

To validate the NMI analyses, we build a null model by fixing number and size of the communities detected in each instance, and randomising the community labels assigned to each node. We compute the NMI matrix averaged over 100 randomizations, and compare it with the one presented in the main text. This allows us to verify whether our results can be attributed to randomness, or they represent an effect present in the data. The results, shown below, bear no resemblance to those in the main text. Also, in all null models we observe that the patterns characterizing the NMI matrices of the original data are absent. This shows that it is very highly unlikely that the structures detected arose due to random fluctuations.

