https://doi.org/10.2352/ISSN.2169-2629.2017.32.20

Visibility of spatiotemporal noise in digital video

Tamara Seybold, Betina L. Koelln, Aynur Pasha, Harald Brendel;
Arnold & Richter Cine Technik, Türkenstraße 89, 80799 München, Germany

Abstract

Motivated by the visual appearance of spatially denoised video sequences, we study the visibility of dynamic (temporal) noise. We investigate the visibility of noise for different spatial frequency bands. We conduct a subjective test with 22 observers. Included are two types of test patterns in the test: static (spatial) noise patterns and dynamic (spatiotemporal) noise. Eight spatial frequency bands are used for each pattern type. We obtain two main results: First, the contrast sensitivity of spatially low-frequency noise is significantly higher with temporal variation. Second, the noise visibility also depends on the content of the image or video. As the noise is masked by the content of the image, it becomes less perceivable. As higher frame rate might be used in future, a second test was performed comparing 24fps and 48 fps. Results show that the noise visibility is very similar. The significant increase of visibility with the temporal variation of spatially low-frequency noise should be respected in the design of future video processing methods.

Introduction

Noise is an unavoidable effect in real-world captured video sequences. Especially in low-light conditions it can seriously reduce the visual quality of video data. With today's high resolution sensors, having a smaller pixel pitch the need for algorithmic solutions is prominent.

In this work, we study the visibility of noise. Our motivation originates from observations in motion picture denoising. An extensive set of denoising methods was developed to reduce the noise in image and video data after capture. Of course none of the methods can perfectly reconstruct the true image, an estimation error always remains. The challenge is to hide this estimation error as good as possible. We observed that this estimation error is especially severe, when a spatial image denoising method is applied to video data. The estimation error, or remaining noise, in this case is usually very low-frequency noise. This low-frequency noise is not visible in still images, due to the fall-off in the contrast sensitivity for low spatial frequencies.

If the images are part of a video sequence, this low-frequency noise reappears as flickering in the video sequences and this effect is very disturbing. As the noise in the low-frequency bands is very difficult to separate from real image content this is still a common quality issue in video denoising.

While temporal denoising methods can better reduce flickering, they can introduce new artifacts, especially motion artifacts, and they come with high computational cost, especially memory requirements are extremely high for the current high resolution data (4k and beyond).

To the best of our knowledge no study is available that gives details about the visibility of noise in video sequences depending on the spectral distribution of the digital video noise. Winkler and Süsstrunk[33] presented a subjective study examining noise visibility in still images. Besides white noise, they also used mid-frequency noise and high-frequency noise. The results show the lower noise visibility for the high-frequency noise compared to the mid-frequency noise, as it would be expected from the contrast sensitivity function. We expect that the same effect could have been shown for low-frequency noise, as the contrast sensitivity also falls off towards low frequencies. Unfortunately this has not been subject of study.

To obtain more detailed information about noise visibility, we study the visibility of noise for eight different frequency bands in this paper. We evaluate both the visibility of static noise, as it occurs in still images, and dynamic noise, as it is present in video data. We start with giving a short overview on the literature about spatial and temporal contrast sensitivity of human vision. We then describe our approach: We explain how we obtain the different noise patterns, which can be displayed on a standard monitor. The results of a test with 22 observers for a static noise pattern and a dynamic (spatiotemporal) noise are subsequently presented. We discuss and compare our results of this test. Subsequently, we present an additional test, allowing us to compare the noise visibility for a standard frame rate (24fps) and a higher frame rate (48fps). Finally we conclude in the last section.

Related Work

Early technical achievements like movies and discharge lamps provoked early experiments on temporal effects of human vision. Temporal contrast sensitivity and explaining the effects using Fourier analysis was already studied almost sixty years ago, by de Lange ([2, 3, 4, 5, 6]), Kelly ([15, 17]) and Roufs ([26, 27, 28, 29, 30]). An overview of the early experiments is given by Kelly in 1977 [19].

The first measurements for combined spatiotemporal sensitivity were presented by Robson in 1966 [25]. He determined the thresholds for four spatial frequencies and for four temporal frequencies. The lowest frequency was 1 Hz, which is considered to be equivalent to static results (Van Nes result's [31] indicate thresholds a little bit above 0 Hz). While the spatial CSF measured for static patterns shows a band-pass characteristic, the sensitivity function for the same spatial frequencies, measured with spatial patterns that are temporally varying, is a low pass. The spatiotemporal contrast sensitivity is hence not separable, it shows a clearly more complicated shape than could be obtained by the product of spatial and temporal CSFs and Kelly in 1966 mentioned effects not explainable by a separable model [16]. He measured the CSF for spatiotemporal stimuli (travelling waves) [20] and the results are similar to Robson's: for 2 Hz the band-pass shape holds, the frequencies 13.5 Hz, 17 Hz and 23 Hz suggest a low pass structure; Chromatic CSF curves in [18, 21].

The experiments for CSF measurements are conducted with sine waves. While this has become the standard procedure to measure CSF curves, we, in this paper, want to evaluate the sensitivity to noise of different spatial frequencies, as it can occur in video sequences. Thus, our test setup is not directly comparable to former CSF measurements.

Some studies examined the visibility of signals when noise is added [8, 9, 1]. However, for the application of video quality assessment and denoising evaluation, it is more important to study the visibility of the noise.

Most applications that use spatiotemporal CSF models still rely on this data, e.g. the perceptual quality metric by Winkler [32]; other perception-based metrics are designed for still images and therefore only rely on the spatial CSF , e.g. [23, 14] and the image difference metric by Perdersen [24]. Nadenau et. al included a masking model for perception-based image compression, however, the algorithm is designed for stilll images and not for video [22].

A masking experiment was conducted by Winkler and Ssstrunk; they measured noise visibility on 30 test images. The visibility of noise in natural images was evaluated for three types of noise: white noise and two bandpass noise patterns of medium and high frequency bands [33]. Their work provides details for noise visibility in still images, but this type of experiment is still missing for temporal noise.

Noise visibility test

We investigate the visibility of noise of different spatial frequency bands in still images and in video sequences (we use 24 fps video sequences). To that end, we conduct a subjective test. Since the content of the background image may have a significant impact on noise perception, we select a homogeneous grey sequence and a rotating rose sequence. Both are displayed in Fig.1. We include two types of test patterns in the test: static (spatial) noise patterns and dynamic (spatiotemporal) noise patterns. Eight spatial frequency bands were used for each pattern type.

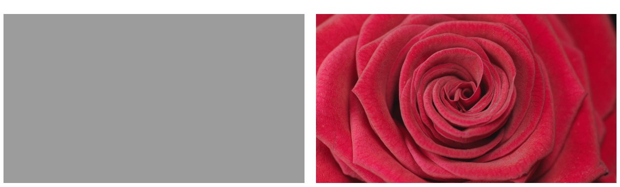

(a) Homogeneous grey image (b) One frame from the rotating flower sequence

Figure 1: Chosen sequences for the experiment.

Test pattern generation

The test patterns were obtained by first generating white noise and subsequent band-pass filtering, which is done by cutting the desired frequency band in the Fourier domain. To test the perception of luminance noise, we used the IPT color space and added the noise to the luminance channel (I channel). IPT is an opponent color space that was developed by Ebner and Fairchild [7] to create a space that is perceptually uniform. The transformation includes the monitor model (gamma transformation).

The workflow for the noise pattern generation is shown in Fig. 2 and described in the following.

1. First, 2-D zero-mean white noise is generated by Matlab's randn-function.
2. The noise is transformed to the Fourier domain and a band-pass filter is used to cut a defined band from the spectrum.
3. Each pixel row of the image is multiplied by a factor that increases logarithmically from top to bottom, to obtain noise with increasing variance (i.e. contrast).
4. A uniform grey image, with a constant pixel value of 0.6104 (in range [0-1]) is generated in sRGB space.
5. The grey image is transformed to the IPT space.
6. The zero-mean noise is finally added to one of the channels of the IPT image.
7. The noisy image is transformed back to sRGB.

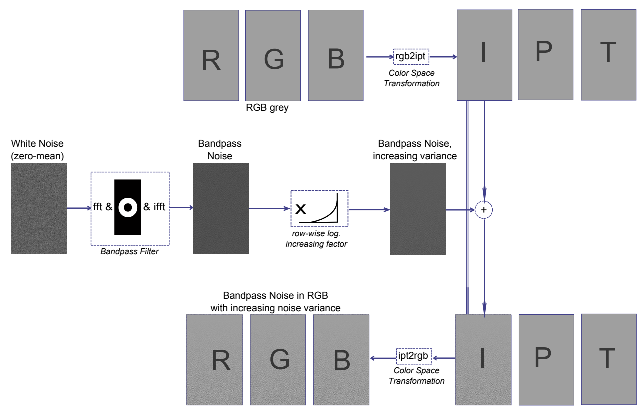

Figure 2: Workflow for generating the test stimuli for the noise sensitivity experiments.

The static noise patterns correspond to the noise in still images while the dynamic noise patterns correspond to the noise in video sequences. For the dynamic noise measurements we generated 360 static noise patterns, which are displayed as a sequence with 24 frames per second in the dynamic noise experiments.

For each frequency band one sequence is generated. We did the measurements for eight frequency bands between 0.14 and 27.62 cpd. Table 1 shows the frequency bands of the noise test patterns.

Table 1: The frequency bands used in the test are given in pixels per period and cycles per degree (cpd).

Nr. of pixels per sine period		Spatial frequency in cpd	
range	mean	range	mean
182.4 - 403.93	293.17	0.14 - 0.31	0.23
91.2 - 182.4	136.81	0.31 - 0.62	0.47
45.6 - 91.2	68.41	0.62 - 1.24	0.93
22.8 - 45.6	34.20	1.24 - 2.48	1.86
10.9 - 22.8	16.86	2.48 - 5.18	3.83
5.4 - 10.9	8.19	5.18 - 10.36	7.77
2.7 - 5.4	4.09	10.36 - 20.7	15.53
2.0 - 2.7	2.38	20.7 - 27.62	24.16

Test session

We conduct a subjective test to find out the noise level at which noise is visible in a digital video. The experiment is divided in four test parts:

- Uniform grey image sequence with static noise patterns.
- Static rose image sequence with static noise patterns.
- Uniform grey image sequence with dynamic noise patterns.
- Rotating rose image sequence with dynamic noise patterns.

The noise level in each sequence is gradually increasing for each frame and each sequence lasts 15 seconds. We include numbers into the frames corresponding to the duration from 1 to 15 seconds.

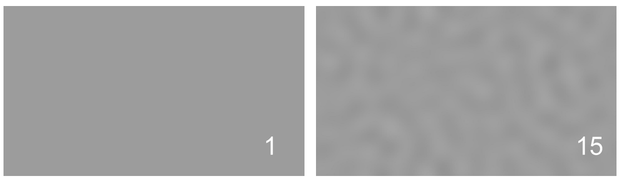

(a) First frame from a sequence (b) Last frame from a sequence

Figure 3: The visibility of noise on first and last step.

Before starting the experiment a short test for visual acuity and color blindness was conducted, for visual acuity the "tumbling E" chart and for color blindness the Ishihara plates were used. After this test, written and oral instructions were provided. Additionally one dummy sequence was presented as a demonstration and the distance between the observers and the screen was corrected. The participants of the test watch the sequences and observe at which second the noise first becomes visible. The observers were free to choose where to look at, because this is the realistic use case for video viewing.

Each sequence is repeated once to let the participant get a more precise observation. Between those repetitions there is a short break (three seconds) during which noiseless grey sequence with the text "Repeating..." is displayed. In order to avoid leaving an afterimage the text is placed in the bottom-right corner of image. In addition, there is another break lasting ten seconds before every new noise type, allowing the participants to write down their observation and preventing the noisy afterimage from damaging the vision of the next sequence. Furthermore we added a progress bar and the number of the next sequence into this break sequence to make sure the observer does not miss the beginning of the new noise type. Thus, for each frequency band the observer will watch a test sequence of duration 43 seconds (10s + 15s + 3s + 15s). There are eight frequency bands for static and dynamic noise types, a total of 32 test sequences.

According to the ITU-R recommendations for the subjective assessment of quality of television pictures [11], visual experiments should not last more than 30 minutes, since the experiment can be very tiring. In addition, at the beginning of the each test part about three training presentations should be introduced to stabilize the observer's opinion.

We introduce two training sequences at the beginning of each test part and each sequence is repeated once. Between the first play and the repetition there is a short break of three seconds. After the repetition there is a break lasting ten seconds before the next noise type, allowing the participants to have enough reaction time. Considering the training sequences one test part lasts about seven minutes, therefore the complete experiment lasts approximately 28 minutes.

Furthermore a random order was used for the presentations in order to eliminate contextual effects.

Test setup settings

The experiment was conducted in the video quality evaluation laboratory of the Institute for Data Processing at the Technical University of Munich in a room compliant with recommendation ITU-R BT.500 [11]. The tests where done with a professional broadcast monitor (Sony BVM-L230) set to ITU-R BT.709 color space. The monitor has a 10 bit serial-digital interface (SDI). The images used for the experiment had a resolution of 1920 x 1080 pixels.

According to the ITU-R recommendations BT.2022 [10] the distance between the screen and the observers should be about three times the picture height for full HD data in BT.709 [12] color space. The height of the reference display is 30 cm, therefore the distance used in the experiment was 90 cm.

Additionally the contrast of the display should be adjusted using a photometer. The display luminance value was 70 cd/m^2, according to the ITU-R BT.814 recommendation [13]. The ratio of background luminance behind the monitor to peak luminance of the picture was around 0.15, as recommended by ITU-R BT.500 [11].

Results

As the peak signal-to-noise ratio (PSNR) is a widely used metric in video applications, we use threshold PSNR levels to illustrate our results. We calculate the average PSNR value for each step (15 steps corresponds to 15 seconds). After that, we equate each observer's input for each sequence with the corresponding step's PSNR value.

Based on the inputs of 22 observers, we calculate the average PSNR and compare the visibility of static and dynamic noise on uniform and natural (content) images. We additionally plotted error bars indicating the standard deviation of the results.

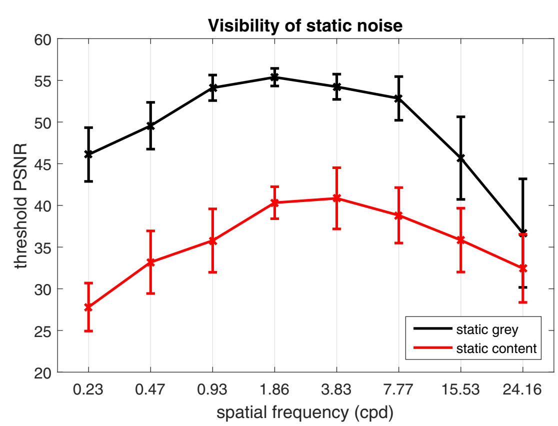

Figure 4: Result of the test with the static noise patterns, error bars show the standard deviation of the results.

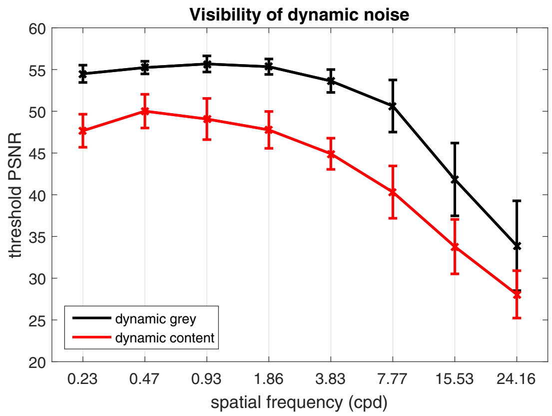

Figure 5: Result of the test with the dynamic noise, which is a sequence of newly generated static noise patterns with 24 frames per second. Error bars show the standard deviation of the results.

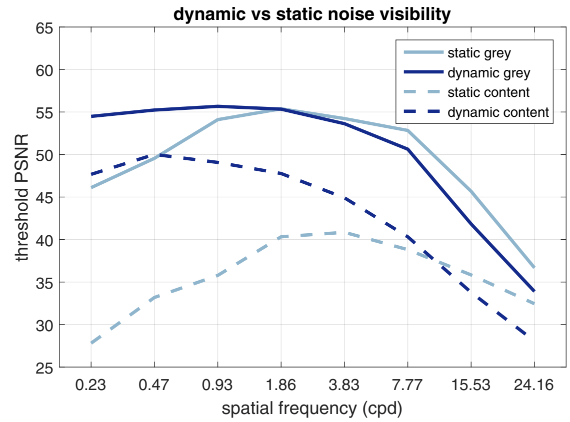

Figure 6: Comparison of test results for dynamic and static noise in grey and content sequences.

Fig. 4 shows the curves of the PSNR threshold for the eight spatial frequencies used in the experiment. The observed sensitivity in grey images shows a peak in the mid frequency 1.86 cpd and in content image in the mid frequency 3.83 cpd. Both curves show bandpass characteristics. There is, however, a significant difference between both curves, in natural (content) images the threshold level of noise visibility is much lower than in the homogeneous grey image sequences. This discrepancy between noise visibility in natural and uniform image sequences can also be observed in the results for visibility of dynamic noise.

The results for the dynamic noise, shown in Fig. 5, are significantly different from the static noise. Both curves show a rather lowpass shape than a bandpass. The peak is at lower spatial frequencies for both sensitivity curves. We see a smooth peak at 0.93 cpd for the grey sequences and a peak at 0.47 cpd for the flower sequences.

Discussion

Comparison of static and dynamic noise visibility

First, we will discuss the test results for static and dynamic noise visibility in grey and flower test parts.

The results for all four test parts are shown in Fig. 6. As already described the visibility curves of dynamic and static noise are significantly different, for the grey images as well as for the rose images. The threshold level of noise visibility for the content images is much lower compared to the plain grey images. This confirms that the content is masking the noise, this means that noise is significantly less perceivable depending on the background it is applied on. We conclude that noise is more perceivable in the uniform image sequences than in the natural image sequences.

Comparing the visibility of static and dynamic noise we observe that low-frequency noise is more perceivable when the noise is dynamic, as in video sequences. This difference is significant for the uniform image, but it is considerably larger in content image sequences. This means our results can be expected to be very relevant for real video processing applications.

Comparison of test results to the contrast sensitivity function

Due to the clear fall-off of the temporal contrast sensitivity above 10 Hz (Robson [25] and Kelly [20]), a flickering grating displayed at 24 Hz should lead to lower sensitivity than a static pattern. However, our results for the dynamic noise show the higher sensitivity, hence the contrary. In Fig. 7 and Fig. 8 we replotted our previously shown results in linear luminance to make them better comparable to other publications. The contrast is calculated using RMS contrast, i.e. the standard deviation.

First, we will discuss the results of the homogeneous grey test part in Fig. 7. In the low frequency range from 0.23 cpd to 0.93 cpd the sensitivity is higher for dynamic noise compared to static noise. This matches the severe differences we see in the visual quality of spatial denoising algorithms when they are applied to motion picture data compared to their application on still images. In the mid and high frequency bands from 1.86 cpd to 24.16 cpd the sensitivity to the dynamic noise is lower than to the spatial noise, but the difference is not large.

Whereas grey background allows study of noise visibility without influence of image content, the question of noise visibility in real world video sequences is even more relevant for practical applications.

Fig. 8 shows the contrast sensitivity in the XYZ color space for the flower test part. We observe that the absolute values of contrast sensitivity function in the flower images are lower compared to the contrast sensitivity function in the grey images.

In the low and mid-frequency range from 0.23 cpd to 7.77 cpd the sensitivity is clearly higher for dynamic noise compared to static noise. As for the grey sequences, this matches the severe differences in the visual quality of spatial denoising algorithms when they are applied to motion picture data compared to their application on still images. In the high frequency bands around 15.53 cpd and 24.16 cpd the sensitivity to the dynamic noise is lower than to the spatial noise, the difference gets larger for the highest spatial frequency.

The main result, that the low-frequency band noise is clearly more visible in the dynamic noise test, matches our visual impres-

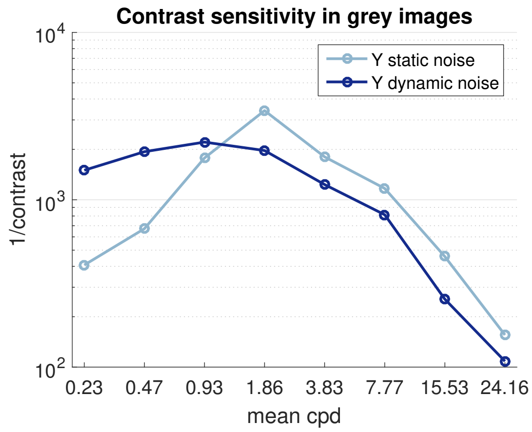

Figure 7: The contrast sensitivity was calculated in XYZ color space for better comparison with other results. Plotted here, is the contrast sensitivity of luminance noise (Y-channel) results of the static and dynamic noise patterns in grey images.

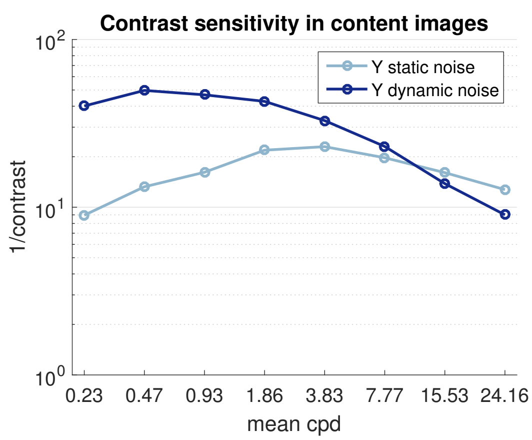

Figure 8: Constrast sensitivity of luminance noise (Y-channel) results of the static and dynamic noise patterns in flower images.

sion, which is that the dynamic noise intensity seems to increase towards lower spatial frequencies.

Our findings show the significant differences in noise visibility from still to video data, which up to now are not respected in video processing algorithms as denoising, or video quality assessment. While our findings show a tendency that might explain quality issues and help improving some of the image processing algorithms for video data, a detailed model would be needed to include precise noise visibility information for perception based video processing algorithms. This would require more tests and a more detailed study on spatiotemporal masking, which is beyond the scope of this paper.

Video frame rate

In future, motion-picture frame rates might be higher than 24 fps. We therefore performed an additional experiment evaluating the difference in noise visibility of 24 fps and 48 fps sequences.

We used the same test setup as described above. As the video frame rate could not be switched during the experiment, we displayed the complete test in 48 fps. The 24 fps content was simulated using a 48 fps sequence showing the exact same image twice. A different monitor had to be used to display 48 fps sequences in full HD resolution (EIZO CG318-4K). The monitor was calibrated before the experiment and the luminance levels were measured in order to meet the requirements described above. To reduce the length of the test, this time three spatial frequencies for the noise were selected (low, mid and high). 20 observers completed the test.

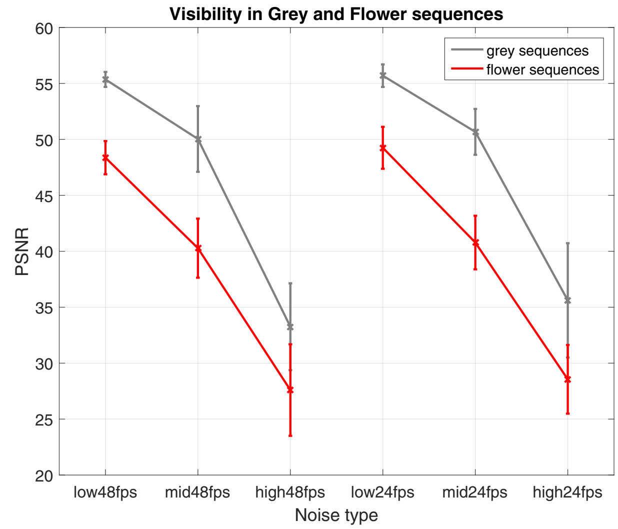

Figure 9: Comparison of test results for dynamic noise in grey and content sequences displayed with 24 fps and 48fps. Three frequency bands were used: "low" corresponds to 0.5 cpd, "mid" corresponds to 8.35 cpd and "high" corresponds to 26.25 cpd.

Fig. 9 shows the results. For all three frequencies the noise visibility is very similar for 48 fps and 24 fps sequences. The noise visibility hence does not decrease significantly with higher frame rates than 24 fps. As stated in the introduction, higher resolution of image sensors lead to higher noise, because less light is trapped by the sensor. As higher frame rates require shorter exposure time, this additionally reduces the light trapped by one pixel, and hence increases noise. Therefore we can conclude that noise will continue to limit video quality in high resolution and high frame rate video data shot in low light conditions.

Conclusion

We measured the visibility of noise for static (spatial) noise, which occurs in still images, and for dynamic (temporal) noise, which occurs in video data. We obtain three main results.

First, the contrast sensitivity of spatially low-frequency noise is significantly higher when the noise is temporally varying. This can explain why algorithms designed for still images might not show high quality results on video data, e.g. denoising algorithms for still images that do not eliminate low-frequency noise, because it is not visible in still images.

Second, we showed that the noise visibility strongly depends on the image or video content. Our results show that the noise is significantly more perceivable in uniform images than in natural images, which can be explained by masking. In addition to that, the above-mentioned difference between noise visibility of static and dynamic noise is considerably larger for our natural image content than for the grey image. That means, that the observed difference in noise visibility is important to consider for improving video processing algorithms.

Third, an additional experiment evaluated the influence of the frame rate on noise visibility by comparing 24 and 48 fps. The results show that the noise visibility does not decrease significantly for 48, compared to 24 fps. As stated in the introduction, higher resolution of image sensors lead to higher noise, because less light is trapped by the sensor. Higher frame rates require shorter exposure time, which also reduces the light trapped by one pixel, and hence increases noise. Therefore we can conclude that noise will continue to limit video quality in high resolution and high frame rate video data shot in low light conditions.

Further research on noise visibility is therefore very important to allow developing and improving perception-based video processing algorithms.

References

[1] K. T. Blackwell. The effect of white and filtered noise on contrast detection thresholds. *Vision Research*, 38(2):267–280, January 1998.

[2] H. De Lange Dzn. Experiments on flicker and some calculations on an electrical analogue of the foveal systems. *Physica*, 18(11):935–950, November 1952.

[3] H. De Lange Dzn. Relationship between Critical Flicker-Frequency and a Set of Low-Frequency Characteristics of the Eye. *J. Opt. Soc. Am.*, 44(5):380–388, May 1954.

[4] H. De Lange Dzn. Research into the Dynamic Nature of the Human Fovea -¿ Cortex Systems with Intermittent and Modulated Light. I. Attenuation Characteristics with White and Colored Light. *J. Opt. Soc. Am.*, 48(11):777–783, November 1958.

[5] H. De Lange Dzn. Research into the Dynamic Nature of the Human Fovea -¿ Cortex Systems with Intermittent and Modulated Light. II. Phase Shift in Brightness and Delay in Color Perception. *J. Opt. Soc. Am.*, 48(11):784–787, November 1958.

[6] H. De Lange Dzn. Eye's Response at Flicker Fusion to Square-Wave Modulation of a Test Field Surrounded by a Large Steady Field of Equal Mean Luminance. *J. Opt. Soc. Am.*, 51(4):415–421, April 1961.

[7] F. Ebner and M. D. Fairchild. Development and Testing of a Color Space (IPT) with Improved Hue Uniformity. *Color and Imaging Conference*, 1998(1):8–13, January 1998.

[8] R. E. Fredericksen and R. F. Hess. Temporal detection in human vision: dependence on stimulus energy. *J. Opt. Soc. Am. A*, 14(10):2557–2569, October 1997.

[9] R. E. Fredericksen and R. F. Hess. Estimating multiple temporal mechanisms in human vision. *Vision Research*, 38(7):1023–1040, April 1998.

[10] ITU-R Recommendation BT.2022. General viewing conditions for subjective assessment of quality of SDTV and HDTV television pictures on flat panel displays. Geneva, Switzerland, 2012. International Telecommunication Union.

[11] ITU-R Recommendation BT.500-13. Methodology for the subjective assessment of the quality of television pictures. Geneva, Switzerland, 2012. International Telecommunication Union.

[12] ITU-R Recommendation BT.709-5. Parameter values for the HDTV standards for production and international programme exchange. Geneva, Switzerland, 2009. International Telecommunication Union.

[13] ITU-R Recommendation BT.814-1. Specifications and alignment procedures for setting of brightness and contrast of displays. International Telecommunication Union, 1992-1994.

[14] G. Ivkovic and R. Sankar. An algorithm for image quality assessment. In *ICASSP '04*, volume 3, pages iii–713–16, May 2004.

[15] D. H. Kelly. Visual Responses to Time-Dependent Stimuli. I. Amplitude Sensitivity Measurements. *J. Opt. Soc. Am.*, 51(4):422–429, April 1961.

[16] D. H. Kelly. Frequency Doubling in Visual Responses. *J. Opt. Soc. Am.*, 56(11):1628–1632, November 1966.

[17] D. H. Kelly. Theory of Flicker and Transient Responses, II. Counterphase Gratings. *J. Opt. Soc. Am.*, 61(5):632–640, May 1971.

[18] D. H. Kelly. Spatio-temporal frequency characteristics of color-vision mechanisms. *J. Opt. Soc. Am.*, 64(7):983–990, July 1974.

[19] D. H. Kelly. Visual contrast sensitivity. *Journal of modern optics*, 24(2):107–129, 1977.

[20] D. H. Kelly. Motion and vision. II. Stabilized spatio-temporal threshold surface. *J. Opt. Soc. Am.*, 69(10):1340–1349, October 1979.

[21] D. H. Kelly. Spatiotemporal variation of chromatic and achromatic contrast thresholds. *J. Opt. Soc. Am.*, 73(6):742–749, June 1983.

[22] M. J. Nadenau, J. Reichel, and M. Kunt. Wavelet-based color image compression: exploiting the contrast sensitivity function. *IEEE Transactions on Image Processing*, 12(1):58–70, Jan 2003.

[23] Lszl Neumann, Kresimir Matkovic, and Werner Purgathofer. Perception based color image difference. *Computer Graphics Forum*, 17(3):233–241, 1998.

[24] Marius Pedersen. An image difference metric based on simulation of image detail visibility and total variation. In *Color and Imaging Conference*, pages 37–42. Society for Imaging Science and Technology, 2014.

[25] J. G. Robson. Spatial and temporal contrast sensitivity functions of the visual system. *Journal of the Optical Society of America*, 56(8):1141–1142, 1966.

[26] J. A. J. Roufs. Dynamic properties of vision–I. Experimental relationships between flicker and flash thresholds. *Vision Research*, 12(2):261–278, February 1972.

[27] J. A. J. Roufs. Dynamic properties of vision–II. Theoretical relationships between flicker and flash thresholds. *Vision Research*, 12(2):279–292, February 1972.

[28] J. A. J. Roufs. Dynamic properties of vision–III. Twin flashes, single flashes and flickerfusion. *Vision Research*, 13(2):309–323, February 1973.

[29] J. A. J. Roufs. Dynamic properties of vision–IV: Thresholds of decremental flashes, incremental flashes and doublets in relation to flicker fusion. *Vision Research*, 14(9):831–851, September 1974.

[30] J. A. J. Roufs. Dynamic properties of vision–V: Perception

lag and reaction time in relation to flicker and flash thresholds. *Vision Research*, 14(9):853–869, September 1974.

[31] F. L. Van Nes, J. J. Koenderine, H. Nas, and M. A. Bouman. Spatiotemporal Modulation Transfer in the Human Eye. *J. Opt. Soc. Am.*, 57(9):1082–1087, September 1967.

[32] S. Winkler. A Perceptual Distortion Metric for Digital Color Video. In *Electronic Imaging '99*, pages 175–184, San Jose, California, 1999.

[33] S. Winkler and S. Süsstrunk. Visibility of noise in natural images. In *Electronic Imaging 2004*, pages 121–129. International Society for Optics and Photonics, 2004.

