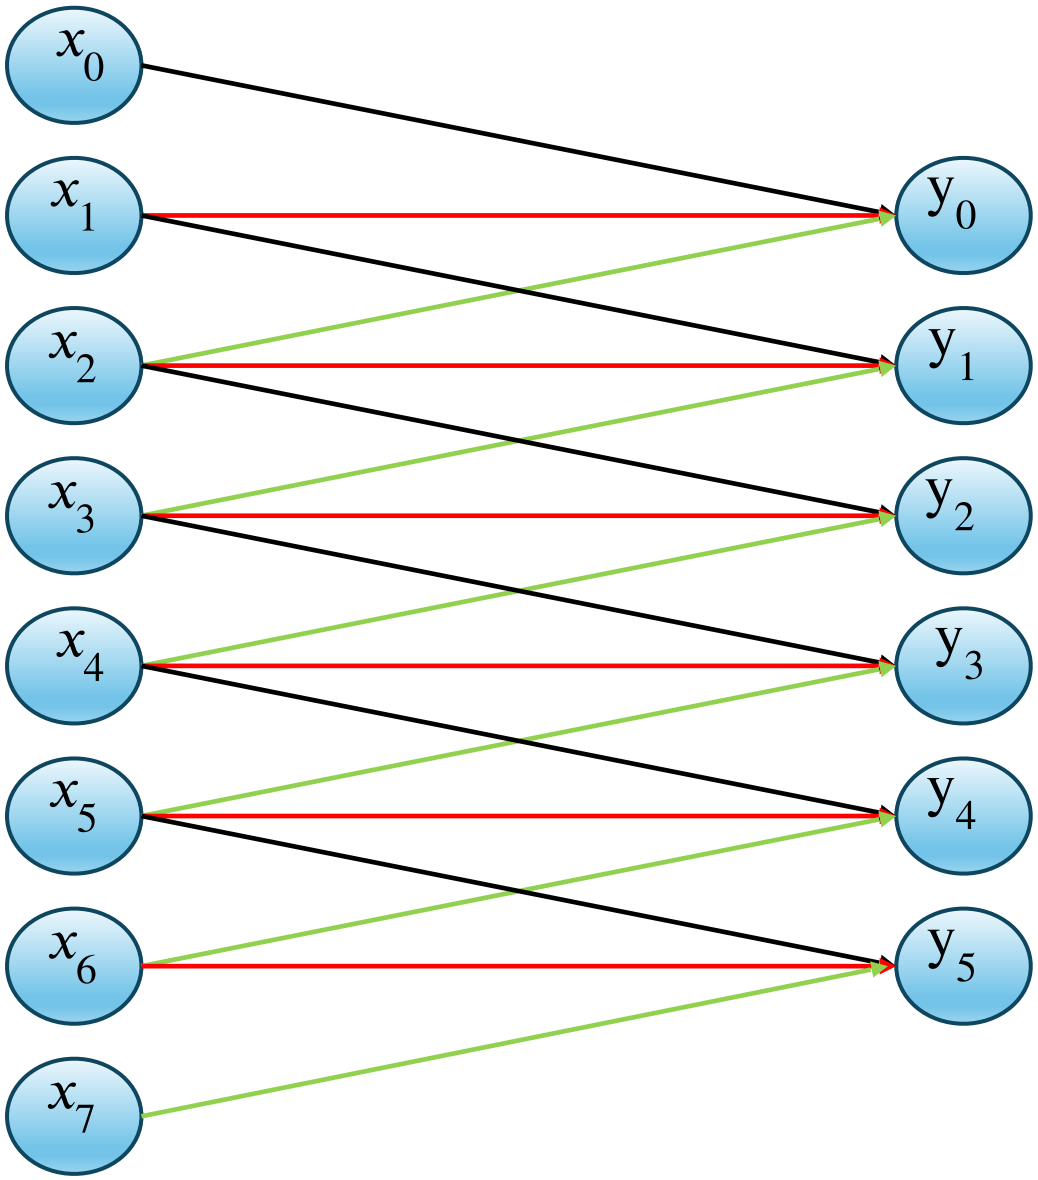

$$\begin{bmatrix} a & b & c & 0 & 0 & 0 & 0 & 0 \\ 0 & a & b & c & 0 & 0 & 0 & 0 \\ 0 & 0 & a & b & c & 0 & 0 & 0 \\ 0 & 0 & 0 & a & b & c & 0 & 0 \\ 0 & 0 & 0 & 0 & a & b & c & 0 \\ 0 & 0 & 0 & 0 & 0 & a & b & c \end{bmatrix}$$

Is it even linear?

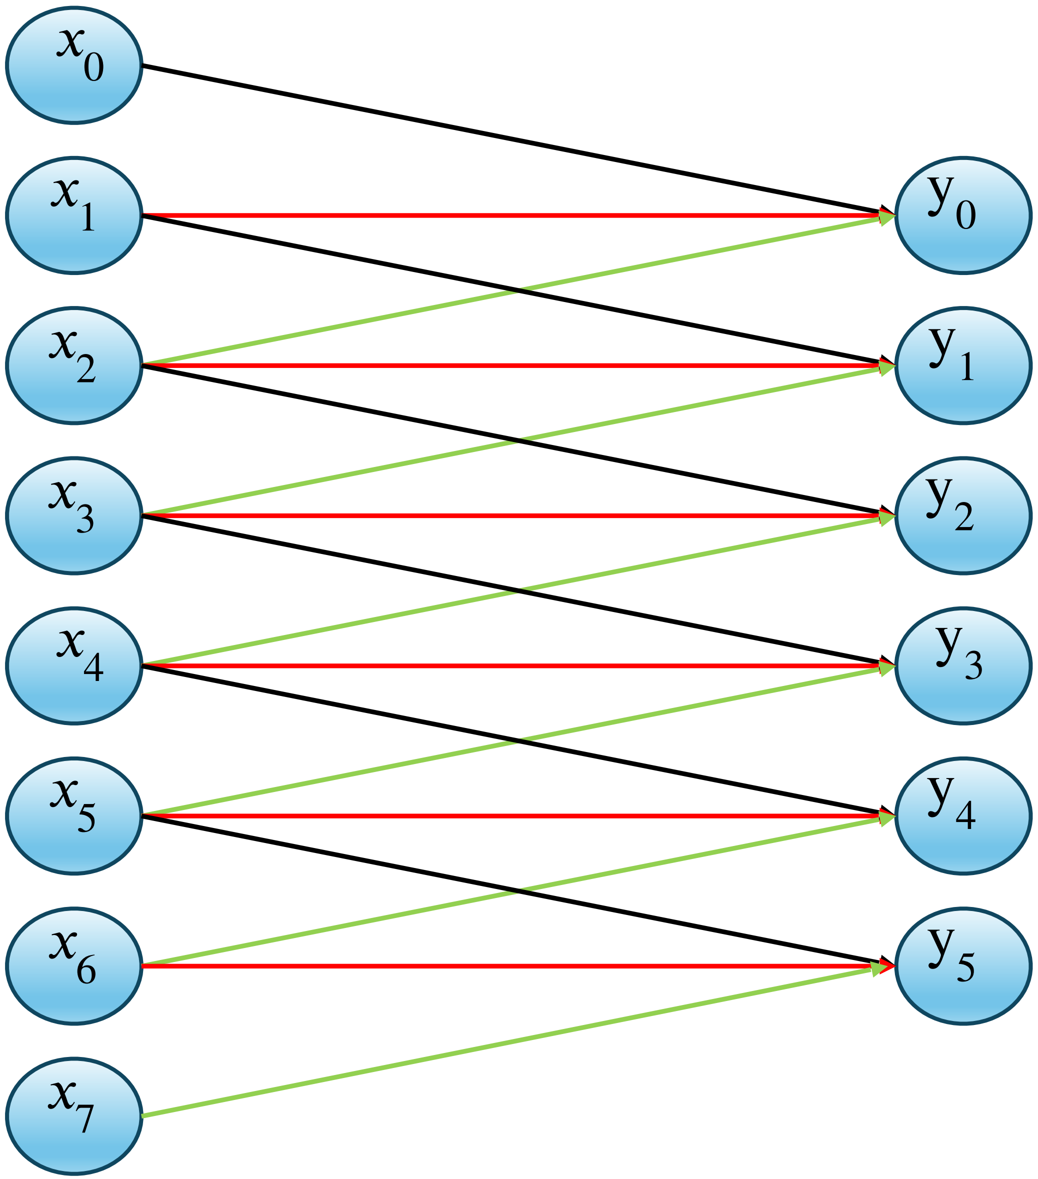

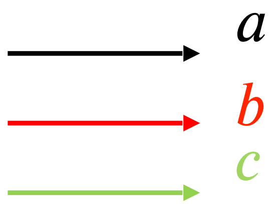

$$\boldsymbol{Y}_t = \{f * g\}[t] \qquad g = [c,\ b,\ a]$$

$$\begin{bmatrix} a & b & c & 0 & 0 & 0 & 0 & 0 \\ 0 & a & b & c & 0 & 0 & 0 & 0 \\ 0 & 0 & a & b & c & 0 & 0 & 0 \\ 0 & 0 & 0 & a & b & c & 0 & 0 \\ 0 & 0 & 0 & 0 & a & b & c & 0 \\ 0 & 0 & 0 & 0 & 0 & a & b & c \end{bmatrix}$$

Toeplitz matrix

Is it even linear?

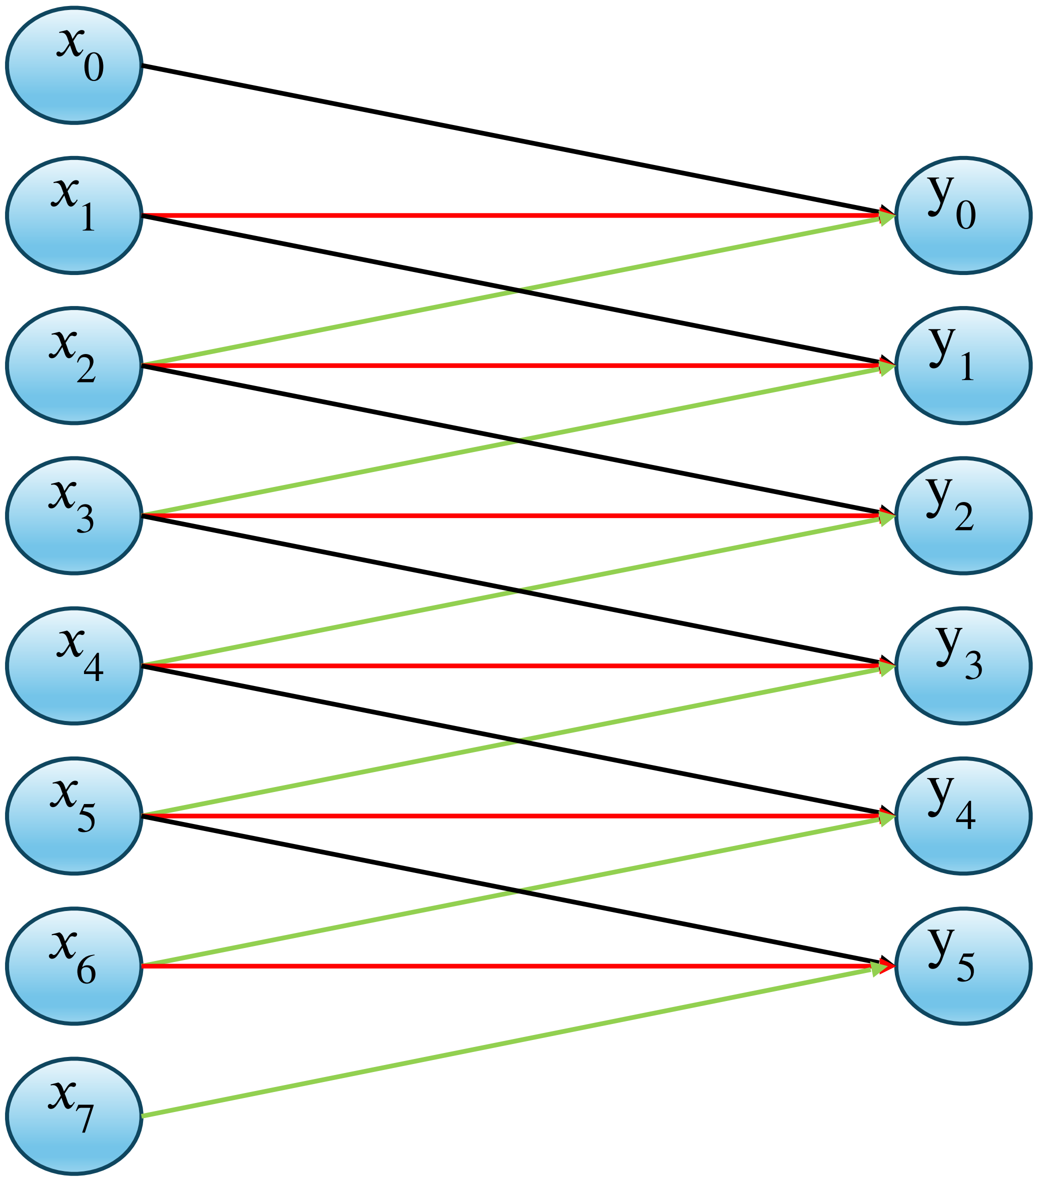

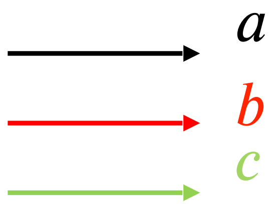

$$Y_t = \{f * g\}[t] \qquad g = [c,\ b,\ a]$$

$$\begin{bmatrix} a & b & c & 0 & 0 & 0 & 0 & 0 \\ 0 & a & b & c & 0 & 0 & 0 & 0 \\ 0 & 0 & a & b & c & 0 & 0 & 0 \\ 0 & 0 & 0 & a & b & c & 0 & 0 \\ 0 & 0 & 0 & 0 & a & b & c & 0 \\ 0 & 0 & 0 & 0 & 0 & a & b & c \end{bmatrix}$$

Toeplitz matrix

Common: cyclic conv. Not covering it at this point

Demo:

Continuous Convolution

$$\{f * g\}(t) = \int_{-\infty}^{\infty} f(t)g(t-\tau)d\tau$$

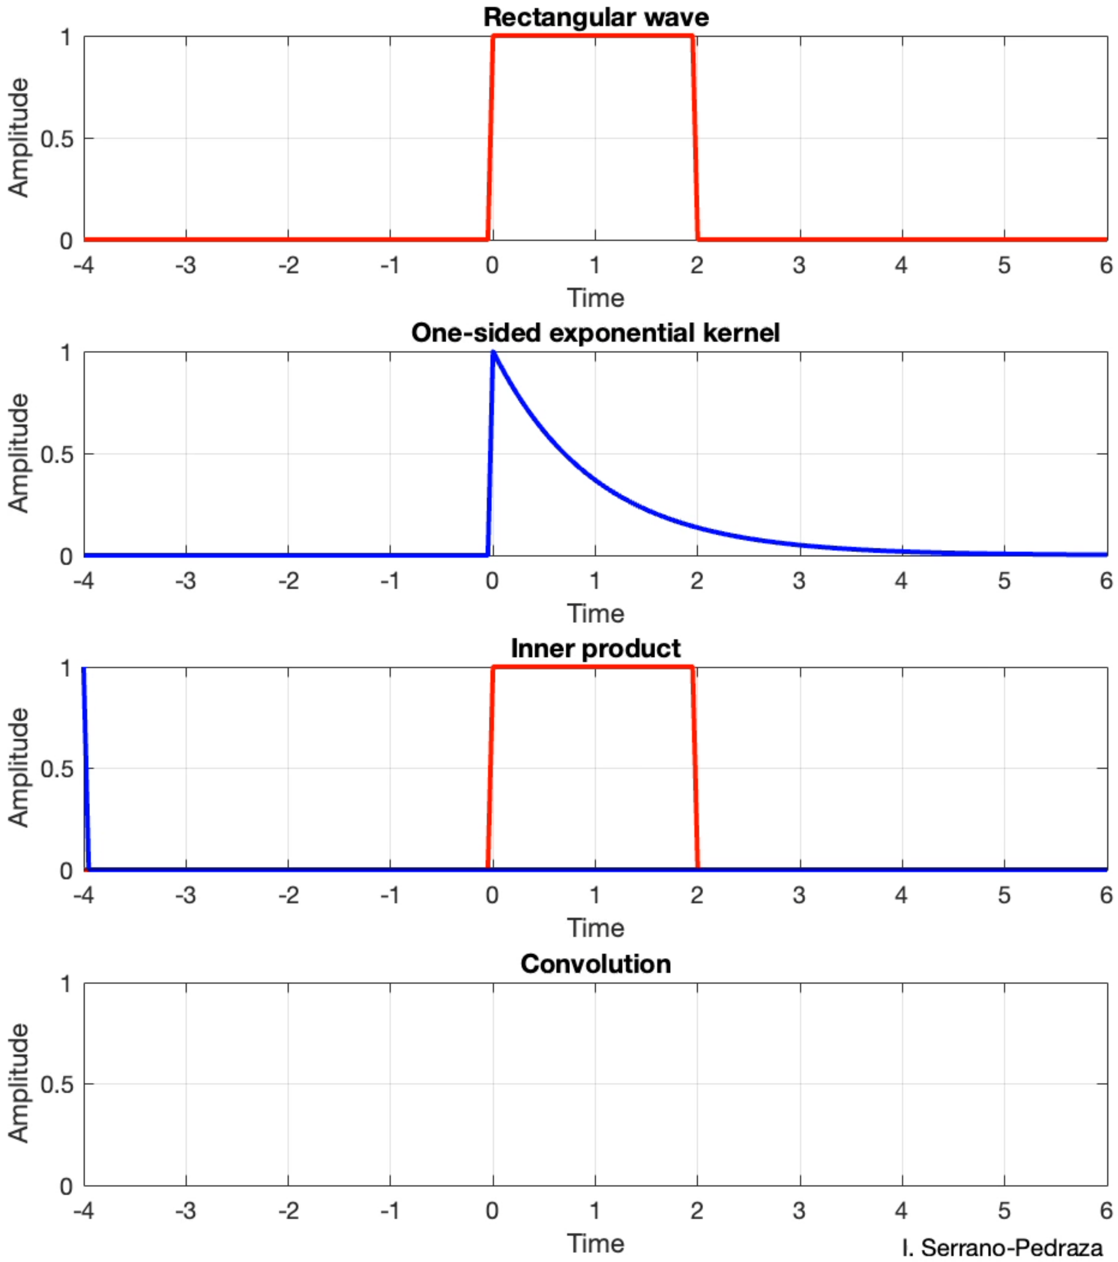

Continuous Convolution

$$\{f * g\}(t) = \int_{-\infty}^{\infty} f(t)g(t-\tau)d\tau$$

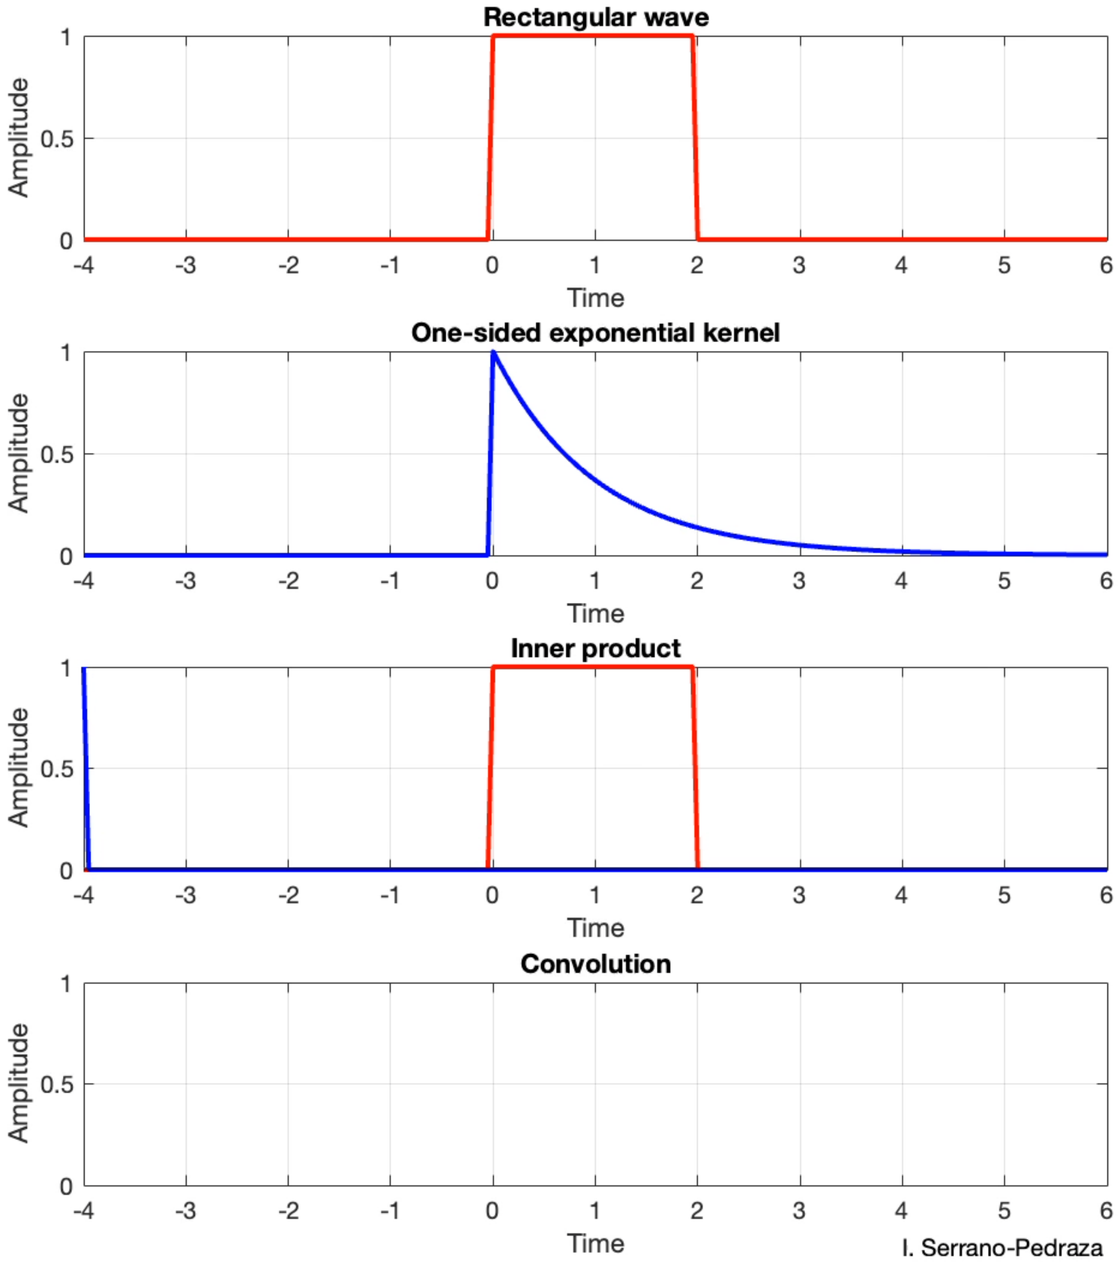

Impulse Response

Impulse Response

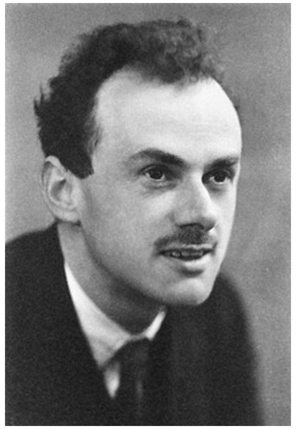

Dirac Delta

$$\delta(x) = \begin{cases} \infty \; ; & 0 \\ 0 \; ; & else \end{cases}$$

Impulse Response

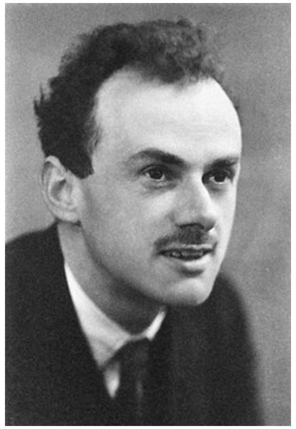

Dirac Delta

$$\delta(x) = \begin{cases} \infty \; ; & 0 \\ 0 \; ; & else \end{cases}$$

$$\int_{-\infty}^{\infty} \delta(x) dx = 1$$ 🤯

Impulse Response

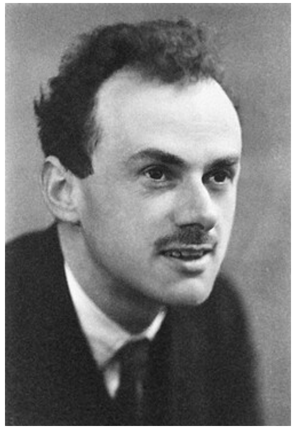

Dirac Delta

$$\delta(x) = \begin{cases} \infty \ ; & 0 \\ 0 \ ; & else \end{cases}$$

$$\int_{-\infty}^{\infty} \delta(x)dx = 1$$ 🤯

$$\delta(x) = \lim_{\varepsilon \to 0} \frac{1}{\varepsilon} \sqcap \left(\frac{x}{\varepsilon} \right)$$

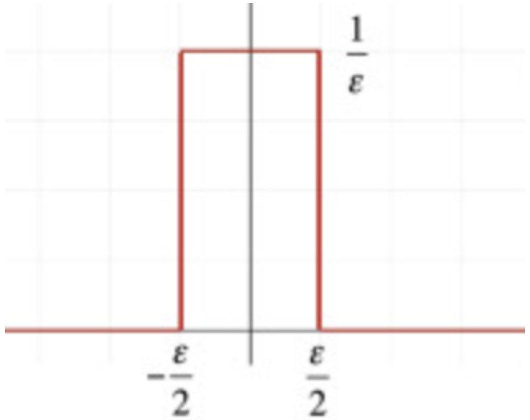

Impulse Response

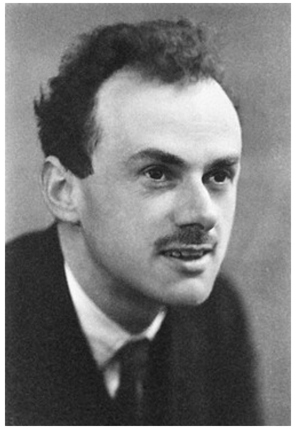

Dirac Delta

$$\delta(x) = \begin{cases} \infty\ ; & 0 \\ 0\ ; & else \end{cases}$$

$$\int_{-\infty}^{\infty} \delta(x)dx = 1$$ 🤯

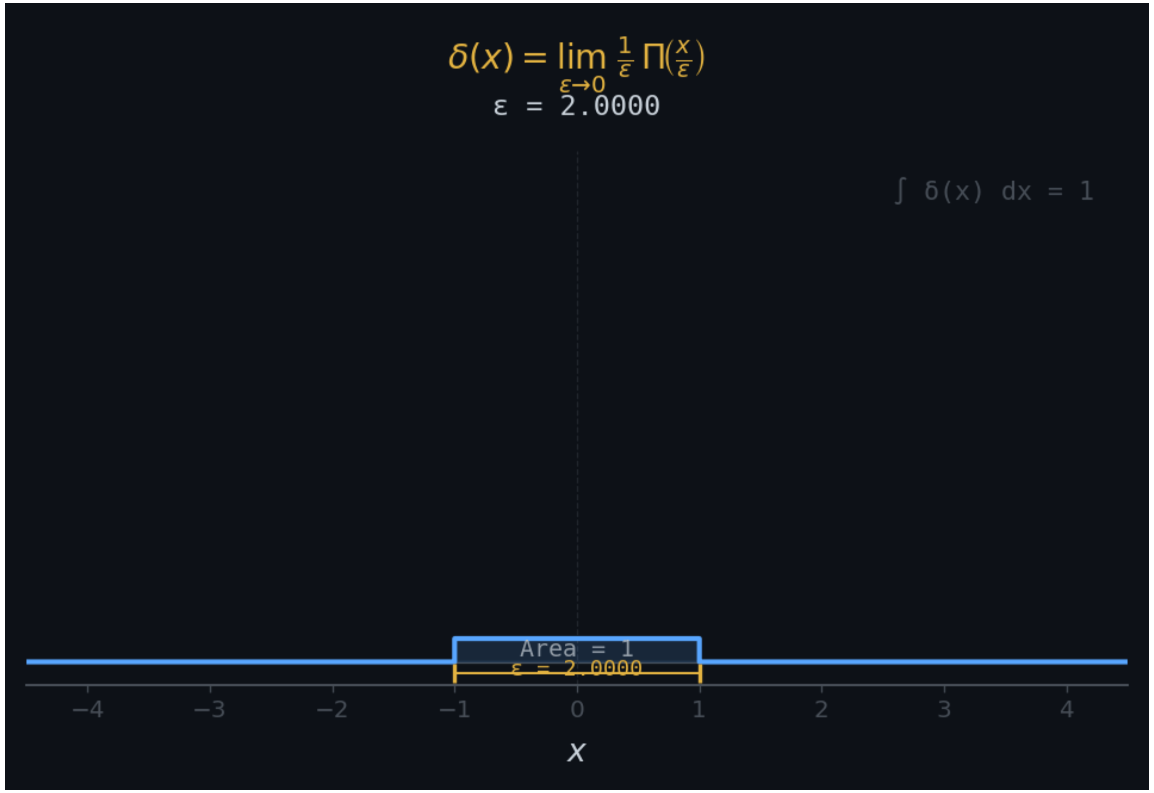

Impulse Response

Q: $\{f * \delta\} = ?$

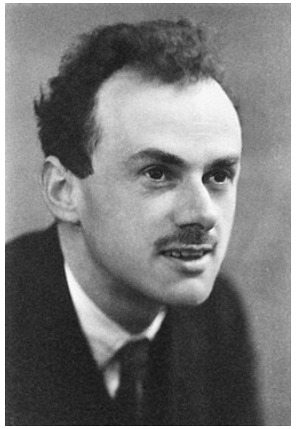

Dirac Delta

$$\delta(x) = \begin{cases} \infty \, ; & 0 \\ 0 \, ; & else \end{cases}$$

$$\int_{}^{\infty} \delta(x)dx = 1$$ 🤯

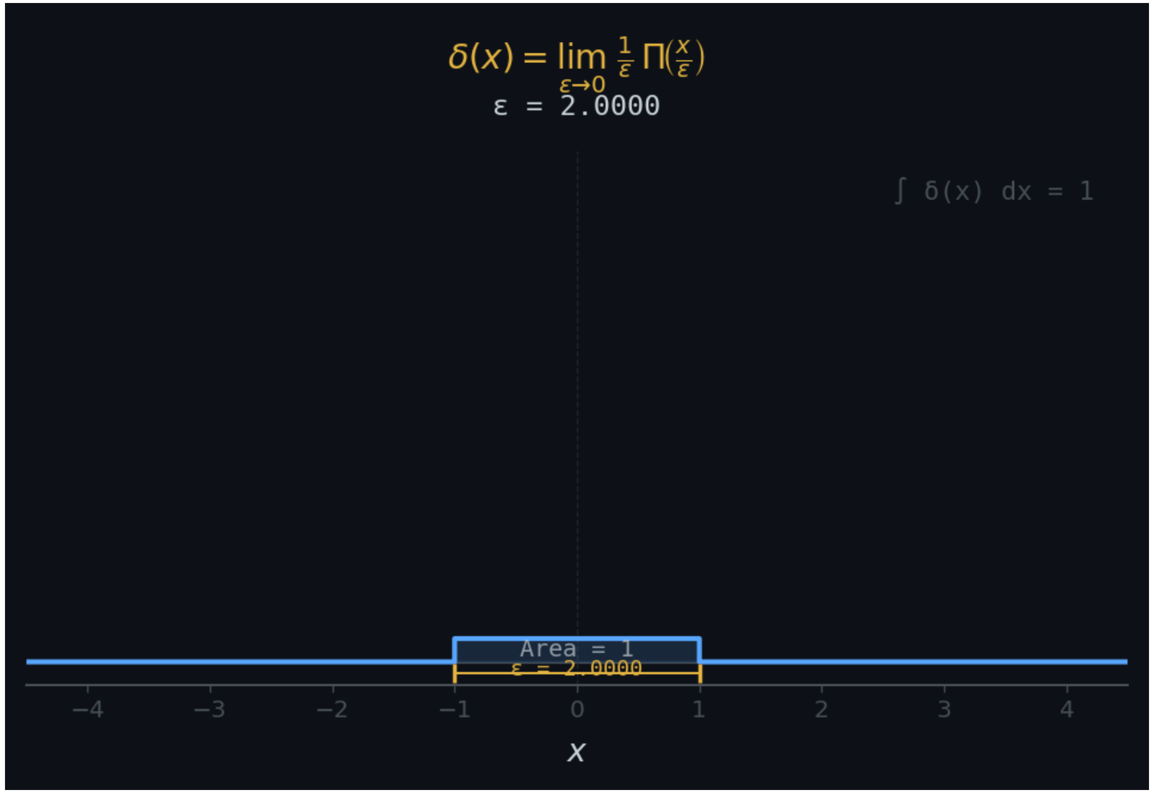

Impulse Response

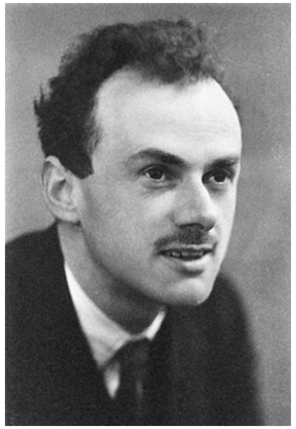

Dirac Delta

$$\delta(x) = \begin{cases} \infty \,; & 0 \\ 0 \,; & else \end{cases}$$

$$\int^{\infty} \delta(x)dx = 1$$

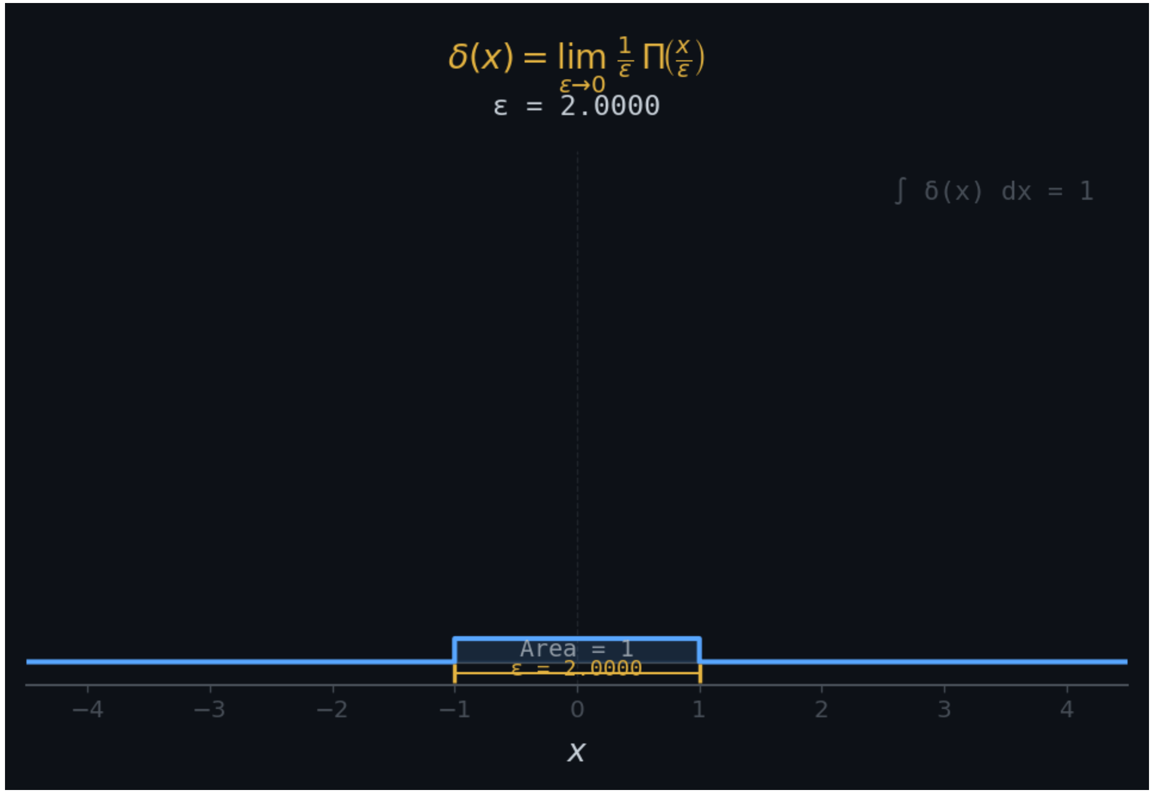

Q: $\{f * \delta\} = ?$

A: f

Proof: HW

Impulse Response

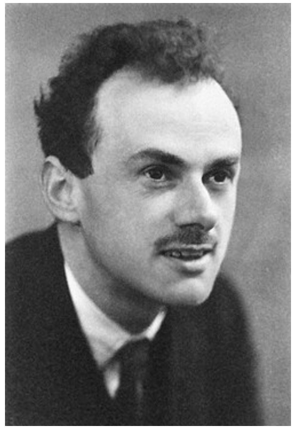

Dirac Delta

$$\delta(x) = \begin{cases} \infty \,; & 0 \\ 0 \,; & else \end{cases}$$

$$\int^{\infty} \delta(x)dx = 1$$ 🤯

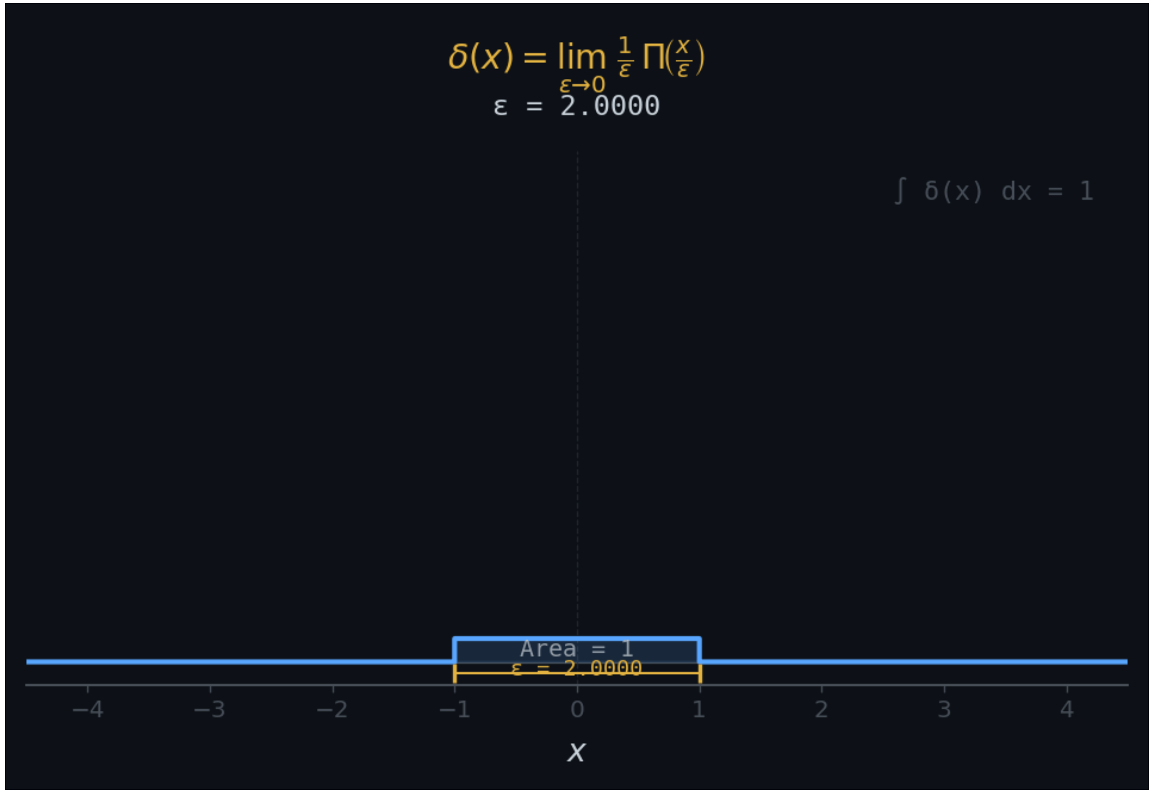

Q: $\{f * \delta\} = ?$

A: f

Proof: HW

Is it a function?

Impulse Response

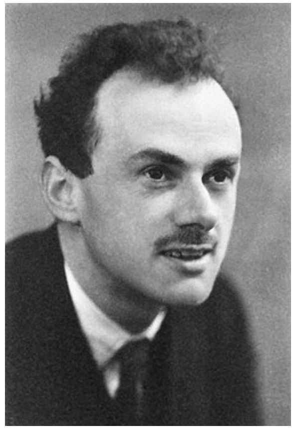

Dirac Delta

$$\delta(x) = \begin{cases} \infty\ ; & 0 \\ 0\ ; & else \end{cases}$$

$$\int^{\infty} \delta(x)dx = 1$$

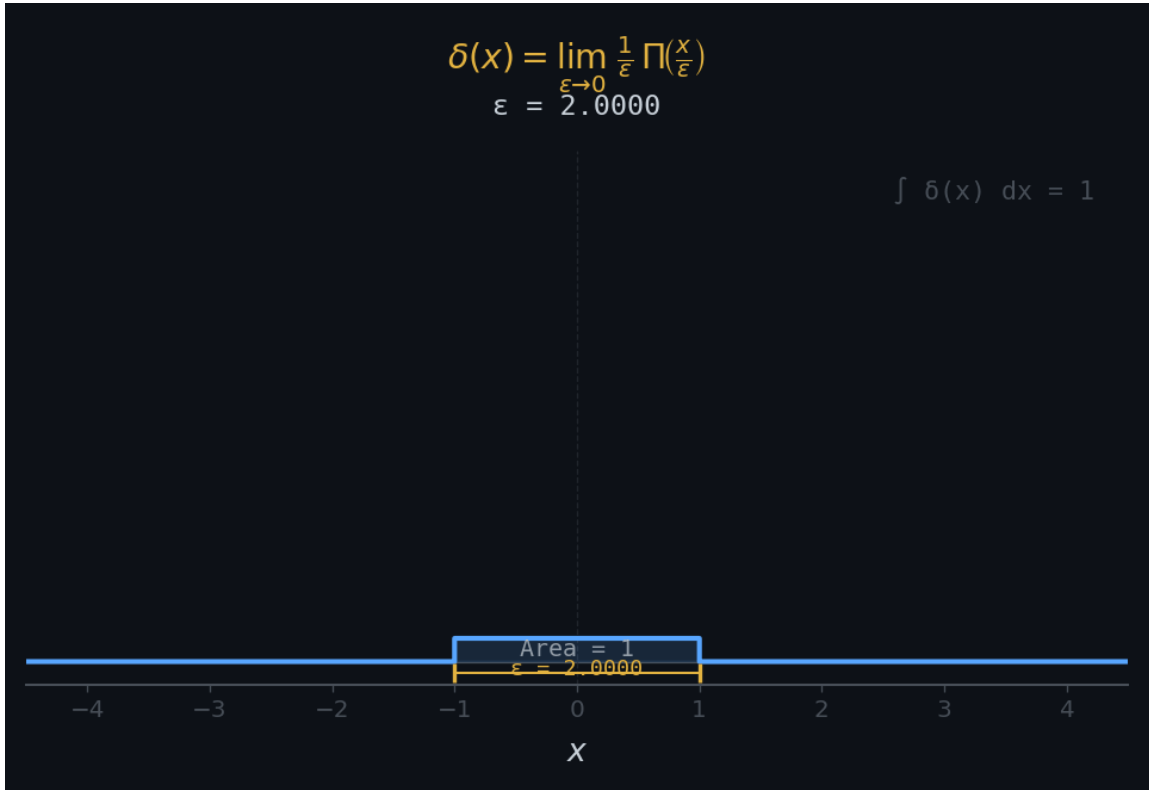

Q: $\{f * \delta\} = ?$

A: f

Proof: HW

Is it a function?

Distribution / Generalized function

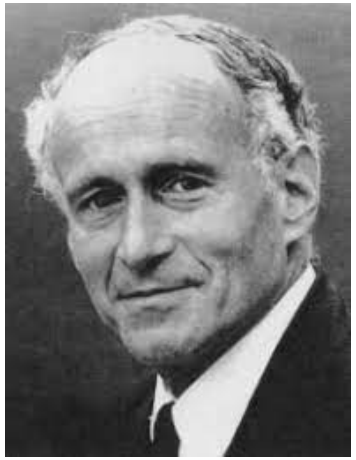

Laurent-Moïse Schwartz

Impulse Response

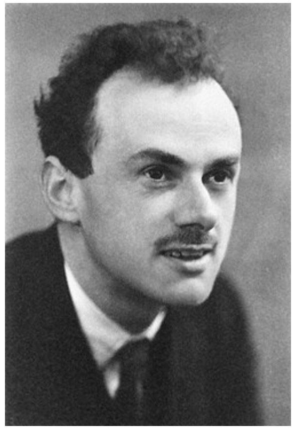

Dirac Delta

$$\delta(x) = \begin{cases} \infty \; ; & 0 \\ 0 \; ; & else \end{cases}$$

$$\int^{\infty} \delta(x) dx = 1$$

Q: $\{f * \delta\} = ?$

A: f

Proof: HW

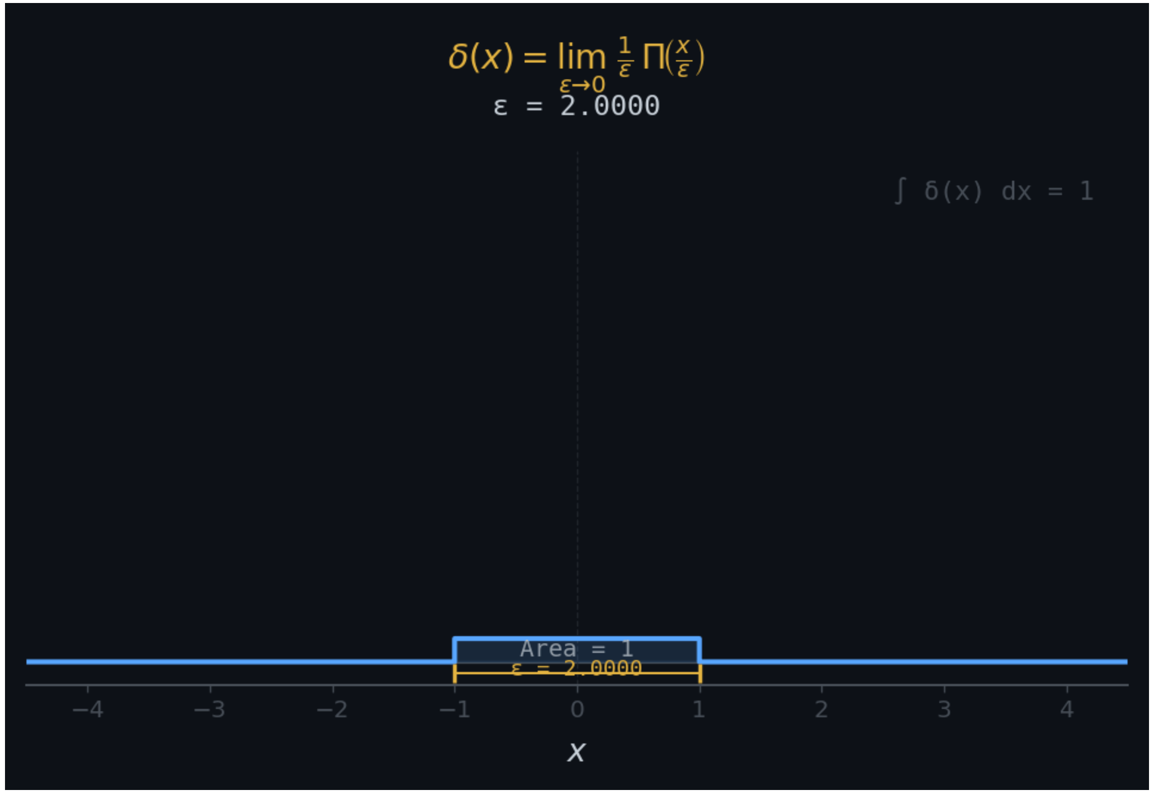

Is it a function?

Distribution / Generalized function

Don’t hate me, I’m not that Schwartz!

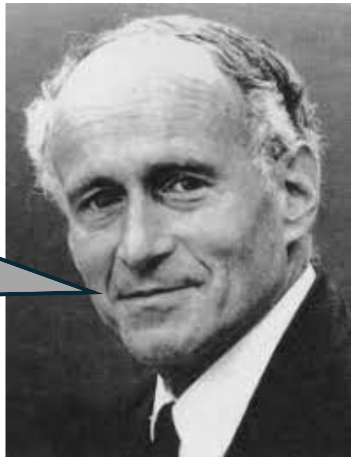

Laurent-Moïse Schwartz

Properties

Commutativity	$f * g = g * f$
Associativity	$(f * g) * h = f * (g * h)$
Distributivity	$f * (g + h) = f * g + f * h$
Linearity	$f * (ag + bh) = a(f * g) + b(f * h)$
Identity	$f * \delta = f$
Shift equivariance	$f * g(\cdot - \tau) = (f * g)(\cdot - \tau)$
Differentiation	$\frac{d}{dt}(f * g) = f' * g = f * g'$
Support addition	$\mathrm{supp}(f * g) \subseteq \mathrm{supp}(f) + \mathrm{supp}(g)$

Properties

Commutativity	$f * g = g * f$
Associativity	$(f * g) * h = f * (g * h)$
Distributivity	$f * (g + h) = f * g + f * h$
Linearity	$f * (ag + bh) = a(f * g) + b(f * h)$
Identity	$f * \delta = f$
Shift equivariance	$f * g(\cdot - \tau) = (f * g)(\cdot - \tau)$
Differentiation	$\frac{d}{dt}(f * g) = f' * g = f * g'$
Support addition	$\mathrm{supp}(f * g) \subseteq \mathrm{supp}(f) + \mathrm{supp}(g)$

Properties

Commutativity	$f * g = g * f$
Associativity	$(f * g) * h = f * (g * h)$
Distributivity	$f * (g + h) = f * g + f * h$
Linearity	$f * (ag + bh) = a(f * g) + b(f * h)$
Identity	$f * \delta = f$
Shift equivariance	$f * g(\cdot - \tau) = (f * g)(\cdot - \tau)$
Differentiation	$\frac{d}{dt}(f * g) = f' * g = f * g'$
Support addition	$\mathrm{supp}(f * g) \subseteq \mathrm{supp}(f) + \mathrm{supp}(g)$

Properties

Commutativity	$f * g = g * f$
Associativity	$(f * g) * h = f * (g * h)$
Distributivity	$f * (g + h) = f * g + f * h$
Linearity	$f * (ag + bh) = a(f * g) + b(f * h)$
Identity	$f * \delta = f$
Shift equivariance	$f * g(\cdot - \tau) = (f * g)(\cdot - \tau)$
Differentiation	$\frac{d}{dt}(f * g) = f' * g = f * g'$
Support addition	$\operatorname{supp}(f * g) \subseteq \operatorname{supp}(f) + \operatorname{supp}(g)$

Properties

Commutativity	$f * g = g * f$
Associativity	$(f * g) * h = f * (g * h)$
Distributivity	$f * (g + h) = f * g + f * h$
Linearity	$f * (ag + bh) = a(f * g) + b(f * h)$
Identity	$f * \delta = f$
Shift equivariance	$f * g(\cdot - \tau) = (f * g)(\cdot - \tau)$
Differentiation	$\frac{d}{dt}(f * g) = f' * g = f * g'$
Support addition	$\operatorname{supp}(f * g) \subseteq \operatorname{supp}(f) + \operatorname{supp}(g)$

Properties

Commutativity	$f * g = g * f$
Associativity	$(f * g) * h = f * (g * h)$
Distributivity	$f * (g + h) = f * g + f * h$
Linearity	$f * (ag + bh) = a(f * g) + b(f * h)$
Identity	$f * \delta = f$
Shift equivariance	$f * g(\cdot - \tau) = (f * g)(\cdot - \tau)$
Differentiation	$\frac{d}{dt}(f * g) = f' * g = f * g'$
Support addition	$\mathrm{supp}(f * g) \subseteq \mathrm{supp}(f) + \mathrm{supp}(g)$

Properties

Commutativity	$f * g = g * f$
Associativity	$(f * g) * h = f * (g * h)$
Distributivity	$f * (g + h) = f * g + f * h$
Linearity	$f * (ag + bh) = a(f * g) + b(f * h)$
Identity	$f * \delta = f$
Shift equivariance	$f * g(\cdot - \tau) = (f * g)(\cdot - \tau)$
Differentiation	$\frac{d}{dt}(f * g) = f' * g = f * g'$
Support addition	$\mathrm{supp}(f * g) \subseteq \mathrm{supp}(f) + \mathrm{supp}(g)$

Properties

Commutativity	$f * g = g * f$
Associativity	$(f * g) * h = f * (g * h)$
Distributivity	$f * (g + h) = f * g + f * h$
Linearity	$f * (ag + bh) = a(f * g) + b(f * h)$
Identity	$f * \delta = f$
Shift equivariance	$f * g(\cdot - \tau) = (f * g)(\cdot - \tau)$
Differentiation	$\frac{d}{dt}(f * g) = f' * g = f * g'$
Support addition	$\mathrm{supp}(f * g) \subseteq \mathrm{supp}(f) + \mathrm{supp}(g)$

Properties

Commutativity	$f * g = g * f$
Associativity	$(f * g) * h = f * (g * h)$
Distributivity	$f * (g + h) = f * g + f * h$
Linearity	$f * (ag + bh) = a(f * g) + b(f * h)$
Identity	$f * \delta = f$
Shift equivariance	$f * g(\cdot - \tau) = (f * g)(\cdot - \tau)$
Differentiation	$\frac{d}{dt}(f * g) = f' * g = f * g'$
Support addition	$\text{supp}(f * g) \subseteq \text{supp}(f) + \text{supp}(g)$

Properties

Commutativity	$f * g = g * f$
Associativity	$(f * g) * h = f * (g * h)$
Distributivity	$f * (g + h) = f * g + f * h$
Linearity	$f * (ag + bh) = a(f * g) + b(f * h)$
Identity	$f * \delta = f$
Shift equivariance	$f * g(\cdot - \tau) = (f * g)(\cdot - \tau)$
Differentiation	$\frac{d}{dt}(f * g) = f' * g = f * g'$
Support addition	$\text{supp}(f * g) \subseteq \text{supp}(f) + \text{supp}(g)$

Proof: HW

2D

$$(f * g)(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(u, v)\, g(x - u,\, y - v)\, du\, dv$$

$$(f * g)[m, n] = \sum_{k=-\infty}^{\infty} \sum_{\ell=-\infty}^{\infty} f[k, \ell]\, g[m - k,\, n - \ell]$$

Q: What is $x * x * x * x * x * x \ldots$

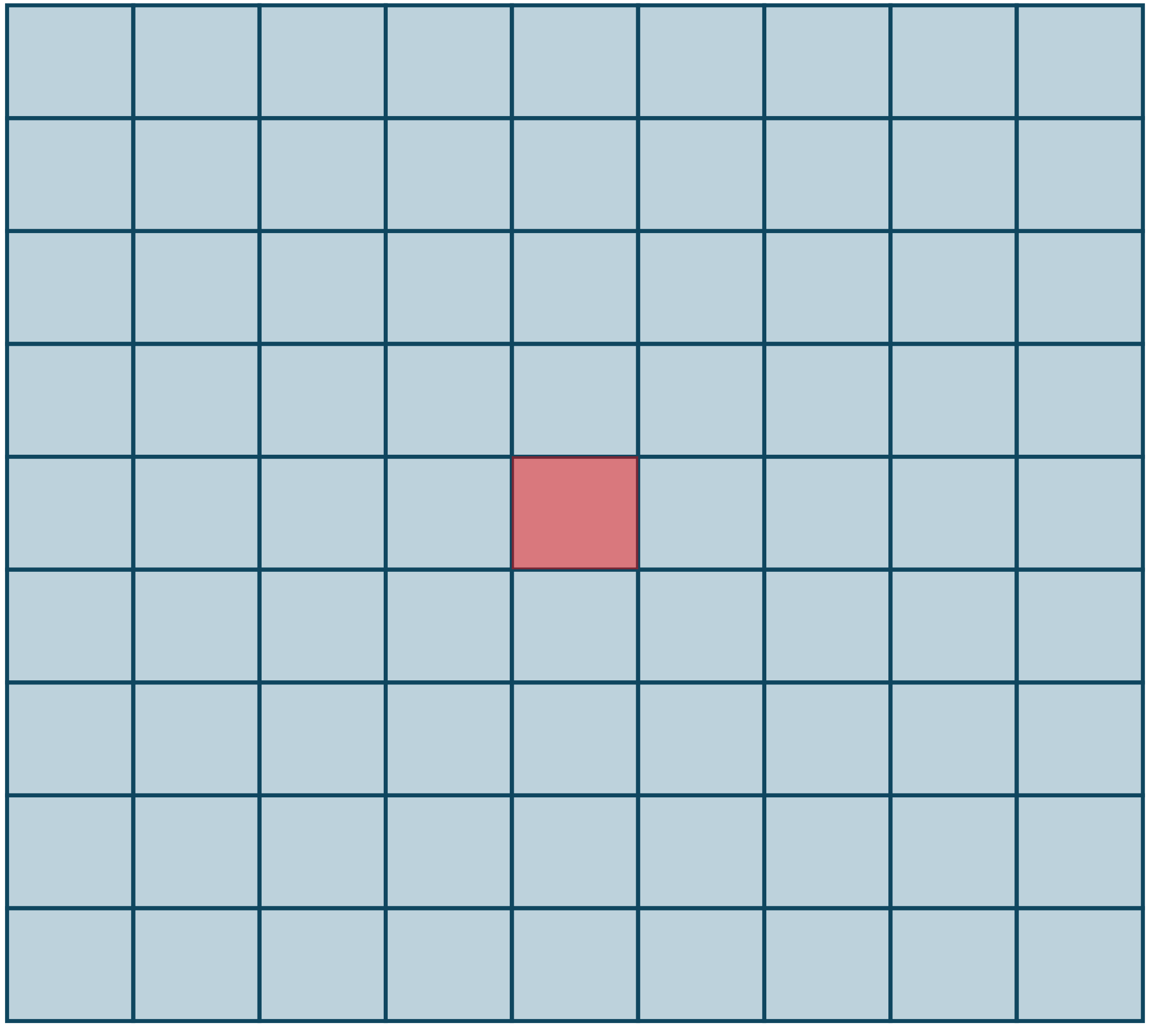

Q: What is $x * x * x * x * x * x \ldots$

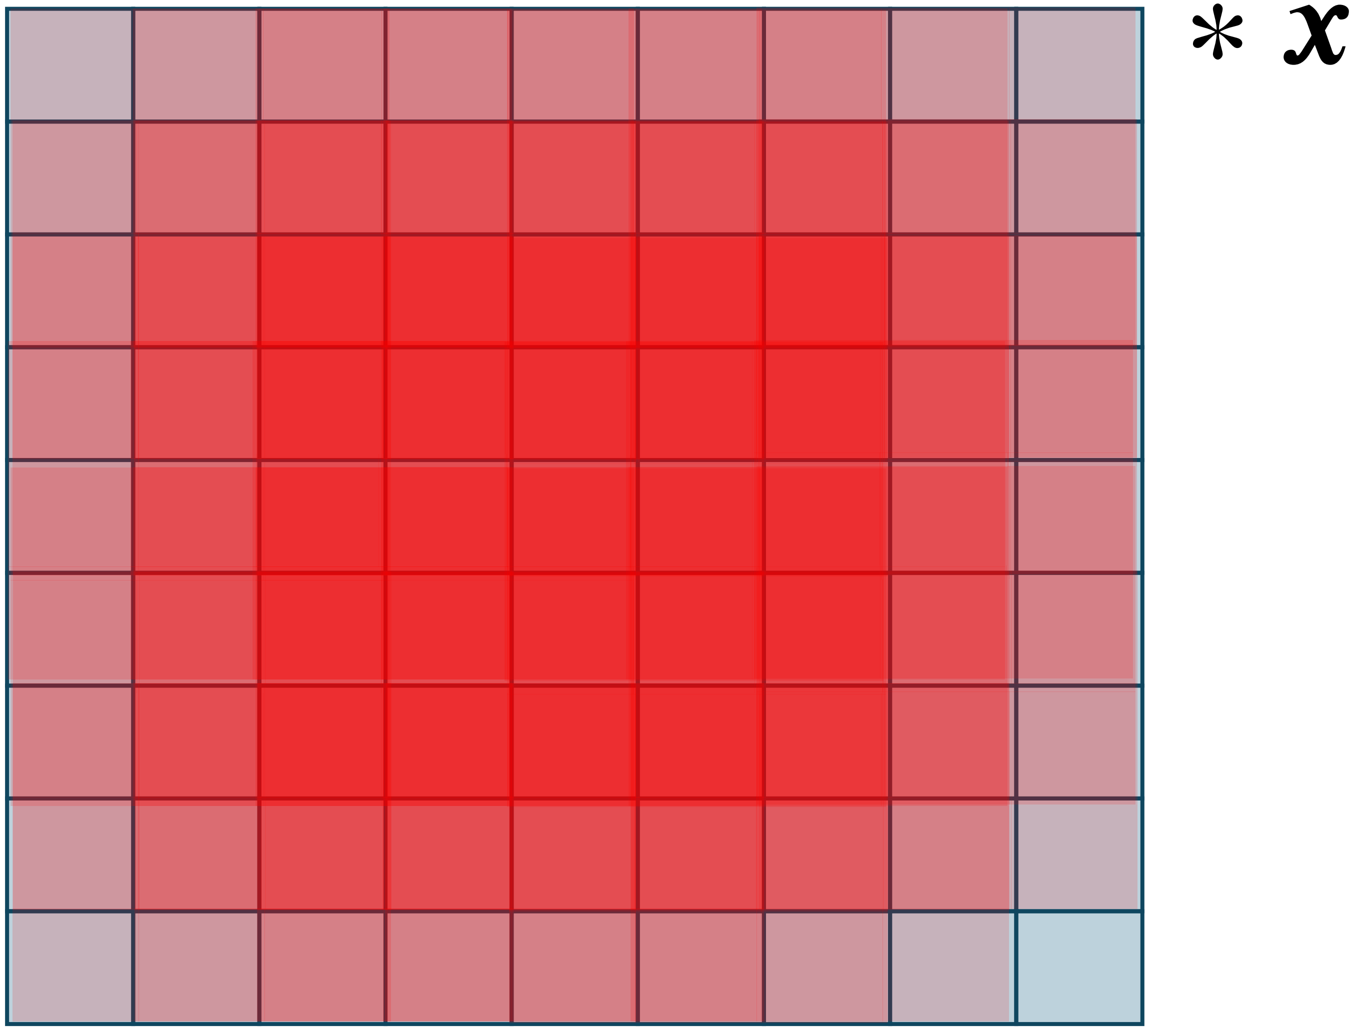

Q: What is $x * x * x * x * x * x \ldots$

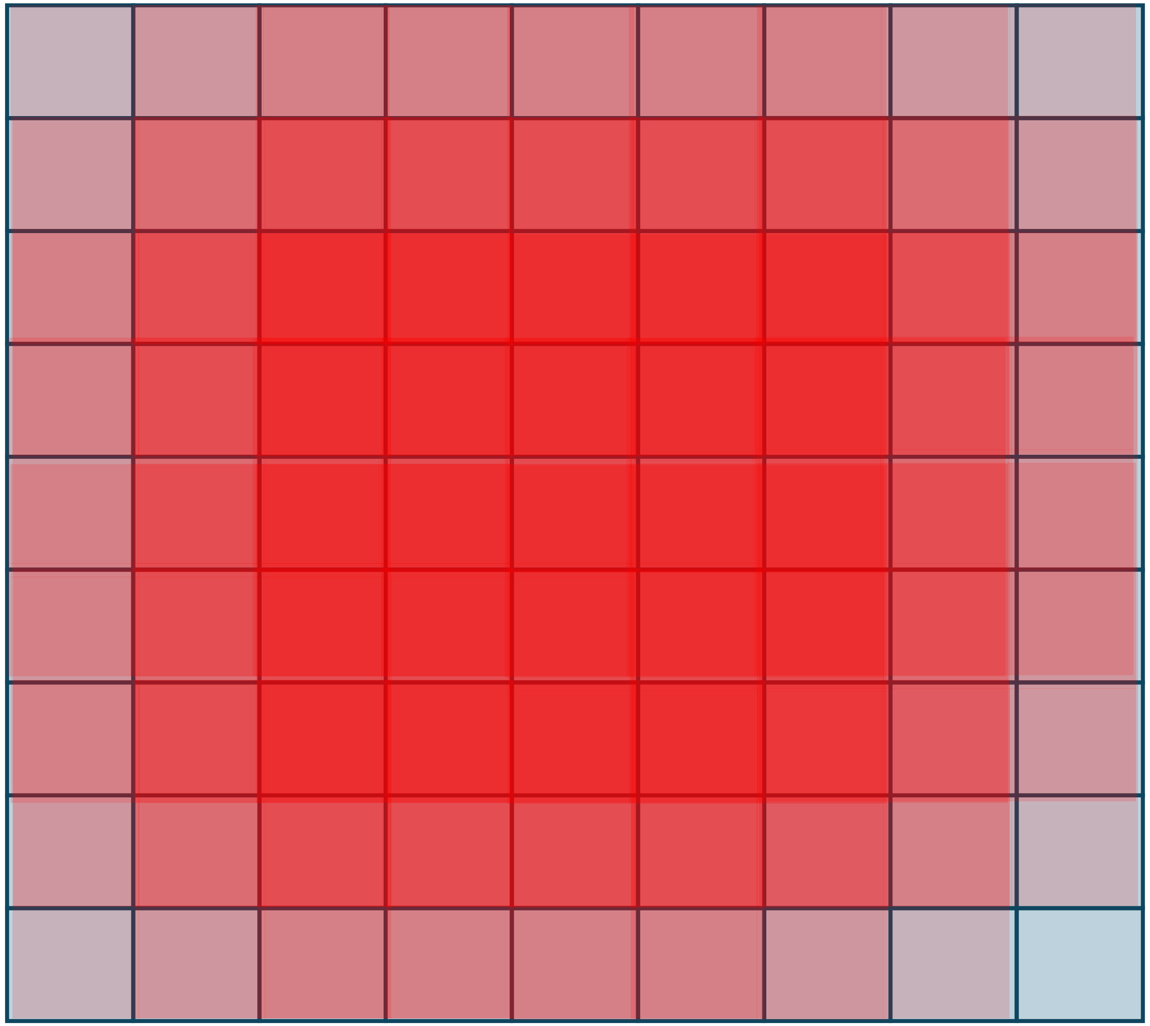

$* \boldsymbol{x} * \boldsymbol{x}$

Q: What is $x * x * x * x * x * x \ldots$

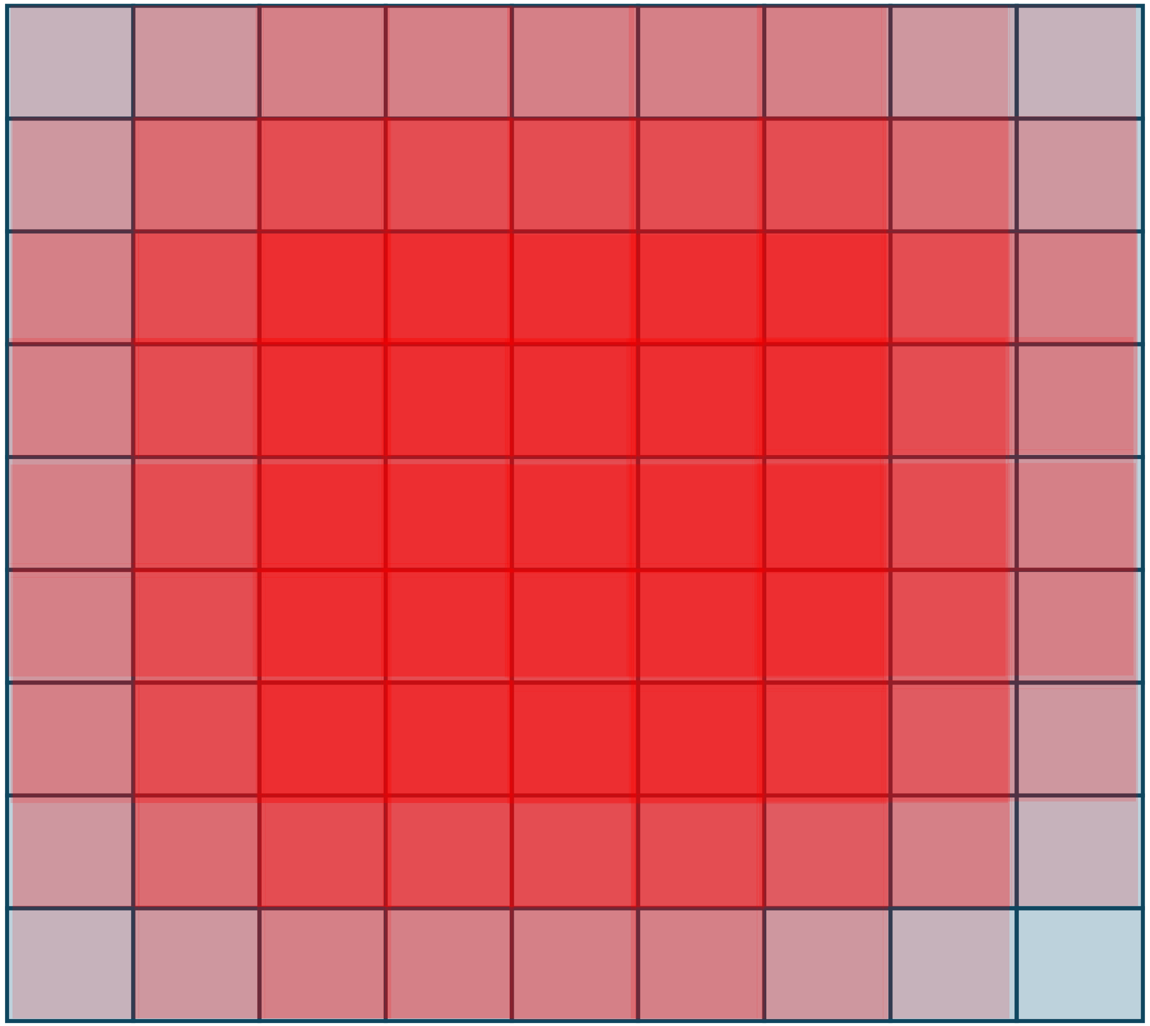

$* \boldsymbol{x} * \boldsymbol{x} * \boldsymbol{x}$

Q: What is $x * x * x * x * x * x \ldots$

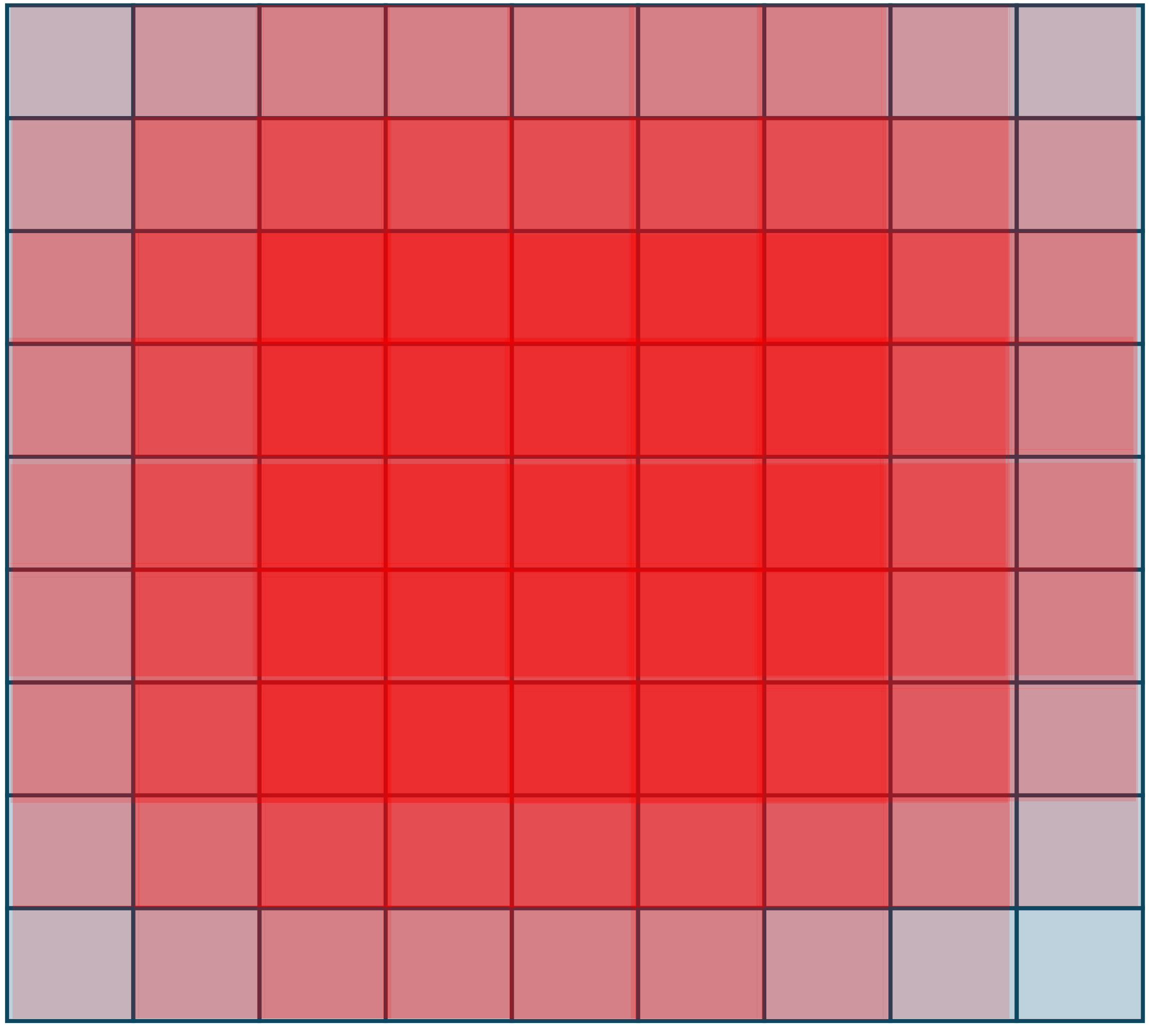

$* \boldsymbol{x} * \boldsymbol{x} * \boldsymbol{x} * \boldsymbol{x}$

Q: What is $x * x * x * x * x * x \ldots$

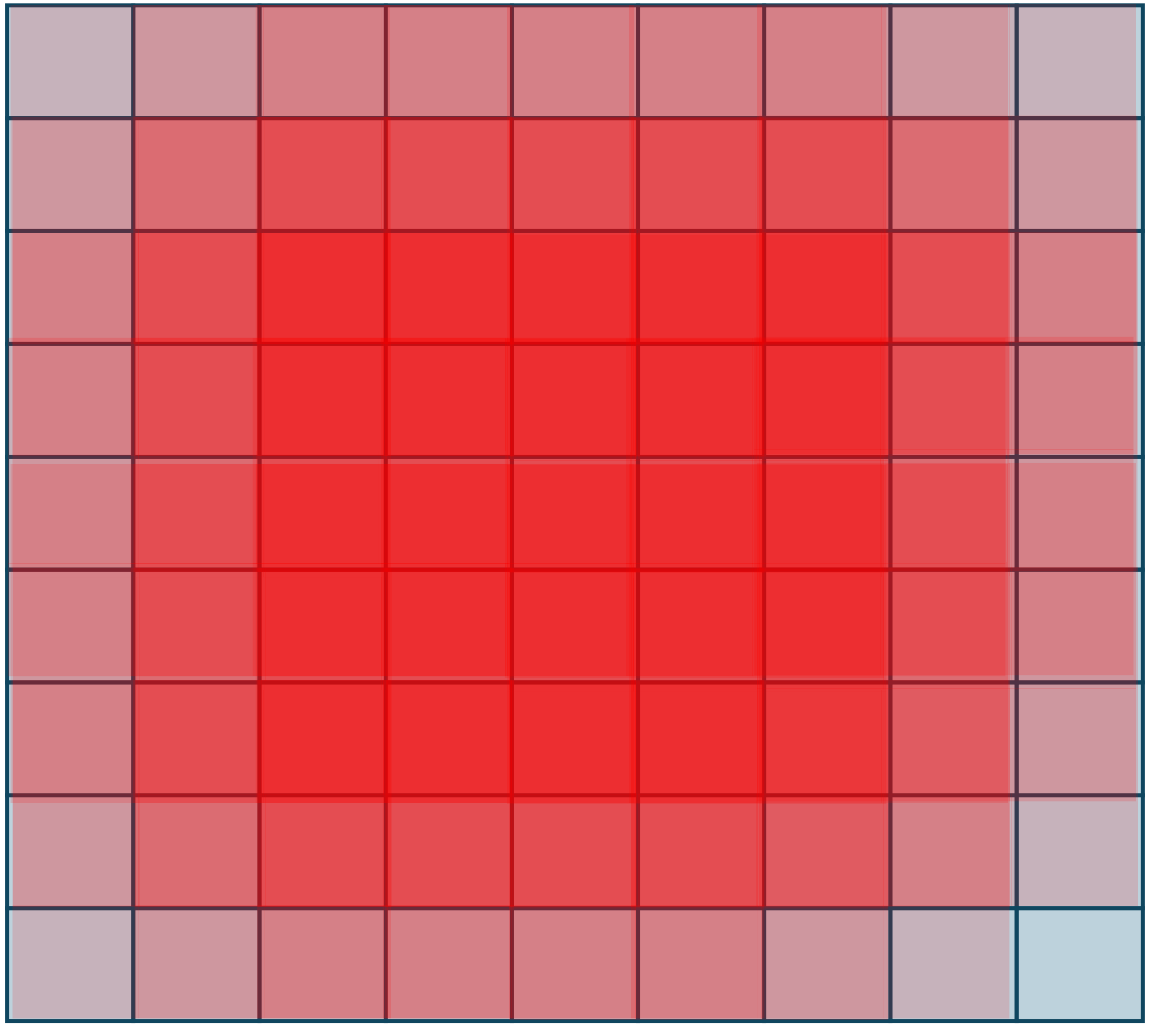

$* \boldsymbol{x} * \boldsymbol{x} * \boldsymbol{x} * \boldsymbol{x}$

Proof : Think CLT

Q: Given that all patches have the same norm (not realistic), where will convolution get the highest value?

Q: Given that all patches have the same norm (not realistic), where will convolution get the highest value?

A: Now it's that Schwartz! (and Cauchy): $\left|\langle f, g\rangle\right|^2 \leq \|f\| \bullet \|g\|,$ max:

$f = g$

Q: Given that all patches have the same norm (not realistic), where will convolution get the highest value?

A: Now it's that Schwartz! (and Cauchy): $\left|\langle f, g\rangle\right|^2 \leq \|f\| \bullet \|g\|,$ max:

$f = g$

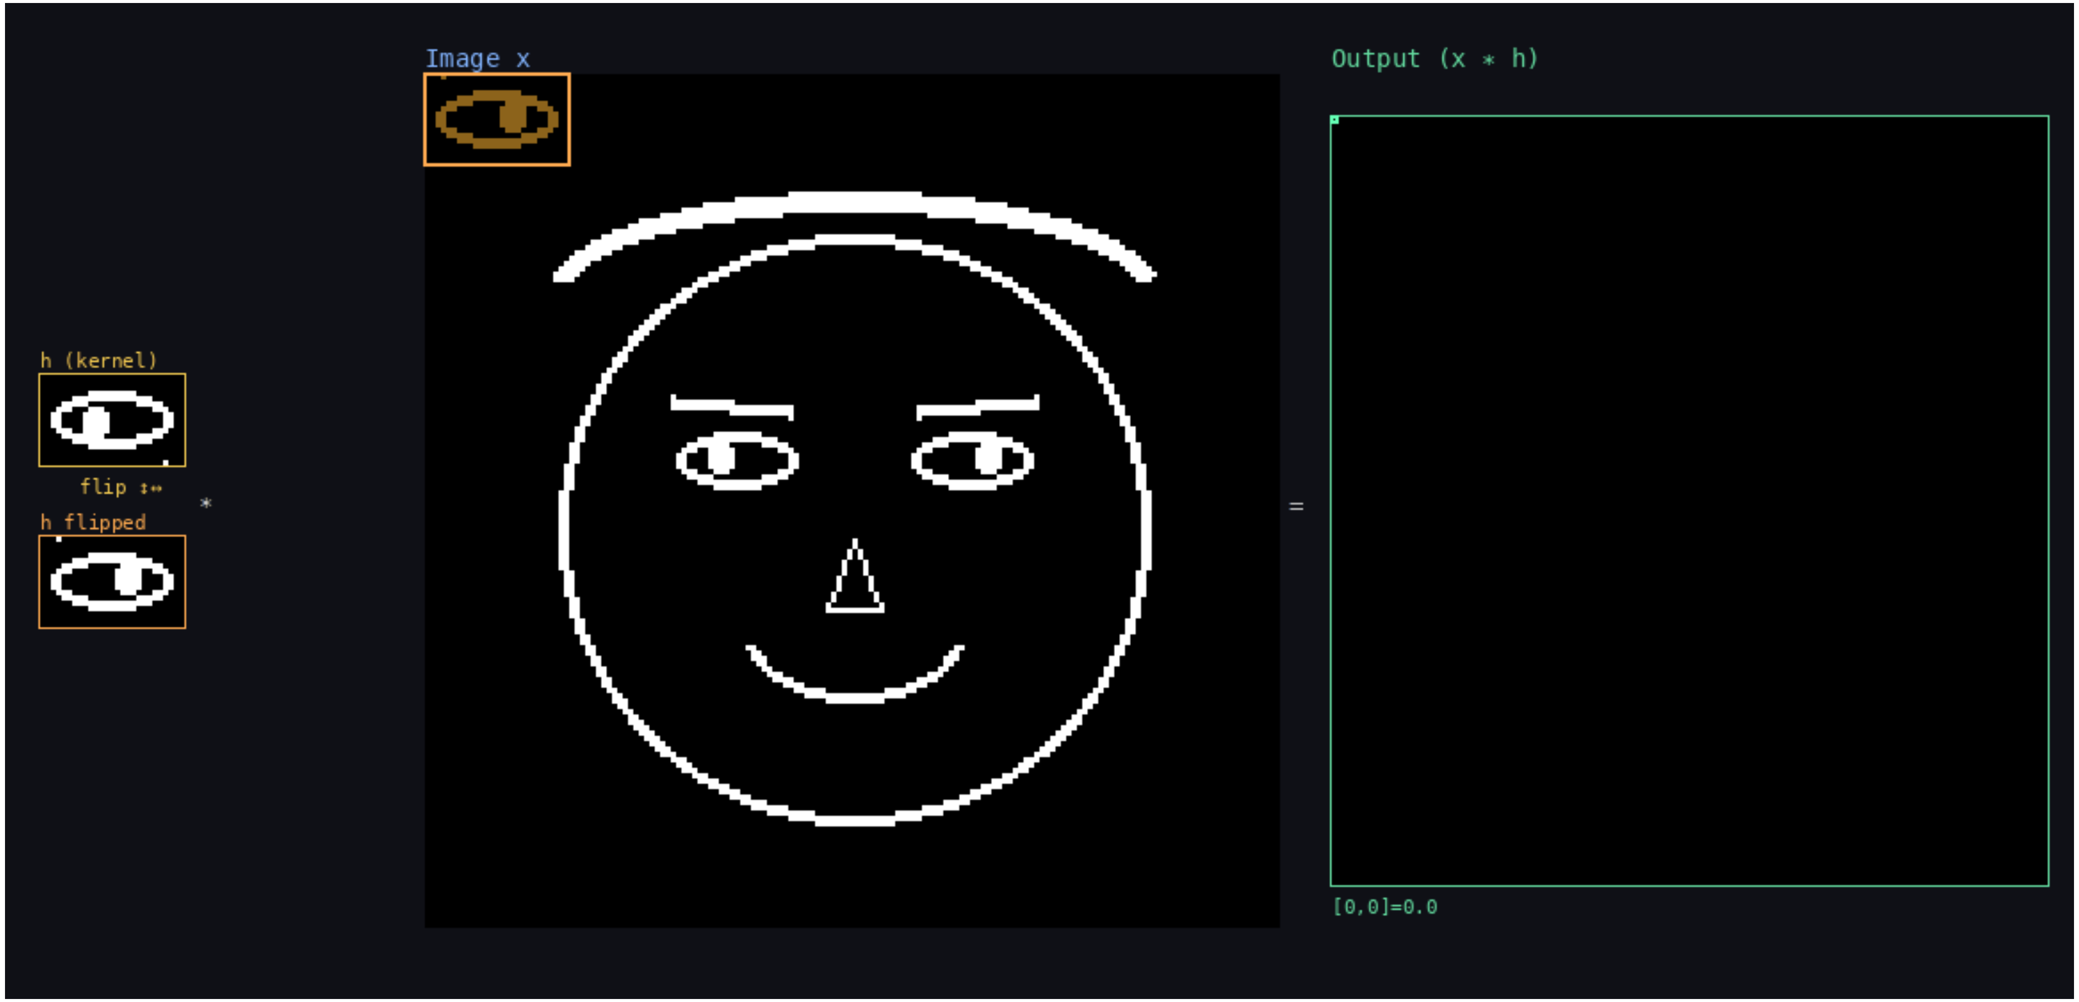

Idea taken from https://compneuro.neuromatch.io/tutorials/W1D5_DeepLearning/student/W1D5_Tutorial2.html

Common Convolution Filters

Smoothing

Box Blur (3×3)

1	1	1
1	1	1
1	1	1

× 1⁄9

Gaussian Blur (3×3)

1	2	1
2	4	2
1	2	1

× 1⁄16

Reduces noise
Averages neighbors

Basic Derivative

∂/∂x (horizontal)

−1	1

∂/∂y (vertical)

−1
1

Simplest finite difference
Sensitive to noise

Smoothed Derivative

Sobel X

−1	0	1
−2	0	2
−1	0	1

Sobel Y

−1	−2	−1
0	0	0
1	2	1

More robust gradient
Smoothing + derivative in one kernel

Basic derivative is noisy → Sobel adds Gaussian smoothing for a more stable gradient

red = negative gray = zero green = weighted

More Filters

Laplacian

$\partial^2/\partial x^2 + \partial^2/\partial y^2$

0	1	0
1	−4	1
0	1	0

2nd derivative → edges at zero-crossings
Isotropic: same in all directions

Sharpening

identity – Laplacian

0	−1	0
−1	5	−1
0	−1	0

center = 1 + 4 = 5

Boosts edges in the original image
"Unsharp mask" idea

Gaussian 5×5

larger kernel = more blur

1	4	6	4	1
4	16	24	16	4
6	24	36	24	6
4	16	24	16	4
1	4	6	4	1

$\times \frac{1}{256}$

Kernel size controls blur scale
Foundation of scale-space theory

Sharpening = original + edges → $I - \alpha\nabla^2 I$ (Laplacian sharpening)

Demo:

Point Spread Function

Func that describes the pattern of zero sized point of light on the optical system.

Point Spread Function

Func that describes the pattern of zero sized point of light on the optical system.

Q: But didn't we see that a point maps to a point?

Point Spread Function

Func that describes the pattern of zero sized point of light on the optical system.

Q: But didn’t we see that a point maps to a point?

A: Small angle approximation!

Point Spread Function

Func that describes the pattern of zero sized point of light on the optical system.

Q: But didn’t we see that a point maps to a point?

A: Small angle approximation!

Q: What does it do to the entire picture?

Point Spread Function

Func that describes the pattern of zero sized point of light on the optical system.

Q: But didn’t we see that a point maps to a point?

A: Small angle approximation!

Q: What does it do to the entire picture?

A: What the hell were we discussing during the last 30 mins?!?!

Point Spread Function

Func that describes the pattern of zero sized point of light on the optical system.

Q: But didn't we see that a point maps to a point?

A: Small angle approximation!

Q: What does it do to the entire picture?

A: What the hell were we discussing during the last 30 mins?!?!

This is exactly the impulse response, so, convolution with the image

Point Spread Function

Func that describes the pattern of zero sized point of light on the optical system.

Q: But didn't we see that a point maps to a point?

A: Small angle approximation!

Q: What does it do to the entire picture?

A: What the hell were we discussing during the last 30 mins?!?!

This is exactly the impulse response, so, convolution with the image

Q: Implications?

Point Spread Function

Func that describes the pattern of zero sized point of light on the optical system.

Q: But didn't we see that a point maps to a point?

A: Small angle approximation!

Q: What does it do to the entire picture?

A: What the hell were we discussing during the last 30 mins?!?!

This is exactly the impulse response, so, convolution with the image

Q: Implications?

A: These are called "Geometric Aberrations" and they induce a blur effect.

Point Spread Function

Func that describes the pattern of zero sized point of light on the optical system.

Q: But didn't we see that a point maps to a point?

A: Small angle approximation!

Q: What does it do to the entire picture?

A: What the hell were we discussing during the last 30 mins?!?!

This is exactly the impulse response, so, convolution with the image

Q: Implications?

A: These are called "Geometric Aberrations" and they induce a blur effect.

Q: Is this the only reason for the blur?

Point Spread Function

Func that describes the pattern of zero sized point of light on the optical system.

Q: But didn’t we see that a point maps to a point?

A: Small angle approximation!

Q: What does it do to the entire picture?

A: What the hell were we discussing during the last 30 mins?!?!

This is exactly the impulse response, so, convolution with the image

Q: Implications?

A: These are called “Geometric Aberrations” and they induce a blur effect.

Q: Is this the only reason for the blur?

A: No. Also diffraction. Light is a wave, remember?

Point Spread Function

Func that describes the pat

Q: But didn't we see t

e approximation!

Q: What does it do to

A: Wha

e last 30 mins?!?!

This is

volution with the

image

Q: Implications?

A: These are called "

e a blur effect.

Q: Is this the only rea

A: No. Also diffraction. Light is a wave, Demo:

Pope Benedict
announcement

Slide credit:
Peyman Milanfar

Pope Benedict announcement

Slide credit:
Peyman Milanfar

Pope Benedict announcement

Pope Francis announcement

Slide credit:
Peyman Milanfar

Digital Image Pipeline

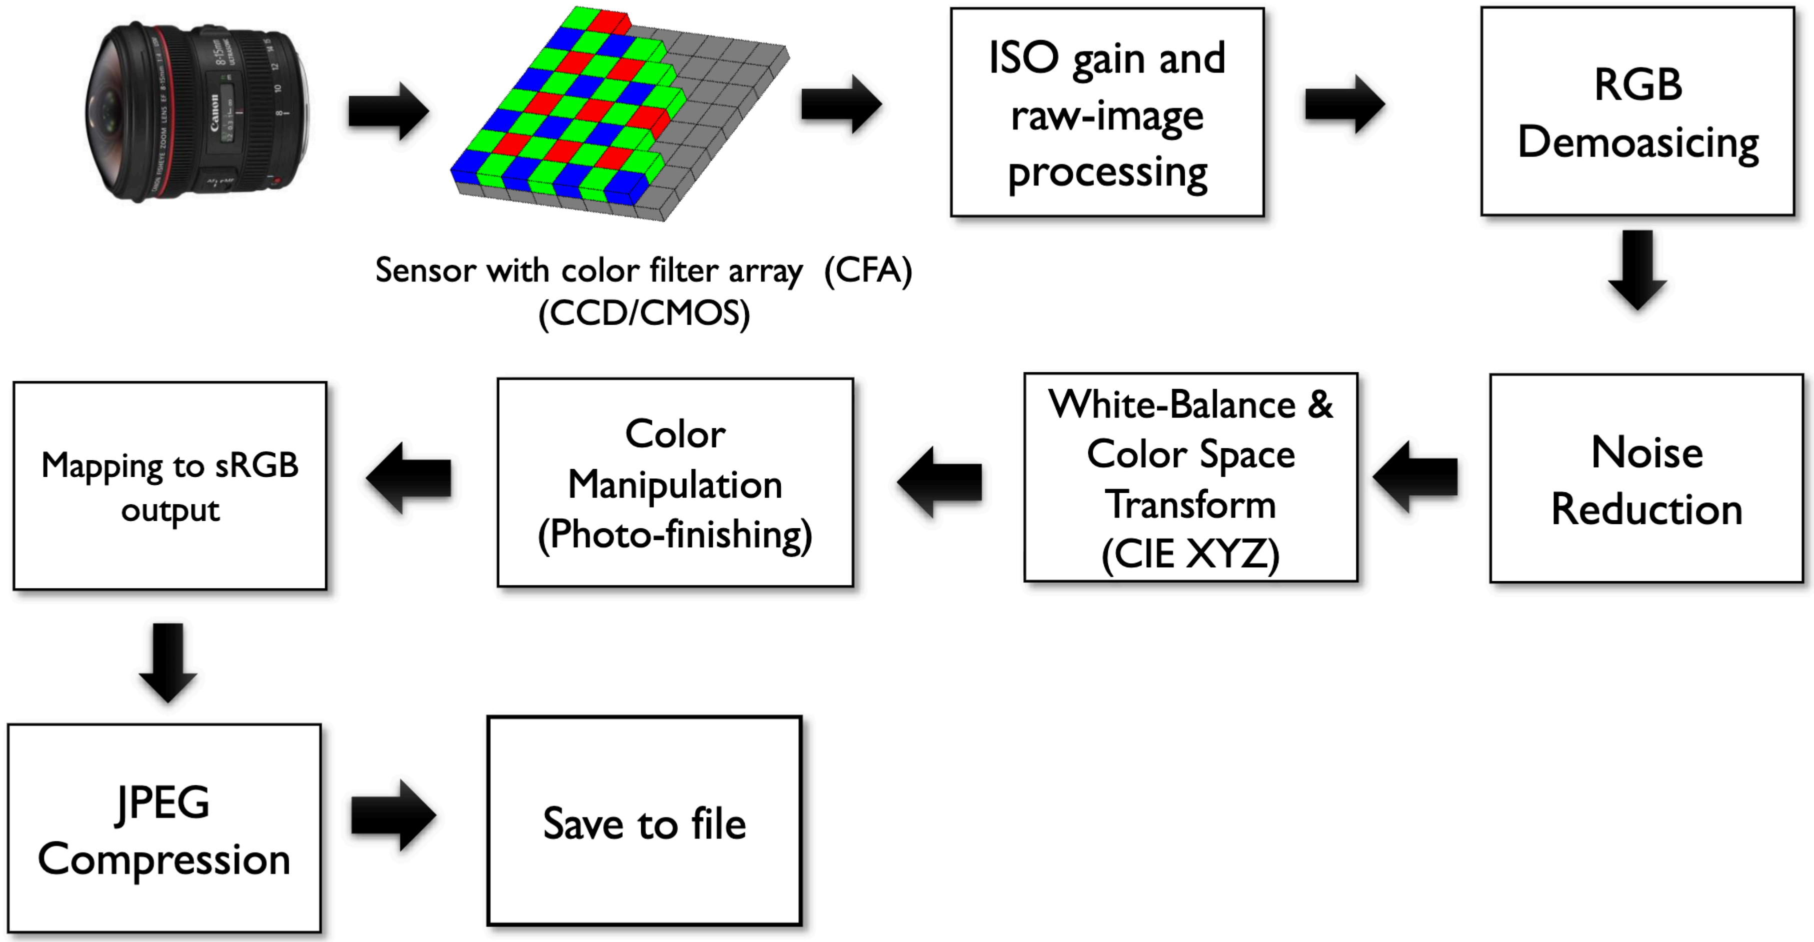

*** Credit: Some figures in this topic taken from Peyman Milanfar and from Michael S Bro**

The Bayer mosaic

How does light become voltage?

The Bayer mosaic

How does light become voltage?

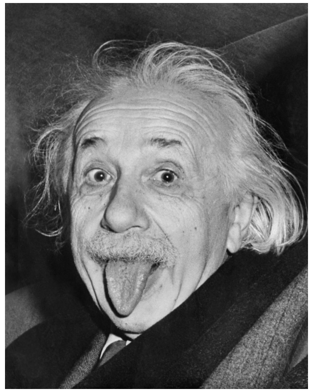

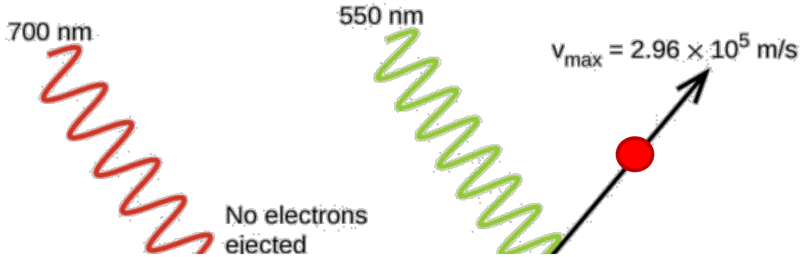

The Bayer mosaic

How does light become voltage?

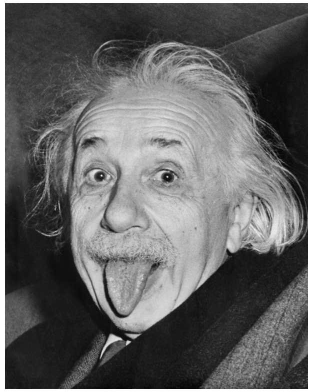

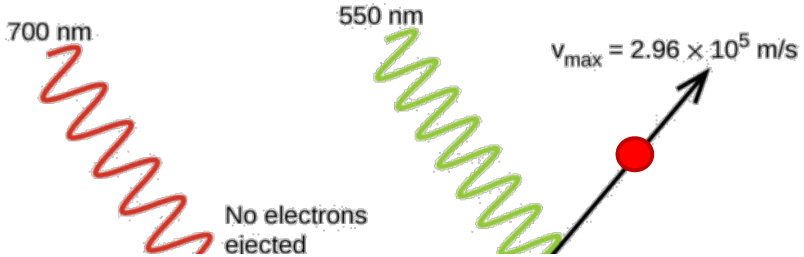

What about color?

The Bayer mosaic

How does light become voltage?

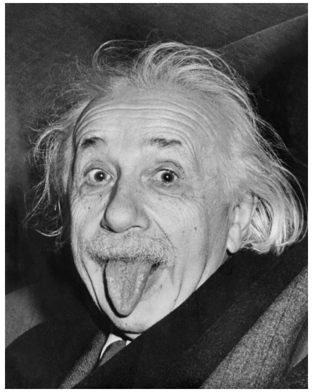

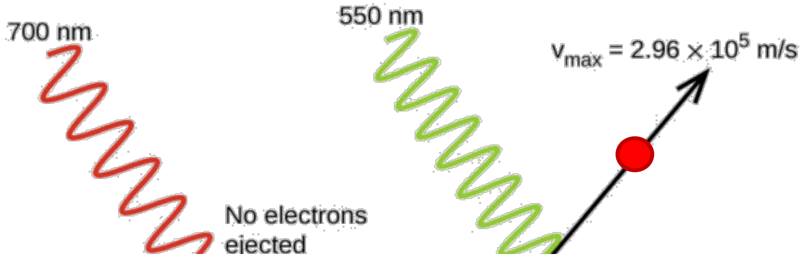

What about color?

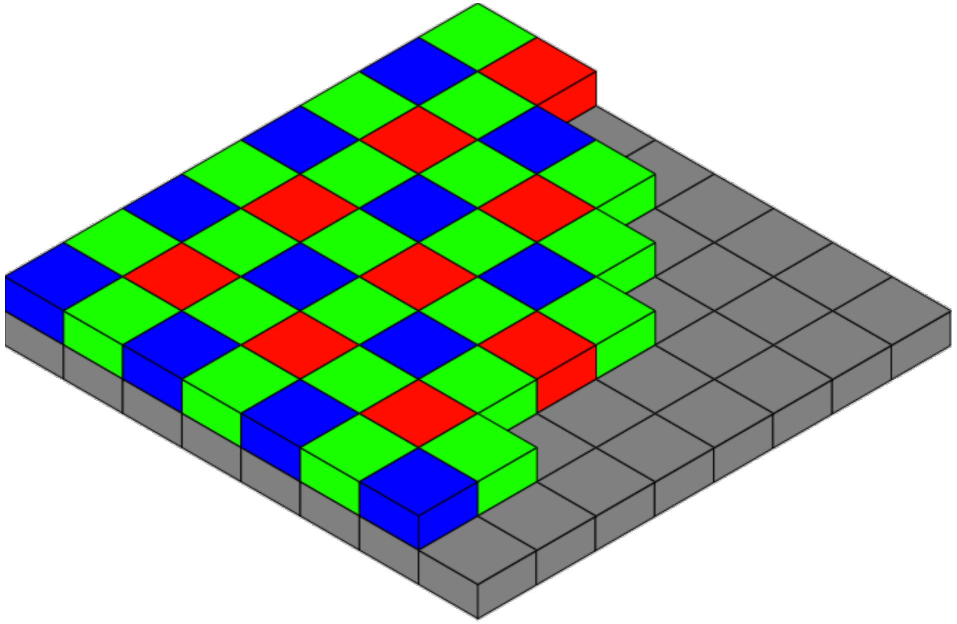

The Bayer mosaic

How does light become voltage?

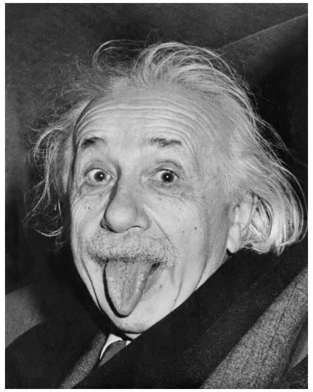

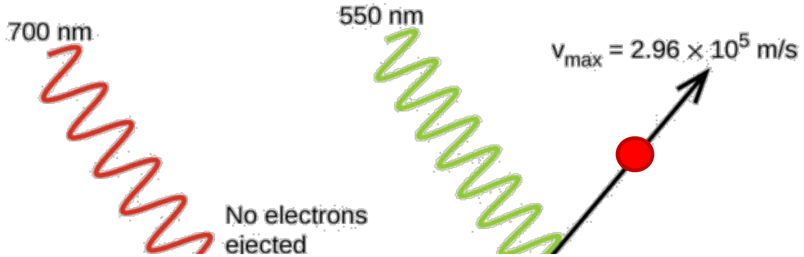

What about color?

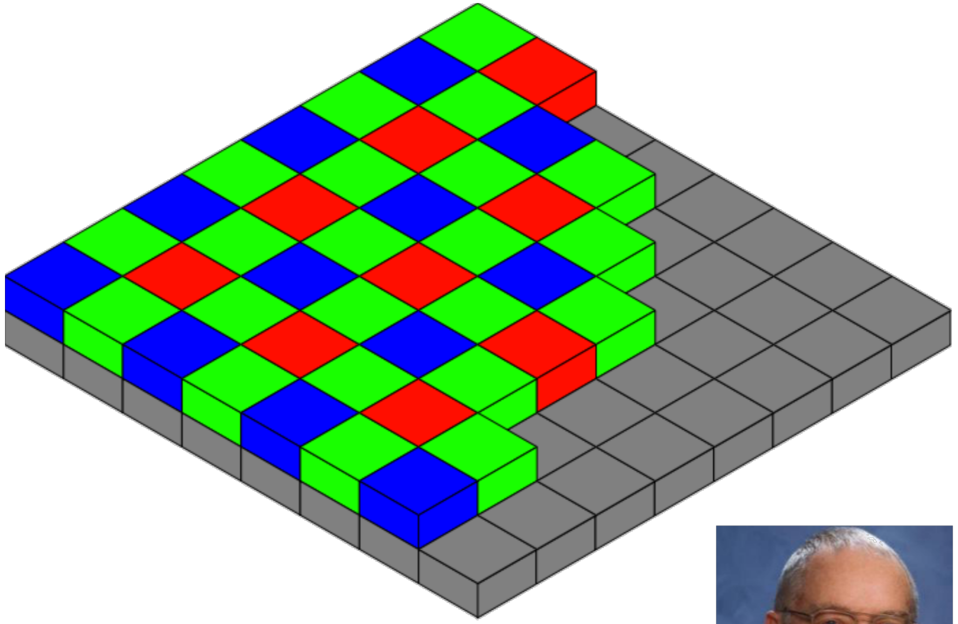

Bryce Bayer
1976

The Bayer mosaic

How does light become voltage?

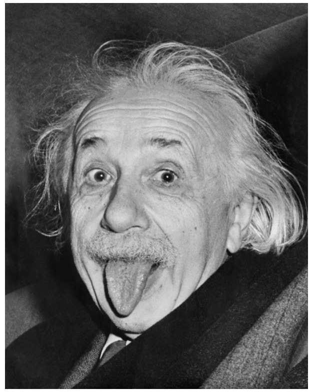

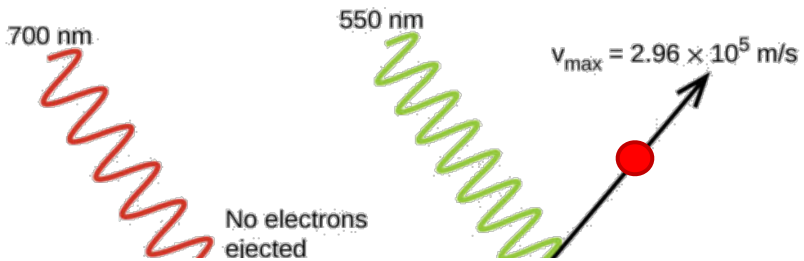

What about color?

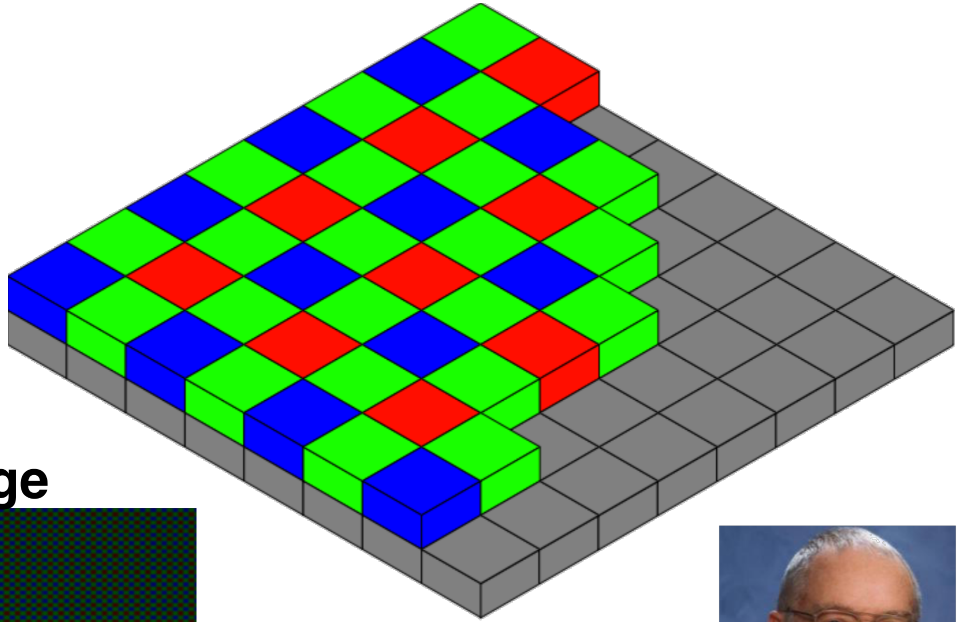

Bryce Bayer
1976

Raw image

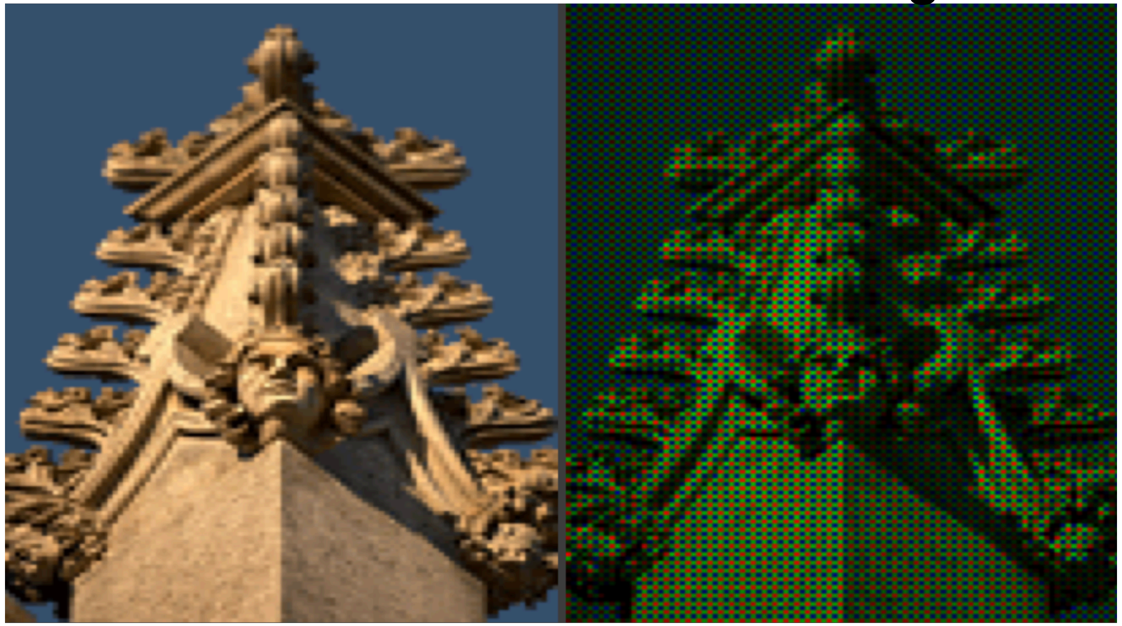

Your eyes do it too!

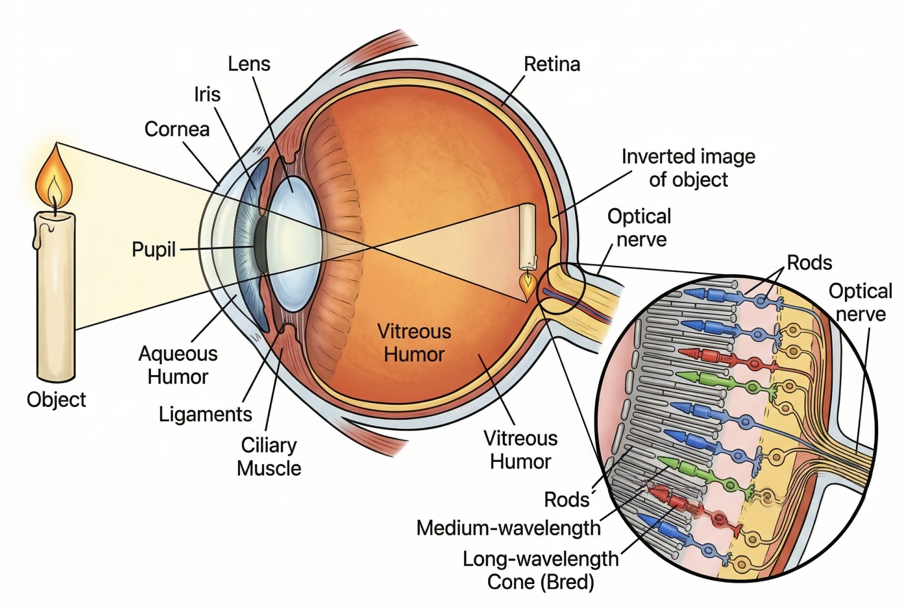

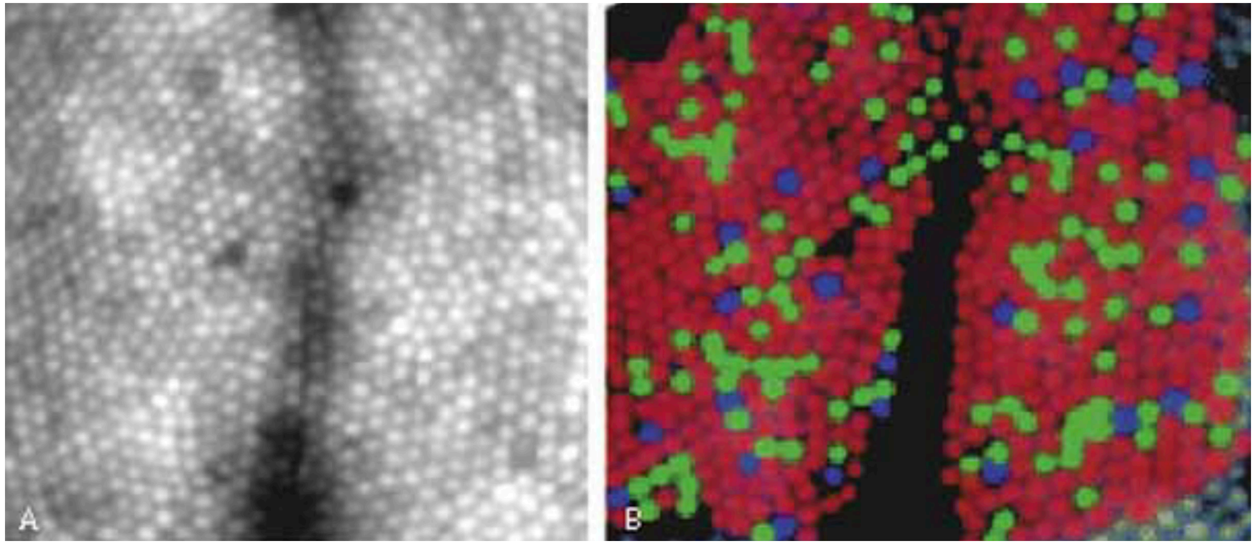

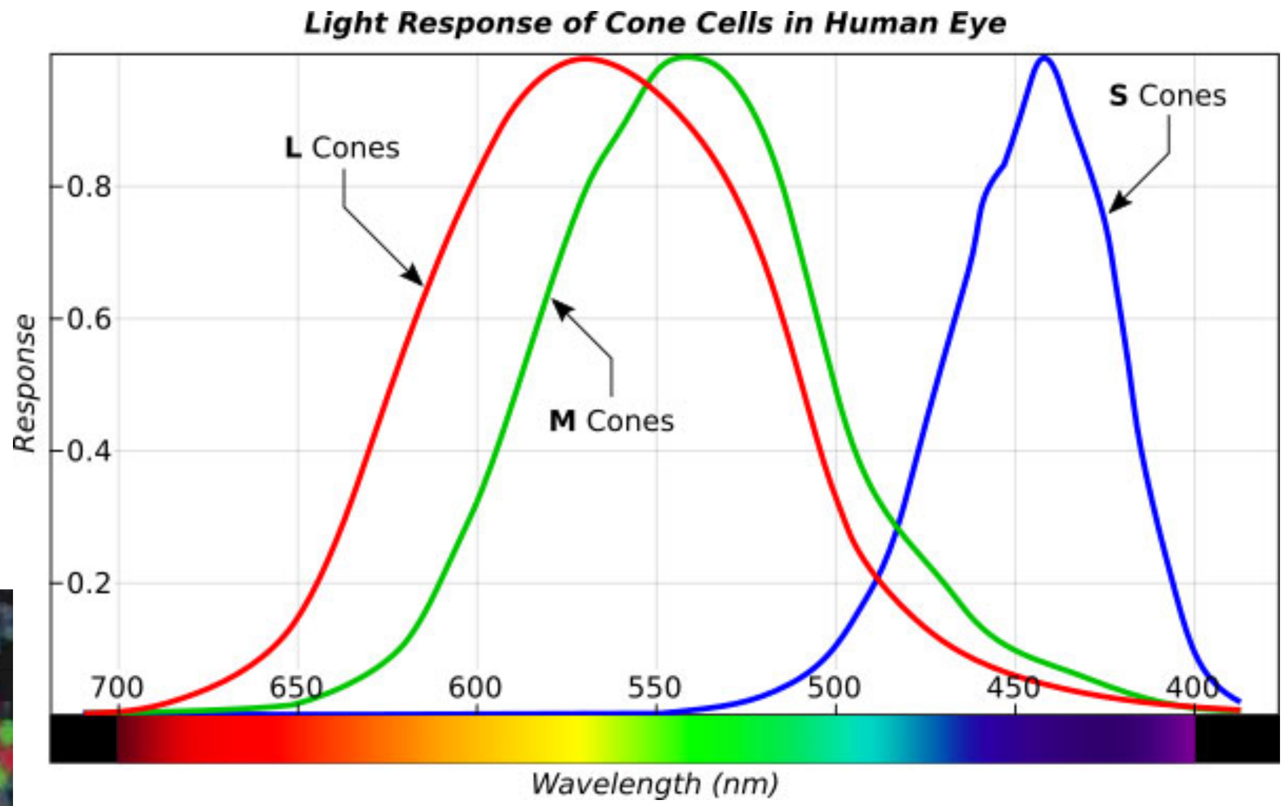

Non-unique

Digital Image Pipeline

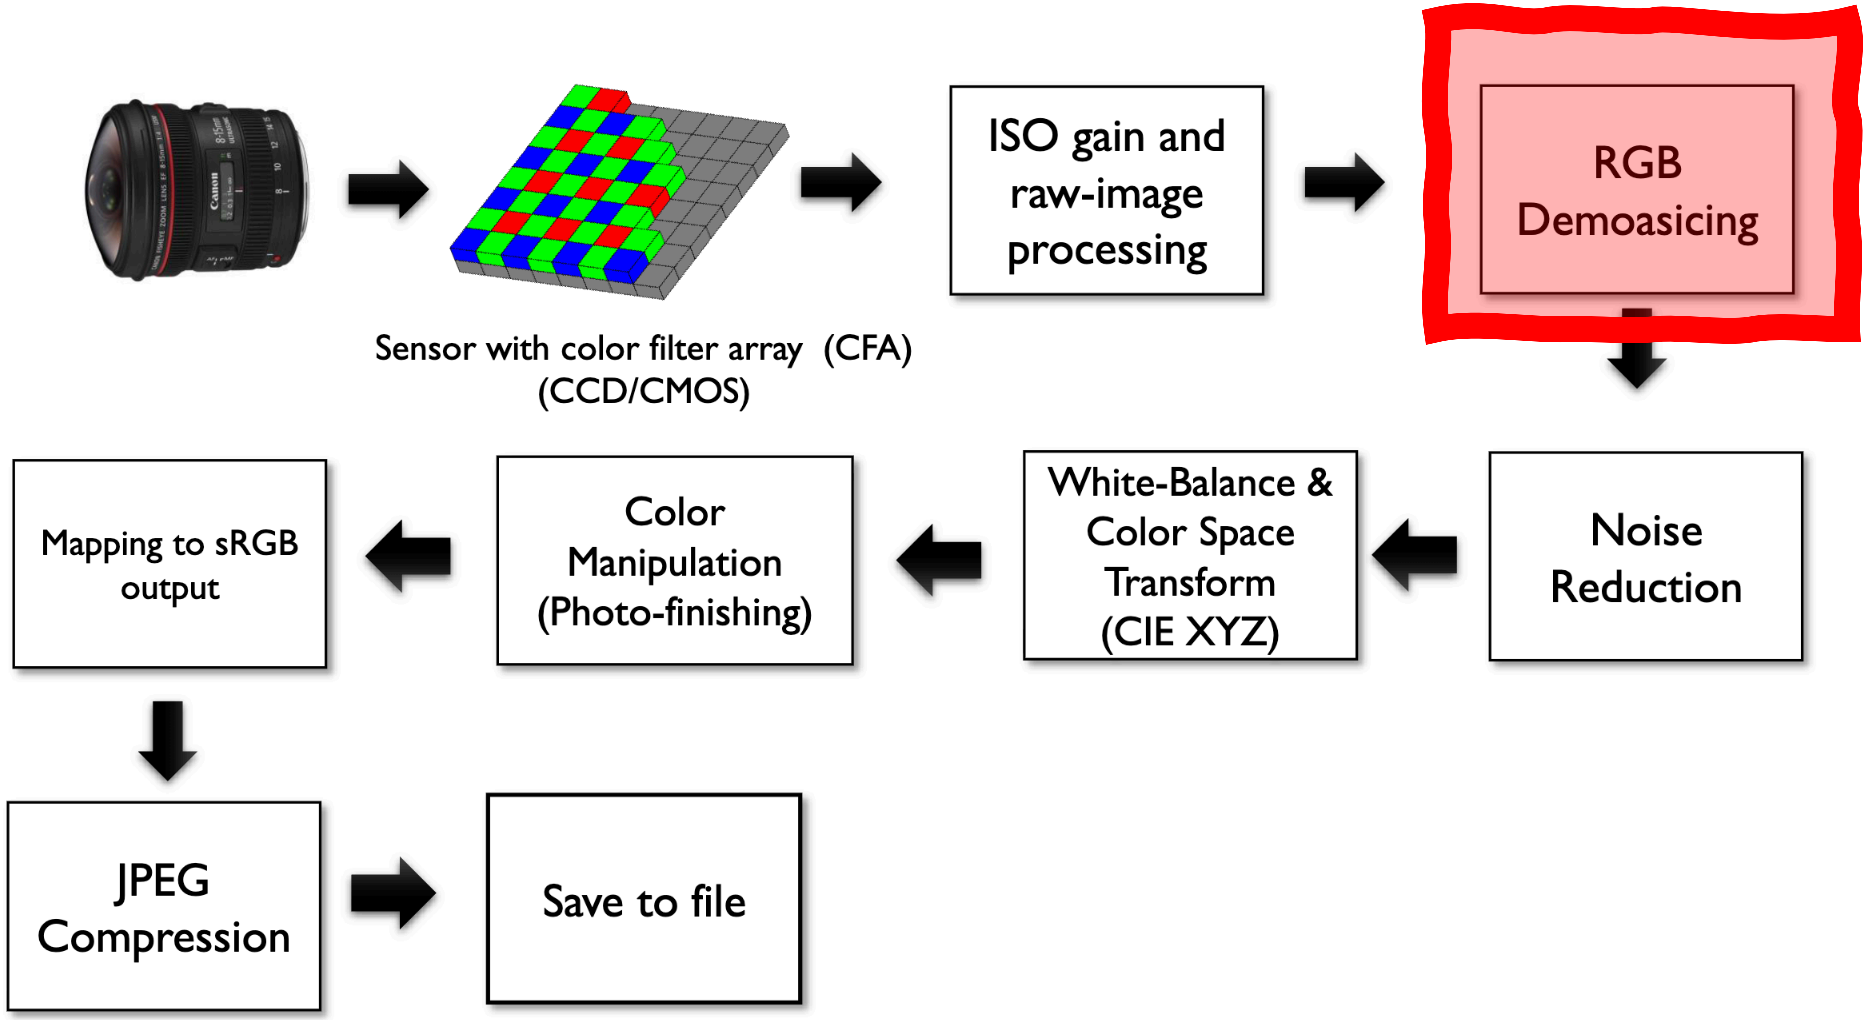

Demosaicing

Demosaicing

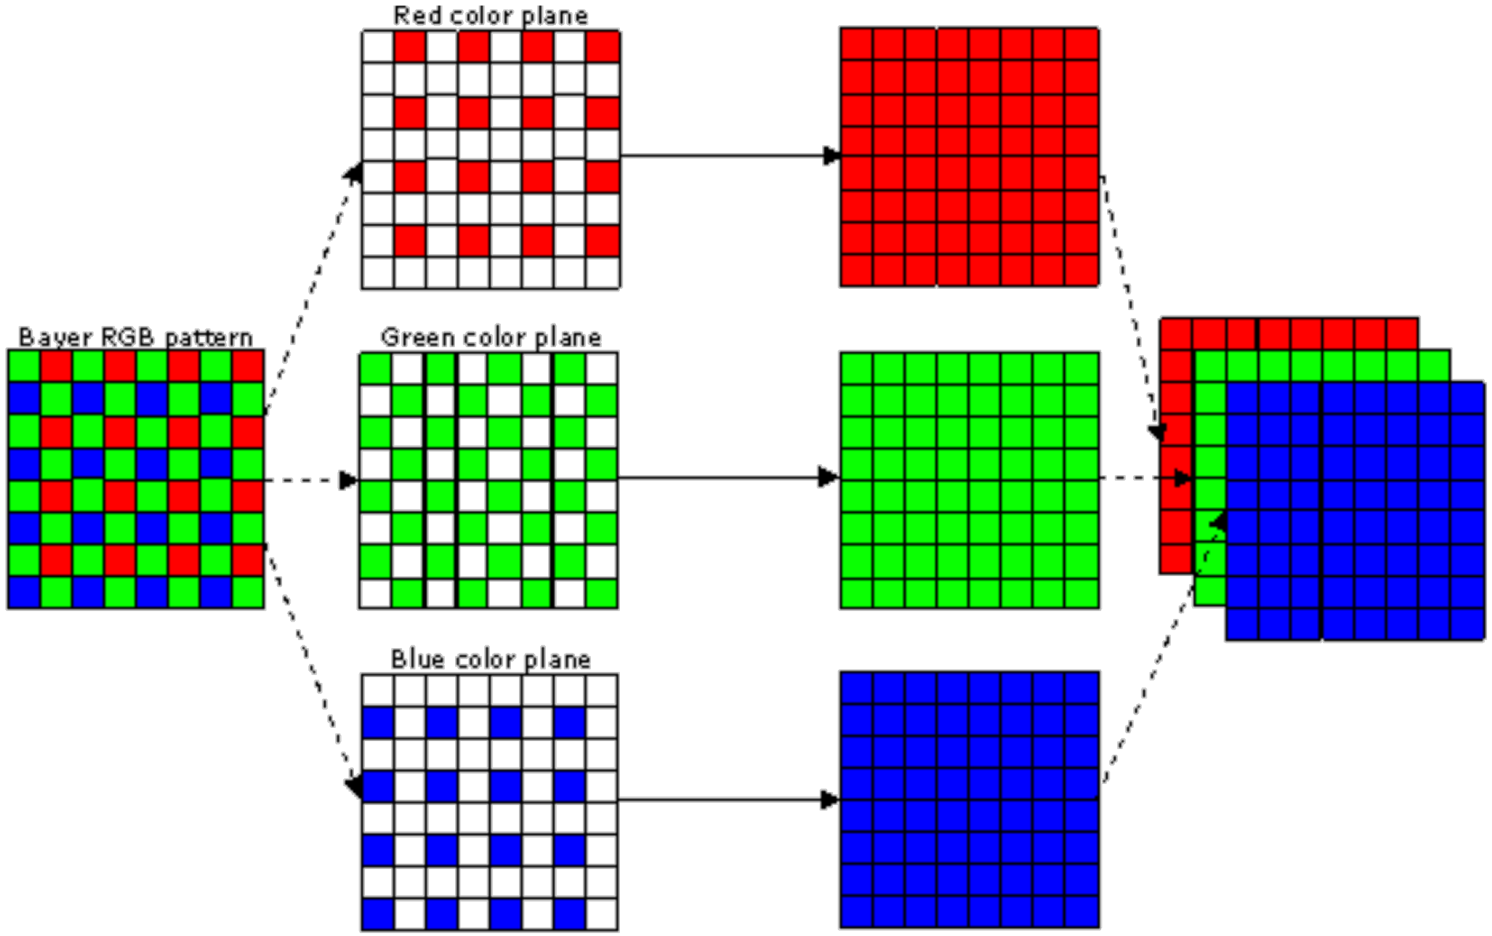

Demosaicing

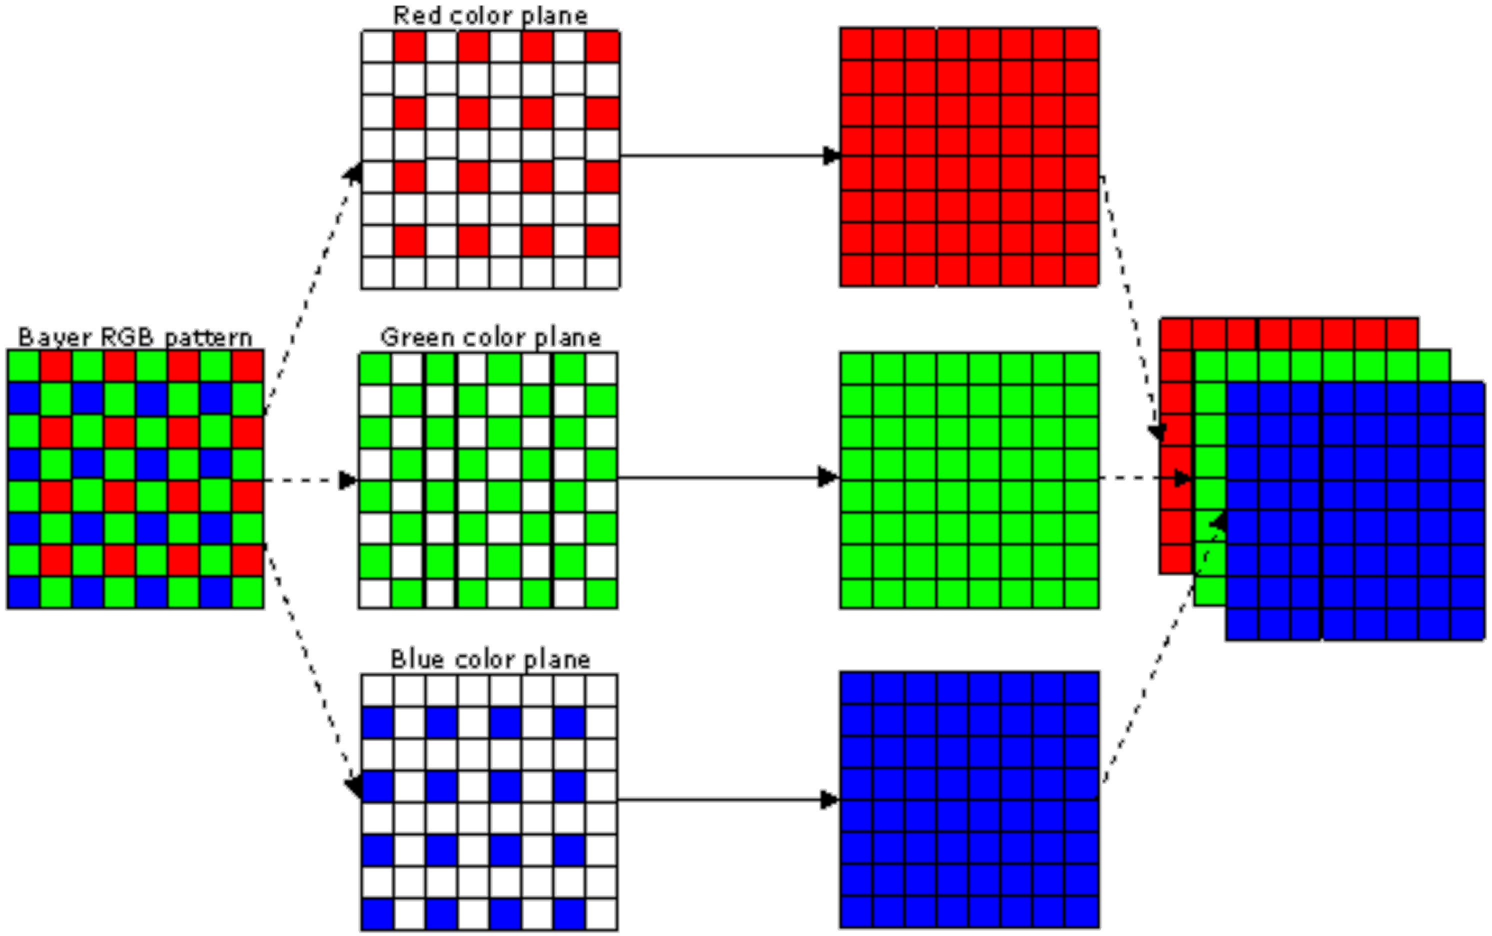

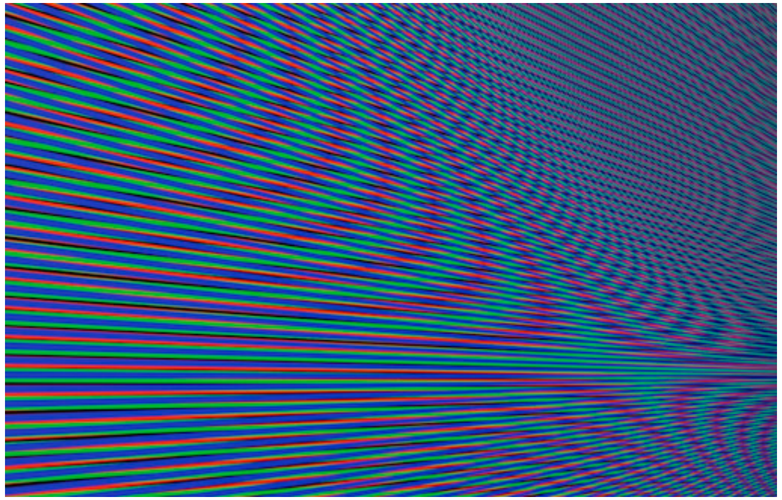

Demosaicing

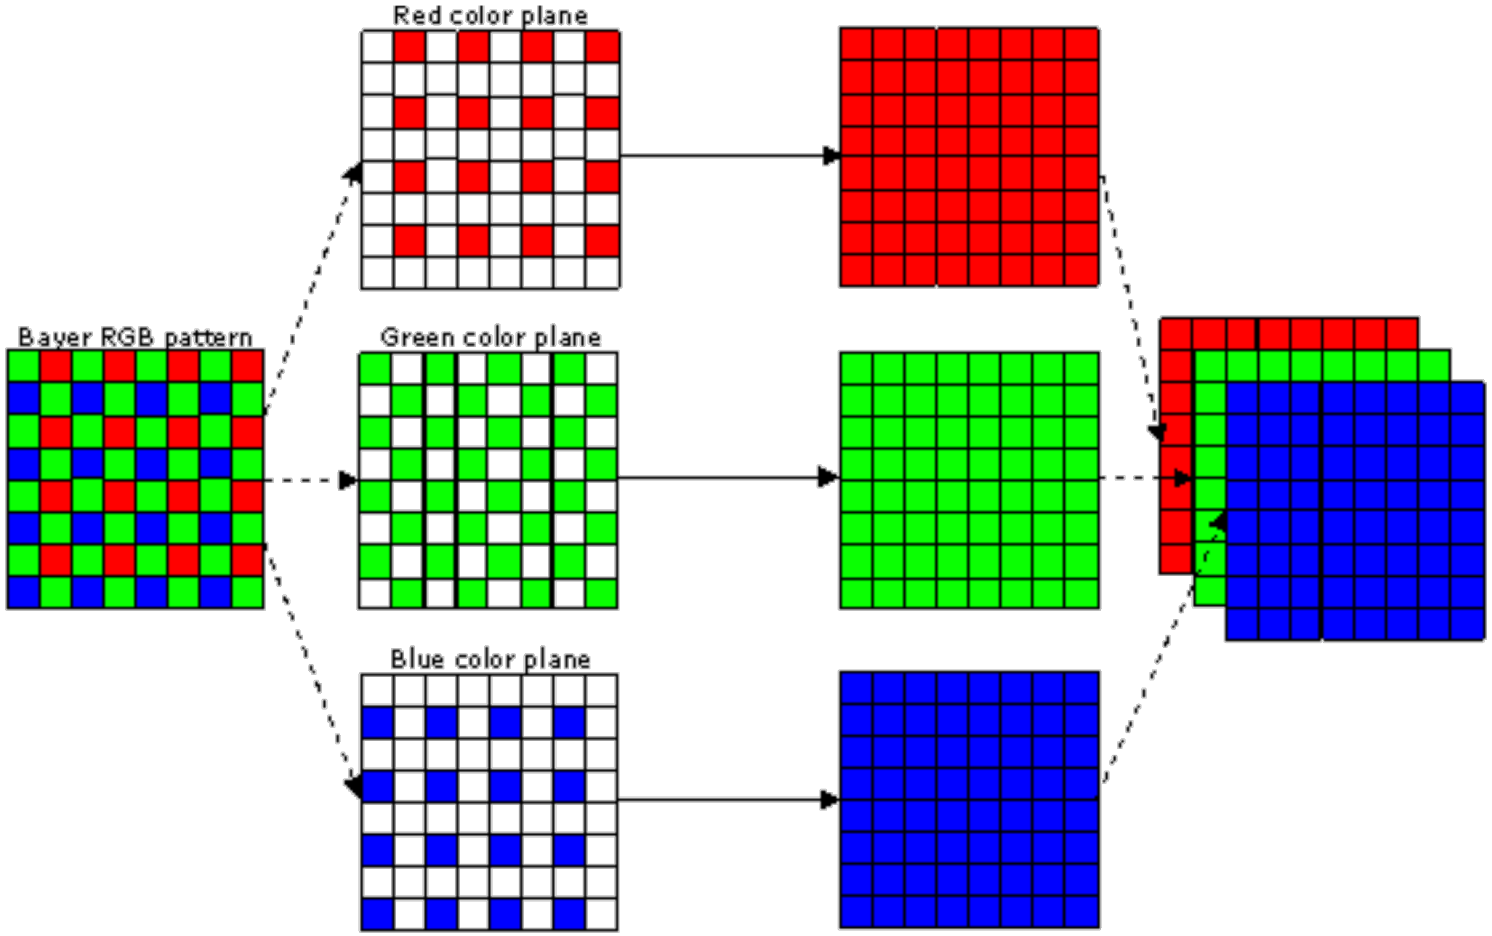

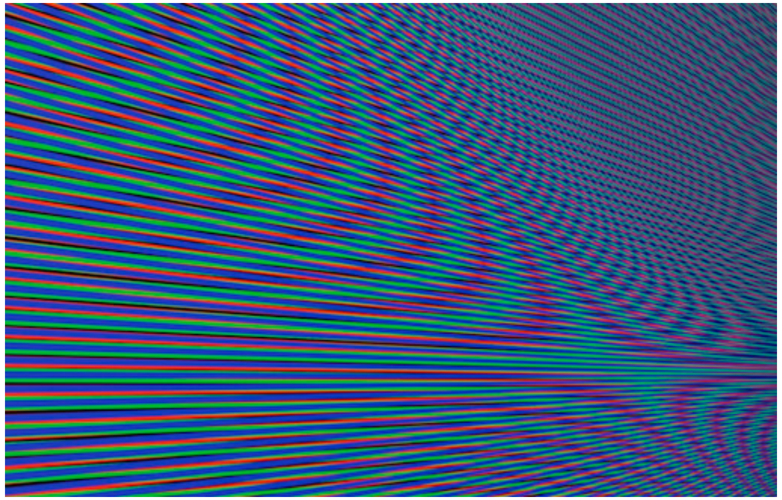

Methods:

Bilinear | Edge-aware | Freq. domain | Deep learning

Digital Image Pipeline

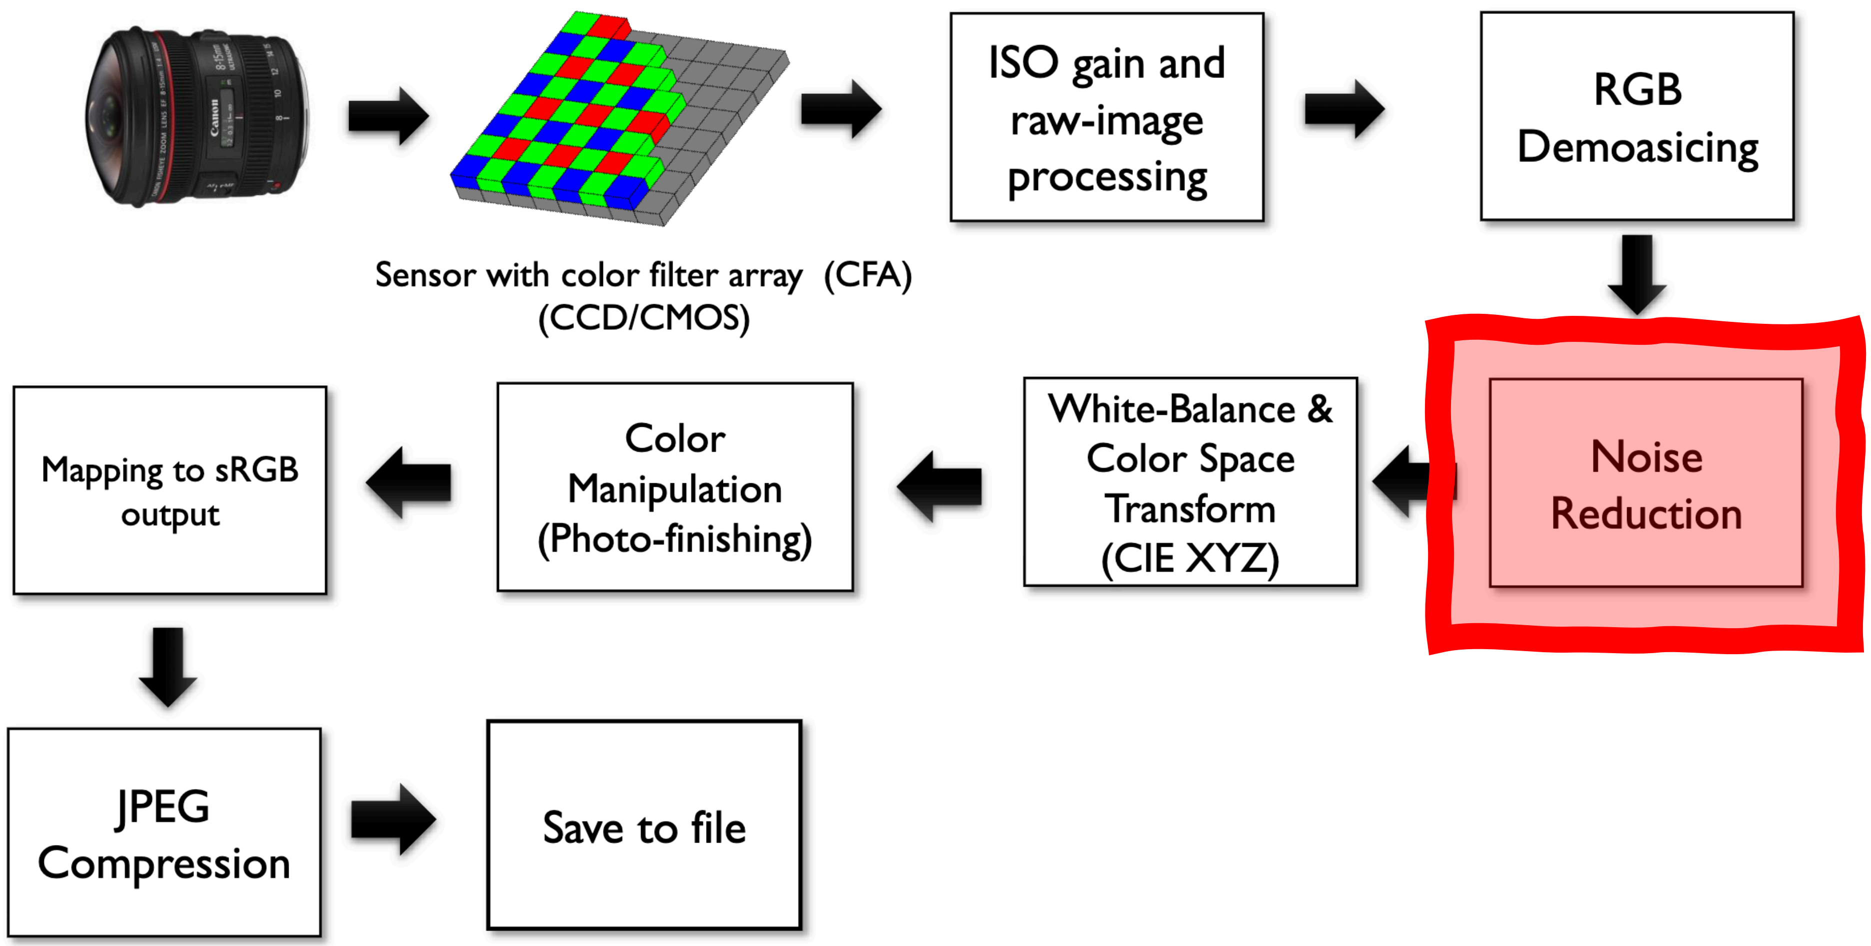

Denoising

Where does noise come from?

Denoising

Where does noise come from?

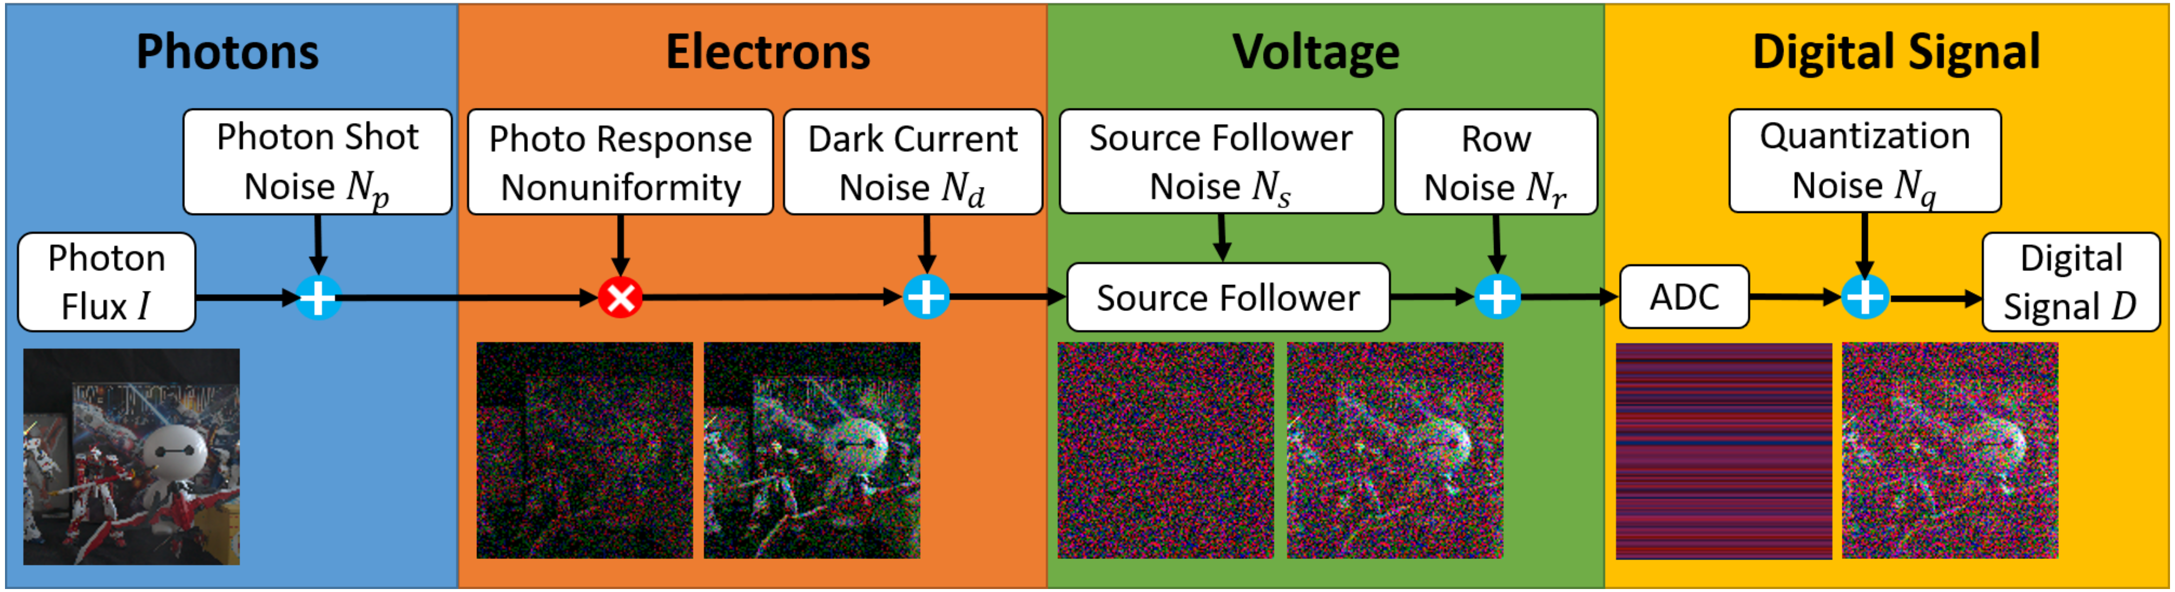

Denoising

Where does noise come from?

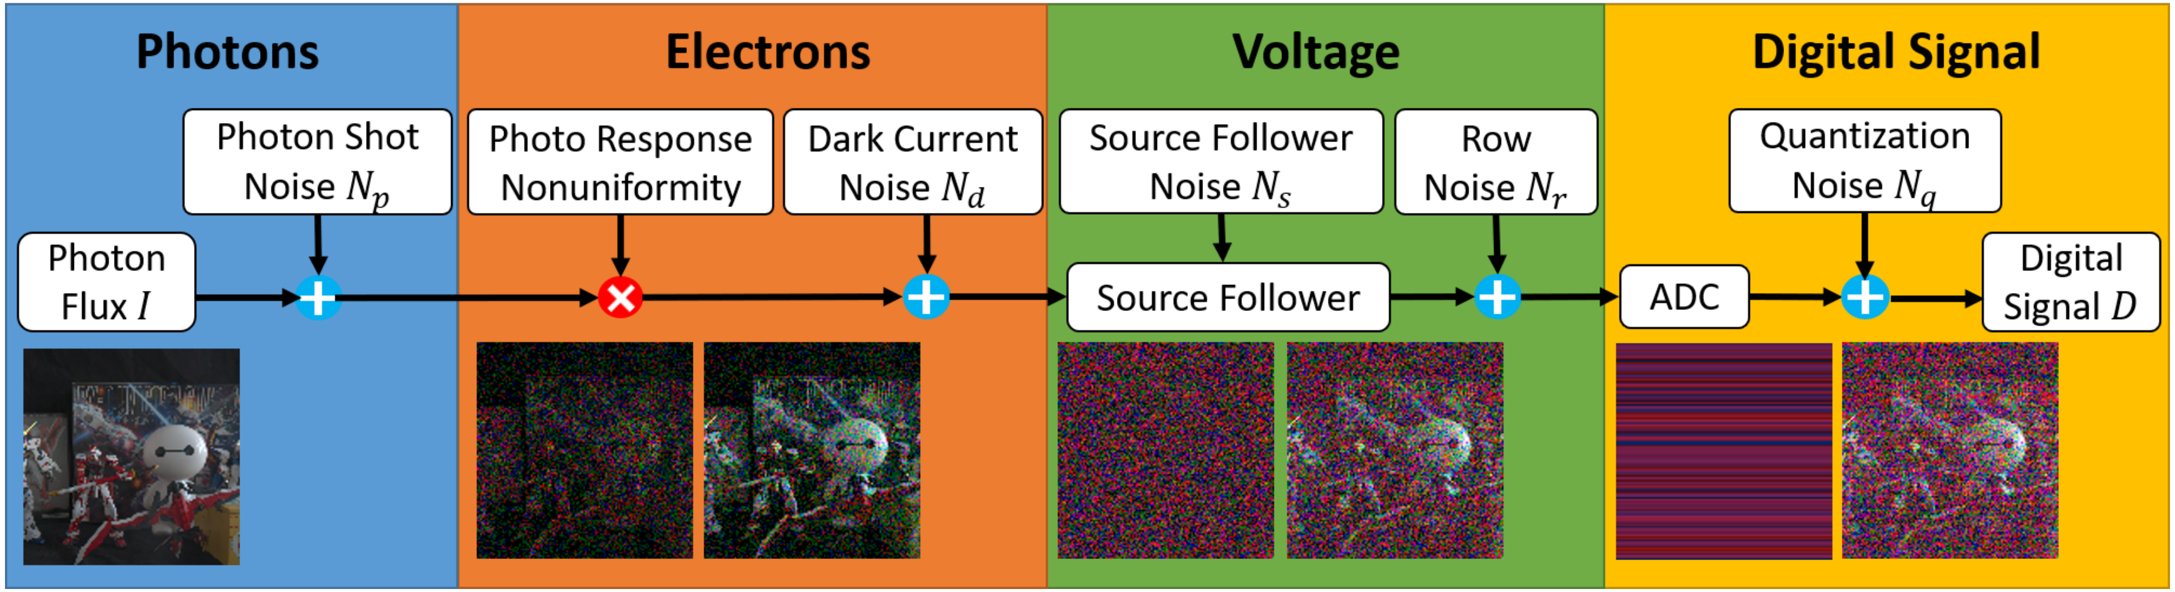

Naïve solution- Gaussian blur

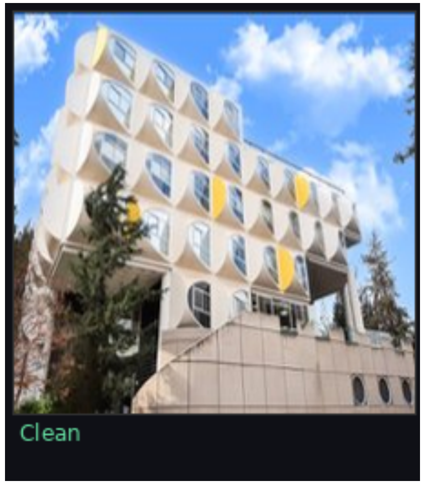

Denoising

Where does noise come from?

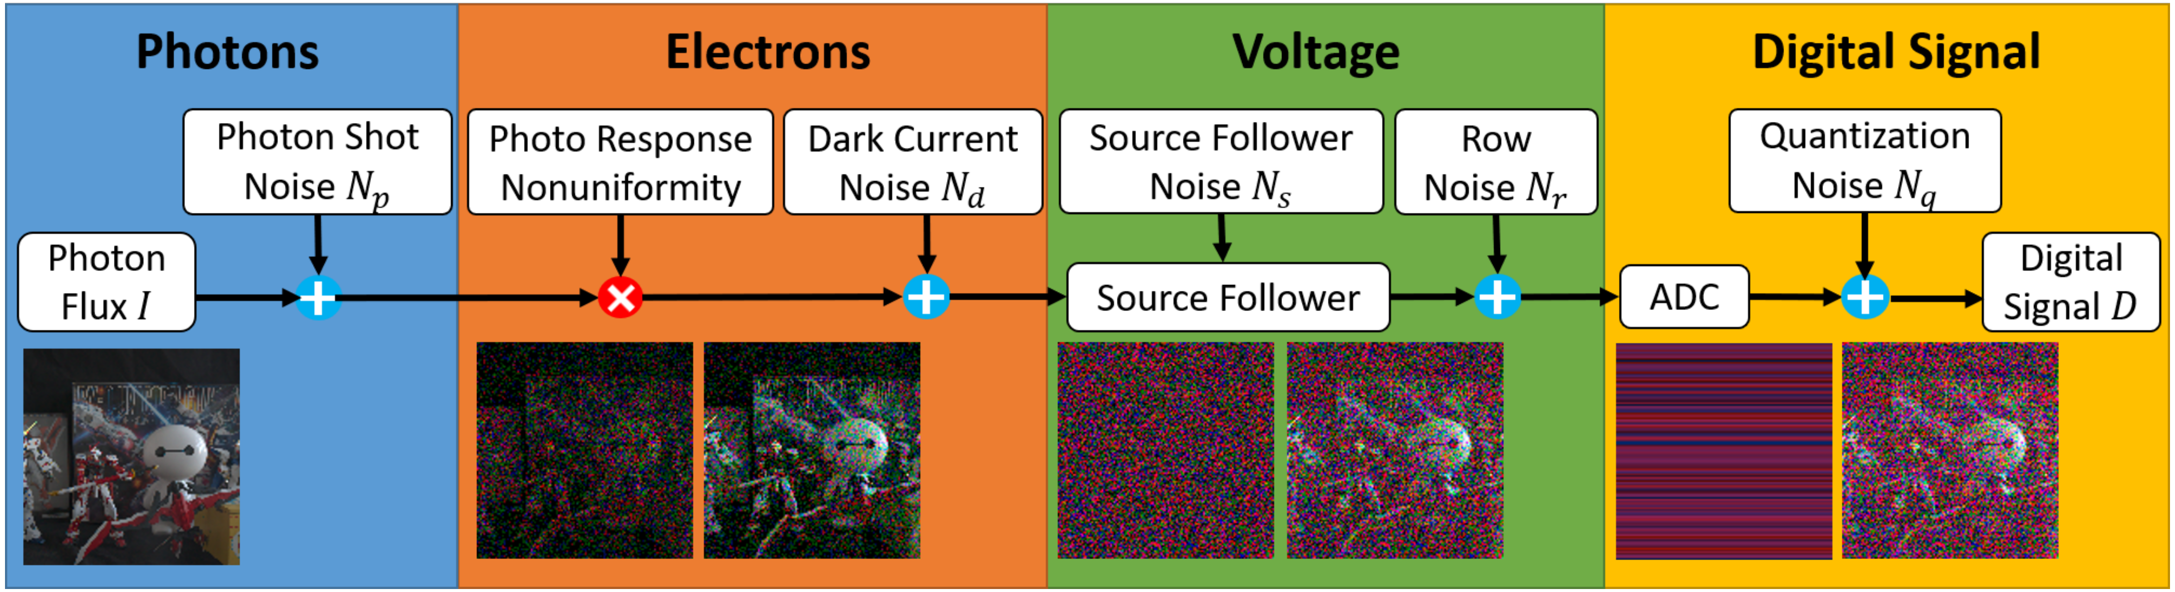

Naïve solution- Gaussian blur

Denoising

Where does noise come from?

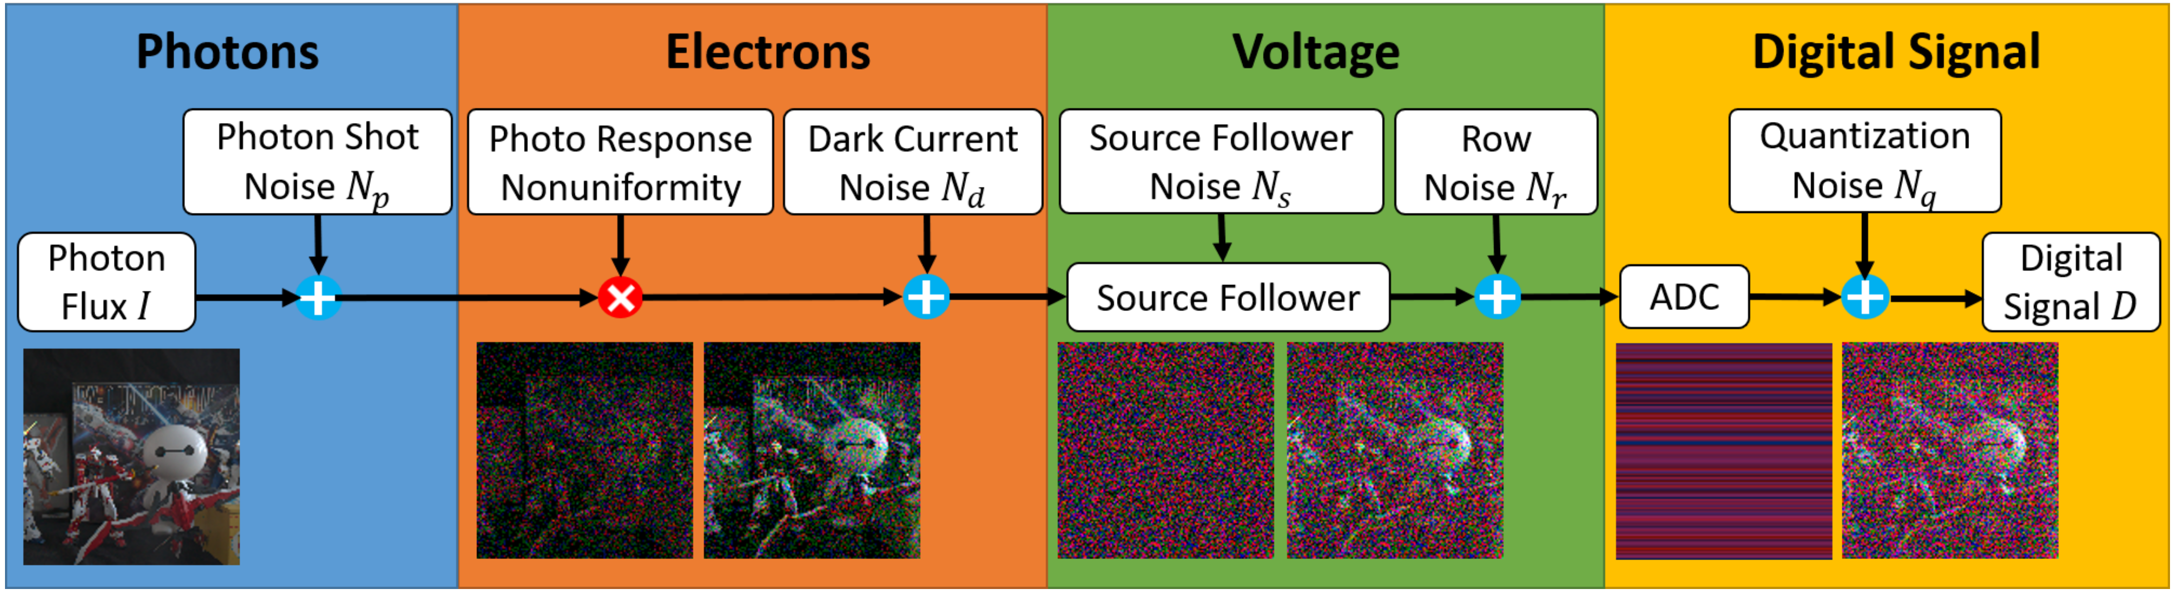

Naïve solution- Gaussian blur

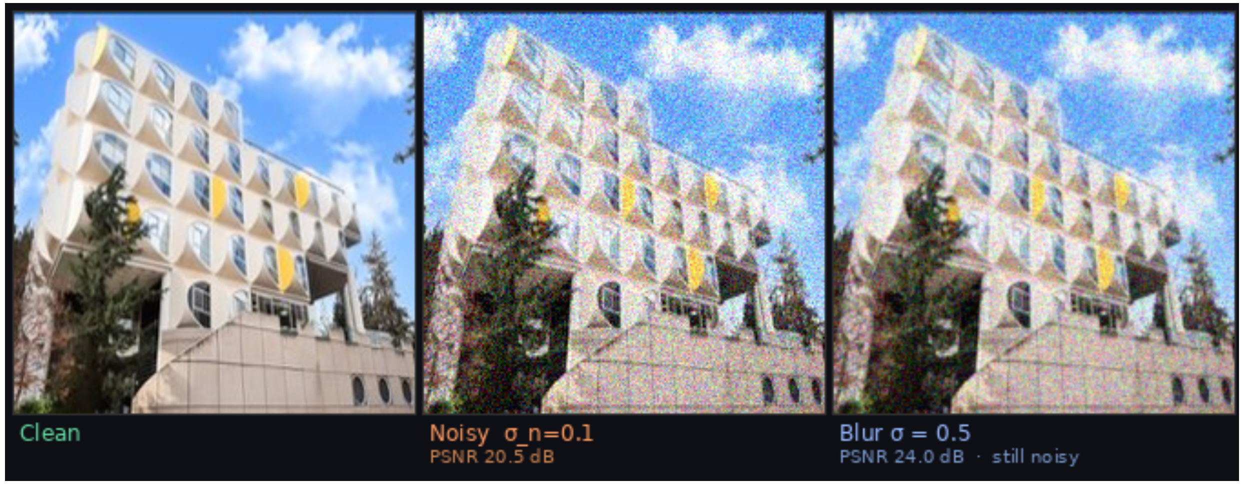

Denoising

Where does noise come from?

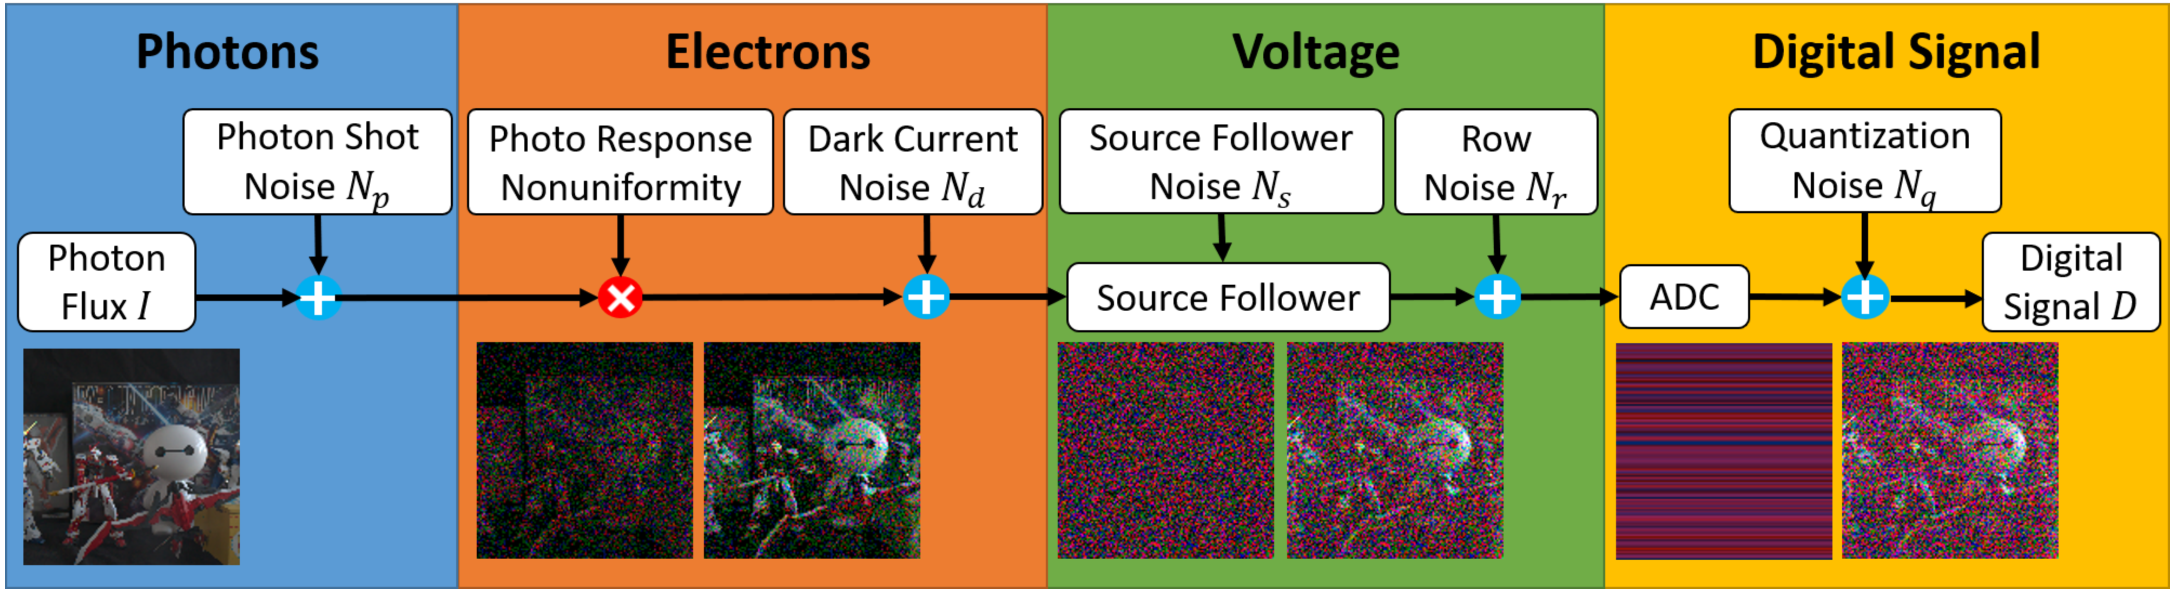

Naïve solution- Gaussian blur

Denoising

Where does noise come from?

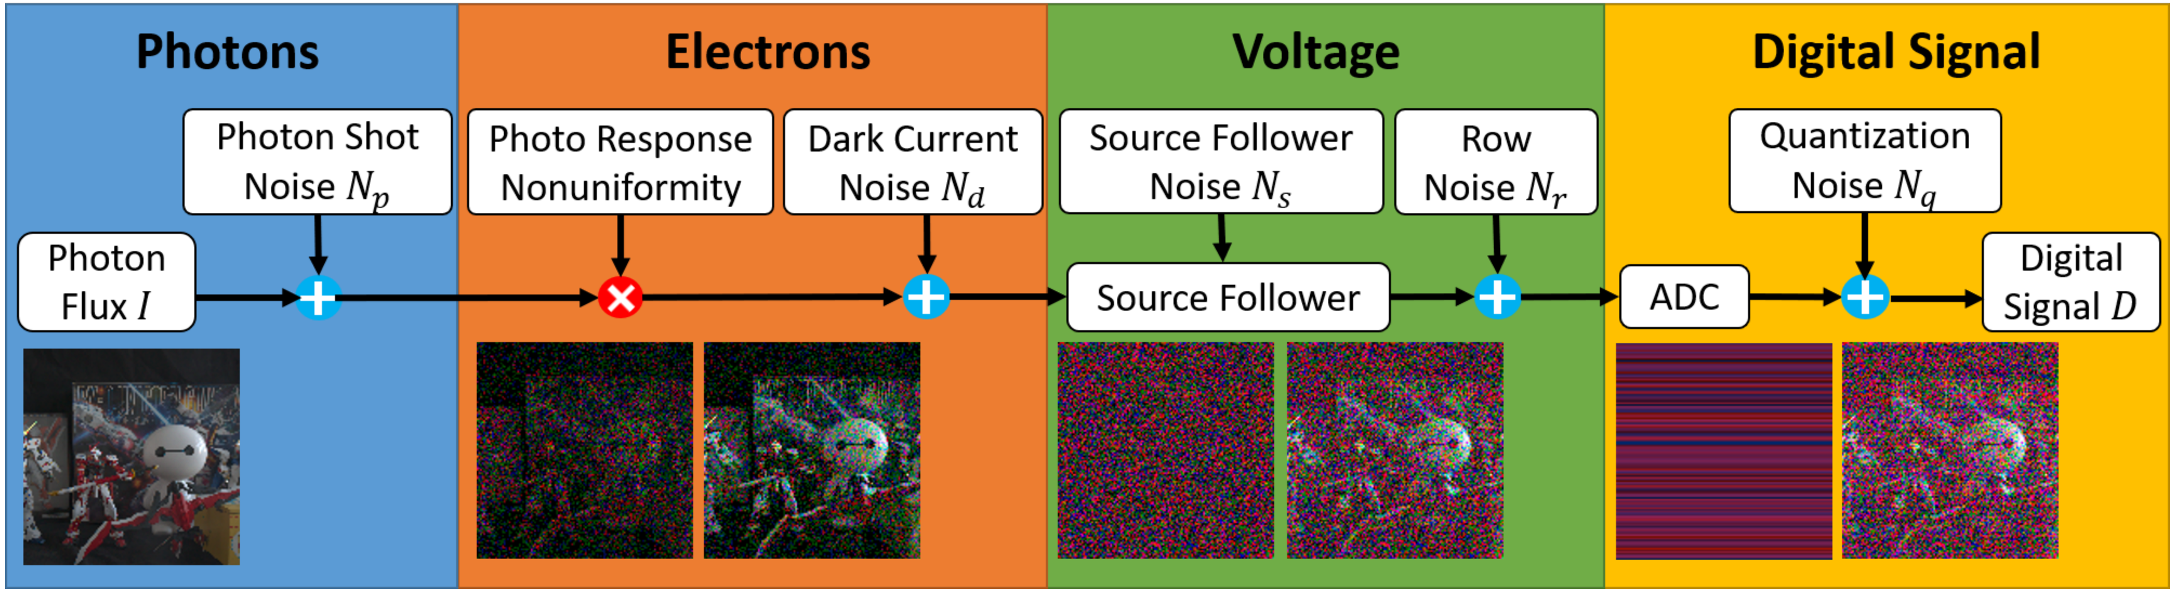

Naïve solution- Gaussian blur

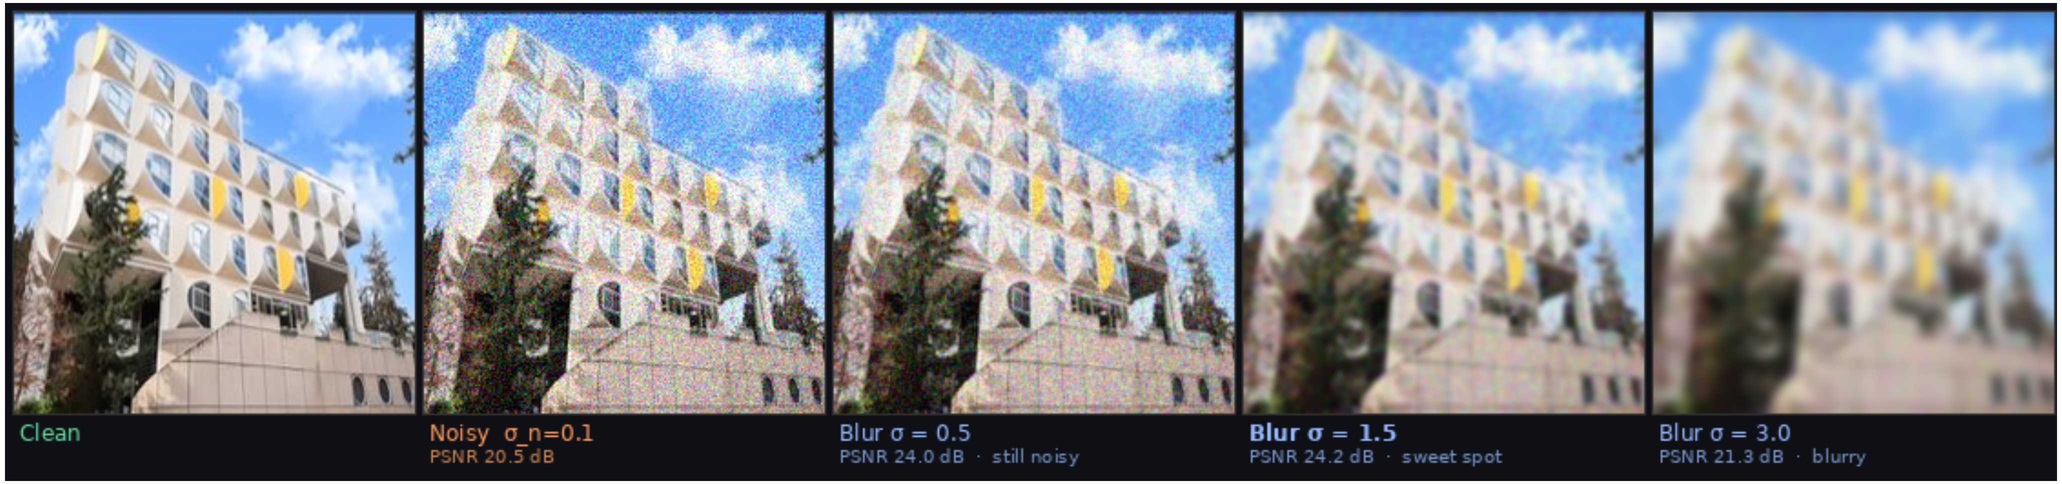

Denoising

Where does noise come from?

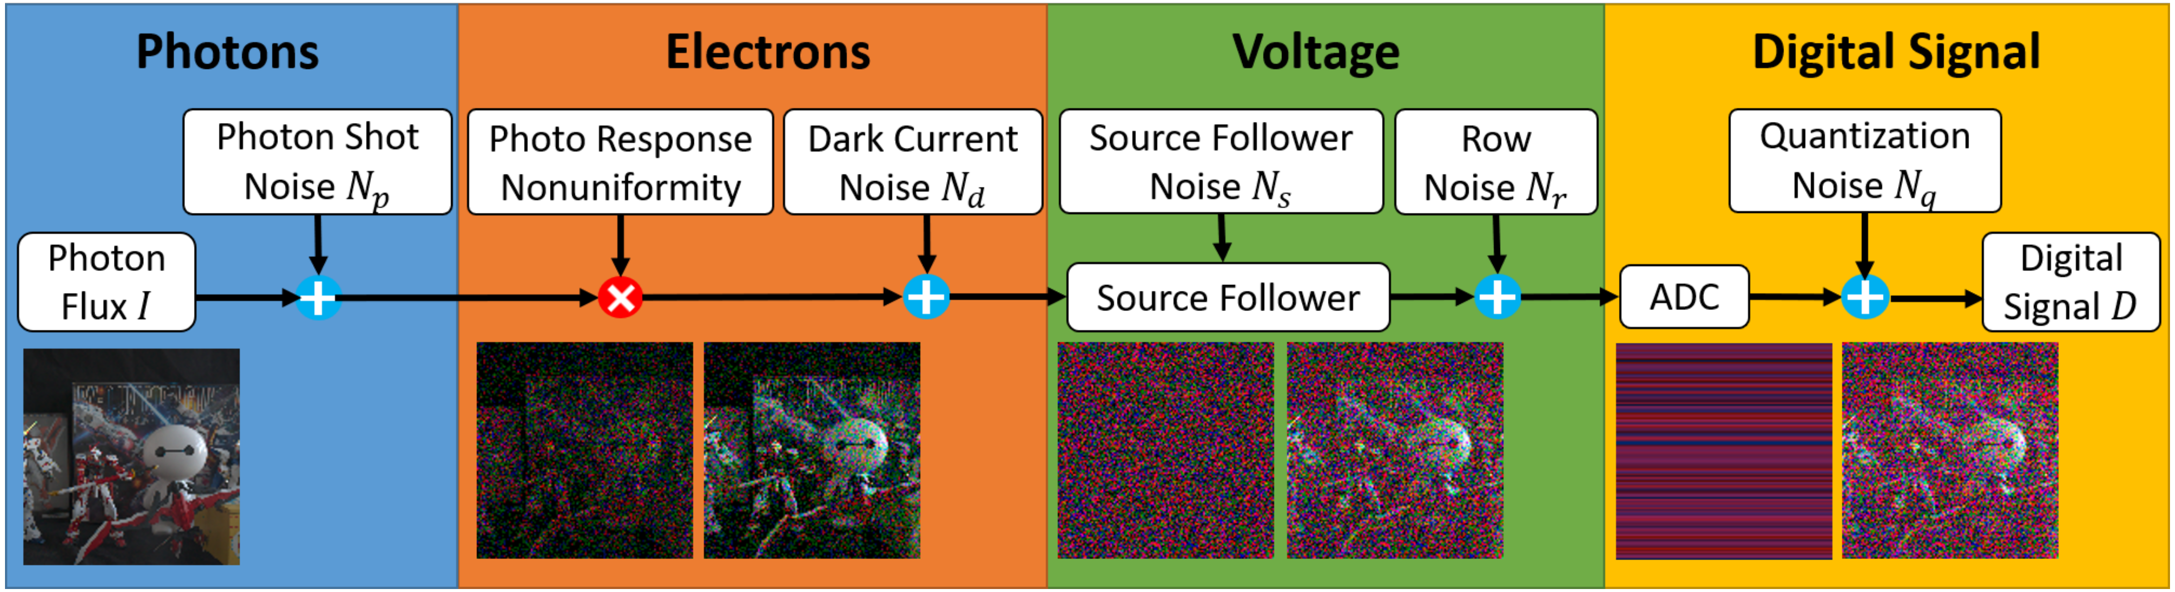

Naïve solution- Gaussian blur

Bilateral filter

A little less naïve

Bilateral filter

A little less naïve

Why did we have this Gaussian trade-off?

Bilateral filter

Why did we have this Gaussian trade-off?

A little less naïve

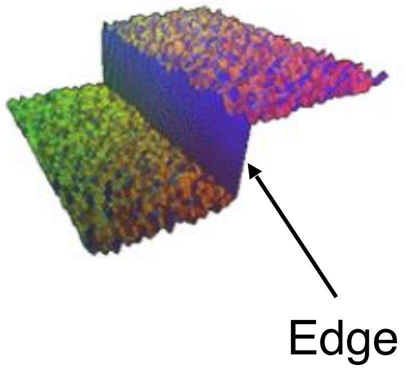

Bilateral filter

Why did we have this Gaussian trade-off?

A little less naïve

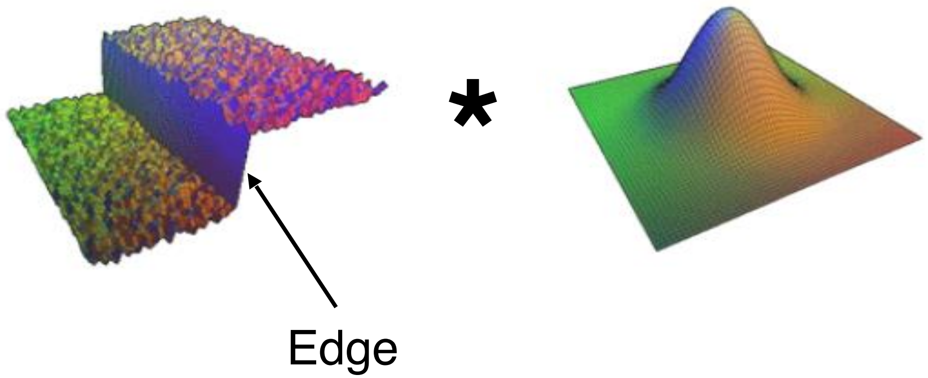

Bilateral filter

A little less naïve

Why did we have this Gaussian trade-off?

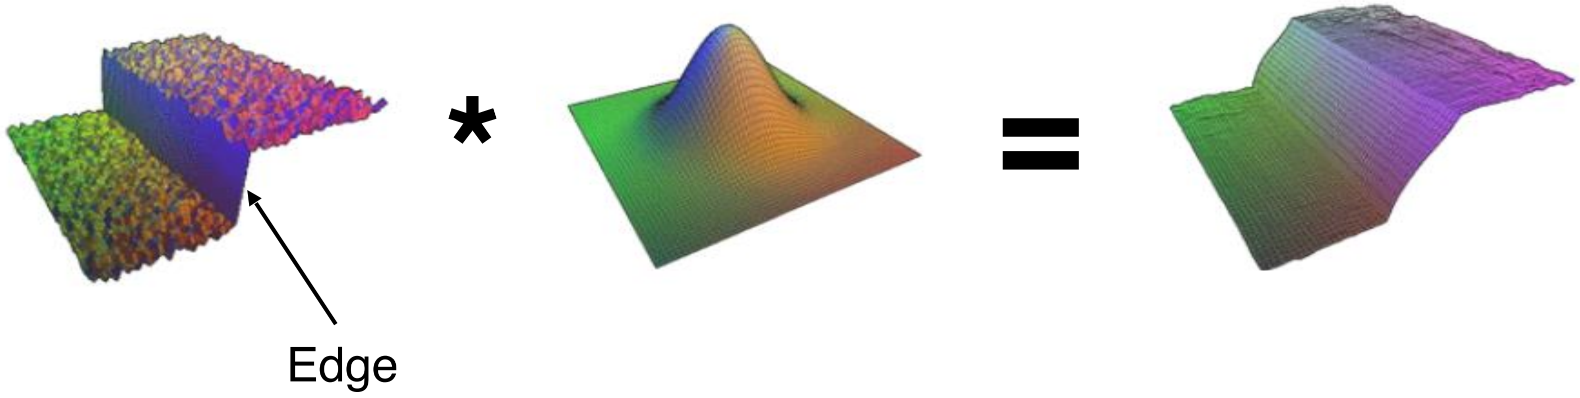

Bilateral filter

A little less naïve

Why did we have this Gaussian trade-off?

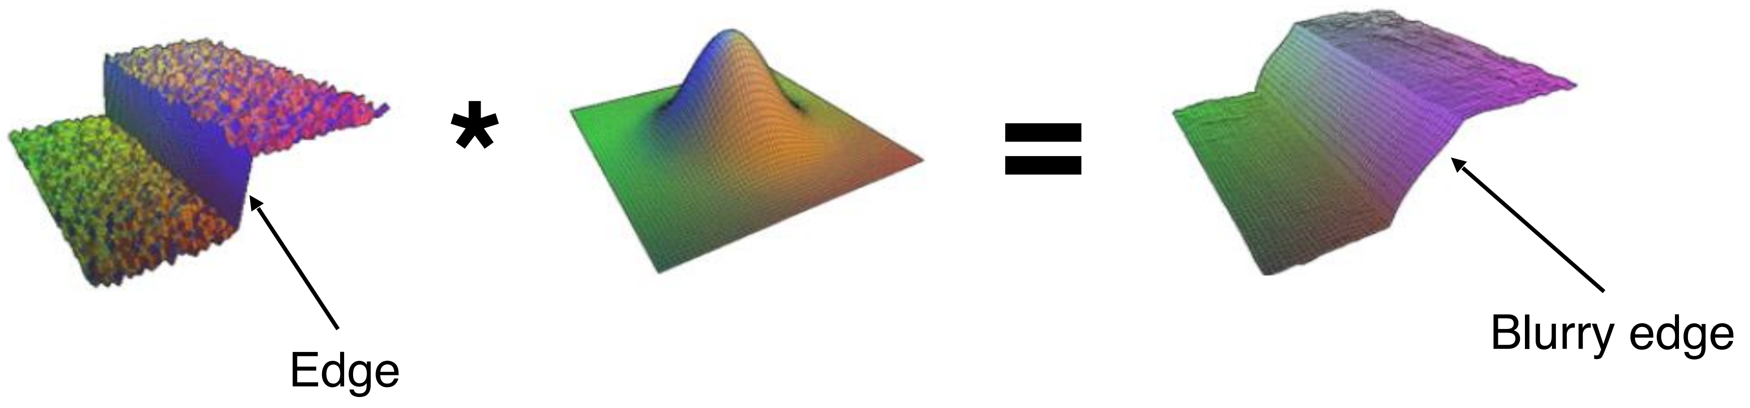

Bilateral filter

Why did we have this Gaussian trade-off?

A little less naïve

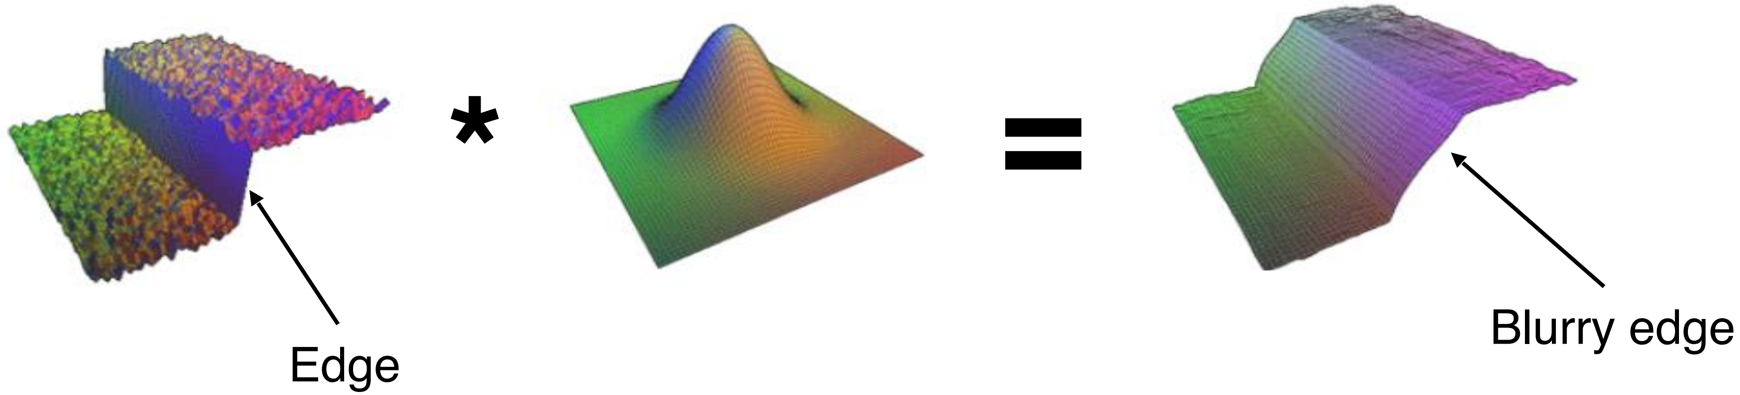

Idea: weigh not only special dist, also value dist (color)!

$$\frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathrm{s}}}(\|\mathbf{p} - \mathbf{q}\|)\, G_{\sigma_{\mathrm{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)\, I_{\mathbf{q}}$$

Bilateral filter

A little less naïve

Why did we have this Gaussian trade-off?

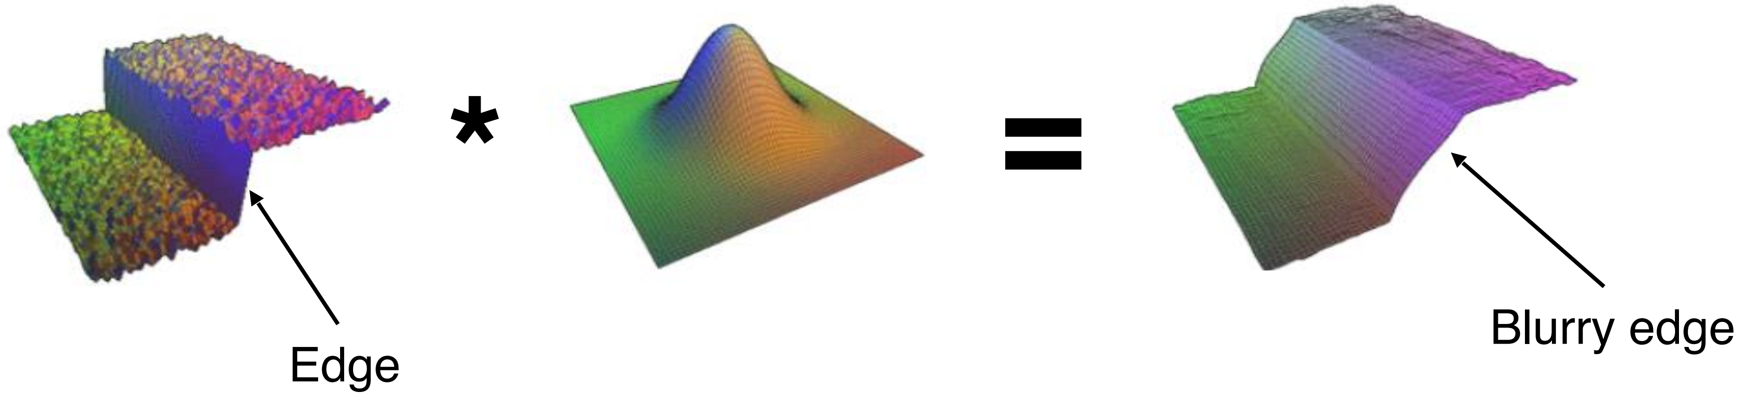

$$\frac{1}{W_{\mathbf{p}}}\sum_{\mathbf{q}\in\mathcal{S}} G_{\sigma_{\mathrm{s}}}(\|\mathbf{p}-\mathbf{q}\|)\, G_{\sigma_{\mathrm{r}}}(|I_{\mathbf{p}}-I_{\mathbf{q}}|)\, I_{\mathbf{q}}$$

Idea: weigh not only special dist, also value dist (color)!

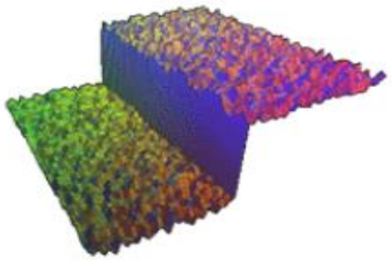

Bilateral filter

Why did we have this Gaussian trade-off?

A little less naïve

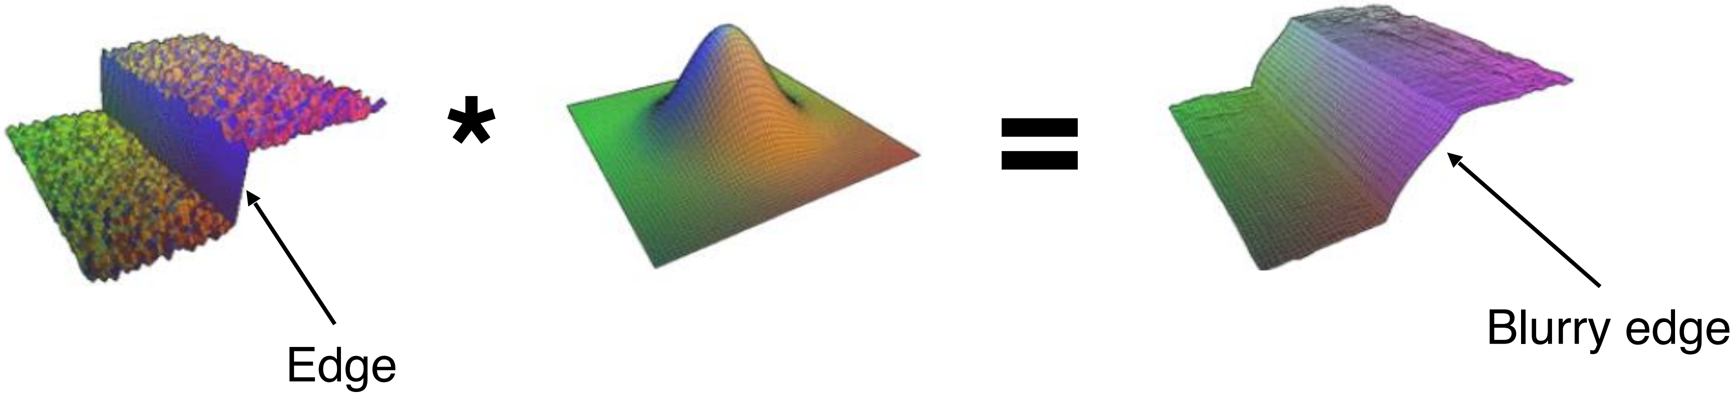

$$\frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathrm{s}}}(\|\mathbf{p}-\mathbf{q}\|)\, G_{\sigma_{\mathrm{r}}}(|I_{\mathbf{p}}-I_{\mathbf{q}}|)\, I_{\mathbf{q}}$$

Idea: weigh not only special dist, also value dist (color)!

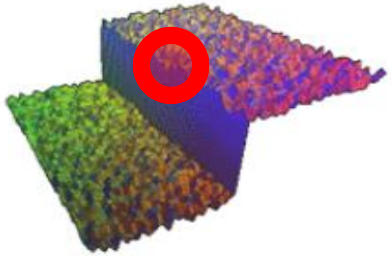

Bilateral filter

Why did we have this Gaussian trade-off?

A little less naïve

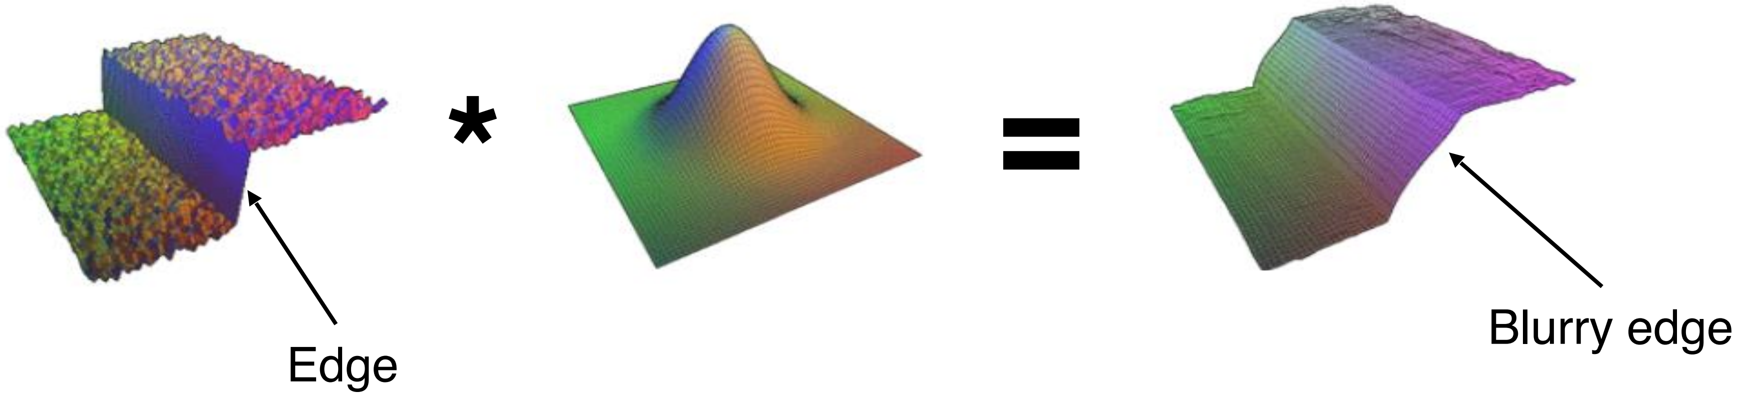

$$\frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathrm{s}}}(\|\mathbf{p} - \mathbf{q}\|)\, G_{\sigma_{\mathrm{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)\, I_{\mathbf{q}}$$

Idea: weigh not only special dist, also value dist (color)!

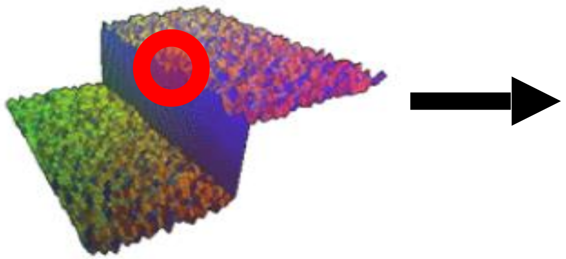

Bilateral filter

A little less naïve

Why did we have this Gaussian trade-off?

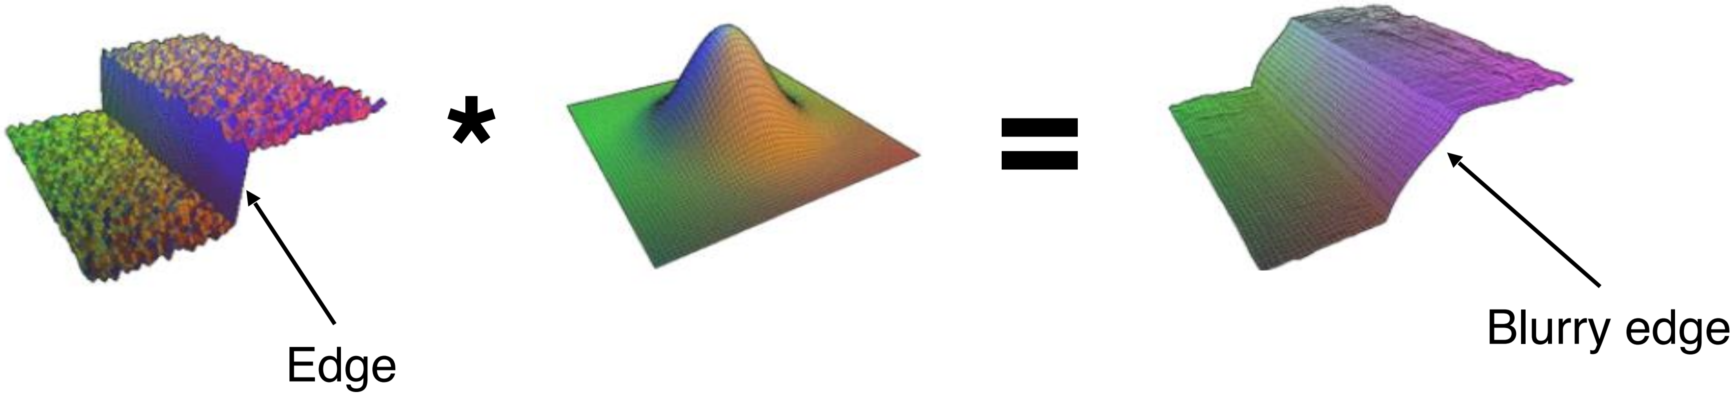

$$\frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathrm{s}}}(\|\mathbf{p} - \mathbf{q}\|)\; G_{\sigma_{\mathrm{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)\; I_{\mathbf{q}}$$

Idea: weigh not only special dist, also value dist (color)!

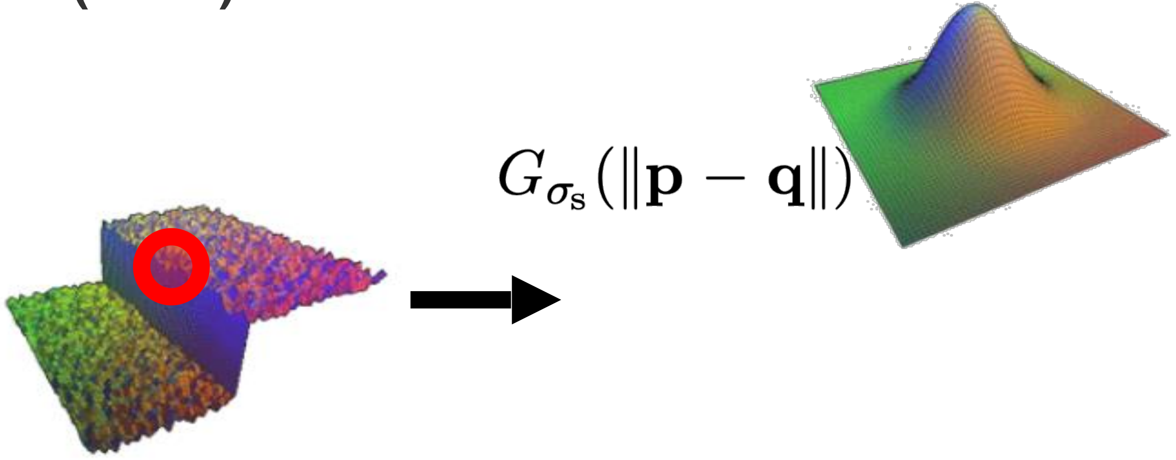

$$G_{\sigma_{\mathrm{s}}}(\|\mathbf{p} - \mathbf{q}\|)$$

Bilateral filter

A little less naïve

Why did we have this Gaussian trade-off?

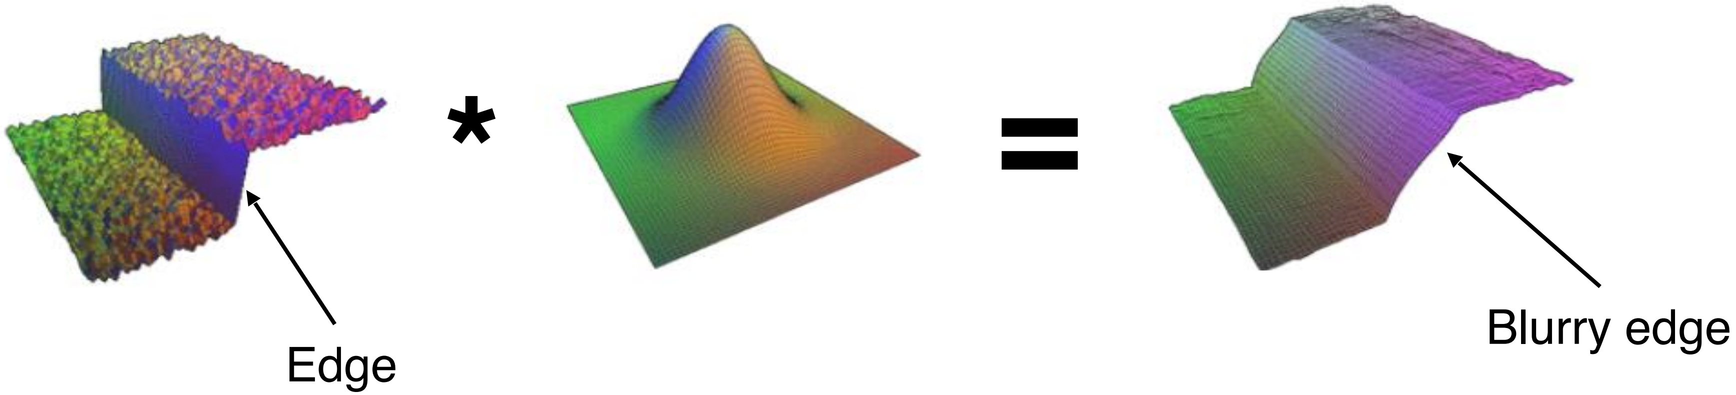

$$\frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathrm{s}}}(\|\mathbf{p} - \mathbf{q}\|)\, G_{\sigma_{\mathrm{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)\, I_{\mathbf{q}}$$

Idea: weigh not only special dist, also value dist (color)!

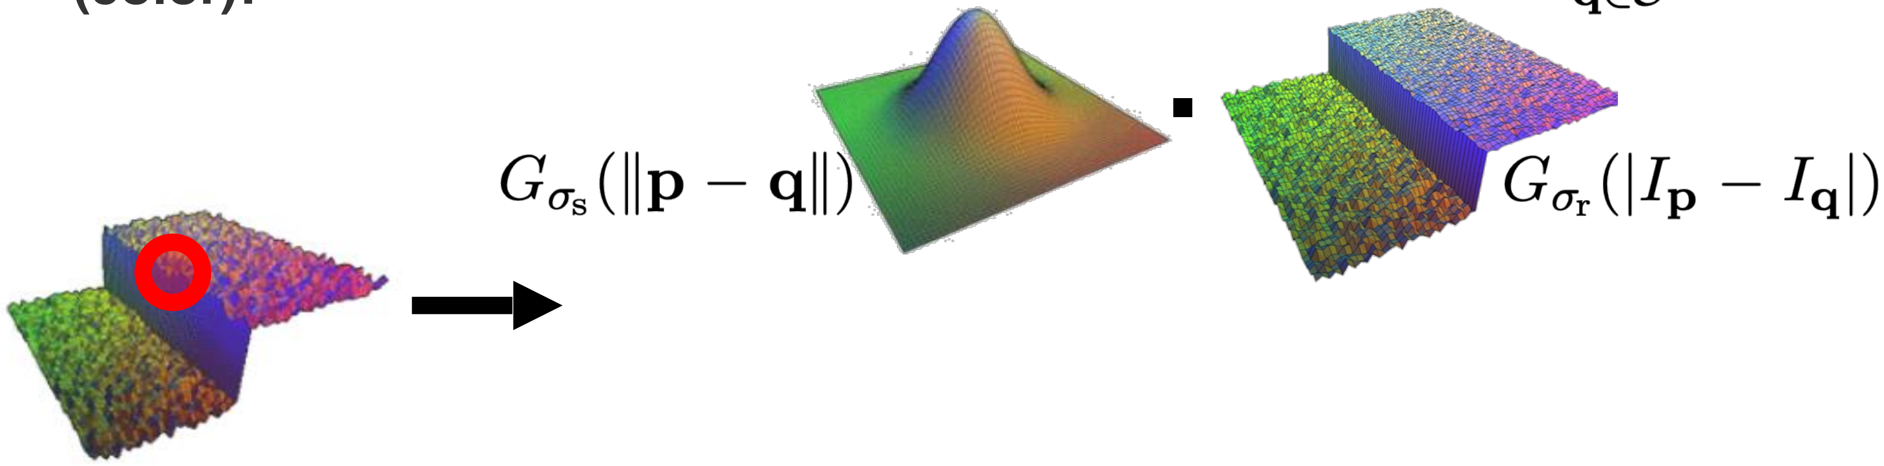

Bilateral filter

Why did we have this Gaussian trade-off?

A little less naïve

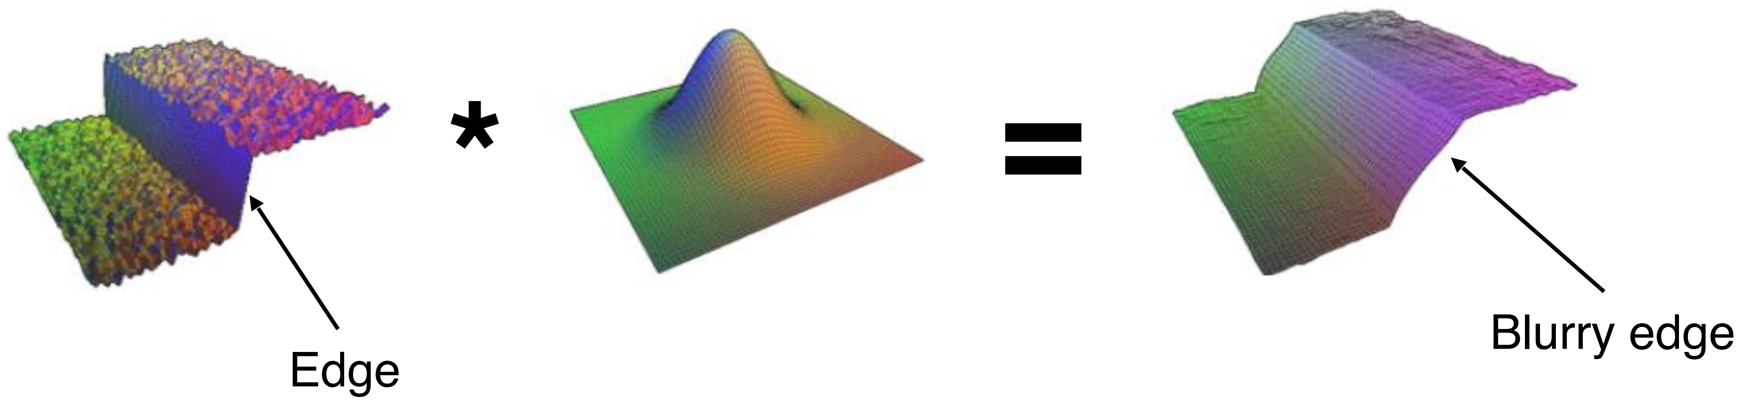

$$\frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_s}(\|\mathbf{p}-\mathbf{q}\|)\, G_{\sigma_r}(|I_{\mathbf{p}}-I_{\mathbf{q}}|)\, I_{\mathbf{q}}$$

Idea: weigh not only special dist, also value dist (color)!

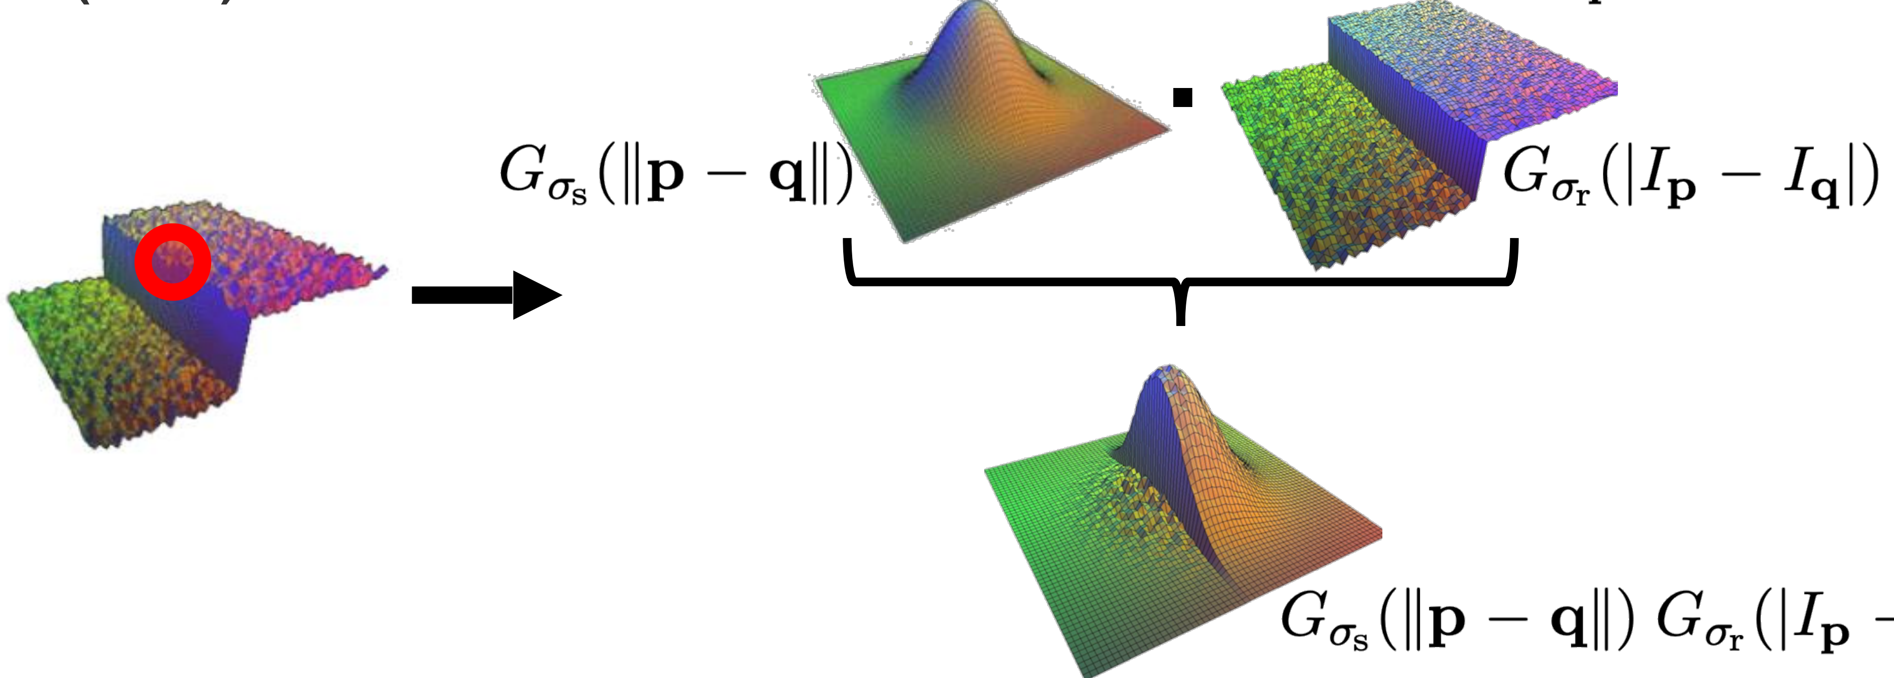

Bilateral filter

A little less naïve

Why did we have this Gaussian trade-off?

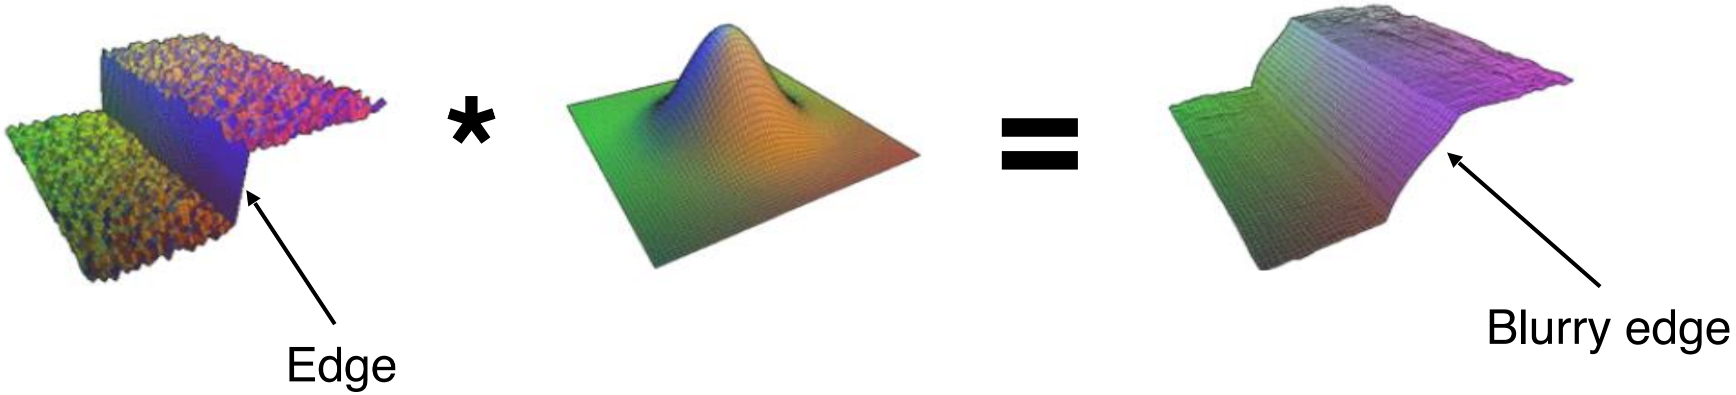

Idea: weigh not only special dist, also value dist (color)!

$$\frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathrm{s}}}(\|\mathbf{p} - \mathbf{q}\|)\, G_{\sigma_{\mathrm{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)\, I_{\mathbf{q}}$$

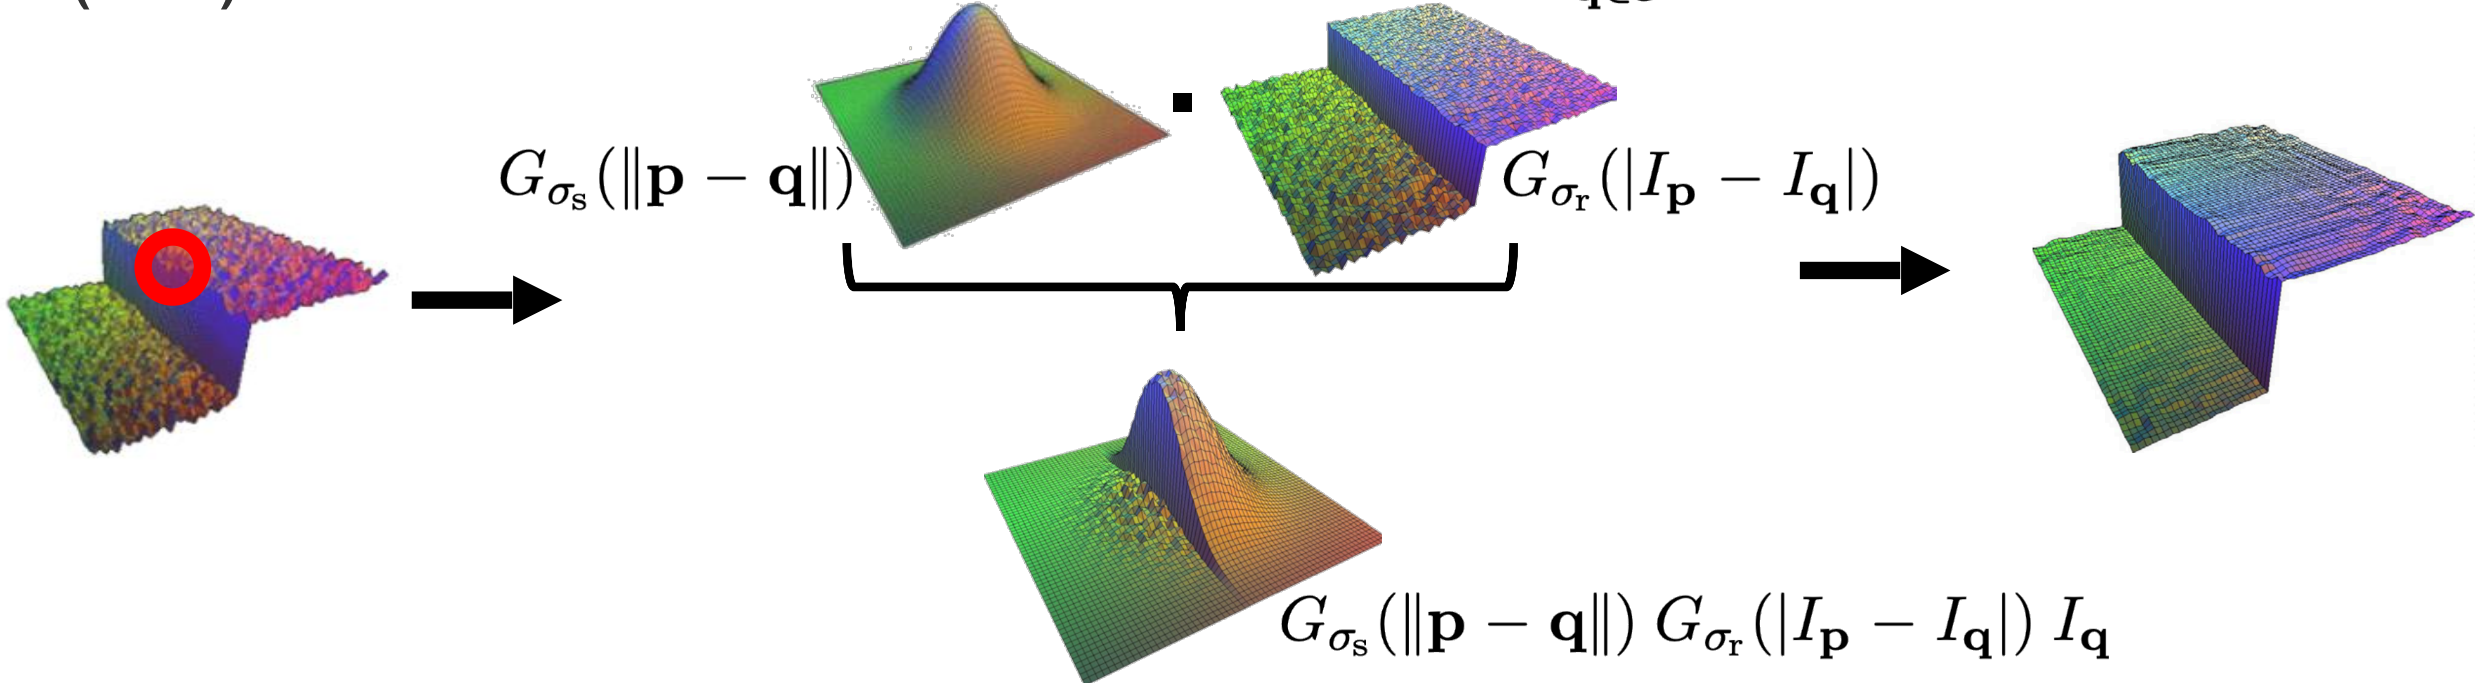

Bilateral filter

A little less naïve

Why did we have this Gaussian trade-off?

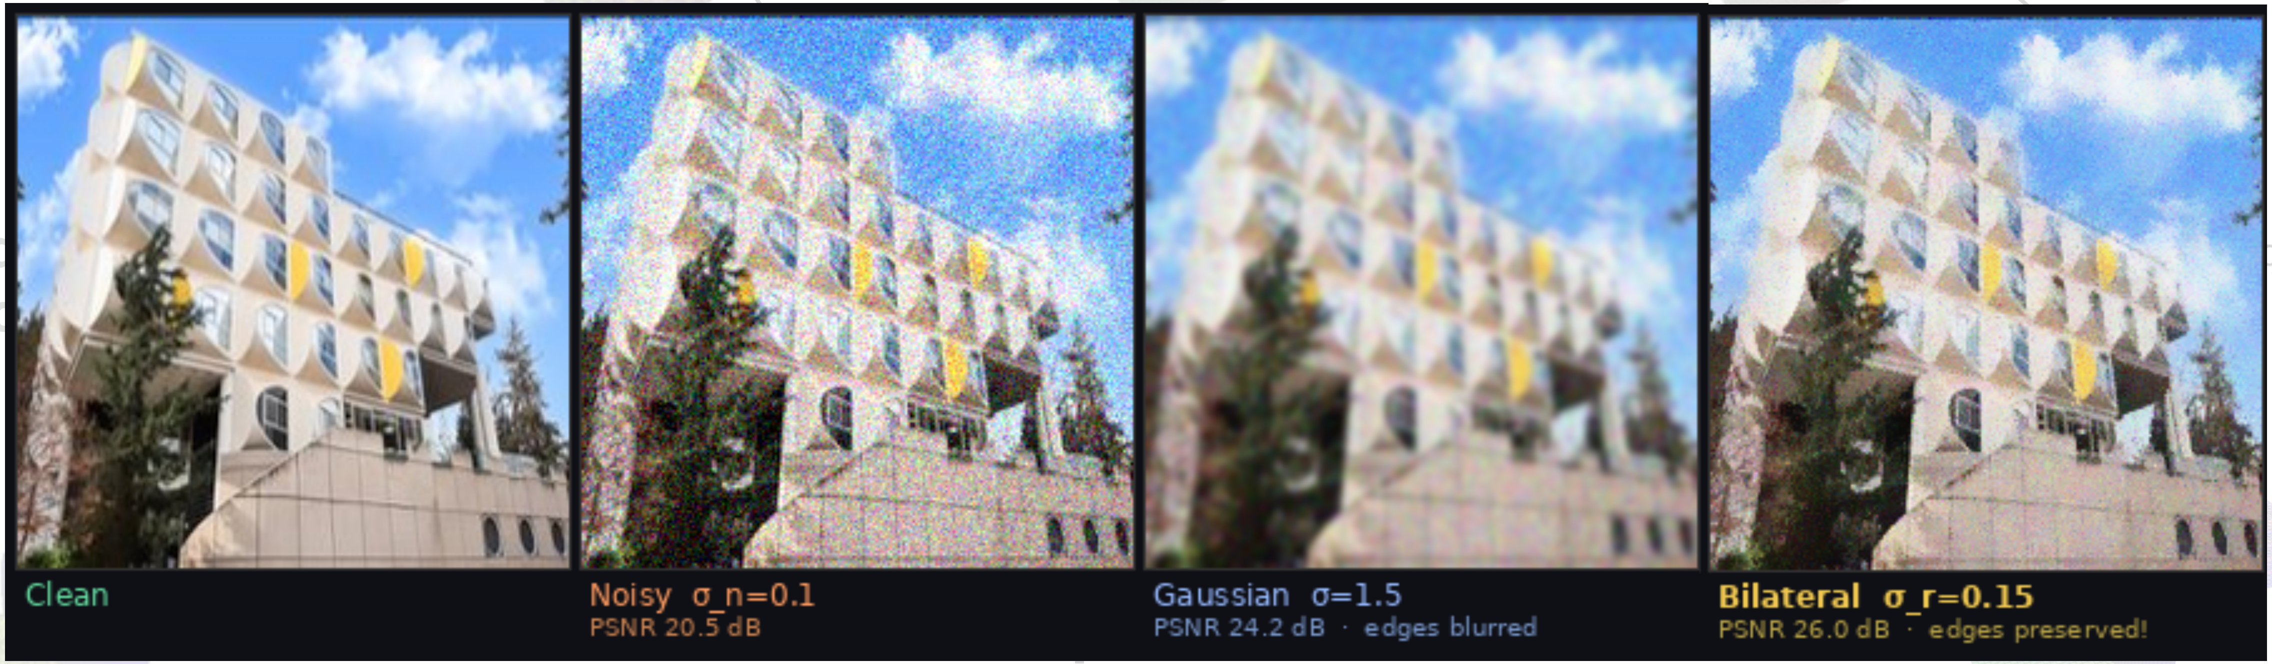

$$G_{\sigma_s}(\|\mathbf{p}-\mathbf{q}\|)\, G_{\sigma_r}(|I_\mathbf{p}-I_\mathbf{q}|)\, I_\mathbf{q}$$

Non-local means

Somewhat less naïve

Non-local means

Somewhat less naïve

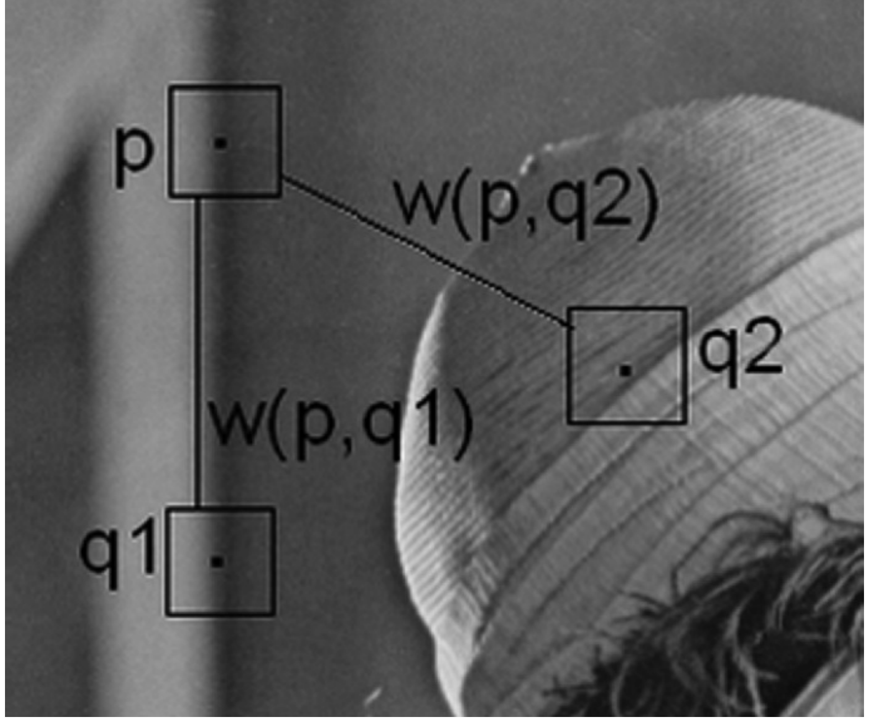

Non-local means

Somewhat less naïve

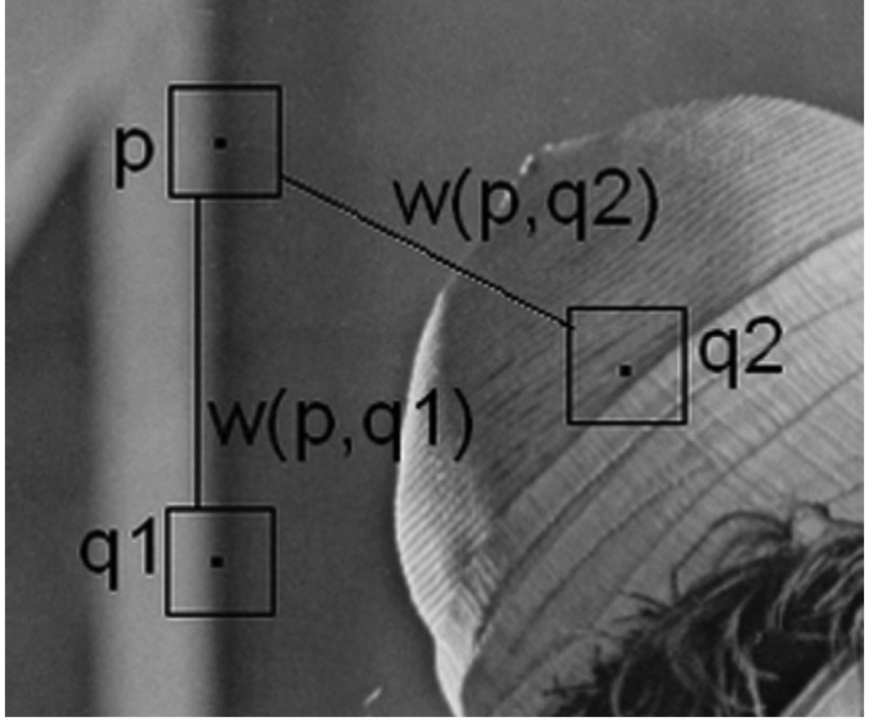

Digital Image Pipeline

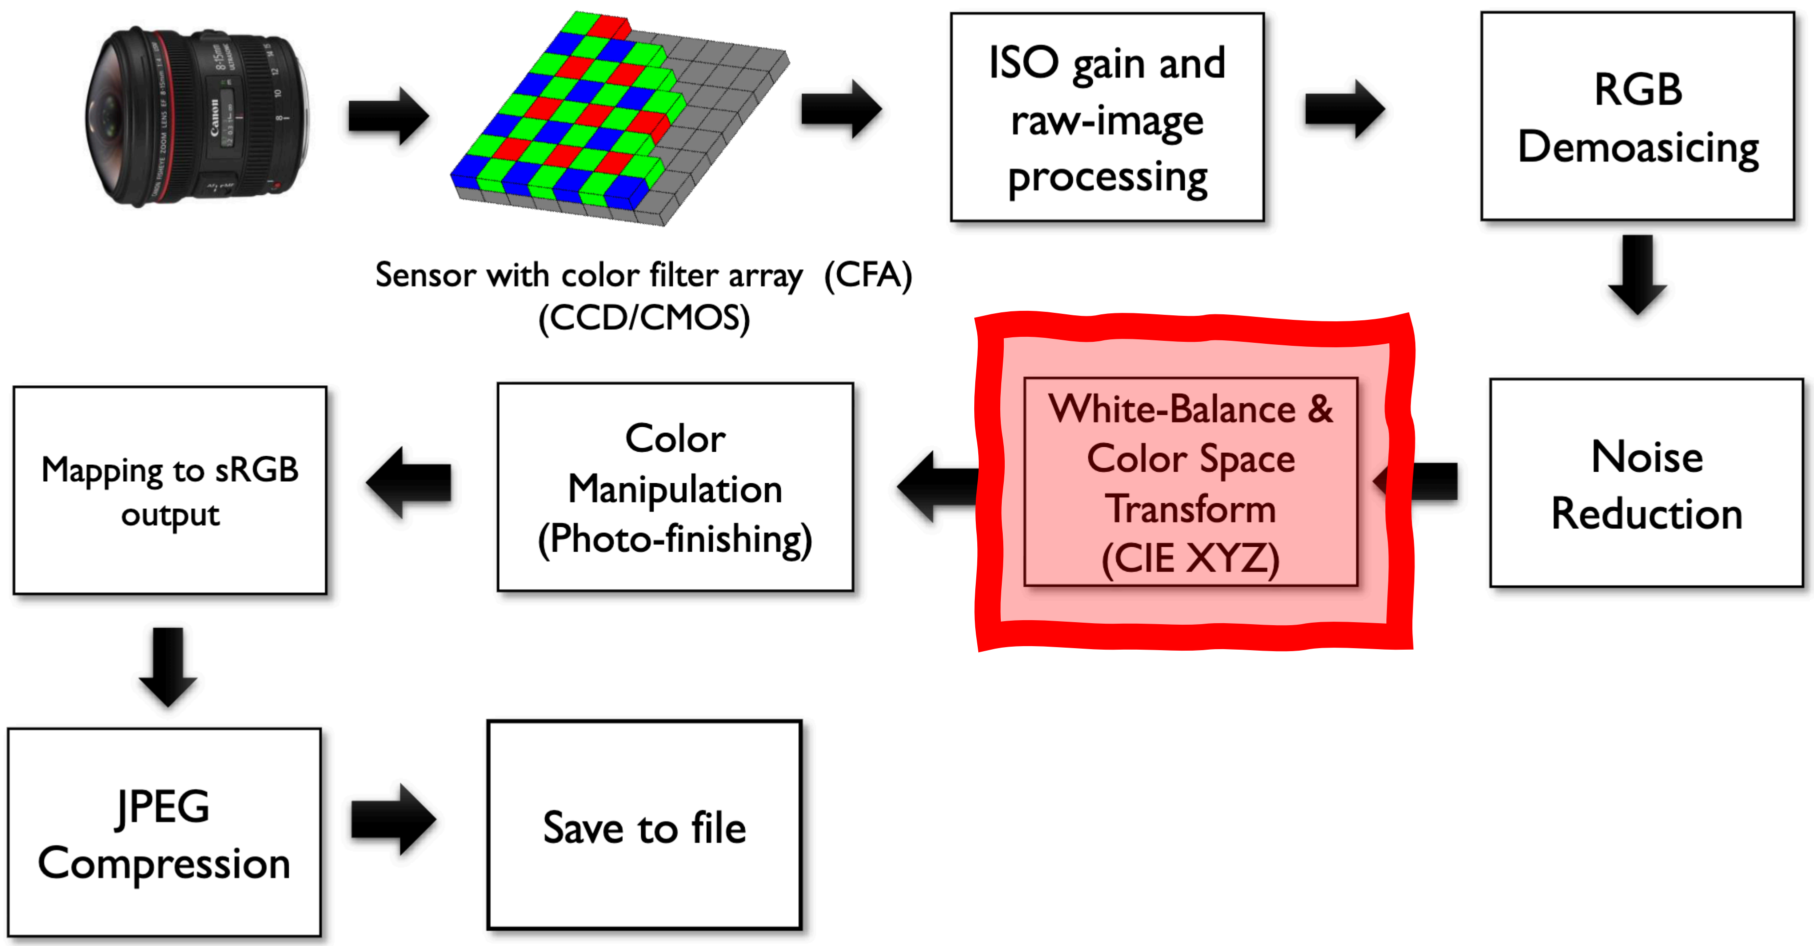

White balance

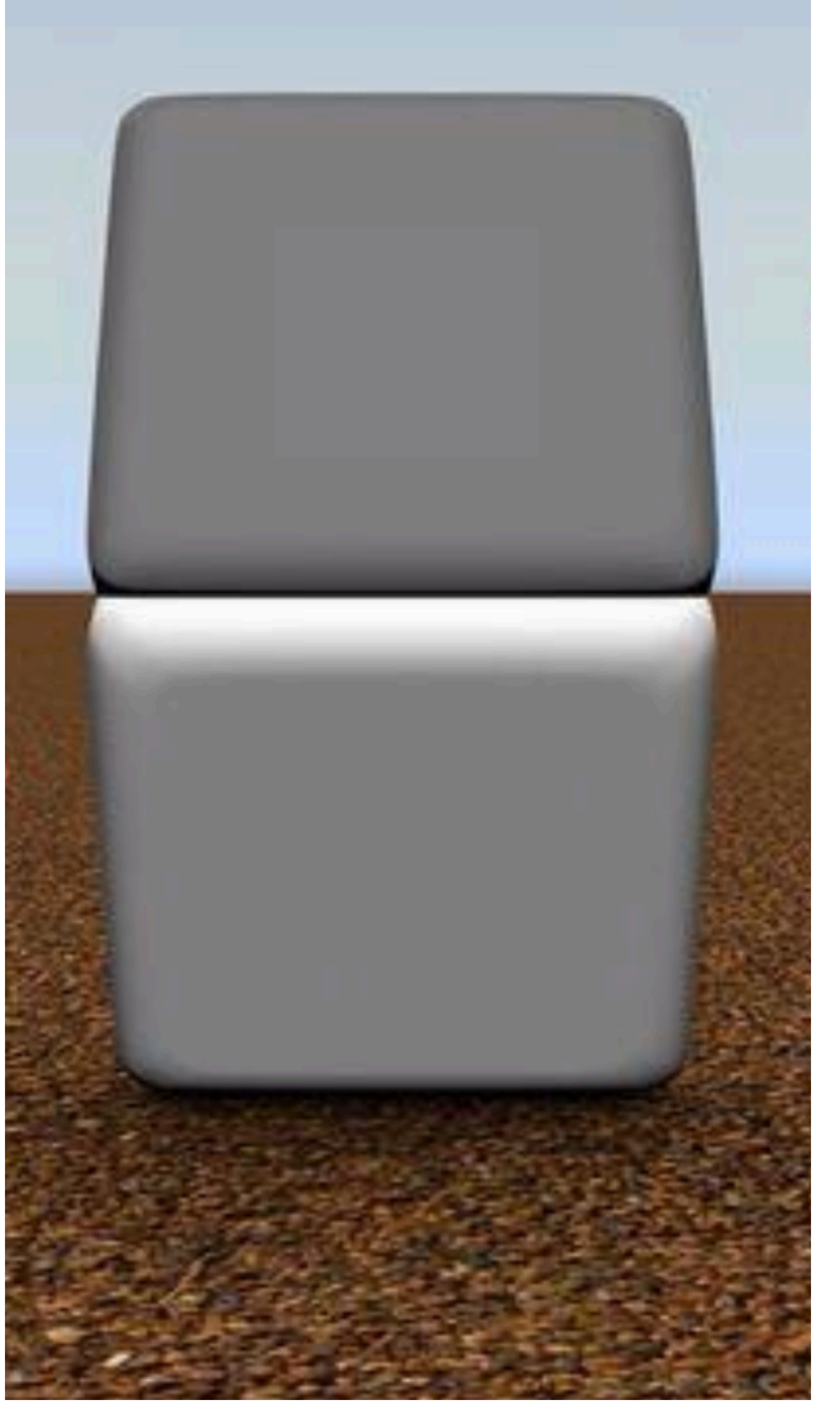

White balance

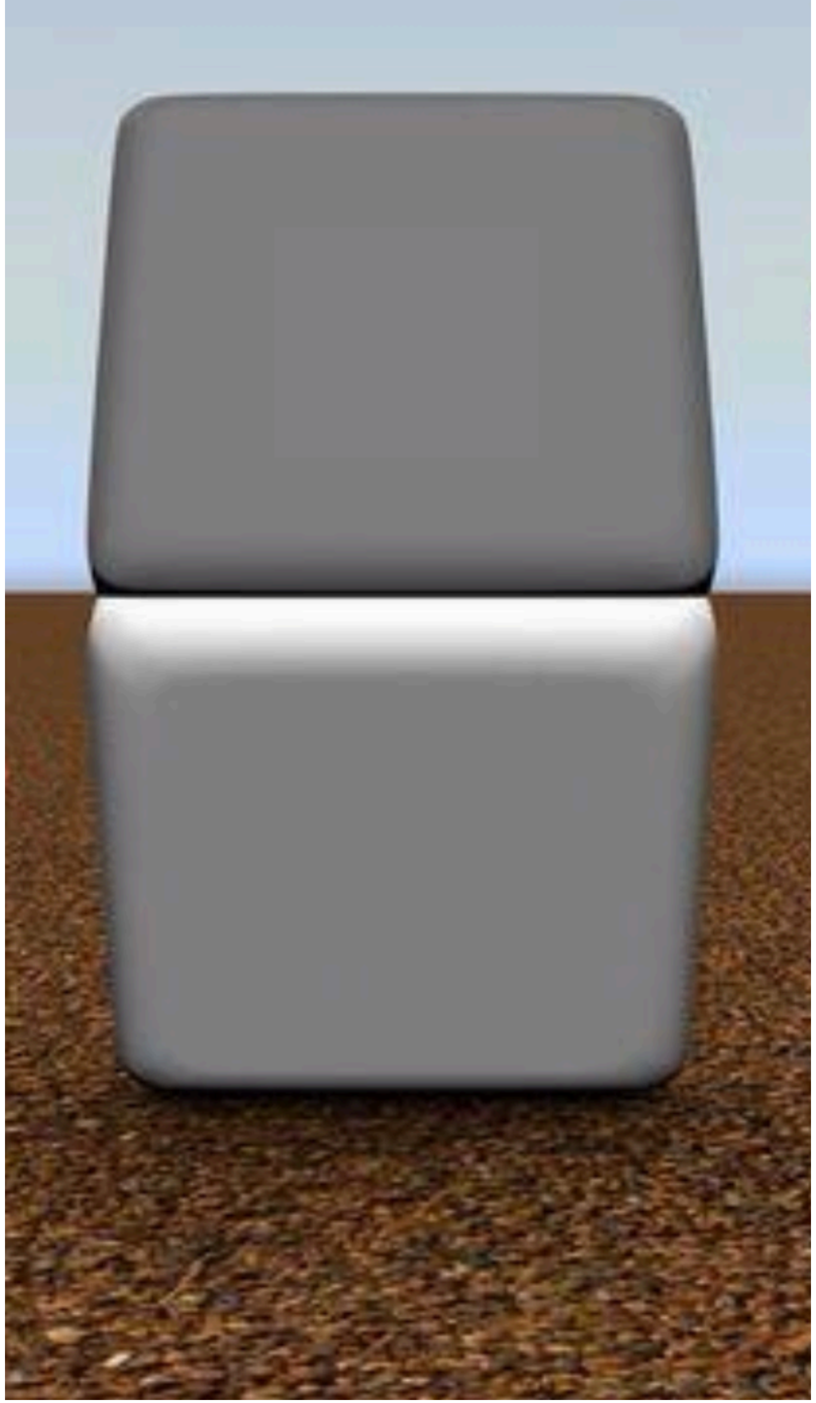

White balance

Same scene different lighting

White balance

Same scene different lighting

White balance

White balance

$$\begin{bmatrix} r_{wb} \\ g_{wb} \\ b_{wb} \end{bmatrix} = \begin{bmatrix} 1/\ell_r & 0 & 0 \\ 0 & 1/\ell_g & 0 \\ 0 & 0 & 1/\ell_b \end{bmatrix} \begin{bmatrix} r \\ g \\ b \end{bmatrix}$$

White balance

$$\begin{bmatrix} r_{wb} \\ g_{wb} \\ b_{wb} \end{bmatrix} = \begin{bmatrix} 1/\ell_r & 0 & 0 \\ 0 & 1/\ell_g & 0 \\ 0 & 0 & 1/\ell_b \end{bmatrix} \begin{bmatrix} r \\ g \\ b \end{bmatrix}$$

Presets

White balance

$$\begin{bmatrix} r_{wb} \\ g_{wb} \\ b_{wb} \end{bmatrix} = \begin{bmatrix} 1/\ell_r & 0 & 0 \\ 0 & 1/\ell_g & 0 \\ 0 & 0 & 1/\ell_b \end{bmatrix} \begin{bmatrix} r \\ g \\ b \end{bmatrix}$$

Presets

Nikon D7000

Sunny

$$\begin{bmatrix} 2.0273 & 0 & 0 \\ 0 & 1.0000 & 0 \\ 0 & 0 & 1.3906 \end{bmatrix}$$

Incandescent

$$\begin{bmatrix} 1.3047 & 0 & 0 \\ 0 & 1.0000 & 0 \\ 0 & 0 & 2.2148 \end{bmatrix}$$

Shade

$$\begin{bmatrix} 2.4922 & 0 & 0 \\ 0 & 1.0000 & 0 \\ 0 & 0 & 1.1367 \end{bmatrix}$$

Canon 1D

Daylight

$$\begin{bmatrix} 2.0938 & 0 & 0 \\ 0 & 1.0000 & 0 \\ 0 & 0 & 1.5020 \end{bmatrix}$$

Tungsten

$$\begin{bmatrix} 1.4511 & 0 & 0 \\ 0 & 1.0000 & 0 \\ 0 & 0 & 2.3487 \end{bmatrix}$$

Shade

$$\begin{bmatrix} 2.4628 & 0 & 0 \\ 0 & 1.0000 & 0 \\ 0 & 0 & 1.2275 \end{bmatrix}$$

Sony A57K

Daylight

$$\begin{bmatrix} 2.6836 & 0 & 0 \\ 0 & 1.0000 & 0 \\ 0 & 0 & 1.5586 \end{bmatrix}$$

Tungsten

$$\begin{bmatrix} 1.6523 & 0 & 0 \\ 0 & 1.0000 & 0 \\ 0 & 0 & 2.7422 \end{bmatrix}$$

Shade

$$\begin{bmatrix} 3.1953 & 0 & 0 \\ 0 & 1.0000 & 0 \\ 0 & 0 & 1.2891 \end{bmatrix}$$

White balance

$$\begin{bmatrix} r_{wb} \\ g_{wb} \\ b_{wb} \end{bmatrix} = \begin{bmatrix} 1/\ell_r & 0 & 0 \\ 0 & 1/\ell_g & 0 \\ 0 & 0 & 1/\ell_b \end{bmatrix} \begin{bmatrix} r \\ g \\ b \end{bmatrix}$$

Presets

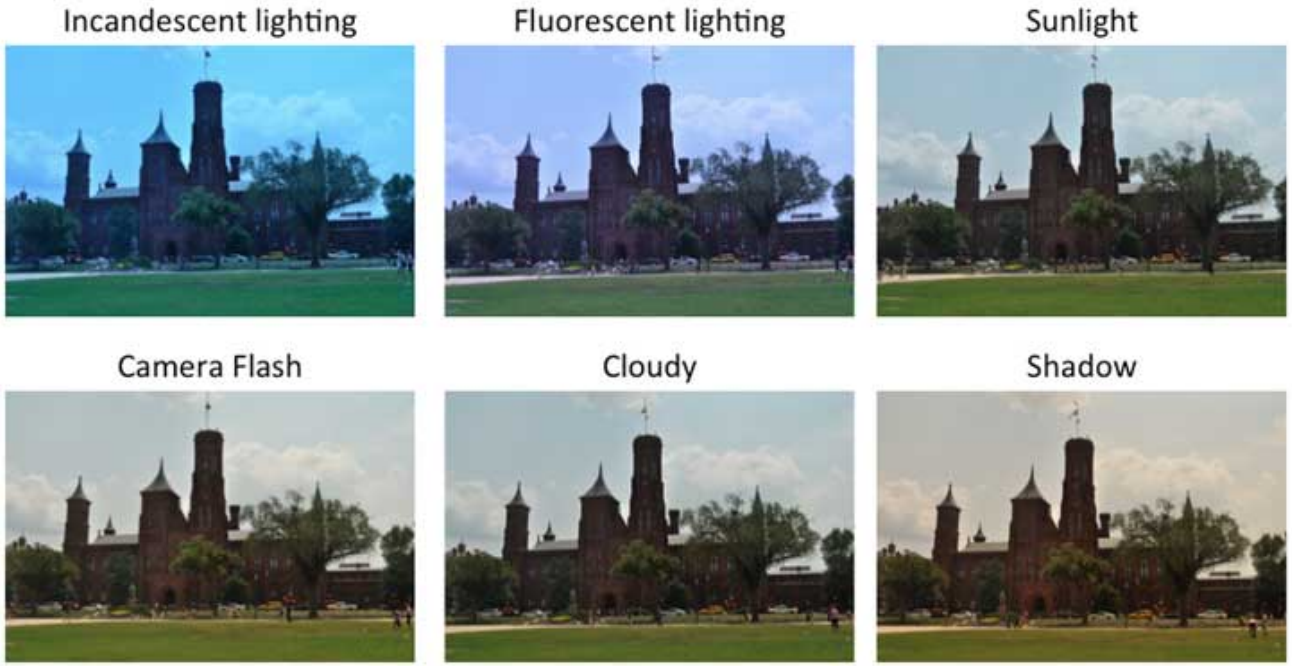

White balance

$$\begin{bmatrix} r_{wb} \\ g_{wb} \\ b_{wb} \end{bmatrix} = \begin{bmatrix} 1/\ell_r & 0 & 0 \\ 0 & 1/\ell_g & 0 \\ 0 & 0 & 1/\ell_b \end{bmatrix} \begin{bmatrix} r \\ g \\ b \end{bmatrix}$$

Presets

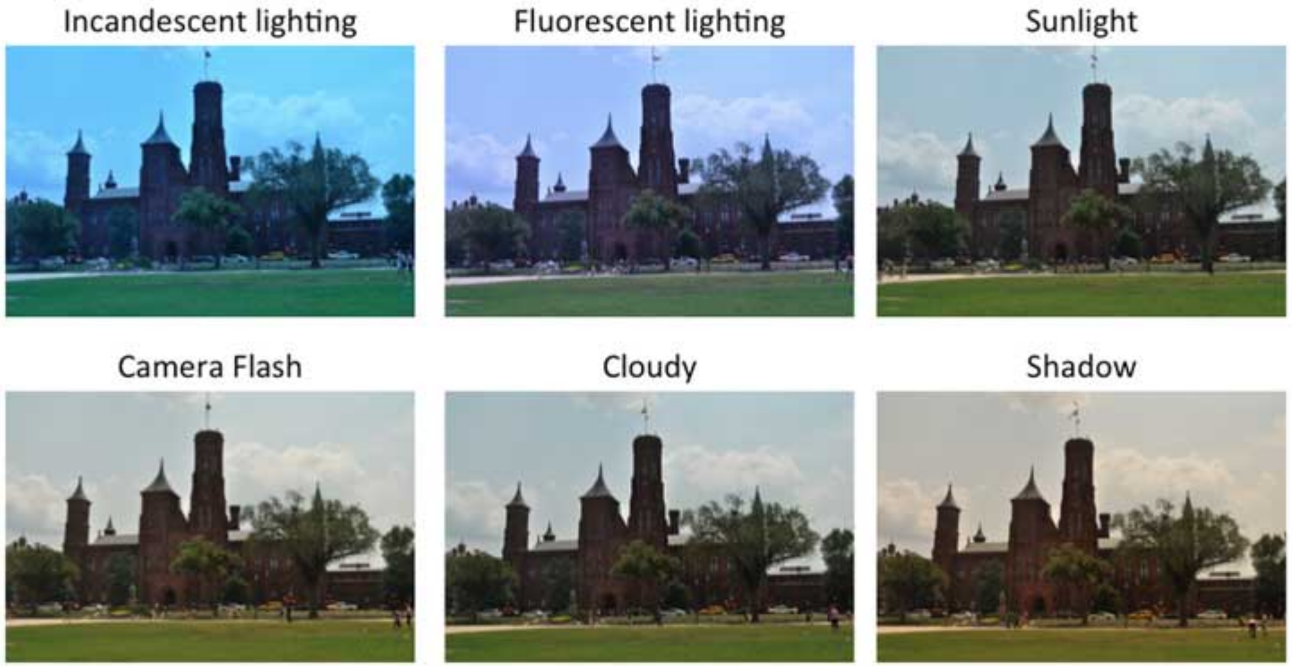

Automatic

Basic: Gray world / White patch

$$\begin{bmatrix} R' \\ G' \\ B' \end{bmatrix} = \begin{bmatrix} G_{avg}/R_{avg} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & G_{avg}/B_{avg} \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

White balance

$$\begin{bmatrix} r_{wb} \\ g_{wb} \\ b_{wb} \end{bmatrix} = \begin{bmatrix} 1/\ell_r & 0 & 0 \\ 0 & 1/\ell_g & 0 \\ 0 & 0 & 1/\ell_b \end{bmatrix} \begin{bmatrix} r \\ g \\ b \end{bmatrix}$$

Presets

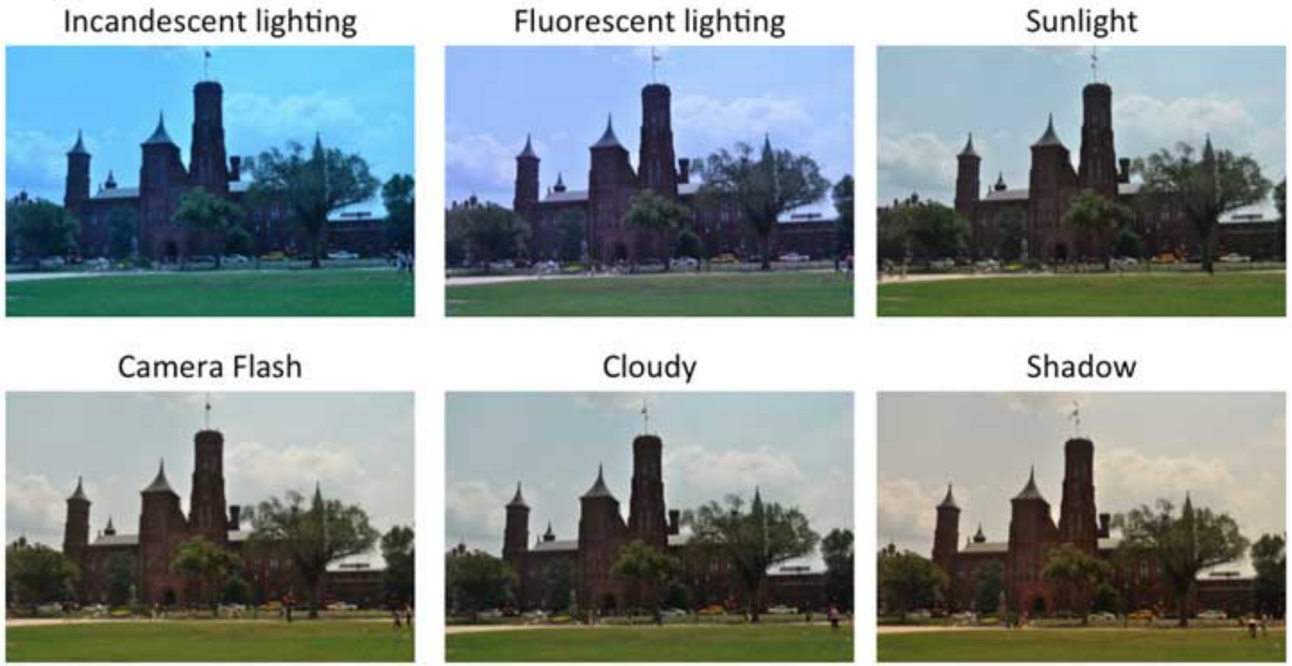

Automatic

Basic: Gray world / White patch

$$\begin{bmatrix} R' \\ G' \\ B' \end{bmatrix} = \begin{bmatrix} G_{avg}/R_{avg} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & G_{avg}/B_{avg} \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

Many fancy algos. Each cam has its own

White balance

$$\begin{bmatrix} r_{wb} \\ g_{wb} \\ b_{wb} \end{bmatrix} = \begin{bmatrix} 1/\ell_r & 0 & 0 \\ 0 & 1/\ell_g & 0 \\ 0 & 0 & 1/\ell_b \end{bmatrix} \begin{bmatrix} r \\ g \\ b \end{bmatrix}$$

Presets

Automatic

Basic: Gray world / White patch

$$\begin{bmatrix} R' \\ G' \\ B' \end{bmatrix} = \begin{bmatrix} G_{avg}/R_{avg} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & G_{avg}/B_{avg} \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

Many fancy algos. Each cam has its own

Photo-Finishing

Tutorial

Photo-Finishing

Gamma / Tone curve

Tutorial

Photo-Finishing

Gamma / Tone curve

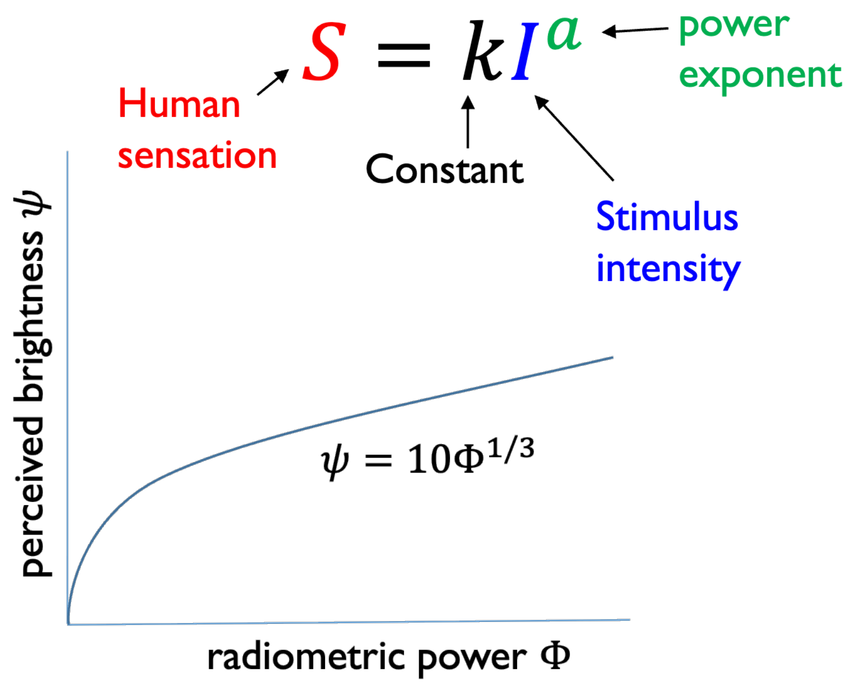

Dr. Stanley Stevens showed that most human sensations follow a power-law relationship between stimuli and sensation.

Tutorial

Photo-Finishing

Gamma / Tone curve

$$S = kI^{a}$$

Human sensation → S; Constant → k; Stimulus intensity → I; power exponent → a

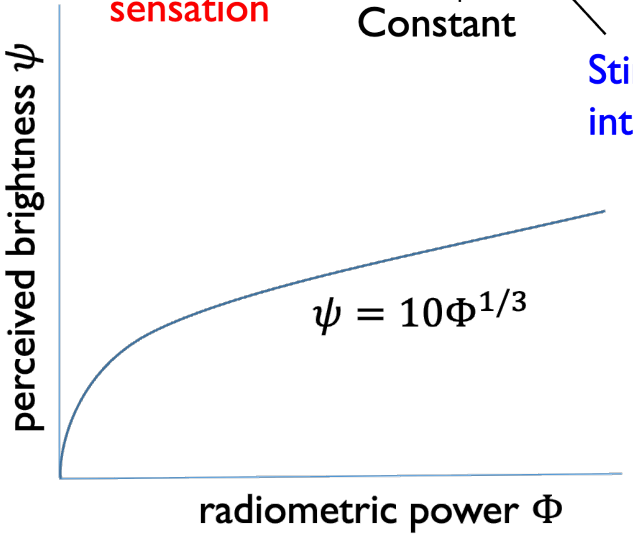

Dr. Stanley Stevens showed that most human sensations follow a power-law relationship between stimuli and sensation.


Tutorial

Photo-Finishing

Gamma / Tone curve

$$S = kI^a$$

Human sensation → S; Constant → k; Stimulus intensity → I; power exponent → a

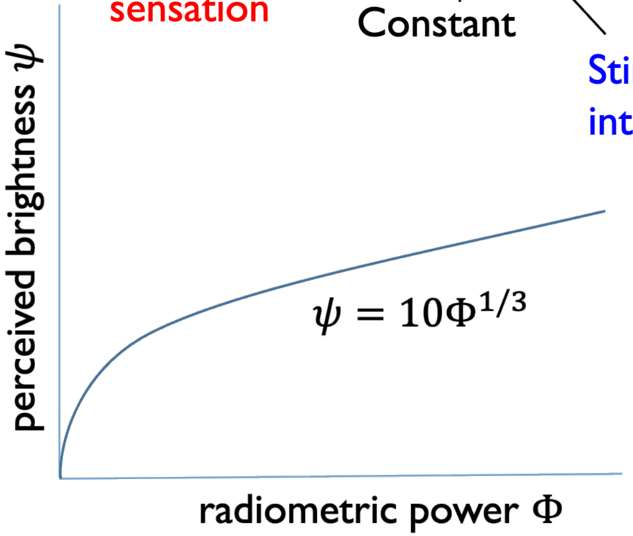

Dr. Stanley Stevens showed that most human sensations follow a power-law relationship between stimuli and sensation.

Contrast enhancement

Tutorial

Photo-Finishing

Gamma / Tone curve

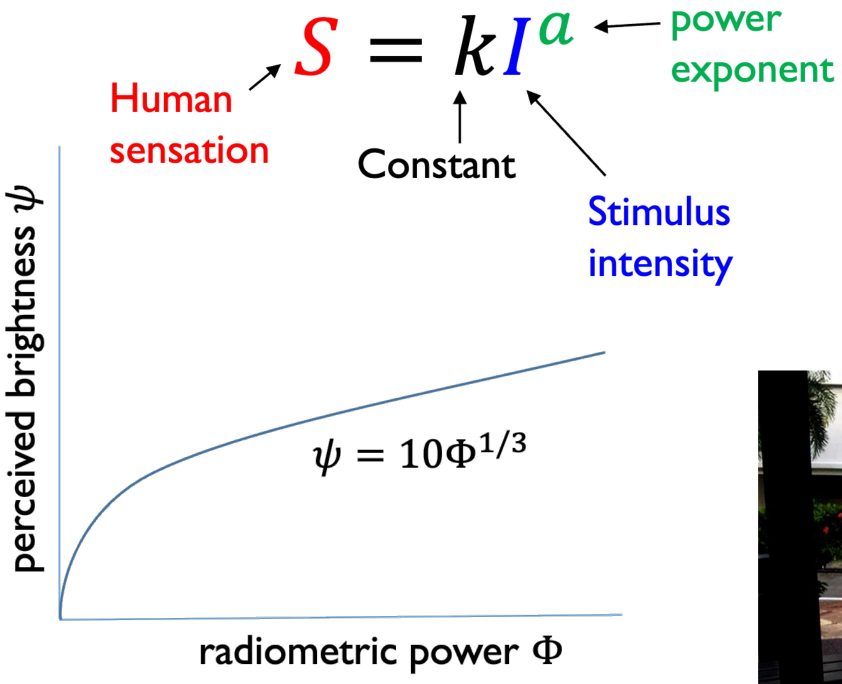

Dr. Stanley Stevens showed that most human sensations follow a power-law relationship between stimuli and sensation.

Contrast enhancement

Tutorial

Sharpening

Photo-Finishing

Gamma / Tone curve

Contrast enhancement

Sharpening

JPEG compression → sRGB

$$S = kI^{a}$$

Human sensation → S; Constant → k; Stimulus intensity → I; power exponent → a

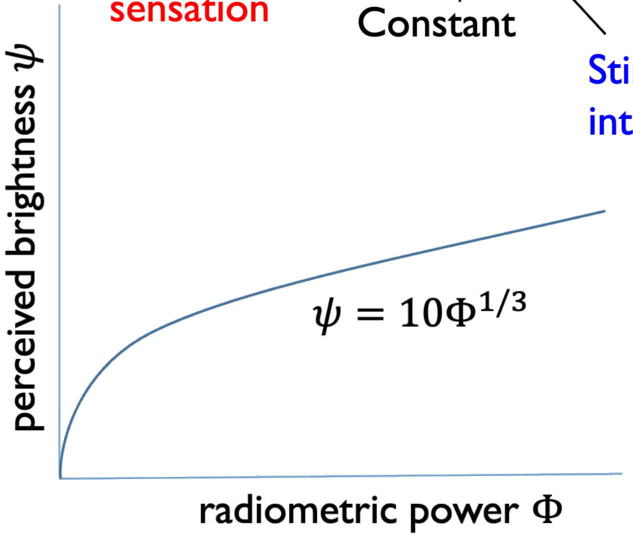

Dr. Stanley Stevens showed that most human sensations follow a power-law relationship between stimuli and sensation.

Tutorial

Digital Image Pipeline

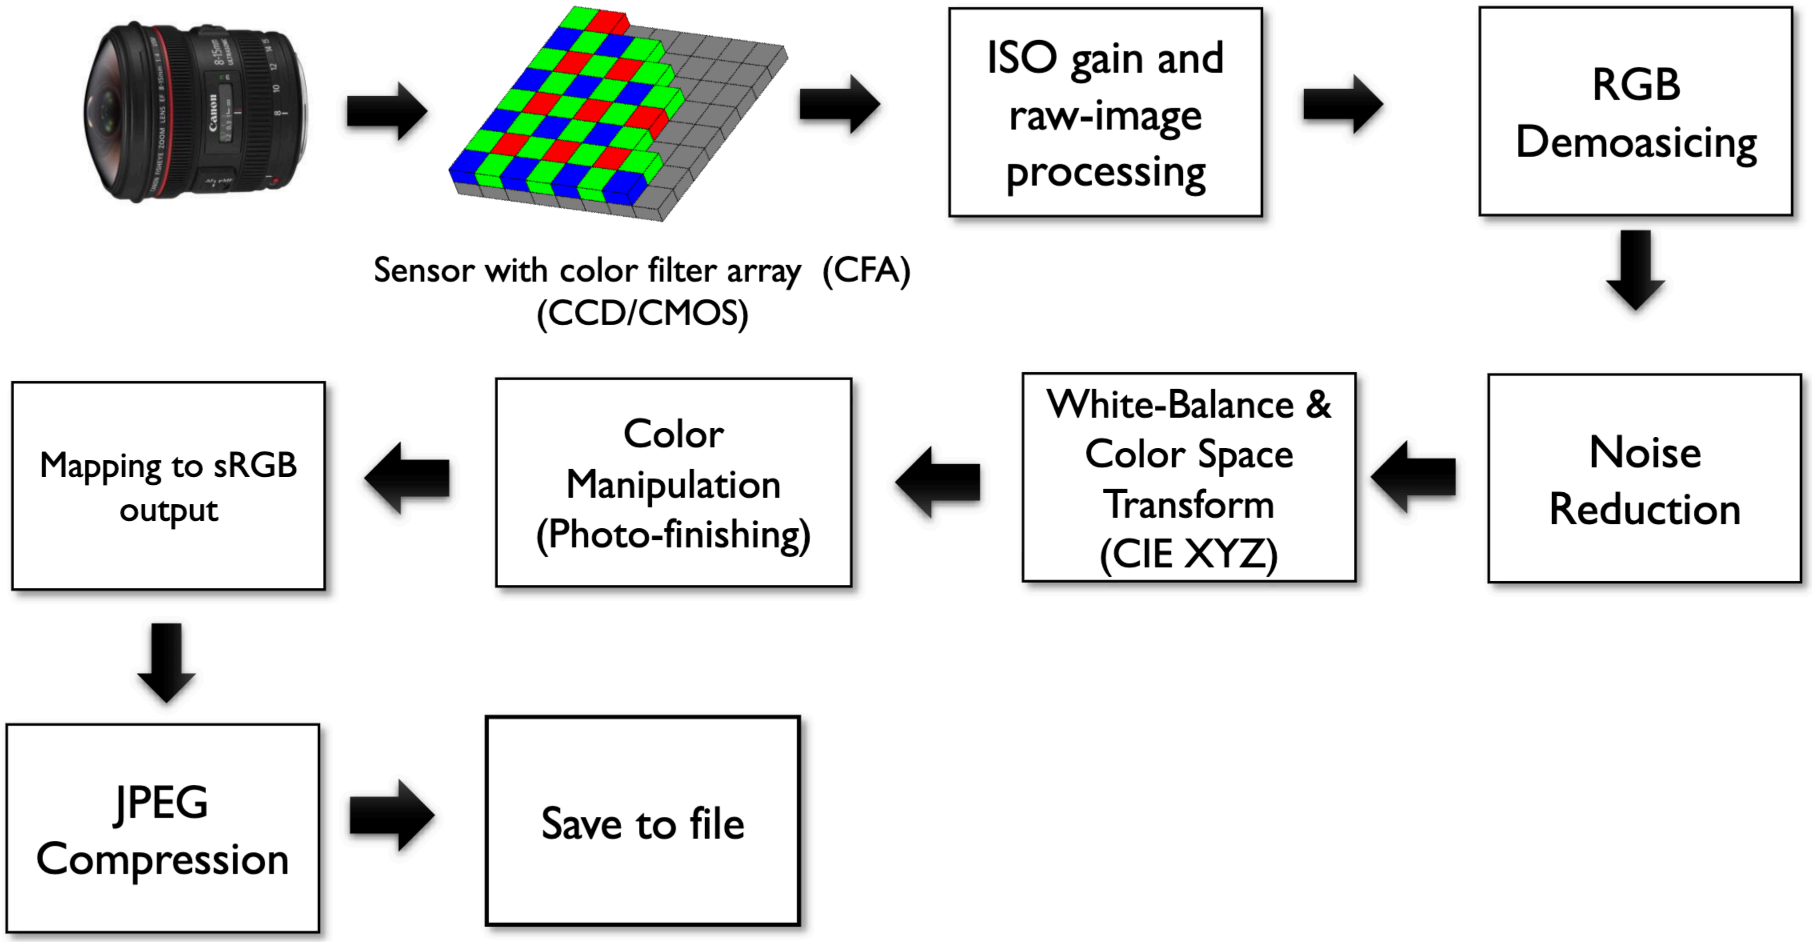

Digital Image Pipeline

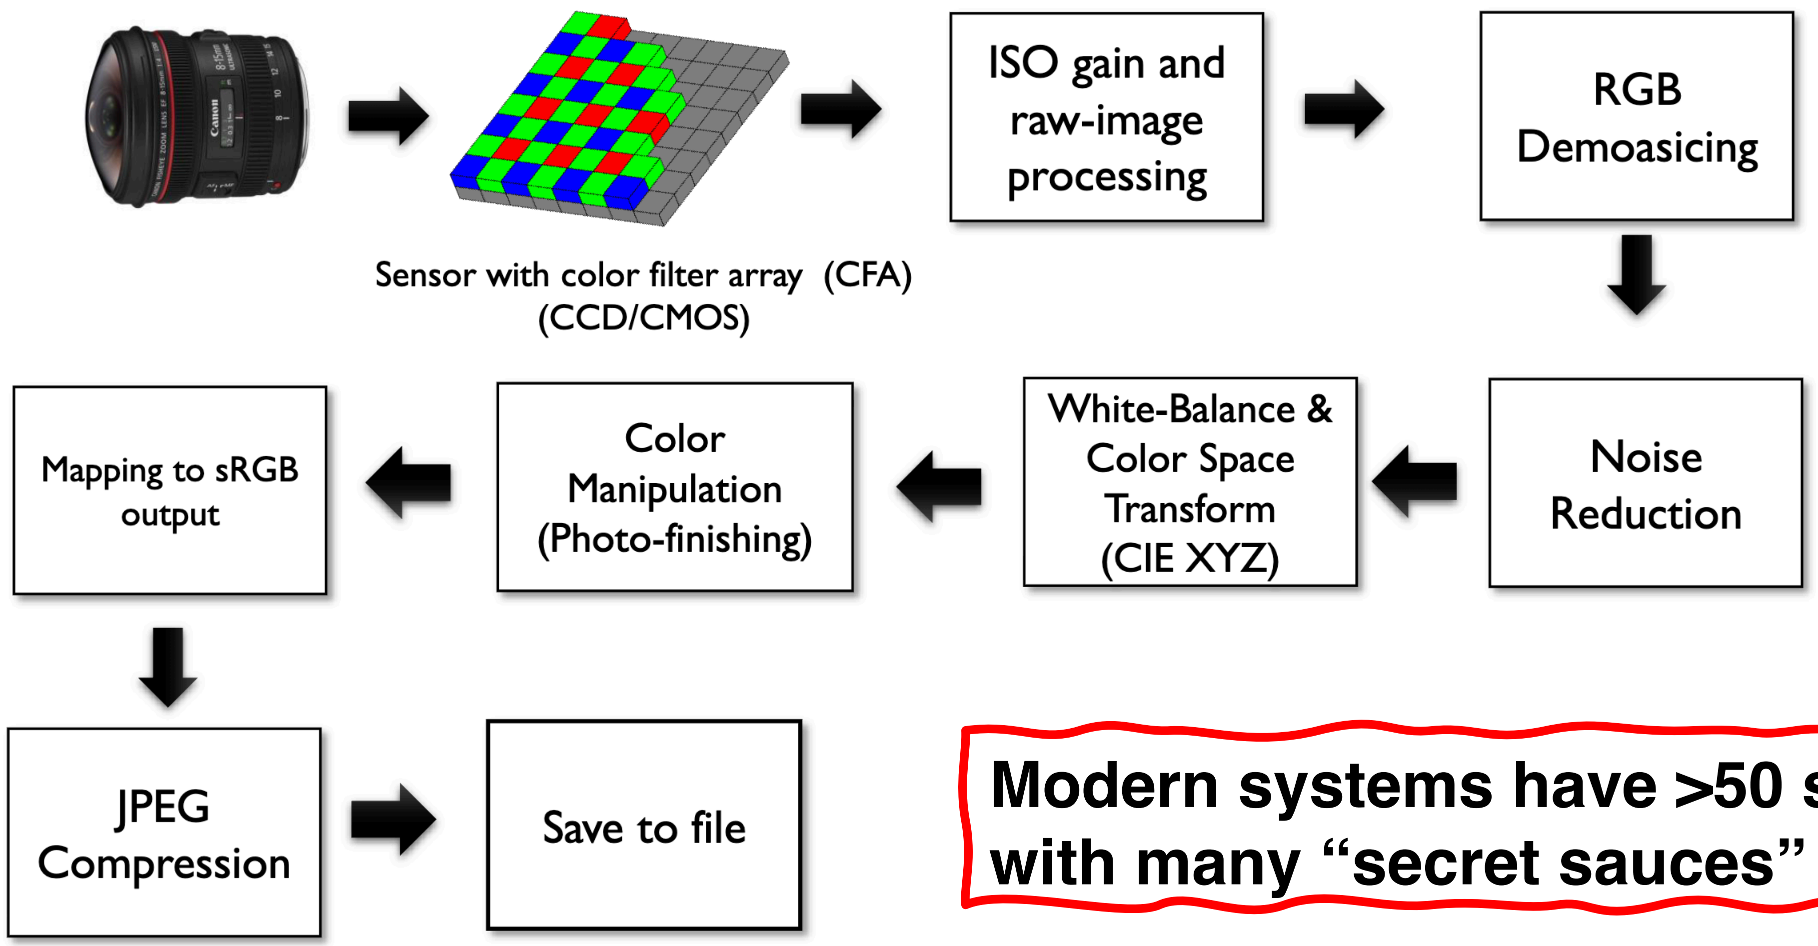

Modern systems have >50 stages with many "secret sauces"

Burst photography

Burst photography

All pipeline elements improve!

- Denoising
- Digital Zoom
- Demosaicing

Burst photography. E.g., Demosaicing

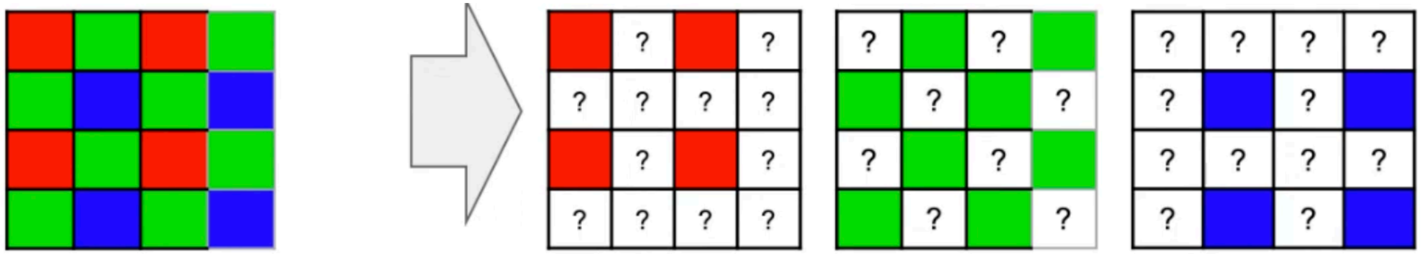

Burst photography. E.g., Demosaicing

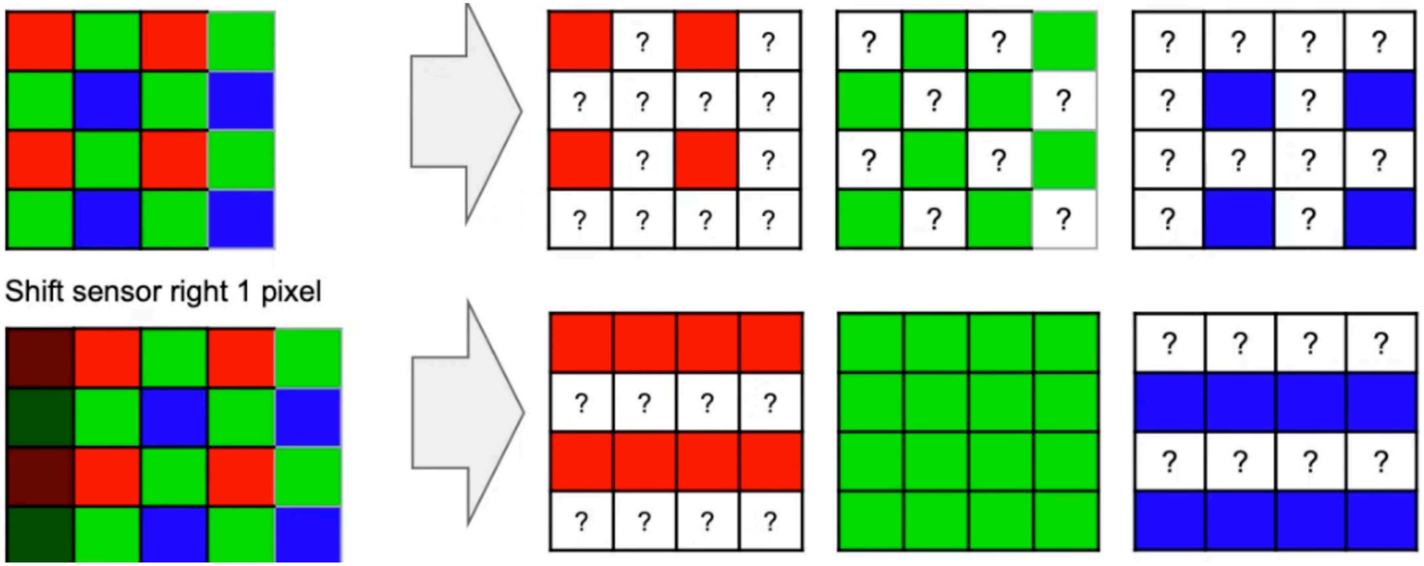

Burst photography. E.g., Demosaicing

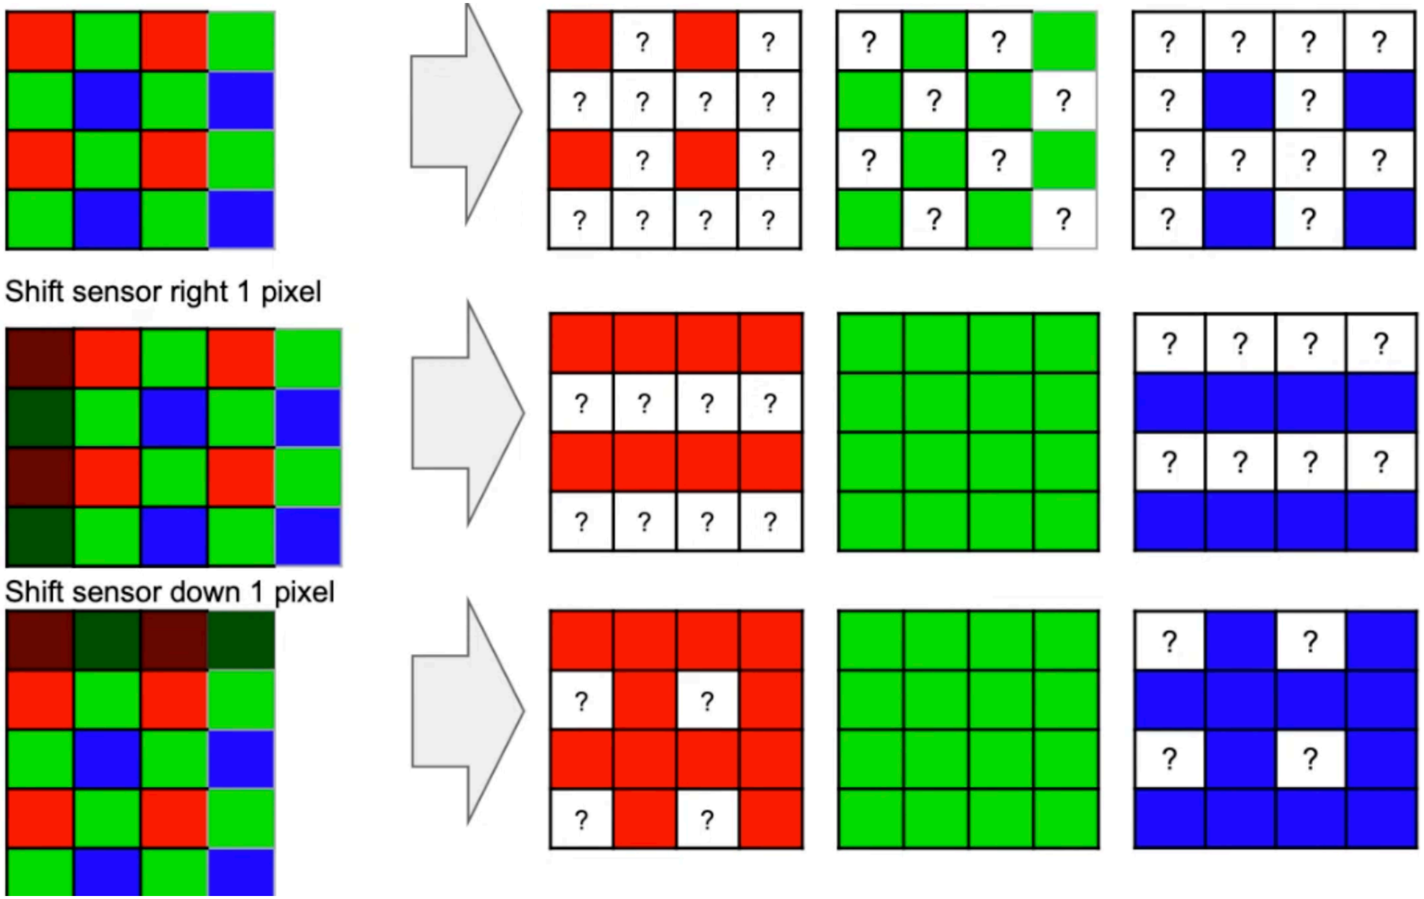

Burst photography. E.g., Demosaicing

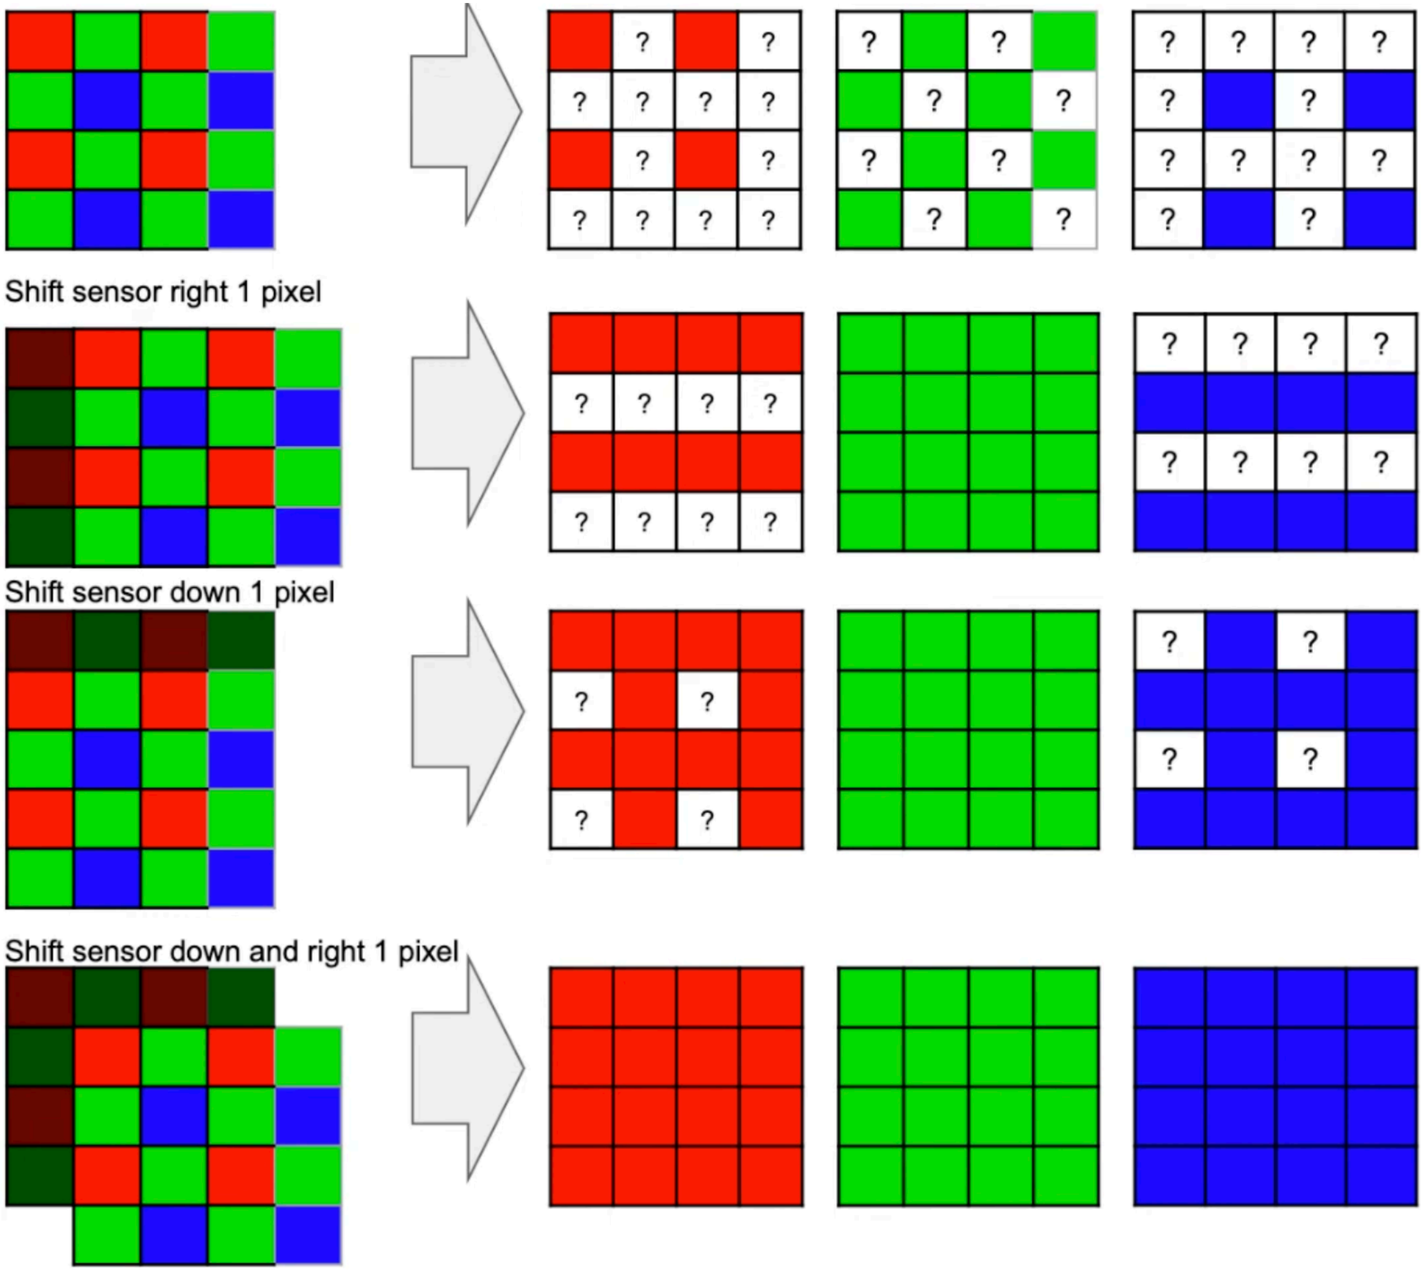

Burst photography. E.g., Demosaicing

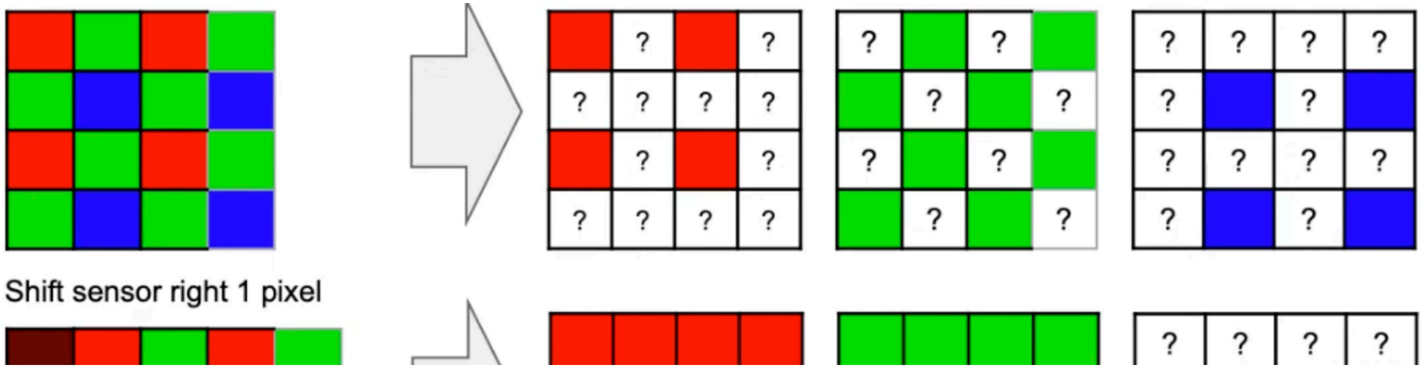

How do we get the motion?

- Hand tremor motion
- Tripod? Google pixel uses lens stabilizer mechanism to induce shifts.
- Our eyes shift too!

Burst photography. E.g., Demosaicing

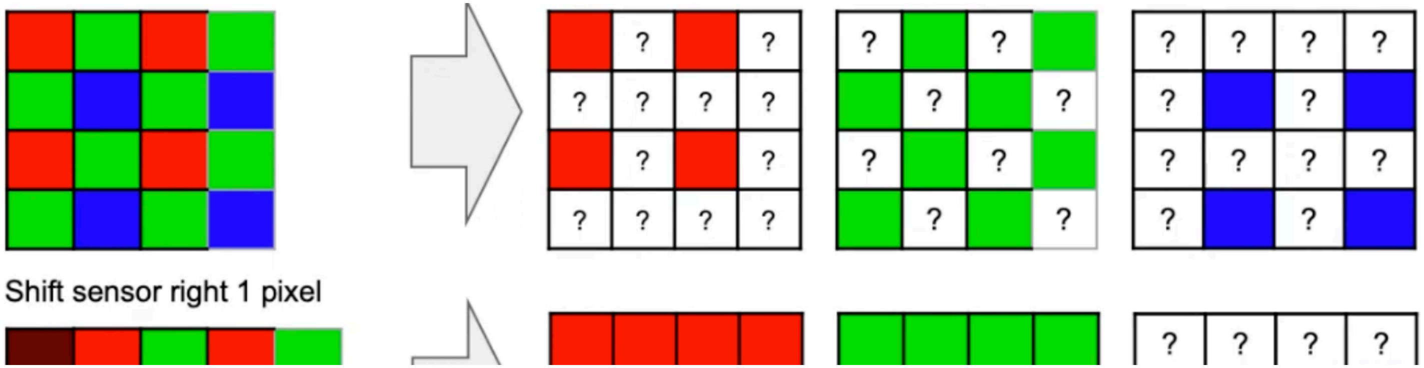

How do we get the motion?

- Hand tremor motion
- Tripod? Google pixel uses lens stabilizer mechanism to induce shifts.
- Our eyes shift too!

That's a wrap!

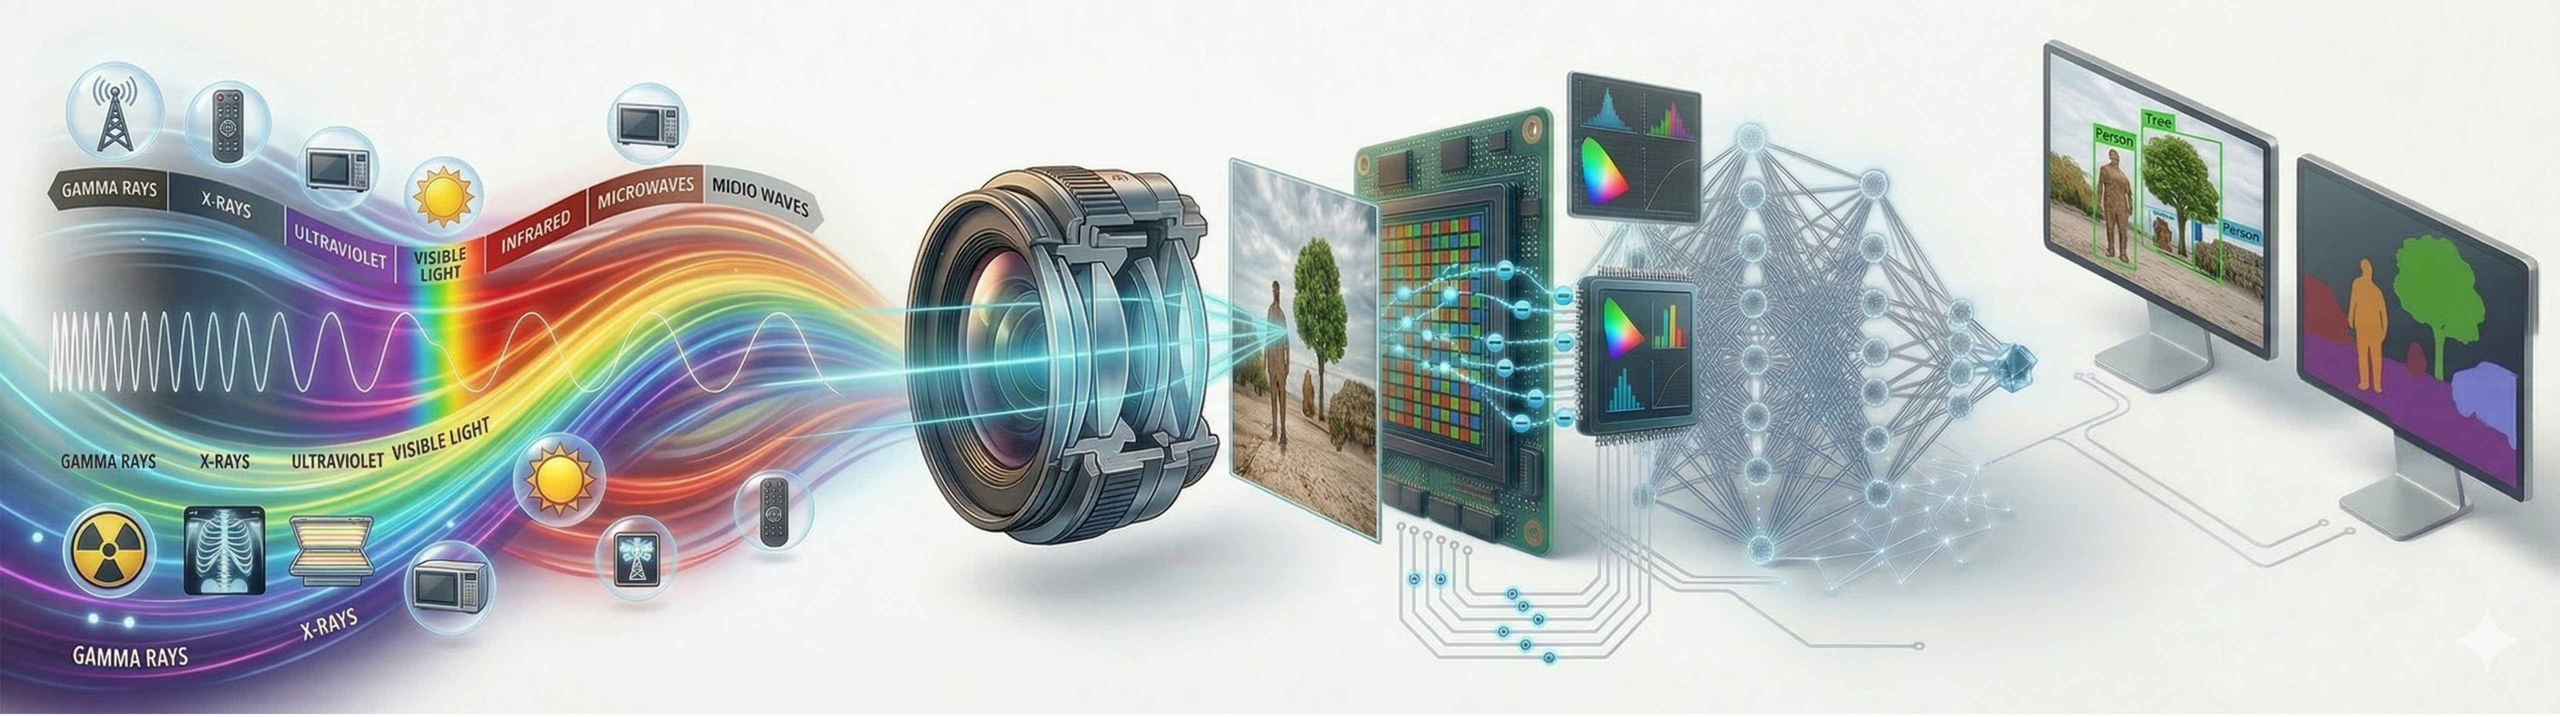
GAMMA RAYS
X-RAYS
ULTRAVIOLET
VISIBLE LIGHT
INFRARED
MICROWAVES
MIDIO WAVES
GAMMA RAYS
X-RAYS
ULTRAVIOLET
VISIBLE LIGHT
X-RAYS
GAMMA RAYS
Person
Tree
Person

That's a wrap!

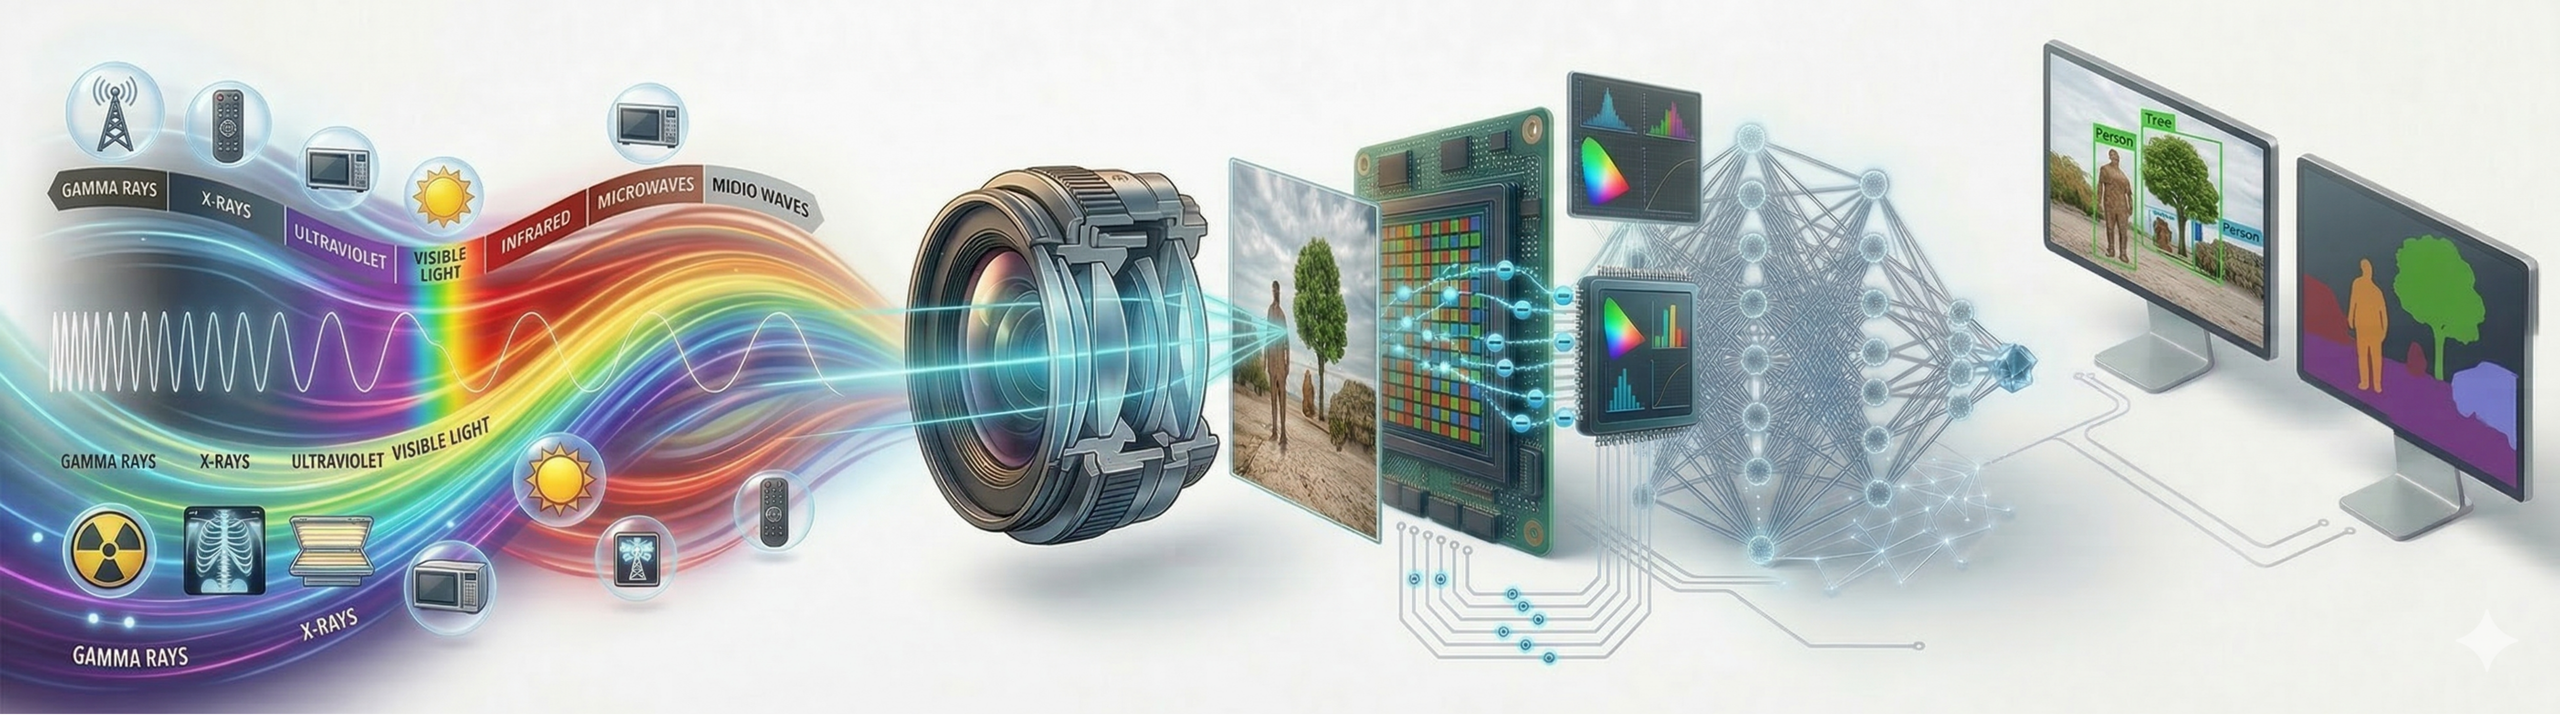

Physics of light

That’s a wrap!

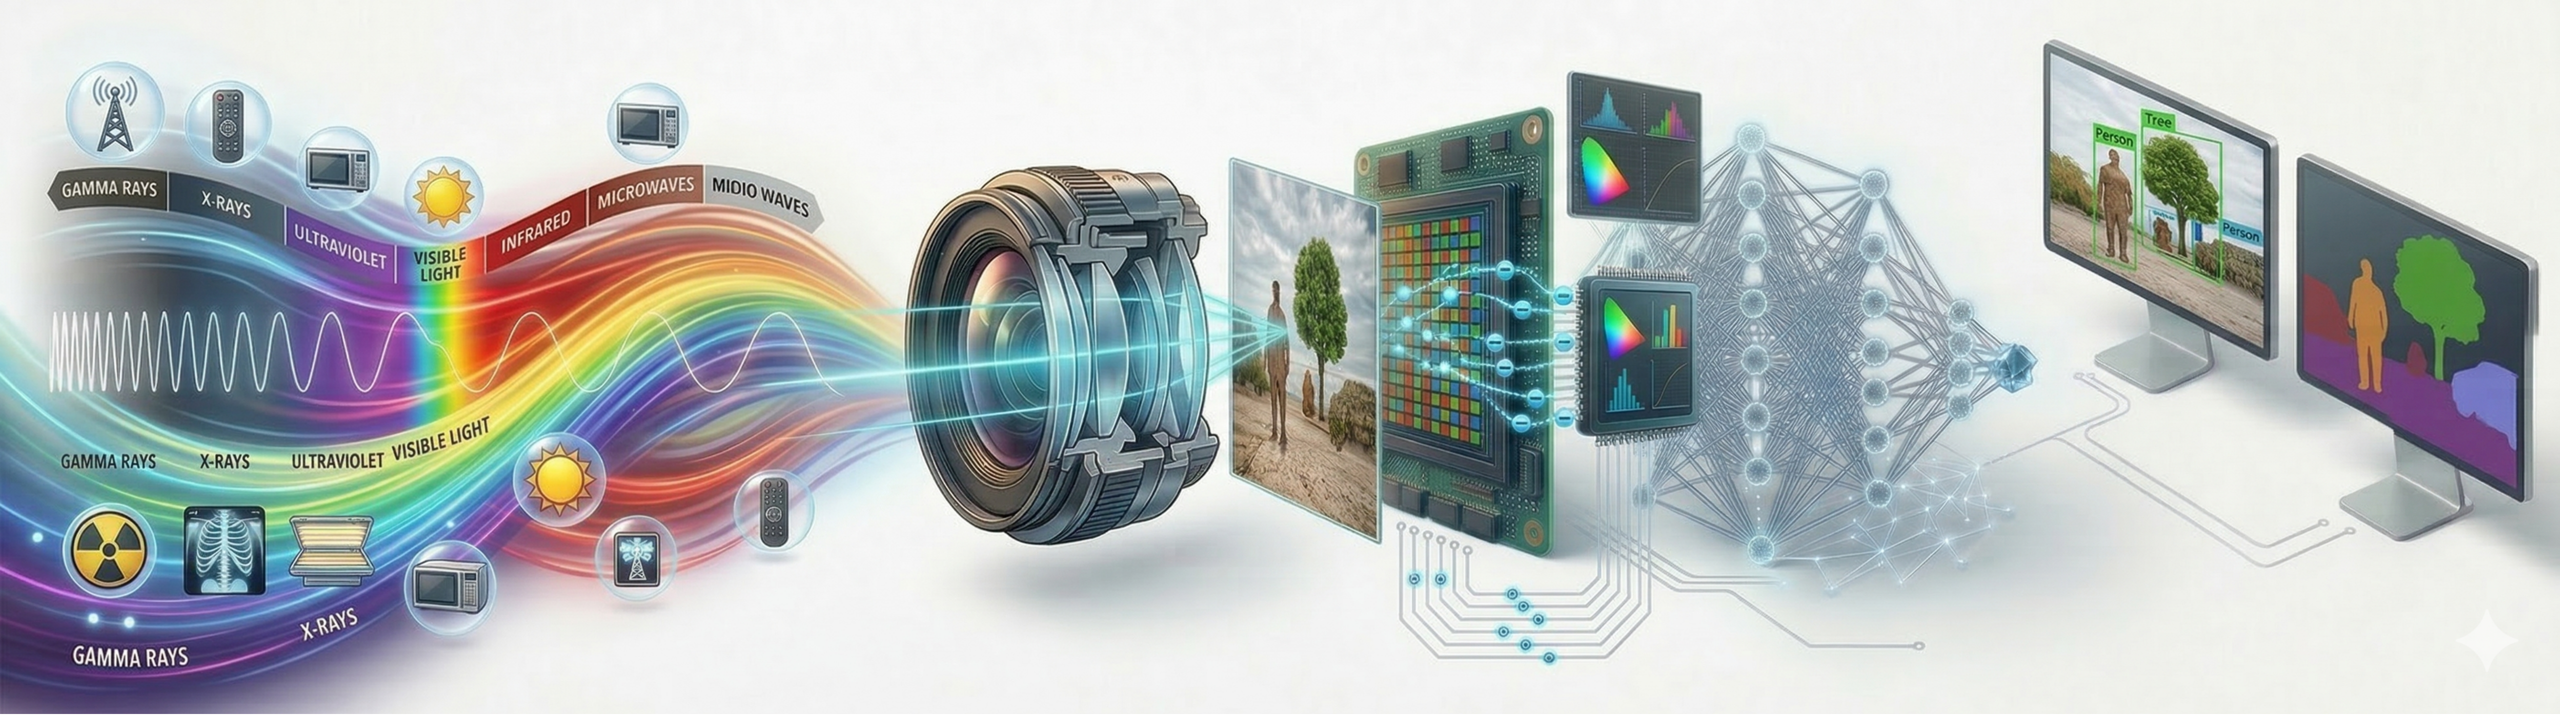

Physics of light

Optics of image formation

That’s a wrap!

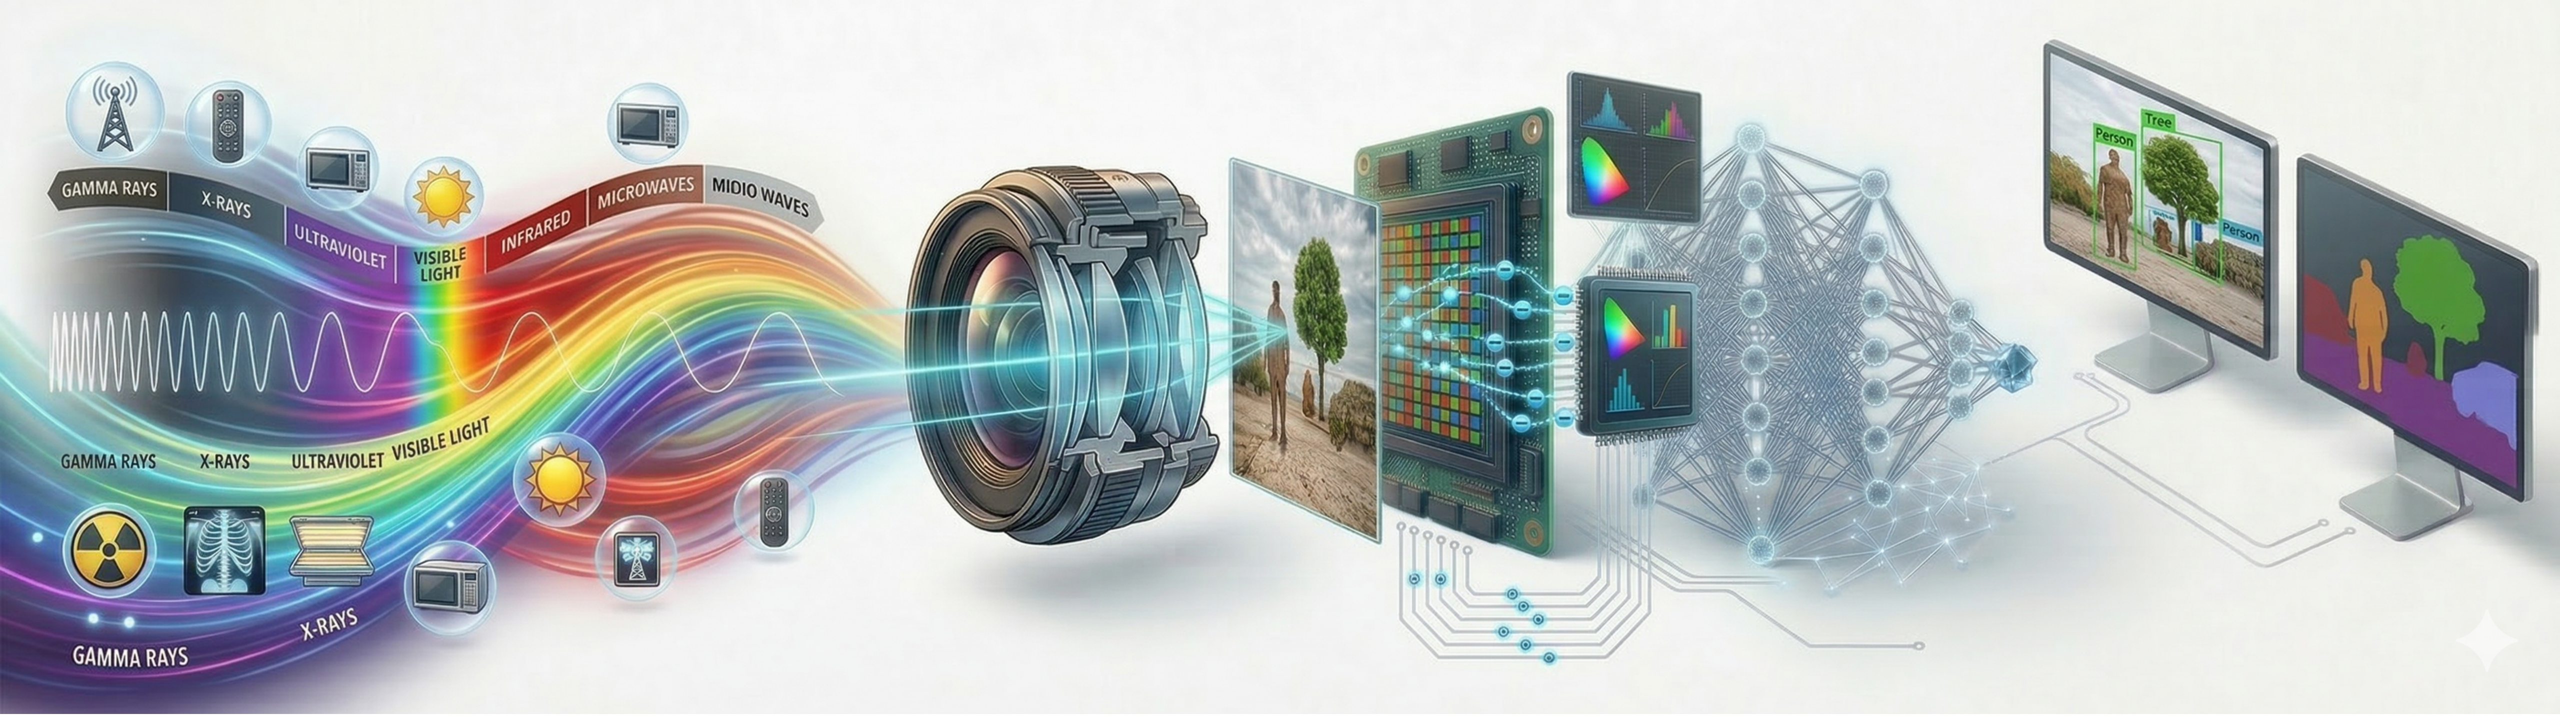

Physics of light

Optics of image formation

Digital image pipeline

That’s a wrap!

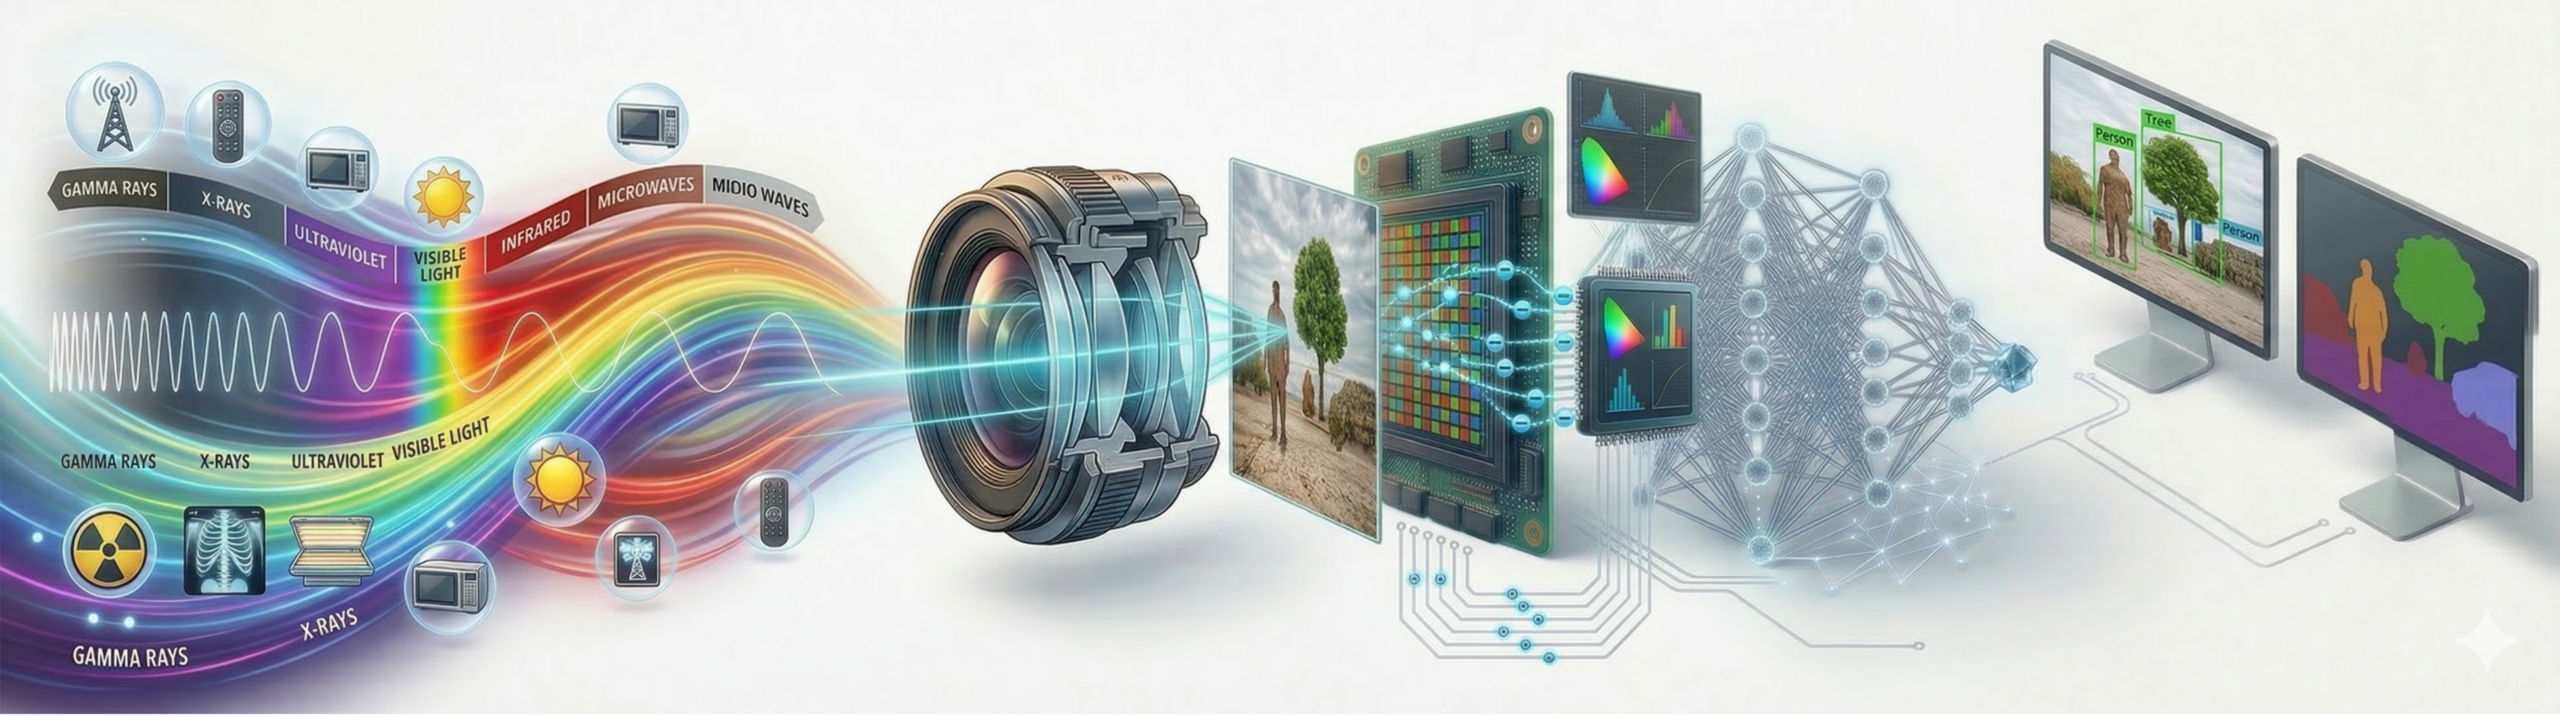

Physics of light

Optics of image formation

Digital image pipeline

Algorithms of computer vision

That's a wrap!

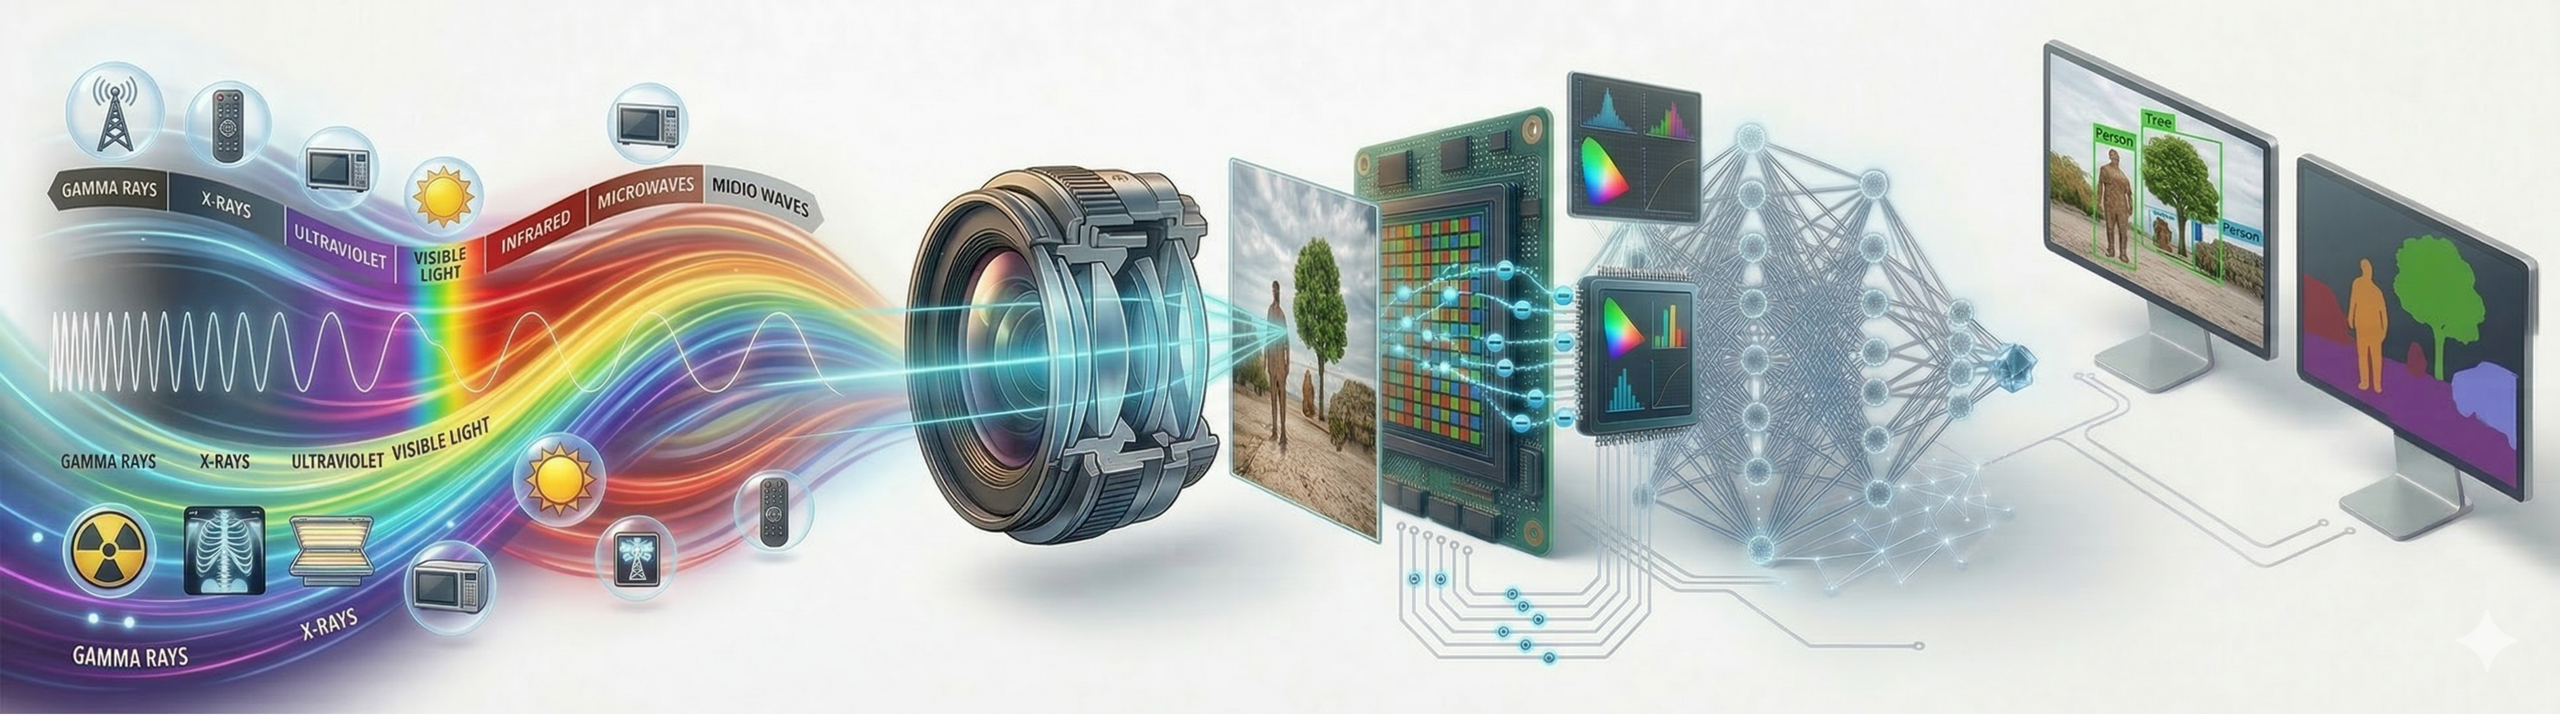

Physics of light

Optics of image formation

Digital image pipeline

Algorithms of computer vision

~~Today~~

That’s a wrap!

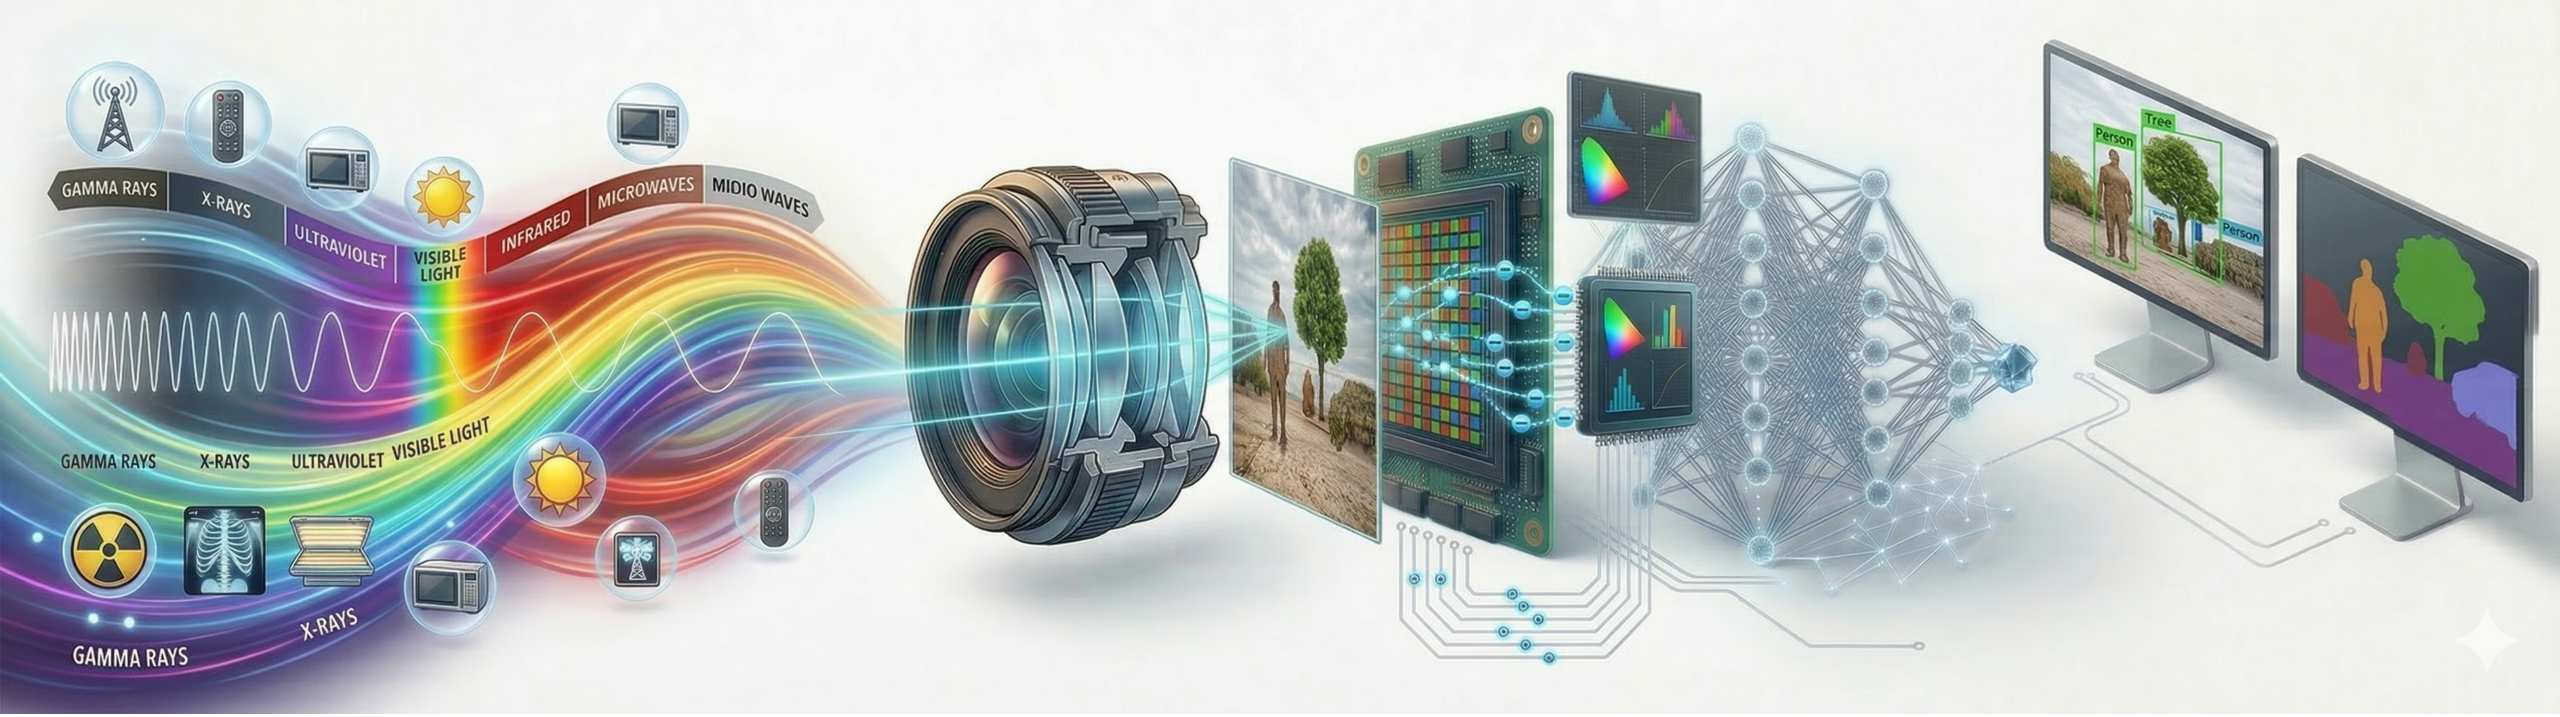

Physics of light | **Optics of image formation** | **Digital image pipeline** | **Algorithms of computer vision**

~~Today~~

Rest of the

Thank you!

A photon checks in to a hotel…

Thank you!

Conclusion from today:

Images are imagined!

A photon checks in to a hotel...

Thank you!

Conclusion from today:

Images are imagined!

Next week:

ML basics from representation learning view

A photon checks in to a hotel…