Digital Image Processing

Ming Jiang

School of Mathematical Sciences
Peking University

ming-jiang@pku.edu.cn

May 18, 2010

Outline

Digital image processing: What, Why and How

What Are the Difficulties

- Poor understanding of human vision system
- Internal representation is not directly understandable
- Why is computer vision difficult?

Image representation and image analysis tasks

Course Overview

Outline

Digital image processing: What, Why and How

What Are the Difficulties

- Poor understanding of human vision system
- Internal representation is not directly understandable
- Why is computer vision difficult?

Image representation and image analysis tasks

Course Overview

Outline

Digital image processing: What, Why and How

What Are the Difficulties

- Poor understanding of human vision system
- Internal representation is not directly understandable
- Why is computer vision difficult?

Image representation and image analysis tasks

Course Overview

Outline

- Digital image processing: What, Why and How
- What Are the Difficulties
  - Poor understanding of human vision system
  - Internal representation is not directly understandable
  - Why is computer vision difficult?
- Image representation and image analysis tasks
- Course Overview

Vision

- Image is better than any other information form for our human being to perceive. Vision allows humans to perceive and understand the world surrounding us.
- Human are primarily visual creatures. Not all animals depend on their eyes, as we do, for 99% or 90% of the information received about the world [Russ, 1995, Zhao and Zhong, 1982].

Vision

- Image is better than any other information form for our human being to perceive. Vision allows humans to perceive and understand the world surrounding us.
- Human are primarily visual creatures. Not all animals depend on their eyes, as we do, for 99% or 90% of the information received about the world [Russ, 1995, Zhao and Zhong, 1982].

Computer Vision

- Computer vision aims to duplicate the effect of human vision by electronically perceiving and understanding an image.
- Books other than this one would dwell at length on this sentence and the meaning of the word *duplicate*
  - whether computer vision is *simulating* or *mimicking* human systems is a philosophical territory,
  - and one very fertile territory, too.

Computer Vision

- Computer vision aims to duplicate the effect of human vision by electronically perceiving and understanding an image.
- Books other than this one would dwell at length on this sentence and the meaning of the word *duplicate*
  - whether computer vision is *simulating* or *mimicking* human systems is a philosophical territory,
  - and one very fertile territory, too.

Computer Vision

- Computer vision aims to duplicate the effect of human vision by electronically perceiving and understanding an image.
- Books other than this one would dwell at length on this sentence and the meaning of the word *duplicate*
  - whether computer vision is *simulating* or *mimicking* human systems is a philosophical territory,
  - and one very fertile territory, too.

Computer Vision

- Computer vision aims to duplicate the effect of human vision by electronically perceiving and understanding an image.
- Books other than this one would dwell at length on this sentence and the meaning of the word *duplicate*
  - whether computer vision is *simulating* or *mimicking* human systems is a philosophical territory,
  - and one very fertile territory, too.

3D vs 2D

- Giving computers the ability to see is not an easy task — we live in a three-dimensional (3D) world.
- When computers try to analyze objects in 3D space, the visual sensors available (e.g., TV cameras) usually give two-dimensional (2D) images.
- This projection from 3D to a lower number of dimensions incurs an enormous loss of information.
- Sometimes, equipment will deliver images that are 3D but this may be of questionable value:
  - analyzing such datasets is clearly more complicated than 2D;
  - sometimes the 'three-dimensionality' is less than intuitive to us;
  - terahertz scans are an example of this.

3D vs 2D

- Giving computers the ability to see is not an easy task — we live in a three-dimensional (3D) world.
- When computers try to analyze objects in 3D space, the visual sensors available (e.g., TV cameras) usually give two-dimensional (2D) images.
- This projection from 3D to a lower number of dimensions incurs an enormous loss of information.
- Sometimes, equipment will deliver images that are 3D but this may be of questionable value:
  - analyzing such datasets is clearly more complicated than 2D;
  - sometimes the 'three-dimensionality' is less than intuitive to us;
  - terahertz scans are an example of this.

3D vs 2D

- Giving computers the ability to see is not an easy task — we live in a three-dimensional (3D) world.
- When computers try to analyze objects in 3D space, the visual sensors available (e.g., TV cameras) usually give two-dimensional (2D) images.
- This projection from 3D to a lower number of dimensions incurs an enormous loss of information.
- Sometimes, equipment will deliver images that are 3D but this may be of questionable value:
  - analyzing such datasets is clearly more complicated than 2D;
  - sometimes the 'three-dimensionality' is less than intuitive to us;
  - terahertz scans are an example of this.

3D vs 2D

- Giving computers the ability to see is not an easy task — we live in a three-dimensional (3D) world.
- When computers try to analyze objects in 3D space, the visual sensors available (e.g., TV cameras) usually give two-dimensional (2D) images.
- This projection from 3D to a lower number of dimensions incurs an enormous loss of information.
- Sometimes, equipment will deliver images that are 3D but this may be of questionable value:
  - analyzing such datasets is clearly more complicated than 2D;
  - sometimes the 'three-dimensionality' is less than intuitive to us;
  - terahertz scans are an example of this.

3D vs 2D

- Giving computers the ability to see is not an easy task — we live in a three-dimensional (3D) world.
- When computers try to analyze objects in 3D space, the visual sensors available (e.g., TV cameras) usually give two-dimensional (2D) images.
- This projection from 3D to a lower number of dimensions incurs an enormous loss of information.
- Sometimes, equipment will deliver images that are 3D but this may be of questionable value:
  - analyzing such datasets is clearly more complicated than 2D;
  - sometimes the 'three-dimensionality' is less than intuitive to us;
  - terahertz scans are an example of this.

3D vs 2D

- Giving computers the ability to see is not an easy task — we live in a three-dimensional (3D) world.
- When computers try to analyze objects in 3D space, the visual sensors available (e.g., TV cameras) usually give two-dimensional (2D) images.
- This projection from 3D to a lower number of dimensions incurs an enormous loss of information.
- Sometimes, equipment will deliver images that are 3D but this may be of questionable value:
  - analyzing such datasets is clearly more complicated than 2D;
  - sometimes the 'three-dimensionality' is less than intuitive to us;
  - terahertz scans are an example of this.

3D vs 2D

- Giving computers the ability to see is not an easy task — we live in a three-dimensional (3D) world.
- When computers try to analyze objects in 3D space, the visual sensors available (e.g., TV cameras) usually give two-dimensional (2D) images.
- This projection from 3D to a lower number of dimensions incurs an enormous loss of information.

- Sometimes, equipment will deliver images that are 3D but this may be of questionable value:
  - analyzing such datasets is clearly more complicated than 2D;
  - sometimes the 'three-dimensionality' is less than intuitive to us;
  - terahertz scans are an example of this.

Video Analysis

- Dynamic scenes such as those to which we are accustomed, with moving objects or a moving camera, are increasingly common and represent another way of making computer vision more complicated.

Figure: A frame from a video of a typical farmyard scene: the cow is one of a number walking naturally from right to left.

Video Analysis: easy for human

- There are many reasons why we might wish to study scenes such as this, which are attractively simple to us
  - the beast is moving slowly;
  - it is clearly black and white;
  - its movement is rhythmic.
- However, automated analysis is very fraught.

Video Analysis: easy for human

- There are many reasons why we might wish to study scenes such as this, which are attractively simple to us
  - the beast is moving slowly;
  - it is clearly black and white;
  - its movement is rhythmic.
- However, automated analysis is very fraught.

Video Analysis: easy for human

- There are many reasons why we might wish to study scenes such as this, which are attractively simple to us
  - the beast is moving slowly;
  - it is clearly black and white;
  - its movement is rhythmic.
- However, automated analysis is very fraught.

Video Analysis: easy for human

- There are many reasons why we might wish to study scenes such as this, which are attractively simple to us
  - the beast is moving slowly;
  - it is clearly black and white;
  - its movement is rhythmic.
- However, automated analysis is very fraught.

Video Analysis: easy for human

- There are many reasons why we might wish to study scenes such as this, which are attractively simple to us
  - the beast is moving slowly;
  - it is clearly black and white;
  - its movement is rhythmic.
- However, automated analysis is very fraught.

Video Analysis: difficult for computer

- The animal's boundary is often very difficult to distinguish clearly from the background;
- the motion of the legs is self occluding;
- (subtly) the concept of *cow-shaped* is not something easily encoded.

Video Analysis: difficult for computer

- The animal's boundary is often very difficult to distinguish clearly from the background;
- the motion of the legs is self occluding;
- (subtly) the concept of *cow-shaped* is not something easily encoded.

Video Analysis: difficult for computer

- The animal's boundary is often very difficult to distinguish clearly from the background;
- the motion of the legs is self occluding;
- (subtly) the concept of *cow-shaped* is not something easily encoded.

Video Analysis: procedures

- The application from which this picture was taken made use of many of the algorithms presented in this book:
  - starting at a low level moving features were identified and grouped;
  - a *training phase* taught the system what a cow might look like in various poses (see the figure on the right), from which a model of a *moving* cow could be estimated.

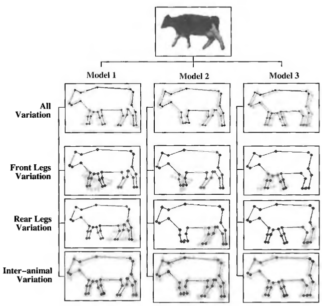

Various models for a cow silhouette: a straight-line boundary approximation has been learned from training data and is able to adapt to different animals and different forms of occlusion.

Video Analysis: procedures

- The application from which this picture was taken made use of many of the algorithms presented in this book:
  - starting at a low level moving features were identified and grouped;
  - a *training phase* taught the system what a cow might look like in various poses (see the figure on the right), from which a model of a *moving* cow could be estimated.

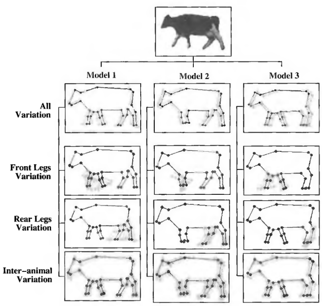

Various models for a cow silhouette: a straight-line boundary approximation has been learned from training data and is able to adapt to different animals and different forms of occlusion.

Video Analysis: procedures

- The application from which this picture was taken made use of many of the algorithms presented in this book:
  - starting at a low level moving features were identified and grouped;
  - a *training phase* taught the system what a cow might look like in various poses (see the figure on the right), from which a model of a *moving* cow could be estimated.

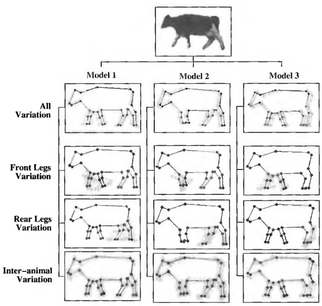

Various models for a cow silhouette: a straight-line boundary approximation has been learned from training data and is able to adapt to different animals and different forms of occlusion.

Video Analysis: procedures

- The application from which this picture was taken made use of many of the algorithms presented in this book:
  - starting at a low level moving features were identified and grouped;
  - a *training phase* taught the system what a cow might look like in various poses (see the figure on the right), from which a model of a *moving* cow could be estimated.

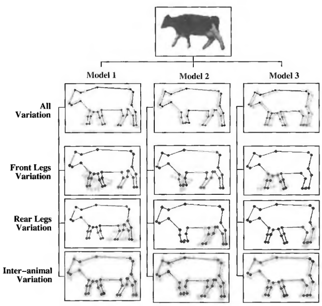

Various models for a cow silhouette: a straight-line boundary approximation has been learned from training data and is able to adapt to different animals and different forms of occlusion.

Video Analysis: procedures

- The application from which this picture was taken made use of many of the algorithms presented in this book:
  - starting at a low level moving features were identified and grouped;
  - a *training phase* taught the system what a cow might look like in various poses (see the figure on the right), from which a model of a *moving* cow could be estimated.

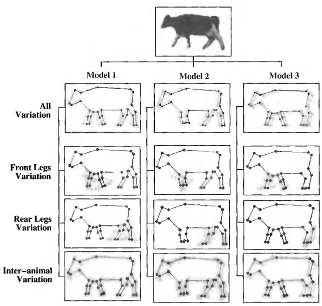

Various models for a cow silhouette: a straight-line boundary approximation has been learned from training data and is able to adapt to different animals and different forms of occlusion.

Video Analysis: procedures

- The application from which this picture was taken made use of many of the algorithms presented in this book:
  - starting at a low level moving features were identified and grouped;
  - a *training phase* taught the system what a cow might look like in various poses (see the figure on the right), from which a model of a *moving* cow could be estimated.

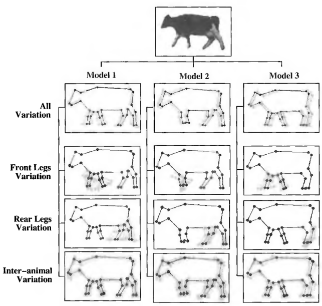

Various models for a cow silhouette: a straight-line boundary approximation has been learned from training data and is able to adapt to different animals and different forms of occlusion.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.

- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.

- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.

- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.

- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.

- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.
- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.

- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.

- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.
- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.
- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.

- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.
- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: operations

- These models could then be fitted to new (*unseen*) video sequences.
- At this stage anomalous behavior such as lameness could be detected by the model failing to fit properly, or well.
- Thus we see a sequence of operations
  - image capture,
  - early processing,
  - segmentation,
  - model fitting,
  - motion prediction,
  - qualitative and/or quantitative conclusion,
- that is characteristic of image understanding and computer vision problems.

Video Analysis: models and cow detection

Figure: Three frames from a cow sequence: notice the model can cope with partial occlusion as the animal enters the scene, and the different poses exhibited.

Video Analysis: disscusions

- Each of these phases (which may not occur sequentially!) may be addressed by a number of algorithms which we shall cover in due course.
- The application was serious; there is a growing need in modern agriculture for automatic monitoring of animal health, for example to spot lameness.
- *A limping cow is trivial for a human to identify, but it is very challenging to do this automatically*.

Video Analysis: disscusions

- Each of these phases (which may not occur sequentially!) may be addressed by a number of algorithms which we shall cover in due course.
- The application was serious; there is a growing need in modern agriculture for automatic monitoring of animal health, for example to spot lameness.
- *A limping cow is trivial for a human to identify, but it is very challenging to do this automatically*.

Video Analysis: disscusions

- Each of these phases (which may not occur sequentially!) may be addressed by a number of algorithms which we shall cover in due course.
- The application was serious; there is a growing need in modern agriculture for automatic monitoring of animal health, for example to spot lameness.
- *A limping cow is trivial for a human to identify, but it is very challenging to do this automatically*.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Video Analysis: disscusions

- This example is relatively simple to explain, but serves to illustrate that many computer vision techniques use the results and methods of
  - mathematics,
  - pattern recognition,
  - artificial intelligence (AI),
  - psycho-physiology,
  - computer science,
  - electronics,
  - and other scientific disciplines.

Difficuties???

- Consider a single gray-scale (monochromatic) image, write down a few reasons why you feel automatic inspection and analysis of it may be difficult.

Outline

Digital image processing: What, Why and How

What Are the Difficulties
- Poor understanding of human vision system
- Internal representation is not directly understandable
- Why is computer vision difficult?

Image representation and image analysis tasks

Course Overview

Human Vision

- How the human perceive process and store the visual information?

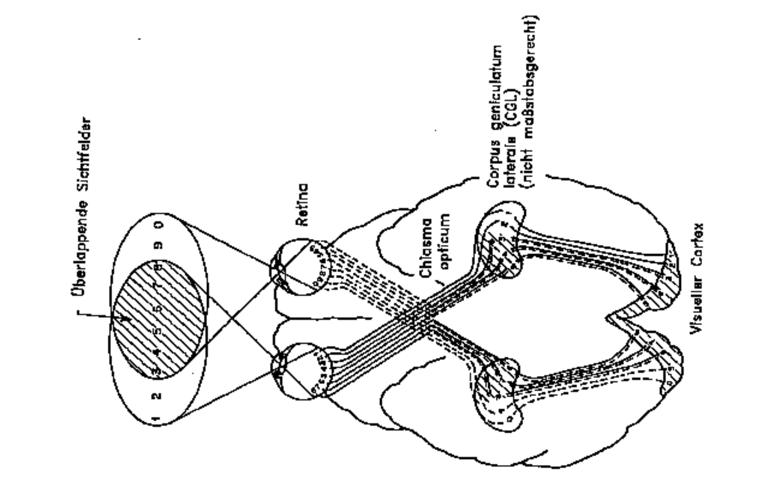

How many legs does this elephant have?

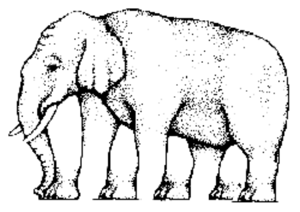

- From the Home of Vision Illusion:
  - http://www.123opticalillusions.com/pages/opticalillusions15.php

How many legs does this elephant have?

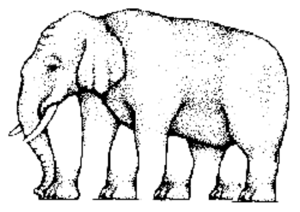

- From the Home of Vision Illusion:
  - http://www.123opticalillusions.com/pages/opticalillusions15.php

What is it?

- Explanation and further information can be found at
  - http://www.psychologie.tu-dresden.de/i1/kaw/diverses%20Material/www.illusionworks.com/html/perceptual_ambiguity.html

What is it?

- Explanation and further information can be found at
  - http://www.psychologie.tu-dresden.de/i1/kaw/diverses%20Material/www.illusionworks.com/html/perceptual_ambiguity.html

Old woman and Young woman: discussions

- Do you see an old woman or a young woman in this illustration?
- They are both present, but you will not be able to see both of them simultaneously.
- Once you perceive both figures, see if you can get them to fluctuate back and forth between the two interpretations.

Old woman and Young woman: discussions

- Do you see an old woman or a young woman in this illustration?
- They are both present, but you will not be able to see both of them simultaneously.
- Once you perceive both figures, see if you can get them to fluctuate back and forth between the two interpretations.

Old woman and Young woman: discussions

- Do you see an old woman or a young woman in this illustration?
- They are both present, but you will not be able to see both of them simultaneously.
- Once you perceive both figures, see if you can get them to fluctuate back and forth between the two interpretations.

Old woman and Young woman: discussions

- This type of reversible figure concerns the meaningful content of what is interpreted by your brain from the same static image.
- Your perception of each figure tends to remain stable until you attend to different regions or contours.
- Certain regions and contours tend to favor one perception, others the alternative.
- Your visual system tends to group like or related regions together.
- It does not present you with some odd mixture of the two alternatives.
- Attending to different regions or contours does tend to initiate a change of perception.

Old woman and Young woman: discussions

- This type of reversible figure concerns the meaningful content of what is interpreted by your brain from the same static image.
- Your perception of each figure tends to remain stable until you attend to different regions or contours.
- Certain regions and contours tend to favor one perception, others the alternative.
- Your visual system tends to group like or related regions together.
- It does not present you with some odd mixture of the two alternatives.
- Attending to different regions or contours does tend to initiate a change of perception.

Old woman and Young woman: discussions

- This type of reversible figure concerns the meaningful content of what is interpreted by your brain from the same static image.
- Your perception of each figure tends to remain stable until you attend to different regions or contours.
- Certain regions and contours tend to favor one perception, others the alternative.
- Your visual system tends to group like or related regions together.
- It does not present you with some odd mixture of the two alternatives.
- Attending to different regions or contours does tend to initiate a change of perception.

Old woman and Young woman: discussions

- This type of reversible figure concerns the meaningful content of what is interpreted by your brain from the same static image.
- Your perception of each figure tends to remain stable until you attend to different regions or contours.
- Certain regions and contours tend to favor one perception, others the alternative.
- Your visual system tends to group like or related regions together.
- It does not present you with some odd mixture of the two alternatives.
- Attending to different regions or contours does tend to initiate a change of perception.

Old woman and Young woman: discussions

- This type of reversible figure concerns the meaningful content of what is interpreted by your brain from the same static image.
- Your perception of each figure tends to remain stable until you attend to different regions or contours.
- Certain regions and contours tend to favor one perception, others the alternative.
- Your visual system tends to group like or related regions together.
- It does not present you with some odd mixture of the two alternatives.
- Attending to different regions or contours does tend to initiate a change of perception.

Old woman and Young woman: discussions

- This type of reversible figure concerns the meaningful content of what is interpreted by your brain from the same static image.
- Your perception of each figure tends to remain stable until you attend to different regions or contours.
- Certain regions and contours tend to favor one perception, others the alternative.
- Your visual system tends to group like or related regions together.
- It does not present you with some odd mixture of the two alternatives.
- Attending to different regions or contours does tend to initiate a change of perception.

Human Vision

- We do not have a clear understanding how the human perceive, process and store the visual information.
- We do not even know how the human measures internally the image visual quality and discrimination.

Human Vision

- We do not have a clear understanding how the human perceive, process and store the visual information.
- We do not even know how the human measures internally the image visual quality and discrimination.

Perception ≡ Description

- If this image is looked at with a steady eye, it will still change, though less often.
- Researchers have stabilized the image directly onto the retina to eliminate any effects that may arise from eye movements.
- Even under these conditions, a perceptual reversal may occur.
- This indicates that higher cortical processing occurs that strives to make meaning out of a stable image presented to the retina.
- This illustrates once more that vision is an active process that attempts to make sense of incoming information.
- As the late David Marr said, "*Perception is the construction of a description*."

Perception $\equiv$ Description

- If this image is looked at with a steady eye, it will still change, though less often.
- Researchers have stabilized the image directly onto the retina to eliminate any effects that may arise from eye movements.
- Even under these conditions, a perceptual reversal may occur.
- This indicates that higher cortical processing occurs that strives to make meaning out of a stable image presented to the retina.
- This illustrates once more that vision is an active process that attempts to make sense of incoming information.
- As the late David Marr said, "*Perception is the construction of a description.*"

Perception ≡ Description

- If this image is looked at with a steady eye, it will still change, though less often.
- Researchers have stabilized the image directly onto the retina to eliminate any effects that may arise from eye movements.
- Even under these conditions, a perceptual reversal may occur.
- This indicates that higher cortical processing occurs that strives to make meaning out of a stable image presented to the retina.
- This illustrates once more that vision is an active process that attempts to make sense of incoming information.
- As the late David Marr said, "*Perception is the construction of a description.*"

Perception ≡ Description

- If this image is looked at with a steady eye, it will still change, though less often.
- Researchers have stabilized the image directly onto the retina to eliminate any effects that may arise from eye movements.
- Even under these conditions, a perceptual reversal may occur.
- This indicates that higher cortical processing occurs that strives to make meaning out of a stable image presented to the retina.
- This illustrates once more that vision is an active process that attempts to make sense of incoming information.
- As the late David Marr said, "*Perception is the construction of a description.*"

Perception ≡ Description

- If this image is looked at with a steady eye, it will still change, though less often.
- Researchers have stabilized the image directly onto the retina to eliminate any effects that may arise from eye movements.
- Even under these conditions, a perceptual reversal may occur.
- This indicates that higher cortical processing occurs that strives to make meaning out of a stable image presented to the retina.
- This illustrates once more that vision is an active process that attempts to make sense of incoming information.
- As the late David Marr said, "*Perception is the construction of a description.*"

Perception $\equiv$ Description

- If this image is looked at with a steady eye, it will still change, though less often.
- Researchers have stabilized the image directly onto the retina to eliminate any effects that may arise from eye movements.
- Even under these conditions, a perceptual reversal may occur.

- This indicates that higher cortical processing occurs that strives to make meaning out of a stable image presented to the retina.
- This illustrates once more that vision is an active process that attempts to make sense of incoming information.
- As the late David Marr said, "*Perception is the construction of a description.*"

History of this illustration

- For many years the creator of this famous figure was thought to be British cartoonist W. E. Hill, who published it in 1915. Hill almost certainly adapted the figure from an original concept that was popular throughout the world on trading and puzzle cards.
- This anonymous dated German postcard (shown at the top of the page) from 1888 depicts the image in its earliest known form.

The 1890 example on the left shows quite clearly its association as “My Wife and Mother-in-Law.” Both of these examples predate the Punch cartoon that was previously thought to serve as the figure’s inspiration.

History of this illustration

- For many years the creator of this famous figure was thought to be British cartoonist W. E. Hill, who published it in 1915. Hill almost certainly adapted the figure from an original concept that was popular throughout the world on trading and puzzle cards.
- This anonymous dated German postcard (shown at the top of the page) from 1888 depicts the image in its earliest known form.

The 1890 example on the left shows quite clearly its association as “My Wife and Mother-in-Law.” Both of these examples predate the Punch cartoon that was previously thought to serve as the figure’s inspiration.

History of this illustration

- For many years the creator of this famous figure was thought to be British cartoonist W. E. Hill, who published it in 1915. Hill almost certainly adapted the figure from an original concept that was popular throughout the world on trading and puzzle cards.
- This anonymous dated German postcard (shown at the top of the page) from 1888 depicts the image in its earliest known form.

The 1890 example on the left shows quite clearly its association as “My Wife and Mother-in-Law.” Both of these examples predate the Punch cartoon that was previously thought to serve as the figure’s inspiration.

History of this illustration

- For many years the creator of this famous figure was thought to be British cartoonist W. E. Hill, who published it in 1915. Hill almost certainly adapted the figure from an original concept that was popular throughout the world on trading and puzzle cards.
- This anonymous dated German postcard (shown at the top of the page) from 1888 depicts the image in its earliest known form.

The 1890 example on the left shows quite clearly its association as “My Wife and Mother-in-Law.” Both of these examples predate the Punch cartoon that was previously thought to serve as the figure’s inspiration.

History of this illustration

- For many years the creator of this famous figure was thought to be British cartoonist W. E. Hill, who published it in 1915. Hill almost certainly adapted the figure from an original concept that was popular throughout the world on trading and puzzle cards.
- This anonymous dated German postcard (shown at the top of the page) from 1888 depicts the image in its earliest known form.

The 1890 example on the left shows quite clearly its association as “My Wife and Mother-in-Law.” Both of these examples predate the Punch cartoon that was previously thought to serve as the figure’s inspiration.

History of this illustration

- The figure was later altered and adapted by others, including the two psychologists, R. W. Leeper and E. G. Boring who described the figure and made it famous within psychological circles in 1930. It has often been referred to as the “Boring figure.”
- Versions of the figure proved to be popular and the image was frequently reprinted; however, perceptual biases started to occur in the image, unbeknownst to the plagiarizing artists and psychologists who were reprinting the images.
- Variations have appeared in the literature that unintentionally are biased to favor one interpretation or another, which defeats its original purpose as a truly ambiguous figure.

History of this illustration

- The figure was later altered and adapted by others, including the two psychologists, R. W. Leeper and E. G. Boring who described the figure and made it famous within psychological circles in 1930. It has often been referred to as the "Boring figure."
- Versions of the figure proved to be popular and the image was frequently reprinted; however, perceptual biases started to occur in the image, unbeknownst to the plagiarizing artists and psychologists who were reprinting the images.
- Variations have appeared in the literature that unintentionally are biased to favor one interpretation or another, which defeats its original purpose as a truly ambiguous figure.

History of this illustration

- The figure was later altered and adapted by others, including the two psychologists, R. W. Leeper and E. G. Boring who described the figure and made it famous within psychological circles in 1930. It has often been referred to as the “Boring figure.”
- Versions of the figure proved to be popular and the image was frequently reprinted; however, perceptual biases started to occur in the image, unbeknownst to the plagiarizing artists and psychologists who were reprinting the images.
- Variations have appeared in the literature that unintentionally are biased to favor one interpretation or another, which defeats its original purpose as a truly ambiguous figure.

History of this illustration

- In the three versions shown above, can you tell which one is biased toward the young girl, the old woman?

History of this illustration

- In 1961, J, Botwinick redesigned this figure once again, and entitled it, "Husband and Father-in-Law."

Outline

Digital image processing: What, Why and How

What Are the Difficulties

- Poor understanding of human vision system
- Internal representation is not directly understandable
- Why is computer vision difficult?

Image representation and image analysis tasks

Course Overview

Images as 2D functions

- Images are usually represented as a two dimensional function.
- Digitized images are usually represented by two dimensional array.
- However, those representations are not suitable for machine understanding, while the computer is able to process those representations.
- General knowledge, domain-specific knowledge, and information extracted from the image will be essential in attempting to **understanding** those arrays of numbers.

Images as 2D functions

- Images are usually represented as a two dimensional function.
- Digitized images are usually represented by two dimensional array.
- However, those representations are not suitable for machine understanding, while the computer is able to process those representations.
- General knowledge, domain-specific knowledge, and information extracted from the image will be essential in attempting to **understanding** those arrays of numbers.

Images as 2D functions

- Images are usually represented as a two dimensional function.
- Digitized images are usually represented by two dimensional array.
- However, those representations are not suitable for machine understanding, while the computer is able to process those representations.
- General knowledge, domain-specific knowledge, and information extracted from the image will be essential in attempting to **understanding** those arrays of numbers.

Images as 2D functions

- Images are usually represented as a two dimensional function.
- Digitized images are usually represented by two dimensional array.
- However, those representations are not suitable for machine understanding, while the computer is able to process those representations.
- General knowledge, domain-specific knowledge, and information extracted from the image will be essential in attempting to **understanding** those arrays of numbers.

Experiment with images as 2D functions

- Read and display a image file as a two dimensional function.
- The example matlab script file is here matlab display example.

Experiment with images as 2D functions

- Read and display a image file as a two dimensional function.
- The example matlab script file is here matlab display example.

Images as 2D functions: discussions

- Both presentations contain exactly the same information.
- But for a human observer it is very difficult to find a correspondence between both.
- The point is that a lot of a priori knowledge is used by humans to interpret the images;
- the machine only begins with an array of numbers and so will be attempting to us more likely the first display than the second display.

Images as 2D functions: discussions

- Both presentations contain exactly the same information.
- But for a human observer it is very difficult to find a correspondence between both.
- The point is that a lot of a priori knowledge is used by humans to interpret the images;
- the machine only begins with an array of numbers and so will be attempting to us more likely the first display than the second display.

Images as 2D functions: discussions

- Both presentations contain exactly the same information.
- But for a human observer it is very difficult to find a correspondence between both.
- The point is that a lot of a priori knowledge is used by humans to interpret the images;
- the machine only begins with an array of numbers and so will be attempting to us more likely the first display than the second display.

Images as 2D functions: discussions

- Both presentations contain exactly the same information.
- But for a human observer it is very difficult to find a correspondence between both.
- The point is that a lot of a priori knowledge is used by humans to interpret the images;
- the machine only begins with an array of numbers and so will be attempting to us more likely the first display than the second display.

Outline

Digital image processing: What, Why and How

What Are the Difficulties

- Poor understanding of human vision system
- Internal representation is not directly understandable
- Why is computer vision difficult?

Image representation and image analysis tasks

Course Overview

Why is computer vision difficult?

- This philosophical question provides some insight into the rather complex landscape of computer vision.
- It can be answered in many ways: we offer six.
- Here, we mention the reasons only briefly — most of them will be discussed in more detail later in the book.

Why is computer vision difficult?

- This philosophical question provides some insight into the rather complex landscape of computer vision.
- It can be answered in many ways: we offer six.
- Here, we mention the reasons only briefly — most of them will be discussed in more detail later in the book.

Why is computer vision difficult?

- This philosophical question provides some insight into the rather complex landscape of computer vision.
- It can be answered in many ways: we offer six.
- Here, we mention the reasons only briefly — most of them will be discussed in more detail later in the book.

I. Loss of information

- Loss of information in projections from 3D to 2D is a phenomenon which occurs in typical image capture devices such as a camera or an eye.
- Their geometric properties have been approximated by a pinhole model for centuries (a box with a small hole in it, called in Latin a `camera obscura` [dark room]).

I. Loss of information

- Loss of information in projections from 3D to 2D is a phenomenon which occurs in typical image capture devices such as a camera or an eye.
- Their geometric properties have been approximated by a pinhole model for centuries (a box with a small hole in it, called in Latin a `camera obscura` [dark room]).

Pinhole camera

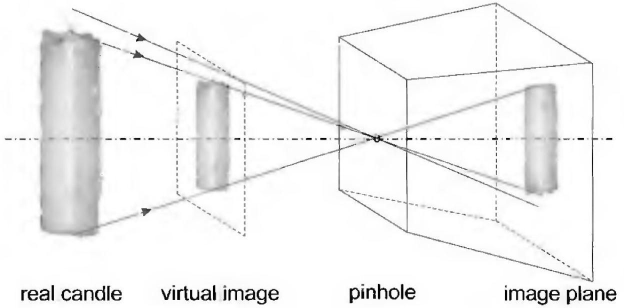

Figure: The pinhole model of imaging geometry does not distinguish size of objects.

- This physical model corresponds to a mathematical model of perspective projection.
- The projective transformation maps points along rays but does not preserve angles and collinearity.

Pinhole camera

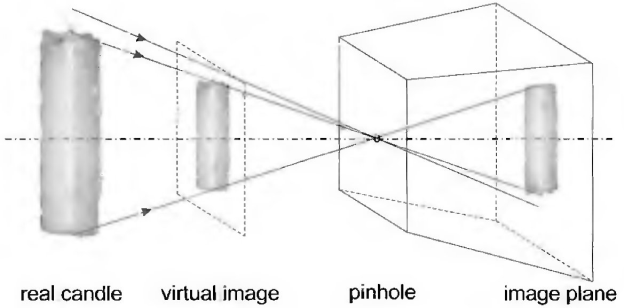

Figure: The pinhole model of imaging geometry does not distinguish size of objects.

- This physical model corresponds to a mathematical model of perspective projection.
- The projective transformation maps points along rays but does not preserve angles and collinearity.

II. Interpretation of images

- Interpretation of image(s) constitutes the principal tool of computer vision to approach problems which humans solve unwittingly.
- When a human tries to understand an image then previous knowledge and experience is brought to the current observation.
- Human ability to reason allows representation of long-gathered knowledge, and its use to solve new problems.
- Artificial intelligence has invested several decades in attempts to endow computers with the capability to understand observations;
- while progress has been tremendous, the practical ability of a machine to understand observations remains very limited.
- Attempting to solve related multidisciplinary scientific problems under the name cognitive systems is seen as a key to developing intelligent machines.

II. Interpretation of images

- Interpretation of image(s) constitutes the principal tool of computer vision to approach problems which humans solve unwittingly.
- When a human tries to understand an image then previous knowledge and experience is brought to the current observation.
- Human ability to reason allows representation of long-gathered knowledge, and its use to solve new problems.
- Artificial intelligence has invested several decades in attempts to endow computers with the capability to understand observations;
- while progress has been tremendous, the practical ability of a machine to understand observations remains very limited.
- Attempting to solve related multidisciplinary scientific problems under the name cognitive systems is seen as a key to developing intelligent machines.

II. Interpretation of images

- Interpretation of image(s) constitutes the principal tool of computer vision to approach problems which humans solve unwittingly.
- When a human tries to understand an image then previous knowledge and experience is brought to the current observation.
- Human ability to reason allows representation of long-gathered knowledge, and its use to solve new problems.
- Artificial intelligence has invested several decades in attempts to endow computers with the capability to understand observations;
- while progress has been tremendous, the practical ability of a machine to understand observations remains very limited.
- Attempting to solve related multidisciplinary scientific problems under the name cognitive systems is seen as a key to developing intelligent machines.

II. Interpretation of images

- Interpretation of image(s) constitutes the principal tool of computer vision to approach problems which humans solve unwittingly.
- When a human tries to understand an image then previous knowledge and experience is brought to the current observation.
- Human ability to reason allows representation of long-gathered knowledge, and its use to solve new problems.
- Artificial intelligence has invested several decades in attempts to endow computers with the capability to understand observations;
- while progress has been tremendous, the practical ability of a machine to understand observations remains very limited.
- Attempting to solve related multidisciplinary scientific problems under the name cognitive systems is seen as a key to developing intelligent machines.

II. Interpretation of images

- Interpretation of image(s) constitutes the principal tool of computer vision to approach problems which humans solve unwittingly.
- When a human tries to understand an image then previous knowledge and experience is brought to the current observation.
- Human ability to reason allows representation of long-gathered knowledge, and its use to solve new problems.
- Artificial intelligence has invested several decades in attempts to endow computers with the capability to understand observations;
- while progress has been tremendous, the practical ability of a machine to understand observations remains very limited.
- Attempting to solve related multidisciplinary scientific problems under the name cognitive systems is seen as a key to developing intelligent machines.

II. Interpretation of images

- Interpretation of image(s) constitutes the principal tool of computer vision to approach problems which humans solve unwittingly.
- When a human tries to understand an image then previous knowledge and experience is brought to the current observation.
- Human ability to reason allows representation of long-gathered knowledge, and its use to solve new problems.
- Artificial intelligence has invested several decades in attempts to endow computers with the capability to understand observations;
- while progress has been tremendous, the practical ability of a machine to understand observations remains very limited.
- Attempting to solve related multidisciplinary scientific problems under the name cognitive systems is seen as a key to developing intelligent machines.

Interpretation of images

- From the mathematical logic and/or linguistics point of view, interpretation of images can be seen as a mapping interpretation:

$$\text{image data} \longrightarrow \text{model} \tag{1}$$

- The (logical) model means some specific world in which the observed objects make sense.
- Examples
  - nuclei of cells in a biological sample,
  - rivers in a satellite image,
  - or parts in an industrial process being checked for quality.
- There may be several interpretations of the same image(s).

Interpretation of images

- From the mathematical logic and/or linguistics point of view, interpretation of images can be seen as a mapping interpretation:

$$\text{image data} \longrightarrow \text{model} \tag{1}$$

- The (logical) model means some specific world in which the observed objects make sense.
- Examples
  - nuclei of cells in a biological sample,
  - rivers in a satellite image,
  - or parts in an industrial process being checked for quality.
- There may be several interpretations of the same image(s).

Interpretation of images

- From the mathematical logic and/or linguistics point of view, interpretation of images can be seen as a mapping interpretation:

$$\text{image data} \longrightarrow \text{model} \tag{1}$$

- The (logical) model means some specific world in which the observed objects make sense.
- Examples
  - nuclei of cells in a biological sample,
  - rivers in a satellite image,
  - or parts in an industrial process being checked for quality.
- There may be several interpretations of the same image(s).

Interpretation of images

- From the mathematical logic and/or linguistics point of view, interpretation of images can be seen as a mapping interpretation:

$$\text{image data} \longrightarrow \text{model} \tag{1}$$

- The (logical) model means some specific world in which the observed objects make sense.
- Examples
  - nuclei of cells in a biological sample,
  - rivers in a satellite image,
  - or parts in an industrial process being checked for quality.
- There may be several interpretations of the same image(s).

Interpretation of images

- From the mathematical logic and/or linguistics point of view, interpretation of images can be seen as a mapping interpretation:

$$\text{image data} \longrightarrow \text{model} \quad (1)$$

- The (logical) model means some specific world in which the observed objects make sense.
- Examples
  - nuclei of cells in a biological sample,
  - rivers in a satellite image,
  - or parts in an industrial process being checked for quality.
- There may be several interpretations of the same image(s).

Interpretation of images

- From the mathematical logic and/or linguistics point of view, interpretation of images can be seen as a mapping interpretation:

$$\text{image data} \longrightarrow \text{model} \tag{1}$$

- The (logical) model means some specific world in which the observed objects make sense.
- Examples
  - nuclei of cells in a biological sample,
  - rivers in a satellite image,
  - or parts in an industrial process being checked for quality.
- There may be several interpretations of the same image(s).

Interpretation of images

- From the mathematical logic and/or linguistics point of view, interpretation of images can be seen as a mapping interpretation:

$$\text{image data} \longrightarrow \text{model} \tag{1}$$

- The (logical) model means some specific world in which the observed objects make sense.
- Examples
  - nuclei of cells in a biological sample,
  - rivers in a satellite image,
  - or parts in an industrial process being checked for quality.
- There may be several interpretations of the same image(s).

Semantics of images

- Introducing interpretation to computer vision allows us to use concepts from mathematical logic, linguistics as syntax (rules describing correctly formed expression), and semantics (study of meaning).
- Considering observations (images) as an instance of formal expressions, semantics studies relations between expressions and their meanings.
- The interpretation of image(s) in computer vision can be understood as an instance of semantics.
- Practically, if the image understanding algorithms know into which particular domain (model in logical terminology) the observed world is constrained, then automatic analysis can be used for complicated problems.

Semantics of images

- Introducing interpretation to computer vision allows us to use concepts from mathematical logic, linguistics as syntax (rules describing correctly formed expression), and semantics (study of meaning).
- Considering observations (images) as an instance of formal expressions, semantics studies relations between expressions and their meanings.
- The interpretation of image(s) in computer vision can be understood as an instance of semantics.
- Practically, if the image understanding algorithms know into which particular domain (model in logical terminology) the observed world is constrained, then automatic analysis can be used for complicated problems.

Semantics of images

- Introducing interpretation to computer vision allows us to use concepts from mathematical logic, linguistics as syntax (rules describing correctly formed expression), and semantics (study of meaning).
- Considering observations (images) as an instance of formal expressions, semantics studies relations between expressions and their meanings.
- The interpretation of image(s) in computer vision can be understood as an instance of semantics.
- Practically, if the image understanding algorithms know into which particular domain (model in logical terminology) the observed world is constrained, then automatic analysis can be used for complicated problems.

Semantics of images

- Introducing interpretation to computer vision allows us to use concepts from mathematical logic, linguistics as syntax (rules describing correctly formed expression), and semantics (study of meaning).
- Considering observations (images) as an instance of formal expressions, semantics studies relations between expressions and their meanings.
- The interpretation of image(s) in computer vision can be understood as an instance of semantics.
- Practically, if the image understanding algorithms know into which particular domain (model in logical terminology) the observed world is constrained, then automatic analysis can be used for complicated problems.

III. Noise

- Noise is inherently present in each measurement in the real world.
- Its existence calls for mathematical tools which are able to cope with uncertainty; an example is probability theory.
- Of course, more complex tools make the image analysis much more complicated compared to standard (deterministic) methods.

III. Noise

- Noise is inherently present in each measurement in the real world.
- Its existence calls for mathematical tools which are able to cope with uncertainty; an example is probability theory.
- Of course, more complex tools make the image analysis much more complicated compared to standard (deterministic) methods.

III. Noise

- Noise is inherently present in each measurement in the real world.
- Its existence calls for mathematical tools which are able to cope with uncertainty; an example is probability theory.
- Of course, more complex tools make the image analysis much more complicated compared to standard (deterministic) methods.

IV. Too much data

- Images and video sequences are huge.
- An A4 sheet of paper scanned monochromatically at 300 dots per inch (dpi) at 8 bits per pixel corresponds to 8.5 MB.
- Non-interlaced RGB 24 bit color video 512×768 pixels, 25 frames per second, makes a data stream of 225 Mb per second.
- If the processing we devise is not very simple, then it is hard to achieve real-time performance; i.e., to process 25 or 30 images per second.

IV. Too much data

- Images and video sequences are huge.
- An A4 sheet of paper scanned monochromatically at 300 dots per inch (dpi) at 8 bits per pixel corresponds to 8.5 MB.
- Non-interlaced RGB 24 bit color video 512×768 pixels, 25 frames per second, makes a data stream of 225 Mb per second.
- If the processing we devise is not very simple, then it is hard to achieve real-time performance; i.e., to process 25 or 30 images per second.

IV. Too much data

- Images and video sequences are huge.
- An A4 sheet of paper scanned monochromatically at 300 dots per inch (dpi) at 8 bits per pixel corresponds to 8.5 MB.
- Non-interlaced RGB 24 bit color video 512×768 pixels, 25 frames per second, makes a data stream of 225 Mb per second.
- If the processing we devise is not very simple, then it is hard to achieve real-time performance; i.e., to process 25 or 30 images per second.

IV. Too much data

- Images and video sequences are huge.
- An A4 sheet of paper scanned monochromatically at 300 dots per inch (dpi) at 8 bits per pixel corresponds to 8.5 MB.
- Non-interlaced RGB 24 bit color video 512 $\times$ 768 pixels, 25 frames per second, makes a data stream of 225 Mb per second.
- If the processing we devise is not very simple, then it is hard to achieve real-time performance; i.e., to process 25 or 30 images per second.

V. Complexity in mage formation

- Brightness measured in the image is given by complicated image formation physics.
- The `radiance` (`brightness, image intensity`) depends on the `irradiance` (light source type, intensity and position), the observer's position, the surface local geometry, and the surface reflectance properties.
- The inverse tasks are ill-posed — for example, to reconstruct local surface orientation from intensity variations.

V. Complexity in mage formation

- Brightness measured in the image is given by complicated image formation physics.
- The `radiance` (`brightness`, `image intensity`) depends on the `irradiance` (light source type, intensity and position), the observer's position, the surface local geometry, and the surface reflectance properties.
- The inverse tasks are ill-posed — for example, to reconstruct local surface orientation from intensity variations.

V. Complexity in mage formation

- Brightness measured in the image is given by complicated image formation physics.
- The `radiance` (`brightness`, `image intensity`) depends on the `irradiance` (light source type, intensity and position), the observer's position, the surface local geometry, and the surface reflectance properties.
- The inverse tasks are ill-posed — for example, to reconstruct local surface orientation from intensity variations.

VI. Local window vs. for global view

- Commonly, image analysis algorithms analyze a particular storage bin in an operational memory (e.g., a pixel in the image) and its local neighborhood;
- the computer sees the image through a keyhole.
- Seeing the world through a keyhole makes it very difficult to understand more global context.
- It is often very difficult to interpret an image if it is seen only locally or if only a few local keyholes are available.

VI. Local window vs. for global view

- Commonly, image analysis algorithms analyze a particular storage bin in an operational memory (e.g., a pixel in the image) and its local neighborhood;
- the computer sees the image through a keyhole.
- Seeing the world through a keyhole makes it very difficult to understand more global context.
- It is often very difficult to interpret an image if it is seen only locally or if only a few local keyholes are available.

VI. Local window vs. for global view

- Commonly, image analysis algorithms analyze a particular storage bin in an operational memory (e.g., a pixel in the image) and its local neighborhood;
- the computer sees the image through a keyhole.
- Seeing the world through a keyhole makes it very difficult to understand more global context.
- It is often very difficult to interpret an image if it is seen only locally or if only a few local keyholes are available.

VI. Local window vs. for global view

- Commonly, image analysis algorithms analyze a particular storage bin in an operational memory (e.g., a pixel in the image) and its local neighborhood;
- the computer sees the image through a keyhole.
- Seeing the world through a keyhole makes it very difficult to understand more global context.
- It is often very difficult to interpret an image if it is seen only locally or if only a few local keyholes are available.

Local parts of an image

Figure: Illustration of the world seen through several keyholes providing only a very local context.

The global view of the image

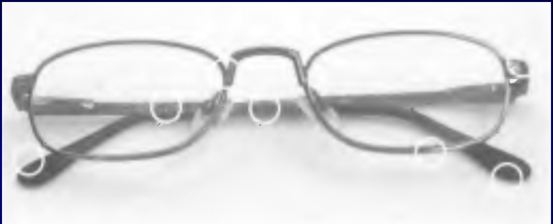

Figure: How context is taken into account is an important facet of image analysis.

Two Approaches

- There are philosophically two approaches: bionics and engineering (that is project attempt coordinated), approaches.
- The bionics approach has not been so successful, since we do have a through understanding about the biological vision system.

Two Approaches

- There are philosophically two approaches: `bionics` and engineering (that is project attempt coordinated), approaches.
- The bionics approach has not been so successful, since we do have a through understanding about the biological vision system.

Image Understanding

- Image understanding by a machine can be seen as an attempt to find a relation between input image(s) and previously established models of the observed world.
- Transition from the input image(s) to the model reduces the information contained in the image to relevant information for the application domain.
- This process is usually divided into several steps and several levels representing the image are used.
- The bottom layer contains raw image data and the higher levels interpret the data.
- Computer vision designs these intermediate representations and algorithms serving to establish and maintain relations between entities within and between layers.

Image Understanding

- `Image understanding` by a machine can be seen as an attempt to find a relation between input image(s) and previously established models of the observed world.
- Transition from the input image(s) to the model reduces the information contained in the image to relevant information for the application domain.
- This process is usually divided into several steps and several levels representing the image are used.
- The bottom layer contains raw image data and the higher levels interpret the data.
- `Computer vision` designs these intermediate representations and algorithms serving to establish and maintain relations between entities within and between layers.

Image Understanding

- `Image understanding` by a machine can be seen as an attempt to find a relation between input image(s) and previously established models of the observed world.
- Transition from the input image(s) to the model reduces the information contained in the image to relevant information for the application domain.
- This process is usually divided into several steps and several levels representing the image are used.
- The bottom layer contains raw image data and the higher levels interpret the data.
- `Computer vision` designs these intermediate representations and algorithms serving to establish and maintain relations between entities within and between layers.

Image Understanding

- `Image understanding` by a machine can be seen as an attempt to find a relation between input image(s) and previously established models of the observed world.
- Transition from the input image(s) to the model reduces the information contained in the image to relevant information for the application domain.
- This process is usually divided into several steps and several levels representing the image are used.
- The bottom layer contains raw image data and the higher levels interpret the data.
- `Computer vision` designs these intermediate representations and algorithms serving to establish and maintain relations between entities within and between layers.

Image Understanding

- Image understanding by a machine can be seen as an attempt to find a relation between input image(s) and previously established models of the observed world.
- Transition from the input image(s) to the model reduces the information contained in the image to relevant information for the application domain.
- This process is usually divided into several steps and several levels representing the image are used.
- The bottom layer contains raw image data and the higher levels interpret the data.
- Computer vision designs these intermediate representations and algorithms serving to establish and maintain relations between entities within and between layers.

Image Representation

- Image representation can be roughly divided according to data organization into four levels.
- The boundaries between individual levels arc inexact, and more detailed divisions are also proposed.

Image Representation

- `Image representation` can be roughly divided according to data organization into four levels.
- The boundaries between individual levels arc inexact, and more detailed divisions are also proposed.

Image Representation & Hierarchy of Computer Vision

- This suggests a bottom up way of information processing, from signals with almost no abstraction, to the highly abstract description needed for image understanding.
- The flow of information does not need to be unidirectional.
- Feedback loops are often introduced to allow the modification of algorithms according to intermediate results.

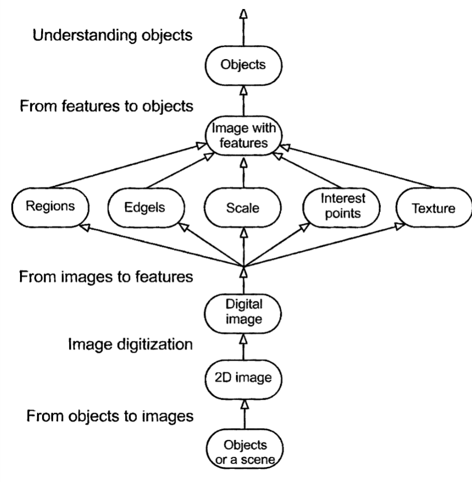

Image Representation & Hierarchy of Computer Vision

- This suggests a bottom up way of information processing, from signals with almost no abstraction, to the highly abstract description needed for image understanding.
- The flow of information does not need to be unidirectional.
- Feedback loops are often introduced to allow the modification of algorithms according to intermediate results.

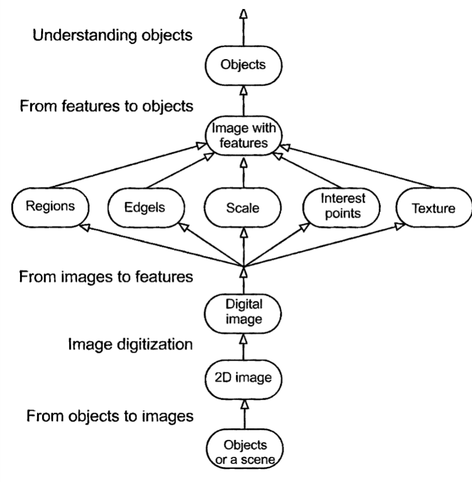

Image Representation & Hierarchy of Computer Vision

- This suggests a bottom up way of information processing, from signals with almost no abstraction, to the highly abstract description needed for image understanding.
- The flow of information does not need to be unidirectional.
- Feedback loops are often introduced to allow the modification of algorithms according to intermediate results.

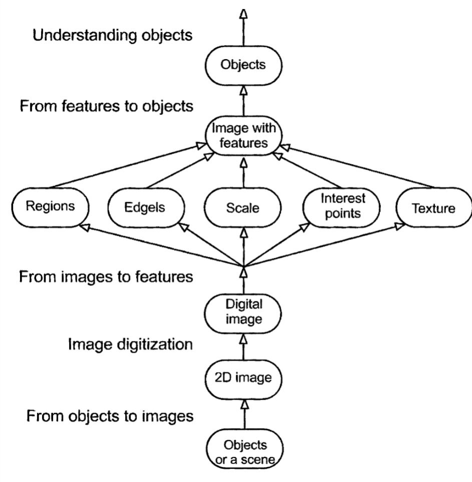

Image Representation & Hierarchy of Computer Vision

- This suggests a bottom up way of information processing, from signals with almost no abstraction, to the highly abstract description needed for image understanding.
- The flow of information does not need to be unidirectional.
- Feedback loops are often introduced to allow the modification of algorithms according to intermediate results.

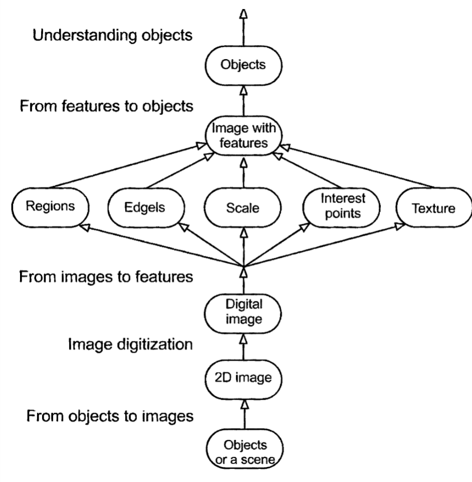

Image Representation & Hierarchy of Computer Vision

- This suggests a bottom up way of information processing, from signals with almost no abstraction, to the highly abstract description needed for image understanding.
- The flow of information does not need to be unidirectional.
- Feedback loops are often introduced to allow the modification of algorithms according to intermediate results.

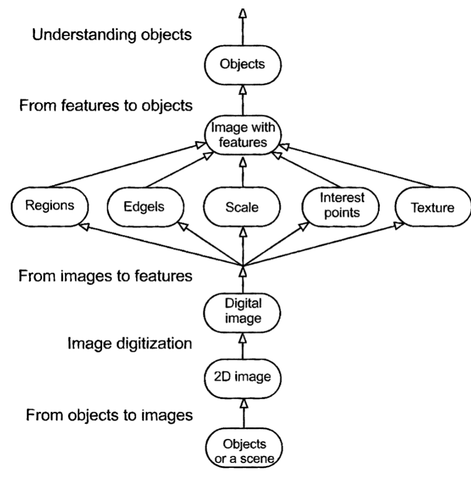

Image Representation & Hierarchy of Computer Vision

- This suggests a bottom up way of information processing, from signals with almost no abstraction, to the highly abstract description needed for image understanding.
- The flow of information does not need to be unidirectional.
- Feedback loops are often introduced to allow the modification of algorithms according to intermediate results.

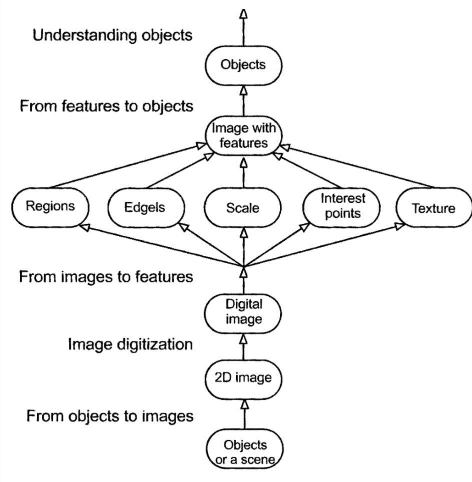

Two Levels

- This hierarchy of image representation and related algorithms is frequently categorized in an even simpler way.
- Two levels are often distinguished:
  - low-level image processing;
  - high-level image understanding.

Two Levels

- This hierarchy of image representation and related algorithms is frequently categorized in an even simpler way.
- Two levels are often distinguished:
  - low-level image processing;
  - high-level image understanding.

Two Levels

- This hierarchy of image representation and related algorithms is frequently categorized in an even simpler way.
- Two levels are often distinguished:
  - low-level image processing;
  - high-level image understanding.

Two Levels

- This hierarchy of image representation and related algorithms is frequently categorized in an even simpler way.
- Two levels are often distinguished:
  - low-level image processing;
  - high-level image understanding.

Low-level processing

- Low-level processing methods usually use very little knowledge about the content of images.
- Low-level methods often include image compression, pre-processing methods for noise filtering, edge extraction, and image sharpening.
- Ee shall discuss in this course.
- Low-level image processing uses data which resemble the input image.
- Very often, such a data set will be part of a video stream with an associated frame rate.
- E.g., an input image captured by a TV camera is 2D in nature, being described by an image function $f(x, y, t)$ whose value, at simplest, is usually brightness depending on parameters x, y and t.

Low-level processing

- Low-level processing methods usually use very little knowledge about the content of images.
- Low-level methods often include image compression, pre-processing methods for noise filtering, edge extraction, and image sharpening.
- Ee shall discuss in this course.
- Low-level image processing uses data which resemble the input image.
- Very often, such a data set will be part of a video stream with an associated frame rate.
- E.g., an input image captured by a TV camera is 2D in nature, being described by an image function $f(x, y, t)$ whose value, at simplest, is usually brightness depending on parameters x, y and t.

Low-level processing

- Low-level processing methods usually use very little knowledge about the content of images.
- Low-level methods often include image compression, pre-processing methods for noise filtering, edge extraction, and image sharpening.
- Ee shall discuss in this course.
- Low-level image processing uses data which resemble the input image.
- Very often, such a data set will be part of a video stream with an associated frame rate.
- E.g., an input image captured by a TV camera is 2D in nature, being described by an image function $f(x, y, t)$ whose value, at simplest, is usually brightness depending on parameters x, y and t.

Low-level processing

- Low-level processing methods usually use very little knowledge about the content of images.
- Low-level methods often include image compression, pre-processing methods for noise filtering, edge extraction, and image sharpening.
- Ee shall discuss in this course.
- Low-level image processing uses data which resemble the input image.
- Very often, such a data set will be part of a video stream with an associated frame rate.
- E.g., an input image captured by a TV camera is 2D in nature, being described by an image function $f(x, y, t)$ whose value, at simplest, is usually brightness depending on parameters x, y and t.

Low-level processing

- Low-level processing methods usually use very little knowledge about the content of images.
- Low-level methods often include image compression, pre-processing methods for noise filtering, edge extraction, and image sharpening.
- Ee shall discuss in this course.
- Low-level image processing uses data which resemble the input image.
- Very often, such a data set will be part of a video stream with an associated frame rate.
- E.g., an input image captured by a TV camera is 2D in nature, being described by an image function $f(x, y, t)$ whose value, at simplest, is usually brightness depending on parameters x, y and t.

Low-level processing

- Low-level processing methods usually use very little knowledge about the content of images.
- Low-level methods often include image compression, pre-processing methods for noise filtering, edge extraction, and image sharpening.
- Ee shall discuss in this course.
- Low-level image processing uses data which resemble the input image.
- Very often, such a data set will be part of a video stream with an associated frame rate.
- E.g., an input image captured by a TV camera is 2D in nature, being described by an image function $f(x, y, t)$ whose value, at simplest, is usually brightness depending on parameters x, y and t.

High-level processing I

- High-level processing is based on knowledge, goals, and plans of how to achieve those goals.
- Artificial intelligence methods are widely applicable.
- High-level computer vision tries to imitate human cognition (although be mindful of the health warning given in the very first paragraph of this chapter) and the ability to make decisions according to the information contained in the image.
- In the example described, high-level knowledge would be related to the shape of a cow and the subtle interrelationships between the different parts of that shape, and their (inter-)dynamics.

High-level processing I

- High-level processing is based on knowledge, goals, and plans of how to achieve those goals.
- Artificial intelligence methods are widely applicable.
- High-level computer vision tries to imitate human cognition (although be mindful of the health warning given in the very first paragraph of this chapter) and the ability to make decisions according to the information contained in the image.
- In the example described, high-level knowledge would be related to the shape of a cow and the subtle interrelationships between the different parts of that shape, and their (inter-)dynamics.

High-level processing I

- High-level processing is based on knowledge, goals, and plans of how to achieve those goals.
- Artificial intelligence methods are widely applicable.
- High-level computer vision tries to imitate human cognition (although be mindful of the health warning given in the very first paragraph of this chapter) and the ability to make decisions according to the information contained in the image.
- In the example described, high-level knowledge would be related to the **shape** of a cow and the subtle interrelationships between the different parts of that shape, and their (inter-)dynamics.

High-level processing I

- High-level processing is based on knowledge, goals, and plans of how to achieve those goals.
- Artificial intelligence methods are widely applicable.
- High-level computer vision tries to imitate human cognition (although be mindful of the health warning given in the very first paragraph of this chapter) and the ability to make decisions according to the information contained in the image.
- In the example described, high-level knowledge would be related to the **shape** of a cow and the subtle interrelationships between the different parts of that shape, and their (inter-)dynamics.

High-level processing II

- High-level vision begins with some form of formal model of the world, and then the “reality” perceived in the form of digitized images is compared to the model.
- A match is attempted.
- When differences emerge, partial matches (or sub-goals) are sought that overcome the mismatches.
- The omputer switches to low-level image processing to find information needed to update the model.
- This process is then repeated iteratively, and “understanding” an image thereby becomes a co-operation between top-down and bottom-up processes.
- A feedback loop is introduced in which high-level partial results create tasks for low-level image processing.
- The iterative image understanding process should eventually **converge** to the global goal.

High-level processing II

- High-level vision begins with some form of formal model of the world, and then the “reality” perceived in the form of digitized images is compared to the model.
- A match is attempted.
- When differences emerge, partial matches (or sub-goals) are sought that overcome the mismatches.
- The omputer switches to low-level image processing to find information needed to update the model.
- This process is then repeated iteratively, and “understanding” an image thereby becomes a co-operation between top-down and bottom-up processes.
- A feedback loop is introduced in which high-level partial results create tasks for low-level image processing.
- The iterative image understanding process should eventually **converge** to the global goal.

High-level processing II

- High-level vision begins with some form of formal model of the world, and then the “reality” perceived in the form of digitized images is compared to the model.
- A match is attempted.
- When differences emerge, partial matches (or sub-goals) are sought that overcome the mismatches.
- The omputer switches to low-level image processing to find information needed to update the model.
- This process is then repeated iteratively, and “understanding” an image thereby becomes a co-operation between top-down and bottom-up processes.
- A feedback loop is introduced in which high-level partial results create tasks for low-level image processing.
- The iterative image understanding process should eventually **converge** to the global goal.

High-level processing II

- High-level vision begins with some form of formal model of the world, and then the "reality" perceived in the form of digitized images is compared to the model.
- A match is attempted.
- When differences emerge, partial matches (or sub-goals) are sought that overcome the mismatches.
- The omputer switches to low-level image processing to find information needed to update the model.
- This process is then repeated iteratively, and "understanding" an image thereby becomes a co-operation between top-down and bottom-up processes.
- A feedback loop is introduced in which high-level partial results create tasks for low-level image processing.
- The iterative image understanding process should eventually **converge** to the global goal.

High-level processing II

- High-level vision begins with some form of formal model of the world, and then the “reality” perceived in the form of digitized images is compared to the model.
- A match is attempted.
- When differences emerge, partial matches (or sub-goals) are sought that overcome the mismatches.
- The omputer switches to low-level image processing to find information needed to update the model.
- This process is then repeated iteratively, and “understanding” an image thereby becomes a co-operation between top-down and bottom-up processes.
- A feedback loop is introduced in which high-level partial results create tasks for low-level image processing.
- The iterative image understanding process should eventually **converge** to the global goal.

High-level processing II

- High-level vision begins with some form of formal model of the world, and then the “reality” perceived in the form of digitized images is compared to the model.
- A match is attempted.
- When differences emerge, partial matches (or sub-goals) are sought that overcome the mismatches.
- The omputer switches to low-level image processing to find information needed to update the model.
- This process is then repeated iteratively, and “understanding” an image thereby becomes a co-operation between top-down and bottom-up processes.
- A feedback loop is introduced in which high-level partial results create tasks for low-level image processing.
- The iterative image understanding process should eventually **converge** to the global goal.

High-level processing II

- High-level vision begins with some form of formal model of the world, and then the “reality” perceived in the form of digitized images is compared to the model.
- A match is attempted.
- When differences emerge, partial matches (or sub-goals) are sought that overcome the mismatches.
- The omputer switches to low-level image processing to find information needed to update the model.
- This process is then repeated iteratively, and “understanding” an image thereby becomes a co-operation between top-down and bottom-up processes.
- A feedback loop is introduced in which high-level partial results create tasks for low-level image processing.
- The iterative image understanding process should eventually **converge** to the global goal.

Low- vs. high-level representations

- Both representations contain exactly the same information.
- But **for a human observer** it is not difficult to find a correspondence between them, and without the second, it is unlikely that one would recognize the face of a child.
- The point is that a lot of a priori knowledge is used by humans to interpret the images.
- A machine only begins with an array of numbers and so will be attempting to make identifications and draw conclusions from data that to us are more uncomprehensible.
- Increasingly, data capture equipment is providing very large data, sets that do not lend themselves to straightforward interpretation by humans.
- We have already mentioned terahertz imaging as an example.
- General knowledge, domain-specific knowledge, and information extracted from the image will be essential in attempting to “understand” these arrays of numbers.

Low- vs. high-level representations

- Both representations contain exactly the same information.
- But **for a human observer** it is not difficult to find a correspondence between them, and without the second, it is unlikely that one would recognize the face of a child.
- The point is that a lot of a priori knowledge is used by humans to interpret the images.
- A machine only begins with an array of numbers and so will be attempting to make identifications and draw conclusions from data that to us are more uncomprehensible.
- Increasingly, data capture equipment is providing very large data, sets that do not lend themselves to straightforward interpretation by humans.
- We have already mentioned terahertz imaging as an example.
- General knowledge, domain-specific knowledge, and information extracted from the image will be essential in attempting to “understand” these arrays of numbers.

Low- vs. high-level representations

- Both representations contain exactly the same information.
- But **for a human observer** it is not difficult to find a correspondence between them, and without the second, it is unlikely that one would recognize the face of a child.
- The point is that a lot of a priori knowledge is used by humans to interpret the images.
- A machine only begins with an array of numbers and so will be attempting to make identifications and draw conclusions from data that to us are more uncomprehensible.
- Increasingly, data capture equipment is providing very large data, sets that do not lend themselves to straightforward interpretation by humans.
- We have already mentioned terahertz imaging as an example.
- General knowledge, domain-specific knowledge, and information extracted from the image will be essential in attempting to “understand” these arrays of numbers.

Low- vs. high-level representations

- Both representations contain exactly the same information.
- But **for a human observer** it is not difficult to find a correspondence between them, and without the second, it is unlikely that one would recognize the face of a child.
- The point is that a lot of a priori knowledge is used by humans to interpret the images.
- A machine only begins with an array of numbers and so will be attempting to make identifications and draw conclusions from data that to us are more uncomprehensible.
- Increasingly, data capture equipment is providing very large data, sets that do not lend themselves to straightforward interpretation by humans.
- We have already mentioned terahertz imaging as an example.
- General knowledge, domain-specific knowledge, and information extracted from the image will be essential in attempting to “understand” these arrays of numbers.

Low- vs. high-level representations

- Both representations contain exactly the same information.
- But **for a human observer** it is not difficult to find a correspondence between them, and without the second, it is unlikely that one would recognize the face of a child.
- The point is that a lot of a priori knowledge is used by humans to interpret the images.
- A machine only begins with an array of numbers and so will be attempting to make identifications and draw conclusions from data that to us are more uncomprehensible.
- Increasingly, data capture equipment is providing very large data, sets that do not lend themselves to straightforward interpretation by humans.
- We have already mentioned terahertz imaging as an example.
- General knowledge, domain-specific knowledge, and information extracted from the image will be essential in attempting to “understand” these arrays of numbers.

Low- vs. high-level representations

- Both representations contain exactly the same information.
- But **for a human observer** it is not difficult to find a correspondence between them, and without the second, it is unlikely that one would recognize the face of a child.
- The point is that a lot of a priori knowledge is used by humans to interpret the images.
- A machine only begins with an array of numbers and so will be attempting to make identifications and draw conclusions from data that to us are more uncomprehensible.
- Increasingly, data capture equipment is providing very large data, sets that do not lend themselves to straightforward interpretation by humans.
- We have already mentioned terahertz imaging as an example.
- General knowledge, domain-specific knowledge, and information extracted from the image will be essential in attempting to "understand" these arrays of numbers.

Low- vs. high-level representations

- Both representations contain exactly the same information.
- But **for a human observer** it is not difficult to find a correspondence between them, and without the second, it is unlikely that one would recognize the face of a child.
- The point is that a lot of a priori knowledge is used by humans to interpret the images.
- A machine only begins with an array of numbers and so will be attempting to make identifications and draw conclusions from data that to us are more uncomprehensible.
- Increasingly, data capture equipment is providing very large data, sets that do not lend themselves to straightforward interpretation by humans.
- We have already mentioned terahertz imaging as an example.
- General knowledge, domain-specific knowledge, and information extracted from the image will be essential in attempting to “understand” these arrays of numbers.

Low-level processing

The following sequence of processing steps is commonly recognized:

- Image Acquisition: An image is captured by a sensor (such as a TV camera) and digitized. Image may come in many **formats** and ways.
- Preprocessing: Image reconstruction or restoration, denoising and enhancement. E.g., computer tomography.
- Image coding and compression: this is important for transferring images.
- Image segmentation: computer tries to separate objects from the image background.
- Object description and classification in a totally segmented image is also understood as part of low-level image processing.

Low-level processing

The following sequence of processing steps is commonly recognized:

- Image Acquisition: An image is captured by a sensor (such as a TV camera) and digitized. Image may come in many **formats** and ways.
- Preprocessing: Image reconstruction or restoration, denoising and enhancement. E.g., computer tomography.
- Image coding and compression: this is important for transferring images.
- Image segmentation: computer tries to separate objects from the image background.
- Object description and classification in a totally segmented image is also understood as part of low-level image processing.

Low-level processing

The following sequence of processing steps is commonly recognized:

- Image Acquisition: An image is captured by a sensor (such as a TV camera) and digitized. Image may come in many **formats** and ways.
- Preprocessing: Image reconstruction or restoration, denoising and enhancement. E.g., computer tomography.
- Image coding and compression: this is important for transferring images.
- Image segmentation: computer tries to separate objects from the image background.
- Object description and classification in a totally segmented image is also understood as part of low-level image processing.

Low-level processing

The following sequence of processing steps is commonly recognized:

- Image Acquisition: An image is captured by a sensor (such as a TV camera) and digitized. Image may come in many **formats** and ways.
- Preprocessing: Image reconstruction or restoration, denoising and enhancement. E.g., computer tomography.
- Image coding and compression: this is important for transferring images.
- Image segmentation: computer tries to separate objects from the image background.
- Object description and classification in a totally segmented image is also understood as part of low-level image processing.

Low-level processing

The following sequence of processing steps is commonly recognized:

- Image Acquisition: An image is captured by a sensor (such as a TV camera) and digitized. Image may come in many **formats** and ways.
- Preprocessing: Image reconstruction or restoration, denoising and enhancement. E.g., computer tomography.
- Image coding and compression: this is important for transferring images.
- Image segmentation: computer tries to separate objects from the image background.
- Object description and classification in a totally segmented image is also understood as part of low-level image processing.

Image Segmentation

- Image segmentation is to separate objects from the image background and from each other.
- Total and `partial segmentation` may be distinguished.
- `Total segmentation` is possible only for very simple tasks, an example being the recognition of dark non-touching objects from a light background.
- Example: optical character recognition, OCR.
- Een this superficially simple problem is very hard to solve without error.
- In more complicated problems (the general case), low-level image processing techniques handle the partial segmentation tasks, in which only the cues which will aid further high-level processing are extracted.
- Often, finding parts of object boundaries is an example of low-level partial segmentation.

Image Segmentation

- Image segmentation is to separate objects from the image background and from each other.
- Total and `partial segmentation` may be distinguished.
- `Total segmentation` is possible only for very simple tasks, an example being the recognition of dark non-touching objects from a light background.
- Example: optical character recognition, OCR.
- Een this superficially simple problem is very hard to solve without error.
- In more complicated problems (the general case), low-level image processing techniques handle the partial segmentation tasks, in which only the cues which will aid further high-level processing are extracted.
- Often, finding parts of object boundaries is an example of low-level partial segmentation.

Image Segmentation

- Image segmentation is to separate objects from the image background and from each other.
- Total and `partial segmentation` may be distinguished.
- `Total segmentation` is possible only for very simple tasks, an example being the recognition of dark non-touching objects from a light background.
- Example: optical character recognition, OCR.
- Een this superficially simple problem is very hard to solve without error.
- In more complicated problems (the general case), low-level image processing techniques handle the partial segmentation tasks, in which only the cues which will aid further high-level processing are extracted.
- Often, finding parts of object boundaries is an example of low-level partial segmentation.

Image Segmentation

- Image segmentation is to separate objects from the image background and from each other.
- Total and `partial segmentation` may be distinguished.
- `Total segmentation` is possible only for very simple tasks, an example being the recognition of dark non-touching objects from a light background.
- Example: optical character recognition, OCR.
- Een this superficially simple problem is very hard to solve without error.
- In more complicated problems (the general case), low-level image processing techniques handle the partial segmentation tasks, in which only the cues which will aid further high-level processing are extracted.
- Often, finding parts of object boundaries is an example of low-level partial segmentation.

Image Segmentation

- Image segmentation is to separate objects from the image background and from each other.
- Total and `partial segmentation` may be distinguished.
- `Total segmentation` is possible only for very simple tasks, an example being the recognition of dark non-touching objects from a light background.
- Example: optical character recognition, OCR.
- Een this superficially simple problem is very hard to solve without error.
- In more complicated problems (the general case), low-level image processing techniques handle the partial segmentation tasks, in which only the cues which will aid further high-level processing are extracted.
- Often, finding parts of object boundaries is an example of low-level partial segmentation.

Image Segmentation

- Image segmentation is to separate objects from the image background and from each other.
- Total and `partial segmentation` may be distinguished.
- `Total segmentation` is possible only for very simple tasks, an example being the recognition of dark non-touching objects from a light background.
- Example: optical character recognition, OCR.
- Een this superficially simple problem is very hard to solve without error.
- In more complicated problems (the general case), low-level image processing techniques handle the partial segmentation tasks, in which only the cues which will aid further high-level processing are extracted.
- Often, finding parts of object boundaries is an example of low-level partial segmentation.

Image Segmentation

- Image segmentation is to separate objects from the image background and from each other.
- Total and `partial segmentation` may be distinguished.
- `Total segmentation` is possible only for very simple tasks, an example being the recognition of dark non-touching objects from a light background.
- Example: optical character recognition, OCR.
- Een this superficially simple problem is very hard to solve without error.
- In more complicated problems (the general case), low-level image processing techniques handle the partial segmentation tasks, in which only the cues which will aid further high-level processing are extracted.
- Often, finding parts of object boundaries is an example of low-level partial segmentation.

Low-level Image Processing

- Low-level computer vision techniques overlap almost completely with digital image processing, which has been practiced for decades.
- Object description and classification in a totally segmented image are also understood as part of low-level image processing.
- Other low-level operations are image compression, and techniques to extract information from (but not understand) moving scenes.
- ....

Low-level Image Processing

- Low-level computer vision techniques overlap almost completely with digital image processing, which has been practiced for decades.
- Object description and classification in a totally segmented image are also understood as part of low-level image processing.
- Other low-level operations are image compression, and techniques to extract information from (but not understand) moving scenes.
- ....

Low-level Image Processing

- Low-level computer vision techniques overlap almost completely with digital image processing, which has been practiced for decades.
- Object description and classification in a totally segmented image are also understood as part of low-level image processing.
- Other low-level operations are image compression, and techniques to extract information from (but not understand) moving scenes.
- ....

Low-level Image Processing

- Low-level computer vision techniques overlap almost completely with digital image processing, which has been practiced for decades.
- Object description and classification in a totally segmented image are also understood as part of low-level image processing.
- Other low-level operations are image compression, and techniques to extract information from (but not understand) moving scenes.
- ....

Low vs High

- Low-level image processing and high-level computer vision differ in the data used.
- Low-level data are comprised of original images represented by matrices composed of brightness (or similar) values.
- High-level data originate in images as well, but only those data which are relevant to high-level goals are extracted, reducing the data quantity considerably.
- High-level data represent knowledge about the image content. —
- E.g., object size, shape, and mutual relations between objects in the image.
- High-level data are usually expressed in symbolic form.

Low vs High

- Low-level image processing and high-level computer vision differ in the data used.
- Low-level data are comprised of original images represented by matrices composed of brightness (or similar) values.
- High-level data originate in images as well, but only those data which are relevant to high-level goals are extracted, reducing the data quantity considerably.
- High-level data represent knowledge about the image content. —
- E.g., object size, shape, and mutual relations between objects in the image.
- High-level data are usually expressed in symbolic form.

Low vs High

- Low-level image processing and high-level computer vision differ in the data used.
- Low-level data are comprised of original images represented by matrices composed of brightness (or similar) values.
- High-level data originate in images as well, but only those data which are relevant to high-level goals are extracted, reducing the data quantity considerably.
- High-level data represent knowledge about the image content. —
- E.g., object size, shape, and mutual relations between objects in the image.
- High-level data are usually expressed in symbolic form.

Low vs High

- Low-level image processing and high-level computer vision differ in the data used.
- Low-level data are comprised of original images represented by matrices composed of brightness (or similar) values.
- High-level data originate in images as well, but only those data which are relevant to high-level goals are extracted, reducing the data quantity considerably.
- High-level data represent knowledge about the image content. —
- E.g., object size, shape, and mutual relations between objects in the image.
- High-level data are usually expressed in symbolic form.

Low vs High

- Low-level image processing and high-level computer vision differ in the data used.
- Low-level data are comprised of original images represented by matrices composed of brightness (or similar) values.
- High-level data originate in images as well, but only those data which are relevant to high-level goals are extracted, reducing the data quantity considerably.
- High-level data represent knowledge about the image content. —
- E.g., object size, shape, and mutual relations between objects in the image.
- High-level data are usually expressed in symbolic form.

Low vs High

- Low-level image processing and high-level computer vision differ in the data used.
- Low-level data are comprised of original images represented by matrices composed of brightness (or similar) values.
- High-level data originate in images as well, but only those data which are relevant to high-level goals are extracted, reducing the data quantity considerably.
- High-level data represent knowledge about the image content. —
- E.g., object size, shape, and mutual relations between objects in the image.
- High-level data are usually expressed in symbolic form.

Image Processing

- Most current low-level image processing methods were proposed in the 1970s or earlier.
- Recent research is trying to find more efficient and more general algorithms, implementations.
- The requirement for better and faster algorithms is fuelled by technology delivering larger images (better spatial resolution), and color.
- A complicated and so far unsolved problem is how to order low-level steps to solve a specific task, and the aim of automating this problem has not yet been achieved.
- It is usually still a human operator who finds a sequence of relevant operations.
- Domain- specific knowledge and uncertainty cause much to depend on this operator's intuition and previous experience.

Image Processing

- Most current low-level image processing methods were proposed in the 1970s or earlier.
- Recent research is trying to find more efficient and more general algorithms, implementations.
- The requirement for better and faster algorithms is fuelled by technology delivering larger images (better spatial resolution), and color.
- A complicated and so far unsolved problem is how to order low-level steps to solve a specific task, and the aim of automating this problem has not yet been achieved.
- It is usually still a human operator who finds a sequence of relevant operations.
- Domain- specific knowledge and uncertainty cause much to depend on this operator's intuition and previous experience.

Image Processing

- Most current low-level image processing methods were proposed in the 1970s or earlier.
- Recent research is trying to find more efficient and more general algorithms, implementations.
- The requirement for better and faster algorithms is fuelled by technology delivering larger images (better spatial resolution), and color.
- A complicated and so far unsolved problem is how to order low-level steps to solve a specific task, and the aim of automating this problem has not yet been achieved.
- It is usually still a human operator who finds a sequence of relevant operations.
- Domain- specific knowledge and uncertainty cause much to depend on this operator's intuition and previous experience.

Image Processing

- Most current low-level image processing methods were proposed in the 1970s or earlier.
- Recent research is trying to find more efficient and more general algorithms, implementations.
- The requirement for better and faster algorithms is fuelled by technology delivering larger images (better spatial resolution), and color.
- A complicated and so far unsolved problem is how to order low-level steps to solve a specific task, and the aim of automating this problem has not yet been achieved.
- It is usually still a human operator who finds a sequence of relevant operations.
- Domain- specific knowledge and uncertainty cause much to depend on this operator's intuition and previous experience.

Image Processing

- Most current low-level image processing methods were proposed in the 1970s or earlier.
- Recent research is trying to find more efficient and more general algorithms, implementations.
- The requirement for better and faster algorithms is fuelled by technology delivering larger images (better spatial resolution), and color.
- A complicated and so far unsolved problem is how to order low-level steps to solve a specific task, and the aim of automating this problem has not yet been achieved.
- It is usually still a human operator who finds a sequence of relevant operations.
- Domain- specific knowledge and uncertainty cause much to depend on this operator's intuition and previous experience.

Image Processing

- Most current low-level image processing methods were proposed in the 1970s or earlier.
- Recent research is trying to find more efficient and more general algorithms, implementations.
- The requirement for better and faster algorithms is fuelled by technology delivering larger images (better spatial resolution), and color.
- A complicated and so far unsolved problem is how to order low-level steps to solve a specific task, and the aim of automating this problem has not yet been achieved.
- It is usually still a human operator who finds a sequence of relevant operations.
- Domain- specific knowledge and uncertainty cause much to depend on this operator's intuition and previous experience.

High-level Vision

- High-level vision tries to extract and order image processing steps using all available knowledge.
- Image understanding is the heart of the method, in which feedback from high-level to low-level is used.
- Unsurprisingly this task is very complicated and computationally intensive.
- David Marr's book [Marr, 1982] influenced computer vision considerably throughout the 1980s.
- It described a new methodology and computational theory inspired by biological vision systems.
- Developments in the 1990s moved away from dependence on this paradigm, but interest in properly understanding and then modeling human visual systems.
- It remains the case that the only known solution to the "vision problem" is our own brain!

High-level Vision

- High-level vision tries to extract and order image processing steps using all available knowledge.
- Image understanding is the heart of the method, in which feedback from high-level to low-level is used.
- Unsurprisingly this task is very complicated and computationally intensive.
- David Marr's book [Marr, 1982] influenced computer vision considerably throughout the 1980s.
- It described a new methodology and computational theory inspired by biological vision systems.
- Developments in the 1990s moved away from dependence on this paradigm, but interest in properly understanding and then modeling human visual systems.
- It remains the case that the only known solution to the "vision problem" is our own brain!

High-level Vision

- High-level vision tries to extract and order image processing steps using all available knowledge.
- Image understanding is the heart of the method, in which feedback from high-level to low-level is used.
- Unsurprisingly this task is very complicated and computationally intensive.
- David Marr's book [Marr, 1982] influenced computer vision considerably throughout the 1980s.
- It described a new methodology and computational theory inspired by biological vision systems.
- Developments in the 1990s moved away from dependence on this paradigm, but interest in properly understanding and then modeling human visual systems.
- It remains the case that the only known solution to the "vision problem" is our own brain!

High-level Vision

- High-level vision tries to extract and order image processing steps using all available knowledge.
- Image understanding is the heart of the method, in which feedback from high-level to low-level is used.
- Unsurprisingly this task is very complicated and computationally intensive.
- David Marr's book [Marr, 1982] influenced computer vision considerably throughout the 1980s.
- It described a new methodology and computational theory inspired by biological vision systems.
- Developments in the 1990s moved away from dependence on this paradigm, but interest in properly understanding and then modeling human visual systems.
- It remains the case that the only known solution to the "vision problem" is our own brain!

High-level Vision

- High-level vision tries to extract and order image processing steps using all available knowledge.
- Image understanding is the heart of the method, in which feedback from high-level to low-level is used.
- Unsurprisingly this task is very complicated and computationally intensive.
- David Marr's book [Marr, 1982] influenced computer vision considerably throughout the 1980s.
- It described a new methodology and computational theory inspired by biological vision systems.
- Developments in the 1990s moved away from dependence on this paradigm, but interest in properly understanding and then modeling human visual systems.
- It remains the case that the only known solution to the "vision problem" is our own brain!

High-level Vision

- High-level vision tries to extract and order image processing steps using all available knowledge.
- Image understanding is the heart of the method, in which feedback from high-level to low-level is used.
- Unsurprisingly this task is very complicated and computationally intensive.
- David Marr's book [Marr, 1982] influenced computer vision considerably throughout the 1980s.
- It described a new methodology and computational theory inspired by biological vision systems.
- Developments in the 1990s moved away from dependence on this paradigm, but interest in properly understanding and then modeling human visual systems.
- It remains the case that the only known solution to the "vision problem" is our own brain!

High-level Vision

- High-level vision tries to extract and order image processing steps using all available knowledge.
- Image understanding is the heart of the method, in which feedback from high-level to low-level is used.
- Unsurprisingly this task is very complicated and computationally intensive.
- David Marr's book [Marr, 1982] influenced computer vision considerably throughout the 1980s.
- It described a new methodology and computational theory inspired by biological vision systems.
- Developments in the 1990s moved away from dependence on this paradigm, but interest in properly understanding and then modeling human visual systems.
- It remains the case that the only known solution to the "vision problem" is our own brain!

3D Vision Problems

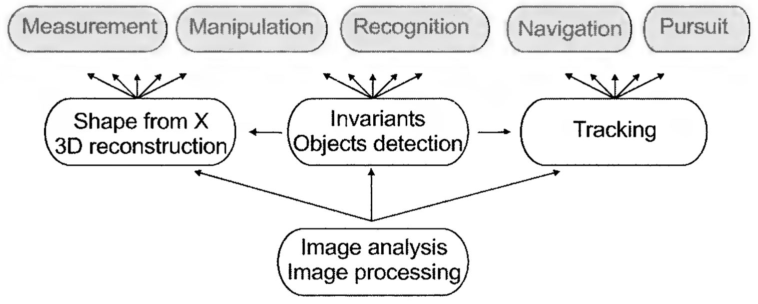

Figure: Several 3D vision tasks and algorithmic components expressed on different abstraction levels. We adopt the user's view, i.e., what tasks performed routinely by humans would be good to accomplish by machines.

3D Vision Problems

- What is the relation of these 3D vision tasks to low-level (image processing) and high- level (image analysis) algorithmic methods?
- There is no widely accepted view in the academic community.
- Links between (algorithmic) components and representation levels are tailored to the specific application solved, e.g., navigation of an autonomous vehicle.
- These applications have to employ specific knowledge about the problem solved to be competitive with tasks which humans solve.
- More general theories are expected to emerge.
- Many researchers in different fields work on related problems.
- There is a belief that research in 'cognitive systems' could be the key.

3D Vision Problems

- What is the relation of these 3D vision tasks to low-level (image processing) and high- level (image analysis) algorithmic methods?
- There is no widely accepted view in the academic community.
- Links between (algorithmic) components and representation levels are tailored to the specific application solved, e.g., navigation of an autonomous vehicle.
- These applications have to employ specific knowledge about the problem solved to be competitive with tasks which humans solve.
- More general theories are expected to emerge.
- Many researchers in different fields work on related problems.
- There is a belief that research in 'cognitive systems' could be the key.

3D Vision Problems

- What is the relation of these 3D vision tasks to low-level (image processing) and high- level (image analysis) algorithmic methods?
- There is no widely accepted view in the academic community.
- Links between (algorithmic) components and representation levels are tailored to the specific application solved, e.g., navigation of an autonomous vehicle.
- These applications have to employ **specific knowledge** about the problem solved to be competitive with tasks which humans solve.
- More general theories are expected to emerge.
- Many researchers in different fields work on related problems.
- There is a belief that research in 'cognitive systems' could be the key.

3D Vision Problems

- What is the relation of these 3D vision tasks to low-level (image processing) and high- level (image analysis) algorithmic methods?
- There is no widely accepted view in the academic community.
- Links between (algorithmic) components and representation levels are tailored to the specific application solved, e.g., navigation of an autonomous vehicle.
- These applications have to employ **specific knowledge** about the problem solved to be competitive with tasks which humans solve.
- More general theories are expected to emerge.
- Many researchers in different fields work on related problems.
- There is a belief that research in 'cognitive systems' could be the key.

3D Vision Problems

- What is the relation of these 3D vision tasks to low-level (image processing) and high- level (image analysis) algorithmic methods?
- There is no widely accepted view in the academic community.
- Links between (algorithmic) components and representation levels are tailored to the specific application solved, e.g., navigation of an autonomous vehicle.
- These applications have to employ **specific knowledge** about the problem solved to be competitive with tasks which humans solve.
- More general theories are expected to emerge.
- Many researchers in different fields work on related problems.
- There is a belief that research in 'cognitive systems' could be the key.

3D Vision Problems

- What is the relation of these 3D vision tasks to low-level (image processing) and high- level (image analysis) algorithmic methods?
- There is no widely accepted view in the academic community.
- Links between (algorithmic) components and representation levels are tailored to the specific application solved, e.g., navigation of an autonomous vehicle.
- These applications have to employ **specific knowledge** about the problem solved to be competitive with tasks which humans solve.
- More general theories are expected to emerge.
- Many researchers in different fields work on related problems.
- There is a belief that research in 'cognitive systems' could be the key.

3D Vision Problems

- What is the relation of these 3D vision tasks to low-level (image processing) and high- level (image analysis) algorithmic methods?
- There is no widely accepted view in the academic community.
- Links between (algorithmic) components and representation levels are tailored to the specific application solved, e.g., navigation of an autonomous vehicle.
- These applications have to employ **specific knowledge** about the problem solved to be competitive with tasks which humans solve.
- More general theories are expected to emerge.
- Many researchers in different fields work on related problems.
- There is a belief that research in 'cognitive systems' could be the key.

Course Overview

- Digital image processing, image analysis, image understanding are related branches of computer vision.
- This course is about digital image processing.
- The following topics are to be covered in this course.

Course Overview

- Digital image processing, image analysis, image understanding are related branches of computer vision.
- This course is about digital image processing.
- The following topics are to be covered in this course.

Course Overview

- Digital image processing, image analysis, image understanding are related branches of computer vision.
- This course is about digital image processing.
- The following topics are to be covered in this course.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

Course Syllabus

- Introduction and Course Overview
- Image representations and properties
  - Images as a stochastic processes or linear systems, etc.
  - Metric and topological properties of digital images
  - Histograms
  - Noise in images
- Data Structures for Image Analysis
- Image Pre-processing
  - Various pre-processing operators
- Image Segmentation
  - Thresholding, edge-based, region growing, segmentation method.
- Scale Space Theory
  - Image processing and partial differential equations.

The textbook and web resource

- Milan Sonka, V. Hlavac, R. Boyle: Image Processing, Analysis and Machine Vision, 3rd edition. Thomson Learning, 2008.
- Image Processing, Analysis, and Machine Vision: A MATLAB Companion, http://visionbook.felk.cvut.cz/

The textbook and web resource

- Milan Sonka, V. Hlavac, R. Boyle: Image Processing, Analysis and Machine Vision, 3rd edition. Thomson Learning, 2008.
- Image Processing, Analysis, and Machine Vision: A MATLAB Companion, http://visionbook.felk.cvut.cz/

References

- Kenneth R. Castleman: Digital Image Processing. Prentice-Hall International, Inc. 1996. Or Tsinghua University Press, 1998.
- Rongchun Zhao: Introduction to Digital Image Processing (in Chinese). Northwestern Polytechnical University Press, 2000.

References

- Kenneth R. Castleman: Digital Image Processing. Prentice-Hall International, Inc. 1996. Or Tsinghua University Press, 1998.
- Rongchun Zhao: Introduction to Digital Image Processing (in Chinese). Northwestern Polytechnical University Press, 2000.

Summary

- Human vision is natural and seems easy; computer mimicry of this is difficult.
- We might hope to examine pictures, or sequences of pictures, for quantitative and qualitative analysis.
- Many standard and advanced AI techniques are relevant.
- “High” and “low” levels of computer vision can be identified.
- Processing moves from digital manipulation, through pre-processing, segmentation, and recognition to understanding — but these processes may be simultaneous and co-operative.
- An understanding of the notions of heuristics, a priori knowledge, syntax, and semantics is necessary.
- The vision literature is large and growing; books may be specialized, elementary, or advanced.
- A knowledge of the research literature is necessary to stay up to date with the topic.

Summary

- Human vision is natural and seems easy; computer mimicry of this is difficult.
- We might hope to examine pictures, or sequences of pictures, for quantitative and qualitative analysis.
- Many standard and advanced AI techniques are relevant.
- “High” and “low” levels of computer vision can be identified.
- Processing moves from digital manipulation, through pre-processing, segmentation, and recognition to understanding — but these processes may be simultaneous and co-operative.
- An understanding of the notions of heuristics, a priori knowledge, syntax, and semantics is necessary.
- The vision literature is large and growing; books may be specialized, elementary, or advanced.
- A knowledge of the research literature is necessary to stay up to date with the topic.

Summary

- Human vision is natural and seems easy; computer mimicry of this is difficult.
- We might hope to examine pictures, or sequences of pictures, for quantitative and qualitative analysis.
- Many standard and advanced AI techniques are relevant.
- “High” and “low” levels of computer vision can be identified.
- Processing moves from digital manipulation, through pre-processing, segmentation, and recognition to understanding — but these processes may be simultaneous and co-operative.
- An understanding of the notions of heuristics, a priori knowledge, syntax, and semantics is necessary.
- The vision literature is large and growing; books may be specialized, elementary, or advanced.
- A knowledge of the research literature is necessary to stay up to date with the topic.

Summary

- Human vision is natural and seems easy; computer mimicry of this is difficult.
- We might hope to examine pictures, or sequences of pictures, for quantitative and qualitative analysis.
- Many standard and advanced AI techniques are relevant.
- “High” and “low” levels of computer vision can be identified.
- Processing moves from digital manipulation, through pre-processing, segmentation, and recognition to understanding — but these processes may be simultaneous and co-operative.
- An understanding of the notions of heuristics, a priori knowledge, syntax, and semantics is necessary.
- The vision literature is large and growing; books may be specialized, elementary, or advanced.
- A knowledge of the research literature is necessary to stay up to date with the topic.

Summary

- Human vision is natural and seems easy; computer mimicry of this is difficult.
- We might hope to examine pictures, or sequences of pictures, for quantitative and qualitative analysis.
- Many standard and advanced AI techniques are relevant.
- "High" and "low" levels of computer vision can be identified.
- Processing moves from digital manipulation, through pre-processing, segmentation, and recognition to understanding — but these processes may be simultaneous and co-operative.
- An understanding of the notions of heuristics, a priori knowledge, syntax, and semantics is necessary.
- The vision literature is large and growing; books may be specialized, elementary, or advanced.
- A knowledge of the research literature is necessary to stay up to date with the topic.

Summary

- Human vision is natural and seems easy; computer mimicry of this is difficult.
- We might hope to examine pictures, or sequences of pictures, for quantitative and qualitative analysis.
- Many standard and advanced AI techniques are relevant.
- “High” and “low” levels of computer vision can be identified.
- Processing moves from digital manipulation, through pre-processing, segmentation, and recognition to understanding — but these processes may be simultaneous and co-operative.
- An understanding of the notions of heuristics, a priori knowledge, syntax, and semantics is necessary.
- The vision literature is large and growing; books may be specialized, elementary, or advanced.
- A knowledge of the research literature is necessary to stay up to date with the topic.

Summary

- Human vision is natural and seems easy; computer mimicry of this is difficult.
- We might hope to examine pictures, or sequences of pictures, for quantitative and qualitative analysis.
- Many standard and advanced AI techniques are relevant.
- "High" and "low" levels of computer vision can be identified.
- Processing moves from digital manipulation, through pre-processing, segmentation, and recognition to understanding — but these processes may be simultaneous and co-operative.
- An understanding of the notions of heuristics, a priori knowledge, syntax, and semantics is necessary.
- The vision literature is large and growing; books may be specialized, elementary, or advanced.
- A knowledge of the research literature is necessary to stay up to date with the topic.

Summary

- Human vision is natural and seems easy; computer mimicry of this is difficult.
- We might hope to examine pictures, or sequences of pictures, for quantitative and qualitative analysis.
- Many standard and advanced AI techniques are relevant.
- "High" and "low" levels of computer vision can be identified.
- Processing moves from digital manipulation, through pre-processing, segmentation, and recognition to understanding — but these processes may be simultaneous and co-operative.
- An understanding of the notions of heuristics, a priori knowledge, syntax, and semantics is necessary.
- The vision literature is large and growing; books may be specialized, elementary, or advanced.
- A knowledge of the research literature is necessary to stay up to date with the topic.

Summary

- Human vision is natural and seems easy; computer mimicry of this is difficult.
- We might hope to examine pictures, or sequences of pictures, for quantitative and qualitative analysis.
- Many standard and advanced AI techniques are relevant.
- “High” and “low” levels of computer vision can be identified.
- Processing moves from digital manipulation, through pre-processing, segmentation, and recognition to understanding — but these processes may be simultaneous and co-operative.
- An understanding of the notions of heuristics, a priori knowledge, syntax, and semantics is necessary.
- The vision literature is large and growing; books may be specialized, elementary, or advanced.
- A knowledge of the research literature is necessary to stay up to date with the topic.

References I

- Marr, D. (1982).
  Vision.
  Freeman.
- Russ, J. C. (1995).
  The Image Processing Handbook, 2nd ed.
  CRC Press.
- Zhao, K. and Zhong, X. (1982).
  Optics.
  Peking University Press, Beijing.
  2 volumes, in Chinese.