An incomplete guide to tapping mode, image postprocessing and life in general

Antti Soininen

January 29, 2013

Send feedback to: antti.soininen@iki.fi

Contents

1 Scanning parameters

1.1 The feedback loop

The feedback loop is the hearth of an SPM. Thus it is essential to understand its basics.

In tapping mode, the feedback loop monitors the amplitude of the tip. The measured amplitude (appears as RMS Amplitude in the Meter window and in the Real time status view when not engaged) is compared to the value defined by the `Amplitude setpoint` scanning parameter. If the measured amplitude differs from the `Amplitude setpoint`, the feedback loop adjusts the tip-to-sample distance by changing the z piezo position.

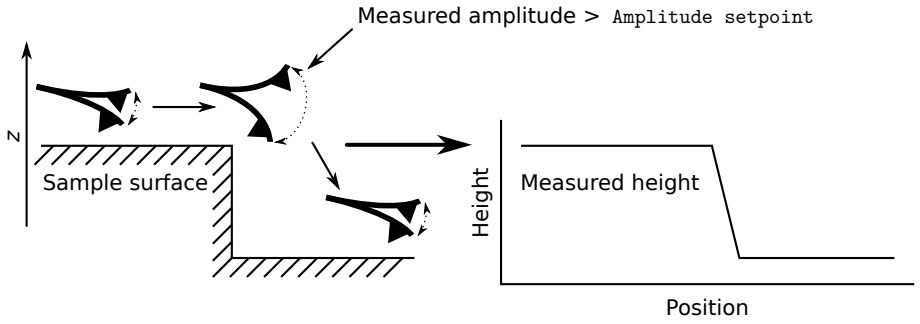

Figure 1: The feedback loop tries to keep the measured amplitude of an oscillating tip at a constant value defined by the `Amplitude setpoint` parameter. This is accomplished by moving the z piezo up or down. The z piezo position is then recorded as the height of the sample.

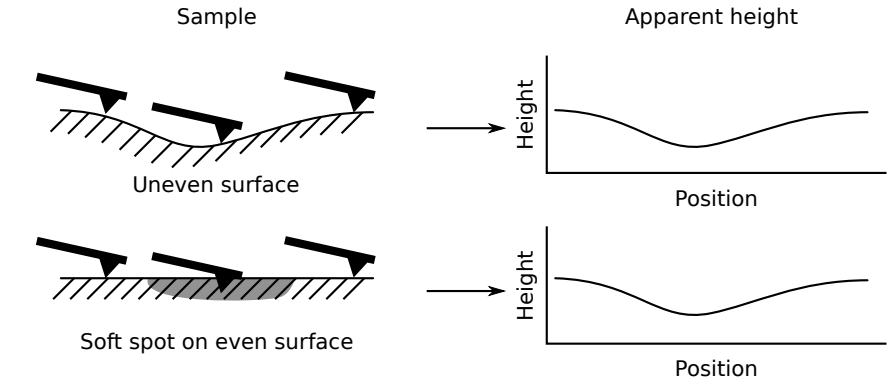

Figure 2: Real height vs. apparent height. A change in the surface properties — such as a soft spot — may result in an artificial hole in the measured height.

The z piezo position is the height that the SPM records. Thus, the height image of the sample is actually a map of the z piezo positions, not the real topography of the surface of the sample! The process is exemplified in Figure 1. The x and y piezos scan the tip horizontally on the sample surface at a constant speed. When the tip moves over a step, its amplitude becomes larger than the `Amplitude setpoint`. To bring the amplitude back to the `Amplitude setpoint` value, the feedback loop moves the tip by the z piezo until its amplitude matches the `Amplitude setpoint` again. This vertical movement is recorded as the height profile of the sample.

Everything that changes the amplitude of the tip makes the feedback loop to change the z piezo position and thus the *apparent* height of the sample. Figure 2 exemplifies the situation: a softer spot on the sample surface may result in an artificial dent in the measured height.

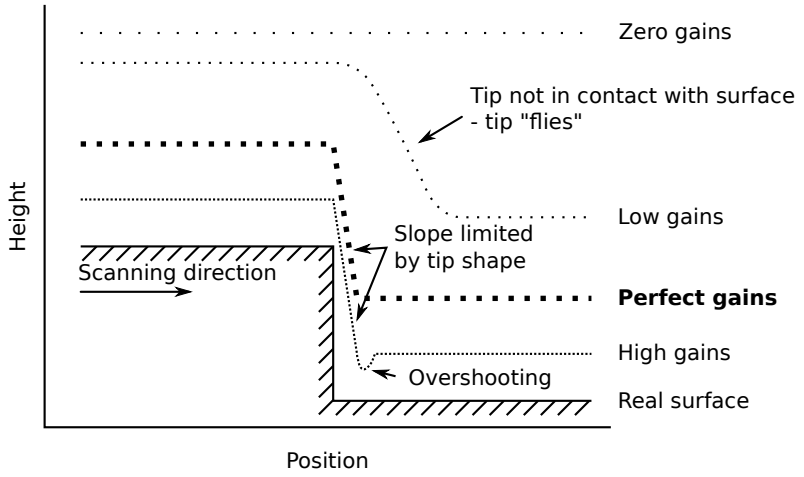

Figure 3: Comparison of measured height profile with different values of `Integral` and `Proportional` gains. At low values, the tip does not follow the surface properly while at too high values, the feedback loop overshoots.

1.2 Integral and proportional gains

The `Integral gain` and `Proportional gain` parameters define how quickly the feedback loop operates.

Basically, the feedback loop is a PI controller, where P stands for proportional and I for integral. Such controllers try to minimize an error signal by applying a control signal which is somehow proportional to the error. In an SPM, the error signal is the difference between the measured amplitude and the `Amplitude setpoint` parameter and is called the amplitude error. The control signal is the z piezo position.

The proportionality between the amplitude error and z piezo position is controlled by two parameters: `Integral gain` and `Proportional gain`. The `Integral gain` tries to account for how the amplitude error has changed in the past while `Proportional gain` accounts for the current amplitude error. In usual imaging conditions, the `Integral gain` solely determines the behavior of the feedback loop and the SPM operator can forget the `Proportional gain`.

Figure 3 shows how different gain values affect the apparent height. When both gains are zero, the feedback loop does not move the z piezo regardless of measured amplitude. When the gains are increased, the feedback loop starts to react. However, at low gain values, the feedback loop is slow to react and the tip ”flies” over rapid changes in the sample height. If the gains are too high, the feedback loop may react too strongly and overshoots.

Another problem with too high gains is resonance. This happens when the feedback loop overshoots and then overshoots again while trying to correct itself resulting in the patterns shown in Figures 4 and 5. The resonating gain limits are strongly dependent on the sample. Sometimes one is able to use surprisingly high values for `Integral gain` but sometimes the highest non-resonating `Integral gain` results in ”flying” tip (see the

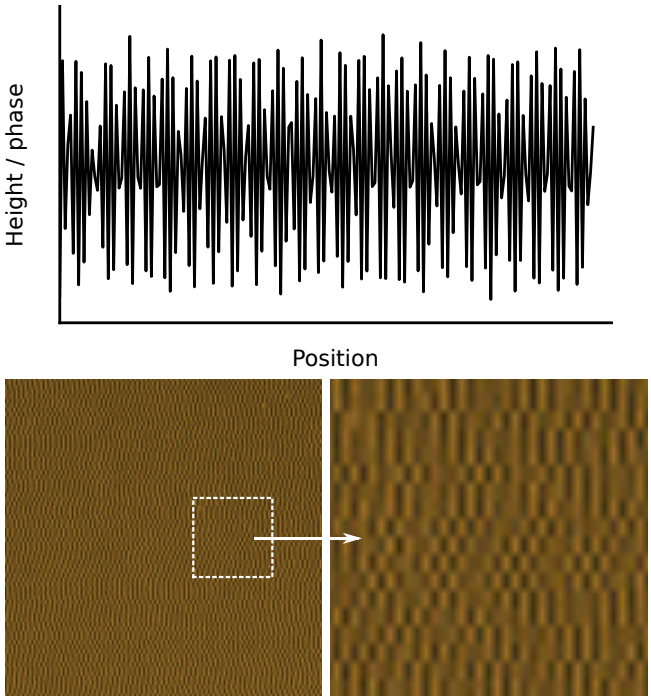

Figure 4: With too high gains, feedback loop may start to resonate. This results in a recognizable pattern in the final image.

Low gains situation in Figure 3). In this case it may be necessary to slow the scanning speed (`Scan rate` parameter) or increase the tapping force (see the next section).

1.3 Amplitude setpoint and drive amplitude

The `Amplitude setpoint` and `Drive amplitude` parameters control the tapping force.

The `Drive amplitude` parameter defines the amplitude of the signal oscillating the tip in the tapping mode. Higher values mean higher amplitude and — with constant `Amplitude setpoint` — higher tapping force. The initial `Drive amplitude` is determined by the `Target amplitude` parameter when the tip is tuned with the auto tune procedure. The auto tune tries to find a `Drive amplitude` value which gives the desired `Target amplitude` as the measured tip amplitude when the tip is far from the surface of the sample.

Commonly the `Drive amplitude` parameters is left as is and the tapping force is controlled by the `Amplitude setpoint` parameter. As discussed in Section 1.1, this is the value the feedback loop tries to catch by adjusting the z piezo position. If the value of `Amplitude setpoint` is decreased, the feedback loop will move the tip closer to the sample surface which results in more tapping force. Increasing the `Amplitude setpoint` has the opposite effect.

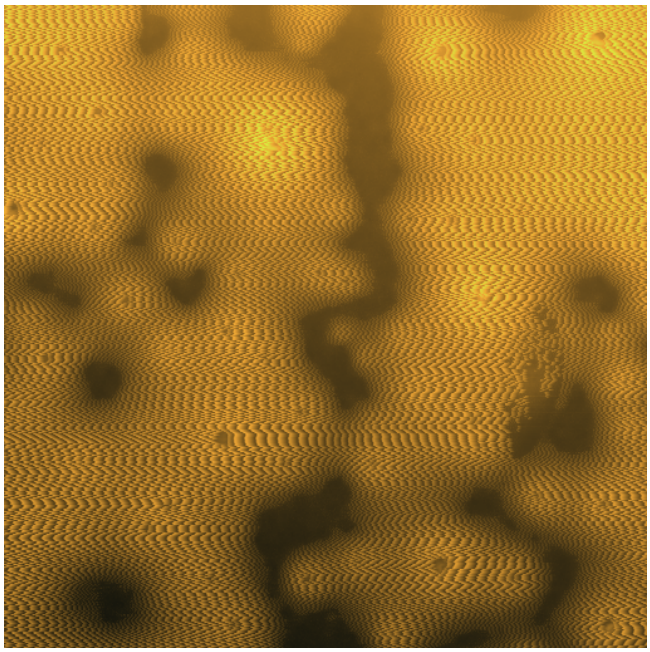

Figure 5: Another example of resonating feedback loop in a phase image.

In summary, the two parameters change the applied tapping force in the following way:

- `Amplitude setpoint` decreases $\rightarrow$ tapping force increases.
- `Drive amplitude` increases $\rightarrow$ tapping force increases.

Note, that increasing the tapping force may potentially damage the sample and the tip!

The initial `amplitude setpoint` is determined by the `Engage setpoint` parameter. When the tip is engaged to the sample, the so-called free amplitude is multiplied with `Engage setpoint` and this value is used as the `Setpoint amplitude`. The free amplitude in turn is ”sniffed” by the SPM during the engagement process. By definition, it is the measured amplitude when the tip is not in contact with the sample. Typical values for `Engage setpoint` are around $0.8-0.9$ which means that while scanning, the tip oscillates at $80-90\%$ of its free amplitude. This is sometimes termed as ”soft tapping” and produces a minimal force on the sample.

Below is a summary of the parameters and terms that affect the amplitude and thus the applied force in tapping mode:

Drive amplitude is the amplitude of the electric signal which drives the piezo that oscillates the tip.

Measured amplitude is the amplitude of the tip oscillations measured by the photodetector.

Target amplitude is the measured amplitude at which the tip should oscillate at the resonance frequency after the auto tune process. Determines the initial `Drive amplitude`.

Free amplitude is the measured amplitude when the tip is not in contact with the sample. Depends on `Drive amplitude`, tip dimensions and laser alignment.

Amplitude setpoint is the setpoint amplitude the feedback loop tries to reach during scanning by adjusting the z piezo.

Engage setpoint determines the initial `Amplitude setpoint` after the engagement process. The initial `Amplitude setpoint` is simply the multiplication of `Engage setpoint` and a free amplitude automatically "sniffed" by the SPM during the engagement process.

1.4 Optimizing the SPM in practice

The optimization of the AFM begins with the laser alignment. To maximize the sensitivity of the amplitude measurement, the laser needs to be aligned to the very free end of the cantilever where the amplitude of the tip oscillations are at maximum. In terms of the amplitudes, this maximizes the measured (free) amplitude at a given `Drive amplitude`. Note, however, that aligning the laser so that it "leaks" over the free end of the tip may result in artifacts during imaging because the laser may reflect from the sample surface and interfere with the laser reflected from the cantilever.

Next, it may be necessary to pay some attention to the `Target amplitude` parameter during the tuning process. For the Dimension 5000 SPM, the default value of 500 mV is a good starting point. A lower value, say, 250 mV is recommended for example when using high resolution tips (tip radius of curvature $\sim$1 nm) or when more tender contact with the sample is required. If the sample is especially tricky, for example in the case of soft and/or sticky polymers, it is possible to try higher `Target amplitude` values, such as 1000 mV.

If you feel like you always need to adjust the `Amplitude setpoint` during scanning for certain types of samples, it is possible to adjust the `Engage setpoint` parameter before engaging. Lowering the value will result in lower initial `Amplitude setpoint` and vice versa. Remember, that smaller values increase the tapping force and values above 0.9 may result in tip losing contact with the sample too easily.

Also, it is generally a good idea to set the `Scan size` to a small value, such as 1 μm, before engaging because it makes parameter optimization easier. Of course, if the optimal parameters are already known, one can directly start with the desired `Scan size`.

After engagement, setting the `Slow scan axis` parameter to disabled makes it easier to optimize the gains and forces because you can follow the effect of each adjustment directly from the Scan view.

In the easiest case, it suffices to adjust the `Integral gain` parameter only. Try to find the highest possible value where you do not observe any

resonances and no overshooting is apparent. This value is very sample dependent. Do not be surprised if you encounter values that feel unusually low or high.

The trace and retrace graphs should follow each other as close as possible.

After you find a satisfactory value for the `Integral gain`, set the `Proportional gain` parameter to the same value.

Sometimes the non-resonating `Integral gain` is so low that the tip does not stay in contact with the sample. In such case the easiest cure is to lower the `Amplitude setpoint` parameter which brings the tip in closer contact with the sample and increases the tapping force. Note, that the tip wears out during the scanning which usually results in more loose and sometimes even lost contact with the sample. Thus, to stay on the safe side it may be good idea to set the `Amplitude setpoint` to a slightly lower than the optimal value in any case. Remember to take into account the usual warnings about applying too much force though.

If you encounter problems with adjusting the `Amplitude setpoint`, another option to make the tip follow the sample surface is to decrease the `Scan rate` parameter. This decreases the `Scan speed` which in turn gives the feedback loop more time to react to changes in the amplitude error. The drawback is the increased scanning time.

When the image looks acceptable, change the `Scan size` parameter to the desired value. If the desired value is larger than initially, it may be necessary to adjust the parameters again since larger `Scan size` equals larger `Scan speed` which gives the feedback loop less time to work.

Remember to enable the `Slow scan axis` if you disabled it during the parameter optimization.

If you do not seem to find the correct parameters, realigning the laser or replacing the whole tip may help. You can also try to tune the tip again with different `Target amplitude`. Sometimes the sample-tip interactions move the tip resonance frequency to an unstable region. This you can quickly check while scanning by the tip tuning window. The current `Drive frequency` should be lower than the resonance frequency. If this is not the case, you can manually adjust the `Drive frequency` or disengage and retune the tip with larger `Peak offset` value.

Sometimes the sample is at fault. Soft substrates, such as mica, may bent when fixed on the stage by the vacuum. Also, exfoliation of mica may lead to air pockets which will give even under the extremely small pressure from the tip. Sometimes the sample may move under the tip. In these cases, observe the video view from the optical microscope. The sample should remain fixed and intact.

The optimal parameters are usually a compromise, because there are always trade offs. You may not get the best contrast in your phase image while taking accurate height data and vice versa. The feedback loop may resonate, but you are imaging so large height differences that the resonance can be considered as small noise. You want to image quickly, but that means your tip will ”fly” from time to time. Thus, the best compromise depends much on the sample you are imaging and a lot of other conditions and need to be considered in a case-to-case basis.

Sometimes nothing helps. Then it is better to go home, have a good

night's sleep and return on the next day, maybe with somebody more knowledgeable in SPM.

2 Notes on imaging

2.1 Scan size and pixel size

The scanning time is determined by two parameter: `Scan rate` and `Lines`. The `Scan rate` parameter tells how many lines will be scanned per second. It takes 512 s or 8 min 32 s to scan the 512 lines for the standard 512×512 pixel image with the usual 1 Hz `Scan rate`. Since maximum `Scan rate` parameter is usually limited by image quality considerations, the time to fully scan an image is controlled in practice by the `Lines` parameter.

The image resolution could be called the pixel resolution. The pixel resolution is the physical size of one pixel in the image. For the standard 512×512 pixel image taken with 1 μm `scan size`, one pixel corresponds to $1\,\mu\text{m}/512 \approx 2\,\text{nm}$. When using the standard tips with 10 nm radius of curvature, the smallest details shown in these images will be about five pixels wide.

There are some rules that may be good to follow to minimize the time spent at the microscope while getting the desired image quality:

- Think how large `Scan size` you need use to catch a comprehensive image of your sample which includes all desired details. After deciding the `Scan size`, optimize the `Scan rate` parameter. Higher `Scan rate` makes scanning faster but lowers image quality.
- It is generally good idea to set the `Lines` parameter to a low value (such as 64) to quickly check the scan area for dust particles etc. before recording the image.
- Think how large the features you are going to image. Remember, that the limiting factor in the lateral dimensions (x and y) is the sharpness of the tip (usually $\sim$10 nm). Set the `Lines` and `Samples/line` parameter accordingly.

For example, you may be imaging fibers which are around 50 nm in diameter and several micrometers long. To catch an entire fiber in one image, you would need, say, $5\,\mu\text{m} \times 5\,\mu\text{m}$ `Scan size`. With the standard 512 `Lines` and 512 `Samples/line` (512×512 pixels), you would get a pixel resolution of $\sim$10 nm/pixel. This means that the fibers would appear about five pixels wide. This should be enough to analyze, for instance, the length of the fibers. If you were interested in the thickness of the fibers instead, you could increase the image size to 1024×1024 pixels, which would increase the pixel resolution to $\sim$5 nm and give better looking statistic. On the other hand, this would double the scanning time and considering the tip radius of 10 nm, would the new pixel resolution really bring any benefits?

For some samples, if time permits it may be a good idea to scan 1024×1024 or 2048×2048 pixel images and later crop the interesting parts. Using the

largest possible `Lines` and `Samples/line` parameters in Dimension 5000 (5120!) is advisable only if your samples have features in several length scales and you need the possibility to really magnify your data. Scanning such large images takes 1 h 25 min and they take up a lot of hard disk space. Additionally, they do not even fit on any computer monitor without scaling.

2.2 Phase images

There are rumors that it is possible to deduce mechanical parameters such as softness of a sample from the phase data. This is generally not possible.

Sometimes taking a high contrast phase image needs different scanning parameters than a good height image. Play with the `Amplitude setpoint` and `Drive amplitude` parameters if you are unhappy with your phase images.

3 Artifacts

3.1 Common artifacts

We have already discussed some artifacts that may appear during scanning in the previous sections, namely the ”flying” tip effect, overshooting and feedback loop resonance. These can be avoided by optimizing the scanning parameters.

The usual suspects for scanning artifacts are things that appear horizontally or vertically in the images. These may be lines, bands or stretched shapes. To find out whether these features are real, change the `Scan angle` parameter by 45°. The features should also rotate by 45°. If this is not the case, you are most certainly dealing with artifacts.

The method below can also be used to detect tip damage on the sample.

1. Enlarge the `Scan size` two-fold.
2. Change the `Scan angle` parameter by 45°.

In a damaged sample, you should be able to see a diamond shaped outline of the damaging scan in the middle of the enlarged scan.

A common artifact is also creep, which happens every time the scanner moves the tip quickly in the lateral dimensions, for example when the `X offset` or `Y offset` parameters have been changed or when a new scan has been started from top or from bottom. Creep can be identified as stretched features in the image. The creeping should cease in a few minutes. However, if it continues for much longer, it may be that the sample is moving. This may be caused by inadequate attachment of the sample to the sample stage.

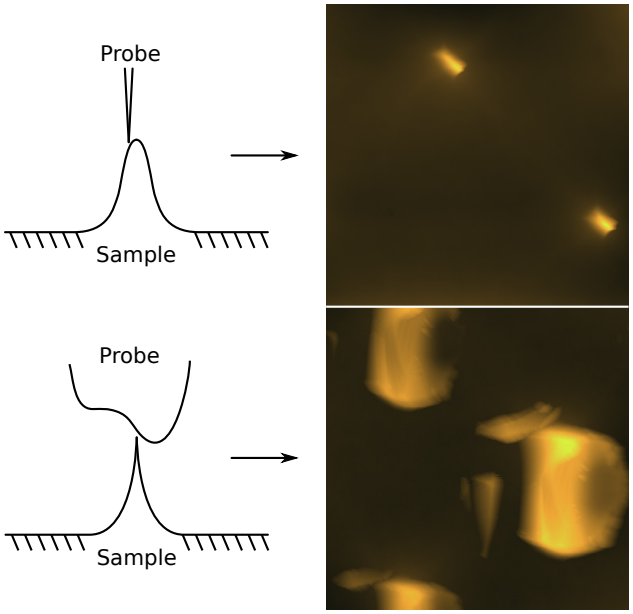

Figure 6: A tip should have smaller radius of curvature than the features it probes. The height images on the right are taken from a reference sample with spikes with radius of curvatures around 10 nm.

3.2 Artifacts due to online plane fitting

3.3 Tip artifacts

The resolution of an SPM is ultimately limited by the tip. Thus, it is important to have a tip which is sharper than the features of interest in a sample. Also, it is important to recognize artifact caused by the tip.

Reference samples can be used to quickly check the sharpness of tip. Figure 6 shows two height images taken from such a sample a sharp and a dull tip. The differences should be mind-opening. It is a good idea to quickly image the reference sample before and after imaging the real sample/samples to make sure the tip stayed sharp throughout the imaging process.

The dull tip image in Figure 6 is also an extreme example of the double tip effect. The sharp spikes in the reference sample actually cause the SPM to image the tip itself. More usual examples of the double tip effect are shown in Figure 7. The repetitive features of a double tip are apparent in the top and middle images. However, from the bottom image of Figure 7, it is difficult to say whether the repeated pattern is due to a tip artifact or not.

It is desirable to recognize the artifacts caused by the double tip effect as early as possible to avoid wasting time. Again, you may check that the tip is really at fault with the reference sample before changing a new tip.

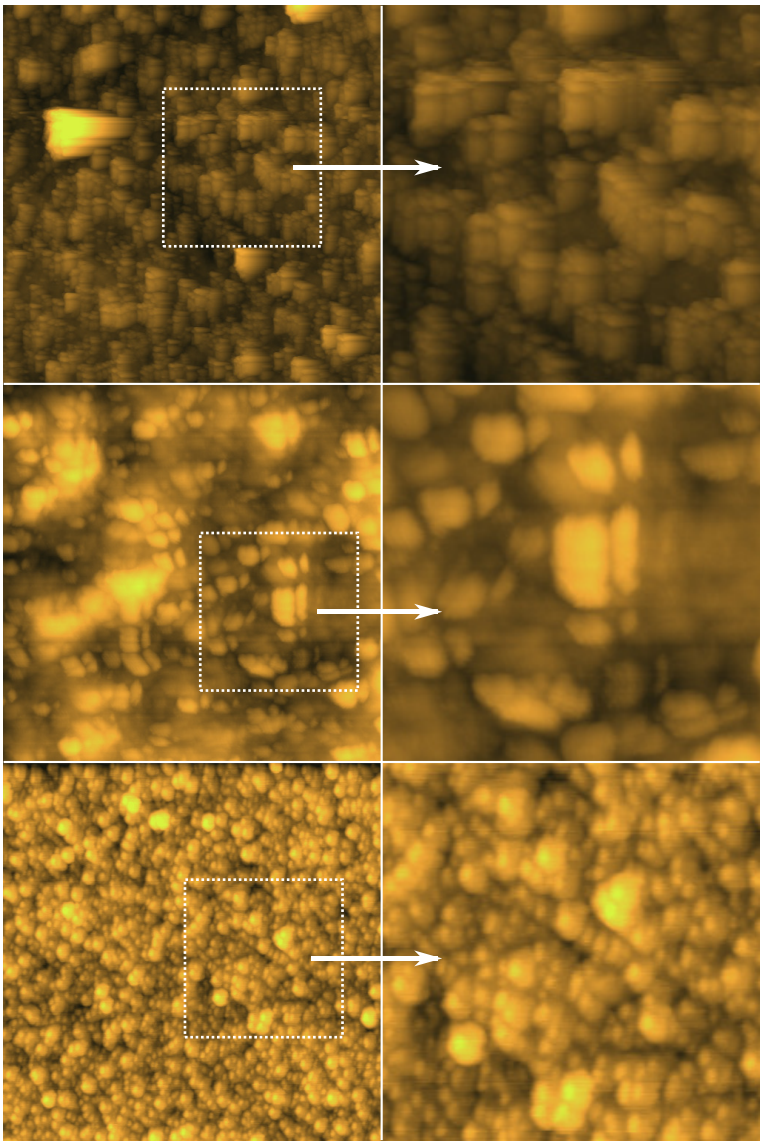

Figure 7: A series of images which may have been imaged with a double tip. In the top and middle images the effect is apparent. However, in the bottom image, it is difficult to say whether the repetitive pattern is due to the tip or if it is the real sample.

Figure 8: Subtracting a first order polynomial background from a raw image does not usually produce flat substrate. Second order polynomial background, however, seems to do the trick in this case.

4 Image post-processing

4.1 Flattening and background subtraction

Usually it is necessary to apply some post-processing corrections to the SPM images. The sample substrate may not have been even or the sample was not in plane with the microscope. Sometimes individual lines or whole horizontal bands are out-of-plane due to tip wear out. A dust particle may stick onto the tip which alters the tip height making the sample appear suddenly at lower heights. And so on.

There are basically two different groups of methods for image correction: line corrections and background subtractions.

Line correction methods (also known as flattening) are used to correct horizontal lines, bands and steps in an image. They are based on some algorithm, which tries to match consecutive lines. For example, a median line matching algorithm could evaluate the median height value of each horizontal line and shift the heights of the lines until the medians match.

Methods based on background subtraction correct the image as a whole. Usually the subtracted background is a simple polynomial plane.

Figure 8 shows examples of the polynomial subtraction. In the raw image, it is clear that the substrate was not in plane with the SPM: the lower part of the supposedly flat substrate is higher than the upper part. To correct for the tilt, a first order (linear) background has been subtracted from the raw image in the middle image of Figure 8. However, the substrate still does not appear completely flat. There is an evident horizontal ”valley” in the middle of the image. This may be due to nonlinear behavior of the AFM and can be corrected by subtracting a second order polynomial background from the raw image. As the rightmost image in Figure 8 shows, this gives satisfactory results.

The polynomial subtraction as well as the line correction methods work well only if there is enough of plain flat substrate in the image. A mask can be applied on features which are not desired to be included in the correction process. Figure 9 exemplifies this. If a second order polynomial background is directly subtracted from the raw image, an even

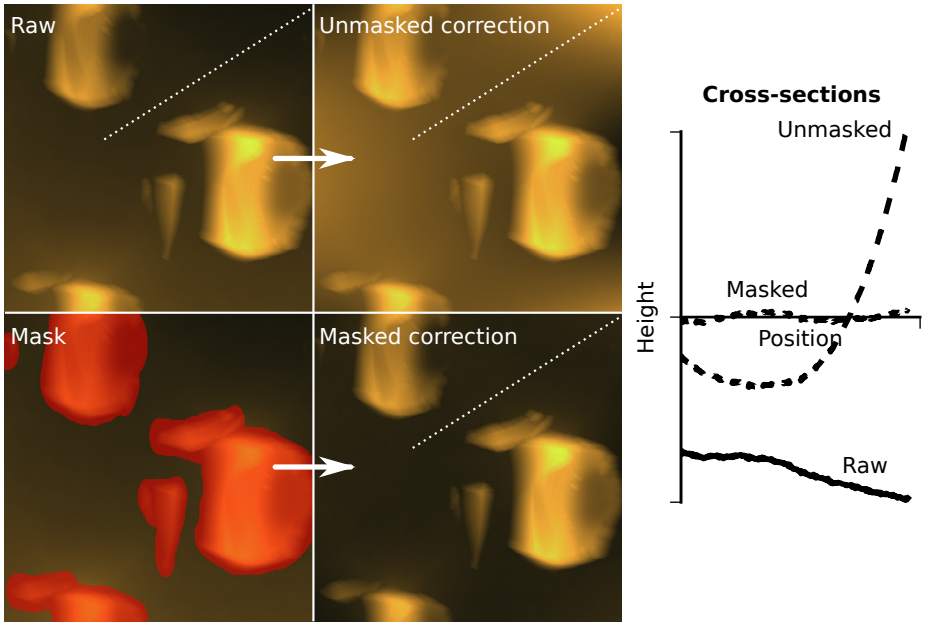

Figure 9: Direct correction (in this case a second order polynomial background subtraction) of a raw image may leave some artifact on the image: a cross-section of supposedly flat substrate is actually more uneven than in the raw image! By masking some features from the image, the corrected image looks much better.

more uneven image results. This should be evident when comparing the cross-sections of the raw and corrected-but-unmasked images in Figure 9. Masking the non-substrate features (the red areas in Figure 9), on the other hand, produces the expected flat substrate.

4.2 Data visualization

A simple rule for data visualization: an optimal image gives the viewer the maximum amount of information in minimum time using minimum amount of ink. In practice, this equals to simple but informative images.

This rule applies to SPM images as well. You may have noticed the absence of scale bars in the SPM images in this guide. I didn't include them because, in my opinion, they are not necessary to get to the point. In scientific publications people often unnecessarily include the phase images alongside the height images even though the phase data really doesn't bring anything new to the discussion. Same goes for the scale bars in phase images since phase is not really a simple physical quantity as, for example, height.

Aesthetics is another point to consider when visualizing your work. Figure 10 shows two SPM images (which were corrected by second order polynomial subtraction) with different color gradients. The simple single-color gradient gives a neutral view on the sample. The two-color version of the same image highlights the lumps on the substrate. Even

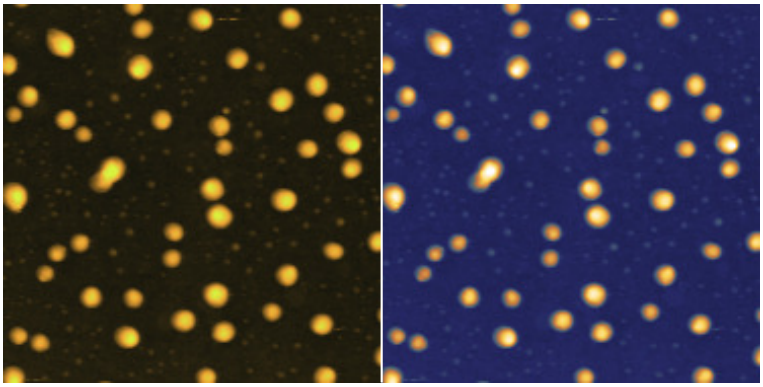

Figure 10: Images can be colored with regular gradients for good readability or with fancier colorings to, for example, highlight some features.

fancier coloring could be used to highlight the smaller bumps on the surface too. Be careful not to overdo it, though. If it takes more than 10 minutes to choose your favorite gradient you should consider reverting back to the software's default.

Some people like cluttering their articles with 3D images. In some cases the 3D images may be a good for visualization. The problem is, however, that usually the z scale is exaggerated by default in the image processing software. This leads to the situation shown in Figure 11. It may seem that the bumps on the sample surface are spike-like while, in fact, they are pretty flat compared to the lateral dimensions of the image.

It is easy to make the 3D images more eye catching by applying some lightning effects as shown in Figure 11. This, of course, renders the images to qualitative only but sometimes that is okay. It is also possible to use for example the phase image as an overlay for the 3D height image thus combining both height and phase information. Some SPM image processing software readily offer this function but it can be accomplished also for example by using a 3D modelling software such as Blender.

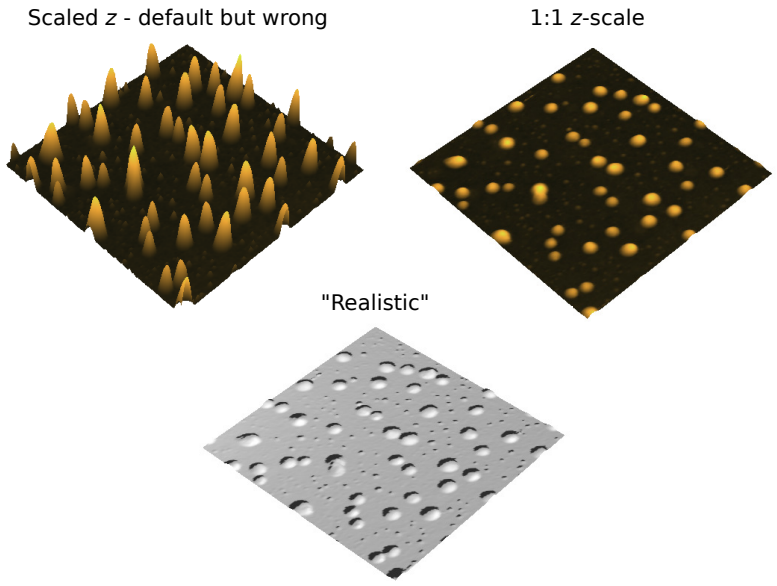

Figure 11: Many SPM software exaggerate the z-scale in 3D views. A 1:1 correspondence between the z and the lateral dimensions gives a more truthful picture of the sample. There are tools to give the image a more real-life look by applying shading to the image.