Lesson 33: Principal Component Analysis

Which of the following three scatterplots (showing x vs. y, x vs. z and y vs. z) for the same three-dimensional data gives us the best picture about the actual layout of the data in space?

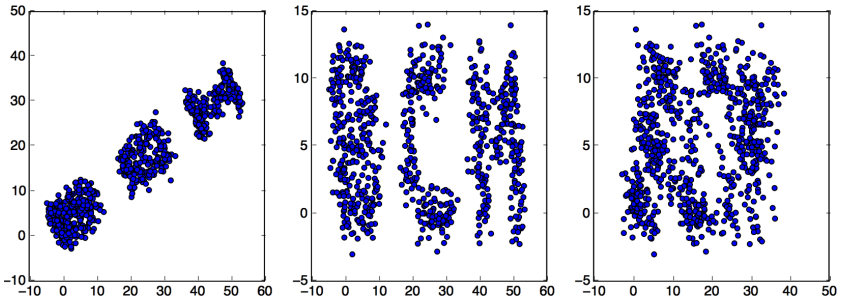

Yes, the first scatter plot looks very useful: it tells us that x and y are highly correlated and that we have three clusters of somewhat irregular shape. But remember: this data is three dimensional. What is we saw it from another, perhaps better perspective?

~~(No spoilers here, but we'll add the figure after the lecture!)~~

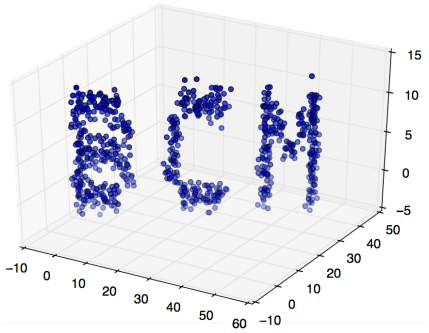

Let's make another experiment. Go to https://in-the-sky.org/ngc3d.php, disable Auto-rotate and Show labels and select Zoom to show Local Milky Way. Now let's rotate the picture of the galaxy to find the layout of the stars.

Think about what we've done. What are the properties of the best projection?

We want the data to be as spread out as possible. If we look from the direction parallel to the galactic plane, we see just a line. We lose one dimension, essentially keeping just a single coordinate for each star. (This is unfortunately exactly the perspective we see on the night sky: most stars are in the bright band we call the milky way, and we only see the outliers.) Among all possible projections, we attempt to find the one with the highest spread across the scatter plot. This projection may not be (and usually isn't) orthogonal to any axis; it may be projection to an arbitrary plane.

We again talk about two dimensional projection only for the sake of illustration. Imagine that we have ten thousand dimensional

data and we would like, for some reason, keep just ten features. Yes, we can rank the features and keep the most informative, but what if these are correlated and tell us the same thing? Or what if our data does not have any target variable: with what should the "good features" be correlated? And what if the optimal projection is not aligned with the axes at all, so "good" features are combinations of the original ones?

We can do the same reasoning as above: we want to find a 10-dimensional (for the sake of examples) projection in which the data points are as spread as possible.

How do we do this? Let's go back to our everyday's three dimensional world and think about how to find a two-dimensional projection.

Imagine you are observing a swarm of flies; your data are their exact coordinates in the room, so the position of each fly is described by three numbers. Then you discover that your flies actually fly in a formation: they are (almost) on the same line. You could then describe the position of each fly with a single number that represents the fly's position along the line. Plus, you need to know where in the space the line lies. We call this line the first principal component. By using it, we reduce the three-dimensional space into a single dimension.

After some careful observation, you notice the flies are a bit spread in one other direction, so they do not fly along a line but along a band. Therefore, we need two numbers, one along the first and one along the — you guessed it — second principal component.

It turns out the flies are actually also spread in the third direction. Thus you need three numbers after all.

Or do you? It all depends on how spread they are in the second and in the third direction. If the spread along the second is relatively small in comparison with the first, you are fine with a single dimension. If not, you need two, but perhaps still not three.

Let's step back a bit: why would one who carefully measured expressions of ten thousand genes want to throw most data away and reduce it to a dozen dimensions? The data, in general, may not and does not have as many dimensions as there are features. Say you have an experiment in which you spill different amounts of

two chemicals over colonies of amoebas and then measure the expressions of 10.000 genes. Instead of flies in a three-dimensional space, you now profile colonies in a 10,000-dimensional space, the coordinates corresponding to gene expressions. Yet if expressions of genes depend only on the concentrations of these two chemicals, you can compute all 10,000 numbers from just two. Your data is then just two-dimensional.

A technique that does this is called Principle Components Analysis, or PCA. The corresponding widget is simple: it receives the data and outputs the transformed data.

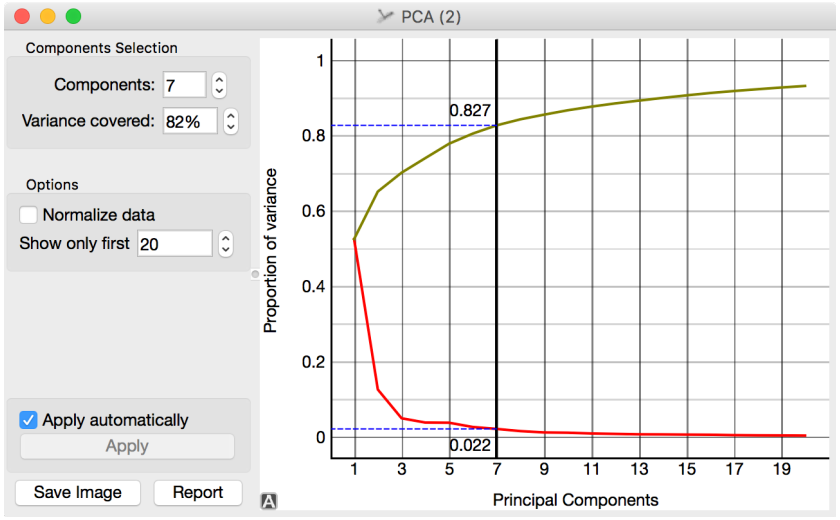

The widget allows you to select the number of components and helps you by showing how much information (technically: explained variance) you retain with respect to the number of components (brownish line) and the amount of information (explained variance) in each component.

The PCA on the left shows the scree diagram for brown-selected data. Set like this, the widget replaces the 80 features with just seven - and still keeping 82.7% of information. (Note: disable "Normalize data" checkbox to get the same picture.) Let us see a scatter plot for the first two components.

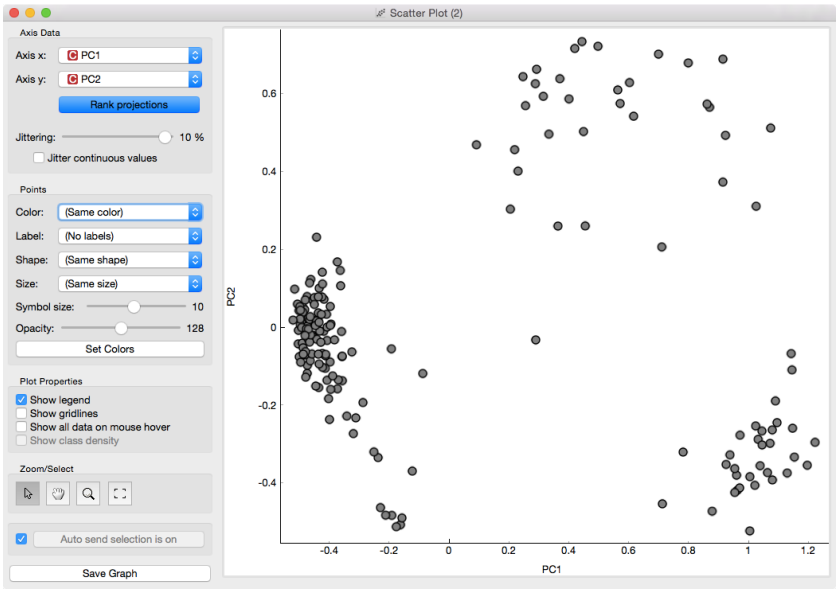

The axes, PC1 and PC2, do not correspond to particular features in the original data, but to their linear combination. What we are looking at is a projection onto the plane, defined by the first two components. When you consider only two components, you can imagine that PCA put a hyperplane into multidimensional space and projecting all data into it.

Note that this is an unsupervised method: it does not care about the class. The classes in the projection may be be well separated or not. Let's add some colors to the points and see how lucky we are this time.

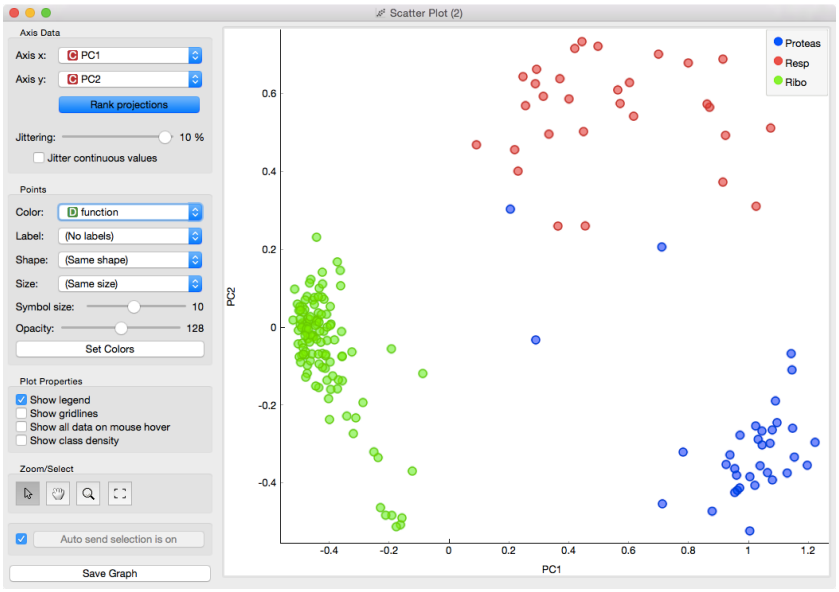

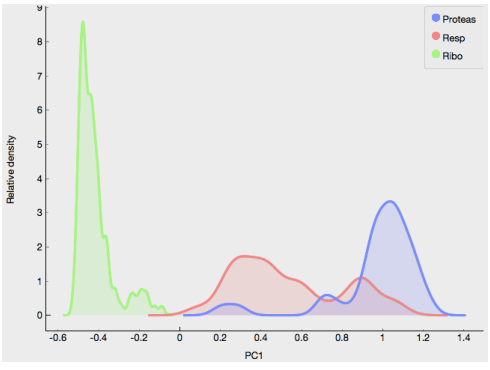

The data separated so well that these two dimensions alone may suffice for building a good classifier. No, wait, it gets even better. The data classes are separated well even along the first component. So we should be able to build a classifier from a single feature!

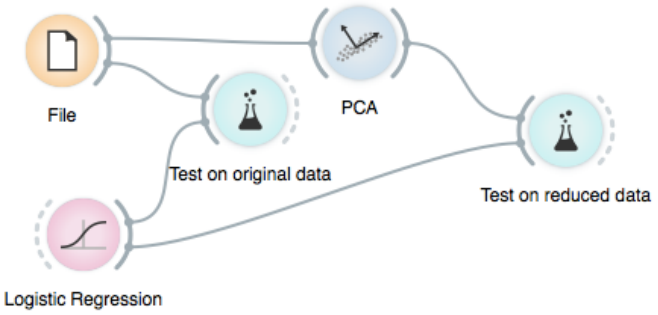

In the above schema we use the ordinary Test & Score widget, but renamed it to "Test on original data" for better understanding of the workflow.

On the original data, Logistic regression gets 98% AUC and classification accuracy. If we select just single component in PCA, we already get a 93%, and if we take two, we get the same result as on the original data.

PCA is thus useful for multiple purposes. It can simplify our data by combining the existing features to a much smaller number of features without losing much data. The directions of these features may tell us something about the data. Finally, it can find us good two-dimensional projections that we can observe in scatter plots.

Lesson 34: Mapping the Data

One thing that PCA does not offer is the sense of distances. It may, but it is not specifically designed for this. Besides, PCA needs features, coordinates. What if we know only distances between objects?

Imagine a foreign visitor to the US who knows nothing about the US geography. He doesn't even have a map; the only data he has is a list of distances between the cities. Oh, yes, and he attended the Introduction to Data Mining.

If we know distances between the cities, we can cluster them.

For this example we retrieved data from http://www.mapcrow.info/united_states.html, removed the city names from the first line and replaced it with "31 labelled".

The file is available at http://bit.ly/1Lm6UEr. To load it, unzip the file and use the File Distance widget from the Prototypes add-on.

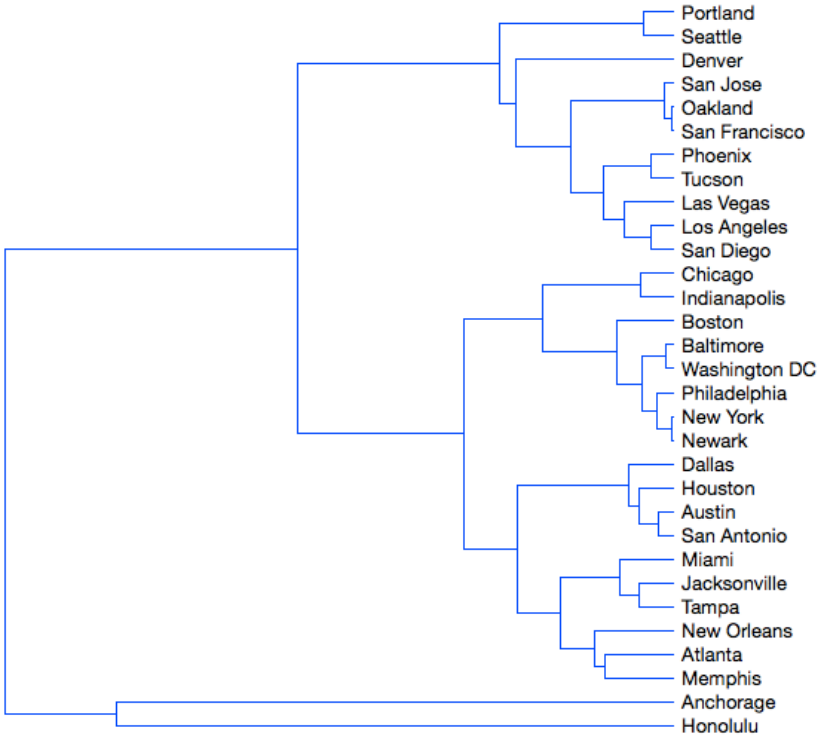

How much sense does it make? Austin and San Antonio are closer to each other than to Houston; the tree is then joined by Dallas. On the other hand, New Orleans is much closer to Houston than to Miami. And, well, good luck hitchhiking from Anchorage to Honolulu.

As for Anchorage and Honolulu, they are left-overs; when there were only three clusters left (Honolulu, Anchorage and the big cluster with everything else), Honolulu and Anchorage were closer to each other than to the rest. But not close — the corresponding lines in the dendrogram are really long.

The real problem is New Orleans and San Antonio: New Orleans is close to Atlanta and Memphis, Miami is close to Jacksonville and Tampa. And these two clusters are suddenly more similar to each other than to some distant cities in Texas.

We can't run k-means clustering on this data, since we only have distances, and k-means runs on real (tabular) data. Yet, k-means would have the same problem as hierarchical clustering.

In general, two points from different clusters may be more similar to each other than to some points from their corresponding clusters.

To get a better impression about the physical layout of cities, people have invented a better tool: a map! Can we reconstruct a map from a matrix of distances? Sure. Take any pair of cities and put them on paper with a distance corresponding to some scale. Add the third city and put it at the corresponding distance from the two. Continue until done. Excluding, for the sake of scale, Anchorage, we get the following map.

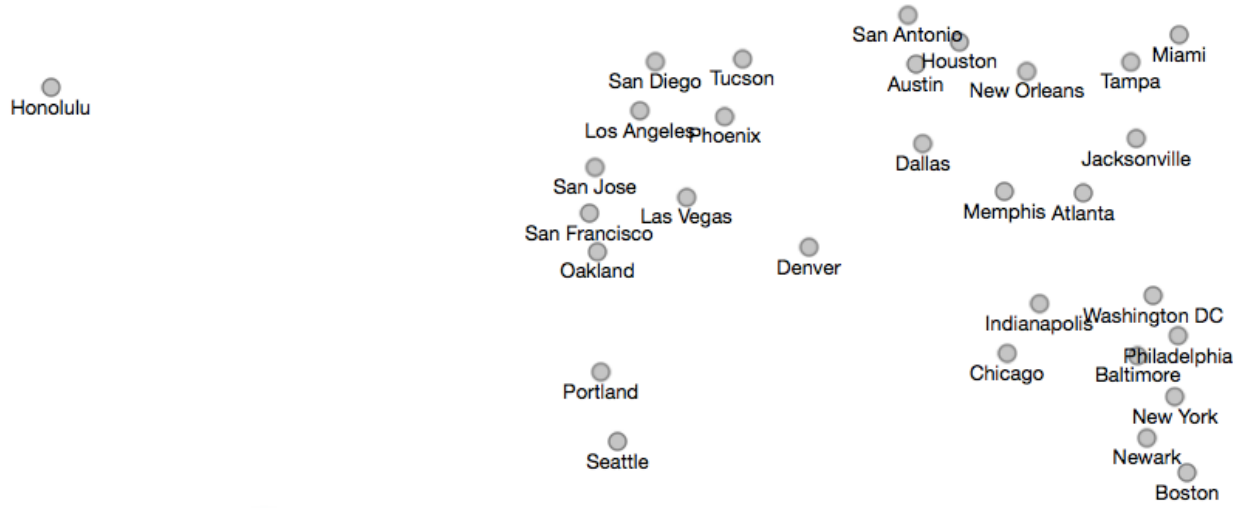

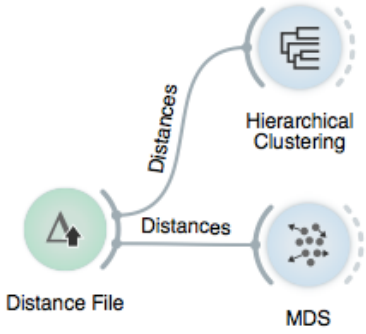

We have not constructed this map manually, of course. We used a widget called MDS, which stands for Multidimensional scaling.

It is actually a rather exact map of the US from the Australian perspective. You cannot get the orientation from a map of distances, but now we have a good impression about the relations between cities. It is certainly much better than clustering.

Remember the clustering of animals? Can we draw a map of animals?

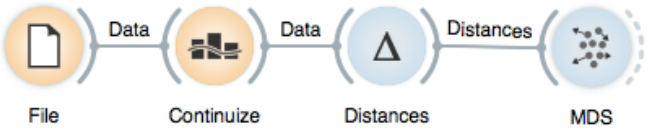

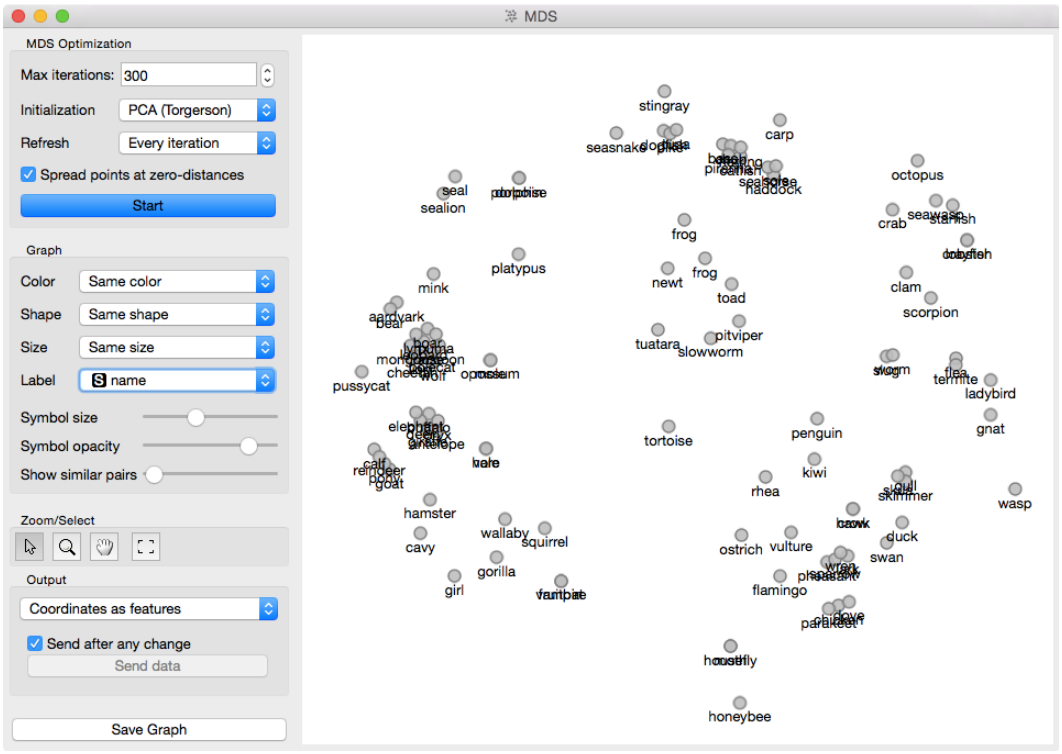

Does the map make any sense? Are similar animals together? Color the points by the types of animals and you should see.

The map of the US was accurate: one can put the points in a plane so that the distances correspond to actual distances between cities. For most data, this is usually impossible. What we get is a projection (a non-linear projection, if you care about mathematical finesses) of the data. You lose something, but you get a picture.

The MDS algorithm does not always find the optimal map. You may want to restart the MDS from random positions. Use the slider "Show similar pairs" to see whether the points that are placed together (or apart) actually belong together. In the above case, the honeybee belongs closer to the wasp, but could not fly there as in the process of optimization it bumped into the hostile region of flamingos and swans.

Lesson 35: Networks

Let us be crisp: given a set of data items, we will impose a distance threshold and claim that items are similar to each if their distance is below the threshold, and different to each other if the distance is above the threshold. We would like to construct a graph, with nodes (dots) in the plane presenting the data instances and edges (lines) between data instances that are similar. We just constructed a network! Here is the corresponding workflow:

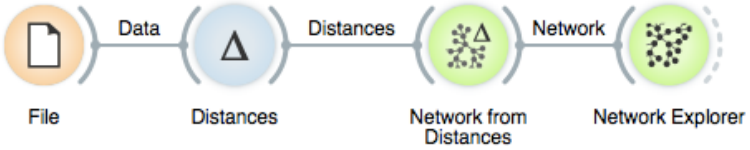

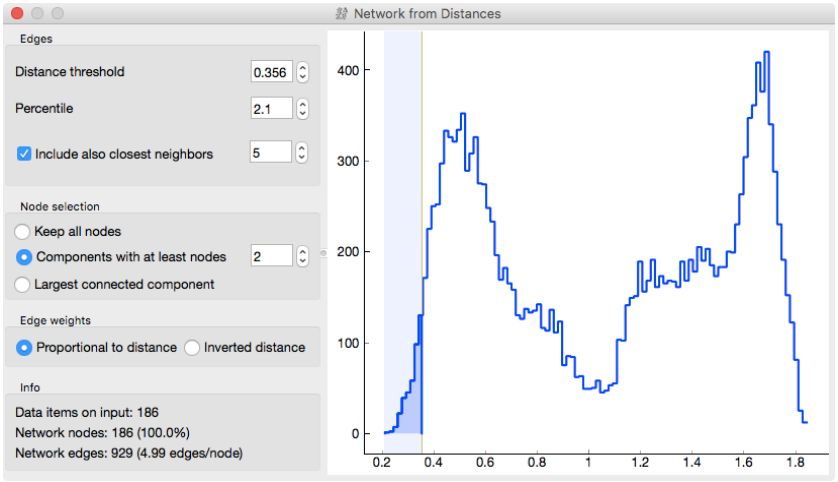

We load the data in the File widget, construct a pairwise-distance matrix in Distances, decide which pairs of data instances (nodes) will be connected by setting the distance threshold in the Network from Distances, and the visualize the network in the Network Explorer. We will use Euclidean distance, and set parameters of other widgets as shown

.

Network from Distances displays distribution of distances between data points. For the network, we have included only the edges in the shaded region, that is those that correspond to most similar pair of nodes. Also interesting for this data set is a bi-modal distribution of distances, perhaps indicating that the data includes a well-defined cluster structure: groups of mutually close points away from the other clusters.

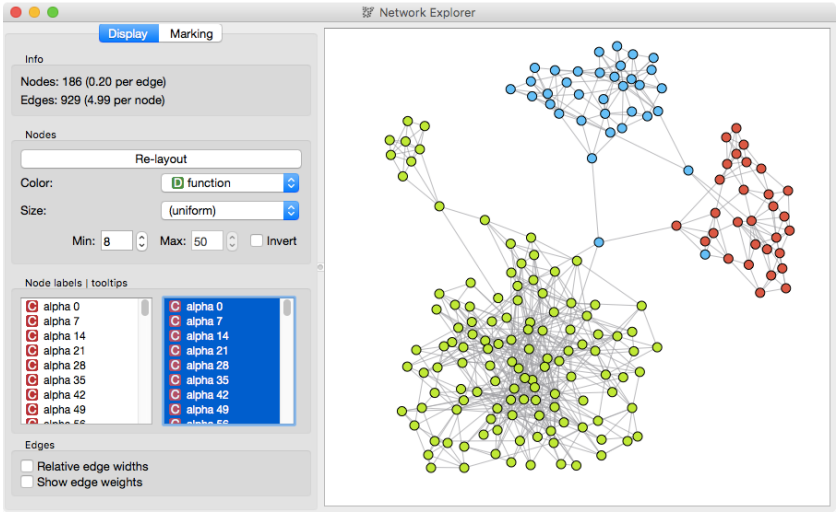

The network looks nice and tidy. But no big surprises here. For the brown-selected data set we already knew that the classes (gene functions) are easy to separate, but it is great to know that this can be done based on the distances from the full gene expression profiles. For our network, it was also helpful to set Network from Distances such that for each node (data instance) we requested to retain edges to at least 5 closest neighbors. Lower the value of this parameter to see how the network breaks down.

Networks, in this way, can complement MDS and PCA. Neighboring data points in the network have, by definition, similar feature profiles. But here the similarity ends: the nodes may be close to each other in rendering of the network, yet they may not be connected and their data profiles could be quite different from each other. Very unlike MDS, which strives to retain original distances in its visual depiction.

Network visualization and analysis is an important current research field, and the topic would require a course of its own (check out Coursera and other MOOCs). Before closing, though, just a remark that networks are often searched for nodes that are central, that may have many connections or even better — the nodes that most of the random network paths would pass through. This notion is called centrality. Orange includes Network Analysis widget that can compute degrees of centrality and many other network statistics. We can then visualize these by, say, the size of network nodes. To do this, we need to modify the workflow and make sure that Network Analysis passes the information on network Items to the Node Data input of Network Explorer.

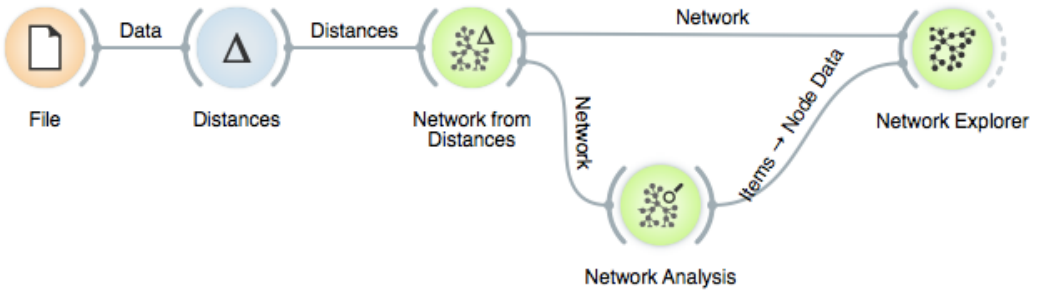

In the Network Analysis, let us choose the Degree centrality from the Node-level indices and make sure that we commit the computed value (Commit button). Now we can choose Degree centrality for the node size of Network Explorer.

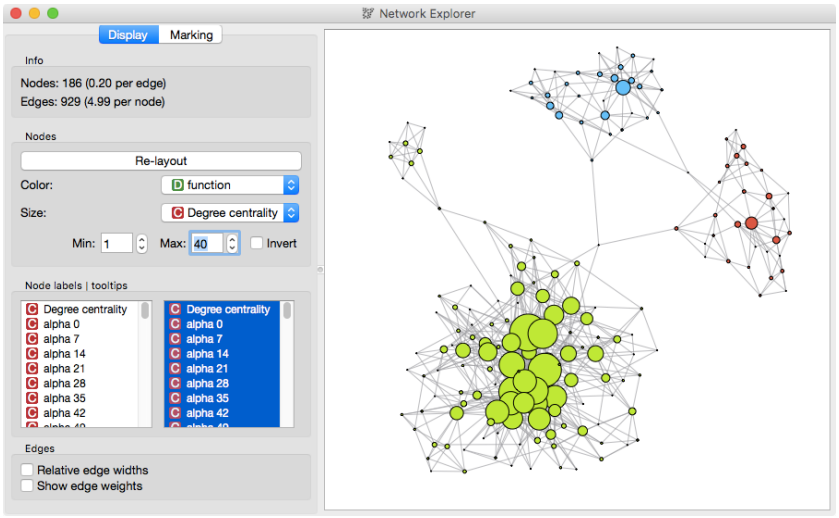

Looks nice and colorful. Central genes are, as expected, those at the core of the clusters. (Did we say clusters? Oh, we can do clustering on the networks. The most popular algorithms for that use label propagation. And there is even widget for that in Orange. But going on like this we will never finish with this course...).

One final word about networks. We have used data distances to construct the network. Such networks are in bioinformatics often referred to as correlation networks, as we can use correlation to measure the similarity of data objects (like genes, proteins, diseases, drugs, or anything else).

Sometimes the natural presentation of data is a network itself. Like protein interaction networks, or gene regulation networks. Or, outside biology, community networks, computer networks, networks of people on social platforms. Some examples of such networks come with the Orange's network add-on. Let us use the Network File widget, load a network of music artists lastfm.net, and visualize it such that we expose which artists are central to the network.

Here is the workflow:

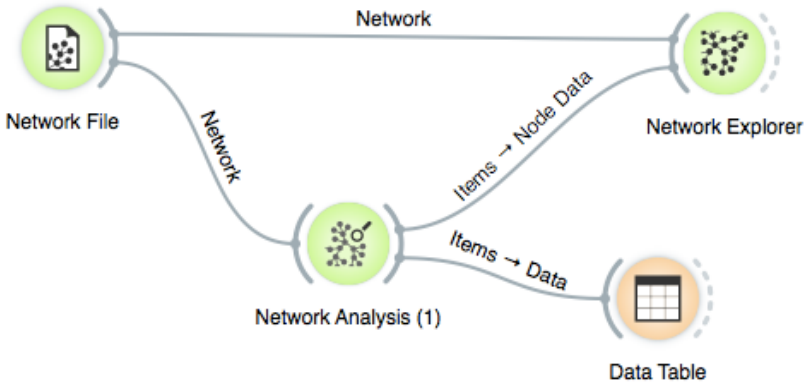

The network (despite being 10 years old — sorry, Beliebers) looks nice and structured, and there are well-known classics within central artists, like Rolling Stone and Duran Duran. There are few surprises as well; authors of this text still need to check out Taking Back Sunday.

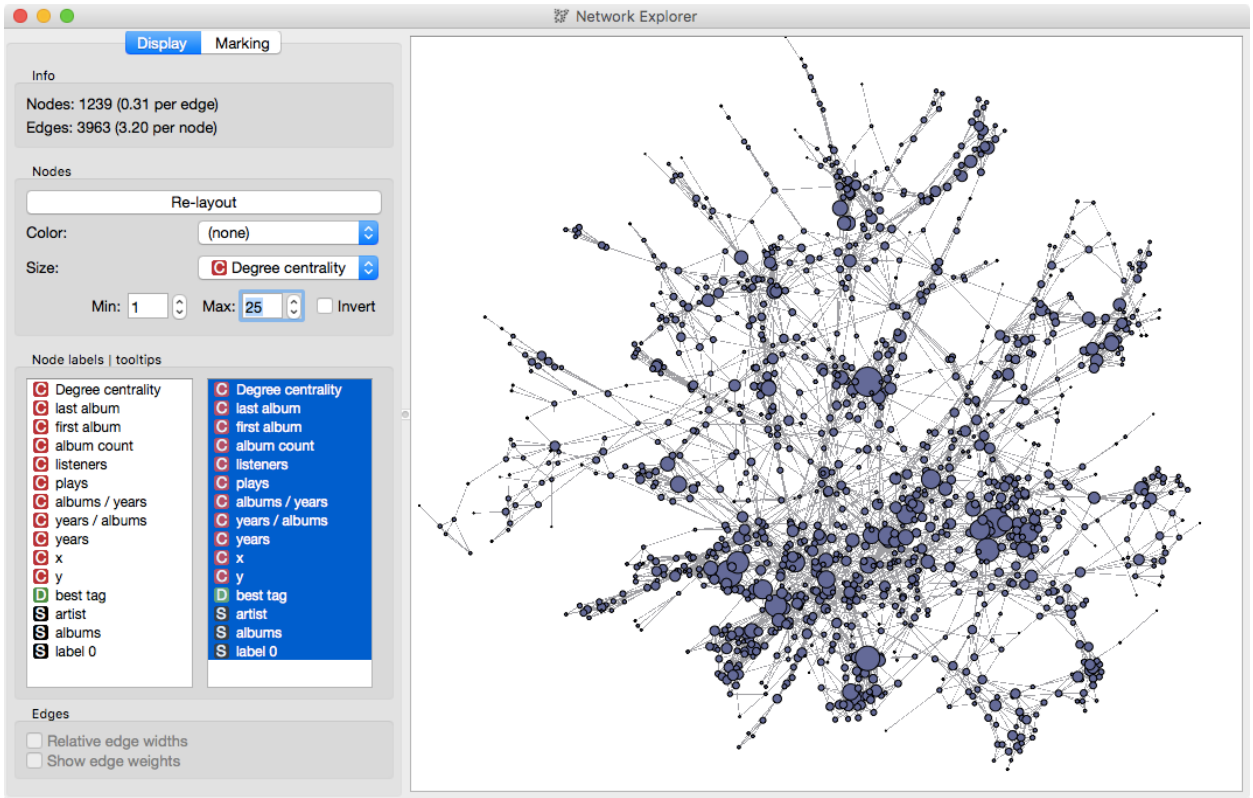

For the End

The Introduction to Data Mining course ends here. We covered quite some milage and hope we have taught you some crucial algorithms that should be on the stack of every data scientists. The goal was not to turn you into one, but to get you familiar with some basic techniques, tools and concepts. Data science is a huge field and it takes years of study and practice to master it. You may never become data scientist, but as an expert in biomedicine it should now be easier to talk and collaborate with statisticians and computer scientists. And for those who want to go ahed with data science, well, you now know where to start.