

Location Prediction



Consiglio Nazionale
delle Ricerche

Content of this lesson

- Mobility prediction: a taxonomy
- Next location prediction
 - (Hidden) Markov Model-based
 - (Frequent) Pattern-based
 - Deep Learning-based

Mobility Prediction

Target of prediction

- Individual targets
 - Trajectory
 - The next location
 - All the trip
 - Destination of trip
 - Events
 - E.g. Crashes
 - All above + time of movement
- Collective targets
 - Aggregate Flows
 - OD matrix
 - Crowd density
 - Events
 - E.g. Crashes

Perspective

- Continuous movement
 - Points expressed as latitude & longitude
 - Prediction means reconstructing the precise points & movement
- Discrete space
 - Points become “areas”
 - Mobile phone cells
 - POIs
 - Voronoi cells
 - Prediction means predicting a cell ID or a sequence of IDs

Prediction Tools

Prediction approaches


- Most typical tools adopted in mobility prediction
 - (Hidden) Markov Models
 - Pattern-based
 - Neural-Networks

Markov Models

- Standard means to model sequential stochastic processes
- Assumption
 - The probability of actual event depends only on the previous event

$$P(s_i | \langle s_1, \dots, s_{i-1} \rangle) = P(s_i | s_{i-1}) \quad \rightarrow \text{all the previous states of the system are irrelevant}$$

- Consequence
 - The model can be represented as a simple “transition matrix” $P_{i,j}$
 - Example:

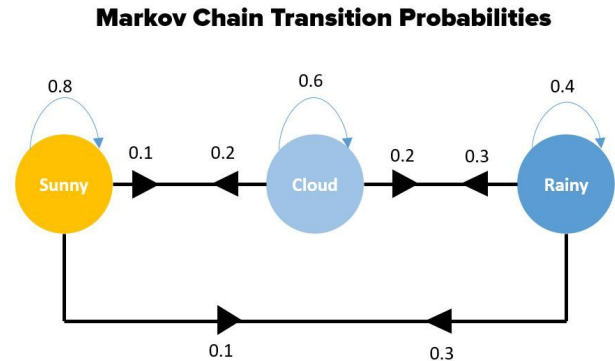


$$P = \begin{matrix} & \begin{matrix} A & B & C \end{matrix} \\ \begin{matrix} A \\ B \\ C \end{matrix} & \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & i \end{pmatrix} \end{matrix}$$

Markov Models


- More formally, we have
 - Set of states:
 $Q = \{q_1, q_2, \dots, q_n\}$
 - Vector of prior probabilities (probabilities of occurrence of each state):
 $\Pi = \{\pi_1, \pi_2, \dots, \pi_n\}$, where $\pi_i = P(S_0 = q_i)$
 - Matrix of transition probabilities
 $P = \{a_{ij}\}$, i, j in $[1..n]$, with $a_{ij} = P(q_j | q_i)$

- **Markov Chains (MC)**
 - The states are directly the observed values
 - Learning Π and P is straightforward: just count!





Markov Models

- **Hidden Markov Models (HMM)**
 - The states are not visible!
 - We can only see some “observables”
 - E.g. states = activity
= { working, leisure }
observables = places visited
= { home, cafeteria, gym, ... }
- Observables depend only on current state
 - Emission probabilities
 - $P(O_j | S_i)$



Observable input sequence



Markov Models

- Three basic problems with HMM
 - **Likelihood:**
 - Given an HMM $\lambda = (A, B)$ and an observation sequence O , determine the likelihood $P(O|\lambda)$
 - **Decoding:**
 - Given an observation sequence O and an HMM $\lambda = (A, B)$, discover the best hidden state sequence Q
 - **Learning:**
 - Given an observation sequence O and the set of states in the HMM, learn the HMM parameters A and B

Markov Models

- Three basic problems with HMM

- **Likelihood:**

- Given an HMM $\lambda = (A, B)$ and an observation sequence O , determine the likelihood $P(O|\lambda)$



- **Decoding:**

- Given an observation sequence O and an HMM $\lambda = (A, B)$, discover the best hidden state sequence Q

- **Learning:**

- Given an observation sequence O and the set of states in the HMM, learn the HMM parameters A and B

Forward Algorithm

- Dynamic programming
- Virtually tries all sequences of states & aggregate probabilities

$$P(\text{😊 😞 😞 😊 😞 😊 😊 😞} \rightarrow) = ?$$

Markov Models

- Three basic problems with HMM

- Likelihood:**

- Given an HMM $\lambda = (A, B)$ and an observation sequence O , determine the likelihood $P(O|\lambda)$

- Decoding:**



- Given an observation sequence O and an HMM $\lambda = (A, B)$, discover the best hidden state sequence Q

- Learning:**

- Given an observation sequence O and the set of states in the HMM, learn the HMM parameters A and B

Viterbi Algorithm

- Dynamic programming
- Quite similar to the Forward Algorithm

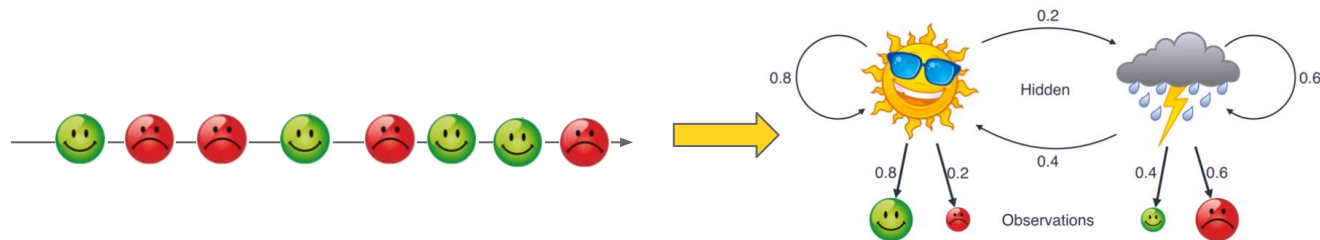


Example on words:
 best("the boat sank") = det → noun → verb

Markov Models

- Three basic problems with HMM
 - **Likelihood:**
 - Given an HMM $\lambda = (A, B)$ and an observation sequence O , determine the likelihood $P(O|\lambda)$
 - **Decoding:**
 - Given an observation sequence O and an HMM $\lambda = (A, B)$, discover the best hidden state sequence Q
 - **Learning:**
 - Given an observation sequence O and the set of states in the HMM, learn the HMM parameters A and B


- Forward-Backward Algorithm
a.k.a. Baum-Welch Algorithm
- Similar to EM or K-means
 - Iterative refinement method



Pattern-based prediction

- Tabular representation of trajectories
- Each pattern can become
 - a Boolean (the pattern occurs/does not occur)
 - a numerical value (how strongly the pattern occurs)
- Standard ML tools can be applied, e.g. decision trees

	A→B→C	A→B→D	B→D→E	B→C→H
Trajectory-1	1	0	0	1
Trajectory-2	0	0	1	0




Pattern-based prediction

- Patterns can be extracted through various criteria
 - Based on frequency
 - General frequent pattern / clustering methods can be applied
 - Frequent patterns expected to express significant features
 - Based on discrimination power
 - Requires ad hoc solutions
 - Might find infrequent yet useful patterns

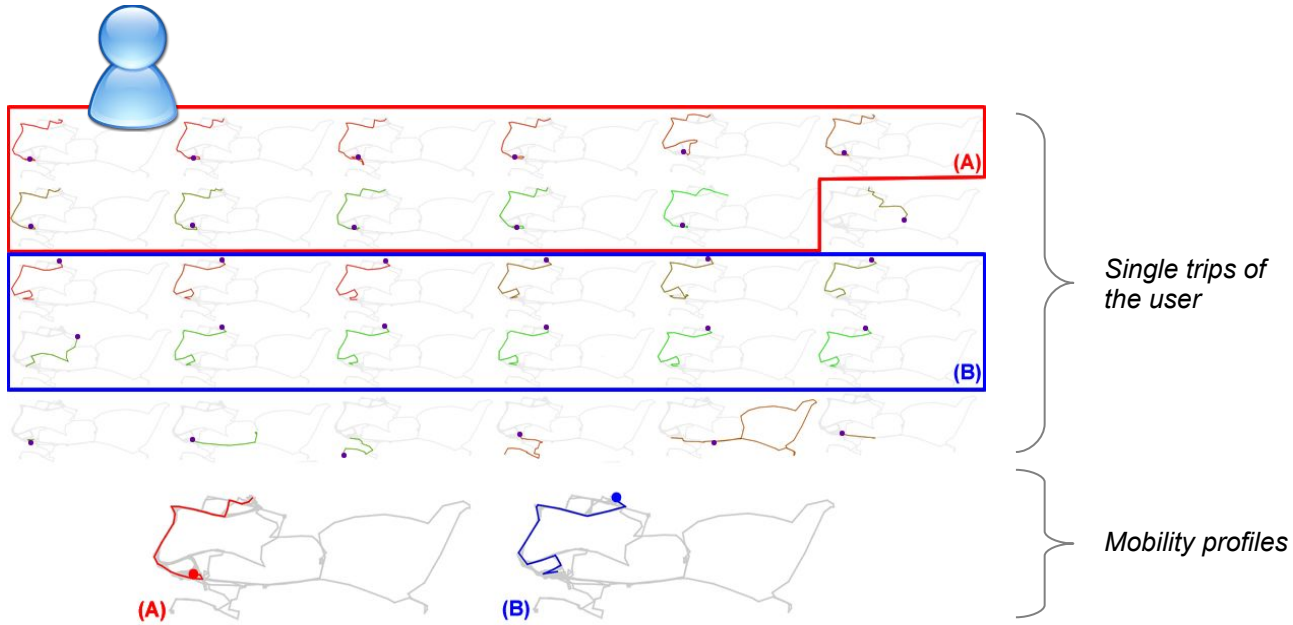
Pattern-based prediction

Mobility profiles as “frequent prediction patterns”



Pattern-based prediction


Mobility profiles as “frequent prediction patterns”



Pattern-based prediction

Mobility profiles as “frequent prediction patterns”

- Step 1: perform “partial match” of a trajectory with a set of systematic trips
- Step 2: check that the best match is similar enough
- Step 3: use the “rest” of the systematic trip as continuation of the trip



Pattern-based prediction

Mobility profiles as “frequent prediction patterns”

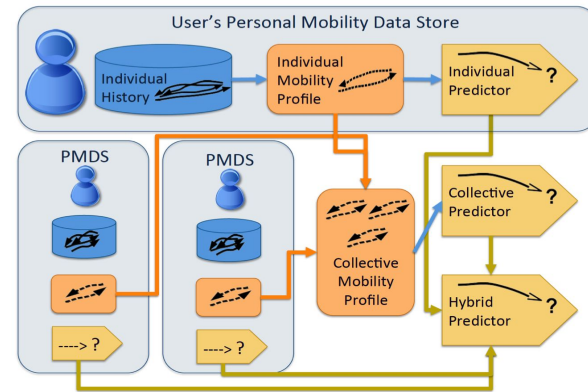
Which systematic trips to use?

- Priority: those of the same user/device/vehicle
- If fails: those of the whole population analyzed

1. Check my personal systematic mobility



2. Check others' systematic mobility




Deep Learning for trajectory prediction

- Three main architectures (+ hundreds of variants)
 - RNN: Recurrent Neural Network / LSTM: Long-Short Term Memory
 - CNN: Convolutional Neural Network
 - GAN: Generative Adversarial Network

RNN


- NN where the hidden layer has a “loop” connection
- The link works as extra input for the next training data instance
 - Makes the model apt for sequential data
- Notice: “A” is always the same




- E.g. $\langle x_0, \dots, x_t \rangle =$ input trajectory, $\langle h_0, \dots, h_t \rangle =$ output trajectory

RNN - LSTM

- RNN has issues with long-term relations between input and output



- Solution: pass a state information that is incrementally updated




(a) RNN

(b) LSTM


RNN - LSTM - Encoder-Decoder schema

- **Encoder step:** during input elaboration the output (W) is a latent representation of the whole sequence
- **Decoder step:** W becomes the input of LSTM, which is fired till generating a special output (stop)




RNN - LSTM - Encoder-Decoder schema

- **Encoder step:** during input elaboration the output (W) is a latent representation of the whole sequence
- **Decoder step:** W becomes the input of LSTM, which is fired till generating a special output (stop)



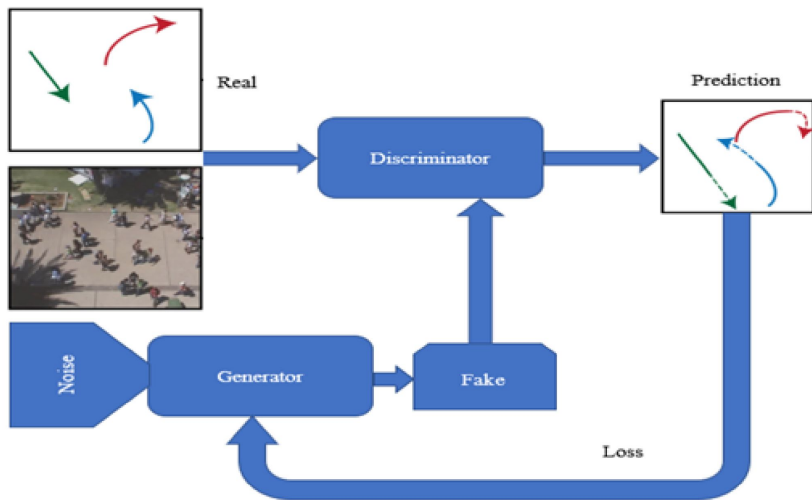
CNN for movement prediction

- Typically used when images are involved, e.g. camera data
- Usually in conjunction with LSTM or other sequence-based models
 - CNN captures spatial relations & identifies objects / features
 - LSTM captures temporal relation & movement



GAN

- Trains 2 models at the same time:
 - Generator: one that generates fake objects
 - Discriminator: one that can distinguish real vs fake objects
- The generator can be seeded with a (representation of a) partial trajectory
- Both G and D can be any suitable model, e.g. LSTM



INTERVALLO

**Baseball: Do Fielders Know Where to Go to Catch
the Ball or Only How to Get There?**

or

**When Psychologists (With Lots of Spare Time)
Meet Sport Analytics**

Peter McLeod & Zoltan Dienes (1996).
J. of Experimental Psychology: Human Perception and Performance.
Vol. 22, No. 3, 531-543.

INTERVALLO

Baseball: Do Fielders Know Where to Go to Catch the Ball...?

Experimental setup





Figure 1. The trajectories of three balls projected at 45° and a velocity (v) of 22.3, 24.0, and 25.7 m/s, respectively, toward a fielder 45 m away. They experienced a deceleration due to aerodynamic drag proportional to v^2 . The constant of proportionality was 0.007 m^{-1} , a value typical of objects such as cricket balls (Daish, 1972).





INTERVALLO

Baseball: Do Fielders Know Where to Go to Catch the Ball...?

Discovery

- Fielders adjust movement in order to keep $d^2(\tan \alpha)/dt^2 = 0$
- α = angle of gaze




INTERVALLO

Baseball: Do Fielders Know Where to Go to Catch the Ball...?


Discovery

- Fielders adjust movement in order to keep $d^2(\tan \alpha)/dt^2 = 0$
- α = angle of gaze



Remark

- Simple solutions like keeping α constant would not work



INTERVALLO

Baseball: Do Fielders Know Where to Go to Catch the Ball...?

Question 1: do fielders predict where the ball will fall?

- TL;DR: most likely no
 - They never reach the spot before the ball
 - They appear to dynamically adjust the movement to ball position

Question 2: do people really evaluate/compute $d^2(\tan \alpha)/dt^2$?

- Answer: most likely no

Homeworks

to be delivered by Thursday, November 9th 2022



Homework 8.1

HMM-based trajectory generation

Select a sample of taxi data in SF and train a HMM (for instance discretize trips to sequences of cells and then use CategoricalHMM). Take 5 random trips T (not used in the training), cut them in two equal length parts $T1$ and $T2$, use the HMM to understand what is the most likely final state of $T1$ and then randomly generate a possible continuation $T3$ of the trip. Finally, compare $T2$ and $T3$.

- Submit a (well commented) python notebook

Homework 8.2

Pattern-based prediction

Randomly select a set MP of 100 trips in SF taxi data (and pretend they are our representative mobility profiles), and another set TS of 5 trips (our test set). (1) Cut all trips T in $MP \cup TS$ in two equal length parts T_1 and T_2 ; (2) for each T in TS find the 3 trips T^ in MP that minimize the $dist(T_1, T_1^*)$, where $dist()$ is a trajectory distance of your choice (e.g. Hausdorff); (3) compare each T_2 with the three corresponding T_2^* “predicted”.*

- Submit a (well commented) python notebook