# Advanced Graphics and Image Processing 

# Computer Science Tripos Part 2 <br> MPhil in Advanced Computer Science 

Michaelmas Term 2022／2023

Department of
Computer Science
and Technology
The Computer Laboratory
William Gates Building
15 JJ Thomson Avenue
Cambridge
CB3 OFD

This handout includes copies of the slides that will be used in lectures and more detailed notes on the selected topics. These notes do not constitute a complete transcript of all the lectures and they are not a substitute for text books. They are intended to give a reasonable synopsis of the subjects discussed, but they give neither complete descriptions nor all the background material.

Material is copyright © Rafał Mantiuk, 2015-2022, except where otherwise noted.

All other copyright material is made available under the University's licence. All rights reserved.


1

Where are graphics and image processing heading?


3


5

## Stride

- Calculating the pixel component index in memory
- For row-major order (grayscale)

$$
i(x, y)=x+y \cdot n_{c o l s}
$$

- For column-major order (grayscale)

$$
i(x, y)=x \cdot n_{\text {rows }}+y
$$

- For interleaved row-major (colour)

$$
i(x, y, c)=x \cdot 3+y \cdot 3 \cdot n_{\text {cols }}+c
$$

- General case

$$
i(x, y, c)=x \cdot s_{x}+y \cdot s_{y}+c \cdot s_{c}
$$

where $s_{x}, s_{y}$ and $s_{c}$ are the strides for the $\mathrm{x}, \mathrm{y}$ and colour dimensions

[^0]6

## Padded images and stride

- Sometimes it is desirable to "pad" image with extra pixels
b for example when using operators that need to access pixels outside the image border
- Or to define a region of interest (ROI)


How to address pixels for such an image and the ROI?
> 7
7

## Pixel (PIcture ELement)

- Each pixel (usually) consist of three values describing the color
(red, green, blue)
- For example
- $(255,255,255)$ for white
- $(0,0,0)$ for black
- $(255,0,0)$ for red
- Why are the values in the 0-255 range?
- Why red, green and blue? (and not cyan, magenta, yellow)
- How many bytes are needed to store 5MPixel image? (uncompressed)
- 9

9

10

- Truecolor $-2^{24}=16,8$ million colors ( 3 bytes)
- Deepcolor - even more colors (>= 4 bytes)

Pixel formats, bits per pixel, bit-depth

- Grayscale - single color channel, 8 bits (I byte)
- Highcolor $-2^{16}=65,536$ colors ( 2 bytes)

- For row-major, interleaved
b $s_{x}=$ ?
- $s_{y}=$ ?
- $s_{c}=$ ?
- 8

8


## Image - 2D function

- Image can be seen as a function $I(x, y)$, that gives intensity value for any given coordinate ( $x, y$ )

- 13

13

## What is a pixel?

```
- A pixel is not
- a box
p a disk
- a teeny light
```



- A pixel is a point
- it has no dimension
- it occupies no area
b it cannot be seen
b it has coordinates

- A pixel is a sample

From: $\mathrm{http}: / /$ groups.csail.mit.edu/graphics/classes/6.837/F01/Lecture05/lecture05.pdf
15

## Resampling

- Some image processing operations require to know the colors that are in-between the original pixels

- What are those operations?
- How to find these resampled pixel values?
> 17

17

## Sampling an image

- The image can be sampled on a rectangular sampling grid to yield a set of samples. These samples are pixels.

- 14

14

## Sampling and quantization

- The physical world is described in terms of continuous quantities
- But computers work only with discrete numbers
- Sampling - process of mapping continuous function to a discrete one
- Quantization - process of mapping continuous variable to a discrete one


- 16

16

Example of resampling: magnification


18

Example of resampling:
scaling and rotation


- 19

19
(Bi)Linear interpolation (resampling)

- Linear - ID
- Bilinear - 2D

- 21

21

Bi-linear interpolation


Given the pixel values:
$I\left(x_{1}, y_{1}\right)=A$
$I\left(x_{2}, y_{1}\right)=B$
$I\left(x_{1}, y_{2}\right)=C$
$I\left(x_{2}, y_{2}\right)=D$

Calculate the value of a pixel $I(x, y)=$ ? using bi-linear interpolation.
Hint: Interpolate first between A and B, and between C and D, then interpolate between these two computed values.

[^1]How to resample?

- In ID: how to find the most likely resampled pixel value knowing its two neighbors?

- 20

20
(Bi)cubic interpolation (resampling)


22


Advanced Graphics \& Image Processing
Introduction to Image Processing
Part 2/2 - Point ops, filters and pyramids
Rafał Mantiuk
Computer Laboratory, University of Cambridge

24


25

Pixel precision for image processing

- Given an RGB image, 8-bit per color channel (uchar)
- What happens if the value of 10 is subtracted from the pixel value of 5 ?
- $250+10=$ ?
- How to multiply pixel values by I.5 ?
a) Using floating point numbers
b) While avoiding floating point numbers

27

27

Image matting and compositing


29

## Point operators

- Modify each pixel independent from one another
- The simplest case: multiplication and addition


26


28

Transparency, alpha channel

- RGBA - red, green, blue, alpha
- alpha $=0$ - transparent pixel
- alpha = I - opaque pixel
- Compositing
- Final pixel value
$P=\alpha C_{\text {pixel }}+(1-\alpha) C_{\text {background }}$
- Multiple layers:
$P_{0}=C_{\text {background }}$
$P_{i}=\alpha_{i} C_{i}+\left(1-\alpha_{i}\right) P_{i-1} \quad i=1 . . N$

- 30

30


31


33


35

## Histogram equalization

- Pixels are non-uniformly distributed across the range of values

- Would the image look better if we uniformly distribute pixel values (make the histogram more uniform)?
- How can this be done?
- 32

32

## Linear filtering

- Output pixel value is a weighted sum of neighboring pixels


Sum over neighboring pixels, e.g. $k=-1,0,1, j=-1,0,1$ for $3 \times 3$ neighborhood
compact notation $g=f * h$
Convolution
operation
34


36

What is the computational cost of the convolution？

$$
g(i, j)=\sum_{k, l} f(i-k, j-l) h(k, l)
$$

－How many multiplications do we need to do to convolve $100 \times 100$ image with $9 \times 9$ kernel ？
－The image is padded，but we do not compute the values for the padded pixels
－ 37
37

## Examples of separable filters

－Box filter：

－Gaussian filter：

$$
G(x, y ; \sigma)=\frac{1}{2 \pi \sigma^{2}} e^{-\frac{x^{2}+y^{2}}{2 \sigma^{2}}}
$$

－What are the corresponding ID components of this separable filter（ $u(x)$ and $v(y)$ ）？

$$
G(x, y)=u(x) \cdot v(y)
$$

－ 39
39
39

## Why＂linear＂filters ？

－Linear functions have two properties：
－Additivity：$f(x)+f(y)=f(x+y)$
＂Homogenity：$f(a x)=a f(x)$（where＂$f$＂is a linear function）
－Why is it important？
－Linear operations can be performed in an arbitrary order $\operatorname{blur}(a F+b)=a \operatorname{blur}(F)+b$
－Linearity of the Gaussian filter could be used to improve the performance of your image processing operation
－This is also how separable filters work：


41

## Separable kernels

－Convolution operation can be made much faster if split into two separate steps：
－I）convolve all rows in the image with a ID filter
2）convolve columns in the result of I）with another ID filter
－But to do this，the kernel must be separable

$$
\left[\begin{array}{lll}
h_{11} & h_{12} & h_{13} \\
h_{21} & h_{22} & h_{23} \\
h_{31} & h_{32} & h_{33}
\end{array}\right]=\left[\begin{array}{l}
u_{1} \\
u_{2} \\
u_{3}
\end{array}\right] \cdot\left[\begin{array}{lll}
v_{1} & v_{2} & v_{3}
\end{array}\right]
$$

$$
\vec{h}=\vec{u} \cdot \vec{v}
$$

－ 38
38

## Unsharp masking

－How to use blurring to sharpen an image ？

$\qquad$

$g_{\text {sharp }}=f+\gamma\left(f-h_{\text {blur }} * f\right)$
－ 40
40

Operations on binary images
－Essential for many computer vision tasks
ままままうう

Binary image can be constructed by thresholding a grayscale image

$$
\theta(f, c)= \begin{cases}1 & \text { if } f \geq c \\ 0 & \text { else }\end{cases}
$$

42


- Set the pixel to the maximum value of the neighboring pixels within the structuring element
- What could it be useful for ?
> 43
43


45

Binary morphological filters: formal definition

$S$ - size of structuring element (number of 1 s in the SI )

- dilation: $\operatorname{dilate}(f, s)=\theta(c, 1)$;
- erosion: $\operatorname{erode}(f, s)=\theta(c, S)$;
$\theta(f, c)= \begin{cases}1 & \text { if } f \geq c, \\ 0 & \text { else },\end{cases}$
- majority: $\operatorname{maj}(f, s)=\theta(c, S / 2)$;
- opening: $\operatorname{open}(f, s)=\operatorname{dilate}(\operatorname{erode}(f, s), s)$;
- closing: $\operatorname{close}(f, s)=\operatorname{erode}(\operatorname{dilate}(f, s), s)$.
- 47

Morphological filters: erosion

b) Structuring

Structur
$\mathrm{x}=$ origin

c) Image after erosion original in dashes

- Set the value to the minimum value of all the neighboring pixels within the structuring element
- What could it be useful for?
> 44
44

Morphological filters: closing

b) Structuring element;
$\mathrm{x}=$ origin
c) Image after closing = dilation followed by dashes.

- Dilation followed by erosion
- What could it be useful for ?


## Multi-scale image processing (pyramids)

- Multi-scale processing operates on an image represented at several sizes (scales)
- Fine level for operating on small details
- Coarse level for operating on large features
- Example:
- Motion estimation
- Use fine scales for objects moving slowly
- Use coarse scale for objects moving fast
- Blending (to avoid sharp boundaries)

```
48
```

48


49


51



50


52


54
$\square$
References

- SZELISKI, R. 20I0. Computer Vision:Algorithms and Applications. Springer-Verlag New York Inc.
- Chapter 3
- http:///szeliski.org/Book
> 55


1


3


5


2


4

How to make the bilateral filter fast?

- A number of approximations have been proposed
- Combination of linear filters [Durand \& Dorsey 2002, Yang et al. 2009]
- Bilateral grid [Chen et al. 2007]
- Permutohedral lattice [Adams et al. 20I0]
- Domain transform [Gastal \& Oliveira 201 I]


6


7


9


11

Joint bilateral filter: Flash / no-flash


8


10

Gradient Domain compositing

- Compositing [Wang et al. 2004]


[^2]

13

## Processing gradient field

- Typically, gradient magnitudes are modified while gradient direction (angle) remains the same

- 15

15


Gradient field reconstruction: derivation

- The minimization problem is given by:
$\underset{I}{\arg \min } \sum_{x, y}\left[\left(I_{x+1, y}-I_{x, y}-G_{x, y}^{(x)}\right)^{2}+\left(I_{x, y+1}-I_{x, y}-G_{x, y}^{(y)}\right)^{2}\right]$
- After equating derivatives over pixel values to 0 we get:
- Derivation done in the lecture
$I_{x-1, y}+I_{x+1, y}+I_{x, y-1}+I_{x, y+1}-4 I_{x, y}=G_{x, y}^{(x)}-G_{x-1, y}^{(x)}+G_{x, y}^{(y)}-G_{x, y-1}^{(y)}$


Laplace operator for $3 \times 3$ image


- 19


21

Weighted gradients - matrix notation (1)

- The objective function:

$$
\underset{I}{\arg \min } \sum_{x, y}\left[w_{x, y}^{(x)}\left(I_{x+1, y}-I_{x, y}-G_{x, y}^{(x)}\right)^{2}+w_{x, y}^{(y)}\left(I_{x, y+1}-I_{x, y}-G_{x, y}^{(y)}\right)^{2}\right]
$$

- In the matrix notation (without weights for now):

$$
\underset{I}{\arg \min }\left\|\left[\begin{array}{c}
\nabla_{x} \\
\nabla_{y}
\end{array}\right] I-\left[\begin{array}{l}
G^{(x)} \\
G^{(y)}
\end{array}\right]\right\|^{2}
$$

- Gradient operators (for $3 \times 3$ pixel image):


- 23

Weighted gradients - matrix notation (2)

- The objective function again:

$$
\underset{I}{\arg \min }\left\|\left[\begin{array}{l}
\nabla_{x} \\
\nabla_{y}
\end{array}\right] I-\left[\begin{array}{l}
G^{(x)} \\
G^{(y)}
\end{array}\right]\right\|^{2}
$$

- Such over-determined least-square problem can be solved using pseudo-inverse:

$$
\left[\begin{array}{ll}
\nabla_{x}^{\prime} & \nabla_{y}^{\prime}
\end{array}\right]\left[\begin{array}{c}
\nabla_{x} \\
\nabla_{y}
\end{array}\right] I=\left[\begin{array}{ll}
\nabla_{x}^{\prime} & \nabla_{y}^{\prime}
\end{array}\right]\left[\begin{array}{l}
G^{(x)} \\
G^{(y)}
\end{array}\right]
$$

- Or simply:

$$
\left(\nabla_{x}^{\prime} \nabla_{x}+\nabla_{y}^{\prime} \nabla_{y}\right) I=\nabla_{x}^{\prime} G^{(x)}+\nabla_{y}^{\prime} G^{(y)}
$$

- With weights:

$$
\left(\nabla_{x}^{\prime} W \nabla_{x}+\nabla_{y}^{\prime} W \nabla_{y}\right) I=\nabla_{x}^{\prime} W G^{(x)}+\nabla_{y}^{\prime} W G^{(y)}
$$

$$
1>24
$$

24

## WLS filter: Edge stopping filter by optimization

- Weighted-least-squares optimization

$$
\underset{u}{\operatorname{argmin}} \sum_{p} \begin{gathered}
\begin{array}{c}
\text { Make reconstructed image } u \\
\text { possibly close to input } g
\end{array}
\end{gathered} \underbrace{\begin{array}{c}
\text { Smooth out the image by making } \\
\text { partial derivatives close to } 0
\end{array}}_{\left(\left(u_{p}-g_{p}\right)^{2}+\lambda\left(a_{x, p}(g)\left(\frac{\partial u}{\partial x}\right)_{p}^{2}+a_{y, p}(g)\left(\frac{\partial u}{\partial y}\right)_{p}^{2}\right)\right)}
$$



$$
a_{x, p}(g)=\frac{1}{\left|\frac{\partial u}{\partial x}(g)\right|^{\alpha}+\epsilon}
$$

- [Farbman, Z., Fattal, R., Lischinski, D., \& Szeliski, R. (2008). Edge-preserving decompositions for multi-scale tone and detail manipulation. ACM SIGGRAPH 2008, I-10.]
- 25

25

27
Spatially varying smoothing - less smoothing near the edges



$$
\min _{f} \iint_{\Omega}|\nabla f-\mathbf{v}|^{2} \quad \text { subject to: }\left.\quad f\right|_{\partial \Omega}=\left.f^{*}\right|_{\partial \Omega}
$$

- Reconstruct unknown values $f$ given a source guidance gradient field $v$ and the boundary conditions $\left.f\right|_{\partial \Omega}=\left.f^{*}\right|_{\partial \Omega}$
- [Pérez, P., Michel Gangnet, \& Blake, A. (2003). Poisson Image Editing. ACM Transactions on Graphics, 3(22), 313-3 I8. https://doi.org/I0.| |45/882262.882269]
- 26


## Gradient Domain: applications

- More applications:
- Lightness perception (Retinex) [Horn 1974]
- Matting [Sun et al. 2004]
, Color to gray mapping [Gooch et al. 2005]
p Video Editing [Perez at al. 2003,Agarwala et al. 2004]
p Photoshop’s Healing Brush [Georgiev 2005]
$-28$
28


## References

- F. Durand and J. Dorsey, "Fast bilateral filtering for the display of high-dynamic-range images," ACM Trans. Graph., vol. 2 I, no. 3, pp. 257-266, Jul. 2002.
- E.S. L. Gastal and M. M. Oliveira, "Domain transform for edge-aware image and video processing," ACM Trans. Graph., vol. 30, no. 4, p. I, Jul. 201 I
- Patrick Pérez, Michel Gangnet, and Andrew Blake. 2003. Poisson image editing. ACM Trans. Graph. 22, 3 (July 2003), 313-318. DOI
http://dx.doi.org/l0.||45/882262.882269
- Zeev Farbman, Raanan Fattal, Dani Lischinski, and Richard Szeliski. 2008. Edgepreserving decompositions for multi-scale tone and detail manipulation. ACM Trans. Graph. 27, 3,Article 67 (August 2008), 10 pages. DOI:
https://doi.org/I0.1145/I3606|2.1360666


1

## Parallel Software - SPMD

- In the vector addition example, each chunk of data could be executed as an independent thread
- On modern CPUs, the overhead of creating threads is so high that the chunks need to be large
- In practice, usually a few threads (about as many as the number of CPU cores) and each is given a large amount of work to do
- For GPU programming, there is low overhead for thread creation, so we can create one thread per loop iteration

$$
3 \quad \text { From: OpenCL } 1.2 \text { University Kit - http://developer.amd.com/partners/university-programs/ }
$$

3

## Parallel programming frameworks

- These are some of more relevant frameworks for creating parallelized code


5

Single Program Multiple Data (SPMD)


Multiple copies of the same program execute on different data in parallel

multiple copies of the multiple copies of the
same program run on different chunks of the data


4

## OpenCL

- OpenCL is a framework for writing parallelized code for CPUs, GPUs, DSPs, FPGAs and other processors
- Initially developed by Apple, now supported by AMD, IBM, Qualcomm, Intel and Nvidia
- Versions
- Latest: OpenCL 3.0
- OpenCL C++ kernel language
- SPIR-V as intermediate representation for kernels
$\square$ Vulcan uses the same Standard Portable Intermediate Representation - AMD, Intel, Nvidia
- Mostly supported: OpenCL I. 2
, OSX, older GPUs

6

## OpenCL platforms and drivers

- To run OpenCL code you need:

- Included in the OS
- Installable Client Driver
- From Nvidia, Intel, etc.
- This applies to Windows and Linux, only one platform on Mac
- To develop OpenCL code you need:
- OpenCL headers/libraries
- Included in the SDKs
$\square$ Nvidia - CUDA Toolkit
$\square$ Intel OpenCL SDK
But lightweight options are also available

Example: Step 1 - Select device


```
//get all platforms (drivers)
std::vector<cl::Platform> all_platforms;
C1::Platform:get(&a11_platforms);
    std::cout << " No platforms found. Check OpenCL installation\\n";
    exit(1);
cl::Platform defau
```



```
//get default device of the default platform
std::vector<cl::Device> all_devices;
default_platform.getDevices(CL_DEVICE_TYPE_ALL, &all_devices)
    l
        exit(1);
cl::Device default device = all devices[日]
std::cout << "Using device: " << default_device.getInfo<CL_DEVICE_NAME>() << "\n";
```

- 

9

Example: Step 3 - Create Buffers and copy memory

| CreateBuffers $\longrightarrow \xrightarrow[\begin{array}{c}\text { Create } \\ \text { Queue }\end{array}]{\longrightarrow$ Enqueue  <br>  Memory Copy $}$ |
| :---: |
|  |
| $\begin{aligned} & \text { int } A[]=\{0,1,2,3,4,5,6,7,8,9\} ; \\ & \text { int } \operatorname{B}[]=\{(1,2, \theta, 1,2,9,1,2, \theta\} ; \end{aligned}$ |
| //create queue to which we will push commands for the device cl::CommandQueue queue(context, default_device); |
| $/ / u r i$ te arrays $A$ and $\theta$ to the device <br> queue. enqueueWriteBuffer(buffer_A, CL_TRUE, $\theta$, sizeof(int) * $10, \mathrm{~A}$ ); queue enqueueWriteBuffer (buffer_B, CL_TRUE, $B$, sizeof(int) * $10, B$ ); |
| - |

Example: Step 2 - Build program

c1: : Context context(\{ default_device \});
c1::Program::Sources sources;
td: : .
"_kernel void simple_add(__global const int* A, _global const int* B, _global int* C) \{" int index $=$ get_global_id $(0)$ )"
$[$ index $]=A[$ index $]+B[$ index $] ;$
$\mathrm{C}[$ index $]=\mathrm{A}[$ index $]+\mathrm{B}$ [index];
sources.push_back(\{ kernel_code.c_str(), kernel_code.length() \});
c1::Program program(context, sources);
cl: $:$ Pr
try
p
pr
try program.build(\{ default_device \})
catch (cl::Error err) i
 $\underset{\text { exit(1); }}{\substack{\text { progra } \\ \hline}}$
\}

Example: Step 4 - Execute Kernel and retrieve the results

| CreateKernel $\longrightarrow \quad$Set Ke <br> Argum | Enqueue Kernel | $\begin{gathered} \text { Enqueue } \\ \text { memory copy } \end{gathered}$ |
| :---: | :---: | :---: |
| c1: Kernel kernel(program, |  |  |
| kernel. setArg ( 0 , buffer_A) ; <br> kernel.setArg(1, buffer_B) <br> kernel.setArg ( 2 , buffer_C) <br> queue. |  |  |
| int C[19]; <br> //read result C from the device to array C queue.enqueuereadibuffer (buffer_c, cL_TRUE, e, sizeof(int) * 10, C); queue. finish(); |  |  |
|  | Our Kernel was |  |
| $3$ <br> std::cout << std::end1; | _read_only const int* $A$, <br> -writeonly inte c) <br> $\mathrm{C}[$ index] A [index] B [index]: |  |

## OpenCL API Class Diagram

- Platform - Nvidia CUDA
- Device - GeForce 1080
- Program - collection of kernels
- Buffer or Image - device memory
- Sampler - how to interpolate values for Image
- Command Queue - put a sequence of operations there
Event - to notify that something has been done


13

## Execution model

- Each kernel executes on ID, 2D or 3D array (NDRange)
- The array is split into work-groups
- Work items (threads) in each work-group share some local memory
- Kernel can querry
, get_global_id(dim)
, get_group_id(dim)
' get_local_id(dim)
- Work items are not bound to any memory entity (unlike GLSL shaders)



## Programming model

- Data parallel programming
, Each NDRange element is assigned to a work-item (thread)
, Each kernel can use vector-types of the device (float4, etc.)
- Task-parallel programming
, Multiple different kernels can be executed in paralle
- Command queue
clCreateCommandQueue( cl context context cl_device_id device,
cl_command_queue_properties properties,
cl_int* errcode_ret)
- Provides means to both synchronize kernels and execute them in paralle


## Memory model

- Host memory
- Usually CPU memory, device does not have access to that memory
, Global memory [__global]
, Device memory, for storing large data
- Constant memory [__constant]
- Local memory [__local]
, Fast, accessible to all work-items (threads) within a workgroup
, Private memory [__private]
, Accessible to a single work-item (thread)



## Platform model

- The host is whatever the OpenCL library runs on , Usually $\times 86$ CPUs for both NVIDIA and AMD
- Devices are processors that the library can talk to
- CPUs, GPUs, DSP,s and generic accelerators


## - For AMD

- All CPUs are combined into a single device (each core is a compute unit
and processing element)
- Each GPU is a separate device

- 14

14



19


21

## Thread Mapping

- Thread mapping I: with an MxN index space, the kernel would be


Mapping for C

| 0 | 4 | 8 | 12 |
| :--- | :--- | :--- | :--- |

for $(\mathrm{i} 3=0$; $\mathrm{i} 3<\mathrm{P}$; $\mathrm{i} 3++$ )
$\mathrm{C}[\mathrm{tx}][\mathrm{ty}]+=\mathrm{A}[\mathrm{tx}][\mathrm{i} 3] * \mathrm{~B}[\mathrm{i} 3][\mathrm{ty}] ;$

| 1 | 5 | 9 | 13 |
| :---: | :---: | :---: | :---: |
| 2 | 6 | 10 | 14 |
| 3 | 7 | 11 | 1 |

- Thread mapping 2: with an NxM index space, the kernel would be

| int $\mathrm{tx}=$ get_global_id | Mapping for C |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
| int ty = get_global_id (1); | 0 | 1 | 2 | 3 |
| for ( $\mathrm{i} 3=0 ; \mathrm{i} 3<\mathrm{P} ; \mathrm{i} 3++$ ) | 4 | 5 | 6 | 7 |
| $\mathrm{C}[\mathrm{ty}][\mathrm{tx}]+=\mathrm{A}[\mathrm{ty}][\mathrm{i} 3] * \mathrm{~B}[\mathrm{i} 3][\mathrm{tx}]$; | 8 | 9 | 10 | 11 |
|  | 12 | 13 | 14 | 15 |

- Both mappings produce functionally equivalent versions of the program

23 From: OpenCL 1.2 University Kit-htip://developeramd. com/parthers/university-proaram

## Thread Mapping

- Consider a serial matrix multiplication algorithm

```
for(il=0; il < M; il++)
    for(i2=0; i2< N; i2++)
    for(i3=0; i3<P; i3++)
        C[i1][i2] += A[i1][i3]*B[i3][i2];
```

- This algorithm is suited for output data decomposition
b We will create $N$ x $M$ threads
- Effectively removing the outer two loops
- Each thread will perform $P$ calculations
- The inner loop will remain as part of the kernel
- Should the index space be MxN or NxM ?

22 From: OpenCL 1.2 University Kit - http://developer.amd. com/partners Suniversity-program
22

## Thread Mapping

- This figure shows the execution of the two thread mappings on NVIDIA GeForce 285 and 8800 GPUs

- Notice that mapping 2 is far superior in performance for both GPUs

[^3]
## Thread Mapping

- The discrepancy in execution times between the mappings is due to data accesses on the global memory bus
- Assuming row-major data, data in a row (i.e., elements in adjacent columns) are stored sequentially in memory
- To ensure coalesced accesses, consecutive threads in the same wavefront should be mapped to columns (the second dimension) of the matrices
b This will give coalesced accesses in Matrices B and C
- For Matrix A, the iterator i3 determines the access pattern for rowmajor data, so thread mapping does not affect it

25 From: OpenCL 1.2 University Kit - htto://developer.amd.com/parters/university-programs

## Reduction

- GPU offers very good performance for tasks in which the results are stored independently
- Process N data items and store in N memory location
- But many common operations require reducing $N$ values into I or few values sum, min, max, prod, min, histogram, ..
- Those operations require an efficient implementation of reduction
, The following slides are based on AMD's OpenCL" Optimization Case Study: Simple Reductions , http://developer.amd.com/resources/articles-whitepapers/opencl-optimization-case-study-simple-reductions/ $>$

27

Reduction tree for the min operation


28

## Multistage reduction

- The local memory is usually limited (e.g. 50kB), which restricts the maximum size of the array that can be processed
- Therefore, for large arrays need to be processed in multiple stages
, The result of a local memory reduction is stored in the array and then this array is reduced


Reduction execution times on CPU/GPU


- Different reduction algorithm may be optimal for CPU and GPU
- This can also vary from one GPU to another
, The results from: http://developer.amd.com/resources/articles-whitepapers/opencl-optimization-case-study-simple-reductions/
$\rightarrow$


## Better way?

Halide - a language for image processing and computational photography
, http://halide-lang.org/

- Code written in a high-level language, then translated to x86/SSE, ARM, CUDA, OpenCL
- The optimization strategy defined separately as a schedule
- Auto-tune software can test thousands of schedules and choose the one that is the best for a particular platform
- (Semi-)automatically find the best trade-offs for a particular platform
- Designed for image processing but similar languages created for other purposes



## OpenCL resources

, https://www.khronos.org/registry/OpenCL/
b Reference cards

- Google:"OpenCLAPI Reference Card"
- AMD OpenCL Programming Guide
http://developer.amd.com/wordpress/media/2013/07/AMD_Accelerated_Parallel_Processing_OC L_Programming_Guide-2013-06-21.pdf


1


3


5

## Baked / precomputed illumination

- We need:
- Geometry + textures + (light maps)
- No need to scan and model materials
- Much faster rendering
- simplified shading

> 4
4


## Light fields + depth

- We need:
, Depth map
, Images of the object/scene
- Camera

- We can use camera-captured images
, View-dependent shading
- Depth-map can be computed using multi-view stereo techniques
, CV methods can be unreliable
- No relighting


## Light fields

- We need:
- Images of the scene
- Or a microlens image
- Camera
- As light fields + depth but
- No geometry, no need for any 3D reconstruction
, Photographs are repprojected on the plane
- Requires massive number of images for good quality
> 7



## Multi-plane images (MPI)

- We need:
- Images of the scene + camera poses
- Each plane: RGB + alpha
- Decomposition formulated as an optimization problem
- Differential rendering
- Only front view
[1] Mildenhall, et al. "Local Light Field Fusion." ACM Transactions on Graphics 38, no. 4 (July 12, 2019): 1-

14. https://doi. orga/10. 1145/3306342. 3322980 with Neural Basis Expansion." In CVPR, 8530-39. IEEE, 2021. https://doi.org/10.1109/CVPR46437.2021.00843 > 8


9

Planar 4D light field


- 11

From a plenoptic function to a light field

- Plenoptic function - describes all possible rays in a 3D space
- Function of position $(x, y, z)$ and ray direction $(\theta, \phi)$
- But also wavelength $\lambda$ and time $t$
p Between 5 and 7 dimensions

- But the number of dimensions can be reduced if
, The camera stays outside the convex hull of the object
- The light travels in uniform medium
- Then, radiance $L$ remains the same along the ray (until the ray hits an object)
, This way we obtain a 4D light field
- 10

10

Refocusing and view point adjustment


12

Depth estimation from light field

- Passive sensing of depth
- Light field captures multiple depth cues
- Correspondance (disparity) between the views
- Defocus
- Occlusions

From: Ting-Chun Wang, Alexei A. Efros, Ravi Ramamoorthi, The IEEE International Conference
on Computer Vision (ICCV), 2015, pp. 3487-3495
 13

Light field image - with microlens array


15


17

Two methods to capture light fields

| Micro-lens array | Camera array |
| :--- | :--- |
| * Small baseline | , Large baseline |
| - Good for digital refocusing | * High resolution |
| - Limited resolution | , Rendering often requires |
|  | approximate depth |



14


16

Raytrix camera

- Similar technology to Lytro
- But profiled for
 computer vision applications


Stanford camera array


96 cameras

19


21


23

PiCam camera array module

- Array of $4 \times 4$ cameras on a single chip
- Each camera has its own lens and senses only one spectral colour band
- Optics can be optimized for that band
- The algorithm needs to reconstruct depth


20


22


24


25

Defocus blur is often desirable


To separate the object of interest from background - 27


Defocus blur is a strong depth cue


Light fields: two parametrisations (shown in 2D)


Position and slope (slope - tangent of the angle)


31


33


35

Lightfield - example



34


36

Light field rendering ( $1 / 3$ )


We want to render a scene (Blender monkey) as seen by camera K. We have a light field captured by a camera array. Each camera in the array has its aperture on plane C .

- 37


## Light field rendering (3/3)

The rays from the camera need to be projected on the focal
plane $F$. The objects on the focal plane will be sharp, and
the objects in front or behind that plane will be blurry (ghosted), as in a traditional camera.


If we have a proxy geometry, we can project on that geometry instead - the rendered image will be less ghosted/blurry

- 39


41

Light field rendering (2/3)
\(\left.$$
\begin{array}{l}\begin{array}{l}\text { From the view point of } \\
\text { camera } \mathrm{K}\end{array}
$$ <br>
\hline Each camera in the <br>
array provides <br>
accurate light <br>
measurements only for <br>
the rays originating <br>
from its pinhole <br>

aperture.\end{array}\right\}\)| The missing rays can |
| :--- |
| be either interpolated |
| (reconstructed) or |
| ignored. |

## Intuition behind light field rendering

- For large virtual aperture (use all cameras in the array)
- Each camera in the array captures the scene
- Then, each camera projects its image on the focal plane F
- The virual camera K captures the projection
- For small virtual aperture (pinhole)
- For each ray from the virtual camera , interpolate rays from 4 nearest camera images
- Or use the nearest-neighbour ray


40


42

## Finding homographic transformation 1/3

- For the pixel coordinates $\boldsymbol{p}_{k}$ of the virtual camera K, we want to find the corresponding coordinates $\boldsymbol{p}_{i}$ in the camera array image
- Given the world 3D coordinates of a point $\boldsymbol{w}$ :
 $\boldsymbol{p}_{i}=K P \boldsymbol{V}_{i} \boldsymbol{w}$ camera matrix $\left.\begin{array}{c}\text { Projection } \\ \text { matrix }\end{array}\right] \quad \begin{gathered}\text { View } \\ \text { matrix }\end{gathered}$

$\left[\begin{array}{l}x_{i} \\ y_{i} \\ w_{i}\end{array}\right]=\left[\begin{array}{ccc}f_{x} & 0 & c_{x} \\ 0 & f_{y} & c_{y} \\ 0 & 0 & 1\end{array}\right]\left[\begin{array}{llll}1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0\end{array}\right]\left[\begin{array}{cccc}v_{11} & v_{12} & v_{13} & v_{14} \\ v_{21} & v_{22} & v_{23} & v_{24} \\ v_{31} & v_{32} & v_{33} & v_{34} \\ 0 & 0 & 0 & 1\end{array}\right]\left[\begin{array}{c}X \\ Y \\ Z \\ 1\end{array}\right]$

Finding homographic transformation $2 / 3$

- A homography between two views is usually found as:

$$
\begin{aligned}
\boldsymbol{p}_{K} & =\boldsymbol{K}_{K} \boldsymbol{P} \boldsymbol{V}_{K} \boldsymbol{w} \\
\boldsymbol{p}_{i} & =\boldsymbol{K}_{i} \boldsymbol{P} \boldsymbol{V}_{i} \boldsymbol{w}
\end{aligned}
$$

hence

$$
\boldsymbol{p}_{i}=\boldsymbol{K}_{i} \boldsymbol{P} \boldsymbol{V}_{i} \boldsymbol{V}_{K}^{-1} \boldsymbol{P}^{-1} \boldsymbol{K}_{K}^{-1} \boldsymbol{p}_{K}
$$

But, $K_{K} P V_{K}$ is not a square matrix and cannot be inverted

- To find the correspondence, we need to constrain 3D coordinates $\boldsymbol{w}$ to lie on the plane:

$$
\begin{aligned}
& \text { coordinates } \boldsymbol{w} \text { to lie on the plane: } \\
& \boldsymbol{N} \cdot\left(\boldsymbol{w}-\boldsymbol{w}_{F}\right)=0 \quad \text { or } \quad d=\left[\begin{array}{llll}
n_{x} & n_{y} & n_{z} & -\boldsymbol{N} \cdot \boldsymbol{w}_{F}
\end{array}\right]\left[\begin{array}{c}
X \\
Y \\
Z \\
1
\end{array}\right]
\end{aligned}
$$

$$
\text { D } 44
$$

43

The plane in
the camera coordinates
(not world coordinates)
3/3

- Then, we add the plane equation to the projection matrix

- Where $d_{i}$ is the distance to the plane (set to 0 )
- Hence

$$
\hat{\boldsymbol{p}}_{i}=\hat{K}_{i} \hat{P} V_{i} V_{K}^{-1} \hat{P}^{-1} \hat{K}_{K}^{-1} \hat{p}_{K}
$$

[^4]
## References

- Light fields
- Micro-lens array
- Ng, Ren and Levoy, Marc and Bredif, M. and D., \& Gene and Horowitz, Mark and Hanrahan, P. (2005). Light field photography with a hand-held plenoptic camera.
, Camera array
, Overbeck, R.S., ERickson, D., Evangelakos, D., Pharr, M., And DEBEVEC, P. 2018.A system for acquiring, processing, and rendering panoramic light field stills for virtual reality. ACM Transactions on Graphics 37,6, I-I5
+ ISAKSEN, A., McMILLAN, L.,AND GortLer, S.J. 2000. Dynamically reparameterized light fields. Proc of SIGGRAPH '00,ACM Press, 297306.


1


3


5

## Correlated colour temperature

- The temperature of a black body radiator that produces light most closely matching the particular source
- Examples:
- Typical north-sky light: 7500 K
- Typical average daylight: 6500 K
- Domestic tungsten lamp ( 100 to 200 W ): 2800 K
, Domestic tungsten lamp ( 40 to 60 W ): 2700 K
, Sunlight at sunset: 2000 K
- Useful to describe colour of the illumination (source of light)


6


7


9


11

## Reflectance

- Most of the light we see is reflected from objects
- These objects absorb a certain part of the light spectrum

Spectral reflectance of ceramic tiles


8


10


12

## Colour vision

- Cones are the photreceptors responsible for colour vision
, Only daylight, we see no colours when there is not enough light
- Three types of cones
, S - sensitive to short wavelengths
- $M$ - sensitive to medium wavelengths
- L - sensitive to long wavelengths


Sensitivity curves - probability that a photon of that wavelengths will be absorbed by a photoreceptor. S,M and $L$ curves are normalized in this plot.

Perceived light
, cone response $=$ sum $($ sensitivity $\times$ reflected light $)$


14

## Practical application of metamerism

- Displays do not emit the same light spectra as real-world objects
- Yet, the colours on a display look almost identical


15

## Metamers

- Even if two light spectra are different, they may appear to have the same colour
- The light spectra that appear to have the same colour are called metamers
- Example:



## Standard Colour Space CIE-XYZ

- CIE Experiments [Guild and Wright, 193I]
- Colour matching experiments
- Group ~I2 people with „normal" colour vision
, 2 degree visual field (fovea only)
, CIE 2006 XYZ
, Derived from LMS colour matching functions by Stockman \& Sharpe
- S-cone response differs the most from CIE I93I
- CIE-XYZ Colour Space
, Goals
- Abstract from concrete primaries used in an experiment
, All matching functions are positive
, Primary „Y" is roughly proportionally to achromatic response (luminance)


## Standard Colour Space CIE-XYZ



19


21

Achromatic/chromatic vision
mechanisms
> 23


## CIE chromaticity diagram

- chromaticity values are defined in terms of $x, y, z$

$$
x=\frac{X}{X+Y+Z}, \quad y=\frac{Y}{X+Y+Z}, \quad z=\frac{Z}{X+Y+Z} \quad x+y+z=1
$$

[^5]can be plotted as a 2D function

- pure colours (single wavelength) lie along the outer curve
, all other colours are a mix of pure colours and hence lie inside the curve
p points outside the curve do not exist as colours
$>20$


20

Achromatic/chromatic vision mechanisms


22
22

Achromatic/chromatic vision mechanisms


24


25


27

## Visible vs. displayable colours

- All physically possible and visible colours form a solid in the XYZ space
- Each display device can reproduce a subspace of that space
- A chromacity diagram is a projection of a slice taken from a 3D solid in XYZ space
- Colour Gamut - the solid in a colour space
- Usually defined in XYZ to be deviceindependent


[^6]
## Standard vs. High Dynamic Range

HDR cameras/formats/displays attempt capture/represent/reproduce (almost) all visible colours
, They represent scene colours and therefore we often call this representation scene-referred

- SDR cameras/formats/devices attempt to capture/represent/reproduce only colours of a standard sRGB colour gamut, mimicking the capabilities of CRTs monitors
, They represent display colours and therefore we often call this representation display-referred
| 30

From rendering to display


31


33

## Why is gamma needed?



- Gamma-corrected pixel values give a scale of brightness levels that is more perceptually uniform
- At least 12 bits (instead of 8 ) would be needed to encode each color channel without gamma correction
- And accidentally it was also the response of the CRT gun
- 35

From rendering to display


32

Display encoding for SDR: gamma

- Gamma correction is often used to encode luminance or tristimulus color values (RGB) in imaging systems (displays, printers, cameras, etc.)


34

Luma - gray-scale pixel value

- Luma - pixel "brightness" in gamma corrected units
$L^{\prime}=0.2126 R^{\prime}+0.7152 G^{\prime}+0.0722 B^{\prime}$
- $R^{\prime}, G^{\prime}$ and $B^{\prime}$ are gamma-corrected colour values
- Prime symbol denotes gamma corrected
- Used in image/video coding
- Note that relative luminance if often approximated with
$L=0.2126 R+0.7152 G+0.0722 B$
$=0.2126\left(R^{\prime}\right)^{\gamma}+0.7152\left(G^{\prime}\right)^{\gamma}+0.0722\left(B^{\prime}\right)^{\gamma}$
- $R, G$, and $B$ are linear colour values
- Luma and luminace are different quantities despite similar formulas
>
36


## Standards for display encoding



37

How to transform between
RGB colour spaces?

- From ITU-R 709 RGB to ITU-R 2020 RGB:

$$
\left[\begin{array}{l}
R \\
G \\
B
\end{array}\right]_{R 2020}=M_{X Y Z t o R 2020} \cdot M_{R 709 t o X Y Z} \cdot\left[\begin{array}{l}
R \\
G \\
B
\end{array}\right]_{R 709}
$$

- From ITU-R 2020 RGB to ITU-R 709 RGB:

$$
\left[\begin{array}{l}
R \\
G \\
B
\end{array}\right]_{R 709}=M_{X Y Z t o R 709} \cdot M_{R 2020 t o X Y Z} \cdot\left[\begin{array}{c}
R \\
G \\
B
\end{array}\right]_{R 2020}
$$

, Where:
$M_{R 709 t o X Y Z}=\left[\begin{array}{lll}0.4124 & 0.3576 & 0.1805 \\ 0.2126 & 0.7152 & 0.0722 \\ 0.019 & 0.1192 & 0.9505\end{array}\right]$ and $M_{X Y Z t o R 709}=M_{R 709 t o X Y Z}^{-1}$
$M_{\text {R2020toXYZ }}=\left[\begin{array}{lll}0.6370 & 0.1446 & 0.1689 \\ 0.2627 & 0.6780 & 0.0593 \\ 0.0000 & 0.0281 & 1.0610\end{array}\right]$ and $M_{X Y Z t o R 2020}=M_{R 2020 t o X Y}^{-1}$

- 39

39

## Representing colour

- We need a way to represent colour in the computer by some set of numbers
- A) preferably a small set of numbers which can be quantised to a fairly small number of bits each
- Gamma corrected RGB, sRGB and CMYK for printers
- B) a set of numbers that are easy to interpret
- Munsell's artists's scheme

HSV, HLS
C) a set of numbers in a 3D space so that the (Euclidean) distance in that space corresponds to approximately perceptually uniform colour differences
, CIE Lab, CIE Luv

41

How to transform between linear RGB colour spaces?


- From ITU-R 709 RGB to XYZ:

> 38
38


40

## $R G B$ spaces

- Most display devices that output light mix red, green and blue lights to make colour
, televisions, CRT monitors, LCD screens
RGB colour space
- Can be linear (RGB) or display-encoded (R'G'B')
, Can be scene-referred (HDR) or display-referred (SDR)
- There are multiple RGB colour spaces
- ITU-R 709 (sRGB), ITU-R 2020, Adobe RGB, DCI-P3
, Each using different primary colours
- And different OETFs (gamma, PQ, etc.)
- Nominally, $R G B$ space is a cube

42

42

## $R G B$ in CIE XYZ space

- Linear RGB colour values can be transformed into CIE XYZ
, by matrix multiplication
- because it is a rigid transformation the colour gamut in CIE XYZ is a rotate and skewed cube
- Transformation into Yxy

- is non-linear (non-rigid)
- colour gamut is more complicated

b 43
43


45

## Munsell's colour classification system

, any two adjacent colours are a standard "perceptual" distance apart
, worked out by testing it on people
p a highly irregular space
e.g. vivid yellow is much brighter than vivid blue


47

## CMY space

- printers make colour by mixing coloured inks
- the important difference between inks (CMY) and lights ( $R G B$ ) is that, while lights emit light, inks absorb light
, cyan absorbs red, reflects blue and green
- magenta absorbs green, reflects red and blue
, yellow absorbs blue, reflects green and red
- $C M Y$ is, at its simplest, the inverse of $R G B$
- CMY space is nominally a cube


44


46

## Colour spaces for user-interfaces

- $R G B$ and $C M Y$ are based on the physical devices which produce the coloured output
- RGB and CMY are difficult for humans to use for selecting colours
- Munsell's colour system is much more intuitive:
- hue - what is the principal colour?
, value - how light or dark is it?
- chroma - how vivid or dull is it?
- computer interface designers have developed basic transformations of $R G B$ which resemble Munsell's humanfriendly system


48

## HSV: hue saturation value

- three axes, as with Munsell
, hue and value have same meaning
p the term "saturation" replaces the term "chroma"
simple conversion from gammacorrected RGB to HSV

- designed by Alvy Ray Smith in 1978
- algorithm to convert $H S V$ to $R G B$ and back can be found in Foley et al., Figs 13.33 and 13.34
- 

49

## $H L S$ : hue lightness saturation

+ a simple variation of $H S V$
- hue and saturation have same meaning
- the term "lightness" replaces the term "value"
+ designed to address the complaint that $H S V$ has all pure colours having the same lightness/value as white
- designed by Metrick in 1979
- algorithm to convert $H L S$ to $R G B$ and back can be found in Foley et al., Figs 13.36 and 13.37


50

## CIE L* $u^{*} v^{*}$ and $u$ 'v'

- Approximately perceptually uniform
- u'v' chromacity

$$
\begin{array}{ll}
u^{\prime}=\frac{4 X}{X+15 Y+3 Z} & =\frac{4 x}{-2 x+12 y+3} \\
v^{\prime}=\frac{9 Y}{X+15 Y+3 Z} & =\frac{9 y}{-2 x+12 y+3}
\end{array}
$$

52


54

Colour - references

- Chapters ,,Light" and „Colour" in

Shirley, P. \& Marschner, S., Fundamentals of Computer Graphics

- Textbook on colour appearance
, Fairchild, M. D. (2005). Color Appearance Models (second.). John Wiley \& Sons
- Comprehensive review of colour research
, Wyszecki, G., \& Stiles,W. S. (2000). Color science: concepts and methods, quantitative data, and formulae (Second ed.). John Wiley \& Sons.

55


1


3


5

Many graphics / display solutions are motivated by visual perception


2

## Steven's power law for brightness

- Stevens (1906-1973) measured the perceived magnitude of physical stimuli
- Loudness of sound, tastes, smell, warmth, electric shock and brightness
p Using the magnitude estimation methods
- Ask to rate loudness on a scale with a known reference
- All measured stimuli followed the power law:

- For brightness ( 5 deg target in dark), a $=0.3$

4


6


7

Threshold versus intensity (t.v.i.) function

- The smallest detectable difference in luminance for a given background luminance

${ }^{L+\Delta L^{\prime}}$

9

9

Psychophysics
Threshold experiments

$>11$

Detection thresholds


- The smallest detectable difference between
b the luminance of the object and
v the luminance of the background

8


10


12


13

How to make luminance (more) perceptually uniform?

- Using "Fechnerian" integration


15


17

## Consequence of the Weber-law

Smallest detectable difference in luminance

$$
\begin{array}{l|l|c|c}
\frac{\Delta L}{L}=k & \text { For } \mathrm{k}=1 \% & \mathbf{L} & \Delta \mathrm{~L} \\
\hline & & 100 \mathrm{~cd} / \mathrm{m}^{2} & 1 \mathrm{~cd} / \mathrm{m}^{2} \\
\hline & \mathrm{Idd} / \mathrm{m}^{2} & 0.0 \mathrm{l} \mathrm{~cd} / \mathrm{m}^{2}
\end{array}
$$

- Adding or subtracting luminance will have different visual impact depending on the background luminance
- Unlike LDR luma values, luminance values are not perceptually uniform!
b 14
14

Assuming the Weber law

$$
\frac{\Delta L}{L}=k
$$

- and given the luminance transducer

$$
R(L)=\int \frac{1}{\Delta L(l)} d l
$$

- the response of the visual system to light is:

$$
R(L)=\int \frac{1}{k L} d L=\frac{1}{k} \ln (L)+k_{1}
$$

- 16

16

## But...the Fechner law does not hold for the full luminance range

- Because the Weber law does not hold either
- Threshold vs. intensity function:

- 18


## Weber-law revisited

- If we allow detection threshold to vary with luminance according to the t.v.i. function:

- we can get a more accurate estimate of the "response":

$$
R(L)=\int_{0}^{L} \frac{1}{t v i(l)} d l
$$

19


21


Resolution and sampling rate

- Pixels per inch [ppi]
- Does not account for vision
- The visual resolution depends on
b screen size
- screen resolution

D viewing distance

- The right measure
- Pixels per visual degree [ppd]
- In frequency space
- Cycles per visual degree [cpd]


24

## Fourier analysis

- Every N-dimensional function (including images) can be represented as a sum of sinusoidal waves of different frequency and phase

- Think of "equalizer" in audio software, which manipulates each frequency
- 25

25

## Nyquist frequency

- Sampling density restricts the highest spatial frequency signal that can be (uniquely) reconstructed
- Sampling density - how many pixels per image/visual angle/...

- Any number of sinusoids can be fitted to this set of samples
- It is possible to fit an infinite number of sinusoids if we allow infinitely high frequency
- 27

27


29

## Spatial frequency in images

- Image space units: cycles per sample (or cycles per pixel)

* What are the screen-space frequencies of the red and green sinusoid?
- The visual system units: cycles per degree
- If the angular resolution of the viewed image is 55 pixels per degree, what is the frequency of the sinusoids in cycles per degree?
- 26

26

## Nyquist frequency

- Sampling density restricts the highest spatial frequency signal that can be (uniquely) reconstructed
- Sampling density - how many pixels per image/visual angle/...

- Any number of sinusoids can be fitted to this set of samples
- It is possible to fit an infinite number of sinusoids if we allow infinitely high frequency
- 28

28

## Nyquist frequency

- Sampling density restricts the highest spatial frequency signal that can be (uniquely) reconstructed
- Sampling density - how many pixels per image/visual angle/...

- Any number of sinusoids can be fitted to this set of samples
- It is possible to fit an infinite number of sinusoids if we allow infinitely high frequency
- 30

30

## Nyquist frequency / aliasing

- Nuquist frequency is the highest frequency that can be represented by a discrete set of uniform samples (pixels)
- Nuquist frequency $=0.5$ sampling rate
- For audio
- If the sampling rate is 44100 samples per second (audio CD), then the Nyquist frequency is 22050 Hz
- For images (visual degrees)

If the sampling rate is 60 pixels per degree, then the Nyquist frequency is 30 cycles per degree

- When resampling an image to lower resolution, the frequency content above the Nyquist frequency needs to be removed (reduced in practice)
- Otherwise aliasing is visible
- 31


33


## Modeling contrast detection



- 32

32


34

CSF as a function of spatial frequency


[^7]36

CSF as a function of background luminance


- 37


39

CSF as a function of spatial frequency and background luminance



- 38

38


40

Spatio-chromatic CSF


42


43


- The same amount of blur was introduced into light-dark, red-green and blue-yellow colour opponent channels
- The blur is only visible in light-dark channel
- This property is used in image and video compression
, Sub-sampling of colour channels (4:2:I)
- 45

45

Mach Bands - evidence for band-pass visual processing

- "Overshooting" along edges
- Extra-bright rims on bright sides
- Extra-dark rims on dark sides
- Due to "Lateral Inhibition"




44


46


48


49

Spatial-frequency selective channels

- The visual information is decomposed in the visual cortex into multiple channels
- The channels are selective to spatial frequency, temporal frequency and orientation
- Each channel is affected by different ,noise" level
- The CSF is the net result of
 information being passed in noiseaffected visual channels
> 51
51



50


52

Predicting visible differences with CSF

- We can use CSF to find the probability of spotting a difference beween a pair of images $X_{1}$ and $X_{2}$ :


54


55


57

Temporal adaptation mechanisms

- Bleaching \& recovery of photopigment
, Slow assymetric (light -> dark, dark -> light)
- Reaction times ( $1-1000 \mathrm{sec}$ )
- Separate time-course for rods and cones
- Neural adaptation
- Fast
- Approx. symmetric reaction times (10-3000 ms)
- Pupil
- Diameter varies between 3 and 8 mm
- About I:7 variation in retinal illumunation
+ 59


56


58


60


61


63


65


62

High-Level Contrast Processing


- 64

64

Shape Processing: Geometrical Clues


66


67


69


68

## References

- Wandell, B.A. (I995). Foundations of vision. Sinauer Associates.
, Available online: https://foundationsofvision.stanford.edu/
- Mantiuk, R. K., Myszkowski, K., \& Seidel, H. (2015). High Dynamic Range Imaging. In Wiley Encyclopedia of Electrical and Electronics Engineering. Wiley.
- Section 2.4
- Available online:
http://www.cl.cam.ac.uk/~rkm38/hdri book.html

70


1

Real-world scenes are more challenging


- The match could not be achieved if the light source in the top of the box was visible
- The display could not reproduce the right level of brightness

3

Dynamic range (contrast)

- As ratio:

$$
C=\frac{L_{\max }}{L_{\min }}
$$

- Usually written as C:I, for example 1000:I.
- As "orders of magnitude"
or $\log 10$ units:

$$
C_{10}=\log _{10} \frac{L_{\max }}{L_{\min }}
$$

- As stops:

$$
C_{2}=\log _{2} \frac{L_{\max }}{L_{\min }} \quad \begin{aligned}
& \text { One stop is doubling } \\
& \text { of halving the amount of light }
\end{aligned}
$$

5

Cornell Box: need for tone-mapping in graphics


Rendering


Photograph
$>2$

2


4


6


Why do we need tone mapping?

- To reduce dynamic range
- To customize the look
- colour grading
- To simulate human vision
, for example night vision

- To adapt displayed images to a display and viewing conditions
- To make rendered images look more realistic
- To map from scene- to display-referred colours
- Different tone mapping operators achieve different goals ${ }^{8} 8$

7

From scene- to display-referred colours

- The primary purpose of tone mapping is to transform an image from scene-referred to display-referred colours

${ }^{\circ} 9$

9
10

## Basic tone-mapping and display coding

- The simplest form of tone-mapping is the exposure/brightness adjustment:

$$
\begin{array}{lll}
\text { Display-referred red value } & R_{d}=\frac{R_{s}}{L_{w h i t e}} & \text { Scene-referred } \\
\text { / R for red, the same for green and blue }
\end{array} \begin{gathered}
\text { Scene-referred } \\
\text { luminance of white }
\end{gathered}
$$

- No contrast compression, only for a moderate dynamic range
- The simplest form of display coding is the "gamma"

| Prime (') denotes a <br> gamma-corrected value |
| :--- |
| ( |$R^{\prime}=\left(R_{d}\right)^{\frac{1}{\gamma}} \quad$ Typically $\gamma=2.2$

For SDR displays only

[^8]
## Techniques

- Arithmetic of HDR images
- Display model
- Tone-curve
- Colour transfer
- Base-detail separation
- Glare
- 13


15


## Arithmetic of HDR images

- How do the basic arithmetic operations
- Addition
- Multiplication
- Power function
affect the appearance of an HDR image?
- We work in the luminance space (NOT luma)
- The same operations can be applied to linear RGB
, Or only to luminance and the colour can be transferred
- 14

14


16

## Techniques

- Arithmetic of HDR images
- Display model
- Tone-curve
- Colour transfer
- Base-detail separation
- Glare


## Display-adaptive tone mapping

- Tone-mapping can account for the physical model of a display
b How a display transforms pixel values into emitted light
- Useful for ambient light compensation


19

## Inverse display model

Symbols are the same as for the forward display model

$$
V=\left(\frac{L-L_{\text {black }}-L_{\text {refl }}}{L_{\text {peak }}-L_{\text {black }}}\right)^{(1 / \gamma)}
$$

Note: This display model does not address any colour issues. The same equation is applied to red, green and blue color channels. The assumption is that the display primaries are the same as for the sRGB color space.
> 21
21


23


20


22

## Example: Ambient light compensation

- We are looking at the screen in bright light

- We assume that the dynamic of the input is 2.6 ( $\approx 400: 1$ )

$$
r_{\text {in }}=2.6 \quad r_{\text {out }}=\log _{10} \frac{L_{\text {peak }}}{L_{\text {black }}+L_{\text {refl }}}=1.77
$$

- First, we need to compress contrast to fit the available dynamic range, then compensate for ambient light


24

## Techniques

- Arithmetic of HDR images
- Display model
- Tone-curve
- Colour transfer
- Base-detail separation
- Glare
- 25


## Tone-curve


log input luminance factor (HDR image)
> 27

## Sigmoidal tone-curves

- Very common in digital cameras
- Mimic the response of analog film
- Analog film has been engineered over many years to produce good tone-reproduction

- Fast to compute
- 29


26

## Tone-curve



[^9]28

## Sigmoidal tone mapping

- Simple formula for a sigmoidal tone-curve:

$$
R^{\prime}(x, y)=\frac{R(x, y)^{b}}{\left(\frac{L_{m}}{a}\right)^{b}+R(x, y)^{b}}
$$

where $L_{m}$ is the geometric mean (or mean of logarithms):

$$
L_{m}=\exp \left(\frac{1}{N} \sum_{(x, y)} \ln (L(x, y))\right)
$$

and $L(x, y)$ is the luminance of the pixel $(x, y)$.



31


33


35

## Histogram equalization

- I. Compute normalized cummulative image histogram

$$
c(I)=\frac{1}{N} \sum_{i=0}^{I} h(i)=c(I-1)+\frac{1}{N} h(I)
$$

- For HDR, operate in the log domain
- 2. Use the cummulative histogram as a tone-mapping function

$$
Y_{\text {out }}=c\left(Y_{\text {in }}\right)
$$

- For HDR, map the $\log -10$ values to the $\left[-d r_{\text {out }} ; 0\right]$ range
- where drout is the target dynamic range (of a display)

> 32
32


## CLAHE: Contrast-Limited Adaptive

 Histogram Equalization

34

CLAHE: Contrast-Limited Adaptive Histogram Equalization

- Truncate the bins that exceed the ceiling;
- Distribute the removed counts to all bins;
- Repeat until converges


36


37

## Colour transfer in tone-mapping

- Many tone-mapping operators work on luminance, mean or maximum colour channel value
- For speed
, To avoid colour artefacts
- Colours must be transferred later form the original image
- Colour transfer in the linear RGB colour space:

- The same formula applies to green (G) and blue (B) linear colour values
> 39
39


## Colour transfer: alternative method

- Colour transfer in linear RGB will alter resulting luminance
- Colours can be also transferred, and saturation adjusted using CIE u'v' chromatic coordinates


41

## Techniques

- Arithmetic of HDR images
- Display model
- Tone-curve
- Colour transfer
- Base-detail separation
- Glare
- 38


40

## Techniques

- Arithmetic of HDR images
- Display model
- Tone-curve
- Colour transfer
- Base-detail separation
- Glare

|  |
| :--- | :--- |
| $\square 42$ |

42


43

## Reflectance \& Illumination TMO

- Hypothesis: Distortions in reflectance are more apparent than the distortions in illumination
- Tone mapping could preserve reflectance but compress illumination


## Illumination

Tone-mapped image $\quad L_{d}=R \cdot T(I)$


- for example:

$$
L_{d}=R \cdot\left(I / L_{w h i t e}\right)^{c} \cdot L_{w h i t e}
$$

- 45

45


47
$L_{d}=R \cdot\left(I / L_{\text {white }}\right)^{c} \cdot L_{\text {white }}$

Illumination and reflectance

Reflectance

- White $\approx 90 \%$
- Black $\approx 3 \%$
- Dynamic range < 100: 1
- Reflectance critical for object \& shape detection

Illumination

- Sun $\approx 10^{9} \mathrm{~cd} / \mathrm{m}^{2}$
- Lowest perceivable luminance $\approx 10^{-6} \mathrm{~cd} / \mathrm{m}^{2}$
- Dynamic range $10,000: 1$ or more
- Visual system partially discounts illumination

How to separate the two?

- (Incoming) illumination - slowly changing
- except very abrupt transitions on shadow boundaries
- Reflectance - low contrast and high frequency variations


46

Bilateral filter $\quad I_{p} \approx \frac{1}{k_{s}} \sum_{t \in \Omega} f(p-t) g\left(L_{p}-L_{t}\right) L_{p}$

- Better preserves sharp edges

- Still some blurring on the edges
- Reflectance is not perfectly separated from illumination near edges
48 [Durand \& Dorsey, SIGGRAPH 2002]

Weighted-least-squares (WLS) filter

- Stronger smoothing and still distinct edges

- Can produce stronger effects with fewer artifacts
- See „Advanced image processing" lecture
[Farbman et al., SIGGRAPH 2008] - 49

49


51

| Techniques |
| :--- |
| - Arithmetic of HDR images |
| - Display model |
| • Tone-curve |
| - Colour transfer |
| - Base-detail separation |
| - Glare |
|  |
|  |
| 53 |

## Retinex

- Retinex algorithm was initially intended to separate reflectance from illumination [Land I964]
- There are many variations of Retinex, but the general principle is to eliminate from an image small gradients, which are attributed to the illumination

| 2 step: compute |
| :--- |
| gradients in log domain |


| $2^{\text {nd }}$ step: set to 0 |
| :--- |
| gradients less than the |
| threshold |


| 3rd step: reconstruct an |
| :--- |
| image from the vector |
| field |

Poisson equation

50


52


54


Scattering of the light in the eye

55


57


59


From: Sekuler, R., and Blake, R. Perception, second ed. McGraw- Hill, New York, 1990 - 56

56


58

Point Spread Function of the eye


60

Selective application of glare


- A) Glare applied to the entire image Glare kernel $I_{g}=I * G \quad$ (PSF)
- Reduces image contrast and sharpness

B) Glare applied only to the clipped pixels
$I_{g}=I+I_{\text {cliped }} * G-I_{\text {cliped }}$ where $I_{\text {cliped }}= \begin{cases}I & \text { for } I>1 \\ 0 & \text { otherwise }\end{cases}$ Better image quality - 61

61

## Glare (or bloom) in games

- Convolution with large, non-separable filters is too slow
- The effect is approximated by a combination of Gaussian filters
, Each filter with different "sigma"
- The effect is meant to look good, not be be accurate model of light scattering
- Some games simulate camera rather than the eye

- 63

63


65

Selective application of glare


62

Does the exact shape of the PSF matter?

- The illusion of increased brightness works even if the PSF is very different from the PSF of the eye


- 64

64

## References

- Comprehensive book on HDR Imaging
, E. Reinhard,W. Heidrich, P. Debevec, S. Pattanaik, G.Ward, and K. Myszkowski, High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting, 2nd editio. Morgan Kaufmann, 2010.
- Overview of HDR imaging \& tone-mapping
, http://www.cl.cam.ac.uk/~rkm38/hdri book.html
- Review of recent video tone-mapping

A comparative review of tone-mapping algorithms for high dynamic range video Gabriel Eilertsen, Rafal K. Mantiuk, Jonas Unger, Eurographics State-of-The-Art Report 2017

- Selected papers on tone-mapping:
G.W. Larson, H. Rushmeier, and C. Piatko, "A visibility matching tone reproduction operator for high dynamic range
 R. Wanat and R. K. Mantiuk, "Simulating and compensating changes in appearance between day and night vision," ACM
Trons. Groph. (Proc. SIGGRAPH), vol. 33, no. 4, p. 147, 2014. r. Clo SPS or Digital Images. Proceedings of SIGGRAPH. (1995), 325-334 Ritschel,T. et al. 2009. Temporal Glare: Real-Time Dynamic Simulation of the Scattering in the Human Eye. Computer Graphics Forum. 28,2 (Apr. 2009), 183-192


66


1


3


5


2


4


6


7

Ivan Sutherland's HMD

- optical see-through AR, including:
- displays ( $2 \times 1$ " CRTs)
- rendering
- head tracking
- interaction
- model generation
- computer graphics
- human-computer interaction


-9
9


11


8


10


12


13


15
(Some) challenges of optical see-through AR

- Transparency, lack of opacity

Display light is mixed with environment light

- Resolution and field-of-view
- Eye-box
, The volume in which the pupil needs to see the image
- Brightness and contrast
b Blocked vision - forward and periphery (safety)
- Power efficiency
- Size, weight and weight distribution
, 50 grams are comfortable for long periods
- Social issues, price, vision correction, individual variability..

More resources: https://kguttag.com/

17


14


16


18


19


## VR/AR challenges

- Latency (next lecture)
- Tracking
- 3D Image quality and resolution
- Reproduction of depth cues (last lecture)
- Rendering \& bandwidth
- Simulation/cyber sickness
- Content creation
- Game engines
- Image-Based-Rendering
$>20$
20


## References

LaValle "Virtual Reality", Cambridge University Press, 2016

- http://vr.cs.uiuc.edu/
- Virtual Reality course from the Stanford Computational Imaging group
, http://stanford.edu/class/ee267/
- KGOnTech blog
- https://kguttag.com/

22

22


1

| Latency in VR |  |
| :---: | :---: |
| - Sources of latency in VR <br> - $1 \mathrm{MU} \sim 1 \mathrm{~ms}$ <br> - Inertial Measurement Unit <br> - sensor fusion, data transfer <br> - rendering: depends on complexity of scene \& GPU - a few ms <br> - data transfer again <br> - Display <br> , $60 \mathrm{~Hz}=16.6 \mathrm{~ms} ;$ <br> , $70 \mathrm{~Hz}=11.1 \mathrm{~ms} ;$ <br> - $120 \mathrm{~Hz}=8.3 \mathrm{~ms}$. | - Target latency <br> - Maximum acceptable: 20 ms <br> - Much smaller (5ms) desired for interactive applications <br> - Example <br> - 16 ms (display) +16 ms (rendering) +4 ms (orientation tracking) $=36$ ms latency total <br> - At $60 \mathrm{deg} / \mathrm{s}$ head motion, IKxIK, I00deg for display: <br> , 19 pixels error <br> - Too much |
| > 3 |  |

3


5

Post-rendering image warp (time warp)

- To minimize end-to-end latency
- The method:
- get current camera pose
( render into a larger raster than the screen buffer

- Original paper from Mark et al. 1997, also Darsa et al. 1997
I997, also Darsa et al. I 997
Meta:Asynchronous Time Warp
- 4

4
, get new camera pose
> warp rendered image using the latest pose, send to the display
, 2D image translation
, 2D image warp
, 3D image warp

Eye movement - basics

Saccade

$160-300 \mathrm{deg} / \mathrm{s}$

Eye movement - basics

Smooth Pursuit Eye Motion (SPEM)

Up to $80 \mathrm{deg} / \mathrm{s}$
The tracking is imperfect

- especially at higher velocities
- and for unpredictable motion

7

## Retinal velocity

- The eye tracks moving objects
- Smooth Pursuit Eye Motion (SPEM) stabilizes images on the retina
- But SPEM is imperfect
- Loss of sensitivity mostly caused by imperfect SPEM
- SPEM worse at high velocities

7

## Motion sharpening

- The visual system "sharpens" objects moving at speeds of 6 deg/s or more
, Potentially a reason why VR appears sharper than it actually is

9


11

## Hold-type blur

- The eye smoothly follows a moving object
* But the image on the display is "frozen" for $1 / 60^{\text {th }}$ of a second


10


12


13


- A similar idea to low-persistence displays in VR
- Reduces hold-type blur

Low persistence displays

- MostVR displays flash an image for a fraction of frame duration
- This reduces hold-type blur
- And also reduces the perceived lag of the rendering



Cathode Ray Tube




21


Liquid Chrystal Displays (LCD)


20
From: htipollcomputer:howstuftworks.com/monitors .htm
20


22

## LCD temporal response

- Experiment on an IPS LCD screen
- We rapidly switched between two intensity levels at 120 Hz
- Measured luminance integrated over Is
- The top plot shows the difference between expected $\left(\frac{I_{t-1}+I_{t}}{2}\right)$ and measured luminance
- The bottom plot: intensity
measurement for the full brightness and half-brightness display settings


24


25


27


29

Liquid Crystal on Silicon (LCoS)

- basically a reflective LCD

- standard component in projectors and head mounted displays
- used e.g.in Google Glass

26


28

## Active matrix OLED

- Commonly used in mobile phones (AMOLED)
- Very good contrast
- But the screen more affected by glare than LCD
- But limited brightness
- The brighter is OLED, the shorter is its live-span

Temporal characteristic


31


33


35


32


34


36


37


39


41

## HDR Display

- Two spatial modulators
, Ist modulator contrast 1000:I
- 2nd modulator contrast 1000:1

Combined contrast 1000,000:I

- Idea: Replace constant backlight of LCD panels with an array of LEDs
, Very few (about IO00) LEDs sufficient
- Every LED intensity can be set individually
- Very flat form factor (fits in standard LCD housing)
- Issue:
- LEDs larger than LCD pixels
, This limits maximum local contrast
- 38


40


42


43


45

HDR rendering algorithm - high level

$\underset{\sim}{\operatorname{argmin}}||I(x, y)-g * D(x, y) L(x, y)||_{2}$
L,D


44


46

## References

- Hainich, R.R.AND BImber, O. 201 I. Displays:Fundamentals and Applications. CRC Press.
- Seetzen, H., Heidrich,W., Stuerzlinger,W., et al. 2004. High dynamic range display systems. ACM Transactions on Graphics 23, 3,760.
- Visual motion test for high-frame-rate monitors:
, https://www.testufo.com/

17


1


3


5

Depth perception

We see depth due to depth cues.
Stereoscopic depth cues: binocular disparity


2

## Depth perception

We see depth due to depth cues.
Stereoscopic depth cues: binocular disparity

Ocular depth cues: accommodation, vergence

Pictorial depth cues:
occlusion, size, shadows...

-
4

Depth perception

We see depth due to depth cues.
Stereoscopic depth cues: binocular disparity

Ocular depth cues: accommodation, vergence

Challenge:
Consistency is required!
Pictorial depth cues:
 occlusion, size, shadows...


7


9


11

Disparity \& occlusion conflict


8

Depth perception

We see depth due to depth cues.


10

Single Image Random Dot Stereograms


- Fight the vergence vs. accommodation conflict to see the hidden image


13


15


17

## Comfort zones

## Comfort zone size depends on:

- Presented content
- Viewing condition


14


16

## Comfort zones

## Comfort zone size

depends on:

- Presented content
- Viewing condition
- Screen distance


## Other factors:

- Distance between eyes
- Depth of field
- Temporal coherence



19


21


23

Stereoscopic displays

- Auto-stereoscopic (without glasses)
- Parallax barrier
- Example: Nintendo 3DS, some laptops and mobile phones
- Switchable 2D/3D
- Lenticular lens
- Better efficiency
, Non-switchable


22


24


25


27


29


26


28

## Anaglyph Stereo - Full Color

- render L \& R images, do not convert to grayscale
- merge into red-cyan anaglyph by assigning $I(r)=L(r), I(g, b)=R(g, b)$ (I is anaglyph)


30


31


33


35





37

Parallax - not well done (vertical parallax = unnatural)


38

## References

- LaValle "Virtual Reality", Cambridge University Press, 2016
- Chapter 6
- http://vr.cs.uiuc.edu/
- Stereoscopic displays:
- Hainich, Rolf R., and Oliver Bimber. Displays: Fundamentals and Applications. 2nd ed. CRC Press, 2016.
- 39


# Advanced Graphics and Image Processing Lecture notes 

Rafał Mantiuk

Lent term 2018/19

## 1 Contrast- and gradient-based methods

Many problems in image processing are easier to solve or produce better results if operations are not peformed directly on image pixel values but on differences between pixels. Instead of altering pixels, we can transform an image into gradient field and then edit the values in the gradient field. Once we are done with editing, we need to reconstruct an image from the modified gradient field.

A few examples of gradient-based methods are shown in Figures 1 and 2.
In one common case such differences between pixels represent gradients: for image $I$, a gradient at a pixel location $(x, y)$ is computed as:

$$
\nabla I_{x, y}=\left[\begin{array}{c}
I_{x+1, y}-I_{x, y}  \tag{1}\\
I_{x, y+1}-I_{x, y}
\end{array}\right]
$$

The equation above is obviously a discrete approximation of a gradient, as we are dealing with discrete pixel values rather than a continous function. This particular approximation is called forward difference, as we take the difference between the next and current pixel. Other choices include backward differences (current minus previous pixel) or central differences (next minus previous pixel).

Once a gradient field is computed, we can start modifying it. Usually better effects are achieved if the magnitude of gradients is modified and the orientation of each gradient remains unchanged. This can be achieved by


Figure 1: Two examples of gradient-based processing. Texture details in the original image were enhanced to produce the result shown in (b). Contrast was removed everywhere except at the edges to produced a cartoonized image in (c).
multiplying gradients by the gradient editing function $f()$ :

$$
\begin{equation*}
G_{x, y}=\nabla I_{x, y} \cdot \frac{f\left(\left\|\nabla I_{x, y}\right\|\right)}{\left\|\nabla I_{x, y}\right\|} \tag{2}
\end{equation*}
$$

where $\|\cdot\|$ operator computes the magnitude (norm) of the gradient.
We try to reconstruct pixel values, which would result in a gradient field that is the closest to our modifed gradient field $G=\left[\begin{array}{ll}G^{(x)} & G^{(y)}\end{array}\right]^{\prime}$. In particular, we can try to minimize the squared differences between gradients in actual image and modified gradients:

$$
\begin{equation*}
\underset{I}{\arg \min } \sum_{x, y}\left[\left(I_{x+1, y}-I_{x, y}-G_{x, y}^{(x)}\right)^{2}+\left(I_{x, y+1}-I_{x, y}-G_{x, y}^{(y)}\right)^{2}\right], \tag{3}
\end{equation*}
$$



Figure 2: Comparison of naive and gradient domain image copy \& paste.


Figure 3: When using forward-differences, a pixel with the coordinates $(x, y)$ is referred to in at moost four partial derivates, two along $x$-axis and two along $y$-axis.
where the summation is over the entire image. To minimize the function above, we need to equate its partial derivatives to 0 . As we optimze for pixel values, we need to compute partial derivates with respect to $I_{x, y}$. Fortunately, most terms in the sum will become 0 after differentiation, as they do not contain the differentiated variable $I_{x, y}$. For a given pixel $(x, y)$, we need to consider only 4 partial derivates: two belonging to the pixel $(x, y)$, x derivative for the pixel on the left $(x-1, y)$ and $y$-derivative for the pixel in the top $(x, y-1)$, as shown in Figure 3. This gives us:

$$
\begin{align*}
\frac{\delta F}{\delta I_{x, y}}= & -2\left(I_{x+1, y}-I_{x, y}-G_{x, y}^{(x)}\right)-2\left(I_{x, y+1}-I_{x, y}-G_{x, y}^{(y)}\right)+  \tag{4}\\
& 2\left(I_{x, y}-I_{x-1, y}-G_{x-1, y}^{(x)}\right)+2\left(I_{x, y}-I_{x, y-1}-G_{x, y-1}^{(y)}\right) . \tag{5}
\end{align*}
$$

After rearanging the terms and equating $\frac{\delta F}{\delta I_{x, y}}$ to 0 , we get:

$$
\begin{equation*}
I_{x-1, y}+I_{x+1, y}+I_{x, y-1}+I_{x, y+1}-4 I_{x, y}=G_{x, y}^{(x)}-G_{x-1, y}^{(x)}+G_{x, y}^{(y)}-G_{x, y-1}^{(y)} \tag{6}
\end{equation*}
$$

In these few steps we derived a discrete Poisson equation, which can be found in many engineering problems. The Poisson equation is often written as:

$$
\begin{equation*}
\nabla^{2} I=\operatorname{div} G \tag{7}
\end{equation*}
$$

where $\nabla^{2} I$ is the discrete Laplace operator:

$$
\begin{equation*}
\nabla^{2} I_{x, y}=I_{x-1, y}+I_{x+1, y}+I_{x, y-1}+I_{x, y+1}-4 I_{x, y} \tag{8}
\end{equation*}
$$

and $\operatorname{div} G$ is the divergence of the vector field:

$$
\begin{equation*}
\operatorname{div} G_{x, y}=G_{x, y}^{(x)}-G_{x-1, y}^{(x)}+G_{x, y}^{(y)}-G_{x, y-1}^{(y)} \tag{9}
\end{equation*}
$$

We can also write the equation using discrete convolution operators:

$$
I *\left[\begin{array}{ccc}
0 & 1 & 0  \tag{10}\\
1 & -4 & 1 \\
0 & 1 & 0
\end{array}\right]=G^{(x)} *\left[\begin{array}{ccc}
-1 & 1 & 0
\end{array}\right]+G^{(y)} *\left[\begin{array}{c}
-1 \\
1 \\
0
\end{array}\right]
$$

Note that the covolution flips the order of elements in the kernel, thus the row and column vectors on the right hand side are also flipped.

When equation 6 is satisfied for every pixel, it forms a system of linear equations:

$$
A \cdot\left[\begin{array}{c}
I_{1,1}  \tag{11}\\
I_{2,1} \\
\ldots \\
I_{N, M}
\end{array}\right]=b
$$

Here we represent an image as a very large column vector, in which image pixels are stacked column-after-column (in an equivalent manner they can be stacked row-after-row). Every row of matrix A contains the Laplace operator for a corresponding pixel. But the matrix also needs to account for the boundary conditions, that is handle pixels that are at the image edge and therefore do not contain neighbour on one of the sides. Matrix A for a tiny

3x3 image looks like this:

$$
A=\left[\begin{array}{ccccccccc}
-2 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0  \tag{12}\\
1 & -3 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 1 & -2 & 0 & 0 & 1 & 0 & 0 & 0 \\
1 & 0 & 0 & -3 & 1 & 0 & 1 & 0 & 0 \\
0 & 1 & 0 & 1 & -4 & 1 & 0 & 1 & 0 \\
0 & 0 & 1 & 0 & 1 & -3 & 0 & 0 & 1 \\
0 & 0 & 0 & 1 & 0 & 0 & -2 & 1 & 0 \\
0 & 0 & 0 & 0 & 1 & 0 & 1 & -3 & 1 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & -2
\end{array}\right]
$$

Obviously, the matrix is enormous for normal size images. However, most matrix elements are 0 , so it can be easily stored using a sparse matrix representation. Note that only the pixel in the center of the image (5th row) contains the full Laplace operator; all other pixels are missing neighbours so the operator is adjusted accordingly. Accounting for all boundary cases is probably the most difficult and error-prone part in formulating gradient-field reconstruction problem. The column vector $b$ corresponds to the right hand side of equation 6 .

## 2 Solving linear system

There is a large number of methods and software libraries, which can solve a sparse linear problem given in Equation 11. The Poisson equation is typically solved using multi-grid methods, which iteratively update the solution at different scales. Those, however, are rarther difficult to implement and tailored to one particular shape of a matrix. Alternatively, the solution can be readily found after transformation to the frequency domain (discrete cosine transform). However, such a method does not allow introducing weights, importance of which will be discussed in the next section. Finally, conjugate gradient and biconjugate gradient [1, sec. 2.7] methods provide a fastconverging iterative method for solving sparse systems, which can be very memory efficient. Those methods require providing only a way to compute multiplication of the matrix $A$ and its transpose with an arbitrary vector. Such operation can be realized in an arbitrary way without the need to store the sparse matrix (which can be very large even if it is sparse). The conjugate gradient requires fewer operations than the biconjugate gradient method, but


Figure 4: The solution of gradient field reconstruction often contain "pinching" artefacts, such as shown in figure (a). The artefacts can be avoided if small gradient magnitudes are weighted more than large magnitudes.
it should be used only with positive definite matrices. Matrix $A$ is not positive definite so in principle the biconjugate gradient method should be used. However, in practice, conjugate gradient method converges equally well.

## 3 Weighted reconstruction

An image resulting from solving Equation 11 often contains undesirable "pinching" artefacts, such as those shown in Figure 4a. Those artefacts are inherent to the nature of gradient field reconstruction - the solution is just the best approximation of the desired gradient field but it hardly ever exactly matches the desired gradient field. As we minimize squared differences, tiny inaccuracies for many pixels introduce less error than large inaccuracies for few pixels. This in turn introduces smooth gradients in the areas, where the desired gradient field is inconsistent (cannot form an image). Such gradients produce "pinching" artefacts.

The problem is that the error in reconstructed gradients is penalized the same regardless of whether the value of the gradient is small or large. This is opposite to how the visual system perceives differences in color values: we are more likely to spot tiny difference between two similar pixel values than the same tiny difference between two very different pixel values. We could account for that effect by introducing some form of non-linear metric, however, that would make our problem non-linear and non-linear problems are in general much slower to solve. However, the same can be achieved by introducing weights to our objective function:

$$
\begin{equation*}
\underset{I}{\arg \min } \sum_{x, y}\left[w_{x, y}^{(x)}\left(I_{x+1, y}-I_{x, y}-G_{x, y}^{(x)}\right)^{2}+w_{x, y}^{(y)}\left(I_{x, y+1}-I_{x, y}-G_{x, y}^{(y)}\right)^{2}\right], \tag{13}
\end{equation*}
$$

where $w_{x, y}^{(x)}$ and $w_{x, y}^{(y)}$ are the weights or importance we assign to each gradient, for horizontal and vertical partial derivatives respectively. Usually the weights are kept the same for both orientations, i.e. $w_{x, y}^{(x)}=w_{x, y}^{(y)}$. To account for the contrast perception of the visual system, we need to assign a higher weight to small gradient magnitudes. For example, we could use the weight:

$$
\begin{equation*}
w_{x, y}^{(x)}=w_{x, y}^{(y)}=\frac{1}{\left\|G_{x, y}\right\|+\epsilon} \tag{14}
\end{equation*}
$$

where $\left\|G_{x, y}\right\|$ is the magnitude of the desired (target) gradient at pixel $(x, y)$ and $\epsilon$ is a small constant (0.0001), which prevents division by 0 .

## 4 Matrix notation

We could follow the same procedure as in the previous section and differentiate Equation 13 to find the linear system that minimizes our objective. However, the process starts to be tedious and error-prone. As the objective functions gets more and more complex, it is worth switching to the matrix notation. Let us consider first our original problem without the weights $w_{x, y}$, which we will add later. Equation 3 in the matrix notation can be written as:

$$
\underset{I}{\arg \min }\left\|\left[\begin{array}{c}
\nabla_{x}  \tag{15}\\
\nabla_{y}
\end{array}\right] I-\left[\begin{array}{l}
G^{(x)} \\
G^{(y)}
\end{array}\right]\right\|^{2} .
$$

In the equation $I, G^{(x)}$ and $G^{(y)}$ are stacked column vectors, representing columns of the resulting image or desired gradient field. The square brackets
denote vertical concatenation of the matrices or vectors. Operator $\|\cdot\|^{2}$ is the $L_{2}$-norm, which squares and sums the elements of the resulting column vector. $\nabla_{x}$ and $\nabla_{y}$ are differential operators, which are represented as $N \times N$ matrices, where $N$ is the number of pixels. Each row of those sparse matrices tells us which pixels need to be subtracted from one another to compute forward gradients along horizontal and vertical directions. For a tiny $3 \times 3$ pixel image those operators are:

$$
\begin{align*}
& \nabla_{x}=\left[\begin{array}{ccccccccc}
-1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & -1 & 0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & -1 & 0 & 0 & 1 & 0 & 0 \\
0 & 0 & 0 & 0 & -1 & 0 & 0 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
\end{array}\right]  \tag{16}\\
& \nabla_{y}=\left[\begin{array}{ccccccccc}
-1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0
\end{array}\right] \tag{17}
\end{align*}
$$

Note that the rows contain all zeros for pixels on the boundary, for which no gradient can be computed: the last column of pixels for $\nabla_{x}$ and the last row of pixels for $\nabla_{y}$.

Equation 15 is in the format $\|A x-b\|^{2}$, which can be directly solved by some sparse matrix libraries, such as SciPy.sparse or the " $\backslash$ " operator in matlab Matlab. However, to reduce the size of the sparse matrix and to speed-up computation, it is worth taking one more step and transform the least-square optimization into a linear problem. For overdetermined systems, such as ours, the solution of the optimization problem:

$$
\begin{equation*}
\underset{x}{\arg \min }\|A x-b\|^{2} \tag{18}
\end{equation*}
$$

can be found by solving a linear system:

$$
\begin{equation*}
A^{\prime} A x=A^{\prime} b . \tag{19}
\end{equation*}
$$

Note that ' denotes a matrix transpose and $A^{\prime} A$ is a square matrix. If we replace $A$ and $b$ with the corresponding operators and gradient values from our problem, we get the following linear system:

$$
\left[\begin{array}{ll}
\nabla_{x}^{\prime} & \nabla_{y}^{\prime}
\end{array}\right]\left[\begin{array}{l}
\nabla_{x}  \tag{20}\\
\nabla_{y}
\end{array}\right] I=\left[\begin{array}{ll}
\nabla_{x}^{\prime} & \nabla_{y}^{\prime}
\end{array}\right]\left[\begin{array}{l}
G^{(x)} \\
G^{(y)}
\end{array}\right],
$$

which, after multiplying stacked matrices, gives us:

$$
\begin{equation*}
\left(\nabla_{x}^{\prime} \nabla_{x}+\nabla_{y}^{\prime} \nabla_{y}\right) I=\nabla_{x}^{\prime} G^{(x)}+\nabla_{y}^{\prime} G^{(y)} . \tag{21}
\end{equation*}
$$

Weights can be added to such a system by inserting a sparse diagonal matrix $W$. For simplicity we use the same weights for vertical and horizontal derivatives:

$$
\begin{equation*}
\left(\nabla_{x}^{\prime} W \nabla_{x}+\nabla_{y}^{\prime} W \nabla_{y}\right) I=\nabla_{x}^{\prime} W G^{(x)}+\nabla_{y}^{\prime} W G^{(y)} . \tag{22}
\end{equation*}
$$

The above operations can be performed using a sparse matrix library (or SciPy/Matlab), thus saving us effort of computing operators by hand.

There is still one problem remaining: our equation does not have a unique solution. This is because the target gradient field contains relative information about differences between pixels, but it does not say what the absolute value of the pixels should be. For that reason, we need to constrain the absolute value, for example by ensuring that a value of a first reconstructed pixel is the same as in the source image $\left(I_{s r c}\right)$ :

$$
\left[\begin{array}{llll}
1 & 0 & \ldots & 0 \tag{23}
\end{array}\right] I=I_{s r c}(1,1)
$$

If we denote the vector on the left-hand side of the equation as $C$, the final linear problem can be written as:

$$
\begin{equation*}
\left(\nabla_{x}^{\prime} W \nabla_{x}+\nabla_{y}^{\prime} W \nabla_{y}+C^{\prime} C\right) I=\nabla_{x}^{\prime} W G^{(x)}+\nabla_{y}^{\prime} W G^{(y)}+C^{\prime} I_{s r c}(1,1) . \tag{24}
\end{equation*}
$$

The resulting equation can be solved using a sparse solver in Python or Matlab.

## References

[1] S. A. Teukolsky, B. P. Flannery, W. H. Press, and W. T. Vetterling. Numerical recipes in C. Cambridge University Press, Cambridge, vol. 2 edition, 1992.

# Advanced Graphics and Image Processing Lecture notes 

Rafał Mantiuk

Lent term 2018/19

## 1 Light field rendering using homographic transformation

This section explains how to render a light field for a novel view position using a parametrization with a focal plane. The method is explained on a rather high level in [1]. These notes are meant to provide a practical guide on how to perform the required calculations and in particular how to find a homographic transformation between the virtual and array cameras.

The scenario and selected symbols are illustrated in Figure 1. We want to render our light field "as seen" by camera $K$. We have $N$ images captured by $N$ cameras in the array (only 4 shown in the figure), all of which have their apertures on the camera array plane $C$. We further assume that our array cameras are pin-hole cameras to simplify the explanation. The novel view is rendered assuming focal plane $F$. The focal plane has a similar function as the focus distance in a regular camera: objects on the focal plane will be rendered sharp, while objects that and in front or behind that plane will appear blurry (in practice they will appear ghosted because of the limited number of cameras). The focal plane $F$ does not need to be parallel to the camera plane; it can be titled, unlike in a traditional camera with a regular lens. Because we have a limited number of cameras, we need to use reconstruction functions $A_{0}, \ldots, A_{1}$ (only two shown) for each camera. The functions shown contain the weights in the range $0-1$ that are used to interpolate between two neighboring views.

To intuitively understand how light field rendering is performed, imagine the following hypothetical scenario. Each camera in the array captures the


Figure 1: Light field rendering for the novel view represented by camera $K$. The pixels $P_{K}$ in the rendered image is the weighted average of the pixels values $p_{1}, \ldots, p_{N}$ from the images captured by the camera array.
image of the scene. Then, all objects in the scene are removed and you put a large projection screen where the focal plane $F$ should be. Then, you swap all cameras for projectors, which project the captured images on the projection screen $F$. Finally, you put a new camera $K$ at the desired location and capture the image of the projection screen. The projection screen (focal plane) is needed to form an image. Ideally, to obtain a sharp image, we would like to project the camera array images on a geometry. However, such a geometry is not readily available when capturing scenes with a camera array. In this situation a single plane is often a good-enough proxy, which has its analogy in physical cameras (focal distance). More advanced light field rendering methods attempt to reconstruct a more accurate proxy geometry using multi-view stereo algorithms and then project camera images on that geometry [3].

This simplified scenario misses one step, which is modulating each projected image by the reconstruction function $A$, as such modulation has no physical counterpart. However, this scenario should give you a good idea what operations need to be performed in order to render a light field from a

Data: Camera array images $J_{1}, J_{2}, \ldots, J_{N}$
Result: Rendered image $I$
for each pixel at the coordinates $\boldsymbol{p}_{K}$ in the novel view do
$I\left(\boldsymbol{p}_{K}\right) \leftarrow 0 ;$
$w\left(\boldsymbol{p}_{K}\right) \leftarrow 0 ;$
for each camera $i$ in the array do
Find the coordinates $\boldsymbol{p}_{i}$ in the $i$-th camera image
corresponding to the pixel $\boldsymbol{p}_{K}$;
Find the coordinates $\boldsymbol{p}_{A}$ on the aperture plane $A$
corresponding to the pixel $\boldsymbol{p}_{K}$;
$I\left(\boldsymbol{p}_{K}\right) \leftarrow I\left(\boldsymbol{p}_{K}\right)+A\left(\boldsymbol{p}_{A}\right) J_{i}\left(\boldsymbol{p}_{i}\right) ;$
$W\left(\boldsymbol{p}_{K}\right) \leftarrow W\left(\boldsymbol{p}_{K}\right)+A\left(\boldsymbol{p}_{A}\right) ;$
end
$I\left(\boldsymbol{p}_{K}\right) \leftarrow I\left(\boldsymbol{p}_{K}\right) / W\left(\boldsymbol{p}_{K}\right) ;$
end
Algorithm 1: Light field rendering algorithm
novel view position.
Now let us see how we can turn such a high-level explanation into a practical algorithm. One way to render a light field is shown in Algorithm 1. The algorithm iterates over all pixels in the rendered image, then for each pixel it iterates over all cameras in the array. The resulting image is the weighted average of the camera images that are warped using homographic transformations. The weights are determined by the reconstruction functions $A_{i}$. The algorithm is straightforward, except for the mapping from pixel coordinates in the novel view $p_{K}$ to coordinates in each camera image $p_{i}$ and the coordinates on the aperture plane $p_{A}$. The following paragraphs explain how to find such transformations.

### 1.1 Homographic transformation between the virtual and array cameras

The text below assumes that you are familiar with homogeneous coordinates and geometric transformations, both commonly used in computer graphics and computer vision. If these topics are still unclear, refer to Section 2.1 in [4] (this book is available online) or Chapter 6 in [2].

We assume that we know the position and pose of each camera in the
array, so that homogeneous 3 D coordinates of a point in the 3D word coordinate space $w$ can be mapped to the 2D pixel coordinates $p_{i}$ of camera $i$ :

$$
\begin{equation*}
\boldsymbol{p}_{i}=K P V_{i} \boldsymbol{w} \tag{1}
\end{equation*}
$$

where $\boldsymbol{V}$ is the view transformation, $\boldsymbol{P}$ is the projection matrix and $\boldsymbol{K}$ is the intrinsic camera matrix. Note that we will use bold lower case symbols to denote vectors, uppercase bold symbols for matrices and a regular font for scalars. The operation is easier to understand if the coordinates and matrices are expanded:

$$
\left[\begin{array}{l}
x_{i}  \tag{2}\\
y_{i} \\
w_{i}
\end{array}\right]=\left[\begin{array}{ccc}
f_{x} & 0 & c_{x} \\
0 & f_{y} & c_{y} \\
0 & 0 & 1
\end{array}\right]\left[\begin{array}{cccc}
1 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 \\
0 & 0 & 1 & 0
\end{array}\right]\left[\begin{array}{cccc}
v_{11} & v_{12} & v_{13} & v_{14} \\
v_{21} & v_{22} & v_{23} & v_{24} \\
v_{31} & v_{32} & v_{33} & v_{34} \\
0 & 0 & 0 & 1
\end{array}\right]\left[\begin{array}{c}
X \\
Y \\
Z \\
1
\end{array}\right] .
$$

The view matrix $\boldsymbol{V}$ translates and rotates the 3D coordinates of the 3D point $\boldsymbol{w}$ so that the origin of the new coordinate system is at the camera centre, and camera's optical axis is aligned with the $z$-axis (as the view matrix in computer graphics). This matrix can be computed using a LookAt function, often available in graphics matrix libraries.

The projection matrix $P$ may look like an odd version of an identity matrix, but it actually drops one dimension (projects from 3D to 2D) and copies the value of $Z$ coordinate into the additional homogeneous coordinate $w_{i}$. Note that to compute Cartesian coordinates of the point from the homogeneous coordinates, we divide $x_{i} / w_{i}$ and $y_{i} / w_{i}$. As $w_{i}$ is now equal to the depth in the camera coordinates, this operation is equivalent to a perspective projection (you can refer to slides 88-92 in the Introduction to Graphics Course).

The intrinsic camera matrix $K$ maps the projected 3D coordinates into pixel coordinates. $f_{x}$ and $f_{y}$ are focal lengths and $c_{x}$ and $c_{y}$ are the coordinates of optical center expressed in pixel coordinates. We assume that the intrinsic matrix is the same for all the cameras in the array.

Besides having all matrices for all cameras in the array, we also have a similar transformation for our virtual camera $K$, which represents the currently rendered view:

$$
\begin{equation*}
\boldsymbol{p}_{K}=\boldsymbol{K}_{K} \boldsymbol{P} \boldsymbol{V}_{K} \boldsymbol{w} \tag{3}
\end{equation*}
$$

Our first task is to find transformation matrices that could transform from pixel coordinates $\boldsymbol{p}_{K}$ in the virtual camera image into pixel coordinates $\boldsymbol{p}_{i}$
for each camera $i$. This is normally achieved by inverting the transformation matrix for the novel view and combining it with the camera array transformation. However, the problem is that the product of $K_{K} P V_{K}$ is not a square matrix that can be inverted - it is missing one dimension. The dimension is missing because we are projecting from 3D to 2D and one dimension (depth) is lost.

Therefore, to map both coordinates, we need to reintroduce missing information. This is achieved by assuming that the 3 D point lies on the focal plane $F$. Note that the plane equation can be expressed as $\boldsymbol{N} \cdot\left(\boldsymbol{w}-\boldsymbol{w}_{F}\right)=0$, where $\boldsymbol{N}$ is the plane normal, and $\boldsymbol{w}_{F}$ specifies the position of the plane in the 3D space. Operator • is the dot product. If the homogeneous coordinates of the point $\boldsymbol{w}$ are $\left[\begin{array}{llll}X & Y & Z & 1\end{array}\right]$, the plane equation can be expressed as

$$
d=\left[\begin{array}{llll}
n_{x} & n_{y} & n_{z} & -\boldsymbol{N} \cdot \boldsymbol{w}_{F}
\end{array}\right]\left[\begin{array}{c}
X  \tag{4}\\
Y \\
Z \\
1
\end{array}\right]
$$

where $d$ is the distance to the plane and $\boldsymbol{N}=\left[\begin{array}{lll}n_{x} & n_{y} & n_{z}\end{array}\right]$. We can introduce the plane equation into the projection matrix from Equation 2:

$$
\left[\begin{array}{c}
x_{i}  \tag{5}\\
y_{i} \\
d_{i} \\
w_{i}
\end{array}\right]=\left[\begin{array}{cccc}
f_{x} & 0 & 0 & c_{x} \\
0 & f_{y} & 0 & c_{y} \\
0 & 0 & 1 & 0 \\
0 & 0 & 0 & 1
\end{array}\right]\left[\begin{array}{cccc}
1 & 0 & 0 & 0 \\
0 & 1 & 0 & 0 \\
n_{x}^{(c)} & n_{y}^{(c)} & n_{z}^{(c)} & -\boldsymbol{N}^{(c)} \cdot \boldsymbol{w}_{F}^{(c)} \\
0 & 0 & 1 & 0
\end{array}\right]\left[\begin{array}{cccc}
v_{11} & v_{12} & v_{13} & v_{14} \\
v_{21} & v_{22} & v_{23} & v_{24} \\
v_{31} & v_{32} & v_{33} & v_{34} \\
0 & 0 & 0 & 1
\end{array}\right]\left[\begin{array}{c}
X \\
Y \\
Z \\
1
\end{array}\right] .
$$

The product of the matrices in above is a full $4 \times 4$ transformation matrix, which is not rank-deficient and can be inverted. Note that the pixel coordinates $\boldsymbol{p}_{K}$ and $\boldsymbol{p}_{i}$ now have an extra dimension $d$, which should be set to 0 (because we constrain 3D point $w$ to lie on the focal plane).

It should be noted that the normal and the point in the plane equation have superscript ${ }^{(c)}$, which denotes that the plane is given in the camera coordinate system, rather than in the world coordinate system. This is because the point $\left[\begin{array}{llll}X & Y & Z & 1\end{array}\right]$ is transformed from the world to the camera coordinates by the view matrix $V_{\boldsymbol{i}}$ before it is multiplied by our modified projection matrix. This could be a desired behavior for the virtual camera, for example if we want the focal plane to follow the camera and be perpendicular to the camera's optical axis. But, if we simply want to specify the focal plane in the
world coordinates, we have two options: either replace the third row in the final matrix (obtained after multiplying the three matrices in Equation 5) with our plane equation in the world coordinate system; or to transform the plane to the camera coordinates:

$$
\begin{equation*}
\boldsymbol{w}_{F}^{(c)}=\boldsymbol{V}_{i} \boldsymbol{w}_{F} \tag{6}
\end{equation*}
$$

and

$$
\begin{equation*}
\boldsymbol{N}^{(c)}=\overline{\boldsymbol{V}}_{i} \boldsymbol{N} \tag{7}
\end{equation*}
$$

$\overline{\boldsymbol{V}}_{i}$ is the "normal" or direction transformation for the view matrix $\boldsymbol{V}_{i}$, which rotates the normal vector but it does not translate it. It is obtained by setting to zero the translation coefficients $w_{14}, w_{24}$, and $w_{34}$.

Now let us find the final mapping from the virtual camera coordinates $\hat{\boldsymbol{p}_{K}}$ to the array camera coordinates $\hat{\boldsymbol{p}}_{i}$. We will denote the extended coordinates (with extra $d$ ) in Equation 5 as $\hat{\boldsymbol{p}_{K}}$ and $\hat{\boldsymbol{p}_{i}}$. We will also denote our new projection and intrinsic matrices in Equation 5 as $\hat{\boldsymbol{P}}$ and $\hat{\boldsymbol{K}}$. Given that, the mapping from $\boldsymbol{p}_{K}$ to $\boldsymbol{p}_{i}$ can be expressed as:

$$
\begin{equation*}
\hat{\boldsymbol{p}}_{i}=\hat{\boldsymbol{K}}_{i} \hat{\boldsymbol{P}} V_{i} \boldsymbol{V}_{K}^{-1} \hat{\boldsymbol{P}}^{-1} \hat{\boldsymbol{K}}_{K}^{-1} \hat{\boldsymbol{p}_{K}} . \tag{8}
\end{equation*}
$$

The position on the aperture plane $\boldsymbol{w}_{A}$ can be readily found from:

$$
\begin{equation*}
\boldsymbol{w}_{A}=\boldsymbol{V}_{K}^{-1} \hat{\boldsymbol{P}}_{A}^{-1} \hat{\boldsymbol{K}}_{K}^{-1} \hat{\boldsymbol{p}_{K}} \tag{9}
\end{equation*}
$$

where the projection matrix $\hat{P}_{A}$ is modified to include the plane equation of the aperture plane, the same way as done in Equation 5 .

### 1.2 Reconstruction functions

The choice of the reconstruction function $A_{i}$ will determine how images from different cameras are mixed together. The functions shown in Figure 1 will perform bilinear-interpolation between the views. Although this could be a rational choice, it will result in ghosting for the parts of the scene that are further away from the focal plane $F$. Another choice is to simulate a wideaperture camera and include all cameras in the generated view ( set $A_{i}=1$ ). This will produce an image with a very shallow depth of field. Another possibility is to use the nearest-neighbor strategy and a box-shaped reconstruction filter (the width of the boxes being equal to the distance between the cameras). This approach will avoid ghosting but will cause the views
to jump sharply as the virtual camera moves over the scene. It is worth experimenting with different reconstruction startegies to choose the best for a given application but also for the given angular resolution of the light field (number of views).

## References

[1] Aaron Isaksen, Leonard McMillan, and Steven J. Gortler. Dynamically reparameterized light fields. In Proc of SIGGRAPH '00, volume 7, pages 297-306, New York, New York, USA, 2000. ACM Press.
[2] Steve Marschner and Peter Shirley. Fundamentals of Computer Graphics. A K Peters/CRC Press, 4 edition edition, 2016.
[3] Ryan S. Overbeck, Daniel Erickson, Daniel Evangelakos, Matt Pharr, and Paul Debevec. A system for acquiring, processing, and rendering panoramic light field stills for virtual reality. ACM Transactions on Graphics, 37(6):1-15, dec 2018.
[4] Richard Szeliski. Computer Vision: Algorithms and Applications. Springer-Verlag New York Inc, 2010.


[^0]:    - 6

[^1]:    > 23

[^2]:    12 images from [Drori at al. 2004]

[^3]:    24 From: OpenCL 1.2 University Kit- hitp://developer.amd.com/patners/university-proarame

[^4]:    - 45

[^5]:    ignores luminance

[^6]:    29

[^7]:    > 36

[^8]:    II

[^9]:    - 28

