













Image captioning



"girl in pink dress is jumping in

air."









'construction worker in orange safety vest is working on road"



"black and white dog jumps over bar."









'young girl in pink shirt is swinging on swing."





'man in blue wetsuit is surfing on wave."





Recurrence is required to capture the representational dynamics of the human visual system

Tim C. Kietzmann^{a,b,1}, Courtney J. Spoerer^a, Lynn K. A. Sörensen^c, Radoslaw M. Cichy^d, Olaf Hauk^a, and Nikolaus Kriegeskorte^e

^aMRC Cognition and Brain Sciences Unit, University of Cambridge, Cambridge CB2 7EF, United Kingdom; ^bDonders Institute for Brain, Cognition and Behaviour, Radboud University, 6525 HR Nijmegen, The Netherlands; ^cDepartment of Psychology, University of Amsterdam, 1018 WD Amsterdam, The Netherlands; ^dDepartment of Education and Psychology, Freie Universität Berlin, 14195 Berlin, Germany; and ^eDepartment of Psychology, Columbia University, New York, NY 10027

PNAS 2019





Kietzmann 2019





Kietzmann 2019

Outline

- Conceptual overview or RNNs
 - •FFN vs RNN
 - Flavours of RNNs
 - Training RNNs
 - Real-world RNN
- Jupyter Notebook
 - Build simple RNN (Keras)
 - Analyse some properties

Multi-layer perceptron



Why they are powerful interpolators



Why they are powerful interpolators



Why they are powerful interpolators



Building piecewise local linear models (even more flexible when stacking multiple layers)







Don't use this for extrapolation!





Recurrent (more common)

(only lateral recurrence)



How do we know h_1 ?

Recurrent (temporal)



 $h_1(t) = W_1^*x(t) + V_1^*h_1(t-1)$

 $!! W_1$ and V_1 don't depend on time !!

Memories

Output depends on all past input through multiple complex transformations



Flavours of RNNs

- Cardinality of input-output
- How the hidden layers are passed forward







 $h(t) = nonlin[W^*x(t) + V^*h(t-1)]$

$$h(t) = nonlin[W^*x(t) + V^*h(t-1)]$$



"Gated" RNN

(examples)

 $h(t) = nonlin[W^*x(t) + V^*h(t-1)]$



GRU



"Gated" RNN

(examples)

 $h(t) = nonlin[W^*x(t) + V^*h(t-1)]$







GRU

$h(t) = nonlin[W^*x(t) + V^*h(t-1)]$





"Gated" RNN







(in practice, smooth sigmoid used instead)



LSTM (approx)



(output depends on input, state, and memory) (memory explicitly changed by input and state)



Training FFN



Training RNNs

Back-propagation through time



A "real-life" RNN



Google Translate

A "real-life" RNN



Google Translate

Word -> vector of numbers

(word "embedding", word piece model)

Word -> vector of numbers

(word "embedding", word piece model)

Sentence -> trajectory

Word -> vector of numbers

(word "embedding", word piece model)

Sentence -> trajectory

Two separate networks (encoding/decoding)

Word -> vector of numbers

(word "embedding", word piece model)

Sentence -> trajectory

Two separate networks (encoding/decoding)



English

Word -> vector of numbers

(word "embedding", word piece model)

Sentence -> trajectory

Two separate networks (encoding/decoding)





https://arxiv.org/pdf/1609.08144.pdf

Time to do hands-on playing with RNNs