Logos: Logic in Computer Science

Michel de Rougemont

University Paris II & IRIF CNRS

Plan

1. Algorithmic Game theory
   - Simultaneous games
     - 0-sum games
     - Cooperative games, Nash equilibria
     - Auctions
   - Sequential games
     - Markov Decion Processes: MDPs
     - 2 players
2. LLM's
   - Tokens
   - Language model
   - RLHF

1. Algorithmic game theory

Cooperative games

http://mdr.free.fr/td/PL14.ppt
http://dup2.free.fr/index.php?n=Main.BD10

- 0-sum games: linear programming, fictitious player
- General games: Nash equilibria, Hemke-Lowson
- Auctions: Vickrey auctions
- Adwords

Algorithmic game theory: MDPs

MDPs: Markov Decision processes, Sequential Games

http://dup2.free.fr/index.php?n=Main.BD10bis

- Games against nature, 1 + ½ players
- 2 players: chess, go
- N-players (robots, autonomous vehicles)

Reinforcement Learning

1. Search for optimal strategies in MDP: Q-learning
   - Value iteration $V^*(s)=\max_{pi} V^{pi}(s)$ Bellman's equation
   - Policy iteration $pi^*(s)= \text{argmax}_a \text{Sum}_{s'} P_{s,a}(s')V^*(s')$
   - Q-learning: $V(s)=\max_a Q(s,a)$

     $a=\arg\max_a Q(s,a)$

2. Learning with experts:

   http://dup2.free.fr/index.php?n=Main.BD10a

2. LLMs

1. Tokens and syntax

1. Language model

1. RLHF: Reinforcement Leaning with Human Feedback

https://stanford-cs324.github.io/winter2022/

https://stanford-cs324.github.io/winter2023/

2.1 Tokens

Morphologie: *re-construc-tion*

Entités Nommées: *First National Bank of Chicago*

Formes: *il y a , tout à coup*

Mots sémantiques: *machine à laver*

Lemmatisation: racine, forme canonique

BPE: Byte Pair Encoding

Tokens: BPE

Phrase d'origine : "abracadabra »

Initialisation :

"a", "b", "r", "a", "c", "a", "d", "a", "b", "r", "a"
Étape 1 : Fusion des paires les plus fréquentes :

"ab", "r", "a", "c", "a", "d", "a", "b", "r", "a"
Étape 2 : Nouvelle fusion :

"abr", "a", "c", "a", "d", "a", "b", "r", "a"
Étape 3 : Nouvelle fusion :

"abra", "c", "a", "d", "a", "b", "r", "a"
Étape 4 : Dernière fusion :
"abrac", "a", "d", "a", "b", "r", "a"

Structure des tokens

"John's cats are chasing butterflies in the garden."

Arbre syntaxique :

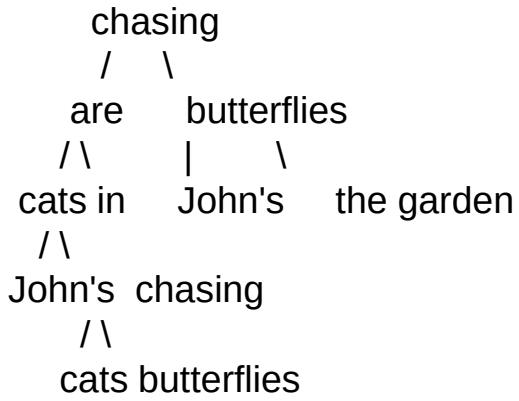

Lemmatisation :

"John's" devient "John" (suppression de la possession et de l'apostrophe).
"cats" reste "cat" (réduction à la forme de base).
"are" reste "are" (verbe auxiliaire).
"chasing" devient "chase" (réduction à la forme de base).
"butterflies" devient "butterfly" (réduction à la forme de base).
"in" reste "in" (préposition).
"the" reste "the" (article défini).
"garden" reste "garden" (nom).

2.2 Language Model

Représentation de la distribution

$Ustat_k$ $(w_1 \ldots w_k)$ et $Next_k(w_1 \ldots w_k)$

Forme compressée, grâce à 3 Embeddings (Word2vec) Query, Key, Value et à l'attention:

John's cats are chasing ?

v_1 v_2 v_3 v_4

Matrice: $M(i,j)= v_i * v_j$

Language Model

Attention Scores:

Butterflies $\text{Sum}_i \; Q_{\text{Butterflies}} * K_i = 5$

Mice: $\text{Sum}_i \; Q_{\text{Mice}} * K_i = 3$

Softmax (3,5): e^5/e^3+e^5 , e^3/e^3+e^5

John's cats are chasing ?

v_1 v_2 v_3 v_4

3. RLHF

"Un coucher de soleil éclatant sur une plage déserte."

Poids : 0.2

"Un paysage urbain avec des gratte-ciels illuminés la nuit. Un paysage urbain avec des gratte-ciels illuminés la nuit.

Poids : 0.8

Supervised fine-tuning 2021,22,23…

Other methods:

1. Best-of-n
2. Expert iteration
3. Conditional tokens
4. Contrastive-based methods

Market for CHatBots

1. ChatGPT
2. Bard
3. Llama
4. Huggingface

Less known:

You: www.you.com

New: Mistral

Citer ses sources
Interface BD

Conclusion

1. Algorithmic Game theory
   - Cooperative games
   - Mechanisms and Equilibria
2. LLM's
   - Tokens
   - Language Model
   - RLHF

Testers and Correctors

Distance: Edit, Edit with moves,….

ε - close

1. Words: regular expressions, U=({1,2,…n},<,A,B)

   w=abababcababacacacab

   r= (ab)*(ac)*

   Corrector: w1=abababababacac

   w2=abababababacacacac

   w3=abababababacac

   w4=abababababacacacac

Q(x): A(x) and B(x+1) then CQA(A)={1,3,5,7}