

Simulation with OpenSim - Best Practices

Overview

There is no scientific consensus on when a simulation is "good enough". Ultimately, you need to be sure you are asking the right question (such as "which muscles can contribute to knee flexion in this posture?") and that your results are robust to variations in model parameters. In general, you should avoid questions with absolute finite answers (e.g., "how much force was generated?") and instead use musculoskeletal simulations to understand relationships between model parameters and components.

Nonetheless, evaluating your results is one of the most important elements of generating and analyzing a simulation of human movement. In the sections below, we provide a set of best practices for using each of the tools in the OpenSim pipeline. These best practices are intended as a starting point to help you make sure you are collecting good data and analyzing and evaluating it effectively. This list is by no means exhaustive and depends highly on the nature of the motion you are studying and the research questions you are asking.

We've also provided a short [checklist](#) to use when evaluating your simulation.

Scaling a Model

Data Collection and Other Preparation:

1. When collecting data, take pictures of your subjects in the static pose. These pictures are valuable for evaluating the results of the Scale Tool.
2. Measure subject specifics, like height, mass, body segment lengths, mass distribution (if DXA is available), and strength (if a Biodex is available). You can use this data, along with marker positions, to best match the generic model to a specific subject.
3. Have your subjects perform movements to calculate functional joint centers at the hip, knee, ankle, and/or shoulders, and append the joint centers to your static trial data (see [Associating Data with a Motion](#)).
4. These are just a few tips. Also review our full guide to [Collecting Experimental Data](#).

Scale Settings:

1. Rely on markers that correspond to anatomical landmarks and functional joint centers (FJC) to position and scale the generic model.
 - a. See [How to Use the Scale Tool](#) for information about defining the measurement set for scaling and setting weights when positioning the model in the static pose.
2. Some segments, like the pelvis and torso, are often best scaled non-uniformly. For example, see the torso scale settings in the section on the Scale Factors Pane in [How to Use the Scale Tool](#).
3. Review [How Scaling Works](#) for more information about Scale Settings.

Evaluating your Results:

1. Scaling a model is an iterative process. Use the "preview static pose" option in the GUI. See the section on "Previewing Scale" in the [How to Use the Scale Tool](#) section for more information. After running preview, perform steps 2 to 5 described below.
2. Check the "Messages" window, which has information about the results of scaling, including the overall RMS marker error and the maximum marker error.
 - a. In general, maximum marker errors for bony landmarks should be less than 2 cm.
 - b. RMS error should typically be less than 1 cm.
 - c. Pay close attention to errors in the bony landmark and FJC markers when assessing the quality of your scaling results.
3. Visualize the scaled model's anatomical marker positions relative to the corresponding experimental markers to see how well the model "fits" the data. Use the pictures you took to assess the results, comparing the joint angles in the "Coordinates" window to the angles you observe in the pictures.
 - a. Do the hip, knee, and ankle angles after scaling match what you observe in the picture?
 - b. Are there any large mismatches between experimental and model markers? Can these mismatches be explained by examining the pictures you took?
 - c. If pictures aren't available, use what you know about a typical static pose capture. For example, the ankle angle is generally less than 5° and hip flexion angle is less than 10°.
 - d. Again, pay close attention to errors in the landmark and FJC markers when assessing the quality of your scaling results.
4. After examining the "Messages" window and performing a visual comparison, adjust the virtual markers and marker weightings to improve your results.
 - a. Again, avoid adjusting the positions of the landmark and FJC virtual markers to match the experimental markers.
5. Once you've adjusted the virtual marker positions and the scale settings, preview the new static pose. Reassess your results using steps 2 to 4 above. Once you are happy with your results, hit "Run" to generate a scaled model and adjust the virtual markers on the model to match all of the experimental markers.

Troubleshooting Tips:

1. It is common to iterate through Scale and Inverse Kinematics to fine-tune segment dimensions and marker positions that yield low marker errors for the task of interest.
2. Use coordinate tasks (Static Pose Weights) to set joint angles for troublesome joints that are very sensitive to how the markers are placed (commonly the ankle joint and lumbar joint). For example, if it is known that the foot is flat, an ankle angle can be provided and then the markers can be adjusted in order to match the known pose.
3. If using coordinates from a motion capture system, make sure that the joint/coordinate definitions match—otherwise, you may cause more harm than good.
4. The model has a built-in assumption that the global Y-axis is up. If your data doesn't fit this assumption, then consider transforming your data. You can use [Previewing Motion Capture \(Mocap\) Data](#) to determine the proper transform to apply.

Inverse Kinematics

Data Collection and Other Preparation:

1. When collecting experimental data, place three non-collinear markers per body segment that you want to track. You need at least three markers to track the 6-degree-of-freedom motion (position and orientation) of a body segment.
2. Place markers on anatomical locations with minimum skin/muscle motion.
3. These are just a few tips. Also review our full guide to [Collecting Experimental Data](#).

Inverse Kinematics Settings:

1. Weight "motion" segment markers (from a triad placed on the thigh segment, for example) more heavily than anatomical markers affixed to landmarks like the greater trochanter and the acromion, which can be helpful for scaling, but are influenced by muscle and other soft tissue movements during motion.
2. Relative marker weightings are more important than their absolute values. Therefore, a weighting of 10 vs. 1 is 10 times more important, whereas 20 vs. 10 is only twice as important. Markers are not necessarily tracked better because they both have higher weightings.
3. See [How Inverse Kinematics Works](#) and [How to Use the IK Tool](#) for more information about IK settings.

Evaluating your Results:

1. Total RMS and maximum marker errors are reported in the "Messages" window. Use these values to guide changes in weightings or, if necessary, to redo marker placement and possibly scaling. Maximum marker error should generally be less than 2-4 cm, and RMS under 2 cm is achievable. These guidelines will vary depending on the nature of the model and the motion being examined.
2. If using coordinates from a motion capture system, make sure that the joint/coordinate definitions match—otherwise, you may cause more harm than good.
3. Compare your results to similar data reported in the literature. Your results from an unimpaired average adult should generally be within one standard deviation.
4. If you are unsatisfied with the results, recheck the results of scaling.

Inverse Dynamics

1. Filter your raw coordinate data, since noise is amplified by differentiation. Without filtering, the calculated forces and torques will be very noisy.
2. Compare your results to data reported in the literature. Your results should be within one standard deviation of reported values.
3. Inspect results from Inverse Dynamics to check whether ground reaction forces were applied correctly or not. Are there large and unexpected forces at the pelvis? For gait, applying ground reaction forces should help reduce the forces computed by Inverse Dynamics at the pelvis.
4. See [How Inverse Dynamics Works](#) and [How to Use the Inverse Dynamics Tool](#) for more information about using the Inverse Dynamics Tool.

Static Optimization

Static Optimization Settings:

1. You can use IK or RRA results as input kinematics. If using IK results, you usually need to filter them, either externally or using the OpenSim analyze/static optimization field; if using RRA results, you usually do not have to filter.
2. For gait and many other motions, you need to add (append) residual actuators to the first free joint in the model (typically the ground-pelvis joint).
 - a. There should be one actuator for each degree-of-freedom (i.e., F_x , F_y , F_z , M_x , M_y , M_z).
 - b. These residual actuators are required because there is dynamic inconsistency between the estimated model accelerations and the measured ground reaction forces. This inconsistency can result from marker measurement error, differences between the geometry of the model and the subject, and inertial parameters.
 - c. Running RRA will reduce—but not eliminate—these residuals. Thus, appending actuators is still necessary.
3. See [How Static Optimization Works](#) and [How to Use the Static Optimization Tool](#) for more information.

Troubleshooting:

1. If the residual actuators or the model's muscles are weak, the optimization will take a long time to converge or will never converge at all.
 - a. If the residual actuators are weak, increase the maximum control value of a residual, while lowering its maximum force. This allows the optimizer to generate a large force (if necessary) to match accelerations, but large control values are penalized more heavily. In static optimization, ideal actuator excitations are treated as activations in the cost function.
 - b. If the muscles are weak, append Coordinate Actuators to the model at the joints in the model. This will allow you to see how much "reserve" actuation is required at a given joint and then strengthen the muscles in your model accordingly.
 - c. If troubleshooting a weak model and optimization is slow each time, try reducing the parameter that defines the maximum number of iterations.
2. Static optimization works internally by solving the inverse dynamics problem, then trying to solve the redundancy problem for actuators/muscles using the accelerations from the inverse dynamics solution as a constraint. If a constraint violation is reported, this could be a sign that the optimizer couldn't solve for muscle forces while enforcing the inverse dynamics solution.
 - a. This likely means that there is noise in the data or there is a sudden jump in accelerations in one frame.
 - b. In this case, you should examine the inverse dynamics solution to determine the problematic frame, and fix/interpolate the data during this portion of the motion.
3. If your model has passive elements (e.g., ligaments or springs), you should use OpenSim version 3.3 or later. Prior to version 3.3, the forces generated by passive elements were not properly accounted for.

Evaluating your Results:

1. Are there any large or unexpected residual actuator forces?

- Find EMG or muscle activation data for comparison with your simulated activations. Does the timing of muscle activation/deactivation match? Are the magnitudes and patterns in good agreement?

Residual Reduction

RRA Settings:

- You should replace the muscles in your model with residual actuators and ideal joint actuators. Residual reduction is a form of forward dynamics simulation that utilizes a tracking controller to follow model kinematics determined from the inverse kinematics. Computed Muscle Control (CMC) serves as the controller, but without muscles, the skeleton of the model can be used to determine a mass distribution and joint kinematics that are more consistent with ground reaction forces.
- Optimal forces for residuals should be low to prevent the optimizer from "wanting" to use residual actuators (an actuator with large optimal force and low excitation is "cheap" in the optimizer cost).
- To help minimize residuals, make an initial pass with default inputs, then check residuals and coordinate errors. To reduce residuals further, decrease tracking weights on coordinates with low error. You can also try decreasing the maximum excitation on residuals or the actuator optimal force.
- Typically, you should "lock" the subtalar and mtp joints in the *.osim file.
- Make sure "use_fast_optimization_target" is false. This allows the kinematics to be slightly adjusted to account for dynamic inconsistencies. This is the default in the settings files distributed with OpenSim or created from the GUI. See [How CMC Works](#) for a comparison of the "slow" and "fast" targets. In the GUI, this option is hidden for RRA and can be viewed by opening the xml settings file.
- The "cmc_time_window" in the settings file should be 0.001 s for RRA. This is the default in the settings files distributed with OpenSim or created from the GUI.
- See [How RRA Works](#) and [How to Use the RRA Tool](#) for more information about RRA settings.

Troubleshooting:

- Check the pelvis COM location in the Actuator files.
- If RRA is failing, try increasing the maximum excitation for residuals by orders of magnitude until the simulation runs, then try working your way back down while also "relaxing" tracking weights on coordinates.
- If residuals are very large (typically, this is greater than 2-3 times body weight, depending on the motion), there is probably something wrong with either (i) the scaled model, (ii) the IK solution, or (iii) the applied GRFs. To double-check that forces are being applied properly, visualize GRFs with IK data (you can use the [Previewing Motion Capture \(Mocap\) Data](#) function in the GUI).
- If there is pelvis drift and/or F_y is not centered around zero, check that the body mass and force calibration are correct.
- When using the example RRA Actuators .xml file, you should note that residual forces are applied to the center of mass (COM) of the unscaled pelvis; however, if you scale the model, the COM of the pelvis can change. Although the effect may be small, you should change the location of the residual force actuators in the RRA Actuators file to correspond to the scaled pelvis COM.

Evaluating your Results:

- RMS difference in joint angle during the movement should be less than 2-5° (or less than 2 cm for translations).
- Peak residual forces should typically be less than 10-20 N. Average residuals should typically be less than 5-10 N.
 - The size of residuals will depend on the type of motion being studied. For example, residuals for high-speed activities like sprinting will typically be larger than walking.
 - Residuals will also be larger if there are external forces that you have not accounted for, such as a subject walking with a handrail.
- Compare the residual moments from RRA to the moments from Inverse Dynamics. You should see a 30-50% reduction in peak residual moments.
- Compare the joint torques/forces to established literature (if available). Try to find data with multiple subjects. Your results should be within one standard deviation of the literature.

The table below shows an example of threshold values used to evaluate RRA results for full-body simulations of walking and running.

Thresholds:	GOOD	OKAY	BAD
MAX Residual Force (N)	0-10 N	10-25N	> 25 N
RMS Residual Force (N)	0-5 N	5-10 N	> 10 N
MAX Residual Moment (Nm)	0-50 Nm	50-75 Nm	>75 Nm
RMS Residual Moment (Nm)	0-30 Nm	30-50 Nm	>50 Nm
MAX pErr (trans, cm)	0-2 cm	2-5 cm	>5 cm
RMS pErr (trans, cm)	0-2 cm	2-4 cm	>4 cm
MAX pErr (rot, deg)	0-2 deg	2-5 deg	> 5 deg
RMS pErr (rot, deg)	0-2 deg	2-5 deg	> 5 deg

Computed Muscle Control

CMC Settings:

- The reserve actuators are torques that are added about each joint to augment the actuator's force in order to enable the simulation to run (reserves turn on when an actuator cannot produce the needed force at a given time point). To help minimize reserve torques, make an initial pass with default inputs, and then check reserves, residuals, and joint angle errors. To reduce reserves further, decrease tracking weights on coordinates with low error.

- Optimal forces for reserves should be low to prevent the optimizer from "wanting" to use reserve actuators (an actuator with a large optimal force and a low excitation is "cheap" in the optimizer cost). If larger forces are needed for a successful simulation, increase the maximum control value of residuals. The residuals will then be able to generate sufficient force, but will be penalized for doing so.
- If you are still getting high reserves for a particular degree-of-freedom during a particular time range in the simulation, it may be useful to examine more closely the muscles which span the degree-of-freedom during that time region (see [TipsForDebuggingMuscleActuatedSimulations.pdf](#)). In particular:
 - Check passive muscle forces (e.g., quadriceps). Large passive forces (from large knee flexion angles) may induce active forces in the antagonistic muscles (e.g., hamstrings, gastrocnemius), which may not be desired. Passive forces cannot be controlled in CMC; they are purely a function of the whole-body kinematics of the motion. Although tempting, do not increase the maximum isometric force excessively unless you know its consequences in the muscle model. Since the passive muscle force is modeled as a function of maximum isometric force, if you increase the maximum isometric force in the hope of making your active muscles stronger, you will also be increasing the passive forces in the muscles as well, thereby not helping the situation. To decrease passive muscle forces, you want to reduce the passive muscle stiffness property of muscle (or, more specifically, increase the `FmaxMuscleStrain` parameter in the `Thelen2003Muscle`).
 - Check normalized fiber length during the motion (using the Plotter Tool). Is the muscle acting at suboptimal fiber lengths (i.e., less than 0.8 or greater than 1.2) at the time where the reserves are being generated? If so, then you may consider modifying either the tendon slack length or the optimal fiber length of the muscle so that it is operating more optimally (and, thus, capable of generating a greater force) during this time in the simulation. Although many cadaver studies report the optimal fiber length of a muscle, the tendon slack length is almost never reported, even though it is especially sensitive to the operating region on the force-length curve. Small adjustments to the tendon slack lengths can, therefore, be justified in reducing the reserves on the basis that we don't have as much confidence in this value to begin with.
- You should almost always use the "fast" target for CMC. Note that this is the default (e.g. if you run CMC through the GUI). You can only change this value in the xml setup file for the tool. The fast target requires the joint accelerations at each time step to be matched to the RRA results. "Fast" target should work for normal subjects; "slow" target may be needed for subjects with pathologies.
- Start CMC at least 0.03 seconds before the point where you want to start analyzing your data, as CMC requires 0.03 seconds to initialize.
- See [How CMC Works](#) and [How to Use the CMC Tool](#) for more information.

CMC Troubleshooting:

- If CMC is failing, try increasing the maximum excitation for reserves and residuals by orders of magnitude until the simulation runs, then try working your way back down while also "relaxing" tracking weights on coordinates.

Evaluating your Results:

- Peak reserve actuator torques should typically be less than 10% of the peak joint torque.
- Peak residual forces should typically be less than 10-20 N; peak residual moments, less than 75 Nm (depending on the type of motion).
- Double-check your kinematics in comparison to RRA. Generally, they should match well as long as you are using the "fast" target.
- If performing an Induced Acceleration Analysis, you should verify that reserves and residuals contribute less than 5% to the net acceleration of interest.
- Compare the simulated activations to experimental EMG data (either recorded from your subject or from the literature). Activations should exhibit similar timing and magnitude to EMG data. You can also compare your muscle activations and/or forces to other simulations from SimTK or the literature.

The table below shows an example of threshold values used to evaluate CMC results for full-body simulations of walking and running.

Thresholds:	GOOD	OKAY	BAD
MAX Residual Force (N)	0-10 N	10-25N	> 25 N
RMS Residual Force (N)	0-10 N	10-25 N	> 25 N
MAX Residual Moment (Nm)	0-50 Nm	50-75 Nm	>75 Nm
RMS Residual Moment (Nm)	0-30 Nm	30-50 Nm	>50 Nm
MAX pErr (trans, cm)	0-1 cm	1-2 cm	>2 cm
RMS pErr (trans, cm)	0-1 cm	1-2 cm	>2 cm
MAX pErr (rot, deg)	0-2 deg	2-5 deg	> 5 deg
RMS pErr (rot, deg)	0-2 deg	2-5 deg	> 5 deg
MAX Reserve (Nm)	0-25 Nm	25-50 Nm	> 50 Nm
RMS Reserve (Nm)	0-10 Nm	10-25 Nm	> 25 Nm

Forward Dynamics

- Forward dynamics simulations are sensitive to initial conditions, and it is good practice to double-check that the initial conditions are appropriate for the desired simulation.
- If the Forward Dynamics Tool fails gracefully (i.e., without crashing OpenSim) or the output of the Tool drifts too much (i.e., the model "goes crazy"), shorten the interval over which the Forward Dynamics Tool runs (i.e., make `initial_time` and `final_time` closer to each other in the Forward Dynamics Tool setup dialog box or setup file). Open-loop forward dynamics tends to drift over time due to the accumulation of numerical errors during integration.
- Forward dynamics will be slow and generate very large results files if the problem is stiff and the integrator has to take many small time steps. Some common cases that result in stiff problems:
 - Bodies with small masses relative to other model bodies make the system stiff. Small masses combined with large forces can also make the problem stiff.

- b. If there's a contact model, check the settings. Try reducing the "transition_velocity" if forward dynamics is slow. Also having lots of damping in the contact model combined with a high stiffness value can also lead to problems.
4. Having many wrapping surfaces in your model can increase computation time. Cylindrical wrapping surfaces tend to have better performance than ellipsoidal wrapping surfaces.

Contributors

Stanford University - Jennifer Hicks, Ajay Seth, Sam Hamner, Matt DeMers

University of Delaware - Jill Higginson, Brian Knarr, Amber Collins, Elisa Schrank, Chris Henderson