Non-Linear Gaussian Filters Performing Edge Preserving Diffusion

Volker Aurich and Jörg Weule

aurich@cs.uni-duesseldorf.de weule@cs.uni-duesseldorf.de
Abteilung für Informatik, Mathematisches Institut
Heinrich-Heine-Universität Düsseldorf, Universitätsstrasse 1, D-40225 Düsseldorf

Abstract: This paper presents a new diffusion method for edge preserving smoothing of images. In contrast to other methods it is not based on an anisotropic modification of the heat conductance equation, rather on a modification of the way the solution of the heat conductance equation is obtained by convolving the initial data with a Gaussian kernel. Hence the method uses simple non-linear modifications of Gaussian filters, thus avoiding iteration steps and convergence problems. A chain of three to five filters with suitable parameters provides excellent smoothing of fine image details without destroying the coarser structures. The size and contrast of the eliminated details can be selected. The choice of the parameters is not critical and the edges are not displaced when changing the scale. The filter stages can be implemented efficiently on almost any parallel hardware architecture.

Keywords: low-level processing, edge preserving smoothing, anisotropic diffusion, robust filters.

1 Introduction

Frequently images contain fine structured details which are not important for further processing, at least if their contrast is below some treshold. Procedures which eliminate these fine structures without destroying the coarser ones are called edge preserving smoothing. Often the problem is considered as a reconstruction problem by modelling the image f as a sum $f_1 + f_2$ of a coarse structured signal f_1 and a fine structured signal f_2. But there is no precise mathematical meaning of coarse and fine structured; rather the notions are defined implicitly by the various smoothing procedures. A definition via the frequency domain is not adequate because a coarse structured image can contain sharp edges and hence high as well as low frequencies. Therefore edge preserving smoothing cannot be performed by linear filtering. The most important nonlinear methods which have been developed are based on numerical approximation of f by predefined coarse signals [BZ87], anisotropic diffusion [ALM92, PM90, Whi93, WP93] or robust filtering [Fle92, NM79, WVL81, Lee83, SMCM91, SVR80]. Methods of robust statistics naturally occur if the fine structured signal f_2 is modelled as realization of a family of random variables Z_p indexed by the pixels, because in real images neither the exact distribution of each Z_p nor the correlation between the Z_p is known.

We present a chain of non-linear filters which combines the aspects of anisotropic diffusion and robust filtering. It possesses the advantages of both, but avoids most of their drawbacks.

2 The non-linear Gaussian filter

A linear Gaussian filter calculates the convolution product $Gf = g \star f$ of the signal f with a Gaussian function g. Since the image signals are only defined on a bounded discrete set $\mathcal{P}$ of pixels, boundary effects have to be taken into account. For our purposes it is convenient to define the linear Gaussian filter by using the following position-dependent renormalization of the weights.

$$\begin{aligned} G_{\sigma_x} f(p) &= \frac{1}{N_p} \sum_{q \in \mathcal{P}} g_{\sigma_x}(\|q-p\|) f(q) \\ &= f(p) + \frac{1}{N_p} \sum_{q \in \mathcal{P}} g_{\sigma_x}(\|q-p\|) \cdot (f(q) - f(p)) \\ \text{with} \qquad & g_{\sigma_x}(t) = \exp(\frac{-t^2}{2\sigma_x^2}) \quad \text{and} \quad N_p = \sum_{q \in \mathcal{P}} g_{\sigma_x}(\|q-p\|). \end{aligned}$$

In order to preserve edges the influence of $f(q)$ should decrease if $|f(q) - f(p)|$ increases. This can be achieved by multiplying $g_{\sigma_x}(\|q-p\|)$ by an additional weight $\psi(f(q) - f(p))$. Because the noise distribution in images is very often bell-shaped we chose $\psi = g_{\sigma_z}$. Thus the non-linear Gaussian filter is defined by

$$\begin{aligned} G_{\sigma_x,\sigma_z}(p) &= f(p) + \frac{1}{N_p} \sum_{q \in \mathcal{P}} g_{\sigma_x}(\|q-p\|) g_{\sigma_z}(f(q) - f(p)) \cdot (f(q) - f(p)) \\ \text{with} \qquad & N_p = \sum_{q \in \mathcal{P}} g_{\sigma_x}(\|q-p\|) g_{\sigma_z}(f(q) - f(p)). \end{aligned}$$

Notice that now the renormalization factor N_p depends not only on the location p, but also on the input data f. Similar operators are used as estimators in robust statistics.

The non-linear Gaussian filters are resistant i.e. the mean value they calculate is insensitive to outliers. If the fine structured signal f_2 is modelled as realization of random variables like in the introduction, they can be shown to be robust estimators [Weu94].

Since the Gaussian kernel is the fundamental solution of the heat conductance equation, the nonlinear Gaussian filter can also be considered to calculate the result of some kind of anisotropic diffusion.

In order to increase the noise reduction, a parameter η is added to the non-linear Gaussian filter:

$$G_{\sigma_x,\sigma_z,\eta}(p) = f(p) + \eta \cdot \frac{1}{N_p} \sum_{q \in \mathcal{P}} g_{\sigma_x}(\|q-p\|) g_{\sigma_z}(f(q) - f(p)) \cdot (f(q) - f(p))$$

Theoretical considerations in [Weu94] imply that the parameter η should be in the range $[1, 1.5]$. To illustrate the effect of η, two filters with $\eta = 1$ and $\eta = 1.3$ were applied to Gaussian noise. The diagrams in figure 1 show the distributions of the input and output noise of these filters. The filter with $\eta = 1.3$ produces a narrower noise distribution.

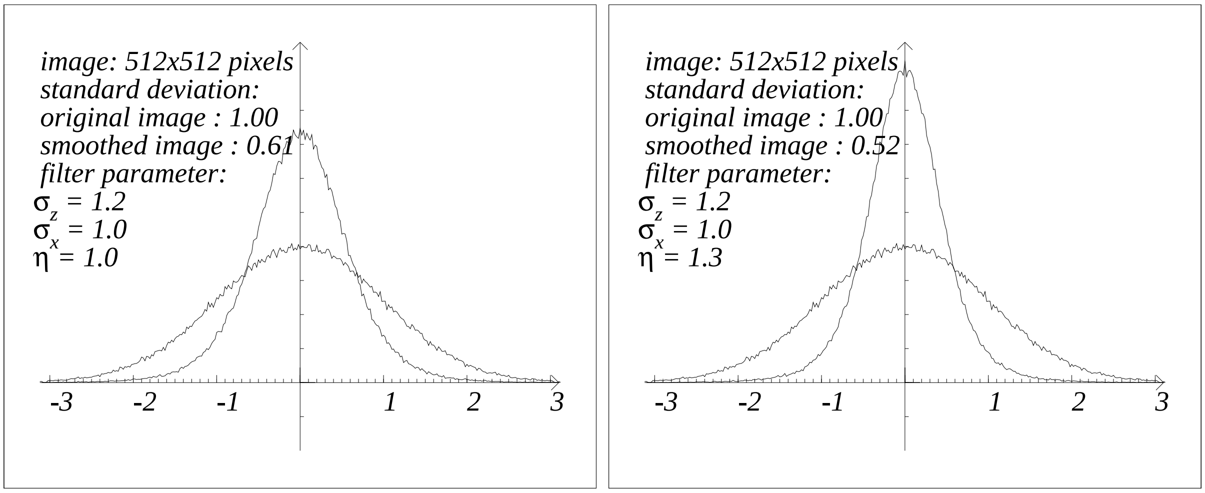

Figure 1: Effect of the parameter η while smoothing Gaussian noise.

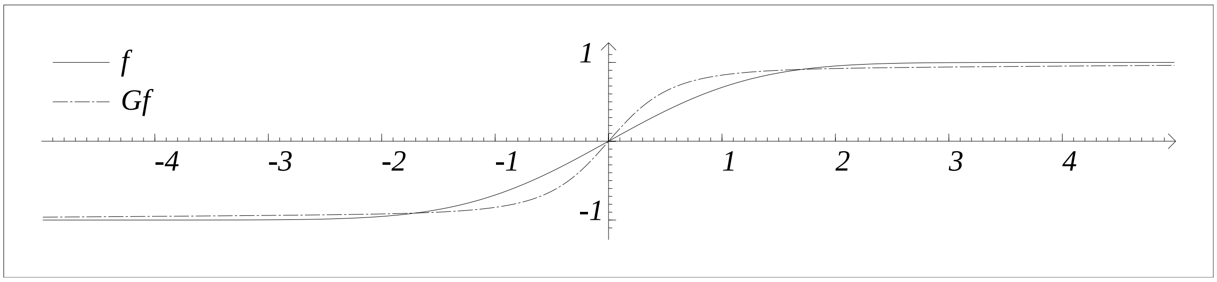

Figure 2: Sharpening of a blurred edge.

3 Sharpening of blurred edges

When the filter is applied to a step edge or ramp edge the shape of the output signal depends strongly on the filter parameters. If σ_z is larger than the maximal grey value the non-linear Gaussian filter tends to work like the linear Gaussian filter, hence edges will be smoothed. If σ_z is small and σ_x is large, the filter can sharpen ramp edges as illustrated in figure 2. The ramp edge in this figure is modelled as a 1-d function $f\colon \mathbb{R} \to \mathbb{R}$ to which continuous 1-d version of the filter is applied:

$$\begin{aligned}
f(s) &:= 2 \cdot \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{x} exp\left(\frac{-s^2}{2}\right) ds - 1 \\
G_{\sigma_x\sigma_z} f(s) &:= f(s) + \frac{1}{N_s} \int_{\infty}^{\infty} g_{\sigma_x}(s-t) \cdot g_{\sigma_z}(f(t) - f(s)) \cdot (f(t) - f(s)) \ dt \\
\text{with} \qquad & N_s = \int_{\infty}^{\infty} g_{\sigma_x}(s-t) \cdot g_{\sigma_z}(f(t) - f(s)) \ dt
\end{aligned}$$

Figure 3 shows the difference of the derivatives $[Gf]'(0) - f'(0)$ for $\sigma_z = 0.5$ as a function of σ_x. Notice that the ramp edge f can be thought of as a step edge blurred by a linear Gaussian filter with $\sigma_x = 1$. The parameter σ_x of the applied non-linear Gaussian filter should be chosen larger than 1 to sharpen the ramp edge f.

4 The filter chain

In general the smoothing effect of a single filter is not satisfactory. Therefore several filter steps with different parameters are used. The first filter step serves mainly for reducing

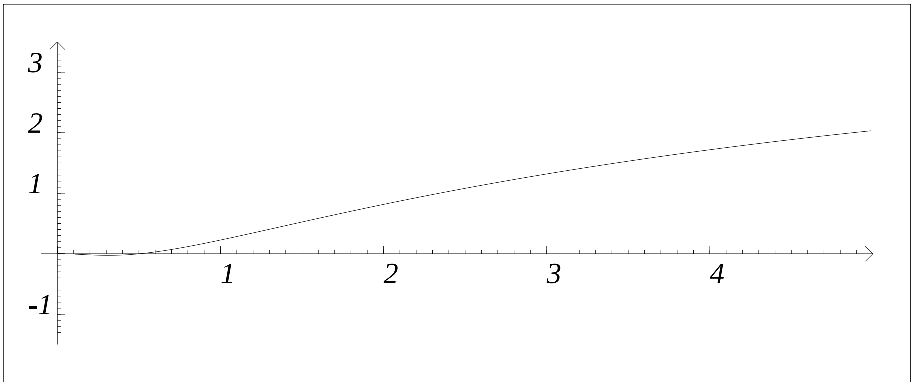

$[G_{\sigma_x,\sigma_z} f]'(0) - f'(0)$ for $\sigma_z = 0.5$ as a function of σ_x.

Figure 3: sharpening of a blurred step edge

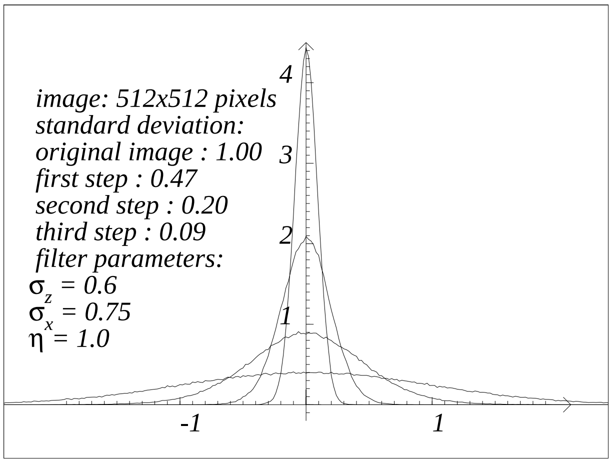

Figure 4: Distribution of noise between filter steps.

the contrast of the fine details. The following steps continue this reduction, but sharpen at the same time the edges of the coarser structures which may have been blurred by the first filter step. The parameter σ_x of the first step should be as small as possible in order to avoid unnecessary blurring of coarser structures, but at least so large as the size of the fine structures. The parameter σ_z of the first step determines approximately the maximal contrast fine details may have in order to be finally eliminated. To make the following filter steps sharpen the edges of the coarser structures as described in the previous section, σ_x has to be increased and σ_z has to be decreased. Extensive experiments showed that almost always the best results are achieved if σ_x is doubled and σ_z is halved before performing the next filter step. A theoretical proof is more or less impossible because the output of a non-linear Gaussian filter consists of random variables which are in general highly dependent in a rather complicated way. But the choice of octave steps can be made plausible by empirically estimating the distributions of the filter outputs when the filter chain is fed by Gaussian noise [WEU94]. The following diagram shows the noise distributions of the output signals of three filter stages. The initial parameters i.e. the parameters of the first filter step are indicated.

When filtering real images typical values for the initial σ_x are 0.5 to 2 times the pixel distance. Usually between three and five filter steps suffice to obtain images which consist of patches with slowly varying grey values. Further filter steps change the image only slightly or not at all because their parameters σ_z drop finally below the quantization step of the grey values.

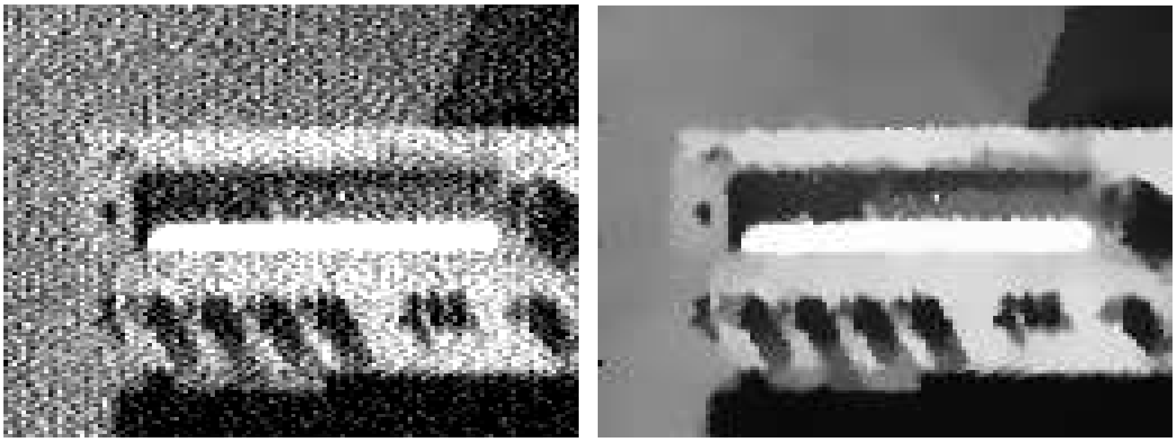

original image $\sigma_x = 0.5, \sigma_z = 128, \eta = 1$

Figure 5: Smoothing a radio by a filter chain with five steps.

The Gaussian filter chain can be considered as an anisotropic diffusion process which is stopped at certain times and restarted with new parameters.

5 Modified Filter Step

Whereas a filter chain does not displace edges it may happen that the values of the filter chain output f_o in the regions between the edges differ slightly from the corresponding local averages of the original image f. This can be avoided by appending a modified filter step which calculates for each pixel p the value

$$\frac{1}{N_p} \sum_{q \in \mathcal{P}} g_{\sigma_x}(\|q-p\|) g_{\sigma_z}(f_o(q) - f_o(p)) \cdot (f(q) - f(p))$$

$$\text{with} \qquad N_p = \sum_{q \in \mathcal{P}} g_{\sigma_x}(\|q-p\|) g_{\sigma_z}(f_o(q) - f_o(p)).$$

This operator smoothes f while respecting the edges present in f_o.

6 Experimental results

In figure 5 a very noisy picture is smoothed by a filter chain with five stages and initial parameters $\sigma_x = 0.5$ and $\sigma_z = 128$.

Figure 6 shows the intensity profile of line 112 in the original lena, 512×512 pixels, image and in the output of a filter chain with 3 stages and initial parameters $\sigma_x = 4$ and $\sigma_z = 36$. Notice that small details are preserved if their contrast is high enough and that the edges are not displaced. If a modified filter stage is appended the original height of the two lobes at the left side is quite exactly reconstructed.

The non-linear Gaussian filters can be efficiently implemented on modern RISC workstations. Moreover there are obvious ways to split up each filter step into several parts which can be executed simultaneously. Hence parallel hardware can be employed. J.Wienbeck implemented the filter on a PowerXplorer of Parsitek, a parallel machine with 8 PowerPC-processors connected by transputers. The image is downloaded from the host, each processor calculates the result of the first filter step within a stripe of the image and sends the result to the neighbours, then the next step is executed. After all filter steps have been performed, the image is uploaded to the host. The overhead caused by the communication between the processors is not very high; N processors need almost only one Nth of the time one processor needs.

The execution times depend heavily on the image size and on σ_x, but not on σ_z. They can be strongly reduced if the summation is not carried out over all pixels $q \in \mathcal{P}$, but only over the pixels in a small neighbourhood of p. We choose a square with center p and side length $2\alpha\sigma_x$; the factor α is gradually lowered from 3.0 in the first filter step to 1.0 in the last step. This reduction of α does not cause significant changes in the filtered image. The following table shows the execution times for a few filter chains with three stages on a Sun SparcStation 20/50 and on a PowerXplorer with 8 processors. The radio image, 128×128 pixels, in figure 5 was filtered by the PowerXplorer in 2.5 seconds (5 filter steps).

	image size	512 × 512	image size	256 × 256
initial σ_x	Sparc 20/50	Xplorer	Sparc 20/50	Xplorer
0.5	11 s	8 s	3 s	2.3 s
1.0	24 s	10 s	6 s	2.5 s
2.0	65 s	15 s	16 s	4 s
4.0	205 s	32 s	52 s	8 s
8.0	817 s	100 s	184 s	24 s

7 Conclusion

The described chain of non-linear Gaussian filters is an edge preserving smoothing method which can even sharpen edges. It avoids large number of iterations and convergence problems. Edges are not displaced when the scale is changed. Only two filter parameters with an intuitive meaning have to be chosen; the choice is not critical. The filters can easily and efficiently be implemented on parallel hardware.

The method is not restricted to two-dimensional image data. It can also be applied to signals $f : I\!R^n \to I\!R^m$, especially to color images and to volume data.

References

[ALM92] Luis Alvarez, Pierre-Louis Lions, and Jean-Michel Morel. Image selective smoothing and edge detection by nonlinear diffusion. ii. *SIAM J. Numer. Anal.*, 29(3):845–866, 1992.

[BZ87] Andrew Blake and Andrew Zisserman. *Visual Reconstruction.* MIT Press, Cambridge (Massachusette), London (England), 1987.

[Fle92] Margaret M. Fleck. Plus ca change, ...*. In *Proc. Second European Conference on Computer Vision*, pages 151–159, 1992.

[Lee83] Jong-Sen Lee. Digital image smoothing and the sigma filter. *Computer Vision, Graphics and Image Processing*, 24:255–269, 1983.

[NM79] Makoto Nagao and Takashi Matsuyama. Edge preserving smoothing. *Computer Graphics and Image Processing*, 9:394–407, 1979.

[PM90] Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.

[SMCM91] Phillipe Saint-Marc, Jer-Sen Chen, and Gerard Medioni. Adaptive smoothing: A general tool for early vision. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(6):514–529, 1991.

[SVR80] Ann Scher, Flavo R. Davis Velasco, and Aziel Rosenfeld. some new image smoothing techniques. *IEEE Transactions on Systes, Man and Cybernetics*, 10(3):153–158, 1980.

[Weu94] Jörg Weule. *Iteration nichtlinearer Gauß-Filter in der Bildverarbeitung*. Dissertation (Thesis), Heinrich-Heine-Universität Düsseldorf, 1994.

[Whi93] Ross. T. Whitaker. Geometry-limited diffusion in the characterization of geometric patches in images. *Computer Vision, Graphics and Image Processing*, 57(1):111–120, 1993.

[WP93] Ross. T. Whitaker and Stephan M. Piter. A multi-scale approach to nonuniform diffusion. *Computer Vision, Graphics and Image Processing*, 57(1):99–110, 1993.

[WVL81] David C. C. Wang, Anthony H. Vagnucci, and C. C. Li. Gradient inverse weighted smoothing scheme and the evaluation of its performance. *Computer Graphics and Image Processing*, 15:167–181, 1981.

Lena

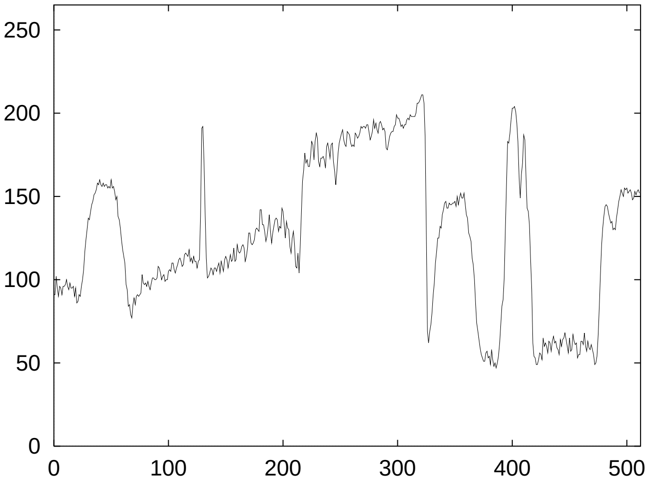

line 112

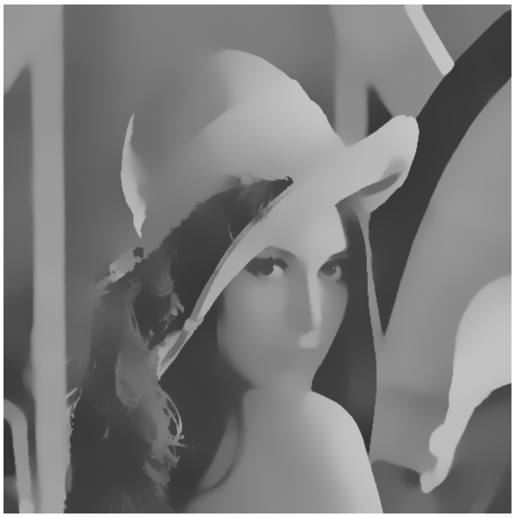

Output after 3 filter steps

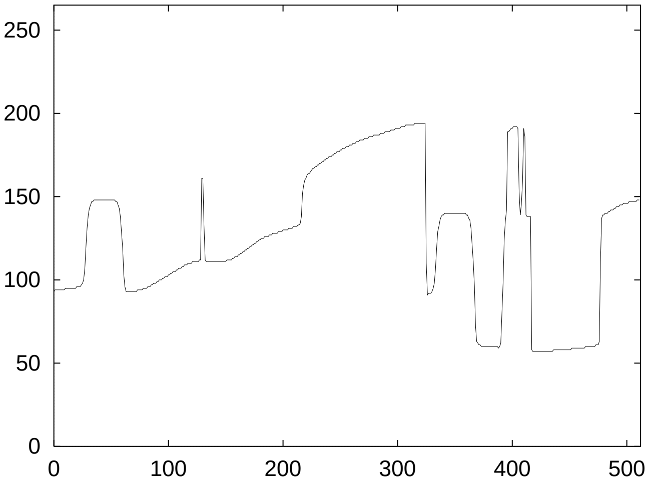

line 112

Output with modified last filter steps

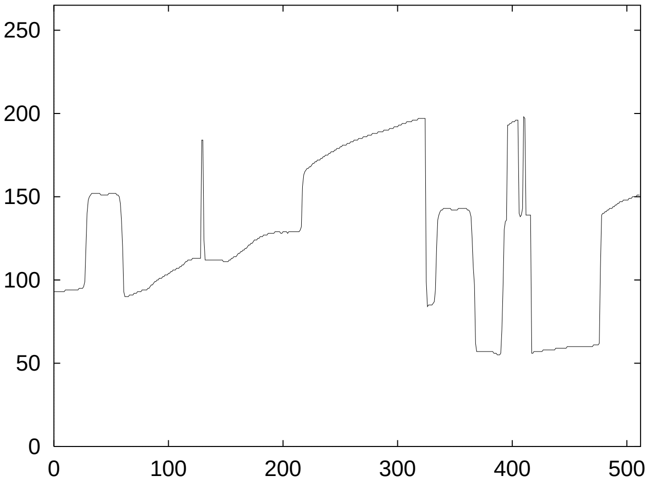

line 112

Figure 6: