Block 2: Hidden Markov models

Below are exercises accompanying the lecture on hidden Markov models (HMMs). In this session, you will implement a GHMM from scratch and use it to compute the probability of observed token sequences – the foundational machinery you will build on in later blocks. Spend ~30 minutes on these exercises; start with the core set, and progress to the extension if you have time.

Core exercises.

1. Using your programming language of choice, implement a function that takes in:
    a. The transition matrices of a GHMM (a three-tensor)
    b. The initial vector
    c. A sequence of tokens

    and outputs the probability of observing that particular sequence of tokens.
2. Test your function on the random-random-XOR (RRXOR) process when initialised in the state $\eta^{(\varnothing)}$ = [1 0 0 0 0]. Make sure you assign zero probability to XOR violations.
3. Convince yourself that that the definition of a GHMM allows one to interpret $\eta^{(\varnothing)}T^{(w1)}T^{(w2)}...T^{(wn)}\phi$ as the probability of emitting the sequence w1 w2 … wn.
4. Derive the transition matrices for the zero-random-random process.

Extension exercises.

1. The RRXOR process can be viewed as a particular instance of a more general process called *random-random-Mod p* (RRModp) where p = 2,3,..., and RRXOR = RRMod2.
    a. What are the transition matrices for the RRMod3 process?
    b. What are the transition matrices for the RRModp process?
    c. Do you notice anything about the matrices sitting in the blocks?
2. The RRModp process can be similarly viewed as a particular instance of a more general process that computes the product of two *group* elements (for RRModp the cyclic group is the relevant one). Identify another group whose multiplication can be expressed as an HMM.