FAQs

1. **Where can I access the code and data from the tutorial?**

   https://github.com/aclouthier/SSM-morphable-model

2. **How do I implement principle component analysis?**

   You can find functions for PCA in various statistical software and coding languages. Matlab has the pca() function and the scikit-learn package in Python has a PCA function as well.

3. **How many shapes do I need to build a shape mode?**

   If you have three, you can generate a statistical shape model, but since it only represents the variation among those three, it probably won't be very useful! Generally more is better. The quality of a statistical shape model can be quantified using metrics like Generalizability, Specificity, and Compactness. There have also been methods proposed for determining sample sufficiency, such as that by Mei et al. (2008).

4. **Where can I find out more about COMAK?**

   There is a SimTK page and GitHub page for this simulation routine (OpenSim-JAM).

5. **Where can I find more information about the knee model used?**

   We used the model developed by Lenhart et al. (2015).

6. **If we create an SSM of healthy joints and diseased joints separately, would the PCs represent different variations? How can we deal with this?**

   Yes, there is a good chance that the modes of variation in separate SSMs would be different. This could be of interest if you want to see what shape features vary most in healthy vs. diseased joints. A common way to handle it would be to make one combined SSM so that you have only one set of modes of variation. Then these modes can be compared between groups to identify differences in geometry.

7. **How do you verify your point correspondence is of high quality and not distorting the mesh so that the SSM captures true shape changes?**

   This may be different depending on the method used to establish point correspondence. In general, you want to verify that the resulting meshes look correct and not distorted and that the corresponding nodes are on the same anatomical locations. This could be done by plotting a selection of node IDs across different meshes to verify they are on the same location, or colouring the mesh based on node ID to make sure this matches across meshes in the data set.

Q8. For your COMAK simulations: do you keep the ligament properties, e.g., stiffness, the same in all morphed models or do you calibrate/tune the parameters based on each geometry?
In this study, we kept the ligament properties the same for all morphed models. We wanted to observe the effects of geometry alone. However, the reference posture was updated based on a settling simulation. This reference posture is used to calculate the ligament slack lengths based on the resting length in the reference posture and a reference strain.

Q9. What difficulties do you see when trying to automatically sample bone geometry and ligament/muscle attachment sites from an SSM (i.e., for Monte Carlo simulations) and create MSK models from it?
It is possible to generate unrealistic geometries using the SSM, especially with large deviations from the mean. In this case, staying within +/-3 standard deviations from the mean kept all the generated geometries looking reasonable. If more than one principal component is varied at once (like in a Monte Carlo approach), it may be necessary to reduce the range. For the COMAK simulation, it is important that the cartilage mesh is smooth and of high quality (consistent triangle size/dimensions). Therefore, we remeshed the cartilage after morphing the model to ensure this.

Unanswered questions from the webinar Q&A

Q1. When analyzing your simulation results, how would you suggest assessing effects of the varying geometry and the resulting kinematic compensations with something like COMAK where these are allowed to vary?

One of the benefits of the simulation approach is that it is possible to isolate the effects of geometry. Since COMAK optimizes the muscle activations and secondary kinematics simultaneously, it is possible to change the geometry and observe the effects on the kinematics, ligament forces, and cartilage contact pressures. The resulting kinematics and forces/pressures will be linked, but both are occurring due to the change in geometry.

Q2. For ligaments, you find the closest node IDs on the reference mesh, but when updating the new mesh you attach the ligaments to those exact node IDs. Would that cause any inaccuracy?

Since the bone meshes are composed of 1000s of nodes, the distance between the ligament attachment site and the closest node ID will likely be very small (<~1mm, depending on the density of the mesh). The error associated with identifying the attachment site location on the reference mesh (either by anatomical texts or from imaging) will likely be much greater than the error associated with “snapping” it to the closest node on the mesh.

The error in attachment site mapping will also largely depend on how well the node correspondence method captures anatomical correspondence. That is why it is important to visualize the same node ID across the corresponding meshes or generated shapes. The same node ID should be located at the same anatomical region for all meshes. If they do not, you may need to adjust your node correspondence method or explore a new method.

Q3. Do we need non-rigid registration for soft tissue like fat around a hard tissue? What’s the best protocol to apply this

It depends on what you intend to do with the soft tissue. For the morphable COMAK model, cartilage was included in the input shape model, but ligament and muscle attachment sites were defined from the bone surfaces. Therefore, ligament and muscle did not need to be segmented or registered across people. Other soft tissues, like fat and skin, were not needed to estimate the outcome measures of interest (i.e., joint contact forces, muscle forces, and contact pressures).

If you wish to create a shape model that accounts for segmented soft tissue like fat or muscle, the registration becomes more complicated because these tissues are highly deformable (i.e., their shape changes depending on a person’s position). I am not aware of one set protocol for this, but you would need to be careful about standardizing the position across participants. In theory, you could use the same non-rigid node correspondence methods that we use to establish correspondence across bones. However, you would need to interpret the resulting shape model with caution, because soft tissue shape is influenced by many factors (i.e., joint position, orientation relative to gravity).

Q4. When segmenting incomplete bones (e.g., femoral condyles without the whole femur), how do you decide the "cut-off" point in the segmentation, since this will create non-meaningful variance in shape?

There are a number of ways you could try to standardize the “cut-off” point when you are analyzing only a portion of a long bone. One option is to crop the bones and maintain a consistent length:width ratio. For example, for the distal femur, you can calculate each bone’s medial-lateral breadth after you have aligned the bones in an anatomical coordinate system. Then, you can define a target length as some ratio of this medial-lateral breadth and crop the shaft to meet this ratio. For example, you could crop all the femurs in your dataset such that the length is 80% of the medial-lateral breadth. Another approach would be to use an anatomical landmark. For example, for the proximal humerus, you could crop the shaft at the deltoid tuberosity.

Each approach has limitations and caveats. For example, people with advanced osteoarthritis (OA) tend to have wider joint surfaces. Therefore, for those with advanced OA, the medial-lateral breadth of their femur may be wider due to their disease, resulting in more of the shaft being included in the shape model relative to people with less advanced disease.

Q5. Addressing a rather basic aspect of MRI sample acquisition with respect to your great work, is it advisable to use the same MRI device for all samples for a statistical model (for a fixed weighting such as T1), or is the modelling process robust enough to produce valid results even for multiple-device / manufacturer data?

If you use meshes from different MRI devices to construct your shape model, your shape model will include ALL sources of variation in shape. That is, it will include the real anatomical variation across people's bone shapes, but it will also include variation due to the different imaging devices or error in segmentation.

Whether or not the shape model is "valid", therefore, would depend on the error (or difference) in shape segmented from scans from the two different MRI devices. For instance, one device or sequence may be less effective at capturing extreme features (e.g., osteophytes) than another device. This could lead to issues if it appears that half the sample has less osteophytes than the other half, when really it is due to measurement error. To determine what is "valid", I would encourage the researcher to assess the error, or difference, between the two MRI devices and/or perform sensitivity analyses to see how the shape model is affected by combining shapes from different MRI devices. If there is a systematic difference in shape between individuals from device A vs. device B, that might suggest that the device error is overshadowing real anatomical variation in shape. Of course, the "validity" of the shape model will also depend on what the model is being used to do. For example, if we want to do an initial exploration of major sources of variation in a dataset, a higher measurement error may be acceptable. If we want to identify people with a specific disease with high accuracy, then the same measurement error may not be acceptable.