

A Fresh Look at Generalized Sampling

Diego Nehab

Instituto Nacional de Matemática Pura e Aplicada (IMPA)

Hugues Hoppe

Microsoft Research

now

the essence of knowledge

Boston — Delft

Foundations and Trends[®] in Computer Graphics and Vision

Published, sold and distributed by:

now Publishers Inc.
PO Box 1024
Hanover, MA 02339
United States
Tel. +1-781-985-4510
www.nowpublishers.com
sales@nowpublishers.com

Outside North America:

now Publishers Inc.
PO Box 179
2600 AD Delft
The Netherlands
Tel. +31-6-51115274

The preferred citation for this publication is

G. Nehab and H. Hoppe. *Using the Foundations and Trends L^AT_EX Macros*.
Foundations and Trends[®] in Computer Graphics and Vision, vol. 8, no. 1,
pp. 1–84, 2012.

*This Foundations and Trends[®] issue was typeset in L^AT_EX using a class file designed
by Neal Parikh. Printed on acid-free paper.*

ISBN: 978-1-60198-729-7
© 2014 G. Nehab and H. Hoppe

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The ‘services’ for users can be found on the internet at: www.copyright.com

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; www.nowpublishers.com; sales@nowpublishers.com

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, www.nowpublishers.com; e-mail: sales@nowpublishers.com

**Foundations and Trends[®] in
Computer Graphics and Vision**
Volume 8, Issue 1, 2012
Editorial Board

Editors-in-Chief

Brian Curless
University of Washington
United States

Luc Van Gool
KU Leuven, Belgium
ETH Zurich, Switzerland

William T. Freeman
Massachusetts Institute of Technology
United States

Editors

Marc Alexa
TU Berlin

Ronen Basri
*Weizmann Institute
of Science*

Peter Belhumeur
Columbia University

Andrew Blake
Microsoft Research

Chris Bregler
New York University

Joachim Buhmann
ETH Zurich

Michael Cohen
Microsoft Research

Paul Debevec
*USC Institute
for Creative Technologies*

Julie Dorsey
Yale University

Fredo Durand
MIT

Olivier Faugeras
INRIA

Mike Gleicher
University of Wisconsin

Richard Hartley
*Australian National
University*

Aaron Hertzmann
University of Toronto

Hugues Hoppe
Microsoft Research

David Lowe
*University of
British Columbia*

Jitendra Malik
UC Berkeley

Steve Marschner
Cornell University

Shree Nayar
Columbia University

James O'Brien
UC Berkeley

Tomas Pajdla
Czech TU

Pietro Perona
Caltech

Marc Pollefeys
UNC Chapel Hill

Jean Ponce
UIUC

Long Quan
*Hong Kong University
of Science and Technology*

Cordelia Schmid
INRIA

Steve Seitz
University of Washington

Amnon Shashua
*Hebrew University
of Jerusalem*

Peter Shirley
University of Utah

Stefano Soatto
UCLA

Richard Szeliski
Microsoft Research

Joachim Weickert
Saarland University

Song Chun Zhu
UCLA

Andrew Zisserman
University of Oxford

Editorial Scope

Topics

Foundations and Trends[®] in Computer Graphics and Vision publishes survey and tutorial articles in the following topics:

- Rendering
- Shape
- Mesh simplification
- Animation
- Sensors and sensing
- Image restoration and enhancement
- Segmentation and grouping
- Feature detection and selection
- Color processing
- Texture analysis and synthesis
- Illumination and reflectance modeling
- Shape representation
- Tracking
- Calibration
- Structure from motion
- Motion estimation and registration
- Stereo matching and reconstruction
- 3D reconstruction and image-based modeling
- Learning and statistical methods
- Appearance-based matching
- Object and scene recognition
- Face detection and recognition
- Activity and gesture recognition
- Image and video retrieval
- Video analysis and event recognition
- Medical image analysis
- Robot localization and navigation

Information for Librarians

Foundations and Trends[®] in Computer Graphics and Vision, 2012, Volume 8, 4 issues. ISSN paper version 1572-2740. ISSN online version 1572-2759. Also available as a combined paper and online subscription.

Foundations and Trends® in
Computer Graphics and Vision
Vol. 8, No. 1 (2012) 1–84
© 2014 D. Nehab and H. Hoppe
DOI: 10.1561/06000000053



A Fresh Look at Generalized Sampling

Diego Nehab
Instituto Nacional de Matemática Pura e Aplicada (IMPA)

Hugues Hoppe
Microsoft Research

Contents

1	Introduction	2
2	Background	8
2.1	Reconstruction kernels	8
2.2	Analysis filters	10
3	Basic notation, definitions, and properties	12
4	Fundamental algorithms	18
4.1	Interpolation	19
4.2	Inverse discrete convolution	20
4.3	Orthogonal projection	22
4.4	Oblique projection	25
5	Translation and scaling	27
5.1	Translation of discretized signals	27
5.2	Scaling of discretized signals	29
6	Approximation of derivatives	37
7	Generalized prefiltering and estimator variance	41
8	Practical considerations	47

	iii
8.1 Grid structure	47
8.2 Efficient use of piecewise-polynomial kernels	48
8.3 Prefiltering, reconstruction, and color spaces	50
8.4 Range constraints	51
9 Theoretical considerations	53
10 Experiments and analyses	57
11 Conclusions	68
Acknowledgements	69
Appendices	70
A Source-code	71
Bibliography	81

Abstract

Discretization and reconstruction are fundamental operations in computer graphics, enabling the conversion between sampled and continuous representations. Major advances in signal-processing research have shown that these operations can often be performed more efficiently by decomposing a filter into two parts: a compactly supported continuous-domain function and a digital filter. This strategy of “generalized sampling” has appeared in a few graphics papers, but is largely unexplored in our community. This survey broadly summarizes the key aspects of the framework, and delves into specific applications in graphics. Using new notation, we concisely present and extend several key techniques. In addition, we demonstrate benefits for prefiltering in image downscaling and supersample-based rendering, and analyze the effect that generalized sampling has on the noise due to Monte Carlo estimation. We conclude with a qualitative and quantitative comparison of traditional and generalized filters.

1

Introduction

Many topics in computer graphics involve digital processing of continuous-domain data, so it is unsurprising that discretization and reconstruction are essential operations. Figure 1.1 shows the traditional sampling and reconstruction pipeline. During discretization (e.g., rasterization of a scene, or capture of a digital photograph), a continuous input signal f is passed through an




Figure 1.1: The traditional signal-processing pipeline is divided into two major stages: discretization and reconstruction. During discretization, a continuous input signal f is convolved with the reflection ψ^\vee of a given analysis filter ψ . The resulting prefiltered signal $f_\psi = f * \psi^\vee$ is then uniformly sampled into a discrete sequence $[f_\psi]$. To obtain the output approximation \tilde{f} , the reconstruction stage computes the mixed convolution between $[f_\psi]$ and a given reconstruction kernel φ , i.e., a sum of shifted copies of φ , where each shifted copy scaled by the corresponding entry in $[f_\psi]$. (Our notation is explained in greater depth in section 3.)




Figure 1.2: A continuous function f is prefiltered with analysis kernel ψ (here the box function β^0 , not to scale). The resulting signal f_ψ is sampled into a discrete sequence $[f_\psi]$. The final output \tilde{f} is obtained by mixed convolution between the discrete sequence $[f_\psi]$ and the reconstruction kernel φ (here the hat function β^1 , not to scale).

analysis filter ψ (a.k.a. *sampling kernel*, *prefilter*, or *antialiasing filter*) before being sampled. The result is a discrete sequence $[f_\psi]$ (e.g., an image). During reconstruction (e.g., interpolation of a texture, or display of an image on a screen), the continuous approximation \tilde{f} of the original signal is obtained by mixed convolution with a *reconstruction kernel* φ (a.k.a. *generating function*, *basis function*, or *postfilter*). Figure 1.2 illustrates each step of the process with a concrete example in 1D.

The roles of the analysis filter ψ and reconstruction kernel φ are traditionally guided by the sampling theorem [Shannon, 1949]. Given a sampling rate $1/T$, the analysis filter $\psi = \text{sinc}(\cdot/T)$ eliminates from the input signal f those frequencies higher than or equal to $1/2T$ so that the bandlimited f_ψ can be sampled without aliasing. And in that case, the reconstruction kernel $\varphi = \text{sinc}(\cdot/T)$ recreates $\tilde{f} = f_\psi$ exactly from the samples.

Sampling may also be interpreted as the problem of finding the function \tilde{f} that minimizes the norm of the residual $\|f - \tilde{f}\|_{\mathcal{L}_2}$. If we restrict our attention




Figure 1.3: The main idea in generalized sampling is to broaden the analysis and reconstruction kernels by expressing these as mixed convolutions ($\mathbf{p} * \psi$ and $\mathbf{r} * \varphi$) with a pair of digital filters (\mathbf{p} and \mathbf{r}) while retaining compact support for the functions ψ and φ .




Figure 1.4: The traditional Keys cubic (Catmull-Rom spline) K has support 4 and a reasonably sharp frequency response $|\hat{K}|$. The cardinal cubic B-Spline β_{int}^3 is a generalized kernel formed from the basic cubic B-spline β^3 and a digital filter. The digital filter acts to widen support to infinity (though with exponential decay) and to significantly sharpen the frequency response.

to the space of bandlimited functions, the ideal prefilter is still $\psi = \text{sinc}(\cdot/T)$. However, functions are often not bandlimited in practice (e.g., due to object silhouettes, shadow boundaries, vector outlines, detailed textures), and for efficiency we desire ψ and φ to be compactly supported.

In addressing these concerns, the signal-processing community has adopted a generalization of the sampling and reconstruction pipeline [Unser, 2000]. The idea is to represent the prefilter and reconstruction kernels as mixed convolutions of *compactly supported* kernels and *digital filters*. As shown in figure 1.3, digital filters \mathbf{p} and \mathbf{r} respectively modify the prefilter ψ and the reconstruction kernel φ . The additional degrees of freedom and effectively larger filter support enabled by \mathbf{p} and \mathbf{r} allow the design of generalized kernels with better approximation properties or sharper frequency response. Figure 1.4 compares a traditional piecewise cubic kernel (the Catmull-Rom spline, or Keys cubic) with a generalized cubic kernel (the cardinal cubic B-spline).




Figure 1.5: Generalized sampling adds a digital filtering stage to the pipeline. The output $[f_\psi]$ of the sampling stage is convolved with a digital transformation filter $q = p^v * r$. It is the result c of this stage (and not $[f_\psi]$) that is convolved with the reconstruction kernel φ to produce the output signal.

Equivalently, the digital filters p and r can be combined as $q = p^v * r$ into a separate filter stage as shown in figure 1.5. The result $[f_\psi]$ of the sampling stage is transformed by the digital filter q (a.k.a. *correction* or *basis change*) to form a new discrete sequence c , which is then convolved with φ as usual to reconstruct \tilde{f} . The key to the efficiency of this generalized sampling framework is that the digital filters p and r that arise in practice are typically compact filters or their inverses [Unser et al., 1991], both of which are parallelizable on multicore CPU and GPU architectures [Ruijters et al., 2008, Nehab et al., 2011]. Thus, the correction stage adds negligible cost.

An important motivation for generalized sampling is improved interpolation [Blu et al., 1999]. As demonstrated in figure 1.6, an image $[f_\psi]$ is processed by a digital filter q resulting in a coefficient array c which can then be efficiently reconstructed with a simple cubic B-spline filter β^3 . The resulting interpolation is sharper and more isotropic (i.e., has higher quality) than that produced by the popular Mitchell-Netravali filter [1988], even though both filters have the same degree and support. The implementation of the digital filtering stage is described in detail in section 4.2. The theory of image upscaling is described in section 5.2, with implementation notes in section 8.2. Source-code is provided in appendix A.

In graphics, careful prefiltering is often necessary to prevent aliasing. McCool [1995] describes an early application of generalized sampling, in which rendered triangles are antialiased analytically by evaluating a prism spline prefilter. The resulting image is then convolved with a digital filter. In this work, we apply generalized sampling to image downscaling and in general




Figure 1.6: Reconstruction example. The top row shows the result of the traditional cubic Mitchell-Netravali filter M . The bottom row uses the generalized sampling approach, first applying a digital filter $r = \left[\frac{1}{6}, \frac{4}{6}, \frac{1}{6}\right]^{-1}$ as a preprocess, and then reconstructing with the cubic B-spline β^3 — which is less expensive to evaluate on a GPU than filter M .

to rendering with supersampling. Figure 1.7 shows an example. The input f is prefiltered using the cubic B-spline basis β^3 . The resulting over-blurred image is then transformed with a digital filter p^\vee that reshapes the antialiasing kernel *a posteriori*. The final low-resolution image is sharper and exhibits less aliasing than with a Catmull-Rom filter, for a similar computational cost. The theory of image downscaling is described in section 5.2, with implementation notes in section 8.2 and source-code in appendix A. Generalized supersampling is described in section 7.

Our aim is to present a concise overview of the major developments in generalized sampling and to extend these techniques to prefiltering in graphics. To facilitate exposition and exploration, we develop a new concise notation for sampling. With this parameter-free notation, key techniques are derived using simple algebraic manipulation. We conclude by comparing a variety of




Figure 1.7: Prefiltering example. The top row shows the result of rendering with the Keys (Catmull-Rom) prefilter K . The bottom row shows rendering using a B-spline β^3 , followed by convolution with a digital filter $p^v = [\frac{1}{6}, \frac{4}{6}, \frac{1}{6}]^{-1}$. The generalized prefilter $p * \beta^3$ equals the cubic cardinal B-spline β_{int}^3 . Kernels K and β^3 have the same support, but the improved frequency response of β_{int}^3 reduces aliasing while maintaining sharpness. (Our notation is explained in section 3.)

traditional and generalized filters, using frequency-domain visualizations and empirical experiments using both \mathcal{L}_2 and SSIM metrics, to identify the best strategies available.

Bibliography

- A. Aldroubi and M. Unser. Sampling procedures in function spaces and asymptotic equivalence with Shannon's sampling theory. *Numerical Functional Analysis and Optimization*, 15(1–2):1–21, 1994.
- U. Alim, T. Möller, and L. Condat. Gradient estimation revitalized. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1495–1504, 2010.
- J. F. Blinn. Return of the jaggy. *IEEE Computer Graphics and Applications*, 9(2):82–89, 1989.
- J. F. Blinn. A ghost in a snowstorm. *IEEE Computer Graphics and Applications*, 18(1):79–84, 1998.
- T. Blu and M. Unser. Quantitative Fourier analysis of approximation techniques: Part I—Interpolators and projectors. *IEEE Transactions on Signal Processing*, 47(10):2783–2795, 1999a.
- T. Blu and M. Unser. Approximation error for quasi-interpolators and (multi-)wavelet expansions. *Applied and Computational Harmonic Analysis*, 6(2):219–251, 1999b.
- T. Blu, P. Thévenaz, and M. Unser. Generalized interpolation: Higher quality at no additional cost. In *Proceedings of the IEEE International Conference on Image Processing*, volume 3, pages 667–671, 1999.
- T. Blu, P. Thévenaz, and M. Unser. MOMS: Maximal-order interpolation of minimal support. *IEEE Transactions on Image Processing*, 10(7):1069–1080, 2001.
- T. Blu, P. Thévenaz, and M. Unser. Linear interpolation revitalized. *IEEE Transactions on Image Processing*, 13(5):710–719, 2004.
- E. Catmull and R. Rom. A class of local interpolating splines. In *Computer Aided Geometric Design*, pages 317–326, 1974.

- L. Condat and T. Möller. Quantitative error analysis for the reconstruction of derivatives. *IEEE Transactions on Image Processing*, 59(6):2965–2969, 2011.
- L. Condat, T. Blu, and M. Unser. Beyond interpolation: optimal reconstruction by quasi-interpolation. In *Proceedings of the IEEE International Conference on Image Processing*, volume 1, pages 33–36, 2005.
- Franklin Crow. Summed-area tables for texture mapping. *Computer Graphics (Proceedings of ACM SIGGRAPH 84)*, 18(3):207–212, 1984.
- M. Dalai, R. Leonardi, and P. Migliorati. Efficient digital pre-filtering for least-squares linear approximation. In *Visual Content Processing and Representation*, volume 3893 of *Lecture Notes in Computer Science*, pages 161–169. 2006.
- C. de Boor. Quasiinterpolants and the approximation power of multivariate splines. Technical Report #99-12, University of Wisconsin-Madison, 1989.
- M. A. Z. Dippé and E. H. Wold. Antialiasing through stochastic sampling. *Computer Graphics (Proceedings of ACM SIGGRAPH 1985)*, 19(3):69–78, 1985.
- N. A. Dodgson. Quadratic interpolation for image resampling. *IEEE Transactions on Image Processing*, 6(9):1322–1326, 1997.
- C. E. Duchon. Lanczos filtering in one and two dimensions. *Journal of Applied Meteorology*, 18(8):1016–1022, 1979.
- I. German. Short kernel fifth-order interpolation. *IEEE Transactions on Signal Processing*, 45(5):1355–1359, 1997.
- R. W. Hamming. *Digital Filters*. Prentice Hall, 1977.
- P. S. Heckbert. Filtering by repeated integration. *Computer Graphics (Proceedings of ACM SIGGRAPH 1986)*, 20(4):315–321, 1986.
- H. S. Hou and H. C. Andrews. Cubic splines for image interpolation and digital filtering. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 25(6):508–517, 1978.
- R. Hummel. Sampling for spline reconstruction. *SIAM Journal on Applied Mathematics*, 43(2):278–288, 1983.
- J. Kajiya and M. Ullner. Filtering high quality text for display on raster scan devices. *Computer Graphics (Proceedings of ACM SIGGRAPH 1981)*, 15(3):7–15, 1981.
- R. G. Keys. Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 29(6):1153–1160, 1981.
- C. Lee, M. Eden, and M. Unser. High-quality image resizing using oblique projection operators. *IEEE Transactions on Image Processing*, 7(5):679–692, 1998.
- M. A. Malcolm and J. Palmer. A fast method for solving a class of tridiagonal linear systems. *Communications of the ACM*, 17(1):14–17, 1974.

- M. D. McCool. Analytic antialiasing with prism splines. In *Proceedings of ACM SIGGRAPH 1995*, pages 429–436, 1995.
- E. H. W. Meijering. A chronology of interpolation: From ancient astronomy to modern signal processing. *Proceedings of the IEEE*, 90(3):319–342, 2002.
- E. H. W. Meijering, W. J. Niessen, J. P. W. Pluim, and M. A. Viergever. Quantitative comparison of sinc-approximating kernels for medical image interpolation. In *Medical Image Computing and Computer-Assisted Intervention*, volume 1679 of *Lecture Notes in Computer Science*, pages 210–217. 1999a.
- E. H. W. Meijering, K. J. Zuiderveld, and M. A. Viergever. Image reconstruction by convolution with symmetrical piecewise n th-order polynomial kernels. *IEEE Transactions on Image Processing*, 8(2):192–201, 1999b.
- E. H. W. Meijering, W. J. Niessen, and M. A. Viergever. Quantitative evaluation of convolution-based methods for medical image interpolation. *Medical Image Analysis*, 5(2):111–126, 2001.
- D. P. Mitchell and A. N. Netravali. Reconstruction filters in computer graphics. *Computer Graphics (Proceedings of ACM SIGGRAPH 1988)*, 22(4):221–228, 1988.
- A. Muñoz, T. Blu, and M. Unser. Least-squares image resizing using finite differences. *IEEE Transactions on Image Processing*, 10(9):1365–1378, 2001.
- D. Nehab, A. Maximo, R. S. Lima, and H. Hoppe. GPU-efficient recursive filtering and summed-area tables. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH Asia 2011)*, 30(6):176, 2011.
- S. K. Park and R. A. Schowengerdt. Image reconstruction by parametric cubic convolution. *Computer Vision, Graphics & Image Processing*, 23(3):258–272, 1983.
- J. A. Parker, R. V. Kenyon, and D. E. Troxel. Comparison of interpolating methods for image resampling. *IEEE Transactions on Medical Imaging*, MI-2(1):31–39, 1983.
- K. Perlin. State of the art in image synthesis, 1985. SIGGRAPH 1985 Course Notes. Pixar. *The RenderMan Interface*, 2005. Version 3.2.1.
- D. Ruijters, B. M. ter Haar Romeny, and P. Suetens. Efficient GPU-based texture interpolation using uniform B-splines. *Journal of Graphics, GPU & Game Tools*, 13(4):61–69, 2008.
- R. W. Schafer and L. R. Rabiner. A digital signal processing approach to interpolation. *Proceedings of the IEEE*, 61(6):692–702, 1973.
- A. Schaum. Theory and design of local interpolators. *Computer Vision, Graphics & Image Processing*, 55(6):464–481, 1993.

- C. E. Shannon. Communication in the presence of noise. *Proceedings of the Institute of Radio Engineers*, 37(1):10–21, 1949.
- C. Sigg and M. Hadwiger. Fast third-order texture filtering. In M. Pharr, editor, *GPU Gems 2*, chapter 20, pages 313–329. Addison Wesley Professional, 2005.
- G. Strang and G. Fix. A Fourier analysis of the finite element variational method. In *Constructive Aspects of Functional Analysis*, pages 793–840, 1973.
- P. Thévenaz, T. Blu, and M. Unser. Interpolation revisited. *IEEE Transactions on Medical Imaging*, 19(17):739–758, 2000.
- M. Unser. Sampling—50 years after Shannon. *Proceedings of the IEEE*, 88(4): 569–587, 2000.
- M. Unser and A. Aldroubi. A general sampling theory for nonideal acquisition devices. *IEEE Transactions on Signal Processing*, 42(11):2915–2925, 1994.
- M. Unser, A. Aldroubi, and M. Eden. Fast B-spline transforms for continuous image representation and interpolation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(3):277–285, 1991.
- M. Unser, A. Aldroubi, and M. Eden. Enlargement or reduction of digital images with minimum loss of information. *IEEE Transactions on Image Processing*, 4(3): 247–258, 1995a.
- M. Unser, P. Thévenaz, and L. Yaroslavsky. Convolution-based interpolation for fast, high-quality rotation of images. *IEEE Transactions on Image Processing*, 4(10): 1371–1381, 1995b.
- Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.