Viewpoint Documentation

This is a work in progress. Each time you download from this link, you will get the latest version of the document.

Overview

1. Getting started

a. Download from here:

i. MSL
   1. OSX: https://updateserver.planetarysciencecommand.com/release/msl/osx/installer/download.html
   2. Windows: https://updateserver.planetarysciencecommand.com/release/msl/windows/installer/download.html

ii. MER
   1. OSX: https://updateserver.planetarysciencecommand.com/release/merb/osx/installer/download.html
   2. Windows: https://updateserver.planetarysciencecommand.com/release/merb/windows/installer/download.html

iii. User: viewpoint

iv. Pass: takeout pizza for breakfast
   1. (spaces included)

b. Installation on OSX

i. Download the Viewpoint.zip file and unzip it in your Downloads directory.

ii. Move the unzipped Viewpoint_MSL app to your Applications directory. (Note: starting it in your Downloads directory can cause it to crash on some systems due to access errors.)

iii. The first time you start it, you'll need to Ctrl-click and select open, and agree to starting an application by an unidentified developer.

iv. You probably want to move the Viewpoint_MSL.app to your /Applications directory like a normal application for easier access. But that's not necessary.

v. If there are multiple copies in the same folder, OSX will add a number to the name, e.g., "Viewpoint_MSL 2.app". This can cause problems for the updater and make it so you can't start the application. So make sure the app name is just "Viewpoint_MSL.app".

c. Installation on Windows

i. As someone with Administrator privileges on your Windows computer, run the installer.

ii. It will install into C:/ProgramData rather than Program Files. This is to allow Viewpoint to self patch itself, similar to how Google Chrome does. (Google Chrome also uses C:/ProgramData). Windows doesn't allow anything in Program Files to change itself.

1. Windows native virus scanning doesn't mind this, but third party virus software packages may complain about this. If so, you can either add an exception in your virus software for Viewpoint, so that it stops complaining. Or you can install Viewpoint somewhere else, such as your Desktop. If you want to do this, you can get a zip file to unzip in the directory of your choice here: https://updateserver.planetarysciencecommand.com/release/msl/windows/installer/Viewpoint.zip

iii. As part of the installation project, shortcuts to Viewpoint should be created in your Start Menu under the name Viewpoint.

2. Login

a. You can change your Viewpoint password at the login window.

i. You will get a username and initial password with your Viewpoint account. (This is not the same as the user/password to the website to download Viewpoint).

ii. You can change your password to something else when you start up Viewpoint at the initial login screen. Click the "Change Password" link next to the username text field.

iii. If you forget your password, email Jon to reset it. jon.proton@planetarysciencecommand.com

b. Choosing the Planning Sol

i. When you log into a planning sol, it's like logging into an IRC channel in that users can only see each other's observations, chat, and interact if logged into the same planning sol.

ii. Is the maximum sol that data (imagery, rover positions, targets, etc) are queried.

iii. Image observations created in a planning sol are only viewable/editable in that planning sol. However, you may copy an observation from one planning sol to another.

c. Troubleshooting login problems.

i. Make sure you are using a VPN split tunnel to JPL, or no VPN at all. Using VPN full tunnel will make it so you can't connect to the Viewpoint server.

ii. You might be behind an extra restrictive firewall that doesn't allow the necessary port connections to the server. See the download website page for more info about which port needs to be opened by your SA to allow Viewpoint to connect to its server.

iii. A slow internet connection can cause a timeout. If the server happens to be under high load, it can also cause a timeout. Retry a couple times and it should work.

iv. There could be a typo in your username or password. You should get a message saying if that's the case.

3. User interface general rules of thumb

a. Hover over UI components to show tooltips for an explanation of what they do.

b. The "?" buttons have some help text related to the items they're positioned next to. Read through the help text for more in depth information than tooltips give.

c. A lot of functionality is available through context menus (aka pop-up menus or right-click menus). So right click (or ctrl-click on single button Macs) on everything in the scene view, as well as Observation and Pointing Draft tabs to show the context menus with possible commands.

4. Quick Reference -- Keyboard Shortcuts

a. General

i. **h** -- hide/show cursor

ii. **m** -- maximize/minimize the scene view

iii. **t** -- create a target

iv. **d** -- enable/disable the measurement tool (d for distance)

v. **delete** -- deletes selected pointings or targets

b. In Camera View

i. **a** -- select all pointings

ii. **c** -- copy selected pointings
iii. **v** -- paste selected pointings
iv. **z** -- undo one step
v. **y** -- redo one step

c. In Rover View

i. **r** -- zoom the view to the rover
ii. **g** -- zoom the view to the cursor (g for ground)

5. View Modes

a. General Info

i. There are a few different view perspectives available. When Viewpoint starts up, you are put in Camera View initially, but you can switch the view using the four view toggles near the upper right corner. The four views are Traverse View, Rover View, Camera View, and Subframe Edit View.
ii. For the View Modes: Rover View, Camera View, and Subframe Edit View, there are additional options and features available by using the toggles on along the right edge. Especially the top four, Scene Options, Rover View Options, Camera View Options, and Subframe Edit View Options. These will be covered in more detail in a subsequent section.
iii. Navigation: For all view modes there are navigation controls near the upper right of the view. You can use these controls to navigate the view. But there are also keyboard shortcuts as well, which I find to be more convenient. If you hover over the navigation control, a tooltip will list the keyboard controls for the current view. Also, clicking the "?" button next to the View Mode toggles will display some help text with information on View Modes and navigation keyboard shortcuts.
iv. Since rendering the view takes a lot of system resources, by default, if you don't move your mouse for awhile, or Viewpoint loses window focus, rendering will be disabled till you wake it up again by selecting Viewpoint or moving the mouse. You can turn this behavior off and/or adjust the timeout length in Settings (gear button on the lower right).

b. Traverse View

i. This is an overhead view of the rover on the HiRise mesh. It displays marker pins at the recent end-of-drive positions with a line drawn between them to show the rover's approximate traverse.
ii. If you click on a marker pin, there are options for moving the rover to that position and/or also loading imagery from that position. This is a quick way to load up data from previous rover positions. (These marker pins are

also available in the other views if you enable “Show Traverse Path” in Rover View Options or Scene Options shortcuts.)

iii. The number of positions shown are controlled by the maximum number of sols of data that are queried. You can change that value in Settings (the gear button on the lower right).

iv. The numbers at the edge of the view are az,el coordinates. By default, they are in LL, but may be changed to RNAV, JOINT, or turned off by changing the coordinate frame of the az-el grid in Camera View Options (see below).

v. If you have made any pointings, they will show in this view. This can be helpful for planning post-drive imaging, since you can see what features will be covered in the HiRise by your mosaic wedge.

c. Rover View

i. This is the view of the scene around the rover using a floating third-person perspective. You can look at the rover and its surroundings from any point in space around the rover.

ii. This view is useful shadow modeling, viewing the terrain in the rover’s work volume, target selection, and viewing playback of rover simulations.

d. Camera View

i. View the scene from the perspective of one of the cameras. (So far just cameras mounted on the mast.)

ii. This is the main view used for planning imaging observations.

iii. You can change which camera you are using for your perspective in Camera View Options.

iv. The numbers along the edge are az/el coordinates. By default these are in LL, but you can change them to RNAV, JOINT, or turn them off under Camera View Options. You can also turn on az/el gridlines in Camera View Options to help see the coordinate frame lines in the scene. This is can be especially helpful when planning mosaics when the rover is on a steep slope.

e. Subframe Edit View

i. To enable this view, you must have a pointing frame to edit.

ii. This view zooms in on a single image pair and allows you to adjust the subframe edge positions by dragging them where you want them.

iii. If you drag outside the frame boundary, you will move the position of the pointing in the scene. (This type of positioning is intended for fine tuning placement of subframes, so overlap snapping is disabled.)

iv. The az/el numbers and grid lines are the same as for Camera View.

6. Observation Bar (the topmost bar)

a. Observation Toggles

i. The topmost bar contains observation toggles. An observation has its own selection of imagery, rover position, and may also contain a collection of pointings drafts. Click the toggles to change which observation is currently being viewed and edited.

ii. **Auto** and **Manual** -- Always present at the left are two special types of observation toggles called Auto and Manual.

   1. **Auto** automatically loads the most recent imagery, puts the rover at the most recent position, and updates everything as new data becomes available.
   2. **Manual** allows you to hand-select which imagery and rover position you want to see. You can select the rover position and imagery using the rover position and imagery dialogs, which are the toggles on the far left corner. Alternatively, you can set the rover position and imagery using the rover position marker pins (viewable in Traverse View and by turning on "Show Traverse Path" in Rover Options and Scene Options shortcuts).
   3. The two types

iii. Pointing Observations -- **Auto** and **Manual** are strictly for scene browsing, and don't allow you to create image pointings. But you can create other observations for modeling pointing by clicking the "+" button.

   1. For each pointing observation, you can switch between Auto and Manual imagery selection modes by choosing which one you want in the imagery or position dialogs, or by right-clicking on the observation toggle and selecting/unselecting "Use Auto Position And Imagery" in the context menu.
      a. An observation with "Use Auto Position And Imagery" turned on will display the same imagery and rover position as the **Auto** observation.
      b. An observation with "Use Auto Position And Imagery" turned off will behave like the **Manual** observation, and allow you to select any imagery or rover position that you wish for that observation.
   2. A pointing observation can contain any number of Pointing Drafts. (See Pointing Draft Bar section below.)

iv. More commands are available by right-clicking (or command clicking on Mac) on an observation toggle to bring up a context menu.

b. Rover Position Dialog

i. Open it using the toggle on the far left.
ii. This controls which position the rover is at for the currently selected observation.
iii. Use the position dropdown selector to select which position you want.
    1. If the selected observation is in Auto mode, you can't adjust the rover position. Click the "Switch to Manual Selection" button to change the observation to Manual mode, and allow you to change the rover position.
    2. The rover's position and heading, pitch, and roll are visible in the input fields. You can only edit these if you have created your own user-defined position.
iv. The Create New Position button allows you to create a user defined position, and input the coordinates you want, such as if you got coordinates for a post-drive position from Rover Planners, and you want to put the rover there for modeling.
v. Alternatively, you can set the rover position using the marker pins in Traverse View, or enabling "Show Traverse Path" in Rover View Options and Scene Options shortcuts.

c. Imagery Selection Dialog

i. Open it using the toggle second from the left.
ii. The imagery selection is only editable when the current observation is in Manual mode. If in Auto mode, you can click the "Switch to Manual Selection" button to change the observation to Manual mode, and allow you to edit the imagery selection.
iii. .The list of toggles on the right shows the imagery that is loaded for the current observation. If you uncheck a toggle there it will unload that imagery.
iv. The list of toggles on the left is imagery available to load. It's organized by rover position. Change the rover position dropdown to display which imagery is available at that position.
v. Alternatively, you can set the imagery using the marker pins in Traverse View, or enabling "Show Traverse Path" in Rover View Options and Scene Options shortcuts.
vi. Types of Imagery
    1. JPL meshes. MIPL/IDS creates these for Rover Planner use, so they have a very high vertex density. They model the terrain very accurately but use a lot of memory and have a high CPU

rendering cost. You probably don't want to load more than one or two of these at a time. Best use: modeling IDD target areas, and rocks that stick up with lots of vertical relief. Worst use: loading many of these at once from multiple rover positions to see the lay of the land.

2. Low Fidelity meshes. These have a much lower vertex density than MIPL meshes. They are intended to model the shape of the terrain on a tens of cm scale, which is good enough for most image pointing needs. These use much less resources than MIPL meshes, so you can load many of them at a time without consuming a lot of memory or CPU. They also have color textures when available. Best use: modeling the lay of the land by loading these from multiple rover positions. Worst use: modeling IDD target areas, or rocks where shape detail is important.

3. Images. These are 2D images that are projected out into 3D space using camera models to determine the pixel vectors. Where the image hits a mesh, it will conform to the shape of the mesh (MIPL mesh first, Lo-Fi mesh second, HiRise mesh third). If there is no mesh in that area, the image will conform to a flat plane at the rover's feet that extends out about 500m, till it curves up into a dome shape. This flat plane approximation works well for creating image pointings from the same perspective (e.g. pointing within a Navcam taken at the same rover position). But doesn't work so well if pointing in imagery from a different perspective (e.g. pointing in Navcam taken from a former drive position). If you want to point within imagery from a previous rover position, it should be projected onto a mesh for best results.

4. In Scene Options, you can toggle the visibility of the various types of imagery. This can be very useful, especially turning off/on 2D imagery, since it often can be distracting and not useful when it’s not projected onto a mesh.

5. Note to keep in mind: The 3D cursor, targets, and the measurement tool, display a red indicator when placed on the flat plane approximation instead of an actual mesh. This is a warning that the position isn't on actual terrain mesh, so the target coordinates cannot be trusted to represent a real location on the terrain surface.

7. Pointing Draft Bar -- The second bar down from the top

a. Each observation can have multiple pointing drafts. A pointing draft is a possible pointing for an observation (like when writing an essay, you have 1st, 2nd, 3rd, drafts of a paper). A pointing draft can be as simple as a single frame, or as complex as a full 360 degree mosaic containing multiple sequences..
b. The Pointing Draft bar contains toggles for the pointing drafts associated with the selected observation (selected in the topmost Observation Bar). When you click on a pointing draft toggle, it selects which pointing draft you will view and are possibly able to edit.
c. Always present at the left is a special type pointing draft called “PUL”. The “PUL” pointing draft is the one that the server translates into sequence RML. Of all pointing drafts in an observation, the PUL pointing draft is the final version that the PUL will use to base their delivered sequence on. This pointing draft is only editable by a user with PUL account status.
d. Any user can create their own pointing draft in any observation. Other users can select it, and see edits being made to the pointing draft in real time. However, one user cannot edit another user's pointing draft. You may instead copy another user's pointing draft to your own and edit that. To do that, right click on the pointing draft toggle to open a context menu and select the "Duplicate" option.
e. When you have a pointing draft of your own selected, you may create image pointings while in Camera View mode. In Camera View mode, click one of the Add Image Pointing buttons in the upper left of the scene view and hold the mouse down and drag the pointing to where you want it in the scene.
f. The visibility of the left and right eye pointings can be turned on or off with the L and R toggles on the far left, and the transparency adjusted with the slider.
g. There are other options for editing pointing drafts in the Camera View Options such as adjusting the overlap snapping, and also by right clicking on frames to bring up a context menu to edit various frame attributes.
h. When an image is selected in a pointing draft, you can use the Subframe Edit View to drag the boundaries of the subframe where you want them. Additional subframe options are also available in Subframe Edit View Options.

8. Assessing Lighting in the Scene

a. Shadow Modeling

i. Shadow modeling may be turned on using the shadow toggle (ninja icon) on the right side of the lower bar. (This toggle is also available in Rover Options.)

1. Since shadow modeling uses a lot of processor power and will cause your computer fan to turn on, it turns off automatically after a period of time. You can disable this behavior and/or adjust its timeout duration in Settings (gear button on the lower right).

ii. To adjust the time of day that shadowing is modeled, use the mini time slider next to the shadowing toggle on the lower bar. The sun position will change, and the time of day in LST/LMST shown next to the slider will update. Alternatively, for more precise control over the time use the Time Dialog (clock toggle next to shadow modeling toggle).

iii. If you make the rover invisible ("Rover is visible" toggle in Rover View Options and Scene Options shortcuts), the rover will not be visible but it will still cast a shadow. This can be useful sometimes for modeling ARM shadowing of target locations.

iv. By default, shadow modeling only models shadows due to the rover body. But you can enable mesh self-shadowing for MIPL terrain meshes, so that features in the scene also cast shadows. Follow these steps:

1. Mesh shadows are only visible on MIPL terrain meshes, and are obscured by other types of imagery in the scene, so first disable the visibility of Image Sets and Lo-fi Meshes in Scene Options.
2. Right-click on a MIPL mesh (you may need to load one if there are none already loaded) and select "Mesh Self-Shadowing" in the context menu to toggle on/off self shadowing on that mesh.
3. Turn on shadow modeling like usual, via the shadow modeling toggle on the lower right bar, and adjust the time of day to see the shadows change on the MIPL mesh.

b. Sun cursor position tool

i. The default compass cursor has an orange ball next to it. This orange ball lies along the vector from the point you clicked on the terrain to the Sun's position at the current time of day. This allows you to assess lighting at an area in a couple ways

1. In Rover View, orbit the cursor position to get a sense of the angle of the Sun relative to the terrain at that point.
2. In Camera View, the shadowing of the ball indicates the shadowing that would occur on a sphere in the scene at that location. So if it's fully lit, you have have fully illuminated rock faces (zero phase), or if highly shadowed, you'll likely have shadowed rock faces.
   a. To find the best time of day for lighting: Put the cursor at the area of interest and view it in Camera View. Adjust the time of day with the time slider on the lower bar back and

forth to find the time when the Sun ball is the most illuminated. That time is the best time to image that location to minimize shadows.

c. Phase angle

   i. Under cursor tools in a toggle “Distance to Camera and Phase”. If you enable it, the phase angle and distance from the cursor point to the camera will be displayed. Phase angle is the angle between the vectors [target to camera eye] and [target to Sun] . So minimizing your phase angle means you’ll have a better lit target.

d. Using existing images

   i. If you have images of the scene already, then you know how the illumination looks at the time they were taken. To find this time, right-click on the image, and select “Image Info” from the context menu. This will display the time it was acquired in LST/LMST and LTST, along with other useful information about the image.

9. Time Dialog

a. You can change the sol being modeled, as well as set (and convert) times in LST/LMST and LTST.
b. Get the sunrise and sunset times for the modeled sol. Sunrise and sunset are being defined as the time when the center of the Sun crosses local level elevation 0 deg.
c. Turn on/off Sun and Earth graphics in the scene, and get local level az el coordinates for them. These are coordinate values are derived from spice kernels, and are then interpolated for performance, but should still be accurate to less than a degree.
d. There is a larger slider for more precise time setting than the mini-slider on the bottom bar.
e. Correct For Clock Drift -- (Currently disabled till I can verify sclk drift calculations) The rover’s clock can drift, which causes the rover’s idea of where the Sun is to be incorrect. This option applies a correction to the conversion from LST/LMST to LTST, so that the Sun’s position is in the right place for a given LST/LMST. This can be important for accurate shadow modeling.

10. Targets

a. Targets are currently imported from the RP sequence rml. So displayed targets are RP targets, and targets created by users in Viewpoint. MSLICE targets are not able to be shown at this point.

b. Users can create targets by positioning the 3D cursor where desired, and pressing Ctrl T or the “Create Target At Cursor Position” button in Cursor Options or Scene Options shortcuts.
c. Targets Dialog
    i. Target dialog, can edit name and coordinates (if its yours)
    ii. Can turn on/off the label visibility
    iii. Select groups by sol, and/or targets individually
    iv. Click select to have the view center on it to find it
d. Right click a target for options, like renaming, IDD fovs
e. Target dialog, can edit name and coordinates (if it’s yours)
f. Can turn on/off the label visibility
g. Select groups by sol, and/or targets individually
h. Click select to have the view center on it to find it
i. Tab on targets to cycle selection with cursor
j. Red circle is indicator of source: warning not from a mesh.

11. Measurement tool

a. Turn it on/off using the ruler toggle on the lower left bar.
b. Click the scene once to define a start position. Click on the scene a second time to define the end position.
c. If the “X” is white, it means you’re over mesh and using mesh coordinates for measurement. If the “X” is red, there is no mesh data there, and it using a flat plane approximation for the coordinates.

12. Scene options

a. Sun Brightness. Use the slider to adjust how bright the illumination is in the scene.
b. Images
    i. Images Are Visible. Toggle the visibility of any image sets that are currently loaded. This can be useful to see which meshes are loaded, since image sets are projected on top of meshes.
    ii. Include Thumbnails. Toggle the visibility of thumbnails in any image sets that are currently loaded.
    iii. Use Decorrelation Stretch. If there are color images in an image set, toggle using a decorrelation stretch on those images instead of the standard false color stretch.
c. Low Fidelity Meshes
    i. Lo-fi Meshes Are Visible. Toggle the visibility of any lo-fi meshes that are currently loaded.
    ii. Level Of Detail Cutoff. Some parts of meshes look better than others. And when you have multiple meshes loaded at once, you don’t want the

blurry stretched out parts of one mesh to obscure the nice high resolution detail in another mesh. This slider allows you to make the blurry stretched out parts of meshes invisible, so that they doesn't interfere with your view.

d. JPL meshes
    i. JPL Meshes Are Visible. Toggle the visibility of any JPL (MIPL / IDS) meshes that are currently loaded.
e. Shortcuts. These controls are also in other places, but are assembled here for convenience.
    i. 3d Cursor Is Visible. Toggle the visibility of the 3d cursor in the scene (also Ctrl-H for hide)
    ii. Rover Is Visible. Toggle the visibility of the rover. Even when it is invisible, it will still cast shadows.
    iii. Show Traverse Path. Toggle the visibility of the rover's previous positions and traverse path.
    iv. Show Pointing Info Overlay. Toggle the visibility of an informational overlay on top of pointing frames that shows details about the settings on each frame.
    v. Target Names Are Visible. Toggle the visibility of the target name labels in the scene.
    vi. Create Target at Cursor Position. This button will create a new target at the current position of the 3d cursor. (Also, Ctrl-T for target.)

13. Rover View options

a. Rover Is Visible. Toggle the visibility of the rover. Even when it is invisible, it will still cast shadows.
b. Rover Shadow Modeling. Toggle enabling shadow modeling in the scene. (The shadow modeling toggle on the right of the lower bar also does this.)
c. Image Pointings Hit Rover. In the case where an image pointing is pointed such that part of it (or all of it) covers some of the rover model, this toggles whether or not image pointings should pass through the rover and be modeled conforming to the terrain, or if they should collide with the rover, and be modeled conforming to the rover model.
d. Show Traverse Path. Toggle the visibility of the rover's previous positions and traverse path.
e. Zoom To Rover. This button will zoom view back to the rover.

14. Camera view options

a. Perspective
    i. Camera View Perspective. This dropdown controls which camera eye is used to view the scene.
    ii. Az El Grid Coordinate Frame. This dropdown controls the az el frame that is used for the coordinate labels at the edge of the view and the az el grid lines.

iii. Az El Grid Transparency. This controls the transparency of grid lines of the az el coordinate frame.
iv. Panoramic Screenshot. Take a 360 degree screenshot of the scene.
    1. It works by rotating the view 360deg while sampling a pixel strip in the center of the view as it rotates. There's a Preview button that takes a low quality fast screenshot so you can see if everything looks how you want it to. Next to it is the Screenshot button, that takes a high quality but slow screenshot. The coordinate frame dropdown controls which coordinate frame the view rotates in while making the screenshot, which can make a difference when the rover is on high slopes. The screenshot will be centered wherever you have the view centered when you press the button, and will also use the same zoom setting you have set. So it's a good idea to take a Preview screenshot to see if the center and zoom look good before taking the high quality one.

b. Pointing
    i. Constrain Frame Movement. Constrains the movement of pointing frames, so that they can move in: 1. any direction 2. azimuth only 3. elevation only. The coordinate frame for movement in azimuth and elevation is set by the “Az El Coordinate Frame” option above.
    ii. Show Pointing Info Overlay. Toggle the visibility of an informational overlay on top of pointing frames that shows details about the settings on each frame.
    iii. View All Mastcam/Pancam Pointings. This will show all the Mastcam/Pancam pointings from all PUL pointing drafts in the scene at the same time.
    iv. View All Navcam Pointings. This will show all the Navcam pointings from all PUL pointing drafts in the scene at the same time.

c. Overlap Snap
    i. Azimuth Overlap Snap
        1. On Left, and On Right. These toggle control whether overlap snapping is applied on the left side or right side of the pointing frame that is being dragging. Having both left and right selected can cause problems when moving frames in the center of a mosaic, since the left and right snapping can fight each other and cause the frame to jump around.
        2. Az Overlap value, with units. This is the amount of azimuth overlap that pointing frames should snap to.
        3. Correct For Elevation Angle. If enabled, this corrects for the effect of the the azimuth coordinate lines getting closer at lower or higher elevations, to give a more consistent overlap at all elevation angles.

4. Eye To Snap. Each pointing frame has a left and right eye. This controls whether the snapping is between left eyes, right eyes, or the region of stereo overlap between the left and right eyes.

15. Subframe options

a. Active Eye To Edit. This controls which eye you're editing in Subframe Edit View for your selected image pair.
b. Auto Subframe Adjustment
    i. None. When you adjust the subframe on one eye, it has no effect on the other eye.
    ii. Eyes Use Same Subframe. When you adjust the subframe on one eye, the other eye is adjusted to have the same subframe window.
    iii. Left/Right Eye Wraps The Right/Left. The non-dominant eye is subframed to cover the same area as the dominant eye. E.g., on Mastcam, the M34 is subframed to cover the same terrain as the M100.
c. Eye Pairing
    i. In the following descriptions, the dominant eye refers to M100 in the case of Mastcam, or the left eye for Pancam and Navcam. The non-dominant eye is M34 for Mastcam, or the right eye for Pancam and Navcam.
    ii. Paired. The left and right eyes both take images without moving the mast.
    iii. Unpaired. The non-dominant eye is adjusted in azimuth to point to the terrain center of the dominant eye. The mast moves in azimuth, but not in elevation. Not moving in elevation improves stereo processing for unpaired stereo pairs (thanks Bob Deen).
    iv. Unlinked. The non-dominant eye can be pointed independently of the dominant eye. This can be useful to get a subframe with good overlap, while avoiding rover hardware. Position the dominant eye where you want, then change the Subframe View Active Eye to the non-dominant eye, and drag its position wherever you want it. The non-dominant eye will be pointed in RNAV REL for MSL or MAST_REL for MER from the dominant eye's position.
d. Subframe Area Fraction
    i. The subframe area fraction is defined as: (number pixels in subframe) / (number of pixels in fullframe). Drag the slider or set the value in the input field to change the subframe area fraction.
    ii. Lock Area Fraction Value. If the subframe area fraction is locked, when you drag the subframe edge the other edges will be adjusted to maintain the same area fraction.
e. Shape Constraint
    i. None. Any type of rectangle is allowed.

ii. Square. The subframe is constrained to be a square.
iii. Horizontal. The subframe is

16. Cursor options
    a. Modeling lighting with Sun tool
    b. Other cursor shapes
    c. Arm FOVs
    d. Red sphere source warning
17. Simulations
    a. Types: downlink vs RP
    b.
18. Mosaic Gen
19. Chat
20. Settings

PUL Training

A. PUL account differences
    a. A PUL is able to write to PUL pointing draft, which is the one that the server turns into sequence RML.
    b. A PUL can delete observations. If a science user deletes an observation, it will send a request message to the PUL to delete that observation. The PUL may or may not be using it, so they make the decision whether to delete it or not.
    c. A PUL can change the Auto imagery selection. A different selection of imagery may be better on some days, and the PUL can make that call.
B. Choosing the right imagery -- there is usually a lot of possible imagery to choose from. Here are some considerations.
    a. For Image Sets, pointing in mast camera images projected on a mesh where both are from the current rover position is the most accurate.
        i. Even in areas where there is no mesh and the 2D images are using a flat plane / dome approximation, the pointing is still pretty accurate since the images were taken from a similar perspective as current camera perspective
    b. Loading meshes from previous positions is generally pretty accurate, but there can be some positioning error due to slip in the drive and imperfect localization. You can gauge the magnitude of the error by blinking 2D images from the current position on the older meshes and seeing how well the features in the terrain match the same features in images from the current position.

c. Loading 2D images from previous positions can be accurate if projected onto a mesh. Similar to above, to gauge the error, blink the images on a mesh from the rover's current position to see how well the features match up. The horizon line may or may not be accurate depending on how far the rover has moved from where the images were taken, and how well the plane/dome approximates the actual terrain. See how well it matches imagery from the current position if available, and use your best judgement to see if it looks reasonable or not.

d. Always be somewhat skeptical of Hazcam 2D images where there is no good mesh data underneath.
    i. The flat plane approximation works well between Navcam, Pancam, Mastcam, since they all have a similar perspective, even when the terrain doesn't model a flat plane very well. But Hazcams are taken from a different point of view than the mast cameras. This means that a rock in a hazcam image projected out onto a flat plane might be show up at a different point in space than the same rock in a mast camera image projected onto the flat plane. This error is especially apparent when the terrain doesn't model well as a flat plane.
    ii. Sometimes when no Navcam or other imagery of the horizon is available, the 2D hazcams can provide a decent estimate of the horizon line. But this can be more or less accurate depending on how the terrain pitches away from the rover at distance, and how well it fits the flat plane/dome model that it is projected on. If possible, compare the Hazcam 2D horizon line with a Navcam image from a current or nearby position to get a sense of how accurate it is.
    iii. Because of the perspective difference, and images are not projected onto the rover model shape,, the wheel location in hazcam images are not where the wheels really are from the mast cameras' perspective. This can be confusing to scientists. The wheel locations are where the rover model wheel locations are.

e. If the rover has been in an area for awhile, you can load navcam and mastcam lo-fi meshes from multiple previous positions to fill out an area with mesh. Use the LOD slider in Scene Options to cull vertices that are too stretched out. This takes less resources than using MIPL meshes, and has the advantage of color meshes when color textures are available.
f. Keeping imagery selection the same between observations reduces memory consumption and loading times.
g. The PUL can change the selection used for the Auto imagery selection for a planning sol. Right click on the "Auto" observation and select "Auto imagery selection" to open the dialog. Depending on the activities for the day, a different selection of imagery might be more useful than the default selection.

h. To improve the view of 2D images, it's sometimes helpful to change how the they overlap. Right click on an image and select "Cycle Image Overlap" to show a "Next" button. Clicking this button iterates through possible overlaps available. This only affects your view of the images, not other users' view.
i. Sometimes a single image in an image set looks bad for some reason and you'd rather not see it. Right click the image and select "Invisible" to make it invisible. You can right-click it again (even though it's invisible) and uncheck "Invisible" to make it visible. This only affects your view of the image, not other users' view.
j. When examining ARM targets in the work volume, it's often useful to increase the number of vertices in a 2D image that is overlaying a mesh so that it conforms to the mesh shape better. To do this, right click on the image and select "Conform Extra Closely". This only affects your projection of the image, not other users' projection.
k. Image Info
l. Open Image in Browser

C. User defined rover positions.
   a. From RP
   b. From Simulations
   c. From a target. (This will go away when Joystick mode is implemented.)
   d. In the future: Joystick mode TBD
   e.
D. More on Observations
   a. Copy Observations from/to other sols
      i. Right click on an observation toggle, and in the context menu select "Copy Observation To Planning Sol" or "Copy Observation From Planning Sol".
      ii. This is useful to:
         1. Add more pointings to a mosaic that was started on a previous sol.
         2. Model how post drive imaging will be pointed in the actual post-drive imagery at the real end of drive position, before we have the images or thumbnails down.
         3. Plan a large mosaic and take it over several sols, copying the observation to the next planning sol as you go.
         4. Retake frames of a mosaic that didn't complete.
      iii. In the case where you copy an observation from a previous planning sol, and the rover has moved since that sol: before selecting the observation, it's usually desirable to right click it and turn on "Use Auto Imagery And Position". Then it will load the pointings within the current imagery and with the rover at the current position, rather than the previous imagery it was planned in and probably an estimated rover end of drive position.
   b. Hierarchy of data storage:
      i. Planning Sol

1. Observations
    a. Rover Position
    b. Imagery Selection
    c. Pointing Drafts
        i. Sequences
            1. Image Pairs
                a. Coordinate Frame
                b. Left Eye
                    i. Coordinates
                    ii. Subframe extents
                    iii. Other Command Parameters
                c. Right Eye
                    i. Coordinates
                    ii. Subframe extents
                    iii. Other Command Parameters

E.

I. Overlap snapping
    A. The settings for overlap snap are adjustable in Camera View Options
    B. Turn on/off snapping on a frame side with the “On Left”, “On Right”, “On Top”, “On Bottom” toggles.
    C. Set the overlap snap amount and units for az and el. The default setting for Mastcam is 20% for az and el. Default for Navcam is 9 deg for az and el. Default for Pancam is 2 deg az, 1.5 deg el.
    D. You need to re-adjust your overlap after you change the elevation angle of your frames when pointing in LL or when “Correct for Elevation Angle” is turned on. Here’s why...

        1. Imagine the az, el sphere like an Earth globe with longitude and latitude lines. A navcam frame covers 45deg at elevation 0. But when it's pointed higher or lower towards a pole, the az lines close together and converge at the pole, which makes the frame effectively cover more azimuth than 45deg.

        2. LL is different from RNAV in that their spheres are rotated relative to each other by the amount of the rover tilt.

        3. If you shoot a 360 at RNAV elevation = 0 deg, you have 10 frames with 9deg overlap. If we're on a steep slope, and you shoot a 360 at LVL_AZ elevation = 0deg, the circle of the mosaic tilts up, so it's high on one side

and low on the other relative to RNAV. But it's distributed evenly along the great circle of LVL_AZ elevation 0, and has an even 9 deg of overlap in LVL_AZ frame. But if you look at the same pointing in RNAV, which has a rotated sphere relative to LVL_AZ, it will not have evenly spaced az overlap. The frames that are nearer to a pole will actually have less overlap in RNAV coordinates than 9deg, since the centers of the frames will have more az (longitude) lines between them.

4. The RSM effectively moves in RNAV space (actually JOINT space which is similar), so when you take some of those higher LVL_AZ frames and move them down, they will spread out as the RNAV lines spread out, and make your actual overlap between the frames smaller.

5. So that's why you need to re-adjust your overlap when you change the elevation of the frames if you pointed your mosaic in LVL_AZ. If you point in RNAV, the PMA, which operates in that frame, should keep a consistent overlap as you move your frames up and down in elevation.

6. The option "Correct for elevation angle" corrects for the effect of the the az lines getting closer at lower or higher elevations, to give a more consistent overlap at all elevation angles. So if that option is turned on, the az overlap will be elevation dependent even when using RNAV, so you need to re-adjust your overlap if you change elevation.

II. Subframing
    A. Active eye, modes, constraints

III. Pointing Type: Paired, Unpaired, Unlinked pointing.
    A. Flight software points paired and unpaired pointing differently. For paired pointing, two cameras are commanded in the same image command. Flight software points the average of their boresights at the specified coordinates, and each eye takes an image with the mast at the same joint angles. For unpaired pointing, one eye is commanded at a time, so flight software points each eye's boresight at the specified coordinates in turn, and images are taken with the mast at different joint angles for each eye. The only exception is when pointing in JOINT frame. JOINT frame specifies the mast joint angles to use directly, so paired and unpaired imaging both point the same (if using the same coordinates).
    B. In Viewpoint, paired pointing has the same meaning as above. Unpaired pointing means the eyes are unpaired, but also that they are centered on the same point on the terrain. Unlinked pointing means that the eyes are unpaired and are not centered on the same point on the terrain.

C. You can set this by right-clicking on selected image pairs (using shift for multiple selection), and hovering over “Pointing Type” and selecting Paired, Unpaired, or Unlinked. Alternatively, this can be set in the Frame Parameters dialog.

D. Paired pointing: The left and right eyes are taken with the mast and camera bar at the same position. Depending on toe-in, elevation angle, and the frame size, you might have more or less overlap between the frames.

    1. This is the best pointing type for stereo correlation, so use it if possible. But if the frame is pointed close enough to the rover that there is little overlap between the eyes, unpaired is a better option.

E. Unpaired pointing: The camera eyes are pointed separately so that the center of each frame is pointed at the same point on the terrain. This is accomplished in the code by raycasting to find the site target in the center of the dominant eye (M100 for Mastcam), then using inverse kinematics to find the joint angles to point the non-dominant eye (M34 for Mastcam) at that target. But adjusted so that the left and right eyes have the same elevation gimbal angle, for improved MIPL/IDS stereo processing (-- advice from Bob Deen).

F. Unlinked pointing. Each eye is pointed separately with the mast, as in unpaired pointing, but they are not centered at the same location on the terrain. This is useful sometimes to get subframes to overlap better when pointing very close to the rover and avoiding coverage of the ARM. When Unlinked pointing is selected, the non-dominant eye (M34 for Mastcam) can be positioned separately from the dominant eye (M100 for Mastcam) in the Subframe Edit View. The non-dominant eye is pointed RNAV_REL from the dominant eye.

IV. More on Pointing Drafts

A. Right click on a pointing draft toggle for a context menu with possible commands.

    1. Copy To PUL -- copies a pointing draft to the PUL pointing draft tab (aka the pointing draft that gets turned into a sequence).
    2. Copy to Observation -- copy this pointing draft to a new or already existing observation.
    3. Duplicate -- Make a copy of this pointing draft. This is useful for making a variation on a pointing draft, keeping the original version.
    4. Delete -- delete this pointing draft

B. As a PUL, I often create different options for an observation and let the scientists discuss what they like best, and/or how to tweak one of the options to be what they want. Keeping a pointing draft for each option is a good way to try out various ideas and keep all of your work.

    1. If you’re making a new possible version (rather than tweaking an existing one) use “Duplicate” to edit a copy of a pointing draft, rather than editing the existing pointing draft.
    2. Alternatively, use the “+” button to create a new empty pointing draft, select the frames you want from an existing pointing draft, Ctrl-C to copy them and Ctrl-V to paste them into the new pointing draft.

3. Use the above commands along with the "Copy To PUL" option to copy pointing drafts created by science team users into your own pointing draft where you can make use of them.

C. Selecting and moving frames
    1. Click on frames to select them
    2. Hold shift down and click on frames to select multiple frames. Or click and drag to use a rubber band box selection.
    3. Ctrl-a to select all frames
    4. Tab key -- points the Camera View to a frame, multiple tab presses select the next image pair in the pointing draft. Shift-tab goes the opposite direction. This can be useful for finding frames in a pointing draft quickly when you don't know where they are.

D. Frame Parameters Dialog
    1. This is where you can set the coordinate frame, compression, downlink priority, etc, for your image pairs.
    2. Select the pointing(s) you want to set parameters for, right click to get the context menu, and select "Frame Parameters".
    3. Coordinate frames
        a) SITE frame allows you to use two types of targets, CENTER and NAMED.
            (1) CENTER takes your current pointing and raycasts onto the scene to find the target currently in the center of the frame. It will use those coordinates in the sequence. That way you don't have to make an explicit target to point a frame at.
            (2) NAMED centers the frame on an actual named target that you can find in the Targets dialog, such as targets created by RPs.
            (3) On MER, if you point in SITE frame and the target is very far away (>100 m or more), the flight software will refuse to point at the target. (We've never triggered this condition that I know of, I've just seen that it exists in the code.) I don't know if this is the case on MSL. But if it is, it might be a good idea to avoid using SITE frame for images of things very far away, and use LL instead.
        b) All non-JOINT frames
            (1) When pointing in these frames with a single eye, usually you will want to set the pointing type to "unpaired". If it's set to paired and there are only filters enabled for a single eye, the sequence will enforce paired pointing by pointing a stereo pair at the coordinates without imaging to point the camera, then it will image the single eye without mast motion. This is because paired and unpaired imaging point

in slightly different directions, so if you have paired pointing specified, the sequence will use paired pointing. This can be useful when you want to reshoot something that was taken paired once, but you're using a single eye now, and you want to have the same pointing as before.

   c) RNAV_REL, LL_REL, JOINT_REL
      (1) Relative pointing requires something that it's relative to. In Viewpoint this is called the Relative Reference. This can be the previous frame (if not the first frame in a sequence), a Site target, or the Sun.
         (a) Viewpoint will prevent the first frame in a sequence from having a relative reference of "Previous Position". Usually you will want to set this frame to some other coordinate frame such as Site or LL. The first frame is chosen by the frame ordering setting in Sequence Parameters dialog. But you can override that and set any frame to be the first one by using "Override Ordering: 1" in Frame Parameters.
         (b) If you use a Site target for the Relative Reference, the sequence RML will start with an image command to point at that target but not take an image.
         (c)

4. Override Frame Order -- Ordering for a sequence is set in the Sequence Parameters dialog. But you can set this to override that order. For example, if you enable this on a frame and set the number to "1", that frame will be the first one in the sequence no matter what. If two frames have the same override value, the order is not determined between them.

5. Set Exposures For Subsequent Frames -- If you enable this on a frame, all subsequent images will use the exposures set in the exposure table by this frame.
   a) Usually, you'd also want to set Override Frame Order to "1" as well, so that is turned on by default when Set Exposures For Subsequent Frames is enabled.
6. Presets --
7.
   a)
   b) Order override
   c) Presets

        d) Coordinates. Concept of dominant eye. And how relates to paired/unpaired/unlinked.

    E.

V. Frame Parameters
    A. Setting params
    B. Coord frames
        1. Site targets, and types, Center, Named
            a) On MER, is flight rule error if target is very far away.
        2. LL paired and unpaired (can have paired with single eye).. Usually want unpaired in that case
        3. Relative frames, and setting the initial frame, and how it will come out as dummy.
        4. Setting subframe for exposure.
        5. Order override
        6. Presets
        7. Coordinates. Concept of dominant eye. And how relates to paired/unpaired/unlinked.

VI. Show Pointing Info Overlay
    A. A good way to check your work to make sure things are how you want them. You can turn it on in Camera View Options and Scene Options shortcuts.
    B. Per frame it shows ordering, seqid, filterset, coordinate frame, paired/unpaired/unlinked, downlink priority. If it's a relative coordinate frame, shows the Relative Reference. If ordering has been overridden, the ordering number is shown in cyan. If the frame is a subframe for exposure, it will indicate that as well.
    C. At top of the view, it also shows the duration and data volume estimates for each sequence in the pointing draft.

VII. Sequence Parameters
    A. Header info
    B. Time of day, this influences the focus settings
    C. Frame ordering. You can override this in Frame Parameters.
    D. Order Around Hardstops. This influences the ordering if a mosaic crosses the hardstops. If enabled, it will break the mosaic ordering up around the hardstops, so that it will not cross the hardstops. For example, say you have a 4x1 with left to right ordering, and the hardstops are in the center of the mosaic. If "Order Around Hardstops" is not enabled, the hardstops will be ignored, and the ordering would be 1, 2, 3, 4. If it's enabled, the ordering would be 3, 4, 1, 2. This can be useful at times for reducing azimuth slewing and saving actuator wear. For instance, I usually enable this for albedo pans to save a 360 slew around the hardstop in the middle of the mosaic.

VIII. Setting multiple sequences, "carcasses"

A. Layout your mosaic how you want it to look.
B. Select image pairs you want to be in a separate sequence, using shift to select multiple, or by using the rubber band box. Holding shift, right click on one of the selected images, hover on “Assign to Sequence” and either select “New Sequence” to create a new sequence, or select an existing sequence to assign them to. If you selected “New Sequence”, fill in a seqid that is different from existing seqids.
C. When a sequence is selected, the frames will have a different color from the other non-selected sequences, so you can see which frames belong to the currently selected sequence. Using “Show Frame Info Overlay” can also show which frames belong to which sequences.
D. If you turn on “Show Frame Info Overlay

IX. Focus setting modeling -- how it's done

X. Vseq, editing your sequence, iterating

XI. Examples.
   A. Point RNAV Rel.
      1.
   B. Make a large mosaic, set paired /unpaired. Set sequence info, frame info,

   C. Site target in workvol : extend ARM, check shadowing, paired, unpaired, or unlinked, obscured by deck or arm?, Post-drive, point LL direction., Post-drive point mosaic at target.
   D. post drive clast survey: get coords, and put rover there, estimate shadows, position frame, set params etc.

XII. Homework problems. Make a test, review sequences.

XIII. Future changes: Joystick mode, add LL slewing mode. Overlap maintenance, mosaic creation.

Light demo:

1. Four View Modes
   a. Traverse View
   b. Rover View
   c. Camera View
   d. Subframe Edit View
2. Imagery types
   a. Load Pos 143
   b. Show meshes and images,

   c. How images use flat plane when no mesh available. -- works well for planning imaging if from the same perspective. Before Viewpoint, 10 years of pancam image planning mostly used flat plane approximation.
   d. conform extra closely
   e. Decor stretch
   f. Image Info
   g. Open in browser
3. IDD motion
4. Shadowing
5. Ruler
6. Cursor
   a. Compass cursor Sun ball for assessing lighting
   b. Create target in scene at cursor position
   c. Also load RP targets
   d. Change cursor shape.
   e. Show IDD FOVs in scene - apxs, mi , rat butterfly
7. Simulations -- show downlink 143.
   a. Tracks
   b. Emily's targets -- no mesh in that area but still useful for planning imaging. So I advocate for not restricting targets to only be on MIPL mesh, but instead having a trustworthiness rating associated with them.
      i. 1. MIPL mesh from current position
      ii. 2. MIPL mesh from previous position, Lo-Fi mesh from current position.
      iii. 3. 3D HiRise Mesh, Lo-Fi mesh from previous position.
      iv. 4. Flat plane approximation
8. Observations
   a. Auto and Manual imagery loading
   b. Contain multiple pointing drafts. Multi-user in real time
   c. Drag and drop frames, with overlap snapping
   d. Mosaic Gen
   e. Set params with dialogs, and check with frame info overlay (resources too).
9. I know chat tools abound, but here's another one in case all others fail.