

Artistic Stylization of Images and Video

John Collomosse¹ Jan Eric Kyprianidis²

¹Centre for Vision Speech and Signal Processing (CVSSP), University of Surrey, Guildford, United Kingdom

²Hasso-Plattner-Institut, University of Potsdam, Germany

Abstract

The half-day tutorial provides an introduction to Non-Photorealistic Rendering (NPR), targeted at both students and experienced researchers of Computer Graphics who have not previously explored NPR in their work. The tutorial focuses on two-dimensional (2D) NPR, specifically the transformation of photos or videos into synthetic artwork (e.g. paintings or cartoons). Consequently the course will touch not only on computer graphics topics, but also on the image processing and computer vision techniques that drive such algorithms. However the latter concepts will be introduced gently and no prior knowledge is assumed beyond a working knowledge of filtering and convolution operations. Some elements of the course will touch upon GPU implementation, but GPU concepts will be described at a high level of abstraction without need for detailed working knowledge of GPU programming.

References

- [Aga02] AGARWALA A.: Snaketoonz: A semi-automatic approach to creating cel animation from video. In *NPAR '02: Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering* (2002).
- [BBPW04] BROX T., BRUHN A., PAPENBERG N., WEICKERT J.: High accuracy optical flow estimation based on a theory for warping. *European Conference on Computer Vision (ECCV)* (2004), 25–36.
- [BCGF10] BÉNARD P., COLE F., GOLOVINSKIY A., FINKELSTEIN A.: Self-similar texture for coherent line stylization. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering (NPAR)* (2010), pp. 91–97.
- [BGH03] BANGHAM J., GIBSON S., HARVEY R.: The art of scale-space. In *Proc. British Machine Vision Conference* (2003), pp. 569–578.
- [BNTS07] BOUSSEAU A., NEYRET F., THOLLOT J., SALESIN D.: Video watercolorization using bidirectional texture advection. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers* (2007), p. 104.
- [BWSS09] BAI X., WANG J., SIMONS D., SAPIRO G.: Video snapcut: robust video object cutout using localized classifiers. *ACM SIGGRAPH 2009 papers* (2009), 1–11.
- [CAS*97] CURTIS C. J., ANDERSON S. E., SEIMS J. E., FLEISCHER K. W., SALESIN D. H.: Computer-generated watercolor. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (1997), pp. 421–430.
- [CGL*08] COLE F., GOLOVINSKIY A., LIMPAECHER A., BARROS H., FINKELSTEIN A., FUNKHOUSER T., RUSINKIEWICZ S.: Where do people draw lines? *ACM Transactions on Graphics* 27, 3 (2008), 1–11.
- [CH03] COLLOMOSSE J., HALL P.: Cubist style rendering from photographs. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 4, 9 (2003), 443–453.
- [CH05] COLLOMOSSE J., HALL P.: Genetic paint: A search for salient paintings. In *Proceedings of EvoMUSART (LNCS)* (March 2005), vol. 3449, Springer, pp. 437–447.
- [CH06] COLLOMOSSE J., HALL P.: Video motion analysis for the synthesis of dynamic cues and futurist art. *Graphical Models* 68, 5–6 (2006), 402–414.
- [CRH03] COLLOMOSSE J. P., ROWNTREE D., HALL P. M.: Video analysis for cartoon-style special effects. In *Proceedings 14th British Machine Vision Conference (BMVC)* (September 2003), vol. 2, pp. 749–758.
- [CRH05] COLLOMOSSE J. P., ROWNTREE D., HALL P. M.: Stroke surfaces: Temporally coherent non-photorealistic animations from video. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* 11, 5 (September 2005), 540–549.
- [DS02] DECARLO D., SANTELLA A.: Stylization and abstraction of photographs. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques* (2002), pp. 769–776.
- [DS10] DECARLO D., STONE M.: Visual explanations. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering* (2010), NPAR '10, pp. 173–178.
- [Dur02] DURAND F.: An invitation to discuss computer depiction. In *NPAR '02: Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering* (2002), pp. 111–124.
- [FBS05] FISCHER J., BARTZ D., STRABER W.: Stylized augmented reality for improved immersion. In *Proc. VR, IEEE Virtual Reality* (2005), pp. 195–202.
- [GCS02] GOOCH B., COOMBE G., SHIRLEY P.: Artistic vision: painterly rendering using computer vision techniques. In *NPAR '02: Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering* (New York, NY, USA, 2002), ACM, pp. 83–ff.
- [GG01] GOOCH B., GOOCH A. A.: *Non-Photorealistic Rendering*. AK Peters, 2001.
- [Goo10] GOOCH A. A.: Towards mapping the field of non-photorealistic rendering. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering* (2010), NPAR '10, pp. 159–164.
- [Gra18] GRAY H.: *Anatomy of the human body*. Lea & Febiger, 1918.
- [GRG04] GOOCH B., REINHARD E., GOOCH A.: Human facial illustrations: Creation and psychophysical evaluation. *ACM Transactions on Graphics* 23, 1 (2004), 27–44.
- [GSS*99] GREEN S., SALESIN D., SCHOFIELD S., HERTZMANN A., LITWINOWICZ P., GOOCH A. A., CURTIS C., GOOCH B.: Non-photorealistic rendering. In *SIGGRAPH '99: ACM SIGGRAPH 1999 courses* (1999).
- [Hae90] HAEBERLI P.: Paint by numbers: abstract image representations. In *SIGGRAPH '90: Proceedings of the 17th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1990), ACM, pp. 207–214.
- [Hag91] HAGGERTY M.: Almost automatic computer painting. *IEEE Computer Graphics and Applications* 11 (1991), 11–12.
- [HCS*07] HALL P., COLLOMOSSE J., SONG Y., SHEN P., LI C.: Rtcams: A new perspective on nonphotorealistic rendering from photographs. *IEEE Transactions on Visualization and Computer Graphics* (2007), 966–979.
- [HE02] HEALEY C. G., ENNS J. T.: Perception and painting: A search for effective, engaging visualizations. *IEEE Comput. Graph. Appl.* 22 (2002), 10–15.
- [HE04] HAYS J., ESSA I.: Image and video based painterly animation. In *NPAR '04: Proceedings of the 3rd international symposium on Non-photorealistic animation and rendering* (2004), pp. 113–120.
- [Her98] HERTZMANN A.: Painterly rendering with curved brush strokes of multiple sizes. In *SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (1998), pp. 453–460.
- [Her01] HERTZMANN A.: Paint by relaxation. In *CGI '01: Computer Graphics International 2001* (2001), pp. 47–54.
- [Her02] HERTZMANN A.: Fast paint texture. In *NPAR '02: Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering* (2002), pp. 91–ff.
- [Her10] .: *Non-Photorealistic Rendering and the science of art* (2010).
- [HJO*01] HERTZMANN A., JACOBS C. E., OLIVER N., CURLESS B., SALESIN D. H.: Image analogies. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (2001), pp. 327–340.
- [HP00] HERTZMANN A., PERLIN K.: Painterly rendering for video and interaction. In *NPAR '00: Proceedings of the 1st international symposium on Non-photorealistic animation and rendering* (2000), pp. 7–12.
- [IEE05] IEEE.: *Synergism in low level vision* (2005), vol. 4.

- [INC*06] ISENBERG T., NEUMANN P., CARPENDALE S., SOUSA M. C., JORGE J. A.: Non-photorealistic rendering in context: an observational study. In *NPAR '06: Proceedings of the 4th international symposium on Non-photorealistic animation and rendering* (2006), pp. 115–126.
- [KD08] KYPRIANIDIS J. E., DÖLLNER J.: Image abstraction by structure adaptive filtering. In *Proc. EG UK Theory and Practice of Computer Graphics* (2008), Eurographics Association, pp. 51–58.
- [KD09] KYPRIANIDIS J. E., DÖLLNER J.: Real-time image abstraction by directed filtering. In *ShaderX7 - Advanced Rendering Techniques*, Engel W., (Ed.). Charles River Media, 2009.
- [KHEK76] KUWAHARA M., HACHIMURA K., EIHO S., KINOSHITA M.: *Digital processing of biomedical images*. Plenum Press, 1976, pp. 187–203.
- [KKD09] KYPRIANIDIS J. E., KANG H., DÖLLNER J.: Image and video abstraction by anisotropic kuwahara filtering. *Computer Graphics Forum* 28, 7 (2009), 1955–1963. Special issue on Pacific Graphics 2009.
- [KKD10] KYPRIANIDIS J. E., KANG H., DÖLLNER J.: Anisotropic kuwahara filtering on the gpu. In *GPUPro - Advanced Rendering Techniques*, Engel W., (Ed.). AK Peters, 2010.
- [KL08] KANG H., LEE S.: Shape-simplifying image abstraction. *Computer Graphics Forum* 27, 7 (2008), 1773–1780. Special issue on the Pacific Graphics 2008.
- [KLC07] KANG H., LEE S., CHUI C. K.: Coherent line drawing. In *Proc. ACM NPAR* (2007), pp. 43–50.
- [KLC09] KANG H., LEE S., CHUI C. K.: Flow-based image abstraction. *IEEE Transactions on Visualization and Computer Graphics* 15, 1 (2009), 62–76.
- [Kol06] :. *Segmentation-based 3d artistic rendering* (2006).
- [KSFC02] KLEIN A., SLOAN P., FINKELSTEIN A., COHEN M.: Stylized video cubes. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2002), pp. 15–22.
- [Lit97] LITWINOWICZ P.: Processing images and video for an impressionist effect. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (1997), pp. 407–414.
- [LSRY10] LEE H., SEO S., RYOO S., YOON K.: Directional texture transfer. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering (NPAR)* (2010), pp. 43–48.
- [LTF*05] LIU C., TORRALBA A., FREEMAN W., DURAND F., ADELSON E.: Motion magnification. *ACM SIGGRAPH 2005 Papers* (2005), 519–526.
- [LYT*08] LIU C., YUEN J., TORRALBA A., SIVIC J., FREEMAN W. T.: Sift flow: Dense correspondence across different scenes. In *Proceedings of the 10th European Conference on Computer Vision (ECCV)* (2008), pp. 28–42.
- [LZL*10] LIN L., ZENG K., LV H., WANG Y., XU Y., ZHU S.-C.: Painterly animation using video semantics and feature correspondence. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering* (2010), NPAR '10, pp. 73–80.
- [McC90] MCCORDUCK P.: *Aaron's Code: Meta-art, Artificial Intelligence and the Work of Harold Cohen*. W.H. Freeman & Co, 1990.
- [Mei96] MEIER B. J.: Painterly rendering for animation. In *SIGGRAPH '96: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1996), ACM, pp. 477–484.
- [MHIL02] MA K.-L., HERTZMANN A., INTERRANTE V., LUM E. B.: Recent advances in non-photorealistic rendering for art and visualization. In *SIGGRAPH '02: ACM SIGGRAPH 2002 courses* (2002).
- [PKTD07] PARIS S., KORNPBST P., TUMBLIN J., DURAND F.: A gentle introduction to bilateral filtering and its applications. In *ACM SIGGRAPH courses* (2007).
- [PPC07] PAPARI G., PETKOV N., CAMPISI P.: Artistic edge and corner enhancing smoothing. *IEEE Transactions on Image Processing* 16, 10 (2007), 2449–2462.
- [PvV05] PHAM T., VAN VLIET L.: Separable bilateral filtering for fast video preprocessing. In *IEEE International Conference on Multimedia and Expo* (2005).
- [Red11] REDMOND N.: *Influencing User Perception Using Real-time Adaptive Abstraction*. PhD thesis, Trinity College Dublin, 2011.
- [SBC06] SHUGRINA M., BETKE M., COLLOMOSSE J.: Empathic painting: Interactive stylization using observed emotional state. In *Proceedings ACM 4th Intl. Symposium on Non-photorealistic Animation and Rendering (NPAR)* (2006), pp. 87–96.
- [SD04] SANTELLA A., DECARLO D.: Visual interest and npr: an evaluation and manifesto. In *NPAR '04: Proceedings of the 3rd international symposium on Non-photorealistic animation and rendering* (2004), pp. 71–150.
- [SS02] STROTHOTTE T., SCHLECHTWEIG S.: *Non-photorealistic Computer Graphics*. Morgan Kaufmann, 2002.
- [Str86] STRASSMANN S.: Hairy brushes. *SIGGRAPH Comput. Graph.* 20 (1986), 225–232.
- [SWHS97] SALISBURY M. P., WONG M. T., HUGHES J. F., SALESIN D. H.: Orientable textures for image-based pen-and-ink illustration. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques* (1997), pp. 401–406.
- [SY00] SHIRAISHI M., YAMAGUCHI Y.: An algorithm for automatic painterly rendering based on local source image approximation. In *NPAR '00: Proceedings of the 1st international symposium on Non-photorealistic animation and rendering* (2000), pp. 53–58.
- [TC97] TREAVETT S., CHEN M.: Statistical techniques for the automated synthesis of non-photorealistic images. In *Proc. 15th Eurographics UK Conference* (1997), pp. 201–210.
- [TKS10] TATZGERN M., KALKOFEN D., SCHMALSTIEG D.: Compact explosion diagrams. In *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering* (2010), NPAR '10, pp. 17–26.
- [TM98] TOMASI C., MANDUCHI R.: Bilateral filtering for gray and color images. In *Proceedings International Conference on Computer Vision (ICCV)* (1998), pp. 839–846.
- [WCS*10] WANG T., COLLOMOSSE J., SLATTER D., CHEATLE P., GREIG D.: Video stylization for digital ambient displays of home movies. In *NPAR '10: Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering* (2010), pp. 137–146.
- [WOG06] WINNEMÖLLER H., OLSEN S. C., GOOCH B.: Real-time video abstraction. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers* (2006), pp. 1221–1226.
- [WXSC04] WANG J., XU Y., SHUM H.-Y., COHEN M. F.: Video tooning. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Papers* (2004), pp. 574–583.
- [ZZXZ10] ZENG K., ZHAO M., XIONG C., ZHU S.-C.: From image parsing to painterly rendering. *ACM Trans. Graph.* 29 (2010), 1–11.

Artistic Stylization of Images and Video


Eurographics 2011

John Collomosse and Jan Eric Kyprianidis

Centre for Vision Speech and Signal Processing (CVSSP)
University of Surrey, United Kingdom

Hasso-Plattner-Institut, University of Potsdam, Germany

■ Texts



Strothotte &
Schlechtweg
ISBN: 1558607870

Gooch & Gooch
ISBN: 1568811330

Romero & Machado
ISBN: 3540728767

■ Web Bibliographies

<http://video3d.ims.tuwien.ac.at/%7Estathis/nprlib/index.php>

<http://isgwww.cs.uni-magdeburg.de/~stefans/npr/nprpapers.html>

<http://www.red3d.com/cwr/npr/> (dated)

■ Tutorials

SIGGRAPH 99 (Green et al.) – 2D/3D NPR
SIGGRAPH 02 (Hertzmann) – 2D NPR
SIGGRAPH 03 (Sousa et al.) – 2D/3D NPR
Eurographics 05,06 and...
SIGGRAPH 06 (Sousa et al) – 3D NPR
SIGGRAPH 10 (McGuire) – 3D NPR for Games

■ Main Publication Forums

NPAR (Symposium on Non-photorealistic Animation)
Held in Annecy even years, at SIGGRAPH odd years.

IEEE Trans Visualization and Comp. Graphics (**TVCG**)
IEEE Computer Graphics and Applications (**CG&A**)

Eurographics and **Computer Graphics Forum**
SIGGRAPH, **SIGGRAPH Asia** and **ACM ToG**

EG Symposium on Rendering (**EGSR**)

ACM/EG Symposium on Computer Animation (**EGSA**)

Non-Photorealistic Rendering (NPR)


Coined by Salesin et al., 1994

Stylized Rendering

Aesthetic Rendering

Artistic Stylization

Artistic Rendering



Anatomy of the Human Body
H. Gray, 1918



Artistic Stylization

■ Why?

Visualization

Comprehension

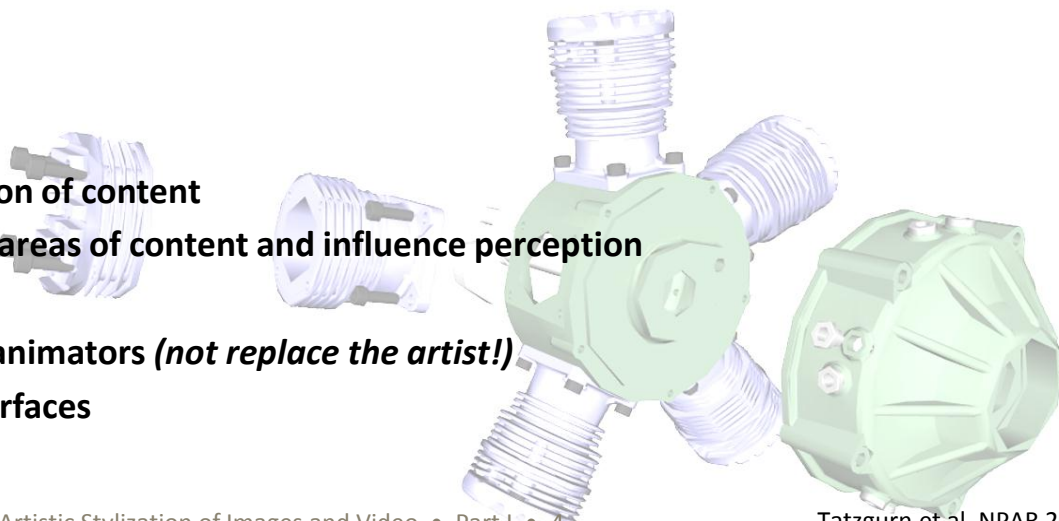
Communication

Aesthetics

Animation

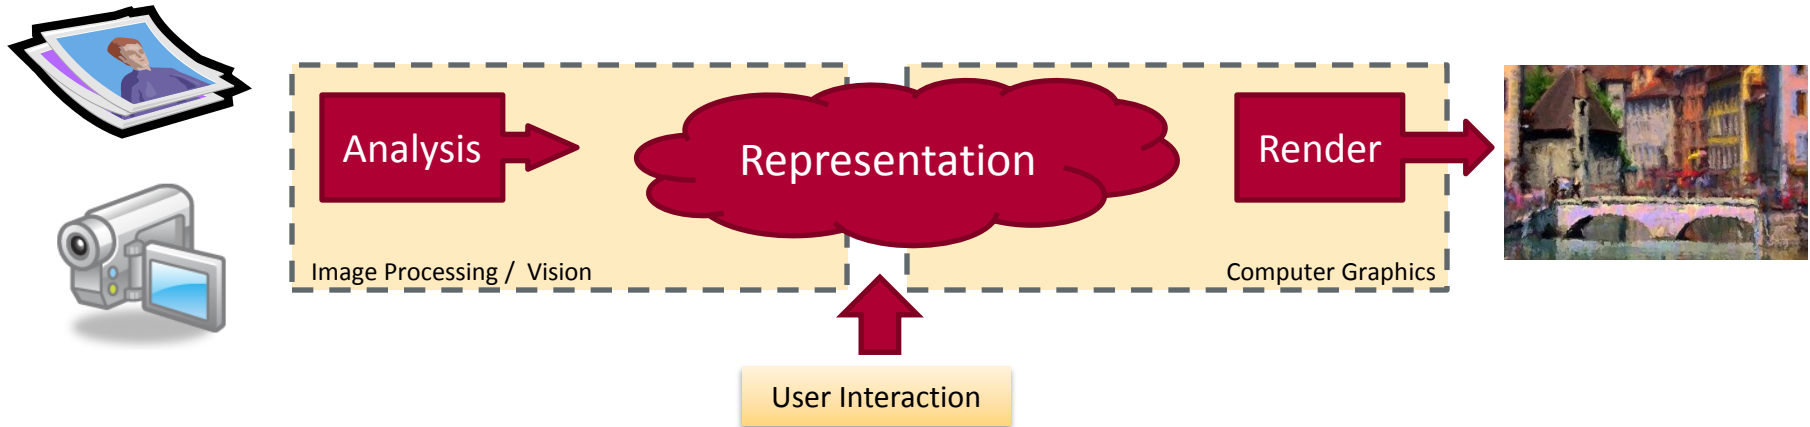
■ Artistic Stylization can

- Simplify and structure the presentation of content
- Selectively guide attention to salient areas of content and influence perception
- Learn and emulate artistic styles
- Provide assistive tools to artists and animators (*not replace the artist!*)
- Help us to design effective visual interfaces

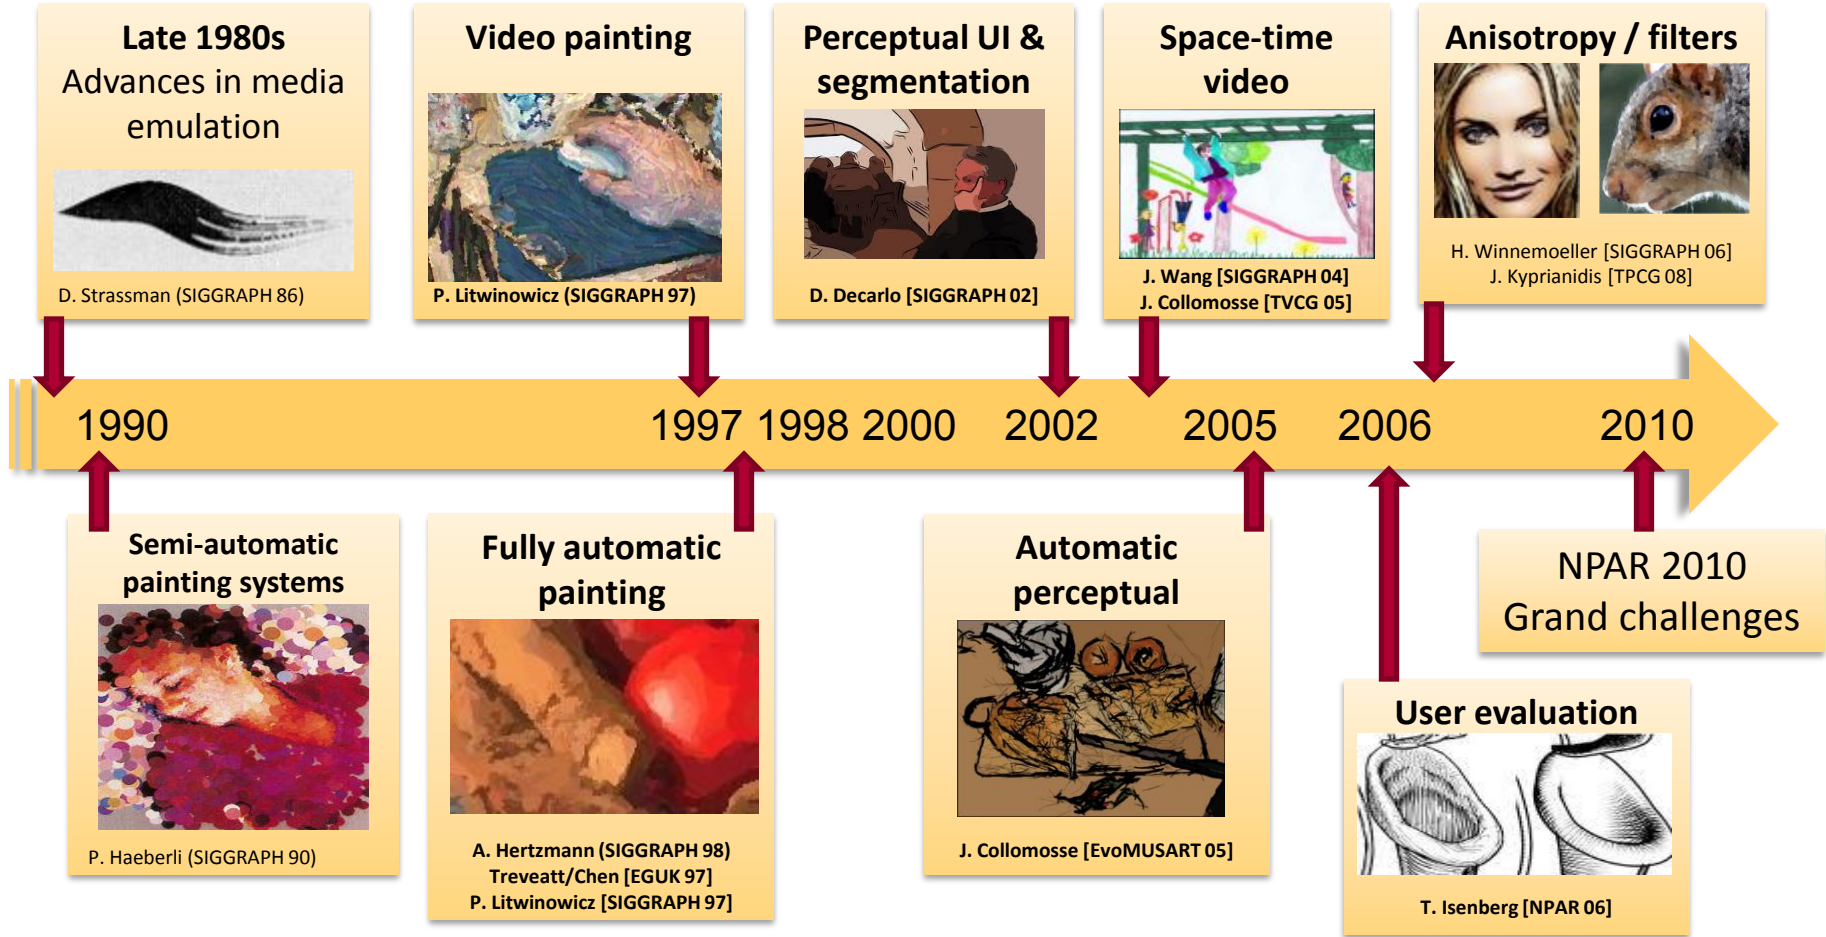


Artistic Stylization

- Rendering real images/video footage in to pseudo-artistic styles
- Convergence of Computer Vision, Graphics (and HCI)



- Visual analysis enables new graphics. Graphical needs motivate new vision.



Rendering process is guided by...

emulation

Perceptual UI & segmentation

User subconscious interaction

Space-time video

Anisotropy / filters

User conscious interaction

Low-level image processing

Higher level computer vision

Direct Anisotropic filtering

1990

1997 1998 2000

2002

2005

2006


2010

Semi-automatic painting systems




P. Haerberli [SIGGRAPH 90]

Fully automatic painting



A. Hertzmann [SIGGRAPH 98]
Treveatt/Chen [EGUK 97]
P. Litwinowicz [SIGGRAPH 97]

Automatic perceptual



J. Collomosse [EvoMUSART 05]

NPAR 2010 Grand challenges

User evaluation



T. Isenberg [NPAR 06]

Rendering process is guided by...

emulation

Perceptual UI & segmentation

User subconscious interaction

Part III: Anisotropy and Filtering (40 min)

User conscious interaction

Low-level image processing

Higher level computer vision

Image filtering



BREAK!

Part I: Classical algorithms (20 min)

Part IV: Future Challenges (10 min)

Part II: Vision for Stylisation (45 min)



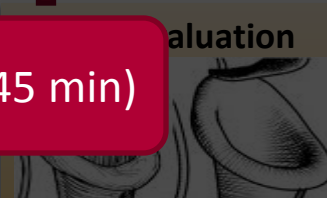
P. Haerberli [SIGGRAPH 90]



A. Hertzmann [SIGGRAPH 98]
Treveatt/Chen [EGUK 97]
P. Litwinowicz [SIGGRAPH 97]



J. Collomosse [EvoMUSART 05]



T. Isenberg [NPAR 06]

Artistic Stylization of Images and Video


Part I – Classical Algorithms /
Stroke Based Rendering
Eurographics 2011

John Collomosse

Centre for Vision Speech and Signal Processing (CVSSP),
University of Surrey, United Kingdom

- **Paint by numbers: Abstract image representations**
P. Haeberli, SIGGRAPH 1990
- **Almost Automatic Computer Painting**
P. Haggerty, IEEE CG & A 1991
- **Orientable Textures for Image based Pen-and-Ink Illustration**
D. Salisbury et al., SIGGRAPH 1997
- **Processing images and video for an impressionist effect**
P. Litwinowicz, SIGGRAPH 1997
- **Statistical techniques for the automated synthesis of non-photorealistic images**
S. Treavett and M. Chen, Eurographics UK 1997.
- **Automatic Painting based on Local Source Image Approximation**
Shiraishi and Yamaguchi, NPAR 2000.
- **Painterly Rendering with Curved Strokes of Multiple Sizes**
A. Hertzmann, SIGGRAPH 1998.
- **Paint by Relaxation**
A. Hertzmann, CGI 2001
- **Fast Paint Texture**
A. Hertzmann, NPAR 2002

- **Stroke based rendering (SBR)**
- **Painting is a manually ordered list of strokes, placed interactively.**
- **Stroke attributes sampled from the photo.**



- Stroke colour and orientation are sampled from the source image
- Stroke order and scale are user-selected
- Addition of RGB noise generates an impressionist effect

Photo credit: Haeberl '90.



Paintings with / without orientable strokes




Sobel Edge
detection

| | | |
|---|----|---|
| 1 | -2 | 1 |
| 0 | 0 | 0 |
| 1 | 2 | 1 |

$$E(I) = \left[\left(\frac{\partial I}{\partial x} \right)^2 + \left(\frac{\partial I}{\partial y} \right)^2 \right]^{\frac{1}{2}}$$

$$\theta(I) = \arctan \left(\frac{\partial I}{\partial y} / \frac{\partial I}{\partial x} \right)$$

Edge Mag.



Edge orient.



Orientation

- **More stylised orientation effects with a manually defined orientation field**



Orientation field



Painterly Rendering





Paint by numbers: Abstract Image Representations

Haeberli. (1990)



Code at <http://www.collomosse.com/EG2011tut/haeberlidemo.zip>

- **Very similar system for pen-and-ink rendering of photos**
- **User defined orientation field.**
 - **Regions manually drawn and marked up with orientation**
- **Stroke (line) placement automatic. Strokes clipped to keep within regions.**



- Stroke colour and orientation are sampled from the source image
 - Stroke order and scale ~~are user selected~~
 - Scale sampled from Sobel edge magnitude
 - Regularly place strokes. Order of strokes randomly generated
- } Fully automated

Photo credit: Haeberl '90.




Interactive (Haeberli)



Pseudo-random (as Haggerty)

**Loss of detail
in important
regions**



No clipping



Clipping





- **Common recipe for SBR in the 1990s**
 - **Sobel edge detection on blurred image**
 - **Regular seeding of strokes on canvas**
 - **Scale strokes inverse to edge magnitude**
 - **Orient strokes along edge tangent**
 - **Place strokes in a specific way using this data**

- **An interesting alternative uses 2nd order moments with local window to orient strokes.**
 - **Extended to multi-scale strokes by Shiraishi and Yamaguchi (NPAR 2000)**



- 2D zero-moments for greyscale image $I(x,y)$

$$M_{lm} = \sum_x \sum_y x^l y^m I(x,y).$$

- 1st order moments provide centre of mass.

$$x_c = \frac{M_{10}}{M_{00}} \quad y_c = \frac{M_{01}}{M_{00}}$$

- 2nd order moments describe grey variance.

$$a = \frac{M_{20}}{M_{00}} - x_c^2,$$

$$b = 2 \left(\frac{M_{11}}{M_{00}} - x_c y_c \right)$$


$$c = \frac{M_{02}}{M_{00}} - y_c^2.$$

- Orient strokes orthogonal to the direction of greatest variance about the centre of mass.


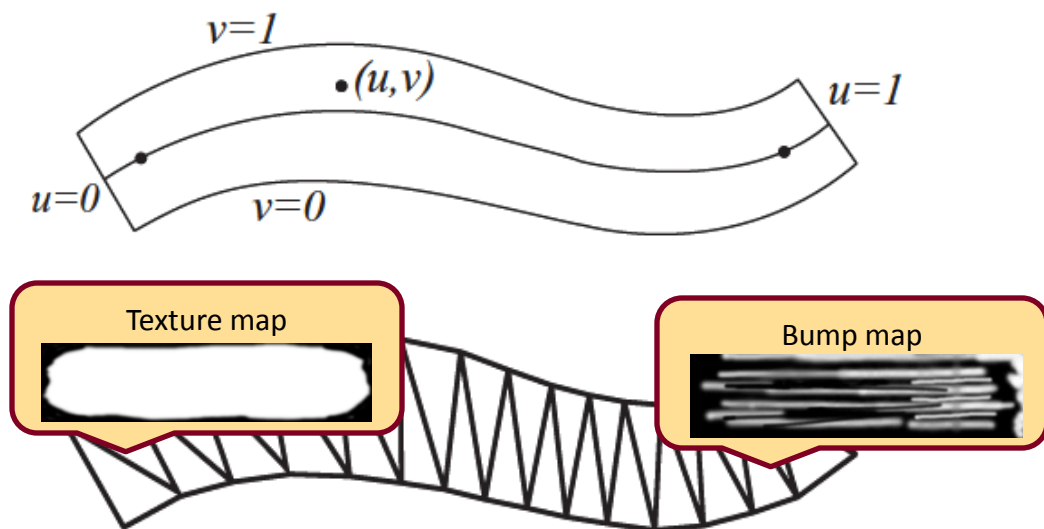
$$\theta = \frac{\tan^{-1} \left(\frac{b}{a-c} \right)}{2}$$

$$w = \sqrt{6 \left(a + c - \sqrt{b^2 + (a-c)^2} \right)}$$

$$l = \sqrt{6 \left(a + c + \sqrt{b^2 + (a-c)^2} \right)}$$




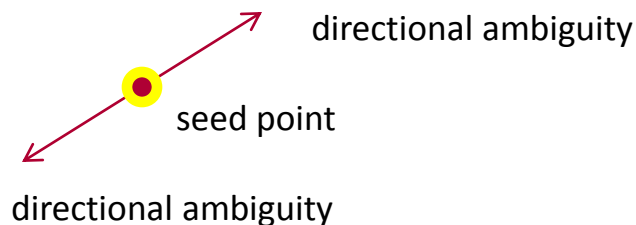
- Artists do not paint with uniformly shaped short strokes (pointillism excepted!)
- Two key contributions (1998)
 - Multi-layer (coarse to fine) painting
 - Painting using β -spline strokes
- Spline strokes can be bump mapped for an improved painterly look (NPAR 2002)





- Greedy algorithm for stroke placement
- Regularly sample the canvas to seed strokes
- Build a list of control point for each stroke by “hopping” between pixels*



- 1) Pick a direction arbitrarily
(some implementations explore both)



* In practice, best to use float coordinates and interpolate edge orientation


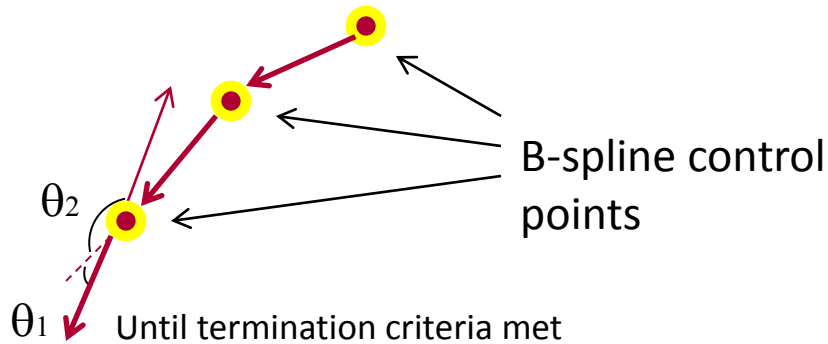
- Greedy algorithm for stroke placement
- Regularly sample the canvas to seed strokes
- Build a list of control point for each stroke by “hopping” between pixels*

2) Make another hop, resolving directional ambiguity by hopping in the direction of min θ



* In practice, best to use float coordinates and interpolate edge orientation


- Greedy algorithm for stroke placement
 - Regularly sample the canvas to seed strokes
 - Build a list of control point for each stroke by “hopping” between pixels*
- 3) Keep hopping until end land on a pixel whose RGB colour differs ($>$ threshold) from mean colour of stroke, or the stroke length is $>$ a second threshold.



* In practice, best to use float coordinates and interpolate edge orientation

- Painting is laid down in multiple layers (coarse to fine)
- Band-pass pyramid (= differenced layers of low-pass)
- Strokes from early layers are visible in final layer

- Paint coarsest layer with large strokes
- Paint next layer with smaller strokes
 - Only paint regions that differ between the layers
 - Use RGB difference



Compositing order





■ Tips and tricks

- Non-linear diffusion* instead of Gaussian blur sharpens the painting – preserves edges and accuracy of edge orientation.
- Build Gaussian pyramid at octave intervals, $\sigma=(1,2,4,8)$. 4 layers sufficient.
- Stroke thickness also at octave intervals
- Low-pass filter the hop direction θ



* “Scale-Space and Edge Detection using Anisotropic Diffusion”. P. Perona and J. Malik. *PAMI* 12:629–639. 1990.

- Global Optimization to Iteratively Produce “Better” Paintings



Hertzmann 1998
(Greedy stroke placement)



Hertzmann 2001
(Global stroke optimization)

- How to define the optimality of a painting 'P' derived from a photo 'G'

$$E(P) = E_{app}(P) + E_{area}(P) + E_{nstr}(P) + E_{cov}(P)$$

$$E_{app}(P) = \sum_{(x,y) \in \mathcal{I}} w_{app}(x,y) \|P(x,y) - G(x,y)\|$$

$$E_{area}(P) = w_{area} \sum_{S \in P} \text{Area}(S)$$

$$E_{nstr}(P) = w_{nstr} \cdot (\text{number of strokes in } P)$$

$$E_{cov}(P) = w_{cov} \cdot (\text{number of empty pixels in } P)$$

Weighted sum of Heuristics

Painting similar to photo - weighted

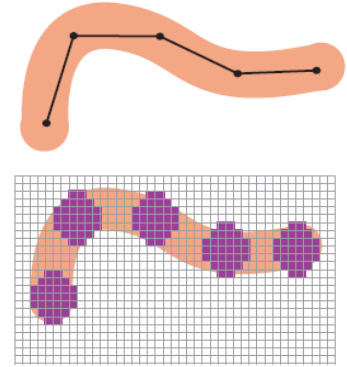
Stroke area ("paint used by artist")

Number of strokes


Fraction of canvas covered by strokes

- Weighting w_{app} is derived from a Sobel edge magnitude (or user defined)
- The right strokes in the right place will minimize the energy function $E(P)$


- Strokes selected at random and modified by local optimization to minimize $E(P)$
- Strokes modelled as active contours (“snakes”)
 - ... but energy is $\sim E(P)$ no 1st/2nd order derivative terms
 - $E(P)$ is approximated under control points
- Dynamic programming solution
 - move each control point to obtain locally optimal position (5x5)
 - $E(P)$ at control point dependent only on current and previous




$$\begin{aligned} s_0(v_1) &= \min_{v_0} e_0(v_0) + e_0(v_1) + e_1(v_0, v_1) \\ s_1(v_2) &= \min_{v_1} s_0(v_1) + e_0(v_2) + e_1(v_1, v_2) \\ &\vdots \\ s_{i-1}(v_i) &= \min_{v_{i-1}} s_{i-2}(v_{i-1}) + e_0(v_i) + e_1(v_{i-1}, v_i) \end{aligned}$$



- Sobel magnitude can be replaced with a manually sketched mask to alter emphasis



- Quick Start: OpenGL **research code** for bump-mapped paint strokes
 - Strokes as Catmull-Rom (interpolating) splines
 - Bump mapping via Multi-texturing (can be disabled)
 - Dependency on OpenCV to load images (can substitute this trivially)
 - Code used in “Empathic Painting” Collomosse et al. NPAR 2006



http://www.collomosse.com/EG2011tut/sbr_opengl.zip

Artistic Stylization of Images and Video

Part II – Vision for Stylisation

Eurographics 2011

John Collomosse

Centre for Vision Speech and Signal Processing (CVSSP),

University of Surrey, United Kingdom

- **Visual Interest and NPR: an Evaluation and Manifesto**
A. Santella and D. DeCarlo, NPAR 2004
- **Stylization and Abstraction of Photographs**
D. Decarlo, A. Santella, SIGGRAPH 2002
- **Segmentation-based 3D Artistic Rendering**
A. Kolliopoulos, J. Wang, A. Hertzmann, EGSR 2006.
- **Synergism in Low Level Vision (EDISON)**
C. Christoudias, B. Georgescu, P. Meer, ICPR 2002.
- **SIFT flow: dense correspondence across difference scenes**
C. Liu, J. Yuen, A. Torralba, J. Sivic, W. Freeman, ECCV 2008.
- **High Accuracy Optical Flow Estimation Based on a Theory for Warping**
T. Brox, A. Bruhn, N Papenberg, J. Weickert, ECCV 2004.
- **What dreams may come (movie)**
Dir. V. Ward. Universal. 1998.
- **Non-photorealistic Rendering SIGGRAPH Course notes**
D. Green, SIGGRAPH 1999
- **Processing Images and Video for Impressionist Effect**
P. Litwinowicz, SIGGRAPH 1997
- **Video Toning**
J. Wang , Y. Xu, H. Shum, M. Cohen, SIGGRAPH 2004

- **Painterly Rendering for Video and Interaction**
A. Hertzmann, K. Perlin. NPAR 2000.
- **Painterly Rendering for Animation**
B. Meier. SIGGRAPH 1996
- **Image Analogies**
A. Hertzmann, C. Jacobs, N. Oliver, B. Curless, D. Salesin. SIGGRAPH 2001
- **Directional Texture Transfer**
H. Lee, S. Seo, S. Ryoo, K. Yoon. NPAR 2010.
- **Empathic Painting: Iterative stylization using observed emotional state**
M. Shugrina, M. Betke, J. Collomosse. NPAR 2006.
- **Genetic Paint: A Search for Salient Paintings**
J. Collomosse, P. Hall. EvoMUSART 2005 (J. IJAIT 2006).
- **The Art of Scale Space**
J. A. Bangham, S. Gibson, R. Harvey. BMVC 2003.
- **Visual interest and NPR: An evaluation and manifesto**
A. Santella, D. DeCarlo. NPAR 2004.
- **Segmentation-based 3D Artistic Rendering**
A. Kolliopoulos, J. Wang, A. Hertzmann. EGSR 2006.
- **Stylized Video Cubes**
A. Klein, P. Sloan, A. Colburn, A. Finkelstein, M. Cohen. EG SCA 2002.
- **Image and Video based Painterly Animation**
J. Hayes and I. Essa, NPAR 2004.

- **Stroke Surfaces: Temporally Coherent Artistic Animations from Video**
J. Collomosse, D. Rowntree, P. Hall. IEEE TVCG 2005.
- **Video Watercolorization using Bidirectional Texture Advection**
A. Bousseau, D. Neyret, J. Thollot, D. Salesin
- **Video Analysis for Cartoon-like Special Effects**
J. Collomosse, D. Rowntree, P. Hall. BMVC 2003.
- **Video Analysis for Dynamic cues and Futurist Art**
J. Collomosse, P. Hall. Graphical Models. 2006.
- **Motion Magnification**
C. Liu, A. Torralba, W. Freeman, F. Durand, E. Adelson. SIGGRAPH 2005
- **Video SnapCut: Robust Video Object Cutout Using Localized Classifiers**
X. Bai, J. Wang, D. Simons, G. Sapiro. SIGGRAPH 2009
- **Stylized Displays of Home Image and Video Collections**
T. Wang, R. Hu, J. Collomosse, D. Slatter, P. Cheatle, D. Greig. NPAR 2010 (CAG 2011)
- **Painterly animation using video semantics and feature correspondence**
L. Liang, K. Zeng, H. Lv, Y. Wang, Q. Xu, S. Zhu. NPAR 2010
- **From Image Parsing to Painterly Rendering**
K. Zeng, M. Zhao, C. Xiong, S. Zhu. ACM ToG 2010.



- **Artistic Stylization pre-2000**
 - Dependent on low-level image processing (e.g. Sobel) to drive preservation of local edge and high frequency content.
- An Artist does not paint a stroke by looking only at the image content under that stroke
- A higher level of visual analysis is needed:

Consider more than local edge information

Global analysis vs. greedy placement


- Computer Vision and Optimisation are solutions

Region-based discrimination



Around the Cake (Thiebaud'62). Markup (Kolloiopoulos '06)

- **Segmentation (EDISON / Mean-Shift)** [Christoudias et al, ICPR 2002]
 - Create a spatial hierarchy of regions
 - Strokes painted in a region have same prominence
 - Or render regions flat with black edges to create 'toon effect
 - Determine prominence of regions interactively
 - ...using an eye tracker



- Segment levels of low-pass (Gaussian) pyramid
 - DeCarlo uses factor of $\sqrt{2}\sigma$ between layers
 - Discard regions < 500 pixels (on 640x480 image)
- Segments grouped into hierarchy from fine to coarse based on overlap and common colour


1. For each* region A at the current level e.g. L₁
2. Find the region B_i in level above e.g. L₂ maximising:


$$\text{overlap}(A, B_i) = \frac{\text{area}(A \cap B_i)}{\|\text{color}(A) - \text{color}(B_i)\| + 1}$$

3. Assign A's parent to B_i, providing $A \cap B_i$ is contiguous+

*At step 1, iterate through regions in order of increasing area.

+ After all levels are processed, any orphan regions become children of root node.






Region detail
identified via fixation

Post-smoothing

- **Painting starts at the coarsest level of region detail**
- **A region is split if more than half its children are fixated upon**
- **The resulting region map is noisy, but aesthetics improve after smoothing and vectorisation**

- Alternative scale-space hierarchy using sieves
 - Morphological operations (closure followed by opening)

```
X=imerode (imdilate (X, ones (1, N)) , ones (1, N)) ;
```



- **Sieves better preserve edges/corners vs. Gaussian**
 - Extended to 2D in [Bangham '99], NPR application [Bangham '03]. Colour sieves (Harvey '04)




- Similar level of detail strategy to Decarlo/Santella can be applied to scale-space tree

■ Painting the regions



Paint via 3rd party algorithm e.g. Hertzmann with constant stroke size [Santella /DeCarlo NPAR'02]





Fill region with strokes in direction of principal axis [Shugrina et al, NPAR '06]




Fill with strokes in directions derived from region exterior contour [Wang et al, NPAR '10]

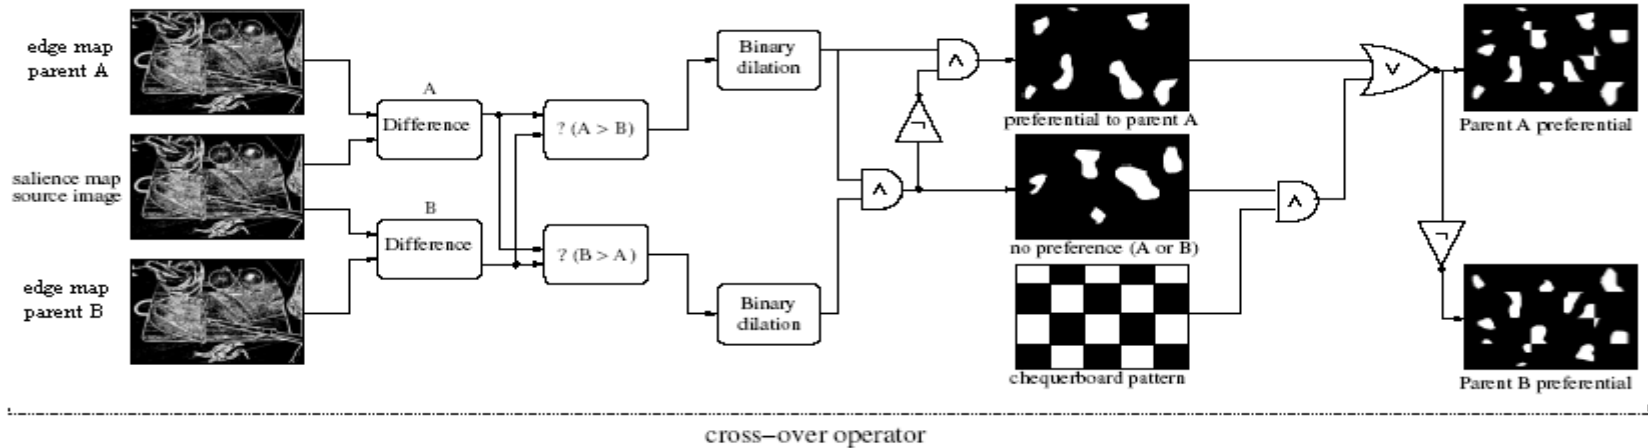
- **Automated Differential Emphasis in Painting**
 - **Prescriptive salience measures [Itti & Koch]**
 - **Not closely correlated to human behaviour [Santella/DeCarlo NPAR'04]**
 - **Salience is subjective and task dependent**
 - **Trainable measure of salience (GMM of radial features)**




- Genetic Optimizaton to find “best” painting
 - The optimal painting preserves detail in salient areas, and removes non-salient detail
 - MSE between salience map and Sobel edge detail in the painting (c.f. Hertzmann '01)




- Paintings are bred by cloning strokes from two individuals
- (Two parent cross-over)
 - fitness proportionate selection with replacement
- Promotion of rapid convergence
 - Top 10% carried over to next gen. automatically
 - Bottom 10% culled



- **Iterative optimization improves detail in salient regions**
 - Population of ~50 paintings
 - Convergence in ~200 iterations
 - Stochastic variation in stroke attributes creates diversity
 - GA combines favourable regions of parent paintings





- **Comparison of Sobel-driven and Saliency-driven painting**
 - **Detail on the sign is preferentially retained (wrt. Leaves of the tree)**
 - **Not all edges / high frequency texture are salient**



Original



Litwinowicz '97



Saliency driven

- **Painting research code available**
 - <http://www.collomosse.com/EG2011tut/summerschool.zip>

- **MATLAB based (experiment with different salience maps)**

- **Code adapted from Collomosse et al. 2005 – single iteration, spline strokes.**

- **Previously released as lab exercise at EPSRC VVG Summer School (2007)**





Style Transfer

- Learning vs Heuristic approach to stylise photos
- Patch based lookup (luminance only)

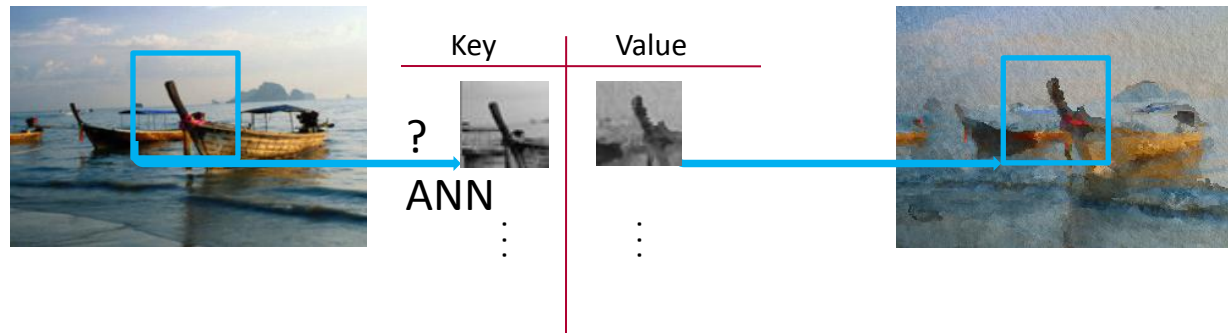
$$Y(p) \leftarrow \frac{\sigma_B}{\sigma_A}(Y(p) - \mu_A) + \mu_B$$

- Similar to Freeman texture synthesis but using external collection of patches
- Learned as lookup table



■ Style Transfer

- Synthesis has 'data' and 'smoothness' terms
 - Data (patch lookup)
 - Pixel-wise luminance comparison (after PCA)
 - Smoothness (derived from Ashikhmin)
 - Minimise MSE between proposed patch and existing neighbours
 - Gaussian weighted distance function (avoids discontinuity)



Style Transfer Examples



Learn



Apply



Learn



Apply



- Other extensions to video [Hors & Essa '02] and to take orientation into account [Lee et al. '10]

- **Video Stylisation**
 - Techniques to create painterly animations or cartoons from video
 - Enabled by automated techniques for image stylization



Stylised Appearance




Stylised Motion

- **Goal of video stylization**
 - Create the desired aesthetic exhibiting good temporal coherence
- **Temporal coherence is here defined as:**
 1. **Absence of distracting flicker**
 2. **Motion of brush strokes (or other component marks) is in agreement with the motion of content**
- **Naïve approaches**
 - Repaint every frame independently
= *Flicker (violates 1.)*
 - Fix strokes in place and change attributes e.g. colour according to video content
= *Motion unmatched (violates 2.)*
“the shower door effect” – Barb Meier



- **Painterly animation using Optical Flow**
 - **Brush strokes are pushed from frame to frame using flow estimate**
 - **Oscar winning visual effects in movie “What Dreams May Come” (1998)**
 - **Manual correction of flow estimate (~1000 person-hours [Green’99])**





No clipping



Clipping




- **Initialisation as per single image (regular seeding)**
 - Randomise rendering order of strokes
- **Strokes translated to next frame via flow field**
- **Greedy approximation to avoid irregular coverage**
 - Delaunay triangulation of seeds (and image corners)
 - **Death.** Seeds too close together are deleted
 - Tested in random order
 - **Birth.** Triangles with area $>$ threshold are subdivided
 - New seeds are randomly place rendering order







▪ Stroke Birth



■ Tips on reducing flicker

- Detect your own scene cuts and reinitialise
- Use a robust Optical Flow algorithm (!)
 - e.g. SIFTFlow or Brox
- Pre-filter heavily (Gaussian). Care with interlaced content.
- Interpolate orientations from strong edges only
 - Smooths out codec noise
 - Litwinowicz uses thin-plate spline (expensive) but can use Poisson filling (fast on GPU) to good effect



Sobel field



Without interpolation



With interpolation



- **Main sources of temporal incoherence**
- **Motion matching**
 - **Optical Flow = visual correspondence problem**
 - **Inevitable inaccuracies in estimate are cumulative**
 - **Content appears to slip below strokes = shower door effect**
 - **Manual correction of OF mitigates this but is expensive**
- **Flicker**
 - **Random order of new strokes disguises regularity**
 - **...but the noise generates flicker**
 - **Sudden disappearance of strokes exposes others = popping**
 - **Sobel edges are noisy at moderate scales**
 - **Strokes are clipped against flicking edge map**



- **Main sources of temporal incoherence**
- **Motion matching**
 - **Optical Flow = visual correspondence problem**
 - **Inevitable inaccuracies in estimate are cumulative**
 - **Content appears to slip below strokes = shower door effect**
 - **Manual correction of OF mitigates this but is expensive**
- **Flicker**
 - **Random order of new strokes disguises regularity**
 - **...but the noise generates flicker**
 - **Sudden disappearance of strokes exposes others = popping**
 - **Sobel edges are noisy at moderate scales**
 - **Strokes are clipped against flicking edge map**

Addressed by

Hertzmann and Perlin
NPAR 2000

Hays and Essa
NPAR 2004

- **Repaint only the areas that change significantly**
 - **Fast** – enables realtime interaction
 - **Limits shower-door** by repainting limited regions of canvas (“paint-over”)
- **RGB Difference to detect regions**
 - **Optical flow *optionally* used to translate strokes**

$$\frac{1}{|M|} \sum_{(i,j) \in M} \|I_{t+1}(i,j) - I_t(i,j)\| > T_V$$

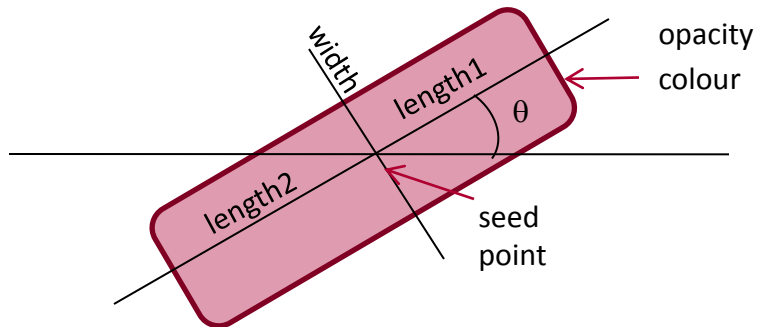
- **Control points shifted under flow**







- **Key Innovation**
 - Temporal smoothing of stroke attributes
- Stroke Opacity for birth/death
- Orientation
 - RBF interpolated field (similar to Litwinowicz)
 - But interpolated from strokes marked “strong”
not from per-frame orientation field
 - Strokes born on strong edges
- Length and orientation are also smoothed








- **Bi-Directional Flow (of textures vs. strokes)**
 - **Adaptation of “Texture advection” from flow visualisation**
 - **Frequent occlusions in video motivated bi-directional flow**
 - **Two textures seeded – one flows forward, one back**



- **Trend towards more global temporal analysis**

- **Temporally local (inter-frame) approaches**
 - No long-view of video structure
 - necessitates averaging of past information
 - Averaging mitigates flicker but exaggerates the shower door effect
- **Spatio-temporal primitives**



timeslice

- **Automated Space-time Analysis**
 - **Goal is coherent segmentation of video into semantic regions**
 - **Coherent space-time regions are smoothed then do not flicker**




Extending Mean-shift to Space-time (3D)




EDISON segment frames and associate (2D+t)

▪ Region association (2D+t)

- Based on a weighted blend of heuristics
 - Shape (Fourier Descriptors)
 - Colour
 - Overlap (as DeCarlo/Santella)
- Associations are filtered by locating
 - Short-time branches
 - Short-time cycles
- Surface voxels between volumes are identified
- Surfaces fragmented into “stroke surfaces” that abut only two volumes




- “Stroke Surfaces” separate volumes
 - Winged edge structure
 - Smoothing the surface smoothes the volumes
 - Generalisation of snakes to 2D surfaces
 - Separate terms for spatial and temporal constraint



$$E = \int_0^1 \int_0^1 (E_{int}[\underline{Q}(s, t)] + E_{ext}[\underline{Q}(s, t)]) ds dt.$$



$$E_{int} = \alpha \left| \frac{\partial \underline{Q}(s, t)}{\partial s} \right|^2 + \beta \left| \frac{\partial \underline{Q}(s, t)}{\partial t} \right|^2 + \gamma \left| \frac{\partial^2 \underline{Q}(s, t)}{\partial s^2} \right|^2 + \delta \left| \frac{\partial^2 \underline{Q}(s, t)}{\partial t^2} \right|^2$$

$$E_{ext} = \eta f(\underline{Q}(s, t)).$$






■ Surface Manipulation

■ Undulation



■ Fragmentation






- **Coherent Segmentation**

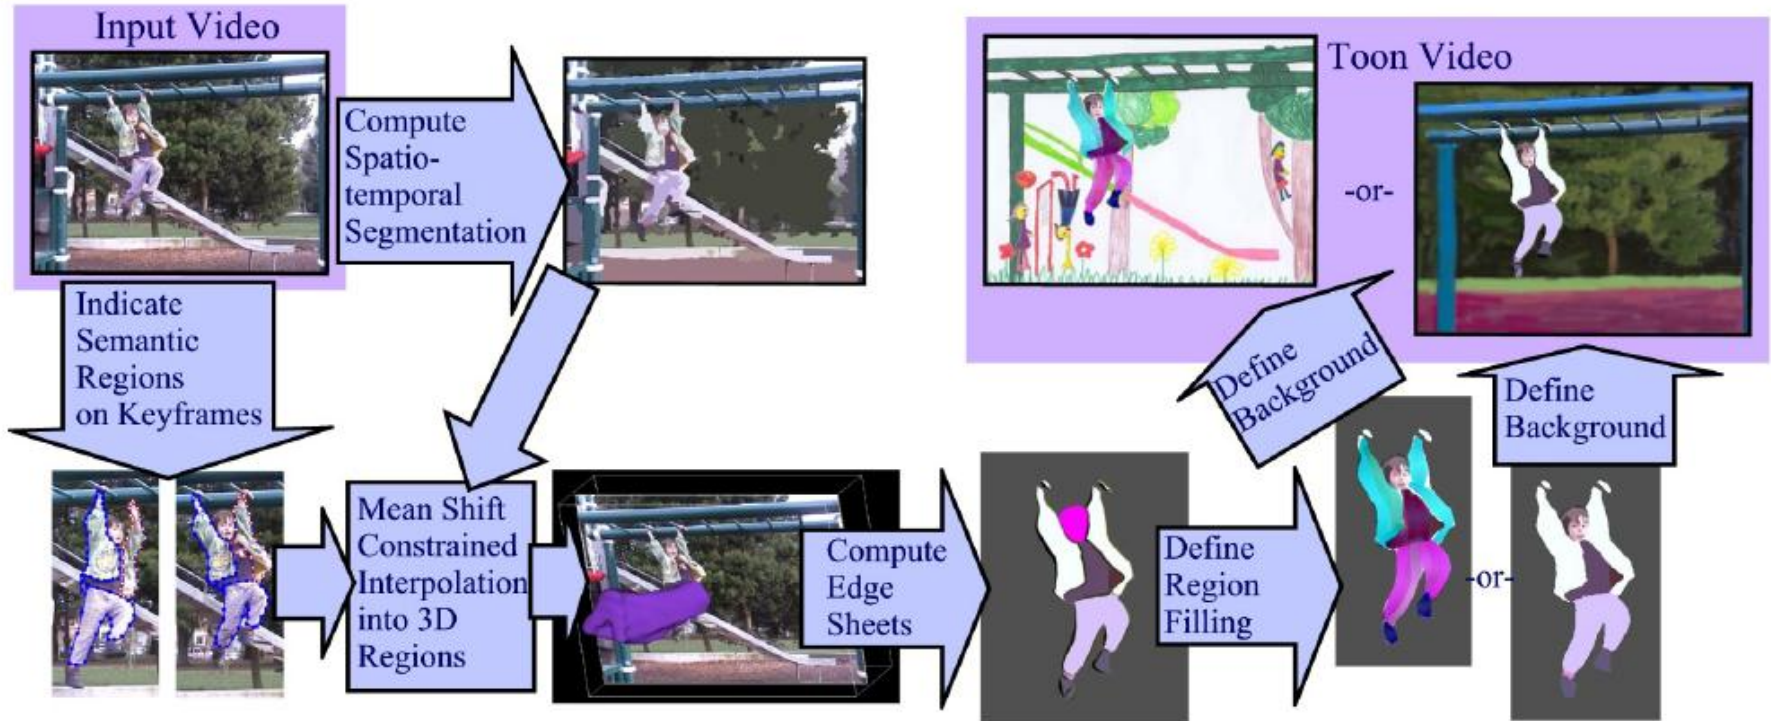


Rotoscoping

- Coherent motion of (groups of) regions can be exploited to paint coherently
- Interpolate internal points (e.g. stroke seeds) from region boundary



■ Rotoscoping



- **Rotoscoping**

- **Coherent motion of (groups of) regions can be exploited to paint coherently**



■ Motion Emphasis

- Augmentation cue (Speed-lines, ghosting)
- Deformation (Squash and stretch, general deformation)
- Time warping (Anticipation/snap)




Augmentation Cues

- Segment trails of corresponded control points into smooth sections.
- Iteratively select smooth sections to maximising:

$$H(0) = 0$$

$$H(i+1) = H(i) + (\alpha v(x) + \beta L(x) - \gamma D(x) - \delta \omega(x, \sigma; w) + \zeta \rho(x))$$

$\alpha, \beta, \gamma, \delta, \zeta$: user weights
 σ : selected sections
 $v(x)$: mean velocity over x
 $\rho(x)$: mean curvature at x
 $L(x)$: temporal extent of x
 $D(x)$: dist from convex hull at x
 $\omega(x, \sigma; w)$: overlap of x and σ



■ Deformation Cues

- A motion dependent curvilinear basis is formed using the trajectory of the region centroid, and its normal.

$$\underline{x} = \underline{G}_c(s) + r\underline{n}(s)$$

$U(\cdot)$ as the transformation from curvilinear space, keep the inverse as a lookup table.


$$\underline{x} = U(\underline{r}) \quad \underline{r} = (s, r)^T$$

$$\underline{x} \leftarrow U\left(\begin{bmatrix} k & \mathbf{0} \\ \mathbf{0} & \frac{1}{k} \end{bmatrix} U^{-1}(\underline{x})\right)$$

$$k = 1 + \frac{K}{2} \left(1 - \cos\left(\pi \frac{v^2 + 1}{2}\right)\right)$$



$$v = \begin{cases} \mathbf{0} & \text{if } |\dot{\underline{u}}| < V_{min} \\ 1 & \text{if } |\dot{\underline{u}}| \geq V_{max} \\ (|\dot{\underline{u}}| - V_{min}) / (V_{max} - V_{min}) & \text{otherwise} \end{cases}$$

Squash and stretch
(after Cheney et al '02)



■ Deformation Cues

- Squash and stretch in a camera motion compensated frame



■ Deformation Cues

- More general motion deformations can be created by specifying a transfer function dependent on a point's local acceleration and position, as well as its speed.

$$\underline{x}' = U(T(U^{-1}(\underline{x}), \underline{\dot{x}}, \underline{\ddot{x}}))$$

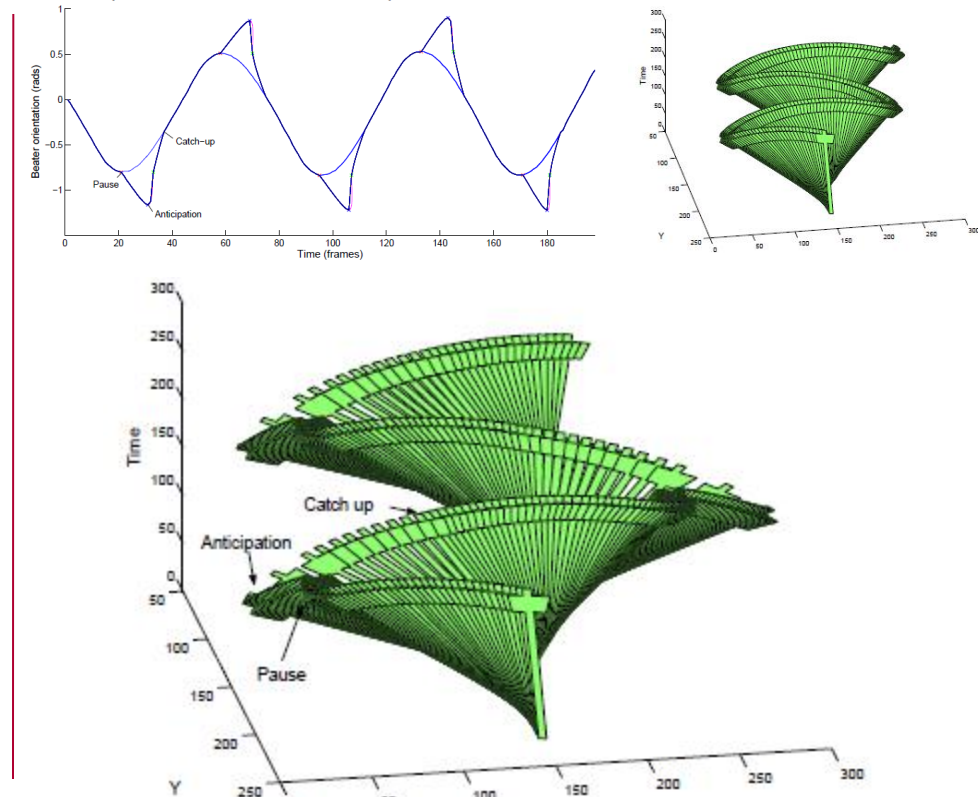

A function can operate on each component of $\underline{r} = (r_1, r_2)$ independently, to create effects suggesting drag we use...

$$r_1 = r_1 - F\left(\frac{2}{\pi} \text{atan}(|\dot{x}_i|)\right)^P \text{sign}(\dot{x}_i)$$



Anticipation (Snap)

- Alter motion timing to introduce a lag then “catch up” prior to changes of motion





■ A Complete Video Paintbox




Segmentation + augmentation + deformation




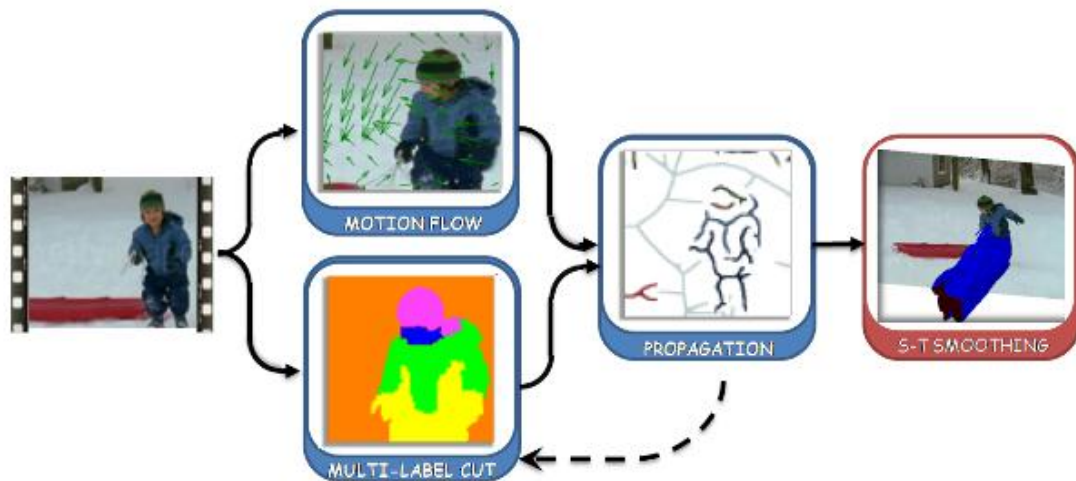
Anticipation + deformation

■ Deformation Cues

- General deformation technique using motion vector clustering to layer video
 - User intervention needed to fix noisy segmentation maps
- Per-pixel flow vector pushes pixels to exaggerate existing motion
- Texture filling compensates for holes



- Combining Segmentation and Optical Flow
 - Multi-label graph-cut segmentation with label prior propagated forward from previous frames
 - Region colour models are learned incrementally



propagation to next...

Combining Segmentation and Optical Flow

- For each pixel $p \in \mathcal{P}$ within frame $I_t(p)$
 - Find best mapping $l: \mathcal{P} \rightarrow \mathcal{L}$
 $\mathcal{L} = (l(1), \dots, l(p), \dots, l(|\mathcal{P}|))$
 - Subset of \mathcal{L} are carried from $t-1$ by flow

$$E(\mathcal{L}, \Theta, \mathcal{P}) = U(\mathcal{L}, \Theta, \mathcal{P}) + V(\mathcal{L}, \mathcal{P}).$$

$$U(\mathcal{L}, \Theta, \mathcal{P}) = \sum_{p \in \mathcal{P}} -\log P_g(I_t(p) | l(p); \Theta).$$

$$P_g(I(p) | l(p) = l_i; \Theta) = \sum_{k=1}^{n_i} w_{ik} \mathcal{N}(I(p); \mu_{ik}, \Sigma_{ik})$$

$$V(\mathcal{L}, \mathcal{P}) = \gamma \sum_{(m,n) \in \mathcal{N}} [l(m) \neq l(n)] e^{-\beta \|I(m) - I(n)\|^2}.$$



Learned colour model

Contrast adaptive



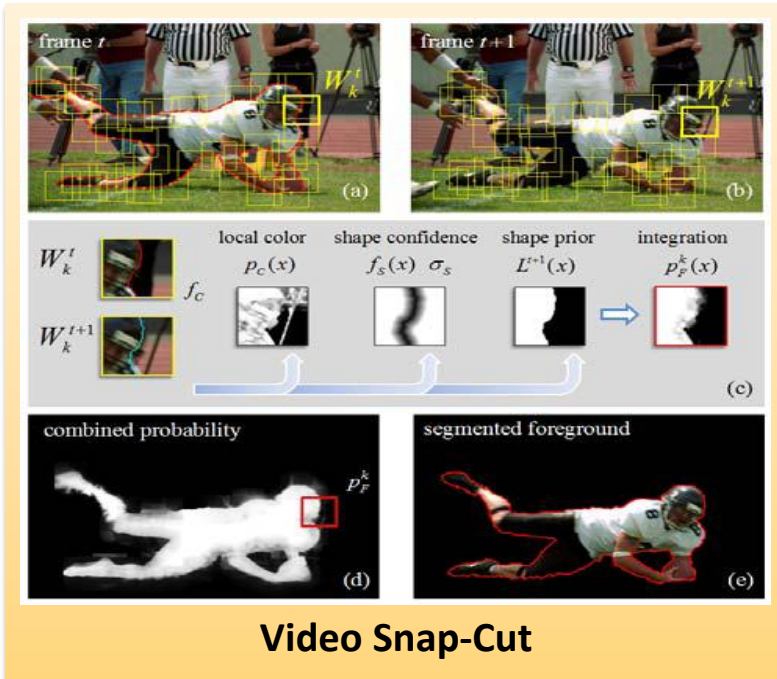
- Colour models are learned over time incrementally via Gaussian Mixtures

- Region colour distribution (GMM) updated with temporal weight
 - Comparison with historic model (χ^2) can detect region birth




Wang et al. '10

- Trends from automatic 'Tooning to interactive tools
 - The necessity of interaction to solve the general segmentation problem



Bai et al. '09





Liang et al. '10

Yet many applications demand automation or real-time. Part III discusses solutions.



Coffee Q & A

- **After the break!**
 - **Part III - Anisotropy and Diffusion**
 - **Part IV - Future Challenges in NPR**



Artistic Stylization of Images and Video

Part III – Anisotropy and Filtering

Eurographics 2011

Jan Eric Kyprianidis

Hasso-Plattner-Institut, University of Potsdam, Germany

- **Stylized Augmented Reality for Improved Immersion**
Fischer et al., 2005
- **Real-time Video Abstraction**
Winnemöller et al., SIGGRAPH 2006
- **Coherent Line Drawings**
Kang et al., NPAR 2007
- **Structure Adaptive Image Abstraction**
Kyprianidis & Döllner, EG Theory and Practice of Computer Graphics 2008
- **Flow-based Image Abstraction**
Kang et al., Transactions on Visualization and Computer Graphics 2009
- **Artistic Edge and Corner Preserving Smoothing**
Papari et al., IEEE Transactions on Image Processing 2007
- **Image and Video Abstraction by Anisotropic Kuwahara Filtering**
Kyprianidis et al., Pacific Graphics 2009
- **Shape-simplifying Image Abstraction**
Kang & Lee, Pacific Graphics 2008



Non-photorealistic display of both the camera image and virtual objects:

- **Abstraction:** Bilateral filter applied to Gaussian pyramid and then upsampled
- **Edges:** Canny edge detector + morphological dilation



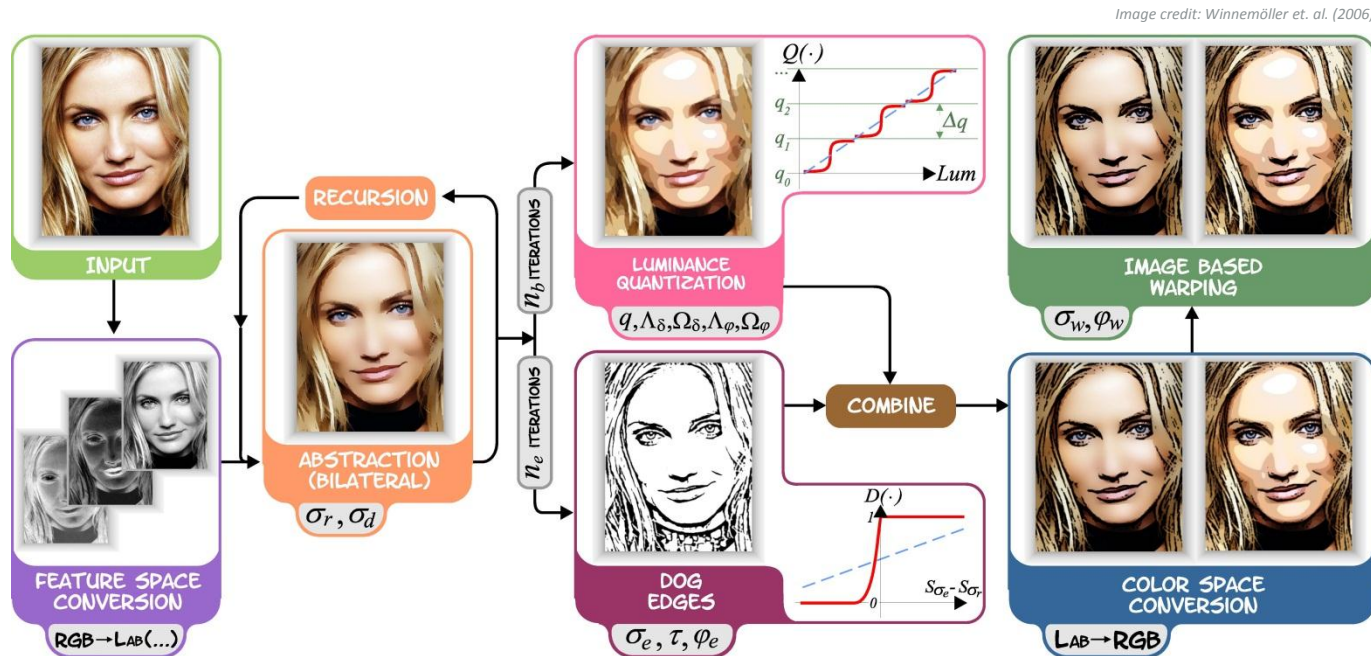
Conventional augmented reality



Stylized augmented reality

Image credit: Fischer et. al. (2005)

- **Abstraction:** Multiple iterations of xy-separable bilateral filter + color quantization
- **Edges:** Difference of Gaussians + thresholding




- **Edges:** 1D difference of Gaussians directed by flow field + flow-guided smoothing and thresholding.



Input image



Edge Tangent Flow



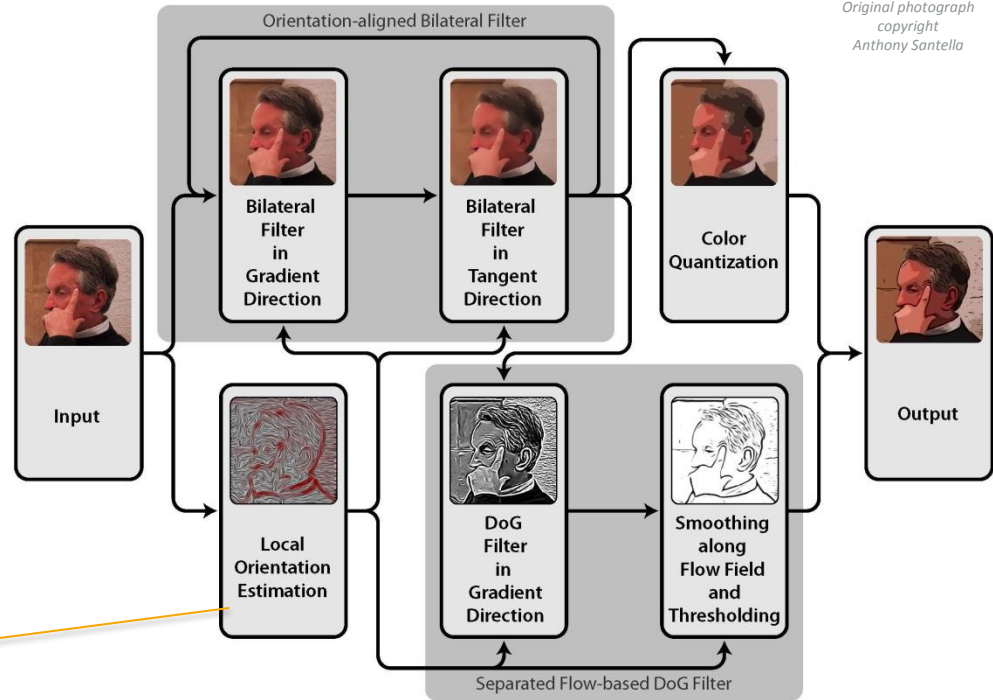
Line drawing

Image credit: Kang et. al. (2007)



Flow-based filtering

- **Abstraction:** Multiple iterations of orientation-aligned bilateral filter
- **Edges:** separable flow-based difference of Gaussians



Local orientation and an anisotropy measure derived from the smoothed structure tensor are used to guide the bilateral and difference of Gaussians filters

- **Abstraction:** Multiple iterations of flow-based bilateral filter
- **Edges:** (separable) flow-based difference of Gaussians
- Local orientation estimation of both techniques is based on the edge tangent flow (ETF)




Image credit: Kang et. al. (2009)


- Generalization of the Kuwahara filter. Creates output with a painterly look.
- Addresses two key issues of the original Kuwahara filter:
 - Rectangular subregions
 - Unstable subregion selection process



Credit for images: Papari et al. (2009)



- Further generalization of the Kuwahara filter.
- Adaptation of the shape, scale and orientation of the filter to the local image structure.



Original image by Paulo Brandão@flickr.com



- PDE-based technique that simultaneously simplifies colors and shape:
 - Constrained mean curvature flow
 - Shock filter

Image credit: Kang & Lee (2008) / original image by Tambako the Jaguar@flickr.com



Input



20 iterations



40 iterations



60 iterations

Difference of Gaussians:

- Laplacian of Gaussian (LoG)
- Isotropic Difference of Gaussians (DoG)
- Flow-based Difference of Gaussians
- Separable Flow-based Difference of Gaussians

DoG Edges vs Canny Edges

Original image from USC-SIPI Image Database




Canny Edges




Flow-based difference of Gaussians

Edge profile without noise:



Edge profile with noise:



In 2D the second derivative corresponds to the Laplacian:

$$L = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}$$



Similar to the second derivative the Laplacian is sensitive to noise. To make the Laplacian less sensitive to noise, apply a Gaussian to the image first:

$$\text{LoG} = L \star G_\sigma ,$$

where G_σ is a 2D Gaussian with standard deviation σ :

$$G_\sigma(x, y) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right)$$

A Laplacian of Gaussian can be approximated by a difference of Gaussians:



$$\text{LoG}(x, y) = L \star G_{\sigma}$$

$$\text{DoG}(x, y) = G_{\sigma_i}(x, y) - G_{\sigma_e}(x, y)$$

Good engineering
solution:
 $\sigma_i = 1.6 \cdot \sigma_e$

Fast to implement
since Gaussian
is separable



Zero-crossing are found by thresholding:



An approach to create smooth edges was proposed by Winnemöller et al.:

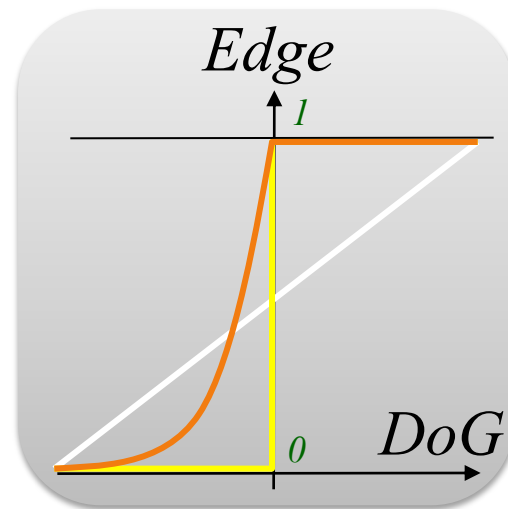
$$D(\sigma_e, \tau, \varphi_e) = \begin{cases} 1 & \text{if } (G_{\sigma_e} - \tau G_{1.6 \cdot \sigma_e}) > 0 \\ 1 + \tanh(\varphi_e \cdot G_{\sigma_e} - \tau G_{1.6 \cdot \sigma_e}) & \text{otherwise} \end{cases}$$

- The parameter τ controls the sensitivity to noise. A typical values are $\tau = 0.98$ or $\tau = 0.99$.
- The falloff parameter φ_e determines the sharpness of edge representations, typical values are $\varphi_e \in [0.75, 5.0]$.




Credit for slide: H. Winnemöller

\tanh
(half-truncated)



Edge Tangent Flow (ETF):

- Smoothly varying vector field
- Feature-preserving flow



Input image



Edge Tangent Flow

Image credit: Kang et al. (2007)

Weighted vector smoothing similar to bilateral filter:

Multiple iterations (≈ 3)

$$t^{n+1}(x) = \frac{1}{k} \sum_{y \in \Omega(x)} \phi(x, y) \cdot t^n(y) \cdot w_s(x, y) \cdot w_m(x, y) \cdot w_d(x, y)$$

t^0 is calculated using the Sobel filter.




Image credit: Kang et al. (2007)

$$t^{n+1}(x) = \frac{1}{k} \sum_{y \in \Omega(x)} \phi(x, y) \cdot t^n(y) \cdot w_s(x, y) \cdot w_m(x, y) \cdot w_d(x, y)$$

Assure different vectors point in the same direction

$$\phi(x, y) = \text{sign}(t^n(x) \cdot t^n(y))$$

Restrict filtering to a predefined radius

$$w_s(x, y) = \begin{cases} 1 & |x - y| < r \\ 0 & \text{else} \end{cases}$$

More weight to vectors with higher gradient magnitude

$$w_m(x, y) = \frac{1}{2} [1 + \tanh(|g(x)| - |g(y)|)]$$

More weight for vectors with direction similar to current filter origin

$$w_d(x, y) = |t^n(x) \cdot t^n(y)|$$



Flow-based Difference of Gaussians

Kang et al. (2007)

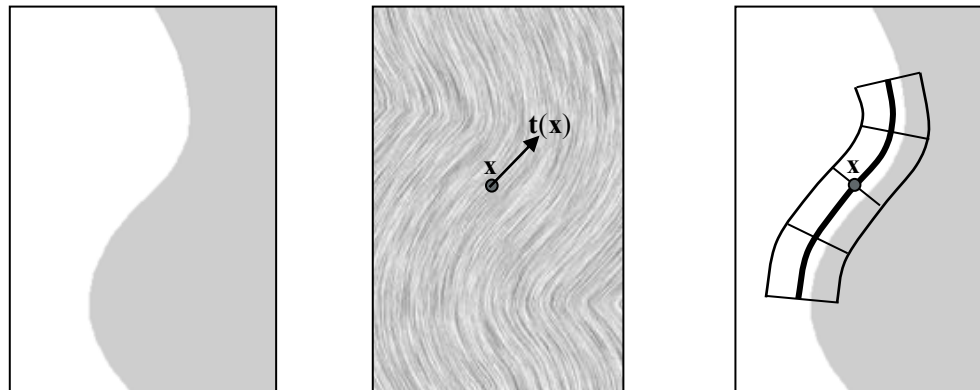

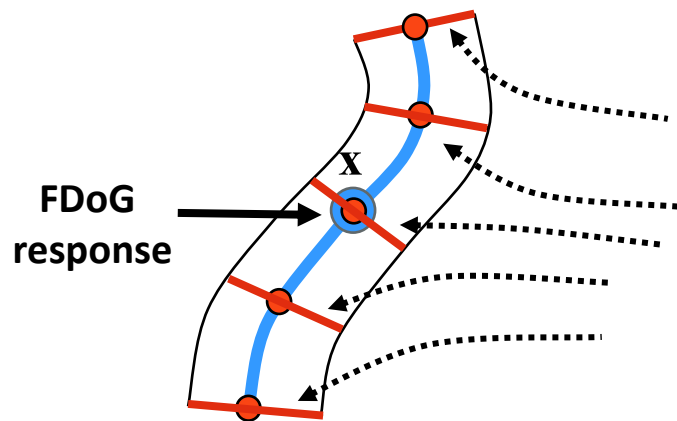


Image credit: Kang et al. (2007)



DoG filter

Let $f: \mathbb{R}^2 \rightarrow \mathbb{R}^3$ denote the input image and let

$$\frac{\partial f}{\partial x} = \begin{pmatrix} \frac{\partial R}{\partial x} & \frac{\partial G}{\partial x} & \frac{\partial B}{\partial x} \end{pmatrix}^t \quad \frac{\partial f}{\partial y} = \begin{pmatrix} \frac{\partial R}{\partial y} & \frac{\partial G}{\partial y} & \frac{\partial B}{\partial y} \end{pmatrix}^t$$

be the partial derivatives of f .

The structure tensor is then defined by:

$$(g_{ij}) = J^t J = \begin{pmatrix} \left\langle \frac{\partial f}{\partial x}, \frac{\partial f}{\partial x} \right\rangle & \left\langle \frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right\rangle \\ \left\langle \frac{\partial f}{\partial x}, \frac{\partial f}{\partial y} \right\rangle & \left\langle \frac{\partial f}{\partial y}, \frac{\partial f}{\partial y} \right\rangle \end{pmatrix} =: \begin{pmatrix} E & F \\ F & G \end{pmatrix}$$

These can be implemented for example using Gaussian derivatives or the Sobel filter.

The structure tensor is a 2×2 symmetric positive semidefinite matrix

In differential geometry the structure tensor is also known as first fundamental form

The induced quadratic form of the structure tensor measures the squared rate of change of f in direction of a vector $n = (n_x, n_y)$:

$$S(n) = En_x^2 + 2Fn_xn_y + Gn_y^2$$

The extremal values of $S(n)$ on the unit circle correspond to the eigenvalues of (g_{ij}) :

$$\lambda_{1,2} = \frac{E + G \pm \sqrt{(E - G)^2 + 4F^2}}{2}$$



The corresponding eigenvectors are:

Eigenvector of major eigenvalue.
Direction of maximum change:
gradient direction.

$$v_1 = \begin{pmatrix} F \\ \lambda_1 - E \end{pmatrix} \quad v_2 = \begin{pmatrix} \lambda_1 - E \\ -F \end{pmatrix}$$


Eigenvector of minor eigenvalue.
Direction of minimum change:
tangent direction.

The eigenvectors corresponding to the minor eigenvalues of the structure define a vector field. Typically this field is not smooth:





Smoothing the structure tensor prior to eigenanalysis with a Gaussian filter removes discontinuities in the vector field:






Eigenvector field of the smoothed structure tensor is similar to the edge tangent flow, but allows a more efficient implementation:



3 iterations of edge tangent flow filter

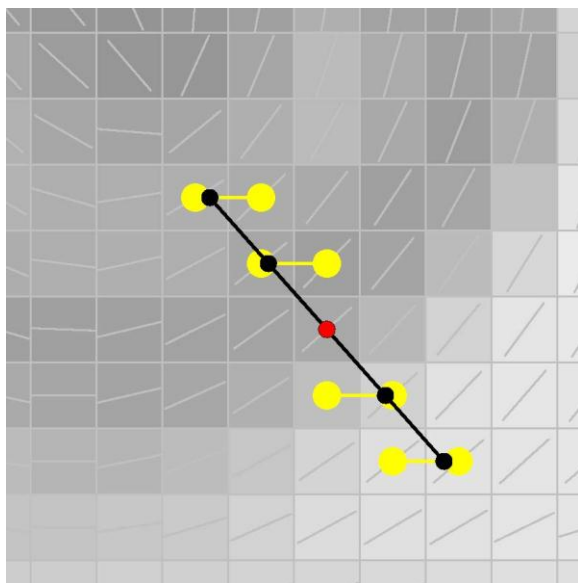


Eigenvector field of the smoothed structure tensor




Split flow-based difference of Gaussians into two passes:

- 1st Pass: one-dimensional DoG in direction of the major eigenvector
- 2nd Pass: smoothing along stream lines defined by minor eigenvector



1st Pass



2nd Pass



Bilateral Filter:

- Classical Bilateral Filter
- xy-Separable Bilateral Filter
- Orientation-aligned Bilateral Filter
- Flow-based Bilateral Filter

The bilateral filter is a nonlinear operation that smooths images while preserving edges:

More weight to closer pixel

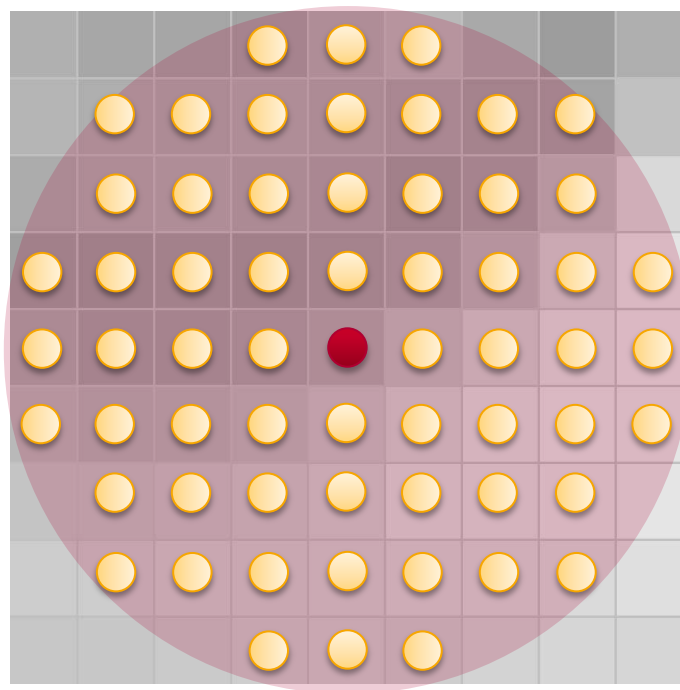
More weight to pixel with similar color

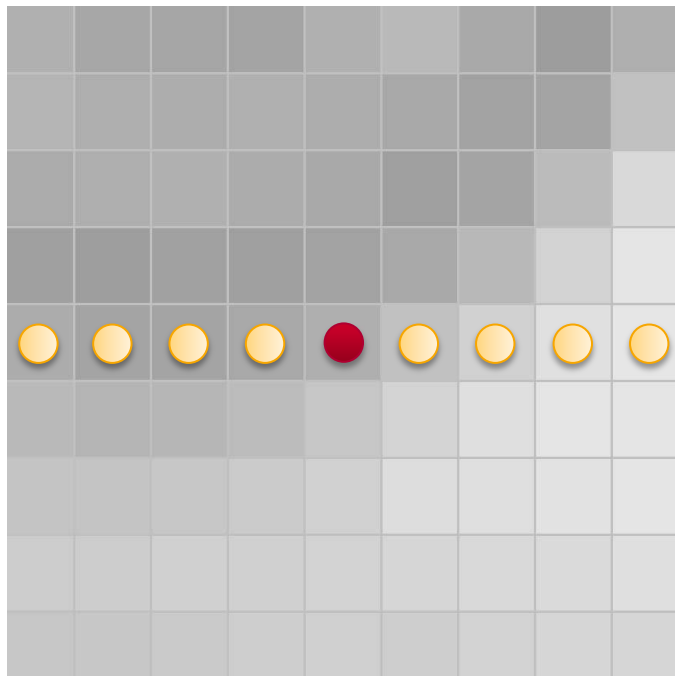
$$\frac{\sum_{x \in B_r(x_0)} f(x) G_{\sigma_d}(|x - x_0|) G_{\sigma_r}(|f(x) - f(x_0)|)}{\sum_{x \in B_r(x_0)} G_{\sigma_d}(|x - x_0|) G_{\sigma_r}(|f(x) - f(x_0)|)}$$

$$G_{\sigma}(t) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{t^2}{2\sigma^2}\right)$$

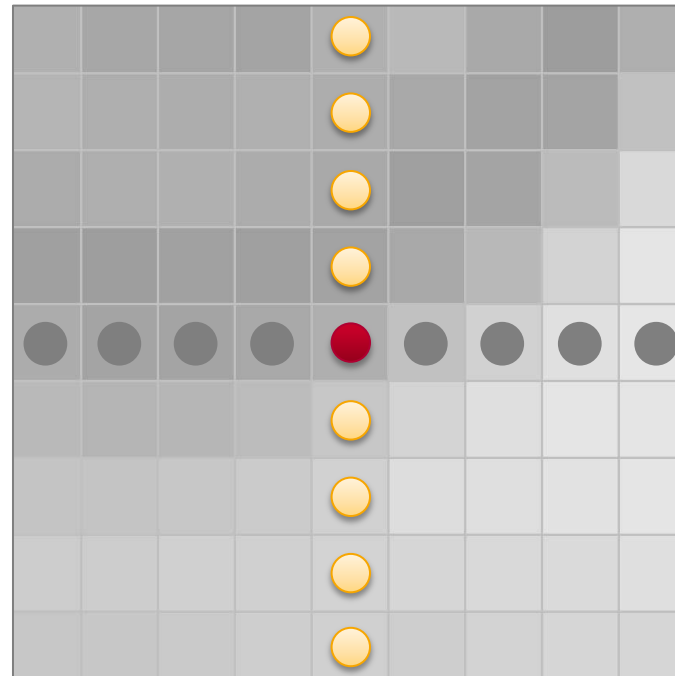


The bilateral filter is a powerful tool, but computationally very expensive ($O(r^2)$ per pixel).






1st Pass




2nd Pass



- Much faster than classical bilateral filter
- But creates noticeable artifacts!



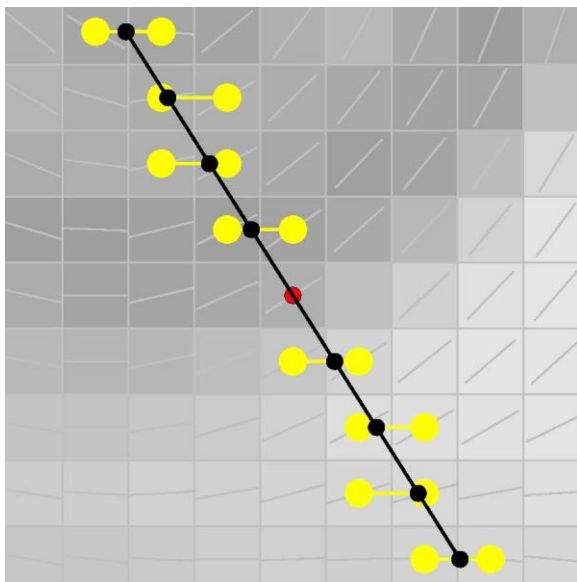
Full kernel bilateral filter




xy-separable bilateral filter



- Align separable bilateral filter to local orientation derived from the smoothed structure tensor




1st Pass




2nd Pass



- Less artifacts. Very well suited for abstraction.



Full kernel bilateral filter






Orientation-aligned bilateral filter




- Linear smoothing of neighboring pixel values creates smooth color boundaries

Unit step size
either along
x- or y-axis




- ETF-guided decomposition of bilateral filter



1st pass in direction of the tangent

2nd pass along stream line defined by the ETF






- Excellent preservation of highly anisotropic image features


Image credit: Kang et al. (2009)



Input image



Bilateral filter



Flow-based bilateral filter

Credit for slide: H. Winnemöller



Input

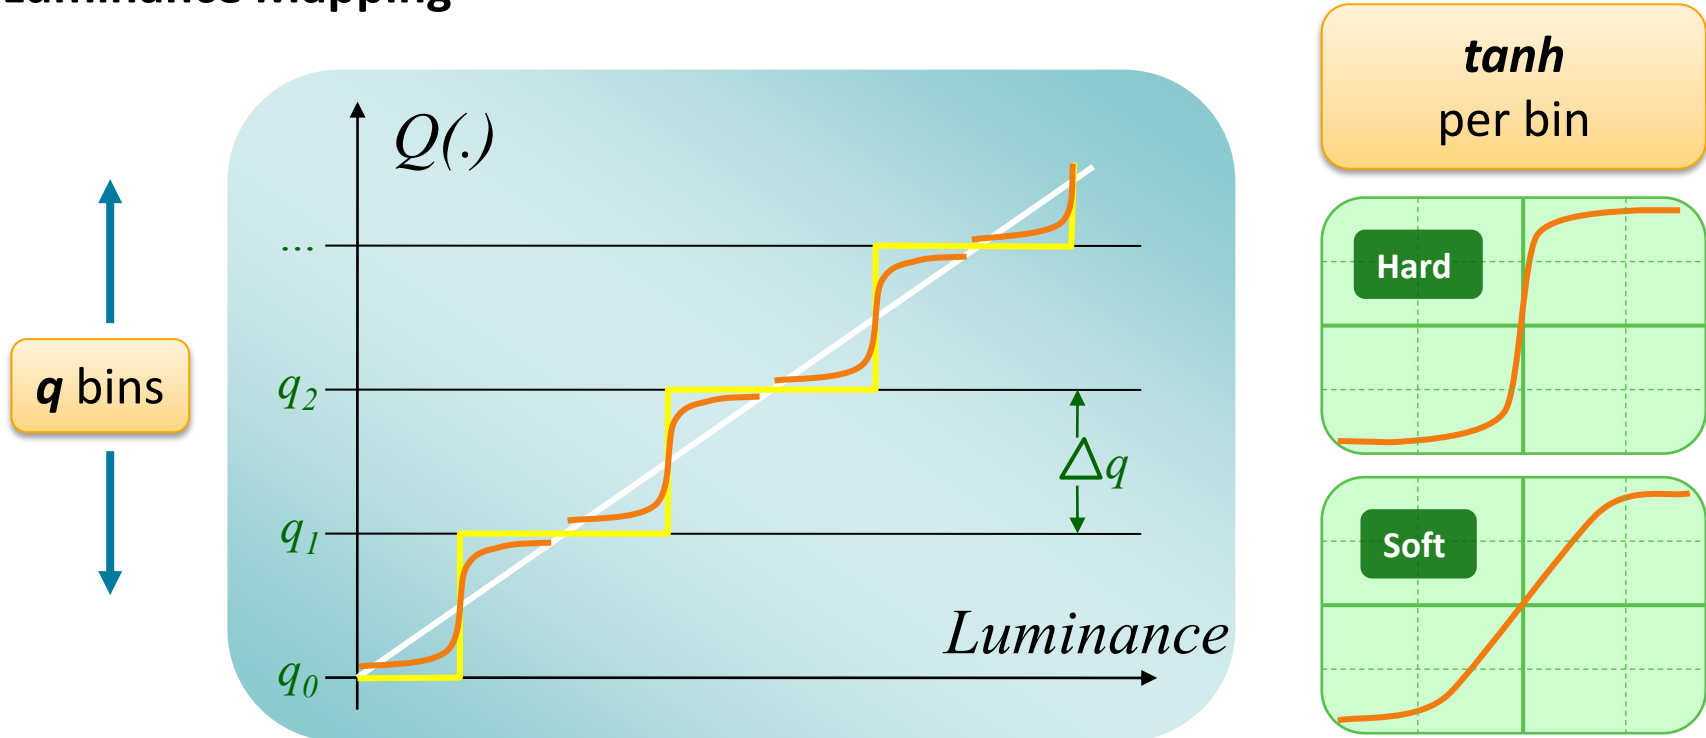


Apply quantization to
luminance channel



Result

Luminance Mapping



Credit for slide: H. Winnemöller




Image credit: Winnemöller et al. (2006)



Abstracted




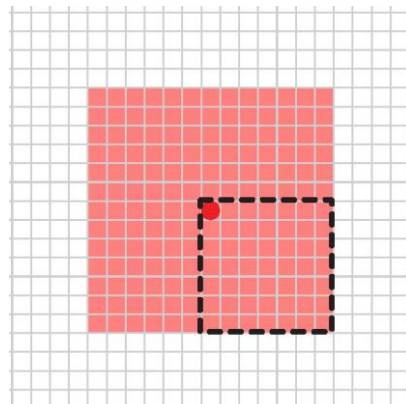


Sharp Quantization
(*Toon-like*)



Smooth Quantization
(*Paint-like*)

Kuwahara Filter:

- Classical Kuwahara Filter
- Kuwahara Filter with Weighting Functions
- Generalized Kuwahara Filter
- Anisotropic Kuwahara Filter
 - Convolution-based Weighting Functions
 - Polynomial Weighting Functions



$$W_0 = [x_0 - r, x_0] \times [y_0, y_0 + r]$$

$$W_1 = [x_0, x_0 + r] \times [y_0, y_0 + r]$$

$$W_2 = [x_0, x_0 + r] \times [y_0 - r, y_0]$$

$$W_3 = [x_0 - r, x_0] \times [y_0 - r, y_0]$$

For every subregion W_i calculate the mean

$$m_i = \frac{1}{|W_i|} \sum_{(x,y) \in W_i} I(x, y)$$

and the variance:

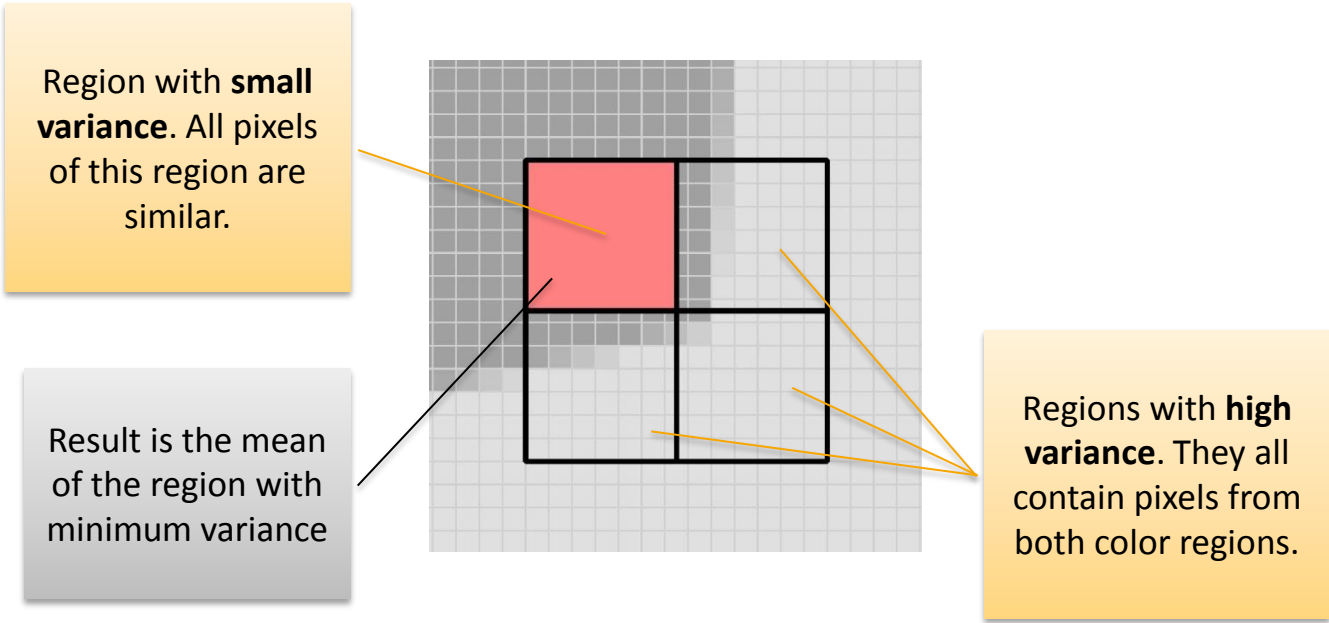
$$s_i^2 = \frac{1}{|W_i|} \sum_{(x,y) \in W_i} (I(x, y) - m_i)^2$$

The output of the Kuwahara filter is then defined as the mean of a subregion with minimum variance:


$$F(x_0, y_0) := m_k, \quad k = \underset{i=0, \dots, 3}{\operatorname{argmin}} s_i$$



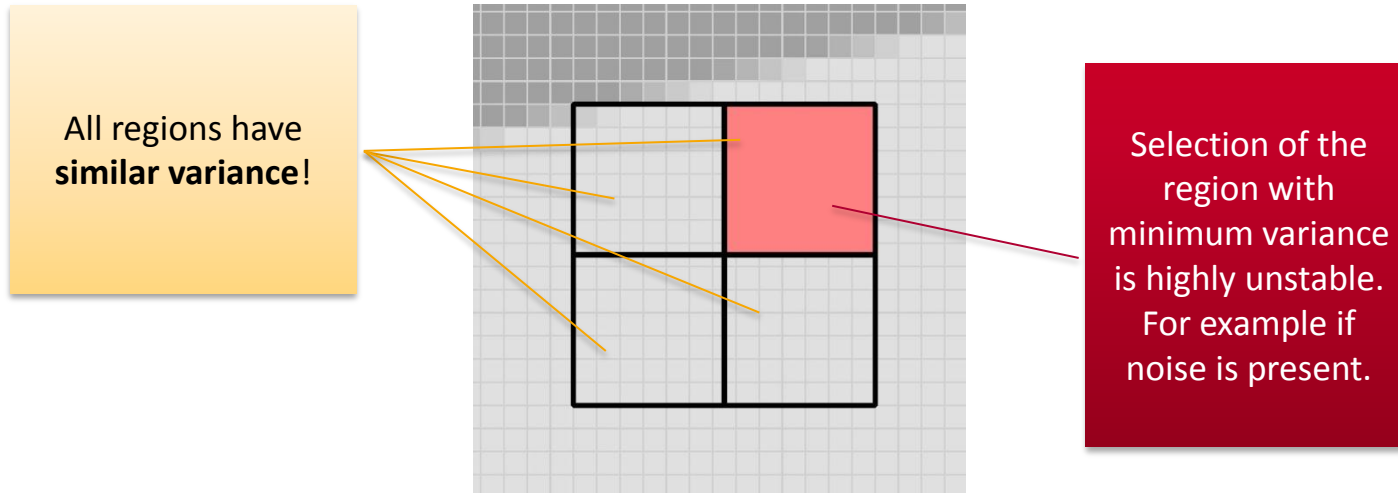
Kuwahara filter for a corner



Kuwahara filter for an edge



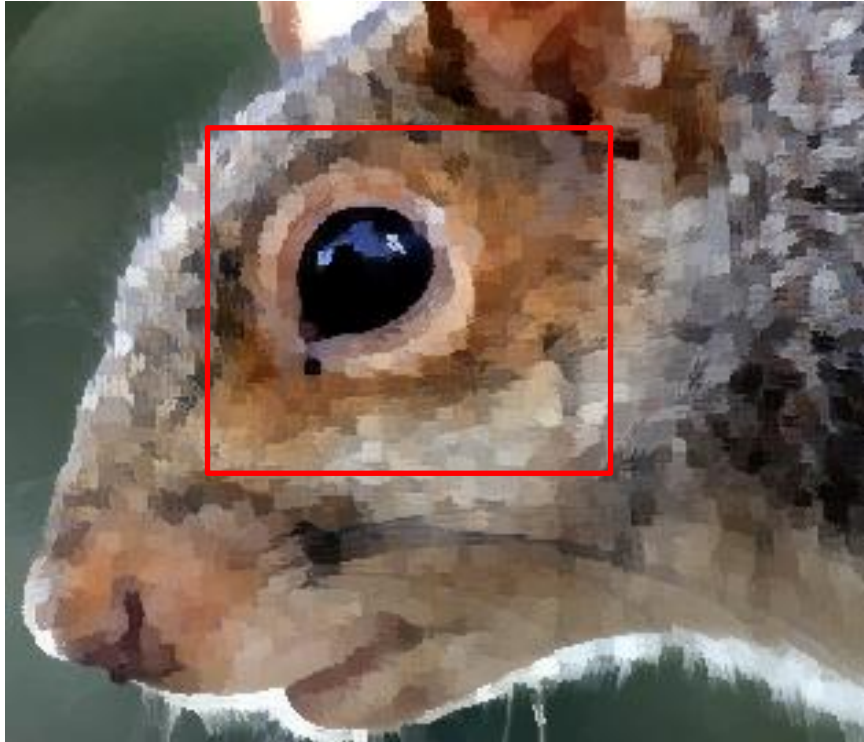
Kuwahara filter for an homogenous neighborhood






Original image by Keven Law@flickr.com







Kuwahara Filter with Weighting Functions




W_0



W_1







W_2



W_3

$$w_i(x, y) = \begin{cases} 1 & \text{if } (x, y) \in W_i \\ 0 & \text{otherwise} \end{cases}$$



Then the mean is given by:


$$\begin{aligned} m_i &= \frac{1}{|W_i|} \sum_{(x,y) \in W_i} I(x, y) \\ &= \frac{1}{|w_i|} \sum_{(x,y) \in \mathbb{Z}^2} I(x, y) \cdot w_i(x - x_0, y - y_0) \end{aligned}$$

And the variance is given by:


$$\begin{aligned} s_i^2 &= \frac{1}{|W_i|} \sum_{(x,y) \in W_i} (I(x, y) - m_i)^2 \\ &= \frac{1}{|w_i|} \sum_{(x,y) \in \mathbb{Z}^2} (I(x, y) - m_i)^2 \cdot w_i(x - x_0, y - y_0) \end{aligned}$$



Idea: Create smooth weighting functions over a disc those sum is a Gaussian




K_0




K_1





...



K_6




K_7




$$\chi_i(x, y) = \begin{cases} 1 & \frac{(2i - 1)\pi}{N} < \arg(x, y) \leq \frac{(2i + 1)\pi}{N} \\ 0 & \text{otherwise} \end{cases} \quad i = 1, \dots, N - 1$$


$$\begin{aligned} K_i &= (\chi_i \star G_{\sigma_s}) \cdot G_{\sigma_r} \\ &= K_0 \circ R_{-2\pi i/N} \end{aligned}$$



χ_0



$\chi_0 \star G_{\sigma_s}$



$(\chi_0 \star G_{\sigma_s}) \cdot G_{\sigma_r} =: K_0$

Then the mean is given by:

$$m_i = \frac{1}{|K_i|} \sum_{(x,y) \in \mathbb{Z}^2} I(x, y) \cdot K_i(x - x_0, y - y_0)$$

And the variance is given by:

$$s_i^2 = \frac{1}{|K_i|} \sum_{(x,y) \in \mathbb{Z}^2} (I(x, y) - m_i)^2 \cdot K_i(x - x_0, y - y_0)$$

The output of the generalized Kuwahara Filter is now defined by:

$$F(x_0, y_0) = \frac{\sum_{i=0}^{N-1} s_i^{-q} m_i}{\sum_{i=0}^{N-1} s_i^{-q}}$$

The parameter q is a tuning parameter that controls the sharpness of color boundaries. A typical value is $q = 8$.

$$F(x_0, y_0) = \frac{\sum_{i=0}^{N-1} s_i^{-q} m_i}{\sum_{i=0}^{N-1} s_i^{-q}}$$


Sectors low high
variance:
 $s_i \rightarrow 0 \Rightarrow s_i^{-q} \rightarrow \infty$
(i.e. most influence
to sum)

Sectors with high
variance:
 $s_i \gg 0 \Rightarrow s_i^{-q} \approx 0$
(i.e. almost no
influence to sum)



Generalized Kuwahara filter for a corner

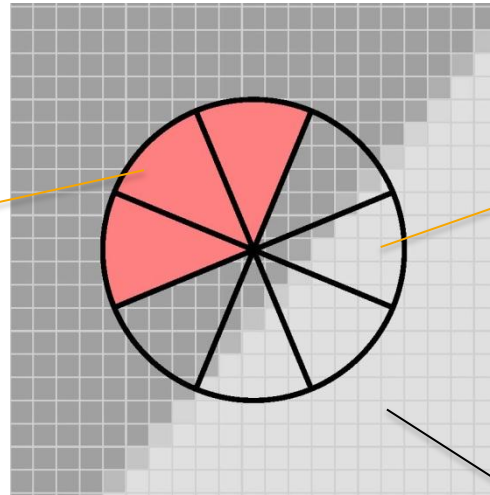
Sector with **small variance**. All pixels of this sector are similar. This sector contributes most to the final result



Sectors with **high variance**. They all contain pixels from both color regions. These sectors have almost no influence.

Generalized Kuwahara filter for an edge

Multiple sectors with **small variance**. All pixels of the sectors lie on the same side of the edge. Result is a weighted sum of the (weighted) mean values of the sectors.

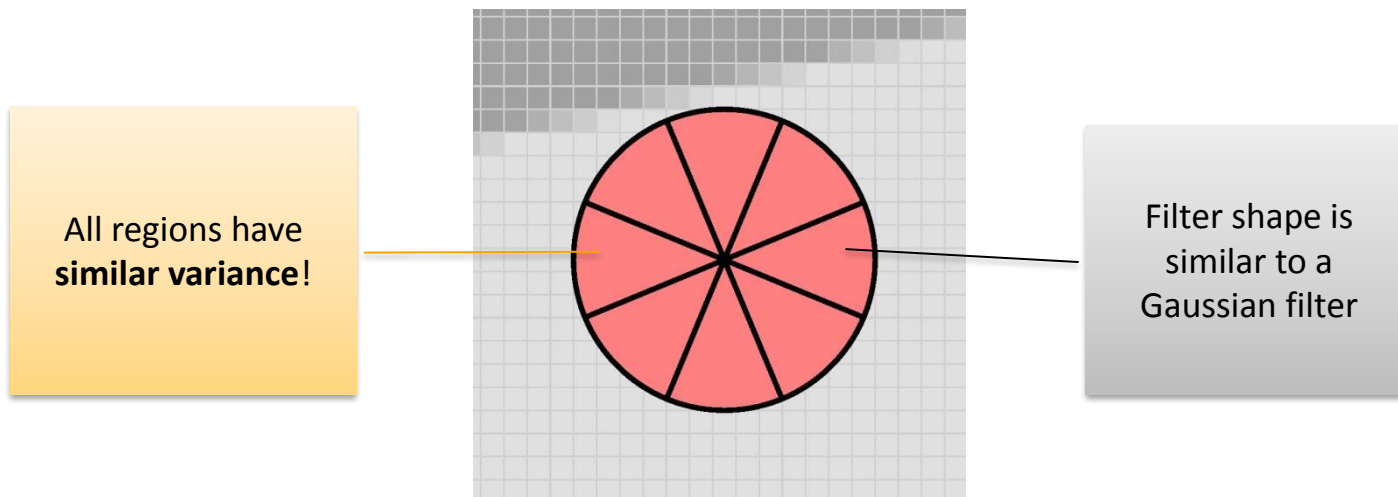


Regions with **high variance**. They all contain pixels from both sides of the edge. These sectors have almost no influence.

Filter shape is similar to a truncated Gaussian



Generalized Kuwahara filter for an homogenous neighborhood







Original image by Keven Law@flickr.com






Generalized Kuwahara Filter






K_0




K_1





...

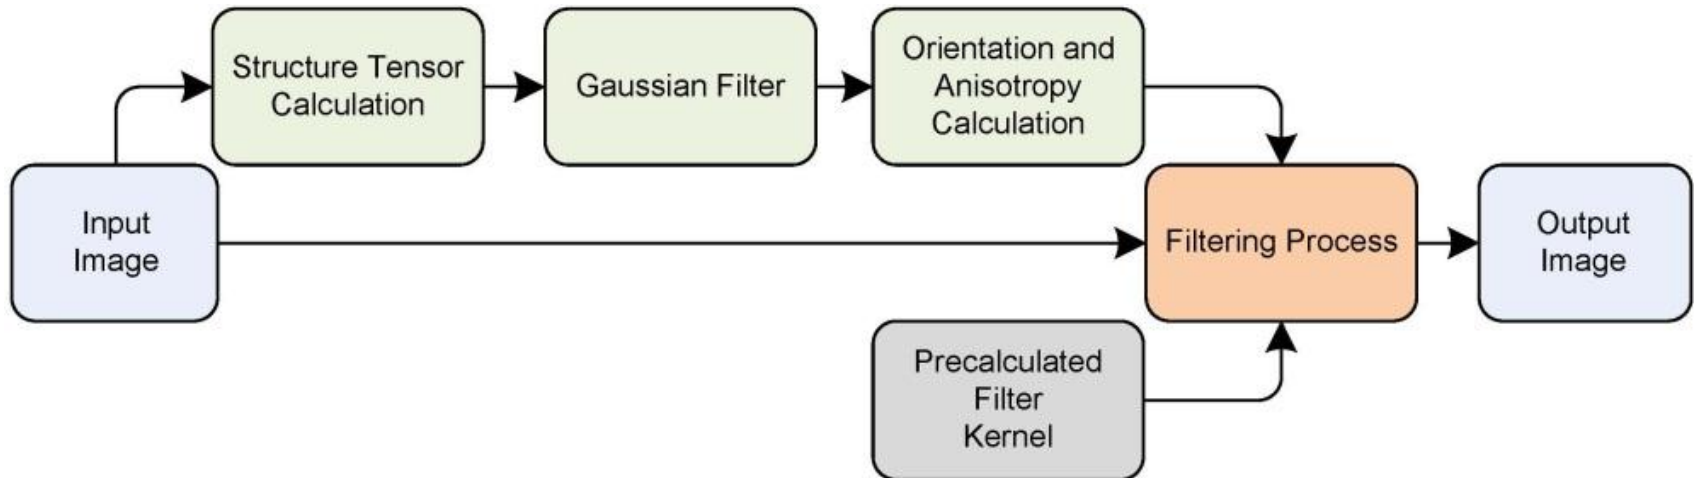


K_6



K_7






Elliptic filter shape (Pham, 2006)

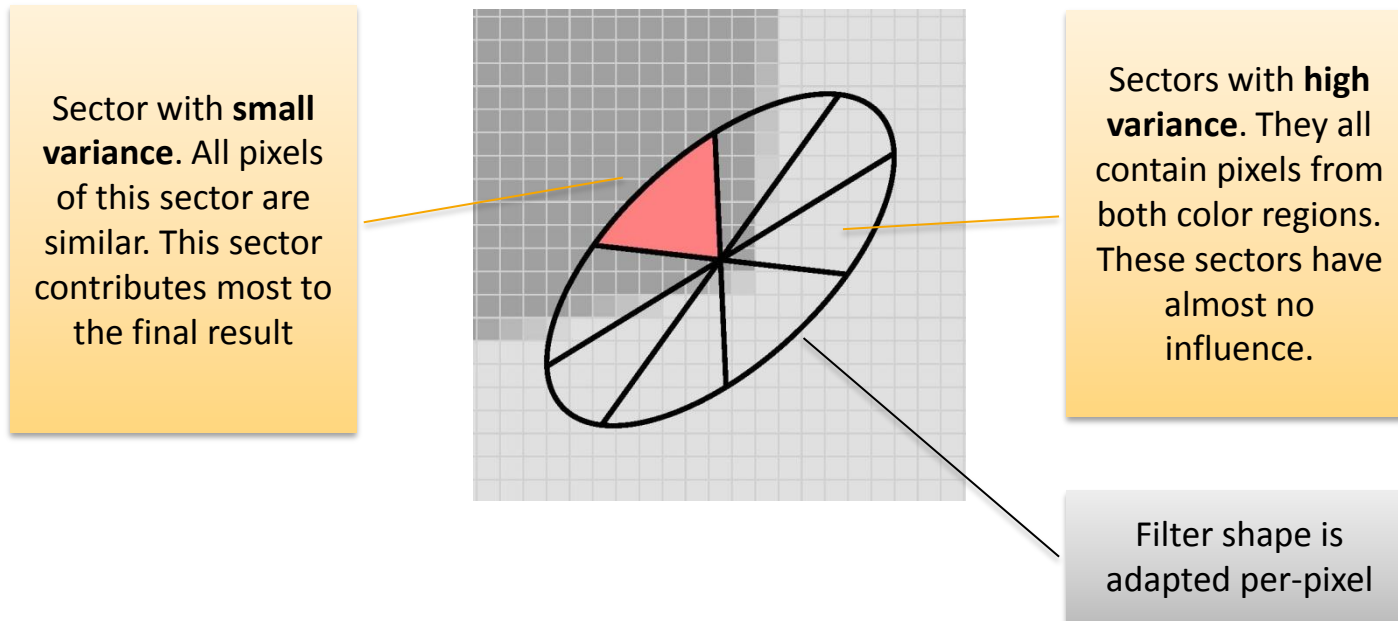
$$a = \frac{v + A}{v} \quad b = \frac{v}{v + A}$$

Here, $A \in [0,1]$ denotes the anisotropy measure derived from the structure tensor.

$v \in (0, \infty)$ is a user parameter that controls the eccentricity of the ellipse. A typical choice is $v = 1$.




Anisotropic Kuwahara filter for a corner






Anisotropic Kuwahara filter for an edge






Anisotropic Kuwahara filter for an homogenous neighborhood





Anisotropic Kuwahara filter for a homogenous neighborhood







Original image by Keven Law@flickr.com






Generalized Kuwahara Filter







$$(x + \zeta) - \eta y^2$$




$$((x + \zeta) - \eta y^2)^2$$



$$\max\{0, ((x + \zeta) - \eta y^2)^2\} := k_0$$




$$\tilde{k}_0(x, y) = \frac{k_0}{\sum_i k_i(x, y)}$$



$$\tilde{K}_0 = \tilde{k}_0 \cdot G_{\sigma_s}$$

Mean curvature flow (Alvarez et al., 1992):






Mean curvature flow


Image credit: Kang & Lee (2008)



Input



20 iterations




60 iterations

Shock filter (Osher and Rudin, 1990; Alvarez and Mazorra 1994):


In the influence zone of a maximum, the Laplacian ΔI is negative and, therefore, a dilation is performed.

$$I_t = -\text{sign}(\Delta G_\sigma \star I) |\nabla I|$$


In the influence zone of a minimum, the Laplacian ΔI is positive, which results in an erosion.



Input



Influence zones



Output

Image credit: Kang & Lee (2008)

Algorithm 1 Image Abstraction by MCF

loop

for 1 to k **do**

$I \leftarrow \text{MeanCurvatureFlow}(I)$

end for

$I \leftarrow \text{ShockFiltering}(I)$

end loop




Image abstraction by mean curvature flow


Image credit: Kang & Lee (2008)



Input



20 iterations



60 iterations

Constrained mean curvature flow:

with

$$I_t = s \cdot \kappa |\nabla I|$$

Stop diffusion if
gradient is not
aligned to edge
tangent flow (ETF)

$$s(x) = |t(x) \cdot \nabla I(x)|$$

Edge tangent
flow (ETF)

Image gradient

Algorithm 2 Image Abstraction by CMCF

loop

for 1 to k **do**

$\mathbf{t} \leftarrow TVF(I)$

$I \leftarrow ConstrainedMeanCurvatureFlow(I, \mathbf{t})$

end for

$I \leftarrow ShockFiltering(I)$

end loop




Image abstraction by constrained mean curvature flow


Image credit: Kang & Lee (2008)



Input



60 iterations image
abstraction by MCF



60 iterations image
abstraction by CMCF

Artistic Stylization of Images and Video

Part IV – Future Challenges

Eurographics 2011

John Collomosse

Centre for Vision Speech and Signal Processing (CVSSP),

University of Surrey, United Kingdom

- **Non-Photorealistic Rendering and The Science of Art**
A. Hertzmann, NPAR 2010.
- **Visual Explanations**
D. DeCarlo, M. Stone. NPAR 2010.
- **Towards Mapping the Field of Non-Photorealistic Rendering**
A. Gooch, NPAR 2010.
- **“Aaron’s Code: Meta-art, Artificial Intelligence and the Work of Harold Cohen”**
P. McCorduck. W.H. Freeman & Co. 1990. ISBN: 0716721732
- **Non-photorealistic Rendering in Context: An Observational Study**
T. Isenberg, P. Neumann, S. Carpendale, M. de Sousa, J. Jorge, NPAR 2006
- **Real-time Video Abstraction**
Salisbury et al., SIGGRAPH 2006
- **Perception and Painting: A search for effective, engaging Visualizations**
Healey, IEEE CG&A 2002.
- **Human Facial Illustrations: Creation and psychophysical evaluation**
B. Gooch, E. Reinhard, A. Gooch. ACM ToG 2004.
- **Influencing User Perception Using Real-time Adaptive Abstraction**
N. Redmond. PhD Thesis, Trinity College Dublin, 2011.

Panel session on
Grand Challenges in NPR
(NPAR 2010)

- **“Cubist-like Rendering from Photographs”**
J. Collomosse and P. Hall. IEEE TVCG 2003.
- **An Invitation to Discuss Computer Depiction**
F. Durand. NPAR 2002
- **“RTCams: A new perspective on non-photorealistic rendering from photos**
P. Hall, J. Collomosse, Y-Z. Song, P. Shen. IEEE TVCG 2007.
- **Self-similar texture for coherent line stylization**
P. Benard, F. Cole, A. Golovinskiy. NPAR 2010.
- **Human Facial Illustrations: Creation and psychophysical evaluation**
B. Gooch, E. Reinhard, A. Gooch. ACM ToG 2004.
- **Where do people draw lines?**
F. Cole, A. Golovinskiy, A. Limpacher, H. Barros, A. Finkelstein. SIGGRAPH 2008.
- **Waking Life (Movie)**
Directed by R. Linklater. Fox Searchlight 2001.
- **The Painting Fool**
www.paintingfool.com (S. Colton, Imperial College)
- **Genetic Paint: A search for salient paintings**
J. Collomosse, P. Hall. EvoMUSART 2005.

Artistic Stylization

■ Why?

Visualization

Comprehension


Communication

Aesthetics

Animation

■ Artistic Stylization can

- Simplify and structure the presentation of content
- Selectively guide attention to salient areas of content and influence perception
- Learn and emulate artistic styles
- Provide assistive tools to artists and animators (*not replace the artist!*)
- Help us to design effective visual interfaces



▪ Challenges in Aesthetics


- Easy to show novelty in style
- ... but today there are few styles remaining to pioneer
- Difficult to show superiority of one style vs. another
- Usually papers include visual comparisons side-by-side




Collomosse et al.
EvoMUSART 2005

▪ Evaluation

- First qualitative study by Isenberg et al. '06
- Compares hand-drawn and NPR images
- Unconstrained pile-sort
 - No prescribed criteria
 - Users manually group images
- Semi-structured Interview




- **Observations (on 3D pen-and-ink renderings)**
 - **Visual “Turing test” not yet passed**
 - **Pure line art mots obviously CG**
 - **Regularities should be avoided**
 - **c.f. more recent work by Cole et al. (SIGGRAPH ‘08 ’09, NPAR ‘10)**
 - **Styles less obviously CG**
 - **Stippling**
 - **Sketchy (Renderbots)**
 - **Simplified forms**
- **Know the goal / audience**
- **Portray materials (c.f. Zhu et al. ACM ToG 09)**



- **Challenges of Communication/Comprehension**
- NPR often claims to be aiming for, or to have achieved:
 - creation of a useful artist / animator's tool
 - simplification of content
 - improvements in of visual communication
- But these are rarely backed up by any form of user study
- No standard methodology has yet been agreed
 - Few have been proposed
 - Task specific:
 - Portrait recognition (Gooch, Winnemoeller)
 - Scientific visualisation (Healey)




Winnemoeller et al. 2006



C. Healey. 2002

- **Early studies measuring the effect of NPR on visual communication**
 - **Recognition speed from caricature (Gooch et al. 2004)**
 - **Recognition speed and image recall speed / accuracy (Winemoeller et al. 2006)**
- **Results:**
 - **Participants recognise stylized celebrities more quickly.**
 - **Abstraction reduces recognition latency (13% reduction)**
 - **Participants can recall stylized images more quickly / accurately.**
 - **Memory “pairs” game faster with stylized images (28% faster)**



Gooch et al. 2004

- **NPR influences perception in real-time interactive environments**
 - **Timed recognition tasks**
 - **Attention measurement (Gaze tracking)**
- **Tasks evaluated**
 - **Person / face recognition**
 - **Shape / object recognition**
 - **Crowds**
 - **Urban Navigation**
 - **Volume Visualisation**
- **Multi-way ANOVA to measure real effect**
- **Newman-Keuls post-hoc analysis**



▪ Challenge of Temporal Coherence

- Reducing flicker in stylized video remains unsolved in the general case
 - Segmentation is stable but content limited
 - Filtering is more general but unstable where texture is absent or poorly expressed
- Flicker most distracting from **6-10Hz** (*typical NPR fps!*)

▪ Twin Challenges

- Defining temporal coherence beyond “*shower door effect*”
 - Objective measures of coherence
 - Community agreement on a preferred definition
- Solving temporal coherence
 - Flicker reduction may take priority over accuracy

Defacto test clips



Hayes & Essa (NPAR 04)
J. Wang [SIGGRAPH 04]

- **Interaction with Creatives**
 - **Most NPR is pitched as a creative tool**
 - **Few are built with users in the loop**
 - **...Even fewer study use of tool in a creative context**
- **Mainstream NPR could collaborate with creative communities**
- **Examples of Computer Science/Artist interaction**
 - **Evolutionary Art Community (EvoMUSART)**
 - **Computational Aesthetics (CAe)**
- **This year NPAR, SBIM and CAe combined workshop (at SIGGRAPH'11)**
 - **Paper submission 25 April**
 - **<http://www.cl.cam.ac.uk/conference/cae-sbim-npar-2011>**

- **Portraits and Caricature**
 - **Common NPR applications are in consumer media**
 - **Mainly people and faces**
 - **Strong perceptual prior and high expectation**
- **Current NPR for portraits**
 - **Caricatures by global non-linear warping (e.g. Gooch '04)**
 - **Generally poor at emphasising salient facial detail**
 - **Higher level models needed**



Waking Life. Linklater. (c) Fox Searchlight. 2001



"Painting Fool". Colton. 2007. paintingfool.com



Gooch et al. 2004

■ Composition and Depiction

- Most NPR still focuses on low-level representation, preserving scene structure
- Artistic projections are common in artwork
 - Depiction of form not sufficiently addressed (Durand, NPAR 2002)
 - Related to “Computational Photography”

RTCams – Artistic views from stereo



Hall et al. 2007

Cubist-like Composition



Collomosse et al. 2003

- Full circle
 - Artistic composition was arguably the first NPR problem tackled (~30 years)
 - Harold Cohen's AARON – heuristic / expert system generative art
 - And it is still unsolved...



```
if (left-arm-posture is "hand-on-hip")  
  (add-upper-arm left -.3 .5 .65)  
else  
if (left-arm-posture is "arms-folded")  
  ...
```

"Aaron's Code". W.H. Freeman & Co. 1990





Open Q & A Session