

Image-quality assessment for full-disk solar observations with generative adversarial networks[★]

R. Jarolim¹, A. M. Veronig^{1,2}, W. Pötzi², and T. Podladchikova³

¹ University of Graz, Institute of Physics, Universitätsplatz 5, 8010 Graz, Austria
e-mail: robert.jarolim@uni-graz.at

² University of Graz, Kanzelhöhe Observatory for Solar and Environmental Research, Kanzelhöhe 19, 9521 Treffen am Ossiacher See, Austria

³ Skolkovo Institute of Science and Technology, Bolshoy Boulevard 30, Bld. 1, Moscow 121205, Russia

Received 17 June 2020 / Accepted 24 August 2020

ABSTRACT

Context. In recent decades, solar physics has entered the era of big data and the amount of data being constantly produced from ground- and space-based observatories can no longer be purely analyzed by human observers.

Aims. In order to assure a stable series of recorded images of sufficient quality for further scientific analysis, an objective image-quality measure is required. Especially when dealing with ground-based observations, which are subject to varying seeing conditions and clouds, the quality assessment has to take multiple effects into account and provide information about the affected regions. The automatic and robust identification of quality-degrading effects is critical for maximizing the scientific return from the observations and to allow for event detections in real time. In this study, we develop a deep-learning method that is suited to identify anomalies and provide an image-quality assessment of solar full-disk H α filtergrams. The approach is based on the structural appearance and the true image distribution of high-quality observations.

Methods. We employ a neural network with an encoder–decoder architecture to perform an identity transformation of selected high-quality observations. The encoder network is used to achieve a compressed representation of the input data, which is reconstructed to the original by the decoder. We use adversarial training to recover truncated information based on the high-quality image distribution. When images of reduced quality are transformed, the reconstruction of unknown features (e.g., clouds, contrails, partial occultation) shows deviations from the original. This difference is used to quantify the quality of the observations and to identify the affected regions. In addition, we present an extension of this architecture that also uses low-quality samples in the training step. This approach takes characteristics of both quality domains into account, and improves the sensitivity for minor image-quality degradation.

Results. We apply our method to full-disk H α filtergrams from the Kanzelhöhe Observatory recorded during 2012–2019 and demonstrate its capability to perform a reliable image-quality assessment for various atmospheric conditions and instrumental effects. Our quality metric achieves an accuracy of 98.5% in distinguishing observations with quality-degrading effects from clear observations and provides a continuous quality measure which is in good agreement with the human perception.

Conclusions. The developed method is capable of providing a reliable image-quality assessment in real time, without the requirement of reference observations. Our approach has the potential for further application to similar astrophysical observations and requires only coarse manual labeling of a small data set.

Key words. atmospheric effects – techniques: image processing – methods: data analysis – Sun: chromosphere

1. Introduction

Modern solar observations are carried out in an autonomous way, covering multiple filter bands and producing a high-cadence output stream that needs to be accessible for monitoring and scientific use within minutes (Harvey et al. 1996; Pesnell et al. 2012; Pötzi et al. 2018). While continuous observation provides a clear benefit, such as for example the permanent monitoring of the Sun, automatic and robust methods are necessary to ensure image quality in the large data streams before they are passed on for further analysis. This is in particular necessary for ground-based observations, which are subject to varying seeing conditions, and are used for event detection (e.g., flares, filament eruptions) in real-time (Pötzi et al. 2015; Veronig & Pötzi 2016).

Ground-based observations confer the advantage that they can be upgraded and are cost efficient compared to space-based

observations, but limiting factors of atmospheric turbulence have to be overcome. Adaptive-optics systems (Rimmele & Marino 2011) and post-facto corrections (Löfdahl et al. 2007) are able to correct for local seeing conditions. For autonomous full-disk observations these methods cannot be applied because they rely to some extent on human supervision (Wöger et al. 2008; Löfdahl et al. 2007). In addition, the daily observation schedule, varying seeing conditions, and the possible presence of clouds lead to unavoidable gaps in the continuous observation series. Ground-based network telescopes can provide a continuous data stream and can mitigate the impact of local effects, as can be seen for example from operating ground-based networks for observation of solar- and stellar oscillation (GONG, Harvey et al. 1996 and SONG, Grundahl et al. 2006). As for future full-disk observation networks, such as the anticipated Solar Physics Research Integrated Network Group (SPRING), the data homogenization between sites poses new challenges because of the increased amount of data, improved instrumentation and automated

[★] Movies are available at <https://www.aanda.org>

algorithms that are sensitive to systematic deviations (Gosain et al. 2018). In order to provide the highest-quality data products, an objective quality assessment is required to remove low-quality observations from the data stream as well as to compare and select between simultaneous observations from different observing sites.

Image quality provides a critical parameter for filtering observation series and frame selection (Popowicz et al. 2017). Overly strong filtering may lead to avoidable gaps in the series, while overly weak filtering can reduce the quality of the data series. Automated detection methods rely on the high quality of the input data. Quality assessment before further processing is important to guarantee the validity of the detection results (Pötzi et al. 2018). Data-driven methods typically improve with the size of the data set, while erroneous samples lead to a performance decrease or even failure of the method (Galvez et al. 2019). The manual cross-checking of tens of thousands of data samples is tedious, prone to errors, and impossible to achieve in quasi real-time. Thus, the development of robust automated methods is essential.

Image-quality metrics (IQMs) have been addressed in several ways and can be categorized into three main groups according to the availability of a reference image. For full-reference IQMs, a distortion-free image exists, and therefore deviations from this image can be quantified. These IQMs range from simple pixel-based metrics, such as the mean-squared-error (MSE), to more advanced quality metrics such as the structural similarity index (SSIM), which shows good agreement with human perception (Wang et al. 2004). In cases where no additional information about a reference image is available, we refer to no-reference IQMs (also known as blind image-quality assessment). Several methods have been proposed for this problem (e.g., CORNIA, Ye et al. 2012 and BRISQUE, Mittal et al. 2012). When information about a reference image is available to some extent, for instance in the form of extracted features, the quality metric is referred to as a reduced-reference IQM (Wang et al. 2004).

In cases of solar observations, where the user intends to quantify the image quality for each observation, no full-reference image exists. Popowicz et al. (2017) reviewed various IQMs for high-resolution solar observations and provide a comparison of 36 different methods. A frequently used quality metric is the root-mean-square contrast, which has a dependence on the solar structure. More recent approaches aim to provide an objective IQM, such as the no-reference metric by Deng et al. (2015). Solar features show strong similarity, which allows for the use of reduced-reference IQMs. In Huang et al. (2019), such a metric has been proposed, and is based on the assumption of the multi-fractal property of solar images.

However, for solar full-disk observations, the problem setting is different as here both global (e.g., large-scale clouds) and local (e.g., contrast, small clouds) effects play a role. Pötzi et al. (2015) developed and implemented an image-quality check for full-disk $H\alpha$ filtergrams as part of the observing and data processing pipeline at Kanzelhöhe Observatory. The method makes use of known properties of the solar images by quantifying the deviations from a circle as fitted to the solar limb, quantifying the large-scale intensity distribution in image quadrants, and estimating the image sharpness by computing the correlation with a blurred version of the original image. The weighting of these different parameters to obtain one combined image-quality parameter was determined empirically, and is thus to some degree subjective.

With the advent of deep-learning methods, two important components can be taken into account, (1) the structural

appearance of solar features and (2) the deviations from the true image distribution. While recent methods try to compare structural similarity over pixel-based estimations (Huang et al. 2019; Deng et al. 2015), to date there exists no IQM for solar observations that directly estimates deviations from the true image distribution. The stability of deep-learning methods relies on the variety and the quantity of available training samples. This means high performance when new data samples are within the domain of the training set, but lower performance if new samples deviate from the training data. In the case of full-disk solar observations, the detection of strong deviations from regular observations is particularly important. As strongly degraded observations are commonly removed from the data archives, a supervised approach cannot fully account for the diversity of regular observation series. Therefore, we use a different approach, applying unsupervised training methods.

Throughout this study, we categorize the full-disk images into three quality classes. Figure 1 shows a representative sample of each image-quality class. (1) We refer to high-quality observations if the image is not affected by clouds, is properly aligned to cover the full solar disk, and provides the sharpness that is attainable under good observing conditions at the observing site (Fig. 1, left panel). Such observations are well suited for scientific applications and for processing by automated algorithms. For the classification, we do not consider the content or scientific importance of the observation (e.g., presence of flares, filaments, active regions). We note that thin clouds often reduce atmospheric turbulence and can lead to exceptionally good image quality. In this study, we only refer to atmospheric effects as clouds if they are visually recognizable in the image. (2) We refer to low-quality observations when a degradation in image quality can be identified. This can be induced by the turbulent atmosphere or other atmospheric effects such as thin clouds or contrails (Fig. 1, middle panel). Low-quality observations can still be used for scientific analysis and visual inspection, but can lead to irregular behavior when applied to automated algorithms. (3) Observations which show strong degradation (e.g., thick clouds, partial occultation, instrumental misalignment) can typically not be used for scientific applications and are therefore removed from the data archives by existing algorithms (Pötzi et al. 2015). As these observations differ from regular observations available in the archive, we refer to them as anomalous (Fig. 1, right panel).

In this paper, we present a novel method for no-reference image-quality assessment for ground-based full-disk solar observations. We employ an unsupervised deep-learning approach which uses the true image distribution of high-quality observations to detect deviations from it. Our method provides an objective image-quality score and reliably detects anomalies in the data. Furthermore, our method can account for the identification of regions that are affected. The model training is performed with high-quality observations and requires no further reference image after training. In addition, we propose a classifier as an extension to our neural network architecture. This classifier also uses low-quality observations to provide increased sensitivity to minor quality degradation. We made our codes publicly available under¹.

2. Data set

In this study, we demonstrate our method for solar full-disk $H\alpha$ filtergrams from Kanzelhöhe Observatory for Solar and

¹ <https://github.com/RobertJaro/SolarImageQualityAssessment>

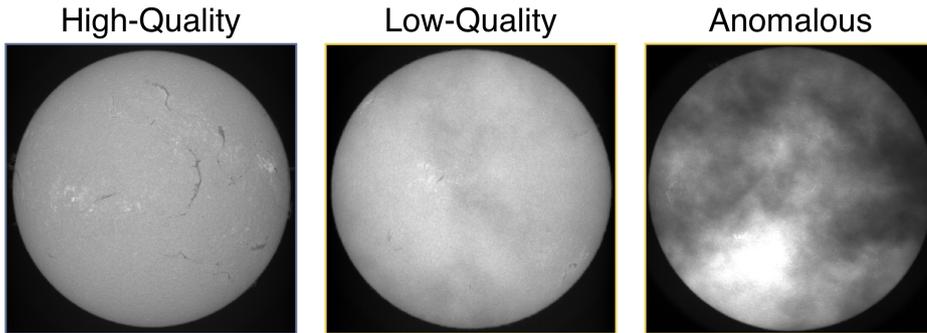


Fig. 1. Representative samples of the three different image-quality classes. High-quality observations are characterized by sharp structures and no degrading effects (*left panel*). Low-quality observations suffer from degrading effects or appear blurred (*middle panel*). Anomalous observations show strong atmospheric influences or instrumental errors which excludes them from further scientific analysis (*right panel*).

Environmental Research (KSO²). The KSO regularly takes H α images at a cadence of 6 s and provides a fully automated data reduction and provisioning, which allows for data access in near real time. The spatial resolution of the telescope is about 2", and the data are recorded by a 2024 \times 2024 pixel CCD corresponding to a sampling of about 1" per pixel (Otruba & Pötzi 2003; Pötzi et al. 2015). The quality assessment is provided by an automated algorithm which separates observations into three classes (Pötzi et al. 2018). Only the highest quality observations (class 1) are used in the pipeline of automated event detection, class 2 observations are still considered for scientific analysis and visual inspections, and the lowest quality observations (class 3) are completely removed from the archive. The image-quality assessment criteria and division into classes were determined empirically, and are specifically adapted to the KSO H α filtergrams. The method has not yet been systematically evaluated. In this study, we account for this evaluation by manually classifying a test set which we compare to the existing quality assignment as well as to the newly developed deep-learning approach that is presented here.

From the KSO data archive³, we randomly sampled solar full-disk H α filtergrams recorded between 2012 and 2019. Here, we alternated between observations labeled as class 1 and class 2. We manually separated high-quality images from observations that contain clouds or blurred solar features (low-quality) until we acquired 2000 images per quality class. Observations with strong quality degradation (anomalous) are sparse in the KSO archive and are not considered for training purposes. From this data set, we separate 1650 observations per quality class and keep them as an independent test set, which we do not use for any of our model training. The remaining 350 observations per quality class are used to automatically create the training set for the primary model (see Sect. 3.5).

Observations with strong degradations (e.g., strong cloud coverage, partial eclipsed observations, overexposure) are removed from the KSO archive (class 3). In order to assert the stability for unfiltered (raw) image time series, which have a larger variety of atmospheric and instrumental effects than the training and test set, we analyze data from an additional five full observing days with varying seeing conditions (2018-09-27 until 2018-09-30 and 2019-01-26). This set has not been pre-filtered with respect to image quality. From the total of 10 050 filtergrams, we manually label all images that show strong degradation. This leads to a total set of 620 images attributed to the “anomalous” class.

We note that our method is not restricted to a specific instrument, wavelength, or observation target and can be applied similarly to new data sets.

3. Method

Neural networks have shown impressive results for classification tasks (He et al. 2016; Simonyan & Zisserman 2014; Chollet 2017; LeCun et al. 2015). Applied to the case of quality classification, this could be solved as a two-class problem. In this setting, a data set of high- and low-quality observations needs to be manually labeled and the network is trained to predict the correct class for a given input image (supervised training). Convolutional neural networks (CNNs) have shown the capability to directly learn from images, by automatically extracting features from edges and shapes within the image (LeCun et al. 2015). This provides an advantage over classical machine-learning approaches that provide a classification based on manually extracted features. In the present case, clouds significantly differ in shape and structure from solar features, while extracted parameters such as the global intensity distribution primarily detect large-scale deviations and are prone to falsely identifying solar features. Even though a CNN can identify small changes in image quality, the classification approach faces two fundamental issues in order to provide a reliable quality assessment: (1) Learning-based algorithms show a high stability for data which are similar to the training set, but the coverage of all possible atmospheric and instrumental effects is not possible. Furthermore, data with strong degradations are often not stored in the archive (Pötzi et al. 2018). Therefore, there is typically a large number of high-quality observations available, while the variety of low-quality observations is sparse. For classification tasks, neural networks can produce unexpected results even for minor deviations from the training set distribution (Goodfellow et al. 2014a; Papernot et al. 2017). (2) The predictions of a neural network classifier are probabilistic scores and typically do not scale with the quality of the images. However, the adequate filtering of solar observations requires a proper quality metric and the identification of affected regions. The lack of information about the reasoning of the neural network is often referred to as the black-box problem, which we are trying to mitigate in this study.

Instead of classification, we use an unsupervised deep-learning approach to learn solely from high-quality images. Our model is composed of two main components: the encoder takes the original image as an input and compresses it into a reduced representation, while the decoder uses the encoded image to reconstruct the input image. In the encoding step, a significant amount of information is truncated. Since we are dealing with a restricted problem of a limited data set, the network can infer information about the high-quality image distribution to recover truncated information during encoding. The quality of the output is therefore determined by the amount of truncation, which is adjusted by the model architecture, and the complexity of the true image distribution. After training the network with high-quality data, it is used to identify low-quality images based on

² <https://kso.ac.at/>

³ <http://cesar.kso.ac.at>

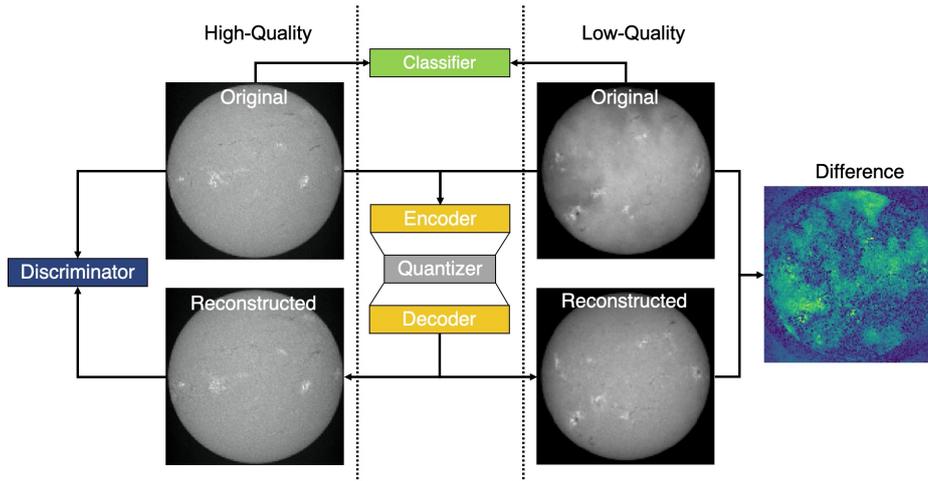


Fig. 2. Overview of the proposed method. The generator consists of an encoder, quantizer, and decoder. The generator is trained with high-quality images (*left*), where the encoder transforms the original image to a compressed representation and further information is truncated by the quantizer. The decoder uses this representation to reconstruct the original image. The discriminator optimizes the perceptual quality of the reconstruction and provides the content loss for the quality metric, which encourages the generator to model the high-quality domain. In addition, an optional classifier can be used which is trained to distinguish between the two image-quality classes. When low-quality images (*right*) are transformed by the pre-trained generator, the reconstruction shows deviations from the original, which allows us to identify the affected regions and to estimate the image quality.

the following property of the network. When translating images of the low-quality domain, the decoder cannot reconstruct the characteristics of the low-quality distribution, which leads to deviations between the original and the reconstructed image. The deviation between reconstruction and original is termed reconstruction loss.

The single optimization for a distortion metric (e.g., mean-squared-error) often shows a lack of performance and leads to blurred rather than sharp structures (Blau & Michaeli 2018). We therefore build upon different concepts: (1) We use an adversarial loss (Sect. 3.2), which allows our network to generate data, learn the characteristics of the high-quality domain, and enhance the perceptual quality of the reconstruction. (2) Instead of a pixel-wise distortion metric, we optimize for feature similarity (or content loss; Sect. 3.3). This leads to a larger deviation of the reconstruction loss for low-quality images and to a better translation of solar features. As an optional component, we introduce a classifier network to our existing architecture. This network is trained in addition to high-quality images also with low-quality samples and provides a probabilistic classification into high- and low-quality data. In addition, this network is used to increase the sensitivity for low-quality observations as estimated by the content loss (Sect. 3.4). An overview of the combined model training is given in Fig. 2.

For our primary model, we build on a multi-scale encoder-decoder architecture (similar to Agustsson et al. 2019; Johnson et al. 2016; Wang et al. 2018), which has shown strong performance for image translation, style transfer, and super-resolution tasks. The detailed model architecture is given in Appendix A. The following sections explain the individual components of our full model setup. In Sect. 3.1, the separation between high- and low-quality observations is discussed. Section 3.2 introduces the components to model the high-quality distribution. In Sect. 3.3 the loss function for feature-based translation is introduced. An optional component to increase the model performance is the classifier network, which is described in Sect. 3.4. The data preparation and evaluation metrics are covered in Sects. 3.5 and 3.6, respectively.

3.1. Feature compression

In order to separate between high- and low-quality observations, we aim at increasing the reconstruction loss for low-quality observations, while keeping the reconstruction of high-quality observations close to the original. We estimate the distance between the quality distributions by the median reconstruction loss of the test set and refer to this distance as margin, which we aim to maximize. The parameter that allows for the adjustment of the reconstruction quality is the amount of truncated information by the encoder. We build upon the image-compression network by Agustsson et al. (2019) which allows for the adjustment of parameters to generate images with a high perceptual quality. Here, the separation of high- and low-quality observations is limited by an upper and lower bound of compression. If the model truncates too much information in the intermediate layers, the performance for both domains suffers and the margin becomes too small to separate the distributions. Similarly, a model which does not truncate a sufficient amount of information will reconstruct the image pixel-wise and does not learn to infer information about the true image distribution, which leads to a similar performance on both data sets. The quantizer uses the latent feature maps of the encoder and maps it to L discrete levels. The information stored in the discretized representation $\hat{\omega}$ is measured by the entropy

$$H(\hat{\omega}) \leq \dim(\hat{\omega}) \log_2(L), \quad (1)$$

which is bound by the model architecture in terms of dimensions of the feature maps $\hat{\omega}$ as provided by the encoder, and the number of discrete levels L (Mentzer et al. 2018; Agustsson et al. 2019).

3.2. Adversarial loss

Training neural networks with a loss based on pixel-wise differences often results in blurred images and a lack of perceptual quality (Isola et al. 2017; Agustsson et al. 2019). This problem can be overcome with the use of generative adversarial

networks (GANs), which optimize the perceptual quality based on an additional neural network. With this setup, the generation of highly realistic synthetic images is possible (Wang et al. 2018; Karras et al. 2017).

As originally proposed by Goodfellow et al. (2014b), GANs are composed of a generating network (generator) that produces a synthetic image from a random input vector (latent space), and a discriminating network (discriminator) that distinguishes between generated and real images. The training is performed in a competitive setup between the generator and discriminator. In the first step, the model parameters of the generator are kept constant and the discriminator is trained to correctly classify images as either synthetic or real. In the second step, the discriminator weights are kept constant and the generator is trained to produce images which lead to a classification as real by the discriminator. In the first iterations the results are arbitrary, but from the iterative repetition of these steps both networks become experts in generating and discriminating between images. In other words, the discriminator learns from the true image distribution and ensures that the generator produces images that are close to real images. By randomly sampling inputs from the latent space, synthetic images can be produced.

We optimize the discriminator D and generator G for the objective proposed by Mao et al. (2017; Least-Squares GAN):

$$\mathcal{L}_D = \min_D \mathbb{E}[(D(x) - 1)^2] + \mathbb{E}[D(G(z))^2], \quad (2)$$

and

$$\mathcal{L}_G = \min_G \mathbb{E}[(D(G(z)) - 1)^2], \quad (3)$$

where the discriminator objective \mathcal{L}_D is given by the minimization of the expectation value of the loss as estimated by the squared difference between the discriminator prediction for the real images $D(x)$ and the generated images $D(G(z))$, and the assigned labels (1 for real images and 0 for generated images). The objective of the generator \mathcal{L}_G is obtained by minimizing the loss of the generated samples $G(z)$ for the inverted labels. In order to synthesize images, z corresponds to a random input vector sampled from a prior distribution. In this way, the network learns to find a mapping between the defined prior distribution and the data distribution (Agustsson et al. 2019; Goodfellow et al. 2014b). In the present case of image transformation, the random input vector z is replaced by an image x which is translated conditionally into a different domain (Isola et al. 2017; Wang et al. 2018; Karras et al. 2017). In our setup, the generating network is given by the encoder and decoder as introduced in Sect. 3. The encoder translates the input image into its latent space representation, while the decoder uses the encoded features to recover the original by generating the missing information from the inferred characteristics of the high-quality image distribution, as enforced by the discriminator.

The expected output of the generator G is the same as the input image x , and therefore we extend the generator loss as given in Eq. (3) by a loss term that ensures that the generated images $G(x)$ are close to the original x (e.g., MSE). In Sect. 3.3 we introduce the corresponding loss function which accounts for this term by estimating the content similarity between the original and reconstruction. We explicitly neglect an additional pixel-based loss, which would benefit the reconstruction of low-quality images. For the adversarial training objective, we follow the implementation by Wang et al. (2018) and use three discriminators trained in parallel and rescale the input by an average pooling layer with pooling window sizes of 1, 2, and 4 (see

also Agustsson et al. 2019). As shown for image compression by Agustsson et al. (2019), the GAN approach is capable of learning common image textures and features, while methods which only use the MSE fail at modeling the true image distribution.

For image-quality assessment, we train the generator only with high-quality observations. Therefore, the network learns only to encode solar features into the latent space representation. When the trained network is used to reconstruct low-quality observations, blurred solar features result in an invalid encoding. Furthermore, unknown features (e.g., clouds) cannot be translated into the compressed representation and will therefore be misinterpreted by the encoder. Here, we make use of this property. From the deviation between the original and reconstruction, we obtain a quality metric which has a reduced sensitivity for solar features and a strong sensitivity for deviations from the high-quality distribution.

3.3. Content loss

For pixel-based metrics (e.g., MSE) small shifts can cause a large increase of the reconstruction loss, which often leads to blurred results. An alternative to the pixel-wise comparison is the evaluation of content similarity between the original and reconstructed image. This can be achieved by comparing the activation of multiple layers of a pre-trained VGG network (Simonyan & Zisserman 2014). The network is hereby trained for a classification task, which extracts patterns at each intermediate layer. By comparing the activation of the generated and original image, a metric which is more sensitive to the content can be obtained. For our application, we define the content loss based on the discriminator, similar to Wang et al. (2018):

$$\mathcal{L}_{\text{Content},j} = \mathbb{E} \sum_{i=0}^4 \frac{1}{N_i} [\|D_j^{(i)}(x) - D_j^{(i)}(G(x))\|_1], \quad (4)$$

where $D_j^{(i)}$ refers to the layer i of the discriminator j , G to the generator, and N_i to the total number of features per layer. For each of our three discriminators we use all intermediate activation layers. Our final generator objective is given by:

$$\mathcal{L}_G = \min_G \sum_{j=0}^3 (\mathcal{L}_{\text{Content},j} + \mathbb{E}[(D_j(G(x)) - 1)^2]). \quad (5)$$

We use the content loss $\mathcal{L}_{\text{Content},j}$ (first term) to ensure that the generated images are close to the original and the adversarial loss (second term) to ensure that the generated images are perceptually similar to images from the high-quality domain.

3.4. Classifier network

In addition to the introduced architecture, we add an optional classifying network, which can be used in cases where a sufficient number of low-quality observations are available. The classifier is trained in parallel to the generator and discriminator and uses additional low-quality observations to provide a probabilistic prediction of the image quality. In the same way as for the discriminator, we use three classifiers at different resolutions. For the combined architecture, the content loss is derived from the feature activation of the classifiers instead of the discriminators. From this setup, we expect a larger margin between high- and low-quality reconstruction loss, since the classifier extracts features from both high- and low-quality images. The generator training is performed in the same way, using only high-quality observations.

3.5. Data preparation

For each image we crop the frame to $[-1000'', 1000'']$, which covers the full solar disk, and resize it to 128×128 pixels. This resolution is in accordance with our detection objective, where we are primarily interested in large-scale degradations. We compare two different data normalizations:

1. Image normalization: The data are rescaled based on the minimum and maximum value to an interval of $[-1, 1]$. In order to reduce the impact of small-scale brightness enhancements, we crop values outside 3σ from the mean of the image prior to normalizing.
2. Contrast normalization: For each image we subtract the median and divide the result by the standard deviation of the image. Values outside $[-2.5, 2.5]$ are cropped and afterwards rescaled linearly to $[-1, 1]$ (Goodfellow et al. 2016).

The contrast normalization centers the data to the mean, which makes the trained network more sensitive to shifts in the image intensity distribution (e.g., induced by partially occulting clouds). For the image normalization, we found a better accordance with the human perception when identifying clouds. A normalization based on the adjustment for exposure time and normalizing to a fixed value range provides less correlation with the apparent image quality. This is primarily due to the dynamic exposure time which compensates for low-intensity observations (e.g., caused by clouds) and high intensities (e.g., ongoing flares).

Deep-learning methods benefit from a larger number of training samples, and therefore we automatically extend our manually labeled data set. We assume that image classification can be used to detect most quality-degrading effects and train the classifier network introduced in Sect. 3.4 for a basic classification task. We use the 700 annotated observations which were not considered for the test set (Sect. 2) and apply the same parameter configuration as for the classifier training as given in Sect. 4. The trained model is used to automatically annotate a new data set of 20 184 high-quality and 2198 low-quality observations. This potentially introduces more misclassified samples in the training set of the primary model, but we found that a larger data set improves the performance and stability of the more challenging translation task.

During model training, we separate 10% of the training set for validation purposes. We note that we do not apply a strict temporal separation of the training and test sets because we are mostly interested in short-term variations. We find that the use of a large random data set is sufficient. Table 1 provides a summary of the considered data sets.

3.6. Quality and evaluation metric

For the evaluation of the reconstruction quality we use four different metrics:

1. Mean-squared-error (MSE): provides a pixel-wise loss which gives a good estimate for larger regions but suffers from small-scale differences.
2. Content-loss (Sect. 3.3): provides a metric optimized for the considered data set. This metric compares image features over pixel-wise differences.
3. Structural-similarity-index (SSIM): provides a good correspondence with human quality estimates and is based on image similarity (Wang et al. 2004).
4. Classification (Sect. 3.4): gives a probabilistic prediction of the image quality, but does not provide a continuous metric and requires a manually annotated data set with low-quality

Table 1. Data sets for training and evaluation of the proposed models.

Data set	Manual	Samples
General test	✓	3300
Classification training	✓	700
General training	✗	22 382
KSO unfiltered test	Partially	10 050 (620)

Notes. “Manual” indicates that a data set was labeled by a human. The number of samples includes both classes (high- and low-quality). The number in brackets refers to the fraction of manually labeled samples. Validation sets are split from the training sets on demand.

images. The classification can better account for minor atmospheric effects than the continuous metrics, but might lead to incorrect predictions for strong deviations from the training set.

The metrics resemble the image quality, where larger losses indicate a stronger deviation from the original and therefore lower quality. In order to separate the images into high- and low-quality, we apply thresholds according to the evaluation of the validation set. The content loss shows the largest margin between the individual classes and is therefore considered as our primary quality measure. We use the MSE and SSIM for additional verification. The result of the classifier cannot account for a quality measure and can lead to unexpected behavior for anomalous data. Therefore, for our classification scheme we combine the classifier predictions with the content loss.

The combined classifier is composed of three networks each predicting at a different resolution, with an output of an 8×8 grid. We take the mean result per grid and sum over the classifiers. We classify images as low-quality above a threshold of 1, which corresponds to the classification as low-quality image by at least one classifier.

In addition, we identify anomalies by the continuous quality metric. To this aim, we use the content loss and scale it according to the results on the validation set. We define the low-quality threshold at 3σ above the mean and scale the data between zero and four times the low-quality threshold. From this scaling, observations with a quality value of 0 refer to a perfect reconstruction, 0.25 defines the low-quality threshold, and values above 1 correspond to observations with strong degrading effects (anomalous observations). For the base architecture without a classifier, we identify anomalies solely on the content loss. Here, we lower the threshold to 2σ above the mean of the high-quality distribution as evaluated on the validation set and leave the threshold for anomalous observations unmodified.

We evaluate the correct predictions of low- and high-quality images in terms of accuracy and the True-Skill-Statistic (TSS; also known as Hanssen & Kuipers Discriminant) (Barnes et al. 2016; Pötzi et al. 2018)

$$\text{TSS} = \frac{\text{TP}}{\text{TP} + \text{FN}} - \frac{\text{FP}}{\text{FP} + \text{TN}}. \quad (6)$$

The variables correspond to the entries of the confusion matrix: number of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN).

4. Results

For the model training we apply different parameter settings and evaluate them according to the metrics introduced in Sect. 3.6. We use the shorthand notation of (CLASS/DISC)-qX-(IMG/CONTR) where DISC refers to the base architecture,

Table 2. Performance of the different model settings, as evaluated on the manually classified test set.

Classifier	Compression	Data normalization	High-quality MSE	Content	MSE	SSIM	TSS	Accuracy
✓	8	ImageNorm	0.0044	0.37	0.013	0.03	0.95	97.7%
✓	8	ContrastNorm	0.0017	0.72	0.042	0.08	0.97	98.5%
✓	1	ImageNorm	0.0044	0.37	0.006	0.02	0.95	97.7%
✓	1	ContrastNorm	0.0019	0.45	0.003	0.01	0.95	97.7%
✗	8	ImageNorm	0.0038	0.12	0.005	0.02	0.82	89.3%
✗	8	ContrastNorm	0.0016	0.18	0.004	0.04	0.92	96.1%
✗	1	ImageNorm	0.0042	0.18	0.006	0.03	0.93	96.5%
✗	1	ContrastNorm	0.0018	0.47	0.028	0.09	0.93	96.5%
Baseline KSO quality classification			–	–	–	–	0.30	64.2%

Notes. Content, MSE, and SSIM refer to the difference between the median losses of the two quality distributions for the according metric (see Sect. 3.6). Compression refers to the number of channels used for the quantized representation. Correct classifications are determined by the criteria introduced in Sect. 3.6 and are quantified in terms of accuracy and TSS. The evaluation of the original KSO labels is given in the last row. The results of the model yielding the best performance (CLASS-q8-CONTR) are marked in bold face.

CLASS to the classifier extension, and X denotes the number of channels in the compressed representation. All our models use five discrete levels for quantization. We compare both data normalizations, where IMG refers to the image normalization and CONTR to the contrast normalization. For example, CLASS-q8-IMG refers to the network that has been extended by a classifier with eight filters in the quantizer and uses images that are normalized by the minimum and maximum value.

We train each of our models for 300 000 iterations until the MSE and content loss of the low-quality samples in our validation set start to converge towards an upper bound. We use the Adam optimizer with a learning rate of 0.0002 and set $\beta_1 = 0.5$ and $\beta_2 = 0.9$ (Kingma & Ba 2014).

4.1. Observation quality metric

For each of our models, we evaluate the performance metrics on the manually labeled test set (3300 samples). The model parameters and results are summarized in Table 2. We vary the architecture, the number of compression channels, and the data normalization. The average error in the reconstruction as measured by the MSE of the high-quality images serves as a performance indicator for configurations with the same normalization. The ability to separate between high- and low-quality observations is estimated by the margin between the mean of the high- and low-quality distribution for each of our metrics. We evaluate the margin in terms of content loss, MSE, and SSIM. From the thresholds defined in Sect. 3.6 we estimate the accuracy and TSS.

The CLASS-q8-CONTR model shows the best performance for each metric despite the SSIM. With an average content loss margin of 0.72, an accuracy of 98.5%, and a TSS of 0.97 it clearly provides the best result, while the other classifier configurations are similar in performance in terms of accuracy (cf. Table 2). The DISC-q8-IMG model shows a significantly lower performance than the other configurations, which all achieve accuracy scores above 96%. As can be seen from the accuracy and TSS, the number of compression channels and normalization have little impact on the ability to separate between high- and low-quality images. In contrast, the classifier architecture results in a performance increase by at least 1.2%. The accuracy of the classifier without a low-quality threshold is 96.9% and 98.1% for the ImageNorm and ContrastNorm, respectively. This is an improvement by about 1% for the models with image normalization, while for the top performing network the addi-

tional quality threshold has only a minor effect (+0.4%). The CLASS-q1-CONTR model has a lower accuracy as compared to the classifier prediction (−0.4%).

Using the same data set, we also compare the results of our deep-learning algorithm with the empirical KSO image-quality assessment method in Pötzi et al. (2015). The accuracy and TSS of the empirical method, listed in the bottom line of Table 2, have values of 0.3 and 64.2%, respectively. This is significantly lower than the new algorithm presented here. We further randomly sampled observations from the full KSO archive and compared the original KSO labels with the deep-learning model predictions, showing an agreement of 72.5% between the two methods.

Figure 3 shows the test set evaluation of our top performing network (CLASS-q8-CONTR) on the defined quality metrics. The low-quality threshold is obtained from the evaluation of the validation set. The plots are centered to the high-quality distribution; quality estimates outside the given range are not included. Panel a shows the evaluation of the classifier, where a large separation between the high- and low-quality class distributions can be seen, with almost no overlap at the low-quality threshold. This is in agreement with our assumption that a simple classification approach can detect most quality-degrading effects. The quality metrics in panels b–d are based on the difference between the original and reconstructed image and show a continuous transition between the two classes. The MSE in panel b provides a distinct separation of the two distributions, while an even larger margin is achieved for the content loss (panel d). The SSIM in panel c shows the weakest performance in separating the two classes, where only a fraction of the low-quality observations show a distinct deviation from the high-quality distribution. The quality value is scaled as discussed in Sect. 3.6 and is indicated by the horizontal bar at the bottom of Fig. 3. Here, the values in brackets indicate the content loss at the given thresholds. From the test set, we randomly selected quality estimates across the scale. The samples are shown at the bottom of Fig. 3. The low-quality threshold is given at 0.25 and images with a quality estimate above 1 are considered to suffer from strong atmospheric degradation or instrumental errors.

We found that the off-limb region can cause high quality scores, therefore we apply an on-disk correction. Especially faint clouds across the full disk can severely impact the reconstruction at the solar limb, which is not in correspondence with the image quality of the solar disk. For that reason, we remove the off-limb

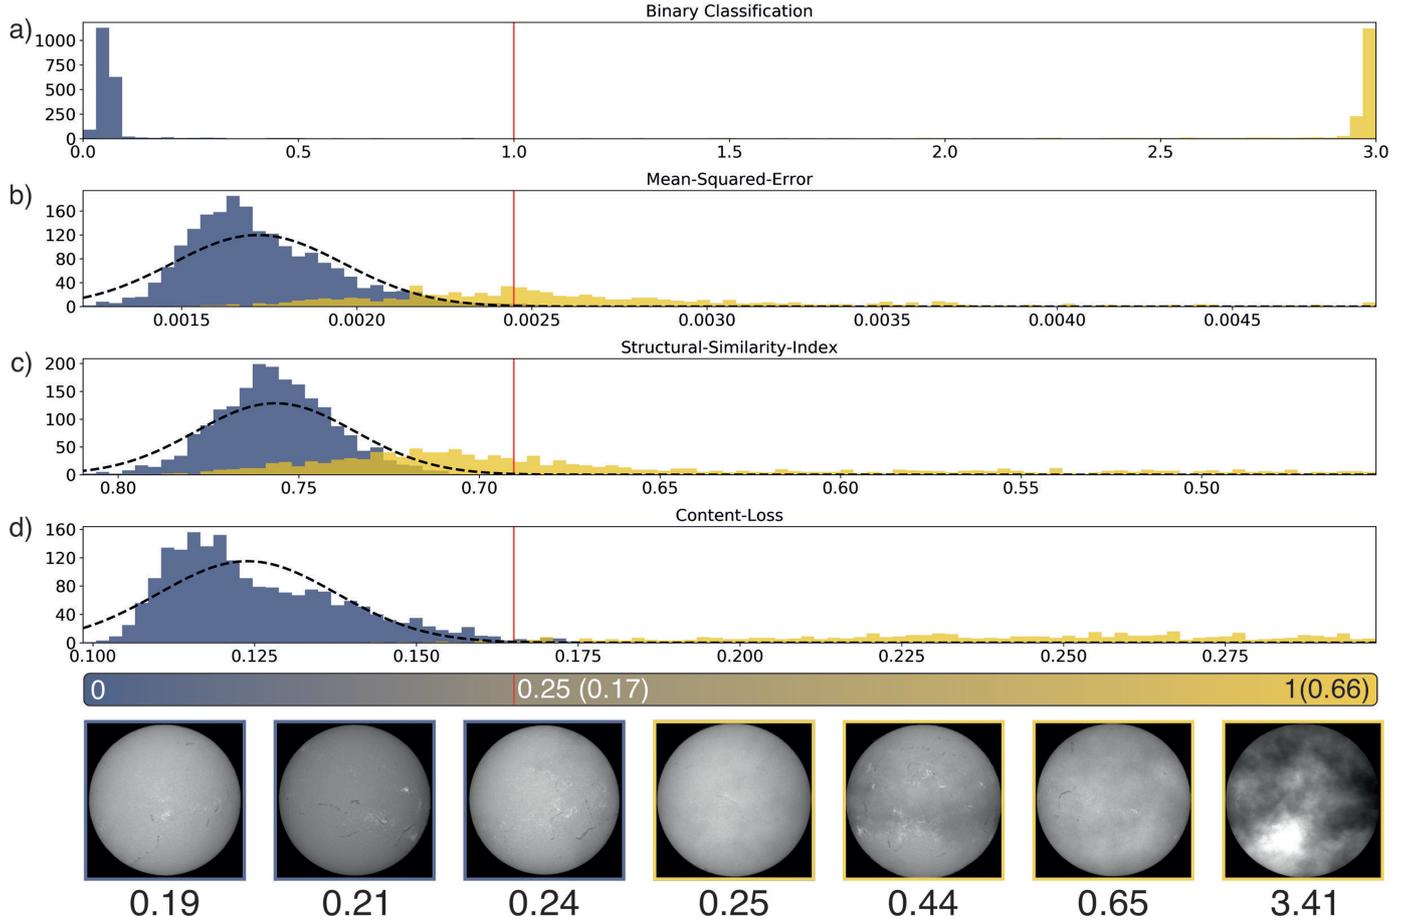


Fig. 3. Evaluation of the test set for our CLASS-q8-CONTR model. High-quality and low-quality samples are shown in blue and yellow, respectively. (a) Distribution of the classifier predictions, (b–d) IQMs between the original and reconstructed image in terms of (b) MSE, (c) SSIM and (d) content loss. The normal distribution (dashed black lines) of the high-quality images as well as the low-quality threshold (red line) are indicated for each metric. Samples of decreasing quality as evaluated by our metric and their corresponding content loss are shown in the *bottom panels*. The image outline is set according to the classifier prediction. An animation of the full test set with increasing image quality is available [online](#) (Movie 1).

region before evaluating the content loss and apply an offset to align the adjusted scale with the original quality measure of the high-quality samples.

We note that the definition of the threshold was selected to suit most applications, but can be adapted for specific demands (e.g., selection of very high-quality observations). A video that visualizes quality samples over the full scale is available [online](#) (Movie 1).

4.2. Region identification

In addition to the quality metric, the reconstructed image is used to identify the affected regions within the image. Convolutional neural networks show a relation between the spatial position of features in the image and the activation within the network (Zhou et al. 2016). Based on this property of the network, we may assume that local atmospheric effects in the original image can only cause deviations in a certain region of the reconstruction. From the difference between the original and reconstruction, regions with degrading effects can be detected. Figure 4 shows three examples of low-quality observations and the regions identified by the CLASS-q8-IMG model. For a first representation, we use the absolute difference map between the original and reconstructed image and visualize it

on a square-root intensity scale (Col. 2 in Fig. 4). In order to obtain the regions affected by strong degradations, we smooth the difference map with a total variation filter with a weight of 0.2 (Chambolle 2004) and apply a threshold of 0.1, which corresponds to the upper limit of the low-quality classification (Col. 3 in Fig. 4).

Our quality metric is optimized for feature similarity, and as features do not necessarily align pixel-wise with the reconstruction, we define a region identification based on the content similarity. To this aim, we use the same networks as for the content loss to obtain the feature activation at each layer, that is, either the classifier or discriminator network, depending on the architecture. From the fully convolutional architecture of the networks, a regional correlation between the feature activation and the position within the image can be drawn (Zhou et al. 2016). From each discriminator or classifier we extract the feature activation of the original and reconstruction, compute the absolute difference at each resolution level, and compute the mean over the channels. We compute the region map by upsampling the individual difference maps to the original image size and summing pixel-wise over all maps. Figure 5 shows the resulting difference maps for two examples and the mean absolute feature differences at a resolution of 8×8 pixel for each of the classifiers.

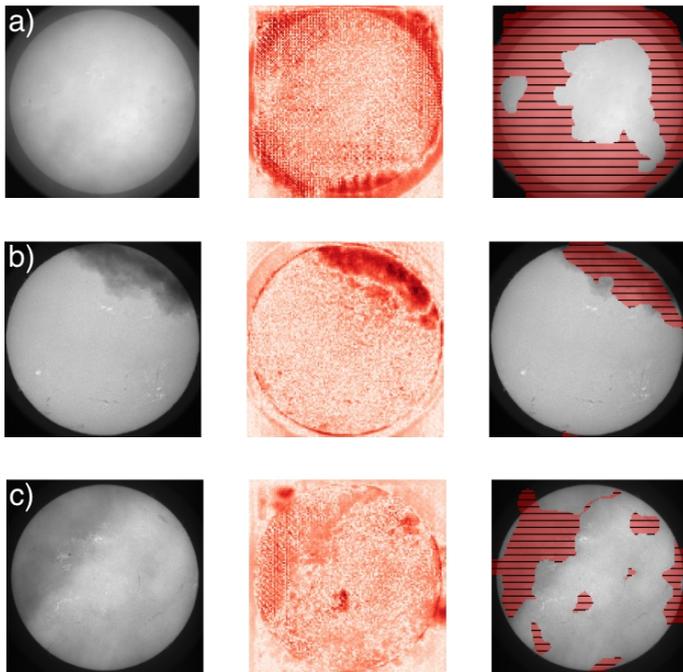


Fig. 4. Identification of affected regions based on the difference between the original and reconstructed image. *First column:* input image. *Second column:* difference between the original and reconstructed image on a square-root scale. *Third column:* original image with an overlay of the identified regions.

4.3. Application to unfiltered time series

In the regular observation mode at KSO, images of very low quality are rejected from further scientific use and are in general also automatically removed from the archive. In order to estimate the stability of our method to identify even strong deviations from regular observations, we use unfiltered data series from five observing days with varying observing and seeing conditions. From the full series, 620 samples with strong quality degradation are manually labeled as described in Sect. 2. From the predictions of the CLASS-q8-CONTR model we find that all samples exceed the low-quality threshold, both in terms of the binary classification and the quality threshold. From our quality scaling, we expect anomalous observations to have a quality value greater than 1. Out of the 620 samples, 617 exceed this threshold, which corresponds to a 99.5% accuracy in terms of identification of anomalous observations for the observing days studied. Figure 6 gives an overview of the individual days as evaluated with the CLASS-q8-CONTR model. The panels show the individual observing days. The series is characterized by smooth transitions for gradual variations in image quality as well as sharp jumps for sudden anomalies (e.g., appearance of clouds or contrails). At the bottom of Fig. 6, examples for various effects are shown (i.e., varying cloud coverage, overexposure, contrails). We note that the overexposed image at the bottom of Fig. 6 is a special observing mode at KSO used to enhance faint off-limb structures like prominences above the limb.

5. Discussion

The primary aim of our method is to derive a continuous quality estimator which correlates with human perception and allows reliable identification of distorted solar full-disk images. This is

accomplished by quantifying the deviation from the high-quality image distribution and by the use of a measure based on feature similarity. The performance of our model configurations is estimated by the ability to separate observations into high- and low-quality classes. Here, a manually labeled test set gives an independent measure of the agreement between the human estimate and the model classification.

As can be seen in Table 2, the margin between the high- and low-quality distribution directly correlates with the model performance in terms of accuracy and TSS. The content loss, MSE, and SSIM margins are in basic agreement with each other, while the content loss shows the best separation between the two distributions, as can be seen from the evaluation of the test set in panels b–d of Fig. 3. All metrics are characterized by a smooth transition between quality classes, with a larger spread of the low-quality distribution. This is in agreement with the samples of the test set, which include a broad range of different types of atmospheric effects. The samples in Fig. 3, as well as the video containing the evaluation of the full test set (included in the supplementary material), show good agreement with the assigned quality scores.

5.1. Model performance

For all eight model configurations, we find high performance overall. Our best-performing network (CLASS-q8-CONTR) achieves an accuracy of 98.5% and a TSS of 0.97 in separating high- and low-quality images. The other configurations provide similar performance with accuracies smaller by about 1–2%. Only the DISC-q8-IMG model shows much lower performance with an accuracy of 89.3%. We note that this is likely due to differences in the training process. With only one significant deviation out of eight trained models, we conclude that our approach achieves stable high performance for various configurations. The performance of these models is significantly better than the present empirical algorithm that is used in the observing pipeline at Kanzelhöhe Observatory (described in Pötzi et al. 2015), for which we obtained an accuracy of 64.2% and a TSS of 0.30 for the same test set.

In the case of available low-quality observations, the classifier architecture can boost the model performance by about 1–2% accuracy and adds additional robustness to the model predictions. This can be seen from the classifier configurations, which provide an accuracy of at least 97.7%, even though the performance varies in terms of content loss margin (Table 2). The classifier provides a probabilistic quality-class assignment, which requires combination with the content loss to account for a continuous quality measure. The low-quality threshold applied for classification has only a minor effect on the performance in terms of accuracy, but shows an increase for most models by 0.4–0.8%.

The different normalization and compression channels show similar results and mostly affect the quality threshold, which has a stronger impact on the architectures without a classifier. Especially at the high- to low-quality threshold, the class assignment becomes more subjective, which leads to expected deviations in model performance.

While the contrast normalization provides a better performance, the specific choice of normalization reveals a low impact on the overall result. The choice of normalization becomes more important for the identification of clouds in the image. From a visual inspection, the image normalization provides a better correspondence to regions covered by dense clouds, while the contrast normalization is centered to the mean and normalized

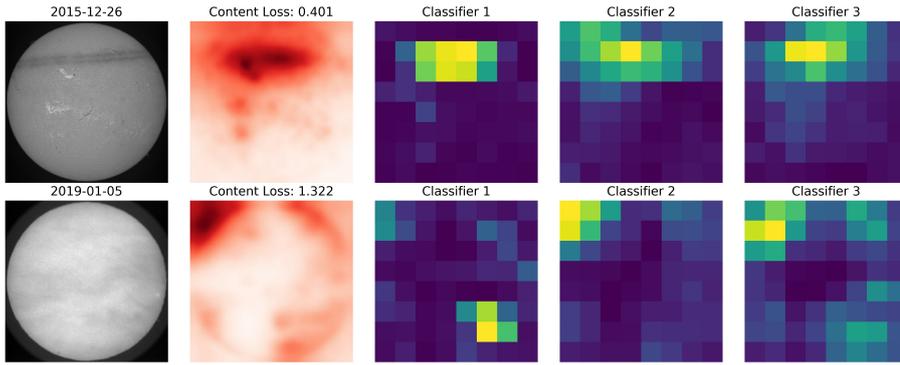


Fig. 5. Two examples of the region identification based on the content loss. *Column 1*: input image. *Columns 3–5*: absolute differences between the feature activation of the original and reconstructed image for the three classifiers. The feature maps shown were taken at a resolution of 8×8 pixels. *Column 2*: averaged difference over all feature maps and corresponding value of the content loss.

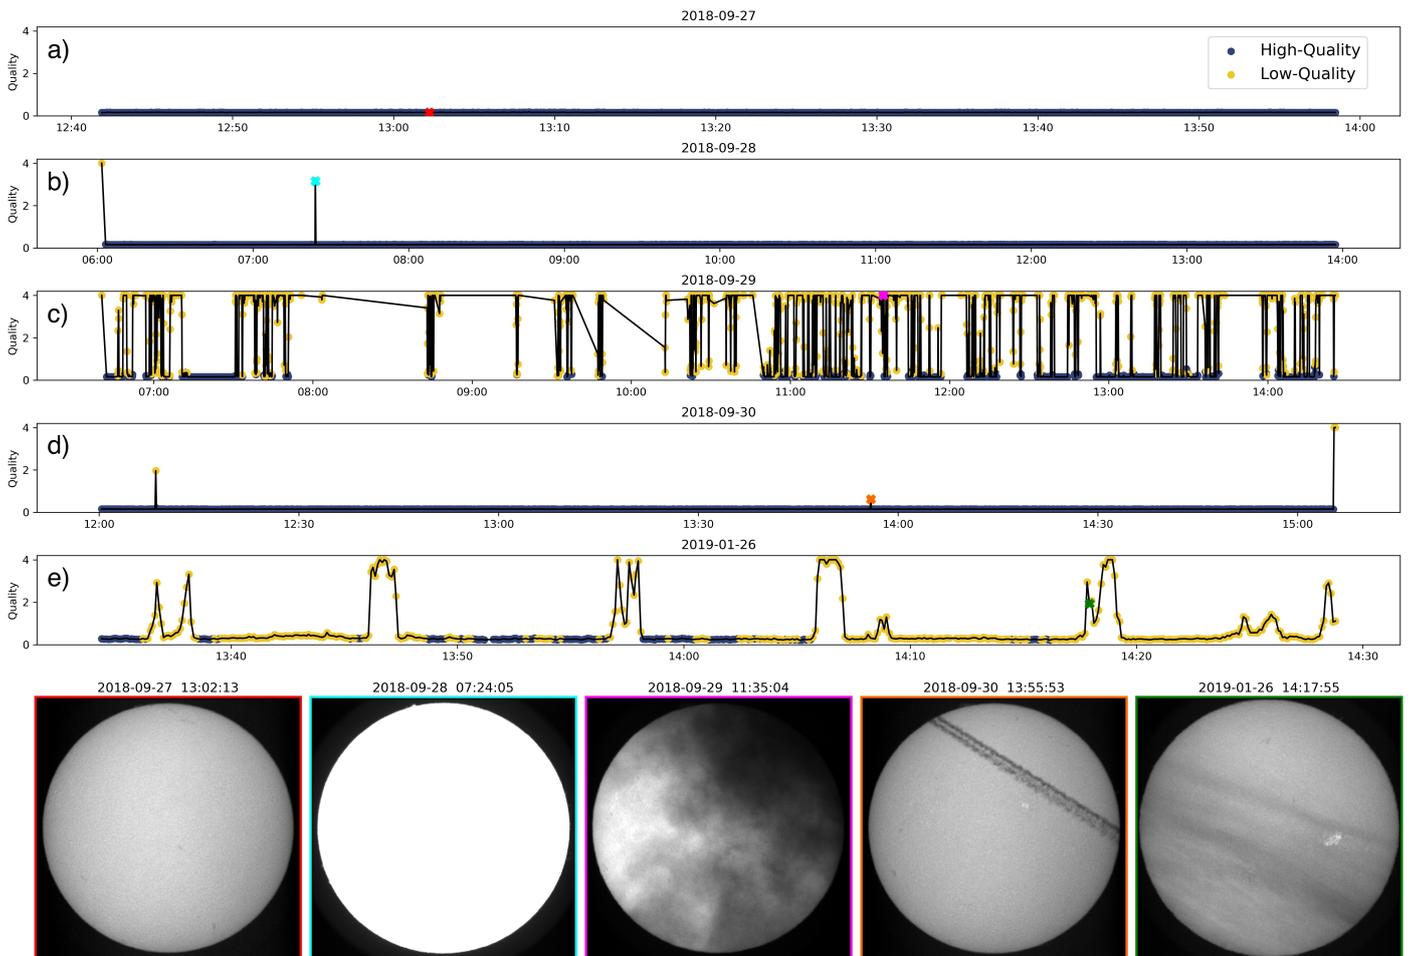


Fig. 6. Overview of five unfiltered observing days of KSO $H\alpha$ full-disk imaging as evaluated with the CLASS-q8-CONTR model. *Panels a–e*: derived image quality as a function of time. Examples across the series are given at the bottom. The first example shows a clear observation, while the subsequent examples show overexposure, strong cloud coverage, a contrail and partial cloud coverage. *Panel a*: a day of clear observing conditions with no quality degradations. *Panel b*: has generally high-quality conditions with a few overexposed images. *Panel c*: a varying cloud coverage leads to frequently changing quality scores throughout the day. *Panel d*: has generally high-quality conditions but a few images reveal degradation due to contrails. *Panel e*: a general quality decrease by clouds, while the quality further gradually decreases when denser clouds transition the disk. A video of *panel e* is available [online](#) (Movie 2).

by the standard deviation, which leads to stronger shifts in the intensity distribution by clouds and tends to identify overexposed regions.

5.2. Region identification

The identification of clouds and other quality-decreasing effects provides additional information to the quality metric. As can be

seen from Figs. 4 and 5, the reconstruction loss aligns with the affected regions in the original image. The network correctly reflects the impact of localized clouds (Fig. 4b) and global quality-degrading effects (Fig. 4a). Due to the dynamic exposure time, observations can show faint regions covered by clouds and overexposed regions simultaneously (Fig. 4c), which is detected by the neural network as deviation from an expected mean intensity value. The model is trained for content similarity, and

therefore the reconstruction does not align pixel-wise with the original. This induces a sensitivity for solar features in the difference masks, as can be seen from Fig. 4. While this can be mitigated with the use of extracted features from the discriminator or classifier (Fig. 5), the detection can only provide a coarse localization. The generated masks based on the content loss are produced by averaging all feature activation differences, which causes a suppression of small deviations, as can be seen in the second row of Fig. 5.

5.3. Stability and training

Often referred to as the black-box problem, it is difficult to extract information on the reasoning within a neural network. Here we are examining the model outputs and training progress in order to obtain information on the functionality and stability of our approach.

From the regions obtained by the difference between the reconstruction and original image in Sect. 4.2 it can be seen that deviations in reconstruction are spatially aligned with the regions of reduced quality in the original image. The extractions at different layers within the network (Fig. 5) reveal that the main contribution to the content loss is due to quality degradation, rather than solar features. From Cols. 3–5 in Fig. 5 it can be seen that each classifier is capable of extracting different features. This finding suggests that our metric provides an objective quality assessment, which covers multiple scales, includes multiple image features, and provides an enhanced sensitivity for atmospheric effects.

Neural networks can reveal large changes in the prediction, even for minor changes to the input (Goodfellow et al. 2014a; Papernot et al. 2017). Our model successfully detected 99.5% of the anomalous observations in the unfiltered time series (Sect. 4.3), which proves the robustness of the chosen approach. The image quality of the unfiltered series correctly reflects the smooth transition in decreasing image quality and also captures the sudden appearance of clouds. As can be seen from days with generally good observing conditions (Figs. 6a,b,d) the model shows high stability over the time series. The increased image-quality value in Fig. 6e is in correspondence with the poor observing conditions at this day. We conclude that the model is not prone to small deviations, as can be commonly observed for neural networks that are applied to classification tasks (Goodfellow et al. 2014a; Papernot et al. 2017).

An important component for the success of our method is the truncation of information during encoding, which increases the margin between the distribution of high- and low-quality observations as evaluated by the proposed distortion metrics (Sect. 3.6). This is controlled by the number of channels in the quantizer. The architecture with eight compression channels reduces the information to approximately 1% of the original input. We find that for lower compression rates the model falls back to a pixel-wise reconstruction, which decreases the sensitivity for faint clouds. Inversely, a higher compression can reduce the reconstruction capability, which results in increased sensitivity to the intrinsic solar features. This can also be seen from Table 2, where the models with eight compression channels show a lower average high-quality MSE than the models with the same normalization and one compression channel. Table 2 also shows that a better reconstruction of the image is not necessarily beneficial for the identification of low-quality images. This can be seen from the DISC-q8-IMG model, which achieved the lowest accuracy (89.3%), while producing the best reconstructions

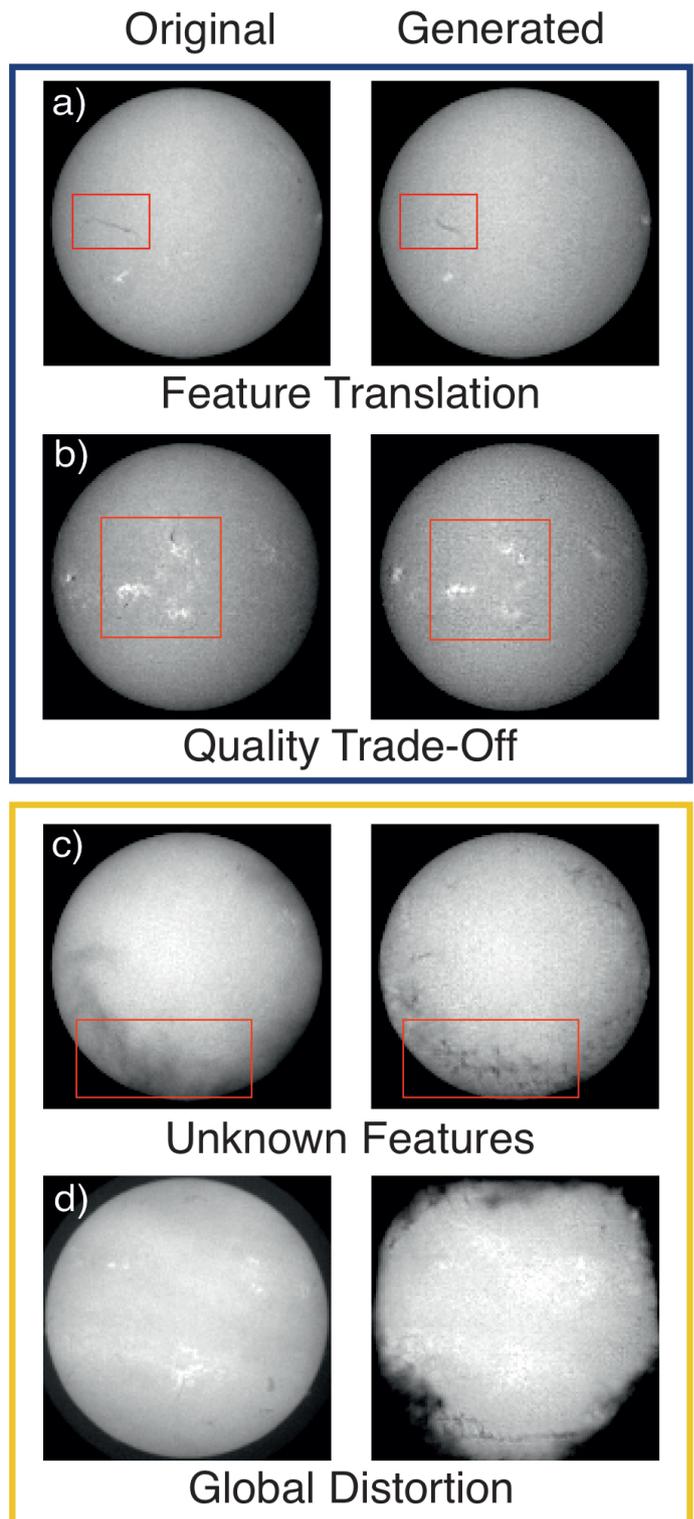


Fig. 7. Pairs of original $H\alpha$ filtergrams (left) as provided to the model and the resulting reconstruction (right). *Top panel:* samples of the high-quality distribution, *bottom panel:* samples from the low-quality domain. (a) Illustration of the feature-based image translation. The reconstructed filament can be clearly identified, but shows differences in appearance from the original. (b) Example of the image-quality decrease of a high-quality image due to strong compression. (c) For low-quality images, the unknown features result in artifacts in the reconstruction. (d) Global atmospheric effects show strong differences in the reconstructed image.

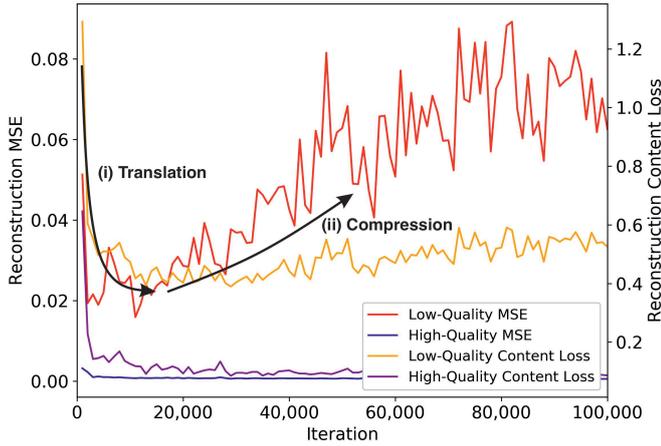


Fig. 8. Evaluation of the validation set during the training progress. From both quality metrics (MSE and content loss), two distinct phases can be identified. The orange and red lines denote the reconstruction performance of the low-quality distribution, while the blue and purple lines are the reconstruction losses of the high-quality distribution.

(0.0038 high-quality MSE) among the models that use image normalization.

As illustrated in the upper row of Fig. 7 the model is able to reconstruct high-quality observations that closely resemble the original image. As a result of the training with the content loss, the reconstructions show a feature-based rather than a pixel-based translation. An example of this behavior is shown in Fig. 7a, where the reconstructed filament can be clearly identified, but appears different in shape. The feature-based reconstruction becomes more evident for low-quality images, where the network fails to translate unknown features. Figure 7c shows a low-quality image with clouds partially occulting the solar disk. As a result, the network reconstruction shows a strong deviation from the original dark structure and produces a structure with a filament-like appearance. For global atmospheric effects (i.e., coverage by faint clouds) the reconstruction yields even stronger deviations, as can be seen from Fig. 7d. While the strong compression increases the sensitivity for unknown features, it leads to a trade-off in reconstruction quality for high-quality images (see Fig. 7b).

From the evaluation of the validation set we identify two characteristic phases during training, which support our assumption of a feature-based translation (Fig. 8): (i) A translation phase, where the network learns to reconstruct the original image by preserving a maximum amount of information in the quantizing layer, which increases the similarity between the original and reconstructed image for both high- and low-quality images; and (ii) a compression phase, where the network starts to learn from the image distribution and truncates information during encoding, which improves the reconstruction of high-quality images, while low-quality images suffer from the learned feature compression.

5.4. Applicability

Our model requires approximately 17 ms for a single image-quality estimation using a Nvidia Tesla M40 GPU. For a CPU-based prediction, the model takes approximately 93 ms per observation when using 15 cores. In both cases, we assume that the observations are already loaded and prepared as described in Sect. 2. Compared to the observing cadence of 6 s of the KSO

H α filtergrams, our image-quality assessment thus requires just a minor fraction of the total acquisition time, which allows for real-time application of our method, even for limited computational resources. For higher-cadence modes, the performance can be linearly increased by parallelizing the computations.

This study presents a first approach to image-quality assessment for solar full-disk observations. Application to different telescopes, different filters, or even multiple filters can be easily accomplished. This involves the labeling of a new data set and training the neural network as proposed by our method. A sufficiently large data set only requires a few hundred coarsely labeled images (binary classification). In the case of sparse low-quality observations available, the basic architecture can be considered, which only requires high-quality observations. The presented quality assessment for KSO observations is only valid for full-disk images; this is especially important for instrumental misalignment, where a partial disk would result in a large reconstruction loss. In order to assess the quality of smaller regions, the network can be trained with image patches, which omits the encoding of the full-disk. Future developments of this method offer great potential for related applications. (1) High-resolution solar observations rely on identifying the best observations for post-facto algorithms (e.g., speckle-interferometry Popowicz et al. 2017). The capability of our method to learn feature appearance from high-quality or even space-based observations can provide a quality metric which objectively estimates the distortion of solar features. (2) The extension to a metric which can operate between observations from different instruments requires further extension of the data set in order to compensate for the instrumental differences, while it still needs to give an objective quality estimation based on the highest image quality attainable. In a first test, the pre-trained model was applied to overlapping data series from KSO and Uccle Solar Equatorial Table (USET), and demonstrated its capability to filter low-quality images between different observation sites. To account for small quality variations between different sites, higher resolutions should be included to allow for the dynamic selection of the highest image quality. (3) A different approach is the detection of solar transient events. Our method already shows some sensitivity for flares. By removing all flaring samples from the training set, a detection module could be obtained. We further note that this concept is not restricted to image data, but could also be applied to 1D or 3D data.

6. Conclusions

We presented a method for stable classification and quantification of image quality in ground-based solar full-disk images. From a set of regular observations, we derived an objective no-reference IQM which accounts for seeing, atmospheric effects, and instrumental errors in H α filtergrams. Our method achieves an almost perfect detection of anomalies (99.5%) for an unfiltered time series of several observing days (covering in total 10 050 images), and shows high performance on an independent test set for the years 2012–2019 (3300 images) covering a broad variety of atmospheric effects and solar activity. Our top-performing neural network achieved an accuracy of 98.5% and a TSS of 0.97 in separating high-quality observations from observations with degrading effects (low-quality), as compared to visual inspection. The proposed IQM shows good agreement with human perception and provides a smooth transition between high-quality and low-quality observations. Once the model is trained, it can be operated without any further reference image. The processing time is short (about 17 ms for inference)

as compared to typical observing cadences and can easily be parallelized, allowing for efficient real time application.

Our method is based on two important concepts that make our approach superior to existing methods. (1) We made use of the true image distribution and quantify deviations from it. (2) We employ an image-quality metric which estimates feature similarity rather than pixel-based variations. As the model can be trained solely with the use of regular observations, it is suitable for many applications in observational astrophysics. With the availability of a sufficient number of low-quality observations, the model performance can be further increased by using the proposed classifier architecture. In addition to the quality score, our method provides an identification of the image regions affected by reduced quality. The presented method can be adapted to other instruments, wavelength channels, and observing targets. Furthermore, our method offers the potential for application to high-resolution solar physics data, homogenization of observation series from telescope networks, time series, and event detection.

Acknowledgements. This research has received financial support from the European Union's Horizon 2020 research and innovation program under grant agreement No. 824135 (SOLARNET). The computational results presented have been achieved using the Vienna Scientific Cluster (VSC) and the Skoltech HPC cluster ARKUDA. This research has made use of SunPy v1.1.4 (Mumford et al. 2020), an open-source and free community-developed solar data analysis Python package (Barnes et al. 2020).

References

- Agustsson, E., Tschannen, M., Mentzer, F., Timofte, R., & Gool, L. V. 2019, *Proceedings of the IEEE International Conference on Computer Vision*, 221
- Barnes, G., Leka, K., Schrijver, C., et al. 2016, *ApJ*, 829, 89
- Barnes, W. T., Bobra, M. G., Christe, S. D., et al. 2020, *ApJ*, 890, 68
- Blau, Y., & Michaeli, T. 2018, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6228
- Chambolle, A. 2004, *J. Math. Imaging Vision*, 20, 89
- Chollet, F. 2017, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1251
- Deng, H., Zhang, D., Wang, T., et al. 2015, *Sol. Phys.*, 290, 1479
- Galvez, R., Fouhey, D. F., Jin, M., et al. 2019, *ApJS*, 242, 7
- Goodfellow, I. J., Shlens, J., & Szegedy, C. 2014a, ArXiv e-prints [arXiv:1412.6572]
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al. 2014b, *Advances in Neural Information Processing Systems*, 2672
- Goodfellow, I., Bengio, Y., & Courville, A. 2016, *Deep Learning* (MIT Press)
- Gosain, S., Roth, M., Hill, F., et al. 2018, in *Ground-based and Airborne Instrumentation for Astronomy VII*, Int. Soc. Opt. Photon., 10702, 107024H
- Grundahl, F., Kjeldsen, H., Frandsen, S., et al. 2006, *Mem. Soc. Astron. It.*, 77, 458
- Harvey, J., Hill, F., Hubbard, R., et al. 1996, *Science*, 272, 1284
- He, K., Zhang, X., Ren, S., & Sun, J. 2016, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770
- Huang, Y., Jia, P., Cai, D., & Cai, B. 2019, *Sol. Phys.*, 294, 133
- Isola, P., Zhu, J. Y., Zhou, T., & Efros, A. A. 2017, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1125
- Johnson, J., Alahi, A., & Fei-Fei, L. 2016, *European Conference on Computer Vision* (Springer), 694
- Karras, T., Aila, T., Laine, S., & Lehtinen, J. 2017, ArXiv e-prints [arXiv:1710.10196]
- Kingma, D. P., & Ba, J. 2014, ArXiv e-prints [arXiv:1412.6980]
- LeCun, Y., Bengio, Y., & Hinton, G. 2015, *Nature*, 521, 436
- Löfdahl, M. G., van Noort, M. J., & Denker, C. 2007, in *Modern Solar Facilities – Advanced Solar Science*, eds. F. Kneer, K. G. Puschmann, & A. D. Wittmann, 119
- Mao, X., Li, Q., Xie, H., et al. 2017, *Proceedings of the IEEE International Conference on Computer Vision*, 2794
- Mentzer, F., Agustsson, E., Tschannen, M., Timofte, R., & Van Gool, L. 2018, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4394
- Mittal, A., Moorthy, A. K., & Bovik, A. C. 2012, *IEEE Trans. Image Process.*, 21, 4695
- Mumford, S. J., Christe, S., Freij, N., et al. 2020, <https://doi.org/10.5281/zenodo.3871057>
- Otruba, W., & Pötzi, W. 2003, *Hvar Obs. Bull.*, 27, 189
- Papernot, N., McDaniel, P., Goodfellow, I., et al. 2017, *Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security*, 506
- Pesnell, W. D., Thompson, B. J., & Chamberlin, P. C. 2012, *Sol. Phys.*, 275, 3
- Popowicz, A., Radlak, K., Bernacki, K., & Orlov, V. 2017, *Sol. Phys.*, 292, 187
- Pötzi, W., Veronig, A. M., Riegler, G., et al. 2015, *Sol. Phys.*, 290, 951
- Pötzi, W., Veronig, A., & Temmer, M. 2018, *Sol. Phys.*, 293, 94
- Rimmele, T. R., & Marino, J. 2011, *Liv. Rev. Sol. Phys.*, 8, 2
- Simonyan, K., & Zisserman, A. 2014, ArXiv e-prints [arXiv:1409.1556]
- Veronig, A. M., & Pötzi, W. 2016, in *Coimbra Solar Physics Meeting: Ground-based Solar Observations in the Space Instrumentation Era*, eds. I. Dorotovic, C. E. Fischer, & M. Temmer, *ASP Conf. Ser.*, 504, 247
- Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. 2004, *IEEE Trans. Image Process.*, 13, 600
- Wang, T. C., Liu, M. Y., Zhu, J. Y., et al. 2018, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8798
- Wöger, F., von der Lühe, O., & Reardon, K. 2008, *A&A*, 488, 375
- Ye, P., Kumar, J., Kang, L., & Doermann, D. 2012, *2012 IEEE Conference on Computer Vision and Pattern Recognition (IEEE)*, 1098
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., & Torralba, A. 2016, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2921

Appendix A: Model architecture

Our primary model is composed of an encoder, quantizer, and decoder. We use the notation from Johnson et al. (2016) and Wang et al. (2018), where $cksj-n$ denotes a $k \times k$ convolution layer with stride j , instance-normalization and ReLU activation with n filters. dn refers to a 3×3 convolution with stride 2, instance-normalization, and ReLU activation with n filters, whereas un denotes the same configuration but with transposed convolution instead of convolution layers. Rn refers to a residual block as proposed in the ResNet50 architecture, with n filters (He et al. 2016). Throughout our network we use reflection padding in order to reduce boundary artifacts (Wang et al. 2018). The last convolution layer of the encoder uses the activation of the quantizer instead of a ReLU activation (see Appendix B). The output is produced by a final convolutional layer where we omit the normalization and use a tanh activation.

The discriminator uses four consecutive stride 2 convolutions with instance normalization and a leaky-relu activation with a slope of 0.2 which we denote Cn , where n refers to the number of filters. No normalization is applied to the first convolutional layer (Wang et al. 2018). Per convolution the number of filters is increased by a factor of two, while the spatial dimension is reduced by a factor of four.

For the classifier network, we apply the same architecture as for the discriminator. We use three discriminators and three classifiers, where we use the full resolution as provided by the generator for the first discriminator/classifier and consecutively reduced resolutions for the second and third. For this task, we apply average pooling to the input images.

Generator (with number of quantization channels y):

Encoder: $c7s1-64,d128,d256,d512,c3s1-y,Q$

Decoder: $9x R512,u256,u128,u64,c7s1-1$

Discriminator/Classifier (3x with Average-Pooling of 1, 2 and 4): $C64,C128,C256,C512$

Appendix B: Quantizer

For the discrete representation $\hat{\omega}$ we use the quantizer Q as proposed in Agustsson et al. (2019) which uses a hard non-differentiable quantization in the forward pass and a differentiable approximation in the backward pass of model training. This is implemented with a gradient stop:

$$\hat{\omega} = \text{tf.stop_gradient}(\hat{z} - \tilde{z}) + \tilde{z}. \quad (\text{B.1})$$

Here the hard quantization \hat{z} is computed by rounding the output of the encoder z to integers and the soft quantization \tilde{z} is obtained by applying the softmax function to the absolute difference between the encoder output z and the discrete centers c (Mentzer et al. 2018):

$$\tilde{z}_i = \sum_{j=1}^L \frac{\exp(-\|z_i - c_j\|_1)}{\sum_{l=1}^L \exp(-\|z_i - c_l\|_1)} c_j. \quad (\text{B.2})$$

The encoder output z is calculated from the last features in the encoder:

$$z_i = \sigma(x_i) * (L - 1), \quad (\text{B.3})$$

where x refers to the output of the last convolutional layer in the encoder, L to the number of centers, and σ to the sigmoid activation function.

Appendix C: Samples

In Fig. C.1 we show examples of different quality degradation. The image quality, as estimated by our CLASS-q8-CONTR model, is indicated on top of the images.

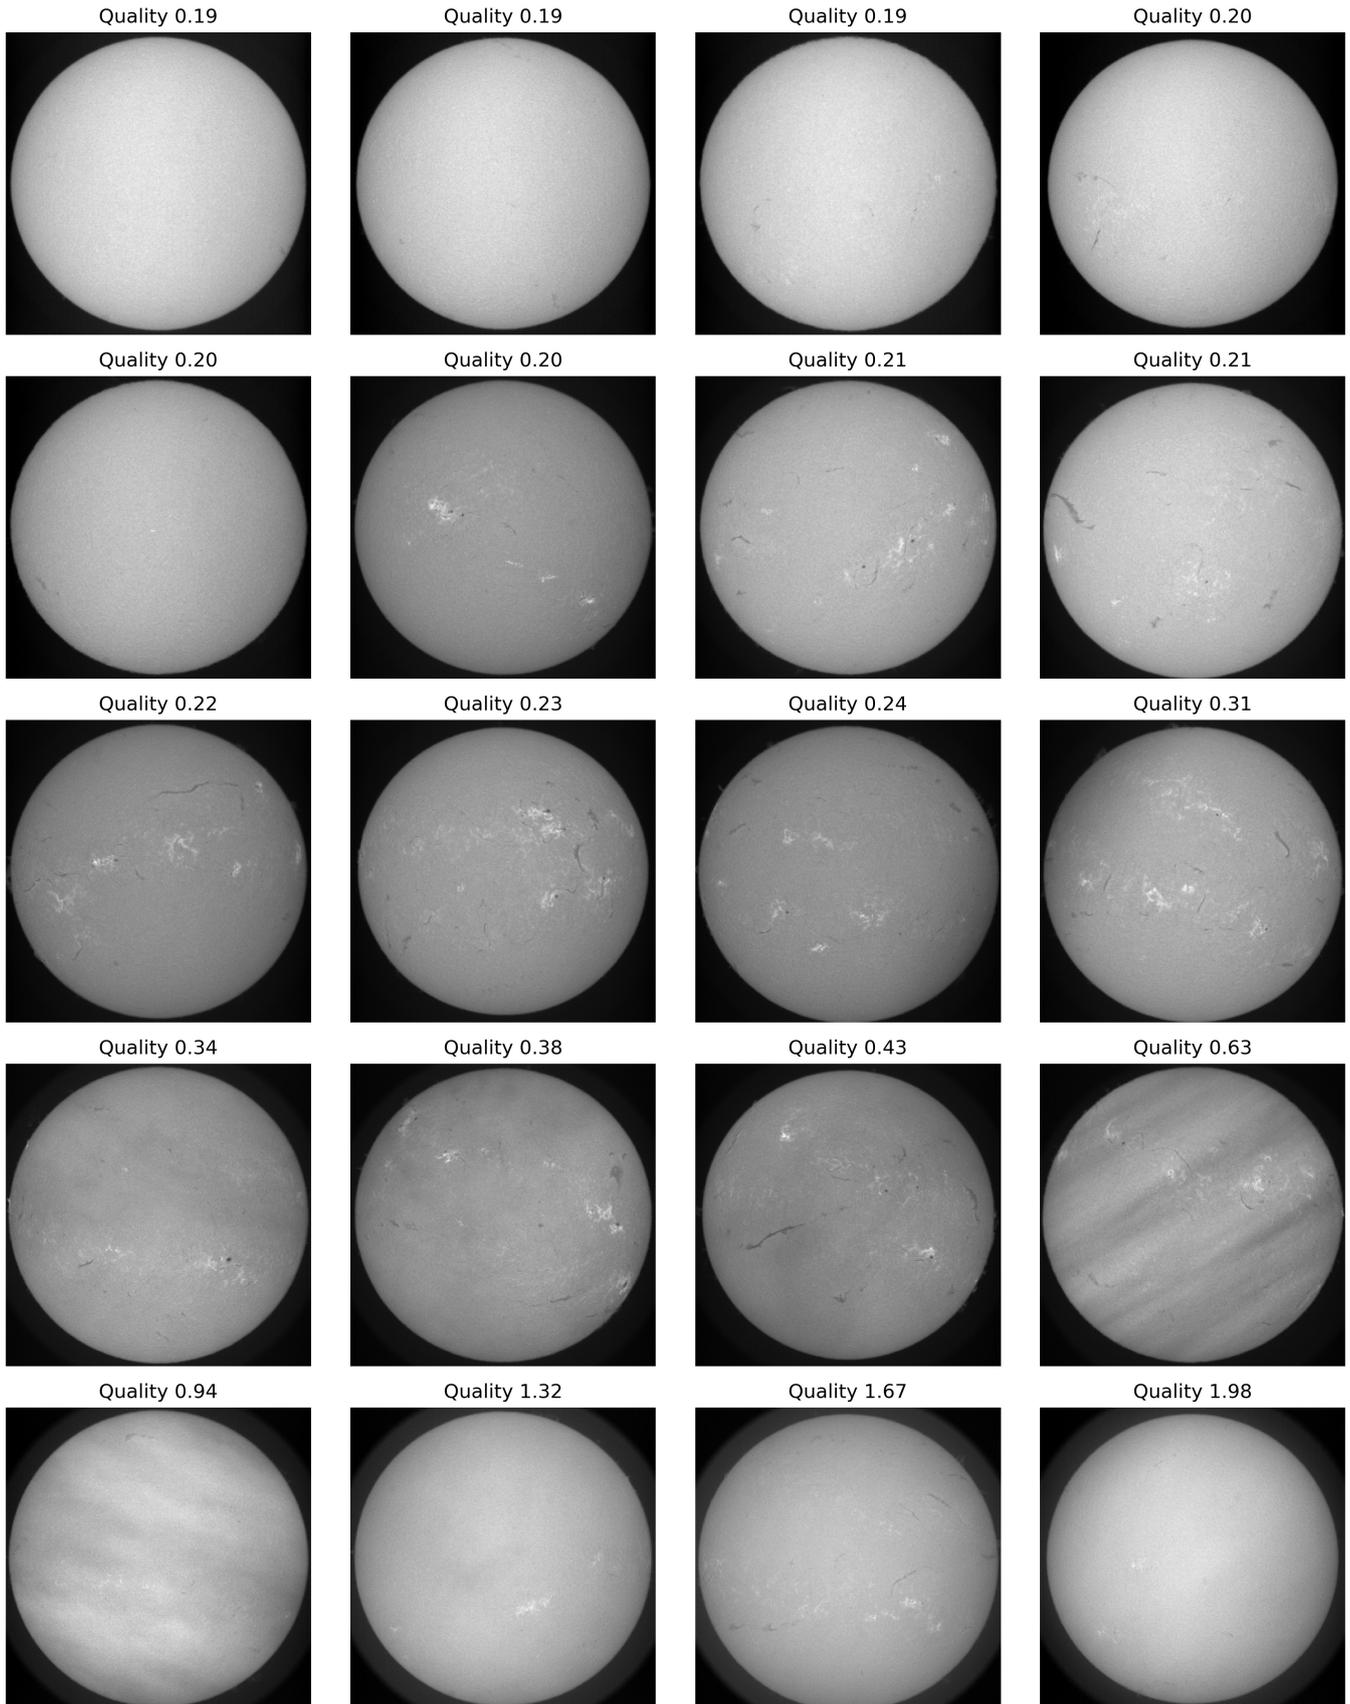


Fig. C.1. Examples from the test set with adjustment of the off-limb region (see Sect. 3.6). The images are sampled across the full test set and sorted with respect to the estimated image-quality value.