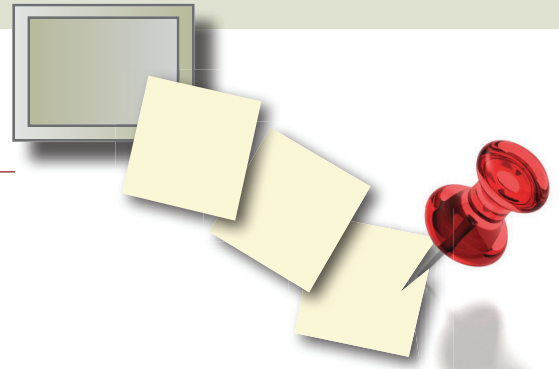


NetNotes

Bob Price

University of South Carolina School of Medicine

Bob.Price@uscmcd.sc.edu



Selected postings are from recent discussion threads included in the Microscopy (<http://www.microscopy.com>), Confocal Microscopy (<https://lists.umn.edu/cgi-bin/wa?A0=confocalmicroscopy>), and 3DEM (<https://mail.ncmir.ucsd.edu/mailman/listinfo/3dem>) listservers. Postings may have been edited to conserve space or for clarity. Complete listings and subscription information can be found at the above websites.

Zemax Simulations and Microscope Objectives

Confocal Listserver

To enable Zemax simulations for a customized two-photon microscope (university research application), I am looking for the prescription for the Olympus 10× 0.6 NA objective, model: XLPLN10XSVMVP 3mm WD. The prescription can usually be found in the patent submitted by Olympus Corporation. Has anyone been able to find this prescription/patent? Thank you. **Nicholas Watson** nicholas.watson@utexas.edu

I have a similar request for the prescription for the Zeiss LD Achromplan 20× 0.40 Korr objective. It is quite an old lens, and I do not know where to start looking for details. The idea is to use it as a water-immersion lens for multiphoton microscopy (so the chromatic aberration is not an issue). The issue is the very long working distance (11 mm in air, some 7 mm in water) making it a bit impractical for some applications. Thanks! **Zdenek Svindrych** zdedenn@gmail.com

You can search <https://sites.google.com/site/danreiley/MicroscopeObjectives> to find the same or similar objective lens. **Lingbo Kong (Bob)** lbkong2@gmail.com

I have also used the website that Bob recommended, and it is pretty good. That said, I have found that objectives are corrected enough that a simple paraxial lens with a radius matching the back aperture of the objective is sufficient in most cases, unless you are trying to do non-sequential modelling. One thing I would love, though, would be some models of immersion objectives. **Benjamin Smith** benjamin.smith@berkeley.edu

For people looking for microscope objective “prescriptions” or design criteria these following links are an incredible resource.

<https://www.degruyter.com/document/doi/10.1515/aot-2019-0002/html>

<https://www.degruyter.com/document/doi/10.1515/aot-2019-0013/html>

<https://www.degruyter.com/document/doi/10.1515/aot-2019-0014/html>

Stan Schwartz sschwartz@saikouoptics.com

Are Lower Magnification Objectives Brighter?

Confocal Listserver

Dear all, for confocal microscopy are lower magnification objectives brighter than higher magnification ones when they have the same

NA, for example, a 40× NA 1.4 objective compared to 63× NA 1.4? Confocal.nl stated this is in a recent webinar and on their website: “A lower magnification allows for a larger field of view and brighter images, since light intensity is inversely proportional to the magnification squared.” (<https://www.confocal.nl/#rcm2>). I would think that this is caused by less light going through the smaller back focal aperture when the illumination is held constant? Most of the light is clipped as explained in figure 1 of <https://www.nature.com/articles/s41596-020-0313-9>. So, the microscope manufacturer could adjust the illumination beam path and laser powers to best suit the objective? Or are lower magnification objectives really brighter? The field of view will obviously be larger for the 40× objective, but I am more interested in understanding the claimed benefit in brightness.

Andreas Bruckbauer a.bruckbauer@imperial.ac.uk


I teach my students that the light gathering power of an objective depends on the NA raised to the power of four divided by its magnification squared. I think the concept comes from Shinya Inoué in his book “Video Microscopy.” Not sure if the confocal setup is changing this principle as it may not always make full use of the NA (for example when using a beam expander). **Christoph Bauer** christoph.bauer@unige.ch

I have always understood this in relationship to a constant detector with a given pixel size (like a camera): lower magnification spreads the same signal over a smaller number of pixels, resulting in higher intensity for the pixels that contain the signal. This is tricky with point-scanning microscopes. **Christophe Letierrier** christophe.letierrier@gmail.com

It has bothered me for many years that people still claim that a CLSM gives brighter images when using a lower magnification objective (for the same NA). Physically, it does not make sense. I have both a 63x/1.4NA and a 40x/1.4NA on the same Zeiss LSM700 confocal. Considering the focused spot on a CLSM, the size of the PSF depends only on the NA of the objective and not its magnification, so the illumination will be identical for a 40x and a 63x objective with the same NA (assuming that overfill of the back aperture is overfilled in both cases to take full advantage of the NA of the lens). Now consider the detection: again, only the NA determines how much light is collected by the lens. So, it would not make sense for a CLSM to give a “brighter” image with a lower mag lens when both lenses have the same NA. But wait! When you look into the eyepieces it looks brighter with the 40x lens. AND, if you keep all of the same settings (laser power percentage and detector gain) you get a brighter image with the 40x objective. So, what’s going on? My relatively new Thorlabs power meter (PM400 console with S170C sensor) is compatible with oil immersion and the difference in brightness with the 40x objective is 100% accounted for by the change in laser power when switching between these objectives. The change in laser power is due to the smaller back aperture of the 63x objective. In other words, when you switch from the 40x to the 63x objective, the edges of the laser beam are blocked by the smaller aperture of the 63x lens, so less excitation reaches the sample. If you



microPREP™ PRO High-Speed Sample Preparation



www.microPREP.pro

- » Free-form laser-based sample preparation
- » Larger-sized samples with micron-level precision
- » Workflows for TEM, XRM, APT, and micromechanical testing
- » Reduces time-to-sample and costs significantly

Contact us now!

Tel: +49 371 40043-222
sales@3d-micromac.com



adjust the % laser power slider so that both the 40x and 63x objectives are reading the same illumination intensity, then you get the exact same image brightness with both lenses.

As you mentioned, I tried to explain this in our *Nature Protocols* paper in Supplementary Figure 1, and I included some of the data there (free download for the Supp Figs - for the full paper if anyone needs it I'm happy to email it to them). <https://www.nature.com/articles/s41596-020-0313-9>. So why is this so broadly misunderstood (I have heard it many, many times!)? When we read the classic textbooks on the brightness of a microscope image, these were originally written with respect to transmitted-light brightfield microscopy: it is not obvious that they should apply to confocal microscopy or even to widefield fluorescence microscopy. On the Microscopy Primer website (<https://www.microscopyu.com/microscopy-basics/image-brightness>), for example, they start with the typical statement that image brightness is proportional to $(NA/M)^2$. They go on to mention that for fluorescence the image brightness should be $\lambda NA^4 / M^2$. However, they fail to mention that the reason for the magnification being in the denominator of the equation is because the size of the back aperture depends on magnification in this way. So even for a widefield fluorescence microscope, the increase in brightness is caused by increased illumination on the sample, not increased detection efficiency, which is not very helpful in this era of over-powered fluorescence lamps. If the confocal manufacturers would specify their laser powers in real-world units instead of percent of maximum, when switching lenses, it would immediately be seen that that for a given excitation power density (in W/cm^2), the same intensity image for 2 lenses with the same NA is obtained, regardless of the magnification of the lens. **James Jonkman** james@aomf.ca

Coatings and lens elements for correcting chroma, flat field, etc., must also be accounted for as each element & coating will necessitate some light loss. I recall the brightest objective in one scope maker's product line (at the time that I asked) was a 40x1.3 with modest chroma correction that hit a sweet spot between magnification, NA, and the number of lens elements & coatings. That makes me wonder whether microscopy has anything like the famous 50mm/1.8 lenses for photography. That length & aperture hit an engineering 'sweet spot' that provided near-perfect optical quality with a minimum number of all-spherical lens elements. Every manual SLR camera used to come with one because the makers could all produce amazing ones in a compact size and basically for free. If such a thing exists in microscopy, I have not heard about it. **Timothy Feinstein** tnf8@pitt.edu

Since you mentioned transmission brightfield, it is also a widely misunderstood topic because brightness is determined mostly by direct light from the condenser and not by diffracted light, so NA of the objective does not matter at all as long as it is larger than the NA of the condenser. In other words, brightness is determined by the smallest NA between the objective and condenser. This can be easily verified using an objective with variable NA. Misstatements on this subject can be found even in some of the classical treatises. **Mike Model** mmodel@kent.edu

Great point, Mike. I feel very comfortable comparing 2 lenses with the same NA and both of them oil immersion, but I'm not certain that the Thorlabs power meter (despite being compatible with oil) necessarily captures 100% of the laser power from the oil objective. The highest angle rays may not hit the sensor – it is difficult to know - and that may account for some of the decrease. Therefore, I try to benchmark the laser powers using a 10x dry objective on all of my microscopes. Some of the decrease may also be because of lower transmission through the oil objective which presumably has more glass. But the bulk of the decrease is because of the

smaller back aperture of the 63x objective. I am also very careful to stop the beam from scanning when I make these measurements. The actual laser power is considerably higher than that measured when the beam is scanning because the confocal AOTFs blank the beam on fly-back and when changing direction. Choosing a higher zoom does not help – it is necessary to do point scanning or point bleaching to get the actual power measurement. **James Jonkman** james.jonkman@uhnresearch.ca

Pina Colarruso and I helped Thorlabs develop that power meter, and our big contribution to the design was ensuring it does, in fact, account for the high angle rays. There is a thin layer of index matching material underneath the glass window that eases the high-angle light into the silicone detector underneath the window. You can test this yourself by measuring a high NA oil lens with and without oil on the sensor. **Craig Brideau** craig.brideau@gmail.com

Thanks for the detailed explanations. This all makes sense to me now. While this was initially only intended for confocal, I did a simple experiment with the widefield microscope, comparing 20x NA 0.8 and 40x NA 0.75 objectives. The images were taken with the same pixel size (2 times binning on the 40x) and the same region (cropping for the 20x). Same LED power and acquisition time settings. Interestingly, the fluorescence intensity of the larger magnification 40x was 1.8x higher!!! When measuring the LED power, it was 2x higher out of the 20x objective. I think the 2x higher LED power is spread over a 4x larger area in case of the 20x objective, so that the power density is half compared to the 40x objective, leading to the lower fluorescence intensity of the image with the 20x objective. The difference between the measured 1.8 and 2.0 could be assigned to the difference in NA^2 and probably slight differences in transmission. **Andreas Bruckbauer** a.bruckbauer@imperial.ac.uk

Thanks for the interesting discussion. Could one then summarize by saying that, when acquiring images with WF, confocal or multiphoton and with 2 objectives of different magnifications but identical NA, coatings and immersion medium, the image is expected to be brighter with the lowest magnification objective, even if it is for different reasons for the 3 types of microscopes? **Sylvie Le Guyader** sylvie.le.guyader@ki.se

Hi, Sylvie. Who would have thought that it could get so complicated - thanks for trying to summarize what has become a lengthy but interesting discussion! Personally, I don't like the word "brighter", because it implies that this is the desired outcome. For laser-scanning confocals, I feel it is misleading to tell a user that a lower mag objective gives you a brighter image, when it's definitely only "brighter" in the sense that you're hitting it with more excitation light, so it is not an advantage at all. With CLSM I prefer to avoid telling people that lower mag=brighter, and instead emphasize that higher NA=more efficient light collection (2x NA gives you 4x more light collection). There have been many other interesting comments on this thread. I liked Andreas's widefield measurements. I had not thought of binning to get the same pixel size. The old formulae that originally inspired people to think that 'lower mag=brighter' were of course developed if you were observing the sample in the binocular; and you cannot bin your photoreceptors, so what is the fairest comparison? As you wish I suppose! **James Jonkman** james.jonkman@uhnresearch.ca

Issues with RELION and cryoSPARC

3D Listserver


I am trying to use RELION to Autopick (Laplacian) a set of elongated particles that are around 150Å long. I managed to get a nice set of 2D classes in cryoSPARC, but, since I have preferential orientation, people have

Show Your FIB System Some Love

High quality extractors, suppressors and aperture strips from Oregon Physics

All parts are made to exacting quality standards for guaranteed performance in your system.

- For FEI/Thermo Fisher Scientific FIB columns
- Reduced costs with volume discounts
- Scheduled deliveries available
- Fast turnaround for urgent requests
- Standard aperture strip configurations
- Made-to-order aperture strips in custom configurations



OREGON
PHYSICS

+1 503 601 0041
info@oregon-physics.com
www.oregon-physics.com


X-Cite[®]
NOVEM



**More Power.
More Wavelengths.
More X-Cite.**

- Full spectral coverage from Fura-2 to IR800
- 9-channel wavelength selection
- LaserLED Hybrid Drive[®] with 4-position motorized filter changer
- Pre-installed clean-up filters
- Whisper quiet operation

Win your own Sam & Ella!



Send us feedback or images generated using any X-Cite system and be entered into a future draw for one of our new plush X-Cite microbes!
Email us at: kavita.aswani@excelitas.com.

X-Cite[®]
Fluorescence Illumination • In Control

www.excelitas.com
x-cite@excelitas.com

EXCELITAS
TECHNOLOGIES[®]

recommended the use of RELION to redo my dataset analysis. But... I'm really stuck, I have no experience with RELION, and I cannot get good autopicking that leads to a 2D classification (even a preliminary one to redo the picking). I just get black images or nonsense in the 2D classification, even when I check the images after the autopicking and they don't look bad. Honestly, I do not know what to put here about the parameters that I am using because I have tried several, but blindly (since I have no experience). I hope someone can provide suggestions because I am desperate. Thanks!
Lucia Torres Sanchez luciatorressanchez3@gmail.com

Have you gone through the RELION tutorial? (see https://www3.mrc-lmb.cam.ac.uk/reliion/index.php/Main_Page). That is usually a good place to start learning RELION. **Sjors Scheres** scheres@mrc-lmb.cam.ac.uk

Don't be desperate. You are in the initial phase of learning how to use RELION and cryoSPARC, so frustration might be normal. Preferential orientation is an interesting topic, and it might not impede your efforts for high-resolution density maps. The suggestion of using RELION to circumvent orientation bias is funny. In RELION you might need to play around with the parameters to pick particles, better you do it with a small set of micrographs. You can also continue with cryoSPARC and maybe change the picking parameters to see whether you really do not have other 2D classes. Go to ab-initio (abinitio.com) with your best 2D-classes to determine what you have. Good luck!
Jacopo Marino jacopo.marino@psi.ch

To me it does not look like you have a severe orientation bias (if you have it at all), and you have plenty of particles! It might be some views are hidden within other classes. How many rounds of 2D-classes and selection have you used? For ab-initio, you might play around with parameters (2 classes, zero similarity, higher number of final iterations, etc.). Look on the cryoSPARC forum for similar discussions. Sometimes you need to play with the ab-initio till you get a volume that makes sense, then take the particles that come with it and go to refinement. If you want help with the parameters for RELION, you need to post a screenshot of your GUI with the parameters. I agree that you have enough views, and the problem may be in the generation of the initial volume. You may try methods that work with class representatives, rather than raw images. In that way you get away from the preferred views. RANSAC (random sample consensus) and Reconstruct Significant from Xmipp are two of these methods. You can access them through Scipion. **Carlos Oscar Sorzano** coss@cnb.csic.es

Sorry to hear about your troubles. I agree with Jacopo that the suggestion to use a different program to overcome orientation bias of your particles is not very helpful without more specific suggestions. I think your initial 3D model in cryoSPARC looks decent. Have you looked at it in a 3D viewer like ChimeraX? Does it make sense? A good initial model should have connected density (multiple disconnected blobs is not good), and if the overall shape corresponds to what you see in the 2D classes (which seems to be the case here, based on what you showed), you should try to use this initial model to run 3D classification ("heterogeneous 3D refinement" in cryoSPARC vocabulary). The number of 3D classes to request depends on how many different species you think are present in the set of particles. I recommend starting with a small number of classes like 4, see if the results make sense, and if not try with a larger number of classes (maybe up to 10). The above suggestions should be easily doable from where you are now in cryoSPARC. That said, if you really are missing orientations, you do need to go back to picking. The most effective way to find particles with rare orientations in the micrographs is to use a neural net particle

picker like crYOLO or Topaz. cryoSPARC has an interface to Topaz (I think you still need to install Topaz, and when you have it indicate in cryoSPARC where to find the program), so it is probably the easiest way to go in your case. You did not say how you performed particle picking in cryoSPARC. I am assuming it was with template matching (manually pick enough particles to get good 2D classes, then use these 2D classes as references for automated picking). In many published cases, and from my own experience, processing a few datasets for 3D reconstructions can improve a lot when picking with Topaz or crYOLO, compared to template matching or LoG. Neural net pickers seem a lot better at not only finding good particles, but also avoiding bad ones, which helps a lot overall (even a small proportion of bad particles sometimes easily messes with 2D and 3D classification, so the most robust way to deal with these is to avoid picking them in the first place). Starting over from picking is a lot of work, so you might as well use the most robust method while you are at it. I hope this helps, good luck with your project.
Guillaume Gaullier guillaume.gaullier@icm.uu.se

From your 2D classification and 3D refinement results in cryoSPARC, I would say your project is very promising. You have so many particles and relatively enough views for the 3D reconstruction. I totally agree with Guillaume's suggestions above. What you can do is to build 4 or more initial models with it in cryoSPARC and use them as references for heterogenous refinement. I guess you will get at least one good class in which the reconstruction is less anisotropic. **Jun Yu** jun.yu@unige.ch

I also think your project is very promising. If preferred orientation is the only problem, you do not have to reprocess the data, instead, you can simply remove the particles in dominant views. Our lab has a nice utility tool that can remove particles in dominant view with a single command: `images2star.py input.star output.star --normEulerDist 5 -1 --verbose 3`. You do have to install jspr to use it but it is pretty easy to install by downloading and unpacking the package at the bottom of the webpage: <https://jiang.bio.purdue.edu/jspr/> **Chen Sun** sun647@purdue.edu

Computer Node for RELION and cryoSPARC 3D Listserver

Hello everyone, I am thinking of acquiring hardware for computer nodes mostly used for cryoSPARC and RELION with Sturm and hope someone can give me advice. I was thinking something like the G292-Z43 servers from Gigabyte with 2-4 RTX 3090, 64-128 GB RAM and a NVMe SSD for cache. I am not sure which CPU to get though I was thinking of the 2x AMD Epyc 7543 32c/64t, 2.8-3.7 GHz, 256 MB L3, (~\$3,700/each) or should it have higher frequency and fewer cores (- 2x AMD Epyc 74F3 24c/48t, 3.2-4.0 GHz, 256 MB L3, ~\$2,900/each)? Or even just a single socket CPU with a different server like the 7532P? I would have to calculate the PCIe lanes I need but would PCIe 4.0 $\times 8$ or $\times 16$ for the GPUs even make a difference? Is it worth going for the 256 MB L3 cache? Or is it worth going for Team Blue CPUs because of AVX 512? Any insight would be highly appreciated. Best wishes. **Kilian Schelle** kilian.schelle@uni-osnabrueck.de

Regarding the number of cores versus frequency, I would suggest going for more cores: some job types in RELION are not GPU-accelerated but scale very well with more MPI processes (motion correction, Bayesian polishing and CTF refinement), so for those you will benefit from having many cores. And 2.8-3.7 GHz is plenty. Other general advice would be: Do not be cheap with RAM and storage, both in terms of speed and amount. Fast CPUs and GPUs are no use if you cannot feed them your data fast enough (they would spend a lot of their time waiting for input). The amount of RAM is also easy to overlook, thinking you will only use RELION and cryoSPARC, which can both

SEM Scintillators & Light Guides

Light Guide Recoating Services too!



Conductive Adhesives, Coatings & Tabs!

Introducing GRAPHENE The New Carbon! Brush or Spray the same product!



SEM Micro
M. E. TAYLOR ENGINEERING INC
SINCE 1972

www.semsupplies.com • (301) 975-9798

Refurbishing or servicing an electron microscope? Scope us out for your ion pump needs.



Enjoy Virtual M&M! See you next year!

New | Rebuilt | Rebuilding Services

DUNIWAY

STOCKROOM CORP.

800.446.8811
www.duniway.com

TGS Technologies, LLC

TEM Sample Holder Service and Modification Specialist

With over 40 years of experience repairing TEM Sample Holders, we provide you with high-quality personalized service. We specialize in repair, modification, and refurbishment of your sample holder as well as custom design if you are seeking specific innovative needs from your equipment.

TGS Technologies, LLC
Ph: 724-453-3865
Fx: 724-453-2968
Email: tom@tgstechnologies.net
<http://www.tgstechnologies.net>

Oxxius
Simply Light

UV, visible and NIR lasers manufacturer

FLEXIBILITY matters

Laser Combiner solutions

COMPACT



Economical version with all embedded functionalities

MODULAR



Upgrade your combiner over time

ADVANCED FUNCTIONALITIES



Get more from your combiner with the extension modules

- Fast switch mirror
- Balanced dual output
- Motorized neutral density

Up to 7 laser lines and 4 optical fiber outputs

MICROSCOPY SPECTROSCOPY
HOLOGRAPHY INDUSTRIAL

www.oxxius.com



use an SSD cache efficiently, but newer programs do not always have this capability (I am thinking about cryoDRGN in particular), and it would be frustrating to have to wait until they implement this or find workarounds (RAM upgrade, or not using your entire dataset, etc.). Last year I bought a workstation with what seemed like a ridiculous amount of RAM, and this year I am thinking it was a good idea since I have been happily running cryoDRGN with large box sizes and many particles without any problem. [Guillaume Gaullier guillaume.gaullier@icm.uu.se](mailto:Guillaume.Gaullier@icm.uu.se)

My apologies for jumping into this conversation with something outside the scope of the question, but Kilian's question about AVX-512 considerations in a build is related to an issue our lab is having with a workstation. We recently purchased a workstation with 2 RTX 3090s, 128GB RAM, and a 10980XE processor on an ASUS x299 SAGE 10GbE. Since the start, it's had issues with running 2D and 3D classification and refinement jobs, primarily in cryoSPARC or programs using a CryoSPARC engine (that is, Scipion). The system will get 3-5 iterations into a job before spontaneously restarting. It does not post a kernel panic error, nor does it post any errors in system logs besides orphaned processes and unrelated process errors from the crash. Normal stress tests of both CPU and GPU have not returned any issues. I have traced the issue, potentially, to an issue with AVX-512 processing causing the CPU to trip some cutoff and shut down the system. Changing some overclocking settings in the BIOS to limit voltage/clock speed has helped, but the issue still happens, just not as soon in the job (iteration ~15). Does anyone have a recommended general BIOS configuration for computers doing cryo-EM processing with CryoSPARC and related programs? [Justas Rodarte jrodarte@fredhutch.org](mailto:Justas.Rodarte@fredhutch.org)

Is a ridiculous amount of RAM 512Gb? [Israel Fernandez israel.elotro@gmail.com](mailto:Israel.Fernandez@icm.uu.se)

In my case it was 768 GB total (64 GB * 12 slots in a Supermicro case). But yes, 512 GB qualifies as a "ridiculous amount" in my book. You could also call it "future proof" (as in future programs implementing more sophisticated analyses than we use now may need more RAM), which is a better vocabulary to use as a justification for your spending. [Guillaume Gaullier guillaume.gaullier@icm.uu.se](mailto:Guillaume.Gaullier@icm.uu.se)

Data Storage

Microscopy Listserver

I am requesting suggestions and cost estimates for off-the-shelf data storage systems to store raw cryo-EM movies and processed data. Our initial target is 150-200 TB with options to expand. We do not have much local IT support for Linux-based systems, which is why I am asking for an off-the-shelf system which should be easier to install and manage. [Krishan Pandey krishan.pandey@health.slu.se](mailto:Krishan.Pandey@health.slu.se)

If you want something requiring minimal IT support but with good performance, I have purchased several of these (we use them mostly with individual workstations, then have traditional rack-mounted RAID's for the larger archives): Synology 12 bay NAS DiskStation DS2419+ (https://www.amazon.com/Synology-Bay-DiskStation-DS2419-Diskless/dp/B07NF987TP/ref=sr_1_3). This costs ~\$1,500 and to add a 10GB network card to it (~\$150) it can easily do ~950 MB/s read speed, and about 600 MB/s write speed (if you have a 10GB network to plug it into). Load it with 16 TB drives (~\$4,500) and you have ~150 TB of usable space (configured as a RAID 6). It is all web-driven and user-friendly. It is also usable in an office environment (fairly quiet, no jet-speed fans). So ~150 TB for \$6,000. It is also possible to supplement it with a second 12-bay JBOD box to double the

capacity. Do not get me wrong, this is not an optimal solution for >500 TB, but it is pretty decent for a good sized cryoEM lab. [Steven Ludtke sludtke@bcm.edu](mailto:Steven.Ludtke@bcm.edu)

I concur with Steven. I have a similar (though 8 bays only) Synology NAS system, equipped with 8 Tb drives and running in RAID 6. I bought it 5 years ago for less than \$3,000. It was easy to configure with almost nothing to do on a daily basis, and you can buy several of them and connect them together very easily. [Sylvain Trepout sylvain.trepout@curie.fr](mailto:Sylvain.Trepout@curie.fr)

RAID-NAS solutions as proposed by Tim and Steven are appropriate for most per-research group storage. Of course, this depends on the use and number of clients; if you have many computers running IO heavy jobs simultaneously, the storage can be a bottleneck. For such cases, a distributed file system is more suitable, but "manual load balancing" to two NASes (for example, data collected on odd days in NAS 1 and even days in NAS 2) also works. Even if you use RAID-6, you should make at least one backup copy of raw movies outside the NAS to protect from "rm -fr" mistakes and malware. LTO tapes are reliable, but drives are expensive and tricky to use. "HDDs on a shelf" are far from ideal, but often the only realistic backup solution. It is better to have a backup than nothing! In any case, one should keep the list of files in each disk/tape somewhere online; otherwise, you must mount and inspect media one-by-one to find a dataset. This is very cumbersome and almost impossible when the student/post-doc who made the backup leaves the lab. I think the above strategy is sufficient to keep data for 3 to 5 years after creation which should be long enough to publish a paper. After publication, you can upload raw movies to EMPIAR. It is free, safe (data are mirrored to Japan and China) and you contribute to open science. [Takanori Nakane tnakane@mrc-lmb.cam.ac.uk](mailto:Takanori.Nakane@mrc-lmb.cam.ac.uk)

Our Chemistry department purchased a 500TB 4U rack from truenas.com, https://www.truenas.com/x-series/?__hstc=216824393.ae734f6409fa54117a8ea5c71d99c7cf.1617304740425.1617304740425.1617304740425.1&__hssc=216824393.1.1617304740425&__hsfp=3483769911. We paid less than 50k Euro, but this was pre-COVID-19. Such systems don't need much maintenance. Once they are set up, they just run. This system runs preinstalled FreeNAS. We just had to configure the network connection on the command line. Everything else, like adding users or backup jobs, can be done through the user-friendly Web-Interface. FreeNAS is based on FreeBSD. BSDs are meant to be stable. [Tim Gruene tim.gruene@univie.ac.at](mailto:Tim.Gruene@univie.ac.at)

Our 60-drive 760TB FreeNAS has been running flawless for more than 1 year. One failed drive was easily hot-swapped. If you want capacity, this does not cost more per TB than inferior hardware, yet you get good redundancy and performance. See my old tweet: [https://urldefense.com/v3/https://twitter.com/hicryoem/status/1223966559976083458?s=21__!!Mih3wA!TpItSV7E6sE1qSOJgCwUfuLYuNhW92PMSZUaIF70zdj2z4BE_zg2izktUsKL5clRQ\\$](https://urldefense.com/v3/https://twitter.com/hicryoem/status/1223966559976083458?s=21__!!Mih3wA!TpItSV7E6sE1qSOJgCwUfuLYuNhW92PMSZUaIF70zdj2z4BE_zg2izktUsKL5clRQ$) [Matthias Wolf matthias.wolf@oist.jp](mailto:Matthias.Wolf@oist.jp)

Our lab is also currently making similar considerations for a moderately large server. So, this email chain has been quite serendipitous. I am wondering whether those with external servers (45drives, trueNAS, Synology, etc.) might comment on their network setup? Did you opt for a special network switch? Has this been necessary, or can these servers directly host the workstations? What kind of local network infrastructure do you use? Connections, SFP+/RJ45/some fancy optical thing? I note Steve has said, in the small 12-drive SynologyNAS that a 10Gbe (RJ45?) card has been sufficient for ~950/600 MB/s IO. I recall the Synology can be

networked directly to workstations without the need for a switch if I am not mistaken. So, what experience with the other larger servers? Perhaps you might also comment on your typical use case? For example, is your server for long term storage of processed or raw data? Or is it a scratch space for analysis with read & write?
Charles Bayly-Jones charles.bayly-jones@monash.edu

It is quite possible to purchase a decent 10 Gbe copper switch nowadays for relatively little money. I am using a Netgear XS512EM in the lab. It can do nonblocking 10G over copper on all 12 ports (it claims). I have not pushed those limits, as I only have a handful of machines with 10G connections. While some institutions are beginning to upgrade internally to 10G, the ones who adopted 1G over a decade ago often used Cat5 cabling internally, which is not capable of reliable 10G connectivity, which means those institutions are faced with the possibility of having to rewire everything they want to provide 10G support. Anyway, my stopgap solution was to just set up a deadnet in the lab for NAS and inter-machine communications. Really no different than a small cluster. I will add that I was not trying to sell Synology on labs with solid IT expertise needing to set up a petabyte of storage. I was saying that for smaller labs Synology presents a very friendly solution, and I dispute the statement that it is “lower quality” in some way. I have been running standard Supermicro storage rackmount units for ~15 years (a common platform for trueNAS; enterprise grade). They are excellent, of course, and I highly recommend them, but over a 5+ year period they are NOT worry-free. They develop hardware issues, such as failed power supplies, failed RAID cards, etc. Further, if you fully populate them at time of purchase that also means that the drives will start approaching end of life all at about the same time (anywhere from 3-8 years depending on how lucky you are with a particular batch of drives), leading to an extremely high risk of data loss even with RAID 6. These boxes are not for labs who say, “we don’t have much IT experience/support”, even if you run something like freeNAS. They will likely be great for 3-4 years before you start running into issues. Before the Synology boxes, the standard solution in the lab was to buy Supermicro workstations with 8 hot-swap drive bays and an internal hardware RAID card. This would give about 800-900 MB/s of bandwidth as a RAID 5 and have a lot of space. However, this storage is all local to the machine, and moving large data over a 1 GB network at ~100 MB/s is so painful for some things. The Synology boxes can sit under the desk and will email you when there is a drive failure or any other issue and can still provide roughly the same performance as the previous CPU-tied storage (with the 10 Gbe card). I will say that one of my biggest issues over the last couple of years has been having RAID cards fail in the Supermicro units, and having to go to eBay, etc., to find equivalent replacements to avoid having to wipe 500TB of data and copy from backup (the newer RAID cards are often not quite compatible enough with the older cards). Even with the Synology boxes, expect that you will likely have to “refresh” the technology in 5-8 years. **Steven Ludtke** sludtke@bcm.edu

We have purchased a 60 drive SuperMicro chassis which I like. Good redundancy and build quality with remote management and KVM included. They are configured as a 6×10 RAIDZ2 ZFS pool, which provides 480TB of usable disk space and decent performance per server. A very fast Intel optane SSD boosts synchronous write performance, +768GB of RAM and 36 CPU cores, for around \$40K each (prices ~2-3 years old). It can hit ~4GB/sec in linear read/writes and do about 15-20K IOPS (not great, but not bad). I got it with four 10GbE ethernet ports but have only used two so far aggregated to the switch. In the last three years we have not had to replace any drives.

For off-site cloud backup, I think Wasabi seems reasonable at \$70/TB/yr. This requires less IT support on your part, and likely better redundancy/availability. I am currently looking into Quobyte and WekaIO to provide an all-encompassing data collection and data processing HPC storage solution, but it is expensive. The dream is to not have to compromise between storage capacity, throughput and high IOPS. Both systems serve primarily from a cluster of all SSD servers but can transparently tier data to a large object store. For large archival storage, I like the combination of performance, space, and cost of using an on-site object store with disks. Tape is still cheaper at scale but seems to have steep buy-in costs and is less convenient. **Craig Yoshioka** yoshiokc@ohsu.edu

We went the opposite direction. Initially we used all software RAID but discovered that at least back then (8-9 years ago), it was easy for something trivial to screw up the RAID. One bad shutdown would put the array out of sync and require hours of painful manual debugging to properly rescue, and swapping out a failed drive was a lot more work than it should be. Eventually we stopped using them. The hardware RAIDs on the other hand are robust. Even when the controller fails, I think the only time I have suffered an actual data loss was when a 3rd drive failed during recovery of 2 other failures. Even then, though, you certainly wind up in occasional cases where you must boot into firmware to debug a problem. IMHO neither solution is a good idea for non-Linux gurus. I had high hopes for ZFS (which Synology generally uses internally), but it still has not reached the sort of stability everyone hoped for. **Steven Ludtke** sludtke@bcm.edu

I also have been using Synology for more than 5 years and for local storage it has been working great. I have 4 units: one 8-bay connected directly to a workstation via 10G SFP+ card, a 12-bay with a 12-bay expansion unit connected to the network via a Synology 10G card, and another 8-bay at home. Its user interface is quite friendly, and Synology frequently releases OS and package updates. **Gökhan Tolun** tolung.bio@gmail.com

My FreeNAS box has a quad port 10G SFP+ baseboard-mounted NIC. I have configured two of its ports for port aggregation with LAGP which is supported by FreeNAS. Two optical fibers are connected to a floor switch managed by IT, which also supports LAGP. The uplink from the floor switch to a core switch in the datacenter has much higher bandwidth (100Gb/s, I think) and this is an all-optical network. My workstations in the lab have 10Gb copper ethernet NICs (Solarflare or Intel) and they are connected to the floor switch with SFP-UTP transceivers on the switch side, because the cables in the building are all ethernet. Fortunately, these are all CAT6a or CAT7, since the buildings are all less than 10 years old. The highest transfer rates between a cluster front end and our FreeNAS box by rsync I have seen was about 1GB/s (~10 Gb/s) disk-to-disk. I have not measured the network speed yet over this route with a tool like iperf. Between this storage server and our lab workstations it is less (200-300Mb/s or so), but that is probably limited by the local RAID5 array on the workstations, or because no network tuning of the workstations that run CentOS8. We have a rack enclosure in a side room to the Titan Krios that hosts the K2 summit processor, FreeNAS box, Warp PC and two FEI Falcon storage servers. All of these have 10G interfaces and are connected to a local optical switch managed by IT (also in that rack). The FreeNAS also has a BMC. Our HPC people did not want the Tyan server in their datacenter, because they think Tyan/Supermicro, etc. are inferior hardware without a service contract. They said that there

were many reports of Supermicro PSUs catching fire and they could not host such hardware in the datacenter. I understand this and I agree, but I bought the Tyan server because it was much cheaper than professional datacenter-grade servers.

Before the FreeNAS, I had bought a DDN storage server (55 drives, which was later expanded with an additional 80-drive expansion chassis), which is administered by IT and hosted directly in the OIST datacenter. Because the DDN has smaller drives, is a couple years older, and has ridiculous redundancy (wasting more than 30% of the raw capacity), the total usable capacity is pretty much the same as the FreeNAS. The annual service contract of the DDN alone is 1/3 of the total cost of the Tyan FreeNAS box. For file access we use NFS and CIFS. But I still have not found enough time to figure out NIS on CentOS8. This was no problem on CentOS7, but for some reason CentOS8 is much harder in this respect. I have one Linux box for user authentication in the lab using LDAP, because I want control over my local segment, and IT demanded giving up root access when authenticating through their central AD domain controller. FreeNAS has all the major network services preconfigured and it is a simple matter of enabling the service. I like it quite a bit. I highly recommend working with your IT team to have them lock down your cryoEM network on the router level. Ours has been a dedicated network segment isolated from the internet and the rest of the university, only allowing routing of the Teamviewer and RAPID ports. **Matthias Wolf** matthias.wolf@oist.jp

There is some great info in this thread! We have gone a bit of a different route, so I thought I would describe that briefly here in case it might be useful. We have implemented a tiered storage strategy:

- Fast scratch for processing: 11–28TB SSD raid zero scratch spaces on several workstations, all cross-mounted via NFSv4. Each user has their own space to use for processing, and all cpu/gpu on this local network can access each scratch. These are truly non-redundant, and users need to offload important results to another location. Note that we had a handful of SAS SSDs which are much more stable for write operations, but unless you are using these for a write cache to a slower disk system (we are not) these drives are IMO a waste of money. For ~1/4 the price, we added several SATA SSDs (for example, Samsung 8xx pro) that in raid 0 are just as fast, and have a similar read life-time. In our case, the cost savings means even if they die in three years (warranted for five) it is still much cheaper than SAS.
- Intermediate (lifetime ~ 1 year) storage: Originally, we planned on a Synology station as several others have mentioned. I recommend this for smaller storage (1–200TB) if your IT permits it. Being in the medical school here at Umass, InfoSec made this a non-starter (That is a different story). We investigated two proprietary options, Netapp and Dell EMC Isilon. The Isilon ended up being roughly 1/2 the cost compared to the NetApp. Notes on the Isilon:
- While more expensive than an off-the-shelf NAS option (like Synology), the failure prevention and snapshots seem more robust than any RAID option. There is a great deal of flexibility in what is backed up (how frequently and how long), and the excess storage used is lower. We have ~480TB raw and ~ 360TB really available based on our snapshot policy and failure policy. That is, there are four nodes, and a full node can be lost w/o data loss.
- Because there is no controller, the system is easily extended to include more nodes, which will increase the raw:usable ratio as a full node is less of the total available. I recall the max efficiency

peaked somewhere around 90 for a large system. You can also combine SSD (F-series) and have a tiered system within your Isilon, though we opted to manage this externally.

- I have not yet benchmarked the performance as it is fast enough for our intermediate (active archive) use.
- Cold storage: we are currently using the good old external HDDs for this. It is not the safest but is certainly the cheapest.
- Network: We had to run 10Gb (cat6A) cables to the wetlab/drylabs to get 10Gb connectivity back to the IT closets, where we also had to add switches to support the speed. The cable run was somewhere around \$7k and each switch was ~\$4k. I am sure there are less expensive options, but these are what IT was willing to support, which is, in the long run, making things affordable from my perspective. We have dual port NICs but did not set up bonding as it seemed like we could just saturate the copper when reading from SSD scratch and dumping into another /dev/null. **Benjamin Himes** himes.benjamin@gmail.com

Connect to LAN or not?

Microscopy Listserver

Obviously, anyone doing microscopy work needs some manner of moving acquired data and images to other computers for analysis, reporting, and such. In the old days, that meant using “Sneaker Net” (the process of walking a floppy disk down the hall). Thankfully, we are long past that with thumb drives storing terabytes, cloud storage and LANs. The pandemic has forced even the most ardent to adopt web meetings, whose numbers have exploded for doing things like microscope demonstrations and remote training. Growing further towards remote support for diagnosis of possible problems or tweaking settings to improve a customer’s use of their microscope or EDS. At the same time, corporations and the government entities have been implementing stricter “traffic cops”. We have recently even seen USB drivers getting blocked. Then comes TCP/IP traffic and IT roadblock police restricting Administrator rights to a local PC that make it feel like George Orwell is running things. Everyone pointing fingers at the source of the problem. Chaos and frustration ensue. Software and hardware for microscopes, or for this matter, any lab instrument is rarely justified to be put through some costly Microsoft certification process to be on some approved list for easy “TSA-like” clearance. 1) Microscope Users need LAN access to move files to their office PC. 2) Internet access is need for microscope User/Supplier support. QUESTION: What do you find is the best solution to achieve these needs when there is suffocating IT overhead that the system is unable to tolerate? Should labs implement wireless access use for temporary access to the web for remote service?

Mike Toalson miketoalson@gmail.com

These are good questions worthy of some discussion. I think, however, an important preface to make is that there will be no one-size-fits-all solution for everyone. Local IT policies vary widely and punishment for breaches of IT policy can vary from a slap on the wrist to being fired and/or having to answer questions from the authorities, so please be very careful about what you do on your network! As a microanalytical consultant, I work with a wide range of users located in labs all around the world. When asked similar questions in the past, my initial advice was to always try and work with your local IT services to achieve the results that you require. Fundamentally, they are providing you with a service and they should be working to help you do your job and provide your users access to their data. The disconnect arises when the IT services are either inexperienced, or they are following a mandate that has been over-optimized for office computing to the detriment of laboratory requirements.

There are two main questions you should have for your IT services: 1) What storage do you provide and how can users access it? The idea here is that IT services will deal with the issues surrounding access to the server, and all you must do is push data from your microscope PC to the local storage server. These types of services should be perfectly adequate for data access, but of course this does not help with remote control of the instrument. 2) What remote-access VPN service do you provide? The idea here is that network services should make it possible for users to access the local network via their VPN. IT services will handle all of the issues and hassle supporting users to get into the network, but once they are in, they can then access services running on your microscope PC directly. This is a nice option because you can use whatever services suit you best without requiring any additional input from IT.

While this approach works very well, it falls over when you have users who are not part of the organization. Some IT departments will be able and willing to issue those users with credentials, while many will not. What can you do when local IT services are unable or unwilling to provide the aforementioned services? As you suggest, you can do an end run around the local network and use a GSM dongle to access data directly. I have used this approach on several occasions, and we have found that it always works extremely well. In one case the access was faster than when we had tried to use the LAN. Rather than leave it on all of the time, my client would set up the dongle in advance of the service window and then remove it afterwards. Permanent use of a GSM dongle may not make much sense if there are data caps on the service, and from a security perspective you may want to think carefully before trying it. For this reason, it will most likely be forbidden by your local IT policy.

Another approach is to use your own VPN service. There are a number of these that are free to use, but there can be a lot of confusion surrounding them since they are often used by “gamers” and by themselves the VPN software does absolutely nothing. You still need to run a service (such as file sharing, FTP server, VNC remote access, etc.) on the microscope PC, while the VPN software simply provides access. You also lose the ability to use DNS names to identify your computers, but this is not much of a drawback, users can just as easily enter a UP like “10.1.2.3” as they can “mysem.company.org”. My personal preference is ZeroTierOne (<https://www.zerotier.com/download/>), but I have also had good success with LogMeIn Hamachi (<https://vpn.net/>) in the past. The benefit of using a VPN like this is that you can often keep using the same services you are already using. The downside is that you now have a VPN to administer, and membership of that VPN will continue to grow over time as you pick up users. To manage this, I would recommend regenerating the network every 6–12 months and then advertising the new network ID to the current pool of users. In terms of hardening access, it is always a good idea to use a second PC as the point of contact for user access to data, and then push your working data to that PC from the microscope PC. This limits the services and software running on the microscope PC. In this case the second PC can be anything that is available, provided it has enough storage. The CPU and RAM requirements are minimal. It can even run an OS different than that on the microscope PC, which is useful in terms of securing a system that will be accessed by multiple users. The challenge is that a second PC is primarily only of use to serve up data. If you require remote access, then while it can be used as an entry point to the network, where users access it via VPN, say, and then initiate a remote desktop connection to the microscope PC, such a setup is quite complicated and not recommended. For pushing data to the server, you have a range of options. FreeFileSync (<https://freefilesync.org/>) is a good one due to the friendly interface and being cross platform, but ultimately the command line program “rsync” (and the many programs based upon it) is fundamentally the best way to sync large data sets across a network. For Windows PCs you will get good

mileage out of the command line program “RoboCopy”. The premise of all these programs is that they only copy updated files to the server, avoiding the need to copy the entire dataset each time. How frequently you push data to the server is up to you; for most practical purposes daily is usually sufficient. If anyone is looking for specific advice on these sorts of setups, please feel free to contact me directly and I will do what I can to help. **Ashley Norris** ashley@norris.org.au

I can relate to the experience, especially for remote troubleshooting. A data server-based system will work well and still be quite secure. Let me explain a simple setup. The main system connected to the instrument has one LAN card to connect to a data server. The data server, which has two LAN cards, has one to connect to the main system and another to connect to the internet. The data server has a shared folder to which the main system can read or write. The data server, however, cannot access the main system. In both systems, all unused ports are blocked, and a basic antivirus and firewall are enabled. The users with appropriate access can SSH to the data server and transfer data to their private computer for analysis.

This way the main system connected to the instrument does not connect to the internet, LAN, or USB lowering the chance of virus or Windows update-based issues. The system however needs manual system updates (both Windows and antivirus). This setup allows aged instruments running older versions of Windows to still be operational and off-network. Generally, IT admins do not like anyone having admin access even temporarily so software- or hardware-based remote access in the presence of IT admin may be okay if it is for rare occasions and short duration. **Rooban Venkatesh K.G.** Thirumalai.rthirumalai@i2at.msstate.edu

Use of Uninterrupted Power Supplies to Protect Confocal Systems

Confocal Listserver

We are unfortunately facing many unplanned (random) electrical outages/voltage variations in our Institute (it is a very old building) and our systems (and particularly lasers) don't like it. I was wondering if anyone is using UPS units to protect confocal/2-P systems? If so, what type of UPS are you using? **Mario Methot** mario.methot@neuro.ulaval.ca

When we designed our new lab, we investigated which equipment is subject to damage. Our multiphoton laser (Coherent chameleon Ultra II) has a built-in safeguard with an internal battery sufficient to cool down the laser. The software reports the state of the battery. This is not written in the laser specifications, so it is worth asking the company if they have such a system or not. The lifetime of argon lasers will be reduced if repeatedly turned off without cooling, but there is no impact if it seldom happens. This is a problem if it happens often as seems to be the case in your institute. The rest of the equipment (cameras, computers, etc.) can be turned off suddenly without damage. However, the main problem is the spike that occurs when one turns the power back on. This is very damaging for a lot of equipment including the computers. It is very important to have a good spike protection system. Alternatively, unplug all equipment before the current returns. **Sylvie Le Guyader** sylvie.le.guyader@ki.se

We have all microscopes and processing computers on UPS for power conditioning and emergency backup. For 120V systems: GE VH 2KVA Tower UPS <https://www.geupssystems.com/ge-vh-series-ups> ~\$1,000 USD. For 209V systems: GE LP-11U UPS <https://www.geupssystems.com/lp-11u-series-ups> ~\$6,000 USD **Doug Richardson** ds.richardson@gmail.com

We have 2 units of 3000 VA (<https://www.riello-ups.es/products/1-sai/44-sentinel-pro>) for all components of our Leica SP2 AOBS, and only one unit for our Nikon AIR+ to protect the argon laser and PC. All conventional microscope PCs have the WinOFF program installed to switch them off after 30 minutes if not being used. **Konstantin Neashnikov** [microscopia-ibis@us.es](mailto:konstantin.neashnikov@ibis.us.es)

We use the following UPS to back up our 2-P workstations: UPS: <https://www.cyberpowersystems.com/product/ups/smart-app-sinewave/pr2200lcdrtxl2u/> Battery extension: <https://www.cyberpowersystems.com/product/ups/extended-battery-modules/bp48v75art2u/> With this setup we can run the workstation for about 15 minutes with the laser on (about 4 minutes without the battery extension), to give enough time to safely save an experiment and shut everything down. Once the laser is on standby, the UPS can run everything else for hours. As others have said, the two major considerations are that you have enough VA (2500 per high current laser system should be sufficient), and that the output is a true sine wave, not an approximated one (approximated sine waves can put unnecessary stress on power supplies for sensitive equipment). **Benjamin Smith** benjamin.smith@berkeley.edu

We have a UPS for every system in our facility. They are primarily 3kVA systems mainly from Socomec but we also have a MGE. They are all online UPS systems, so they smooth as well as protect from spikes/outages. Every microscope user is trained to listen for the beeping and shut the system down if this is ongoing. Also, we have them all plugged into essential power sockets. These sockets are provided with power by generators in the event of a complete campus power failure but do not come on if it is just local, for example, one building/floor. We started adding these to every new system after having an expensive repair bill (\$10k) after our Leica confocal went off and on again. They are included as a necessary expense. We also have a contract with a local supplier to check them and replace batteries regularly. They are also good for the time-lapse systems because a short outage or spike would otherwise terminate the experiment immediately. **Jacqui Ross** jacqui.ross@auckland.ac.nz

As a related aside, most UPS systems use lead-acid batteries for storage. These typically have a lifespan of 4-5 years in a typical UPS. Make sure to replace them after about four years in service. Better quality UPS systems will include an alert to indicate when the batteries are aging out. Batteries also do not fare well in storage, so it is pointless to buy spares ahead of time. **Craig Brideau** craig.brideau@gmail.com

Nice discussion. So, if I summarize what is needed: 1) good UPS with batteries regularly checked and replaced when needed; 2) backup power from a diesel generator that takes a short while to take over (basically it is not enough to have the backup power without UPS); 3) someone on call. This is exactly the conclusion we came to when designing our new building. But all this is necessary only if there are power cuts several times per year. If the power cuts come seldom and there is a good surge protection system, there is no need for these expenses. This is of course to be balanced with the cost (and number) of the wasted experiments. **Sylvie Le Guyader** sylvie.le.guyader@ki.se

Widefield Autofluorescence in Unlabeled Cells – A Filter Mystery

Confocal Listserver

A colleague asked me to look at their widefield microscope. It is an inverted microscope with a 100W Hg source, excitation filter wheel, a

couple of choices for dichroics in the microscope filter changer, and a filter wheel in front of the camera. They are seeing unlabeled cells fluoresce green (FITC/GFP set) with an otherwise black background where there are no cells. The microscope is approximately 15 years old. My guess is that the excitation filters have failed (or are failing) after being on the receiving end of a 100W Hg lamp for all this time. Any other thoughts? **Doug Cromey** [dcromey@arizona.edu](mailto:doug.cromey@arizona.edu)

Unlabeled cells do fluoresce in the yellow-green range. I have always thought it is because they have soluble molecules such as flavonoids in the cytoplasm. If the cells are permeabilized, for example, to use them in immunofluorescence procedures, there is much less of this background. **Carol Heckman** heckman@bgsu.edu

When I first learned immunofluorescence, we added a small amount of glutaraldehyde to the PFA, and this caused autofluorescence. Therefore, we bleached the cells (I think after saponin extraction) with NaBH₄ (which is not compatible with Bodipy). After 4 washes, we would check the cells on a fluorescence microscope to assure that we had bleached them sufficiently. I believe other labs similarly use NaIO₄ and/or glycine. This discussion is particularly apropos to us now because a lab asked for my help with what I first thought was a simple autofluorescence problem but appears to be something more difficult. It also harkens to a discussion I saw elsewhere about imaging DAPI after imaging green because the DAPI may photoconvert. I was going to post a more detailed description of the problem, but the timing is right to post about it now. Rather than tell the story how we got to this point and discovered the problem, I will jump directly to the problem. Cell cultures that are fixed with fresh PFA in PBS and Triton extracted have a very faint background when excited with 488nm light and similarly faint background with secondary antibody or isotype labeling. The background is weak enough to not need further bleaching before the specific labeling. (So far, no problem.) The problem is that exposure to UV light (365 ex DAPI filter with Hg lamp, 370nm DAPI filter with X-Cite LED, or 395nm or 405nm LED or laser) causes the cells to emit brightly when subsequently excited at 488nm. Practically this means that any imaging of DAPI or Hoechst prevents subsequent imaging of green. This is dose-dependent, so more exposure to UV means brighter green (no, we have not plotted a proper curve or looked for saturation, although we have a few data points). This is not due to DAPI photoconverting; we ran controls of completely unlabeled cells mounted both in ProLong Diamond and glycerol (where the response appears to be stronger). They have seen this problem with a few cell types, so it is not something like a line mistakenly expressing a photoconvertible protein. Of course, there are ways around this. Stop using DAPI as a convenient way to find cells. Always take a picture of green first with the last exposure being DAPI. Switch nucleic acid labels to another color. But while these techniques work, when limiting the staining to three colors, this problem eliminates blue as a possible color. By widefield fluorescence, tiling is ruled out because the exposed circle is larger than the rectangular camera FOV. Have other people seen this problem? Any ideas? Thank you!! **Michael Cammer** michael.cammer@med.nyu.edu

The article Michael mentions about UV conversion of DAPI to a green emitting forms is “UV-activated conversion of Hoechst 33258, DAPI, and Vybrant DyeCycle fluorescent dyes into blue-excited, green-emitting protonated forms” by Żurek-Biesiada et al., <https://onlinelibrary.wiley.com/doi/full/10.1002/cyto.a.22260>. There is also a tech note on the Leica website that deals with this topic:

<https://www.leica-microsystems.com/science-lab/learn-how-to-remove-autofluorescence-from-your-confocal-images/>. Fixatives, besides the mounting media, imaging dishes, culture media, etc., also cause autofluorescence. Spectral imaging and unmixing, if available, are an option. [Sathya Srinivasan scitecheditor@gmail.com](mailto:Sathya.Srinivasan.scitecheditor@gmail.com)

Fixing Plant Fluorescence

Confocal Listserver

I have a user who would like to use tetrazolium to label plant mitochondria in both slow- and fast-growing plants. He says that it is fluorescent. Since I know very little about preparing plant tissues for fluorescence microscopy, can they be fixed with aldehydes (formaldehyde or glutaraldehyde)? Our hope is to image them using our Zeiss Elyra (structured illumination). Any "red flags"?

Doug Cromey dcromey@arizona.edu

Plant tissues can be fixed by vacuum or syringe infiltration with 2% glutaraldehyde and 2% paraformaldehyde. Glutaraldehyde can autofluoresce, but depending on the tissue and time before imaging, paraformaldehyde may be sufficient. If they are fixing these for an extended period, then an additional step of 0.2M glycine for a few hours will quench the autofluorescence. What are they saying the spectra of the tetrazolium is? I am only aware of nitroblue tetrazolium as a colorimetric stain for reactive oxygen species labeling. Plants are full of things that autofluoresce, so depending on how sharp and isolated the emission peak is, you may want to do a spectral image for confirmation. Live cell imaging of plant mitochondria can be performed if you infiltrate the leaves with Mitotracker. However, I do not know the stability of this stain post-fixation in plant tissue.

Timothy Chaya tchaya@udel.edu

You might get away with imaging fresh hand-sections of the tissue without fixation. It would be worth imaging the fresh sections and making sure the fluorescence is still there after transferring a section through 2-4% formaldehyde in 25-50mM phosphate buffer, pH 6.5-7.0. **Rosemary White rosemary.white@anu.edu.au**

This is classic stuff. Literature I am aware of uses tetrazolium on isolated mitochondria. This is a functional assay (the tetrazolium gets reduced). In the standard sense, it is unlikely to work in a fixed sample. Or even an intact plant organ. Having said that there could be ways to use tetrazolium as a simple stain for mitochondria. Yes, plant tissues are certainly well fixed with aldehydes. But both formaldehyde and glutaraldehyde can generate autofluorescent compounds. The severity of the background depends on the organ, tissue, and condition of the plants. Many plant organs are also autofluorescent even without fixation: Chlorophyll in photosynthetic tissue and many compounds in cell walls and vacuoles elsewhere. Clearly some of this can be handled with spectral selection. As a rule, roots tend to be easier to work with than leaves. **Tobias Baskin baskin@bio.umass.edu**

This review article might be useful in the plant microscopy community. Autofluorescence in Plants: <https://www.mdpi.com/1420-3049/25/10/2393/htm>. **Timothy Chaya tchaya@udel.edu**

Not sure which Elyra you have, but the red flag on the Elyra PS.1 is that most filter sets have a long pass 750. If you are working with any green tissue, the chlorophyll autofluorescence will blow out any other fluorescence. The solution is to put a SP740 filter in the slider underneath the filter turret. **Jeff Caplan jcaplan@udel.edu**

MT

REGISTER NOW

M&M 2021

PLENARY SPEAKERS

COVID-19 VACCINE DEVELOPERS

Kizzmekia S. Corbett, PhD

Coronavirus Vaccines & Immunopathogenesis Team
Vaccine Research Center
National Institute of Allergy and Infectious Diseases, National Institutes of Health

Jason McLellan, PhD

Department of Molecular Biosciences
Department of Chemistry
College of Natural Sciences
University of Texas at Austin

2020 KAVLI AWARDEE

Ondrej Krivanek, PhD

President, Nion Co.
Affiliate Professor at Arizona State University








www.microscopy.org/MandM/2021

Crossword Puzzle Answers

See puzzle on page 56.

1	P	2	L	3	E	4	N	5	A	6	R	7	8	9	10	11	12	13	S		
14	O	O	O		15	W	A	R	E	S		16	C	A	G	E	Y				
17	S	U	N	18	D	A	Y					19	I	T		20	G	Y	M		
21	T	V		22	E	R	S	23	A	T	Z		24		25	26	I	R	E	P	
27	E	R	28	R	E	D		29	C	E	O		30	V	I	D	E	O			
32	R	E	A	D			33	M		34	D	O	R	I	C		35	M	S		
					I		36	A	A	H			M		38	T	E	M	P	I	
40	T	41	E	L	E	G	R	A	43	M		44	E	Y		45	E	T	A		
46	U	M			47	M	A	G	N	I	F	Y			48		49	N	Y		
					50	D	E	I			51	T	R	E	A	D				54	O
55	O	56	V	E	R	N	I	G	57	H	T			58	F	I	F	T	H		
61	R	E	S	I						E				62	F	L	O	E			
63	I	N	I	T	I	64	A	L	65	S		66		67	F	O	U	R	T	H	68
69	A	U	G	U	S	T					70	A	P	E	R	T	U	R	E		
72	L	E	N	S	O	N					73	P	A	N	D	E	M	I	C		