# Legibility of Condensed Perceptually-Tuned Grayscale Fonts[1]

R. A. Morris[1], R. D. Hersch[2] and A. Coimbra

[1] University of Massachusets, Boston, USA
[2]Ecole Polytechnique Fédérale, Lausanne, Switzerland
ram@cs.umb.edu, hersch@di.epfl.ch

**Abstract.** We analyze the quality of condensed text on LCD displays, generated with unhinted and hinted bilevel characters, with traditional anti-aliased and with perceptually-tuned grayscale characters. Hinted bi-level characters and perceptually-tuned grayscale characters improve the quality of displayed small size characters (8pt, 6pt) up to a line condensation factor of 80%. At higher condensation factors, the text becomes partly illegible. In such situations, traditional anti-aliased grayscale character seems to be the most robust variant. We explore the utility of perceptually-tuned grayscale fonts for improving the legibility of condensed text. A small advantage was found for text searching, compared to bilevel fonts. This advantage is consistent with human vision models applied to reading.

## 1 Introduction

Information is increasingly accessed electronically, either from the World-Wide Web or from CD-ROMs. Often, printable text in book format or in A4 format is directly read from CRT or LCD displays. The most wide-spread document visualization tool, Adobe Acrobat, offers advanced zooming and scrolling functions in order to read printable text at an enlarged size. Nevertheless, when browsing through a certain number of documents, users may first start reading the document at the original size, in order to decide if that document is of interest to them. In that situation, the legibility of very small characters is of high importance.

For very small text on a video display, few pixels are available to represent the characters. The precise placement of these pixels based on the outlines of the font is sometimes improved by additional information ("hints") in the font offering the rendering algorithm advice about the compromises it must make. Furthermore, there is no need to restrict the intensity of the pixels to either black or white ("bilevel rendering"). We can also allow intermediate shades of gray ("grayscale rendering"). In this case, the human visual system may use this additional information to compensate for the inadequate visual data due to the low resolution.

---

[1] The grayscale figures described in the text can be downloaded from the Web and visualized on a display device (http://diwww.epfl.ch/w3lsp/pub/papers/grayscale/index.html).

To be easily readable, i.e to avoid horizontal scrolling at each line, the full text width of a document needs to be displayable on a given CRT or LCD device. In contrast, thanks to vertical scrolling, only a portion of the document page needs to be displayed at one time. Horizontal condensation of fonts is therefore important to be able to fit long text lines, for example printable text lines, within the limited horizontal space and resolution offered by most display devices.

We are therefore interested in the synthesis and legibility of horizontally condensed display fonts. Both bilevel and grayscale display fonts can be generated with or without grid-fitting rules (hints). We therefore present experiments where fonts are condensed by different methods up to the point where they are not legible any more. Condensed text samples are shown, where horizontal condensation results in text lines up to 50% shorter than the original text lines.

We ask ourselves to what extent perceptually-tuned grayscale fonts improve the quality and the legibility of condensed text. Therefore, we discuss the results of reading experiments made with subjects asked to detect specific strings within texts made of scrambled words, both with perceptually tuned grayscale fonts (see section 2) and with bilevel fonts, and at different condensation factors.

We limited our investigations to LCD devices because they have sharper spots than CRTs whose spot blur can confound conclusions about the chosen character generation method. Many argue that due to the difference in spot function, grayscaling is expected to be less effective on CRT than on LCD displays.

The fonts are produced according to 4 different methods[2]

1. Bilevel fonts *without* hinting, but with dropout control, i.e. with a special rasterization algorithm ensuring that character elements (stems, round parts) remain continuous even at very small size [2].
2. Bi-level fonts *with* hinting and with dropout control [2].
3. Traditional grayscale fonts, also called anti-aliased or unhinted grayscale fonts, without grid-fitting (hinting), produced by simple low-pass filtering and resampling of high-resolution master characters [3].
4. Perceptually-tuned grayscale fonts [1], also called hinted grayscale fonts, which are obtained by grid-fitting the character, so as to keep high contrast transitions in stems and bars and to avoid aliasing effects in round parts and terminals (see section 2).

Condensed instances of fonts are obtained by first applying to the original outline font a linear horizontal scaling factor , and then by rasterizing the resulting condensed outline according to the chosen font generation method.

In section 1, we briefly describe the difference between traditional and perceptually-tuned grayscale fonts. In section 2, we describe strategies for font condensation and show text condensed with different condensation factors. In section 3 we give an overview of the results of previous research in respect to the legibility of grayscale fonts. In section 4 we describe the validation experiment, give a brief overview of the human
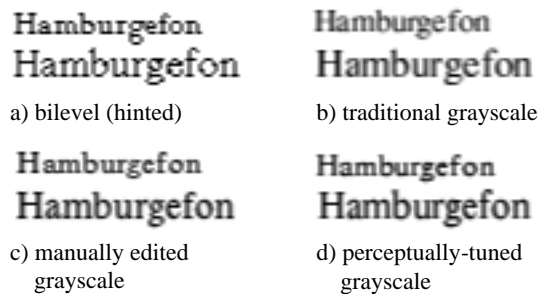
---

[2] Hand-tuned grayscale pixmap fonts have not been used in the present study. Without condensation, their quality is only slightly better than the one of automatically generated perceptually-tuned grayscale fonts [1].

vision research which supports the utility of grayscale fonts on visual grounds, and suggest further research directions. Section 5 summarizes the results.

## 2 Perceptually-tuned grayscale fonts

In order to improve the display of text on limited resolution computer displays, researchers have tried to trade-off the lack of spatial resolution for an increased number of intensity levels [3, 4 and 5]. Especially with respect to typographic text display on resolution limited CRT and LCD display devices, it has been shown that using grayscale makes font dependent character features visible which disappear when displaying text with bilevel characters (Fig. 1a). Though they can more accurately render font differences, traditional grayscale characters generated by filtering and resampling high-resolution bilevel master characters look rather fuzzy (Fig 1b).

Hamburgefon
Hamburgefon

a) bilevel (hinted)

Hamburgefon
Hamburgefon

b) traditional grayscale

Hamburgefon
Hamburgefon

c) manually edited
grayscale

Hamburgefon
Hamburgefon

d) perceptually-tuned
grayscale

**Fig. 1.**  Bilevel and traditional grayscale characters generated by filtering and resampling compared with manually edited and perceptually-tuned grayscale characters

Different instances of similar character structure elements such as vertical bars or curved stems do not have similar intensity profiles (Fig. 1b). This is due to the fact that the original bilevel master character incorporates frequencies well beyond the Nyquist limit (1/2 the resampling frequency) and that LCD displays are capable of displaying individual pixels as constant intensity squares of unit size, thereby enabling high intensity gradients between neighbouring pixels. The phase of character elements with respect to the target grid influences their intensity profiles (Fig. 2a).

In order to automatically generate improved quality grayscale characters similar to the characters which would be designed manually by skilled typographic designers, we have created a software package which applies modifications to the outline of the characters according to information specifying character structure elements such as bars, curved stems and serifs [2]. These modifications ensure that the left edge of vertical bars have a sharp contrast and that similar bars and curved stems have similar intensity profiles [1]. Furthermore, we preserve the structure of characters by enlarging character parts (Fig. 2b) which would be too thin to be displayed correctly.

a)

enlarged
junction

b)

same phase
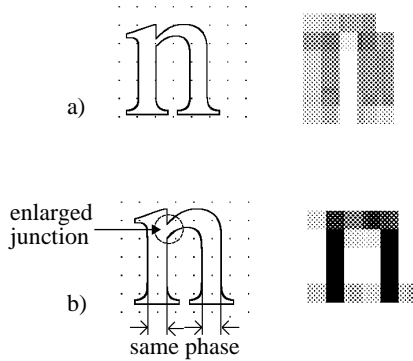
**Fig. 2.** The phase of the vertical bars defines their intensity profile: (a) without outline grid-fitting and (b) with outline grid-fitting

## 3  Strategies for text condensation

In the following analysis, we consider small size characters, i.e. characters which are difficult to visualize on existing CRT and LCD displays having a resolution between 72dpi and 120 dpi, i.e. characters at sizes 10 pt, 8 pt and 6 pt, possibly horizontally condensed by up to 50%.

There are several means to condense the length of a line of text. One may reduce the interword space by up to one pixel. With a two pixel reduction of the interword space, the separation between words tends to disappear. One may also try to reduce the intercharacter space by systematically rounding down the computed optimal visual space [1] between two characters.

Table 1 gives the line condensation factors obtained by reducing the interword space by one pixel and by rounding down the optimal spacing between successive characters as a function of character size.

**Table 1.** Line Condensation factors obtained by reducing the interword spacing or the intercharacter spacing

|  | interword spacing reduced by 1 pixel | intercharacter spacing rounded down |
|---|---|---|
| 10 pt bilevel | 97% | no effect |
| 10 pt grayscale | 97.5% | 98.5% |
| 8 pt bilevel | 97% | no effect |
| 8 pt grayscale | 97% | 97% |
| 6 pt bilevel | 96% | no effect |
| 6pt grayscale | 96% | 98.7% |

The condensation factors obtained by interword or intercharacter space reductions are relatively small, in regard of the necessity of reducing the length of display text lines by up to 50%. We therefore focus our attention on the horizontal scaling of individual characters. The method consists of applying to the original character outline scaling factors of 98%, 96%, 94%, etc.., to generate the target character set according to the selected character generation method (unhinted or hinted bilevel, traditional or perceptually-tuned grayscale), and to measure the resulting mean line length condensation factors. Word and intercharacter spaces are not reduced. Figure 3 gives the line condensation factors obtained for different horizontal character outline scaling factors, for 6pt, 8pt and 10pt hinted bilevel and perceptually-tuned grayscale fonts.
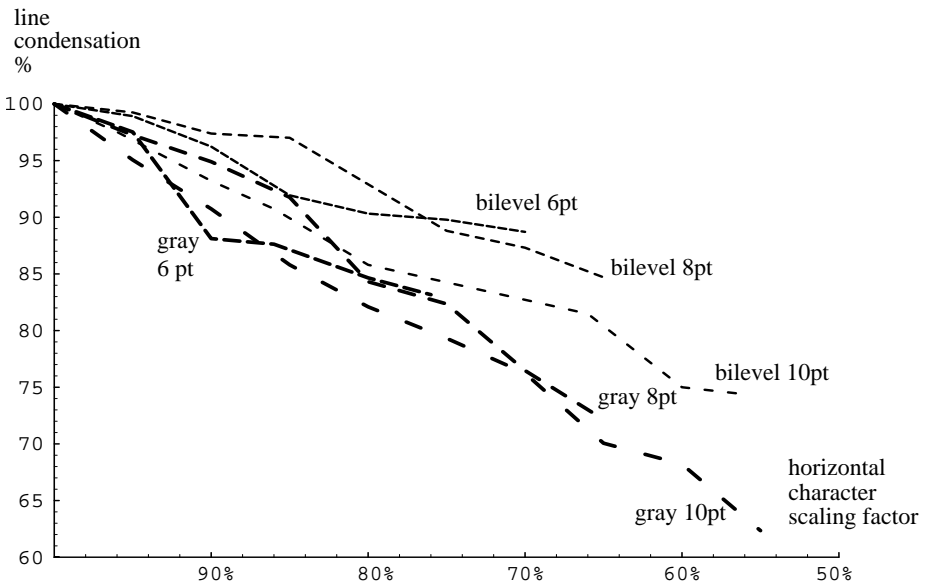


**Fig. 3.** Graph showing horizontal text line condensation as a function of the horizontal character outline scaling factor.

It is interesting to compare the degradation of character quality at line condensations of 90%, 80%, 70% 60% and 50%, as a function of the selected character generation method (unhinted versus hinted bilevel, traditional versus perceptually-tuned grayscale). The appendix shows the results obtained when generating condensed characters having a capital height of 8 pixels, i.e. for 9pt characters on a 95 dpi or 8pt characters on a 108 dpi display. At a line length reduction of up to 80% text remains legible and perceptually tuned grayscale seems to look better than traditional grayscale, which looks somehow fuzzy. However at a line condensation factor of 70% and at small size (8 pixels capital height), hinted bilevel fonts start to be less legible than unhinted bilevel fonts. At a line condensation factor of 60%, hinted bilevel is completety illegible, whereas unhinted bilevel, unhinted and hinted grayscale are still legible. At a line con-

densation factor of 50%, hinted bilevel is completely illegible, unhinted bilevel and hinted grayscale are partly illegible and unhinted grayscale is to some extent legible.

At a capital size of 10 pixels, similar effects are apparent, but at a slightly higher condensation. For example, at a line condensation factor of 40%, only unhinted grayscale text remains partly legible, whereas all other variants are completely illegible.

## 4  Previous comparative grayscale versus bilevel fonts legibility studies

A previous study comparing the legibility of traditional grayscale (unhinted grayscale) characters with that of bilevel characters showed that there was no significant difference in reading speeds between the two typographies at 10 point size [6], but that font identification was much easier with grayscale fonts than with bilevel fonts. Regarding esthetic judgement, there was no preference for the grayscale, compared with the bilevel typography.

A more recent study [7] was conducted in order to compare the legibility of perceptually-tuned grayscale and bilevel display fonts at small and very small sizes. This study confirmed the results of previous studies indicating that reading speed is to a large extent independent of typography (bilevel or grayscale) and font size. Perceptually-tuned grayscale characters performed much better than bilevel characters for an italic string search task in a meaningless text. Regarding the subjective preferences of the test subjects, perceptually-tuned grayscale fonts at 8 and 10 point sizes were given a clearly superior rating compared to bilevel fonts at these sizes. This result is to be compared to the study conducted by Black and Boag [6], where there was no user preference for traditional grayscale fonts versus bilevel fonts at 10 point size. This difference in user preference may be attributed to the improved visual quality of perceptually-tuned grayscale fonts.
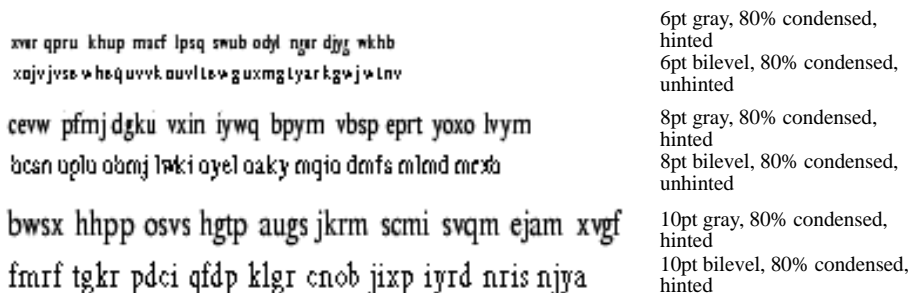
## 5  Experimental results

The results presented in this section show that the perceptually-tuned grayscaling technique offers an advantage in search speed over screen display using bilevel fonts. For small caps height, that advantage is somewhat larger with increased condensation, except at an extreme condensation, where the best bilevel either could not be read at all (6 pixel caps height 60% condensation) or had an insignificant advantage over the grayscale (8 pixel caps height, 60%).  Table 2 shows the ratio of the time for searching with binary fonts compared to searching with grayscale fonts in the same typographic condition (size and condensation).

This section describes the experiment in detail, the statistical analysis leading to the conclusions, and the human vision research which suggests why grayscaling might be expected to facilitate the search task. The search task experiment consisted in finding a specific four character string of randomly chosen letters among lines of 10 distractors similarly constructed. All text was lower-case. The subject was presented with a "paragraph" of 10 such lines and asked to either click on the target (which was always guar-

anteed to be among the distractors) or advance to the next paragraph if she wished to abandon the search in the current one. The target was the same string in each of ten such paragraphs. It is positioned in each paragraph at a random word location (equal positioning probability among 100 possible locations). A software timer was started when the subject initiated display of the first paragraph, and ended when 8 of the possible 10 occurrences of the target were found. The subject received a visual cue when the correct identification was made and could return to previous paragraphs if she wished. In each set of 10 paragraphs, the typographic and filtering conditions were kept constant. The measure of performance was the time taken to find 8 of the 10 targets.

Times-Roman fonts were used with three different sizes (10-, 8-, and 6-pixel caps height) and three different typographic conditions (normal setting, and condensation to 80% and 60% as described above). For bilevel fonts we chose whichever of the hinted or unhinted we deemed more readable. This was generally the hinted font at all but the smalles sizes and highest condensations. Figure 4 shows one line of randomly generated text at condensation factor 80%, at sizes 6pt, 8pt and 10pt, perceptually-tuned gray and bilevel.



6pt gray, 80% condensed, hinted

6pt bilevel, 80% condensed, unhinted

8pt gray, 80% condensed, hinted

8pt bilevel, 80% condensed, unhinted

10pt gray, 80% condensed, hinted

10pt bilevel, 80% condensed, hinted

**Fig. 4.** Randomly generated text strings at 80% condensation, perceptually-tuned gray and bilevel.

Subjects viewed the text on an IBM Thinkpad 560 portable computer with an active matrix LCD display whose resolution is 82 pixels/inch. Viewing was under normal room illumination with subjects selecting their own viewing distance. All subjects had normal or corrected to normal vision.

The data was analyzed by applying a t-test to the logarithm of the ratio of grayscale performance to binary font performance in each size/spacing condition, the data being taken over 10 subjects. Exponentiating the resulting confidence interval provides an estimate for the performance advantage of the task with grayscale fonts to that with binary fonts of the same size and spacing. See [8] for arguments for transforming confidence intervals by concave functions. Table 2 shows the results.

The low-pass filtering algorithm used for the generation of the perceptually-tuned grayscale characters was designed to overcome display limitations caused by low resolution. In particular, it removed high-frequency contrast energy in asymmetric ways (a high black-white contrast was kept at the left edges of character stems) and no attempt was made to control the overall contrast of the resulting stimuli. In addition, subjects

were permitted to choose their own viewing distance and and instructed but not required to keep that distance fixed.

**Table 2.** Size is caps height in pixels. B/G is an estimate of the ratio of search time for bilevel with time for perceptually tuned fonts. p is the probability that the inequality for B/G is in error. Details are described in the text.

| Size | Condensation | B/G | p |
|------|--------------|-----|---|
| 10 | 100% | >1.20 | <.046 |
| 10 | 80% | >1.02 | <.047 |
| 10 | 60% | >1.02 | <.047 |
| | | | |
| 8 | 100% | >1.04 | <.052 |
| 8 | 80% | >1.05 | <.028 |
| | | | |
| 6 | 100% | >1.05 | <.036 |
| 6 | 80% | >1.2 | <.038 |

In the display engineering community it is common wisdom that anti-aliased (i.e. low-pass filtered) fonts overcome limitations of display resolution but recent studies among dyslexic subjects give reason to suspect that low-pass filtering may also enhance readability due to properties of the human visual system. We briefly summarize this argument, which is surveyed in detail in Chapters 9-11 of [9]. First we outline the vision theory on which this rests.

It is now widely agreed that human visual processing proceeds through a number of visual "channels", collections of neurons which are tuned for particular orientation and sizes of features. The channels which process small features (that is, high spatial frequencies) form the "Parvocellular" system (Parvo- or P-system for short) and those processing larger features form the Magnocellular (Magno- or M-system). Magno neurons have large receptive fields. That is, the set of points in visual space to which a given Magno neuron responds is large, whereas Parvo neurons have small receptive fields. Thus, the low-pass filtering which comprises anti-aliasing removes or reduces information which stimulates only Parvo system cells. The visual effect is to blur edges and small features, to the point where neither their exact location or in some cases even their shape can not be accurately perceived. Such filtering, however, has no impact on larger features, such as the general shape of letters and words. For reading and related visual tasks, another important difference between these two subsystems is that they have radically different temporal response and radically different sensitivity to small luminance differences. The Magno system responds quickly to sudden *changes* in luminance at a given visual position, such as that which occurs during reading when a saccade brings a new word into the center of view. Magno cells respond poorly to unchanging stimuli. For this reason, the Magno system is also referred to as the transient system. The Parvo system responds mainly to sustained stimuli and so would be

stimulated by the fixations (typically lasting about 250 ms.) which follow saccades during reading. After a change in stimulus at the center of vision, there is a 10-20 ms latency before any parvo processing takes place, at which time the parvo system tends to inhibit the magno system. There are a number of other properties of these two subsystems which permit their responses to be measured by psychophysical and electrophysiological means in humans.

Lovegrove and others [9] have found that reading disabled subjects have magnosystem deficits and have proposed a model for reading which accounts for the importance of the magno system in ordinary reading in two ways. First it is important in processing peripheral information for saccade planning. (Peripheral vision has more magno input than parvo input: you can't see very small things in your periphery, but larger, especially moving, stimuli are easily detected there.) Secondly, the transient magno system is the first stimulated in the acquisition of letter and words in a top-down way beginning with their gross shapes and proceeding to their fine detail. (This "course-to-fine" recognition strategy is common in vision models, not just reading models.). If the transient magno system is weak, Lovegrave holds, the normal inhibition of the magno system by the parvo system will tend to prevent the initiation of the course-to-fine pattern recognition needed to recover the word being viewed. In particular, both for normal and disabled readers, reducing or delaying the parvo input might facilitate the magno initiation of the recognition. This reduction is exactly what is accomplished by the low-pass filtering which comprises anti-aliasing. Indeed, Williams et al. [10] found that poor readers have visual search time improved by low-pass filtering more than normal readers. The task of Williams et al. is very similar to ours, but used targets and distractors comprising only upper case letters and also did not manipulate the letter size and condensation, and hence they kept the underlying parvo input constant except for the filtering operation. Thus, their experiment would correspond to a fixed size/no condensation bilevel vs. traditional low-pass filtered grayscale test for us. In summary, removing high spatial frequency information may be expected to permit more magno processing thereby permitting quicker recognition of letters and words at the expense of precise location of their edges and of the details of their small features. Our results and those of Williams et al. are consistent with this for searching for nonsense words in nonsense distractors.

These search tasks might not generalize to reading tasks, which is the subject of future research. However, recently, Vitu et al.[11] have shown that eye-movements, at least, are similar in such a task and in reading.

*Future Research*. The algorithm described in this paper reduces the high-frequency image energy selectively: it keeps information which may have relevance for character identification, such as one sharp edge at the boundaries of horizontal and vertical bars. It could therefore, in principle, represent a visual compromise: the important remaining sharp features might facilitate character recognition---a precursor to word recognition known to be carried out in parallel for the entire word and so not requiring any shift of attention---but suppressing unnecessary high frequencies and thereby facilitating magno-system attention shift mechanisms. Therefore, an important followup study we will make will be to compare letter recognition performance for binary fonts, perceptu-

ally tuned grayscale fonts, and standard low-pass filtering. A second line of inquiry will compare perceptually tuned fonts to traditional low-pass filtered fonts in reading tasks.

## 6  Conclusions

Our experiments in condensed text generation show that hinted bi-level characters and perceptually-tuned grayscale characters improve the quality of displayed small size characters (8pt, 6pt) up to a line condensation factor of 80%. At higher condensation factors, the text becomes partly illegible. In such situations, traditional anti-aliased grayscale characters seem to be the most robust variant.

Perceptually-tuned grayscale fonts offer a small advantage over binary fonts when searching in text which is condensed, suggesting that such rendering may be advantageous for Web pages or other presentation where large amounts of text must be searched visually. The perceptually tuned algorithm leaves sharp features which are important to character identification and location, and removes features which are not. The advantage is suggested by human vision models, but it does not seem to accrue at very small caps height (in pixels) and extreme condensation due to insufficient sampling resolution for the representation of a character.

## References

1.  R. D. Hersch, C. Bétrisey, J. Bur, A. Gürtler, "Perceptually-Tuned Generation of Grayscale Fonts", IEEE Computer Graphics and Applications, 15 (6), 1995, 78–89.

2.  R. D. Hersch, "Font Rasterization, the State of the Art", Visual and Technical Aspects of Type, (R.D. Hersch, Ed.), Cambridge University Press, 1993, 78-109.

3.  J. Warnock, "The Display of Characters Using Gray Level Sample Arrays", Computer Graphics (Proc. Siggraph), 14 (3), 1980, 302–307.

4.  F. Crow, "The Use of Grayscale for Improved Raster Display of Vectors and Characters", Computer Graphics (Proc. Siggraph), 12 (3), 1978, 1–6.

5.  A. Naiman, A. Fournier, "Rectangular Convolution for Fast Filtering of Characters", Computer Graphics (Proc. Siggraph), 21 (4), 1987, 233–242.

6.  A. Black, A. Boag, "Choosing Binary or Greyscale Bitmaps: some Consequences for Users", C. Vanoirbeek & G. Coray (Eds.)  Proceedings Electronic Publishing, Cambridge University Press, 1992, 247–260.

7.  K. O'Regan, N. Bismuth, R. D. Hersch, A. Pappas, "Legibility of Perceptually-Tuned Grayscale Fonts", Proc. IEEE Int. Conf. Image Processing, (Ed. P. Delogne), Vol 1, 1996, 537-540.

8.  D. S. Mitrinovic, Elementary Inequalities, Noordhoff, Groningen, 1964.

9.  R. P. Garzia, ed. Vision and Reading, Mosby-Year Book, 1996.

10. M. C. Williams, J. G. May, R. Solman, H. Zhou, The effects of spatial filtering and contrast reduction on visual search times in good and poor readers, Vis Res. 35, 1995, 285-291.

11. F. Vitu, K. O'Regan, A. W. Inhoff and R. Topolski, "Mindless reading: eye-movement characteristics are similar in scanning letter strings and reading texts", Percept. Psychophys., 57(3), Apr. 1995, 352-36.

## Acknowledgement

# Appendix

## Capital letter height: 8 pixels

### 100 %

**Unhinted bilevel**

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler
and die to become one of the world's favor
composers. Shortly after his birth, his fat
distillor, moved the family to Iglau, Bohe

**Traditional grayscale**

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler wo
and die to become one of the world's favorite
composers. Shortly after his birth, his father
distillor, moved the family to Iglau, Bohemia

**Hinted bilevel**

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler
and die to become one of the world's favc
composers. Shortly after his birth, his fa
distillor, moved the family to Iglau, Boh

**Perceptually-tuned grayscale**

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler v
and die to become one of the world's favorit
composers. Shortly after his birth, his fathe
distillor, moved the family to Iglau, Bohem

### 80 %

**Unhinted bilevel   81.56%**

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the world's favorite
composers. Shortly after his birth, his father, a Jew
distillor, moved the family to Iglau, Bohemia.

**Traditional grayscale   79.45%**

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the world's favorite
composers. Shortly after his birth, his father, a Jewish
distillor, moved the family to Iglau, Bohemia.

**Hinted bilevel   80.58%**

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would liv
and die to become one of the world's favorite
composers. Shortly after his birth, his father, a Jew
distillor, moved the family to Iglau, Bohemia.

**Perceptually-tuned grayscale   79.09%**

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the world's favorite
composers. Shortly after his birth, his father, a Jewish
distillor, moved the family to Iglau, Bohemia.

### 70 %

**Unhinted bilevel   70.47%**

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the world's favorite
composers. Shortly after his birth, his father, a Jewish
distillor, moved the family to Iglau, Bohemia.

**Traditional grayscale   71.44%**

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the world's favorite
composers. Shortly after his birth, his father, a Jewish
distillor, moved the family to Iglau, Bohemia.

**Hinted bilevel   69.81%**

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would liv
and die to become one of the world's favorite
composers. Shortly after his birth, his father, a Jewish
distillor, moved the family to Iglau, Bohemia.

**Perceptually-tuned grayscale   70.61%**

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the world's favorite
composers. Shortly after his birth, his father, a Jewish
distillor, moved the family to Iglau, Bohemia.

# Capital letter height: 8 pixels

**60 %**

### Unhinted bilevel  60.28%

Gustav Mahler ( 1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the world's favorite
composers. Shortly after his birth , his father, a Jewish
distiller, moved the family to Iglau, Bohemia.

### Traditional grayscale  60.95%

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the worlds favorite
composers. Shortly after his birth, his father, a Jewish
distiller, moved the family to Iglau, Bohemia.

### Hinted bilevel: 60.24%

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the world's favorite
composers. Shortly after his birth, his father, a Jewish
distiller, moved the family to Iglau, Bohemia.

### Perceptually-tuned grayscale  60.26%

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the worlds favorite
composers. Shortly after his birth, his father, a Jewish
distiller, moved the family to Iglau, Bohemia.

**50 %**

### Unhinted bilevel    51.06%

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the world's favorite
composers. Shortly after his birth, his father, a Jewish
distiller, moved the family to Iglau, Bohemia.

### Traditional grayscale  51.86%

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the worlds favorite
composers. Shortly after his birth, his father, a Jewish
distiller, moved the family to Iglau, Bohemia.

### Hinted bilevel: 57.09%

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the world's favorite
composers. Shortly after his birth, his father, a Jewish
distiller, moved the family to Iglau, Bohemia.

### Perceptually-tuned grayscale   48.70%

Gustav Mahler (1860 - 1911)
Born in Moravia in 1860, Gustav Mahler would live
and die to become one of the worlds favorite
composers. Shortly after his birth, his father, a Jewish
distiller, moved the family to Iglau, Bohemia.