Natural Language Processing

… from a Translational Data Science Perspective in Dutch Healthcare

Name speaker: *Prof. dr. Marco Spruit (LUMC/LIACS)*

Talk: *Research Facility Data Analytics workshop on Free Text Analysis, 24 May 2022*

ZyLAB

1993 1995 1997 2003 2007

2020

Leiden University Medical Center

liacs Leiden Institute of Advanced Computer Science

Universiteit Leiden

Leiden University Campus The Hague

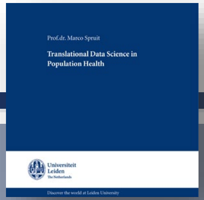

APRIL FOOLS' DAY

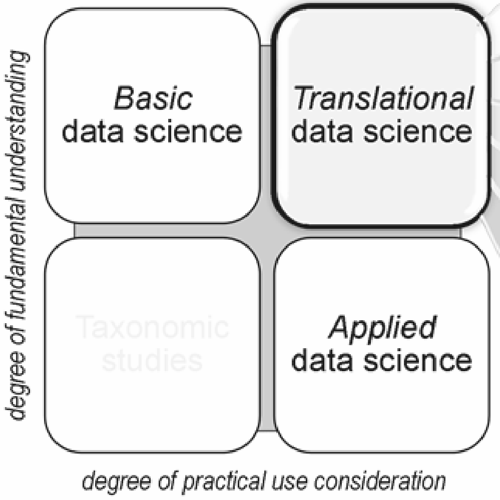

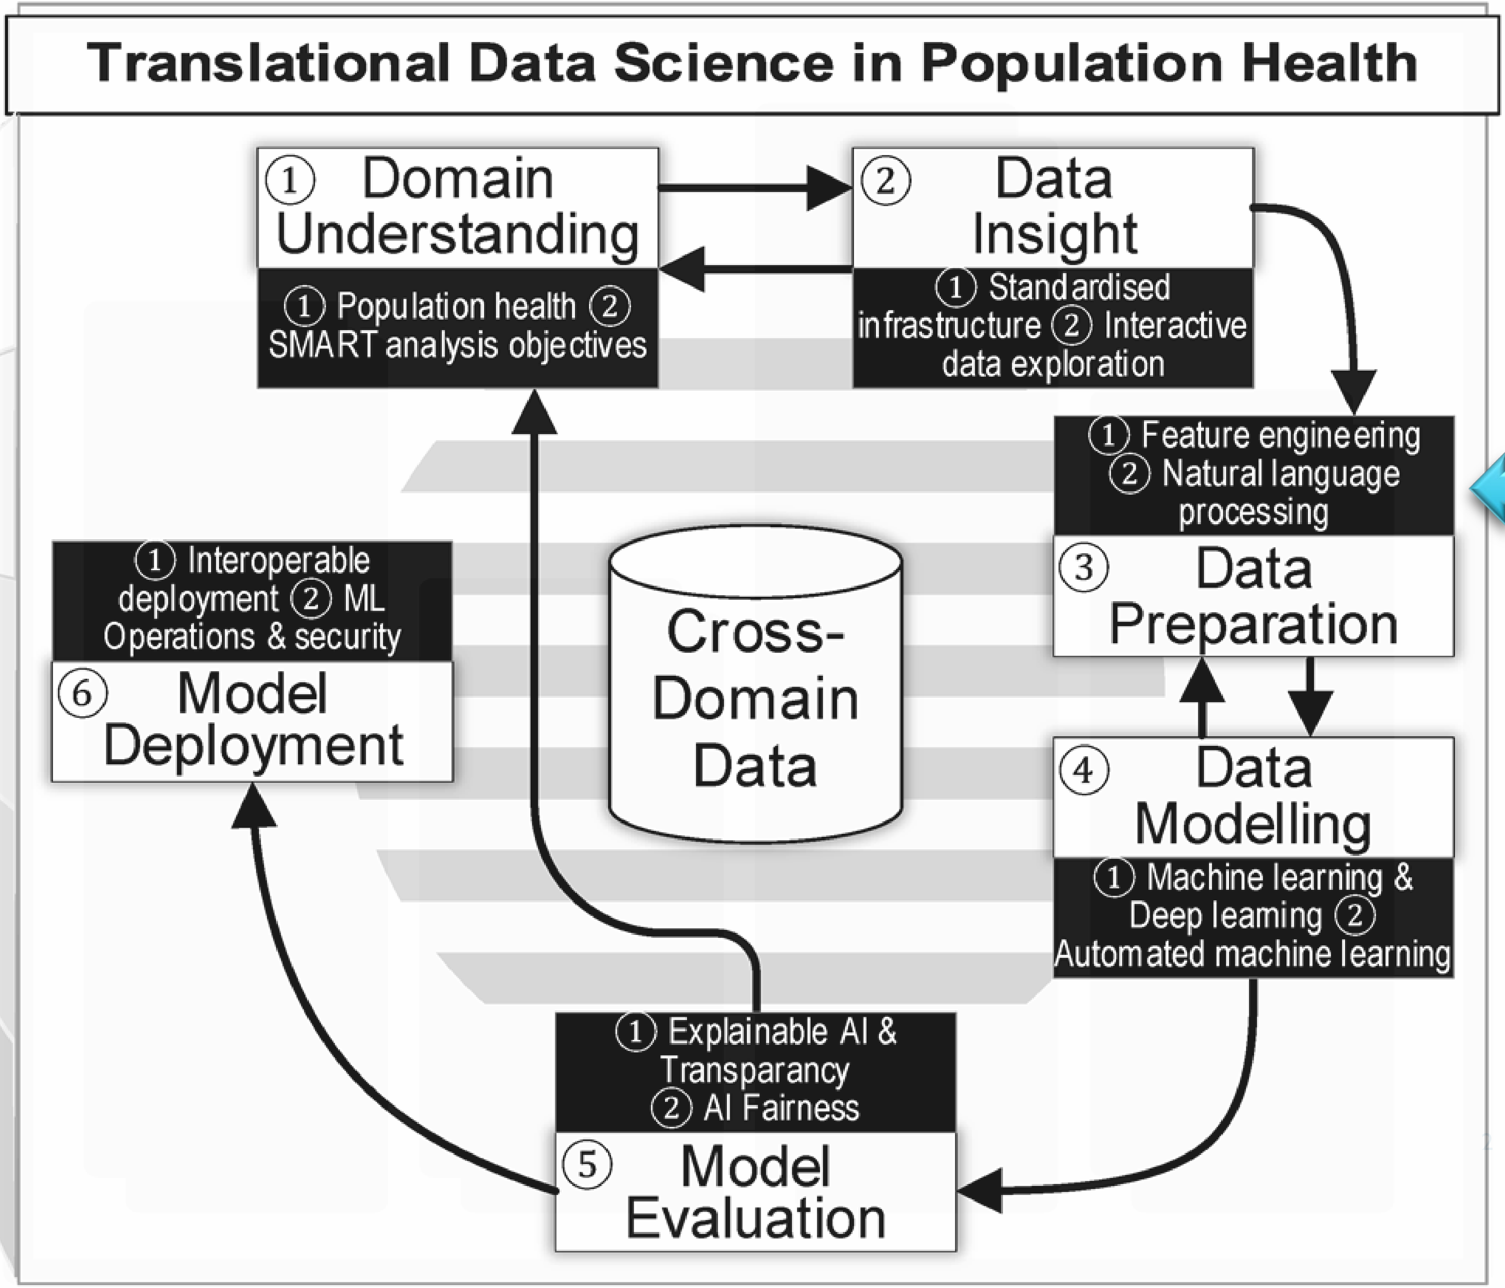

Accompanying slide to inaugural lecture of Prof. Dr. Marco Spruit on 1 April 2022 at Leiden University

Agenda: Setting the NLP Scene with Examples in Healthcare

"Traditional" NLP 101 → Symbolic NLP

- De-identification & ADR extraction

"Modern" NLP → Probabilistic NLP

- Embeddings for Classification (VRA)

"Effective" NLP → Hybrid NLP

- Combining traditional and Modern approaches (e.g. ADRIN)

NLP is a Multidisciplinary Wicked Problem within AI

- Computers are confused by (human) language
  - Specific techniques are needed
  - NLP draws on research in
    - Linguistics,
    - Theoretical Computer Science,
    - Mathematics,
    - Statistics,
    - Artificial Intelligence,
    - Psychology,
    - etc.

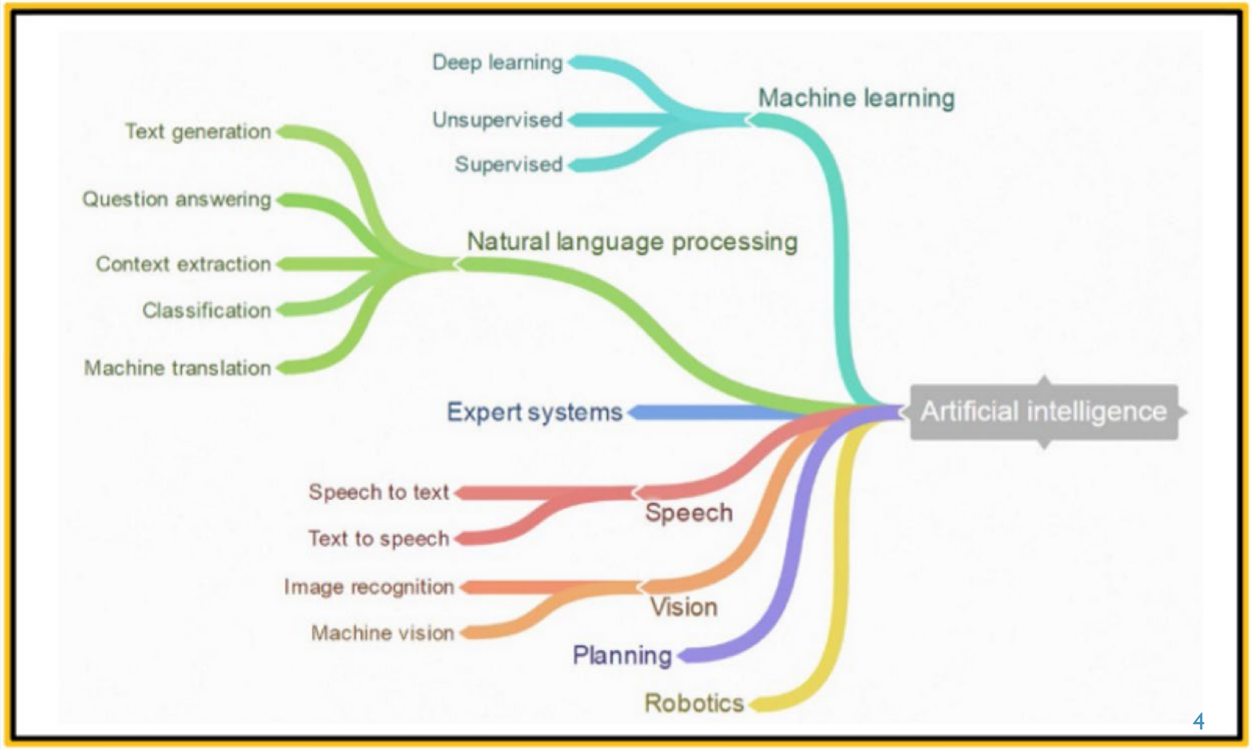

"TRADITIONAL" VERSUS "MODERN" NLP APPROACH

Derivation rules ('50s (Chomsky) →)

A grammar and a parse tree for "the giraffe dreams".

```
s    → np vp

np   → det n
vp   → tv np
     → iv

det  → the
     → a
     → an

n    → giraffe
     → apple

iv   → dreams
tv   → eats
     → dreams
```

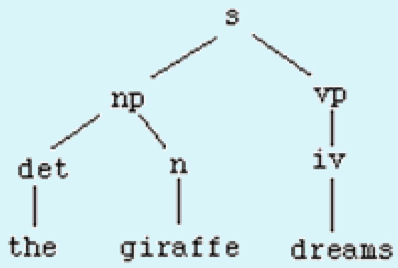

RegEx `^([0|\+[0-9]{1,5})?([0-9]{10})$`

Word embeddings (2013 (Google) →)

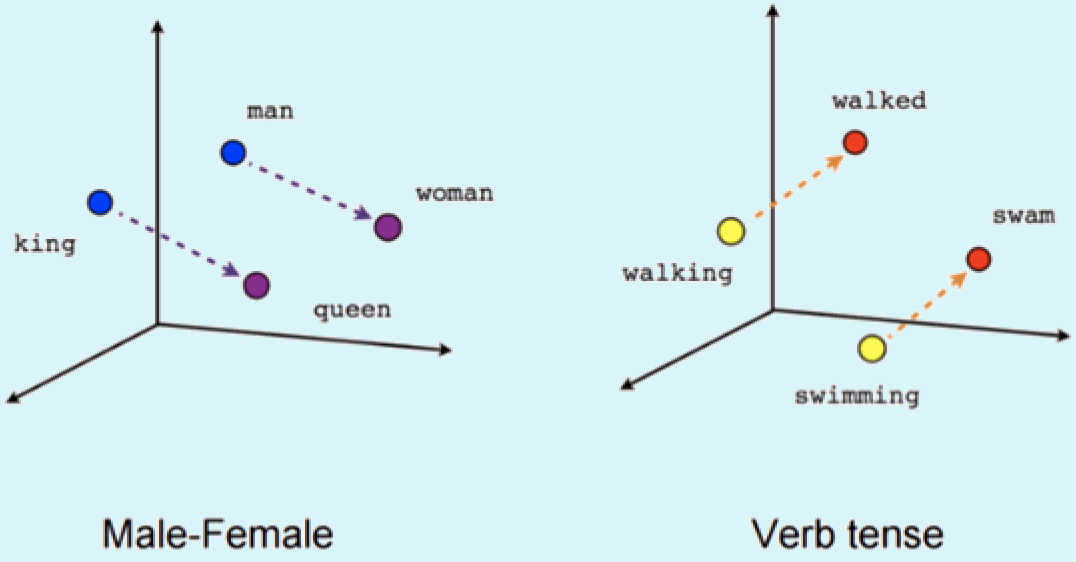

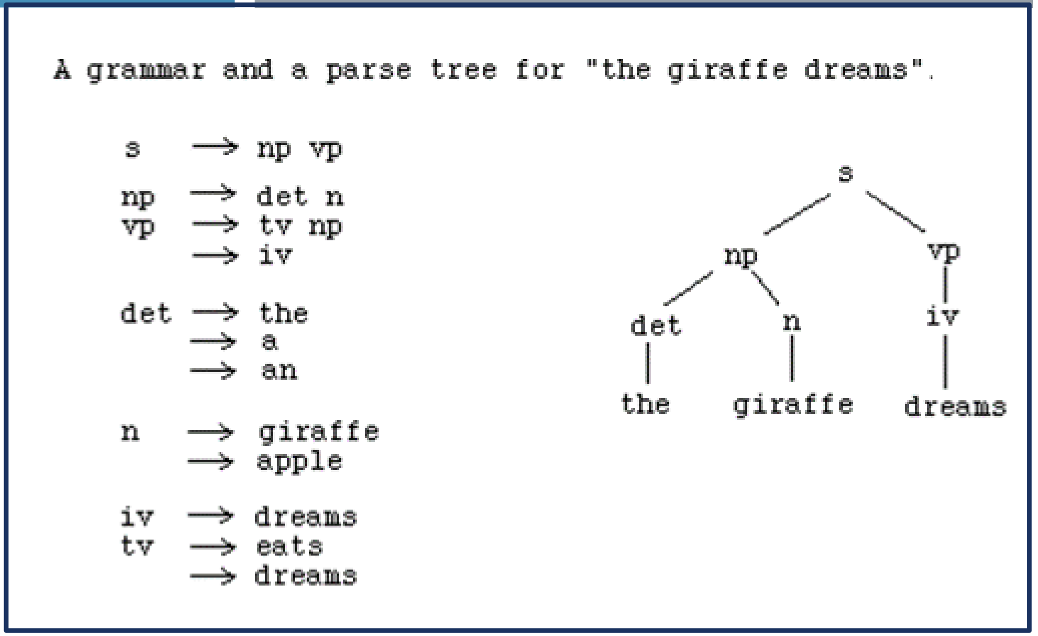

“Traditional” NLP 101 → Symbolic NLP

DE-IDENTIFICATION & INFORMATION EXTRACTION

EXAMPLE #1: DE-IDENTIFICATION IN DUTCH

Menger,V., Scheepers,F., Wijk,L. van, & Spruit,M. (2018). DEDUCE: A pattern matching method for automatic de-identification of Dutch medical text. *Telematics and Informatics, 35*(4), Patient Centric Healthcare, 727–736. [JIF: 6.182] [pdf] [online]

DEDUCE:

- De-identification of Dutch medical text
  - Information extraction of Protected Health Information (PHI) categories
- *Method:* Combines
  - Lookup tables, decision rules, and fuzzy string matching
- https://tdslab.nl/deduce
- `>>> pip install deduce`

[Legend: Patient Persoon Locatie Instelling Datum Leeftijd Patientnummer Telefoonnummer Url]

Annotated text

Intakegesprek met Jan Jansen (e:j.g.jsnen_1966@email.com, t:0612345678, patnr:1243567). Het betreft een 51-jarige man die van 14 maart t/m 31 juli op de polikliniek van het umcu zal worden behandeld i.v.m. somberheidsklachten. Patient is woonachtig aan de Voorstraat 45b in Utrecht en zal hier onder behandeling komen te staan van Peter de Visser.

De-identified text

Intakegesprek met <PATIENT> (e:<URL-1>, t:<TELEFOONNUMMER-1>, patnr:<PATIENTNUMMER-1>). Het betreft een <LEEFTIJD-1>-jarige man die van <DATUM-1> t/m <DATUM-2> op de polikliniek van het <INSTELLING-1> zal worden behandeld i.v.m. somberheidsklachten. Patient is woonachtig aan de <LOCATIE-1> in <LOCATIE-2> en zal hier onder behandeling komen te staan van <PERSOON-1>.

EXAMPLE #1: DE-IDENTIFICATION - METHOD

Menger,V., Scheepers,F., Wijk,L. van, & Spruit,M. (2018). DEDUCE: A pattern matching method for automatic de-identification of Dutch medical text. *Telematics and Informatics, 35*(4), Patient Centric Healthcare, 727–736. [JIF: 6.182] [pdf] [online]

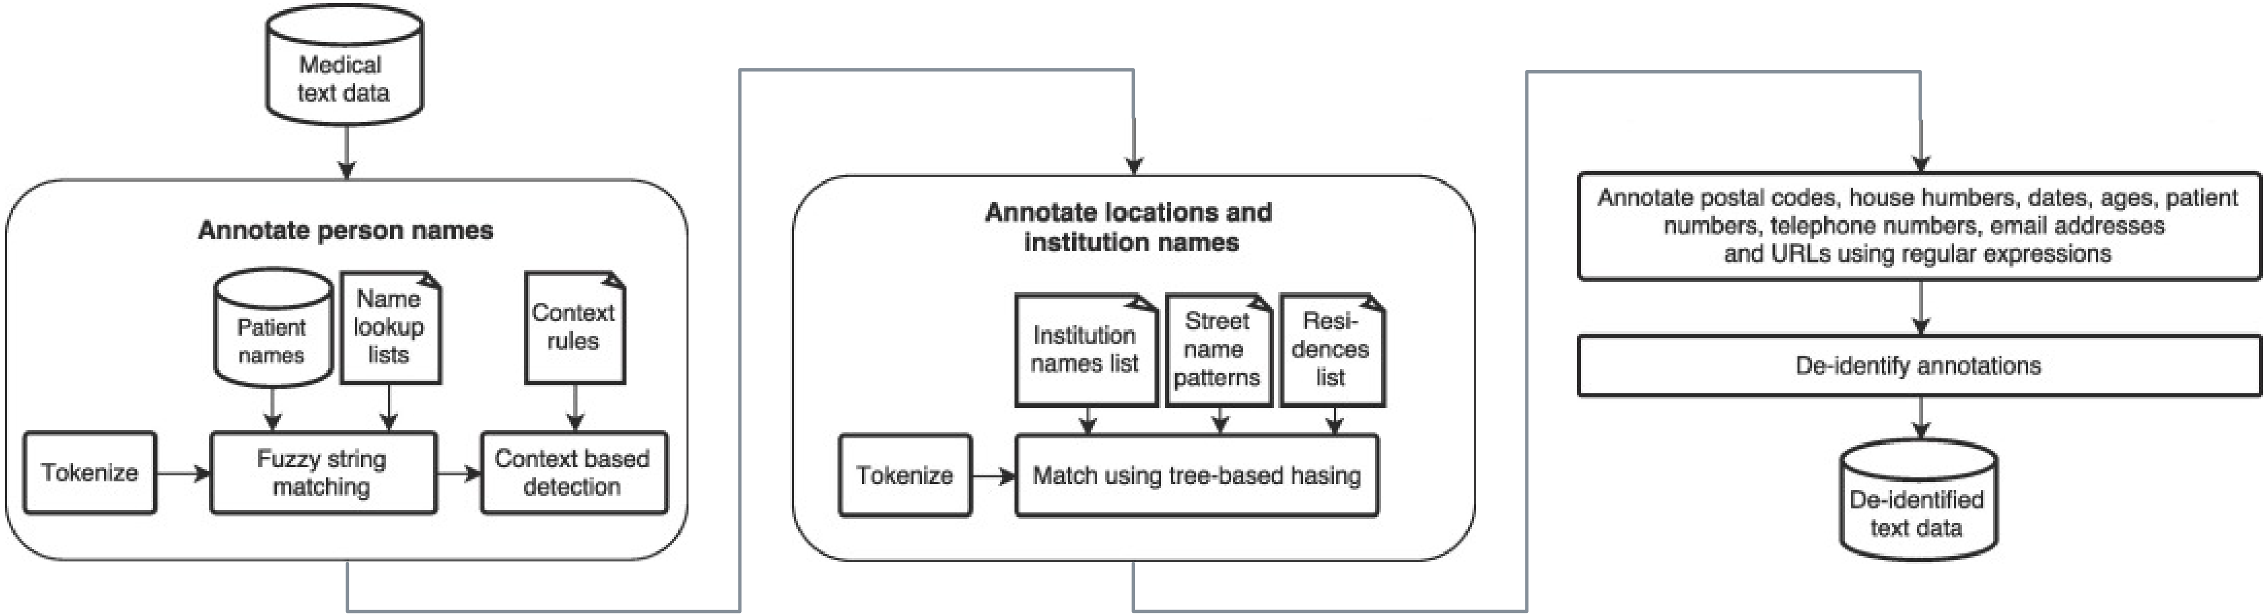

EXAMPLE #1: DE-IDENTIFICATION - DEMO

Menger,V., Scheepers,F., Wijk,L. van, & Spruit,M. (2018). DEDUCE: A pattern matching method for automatic de-identification of Dutch medical text. *Telematics and Informatics, 35*(4), Patient Centric Healthcare, 727–736. [JIF: 6.182] [pdf] [online]

- https://tdslab.nl/deduce

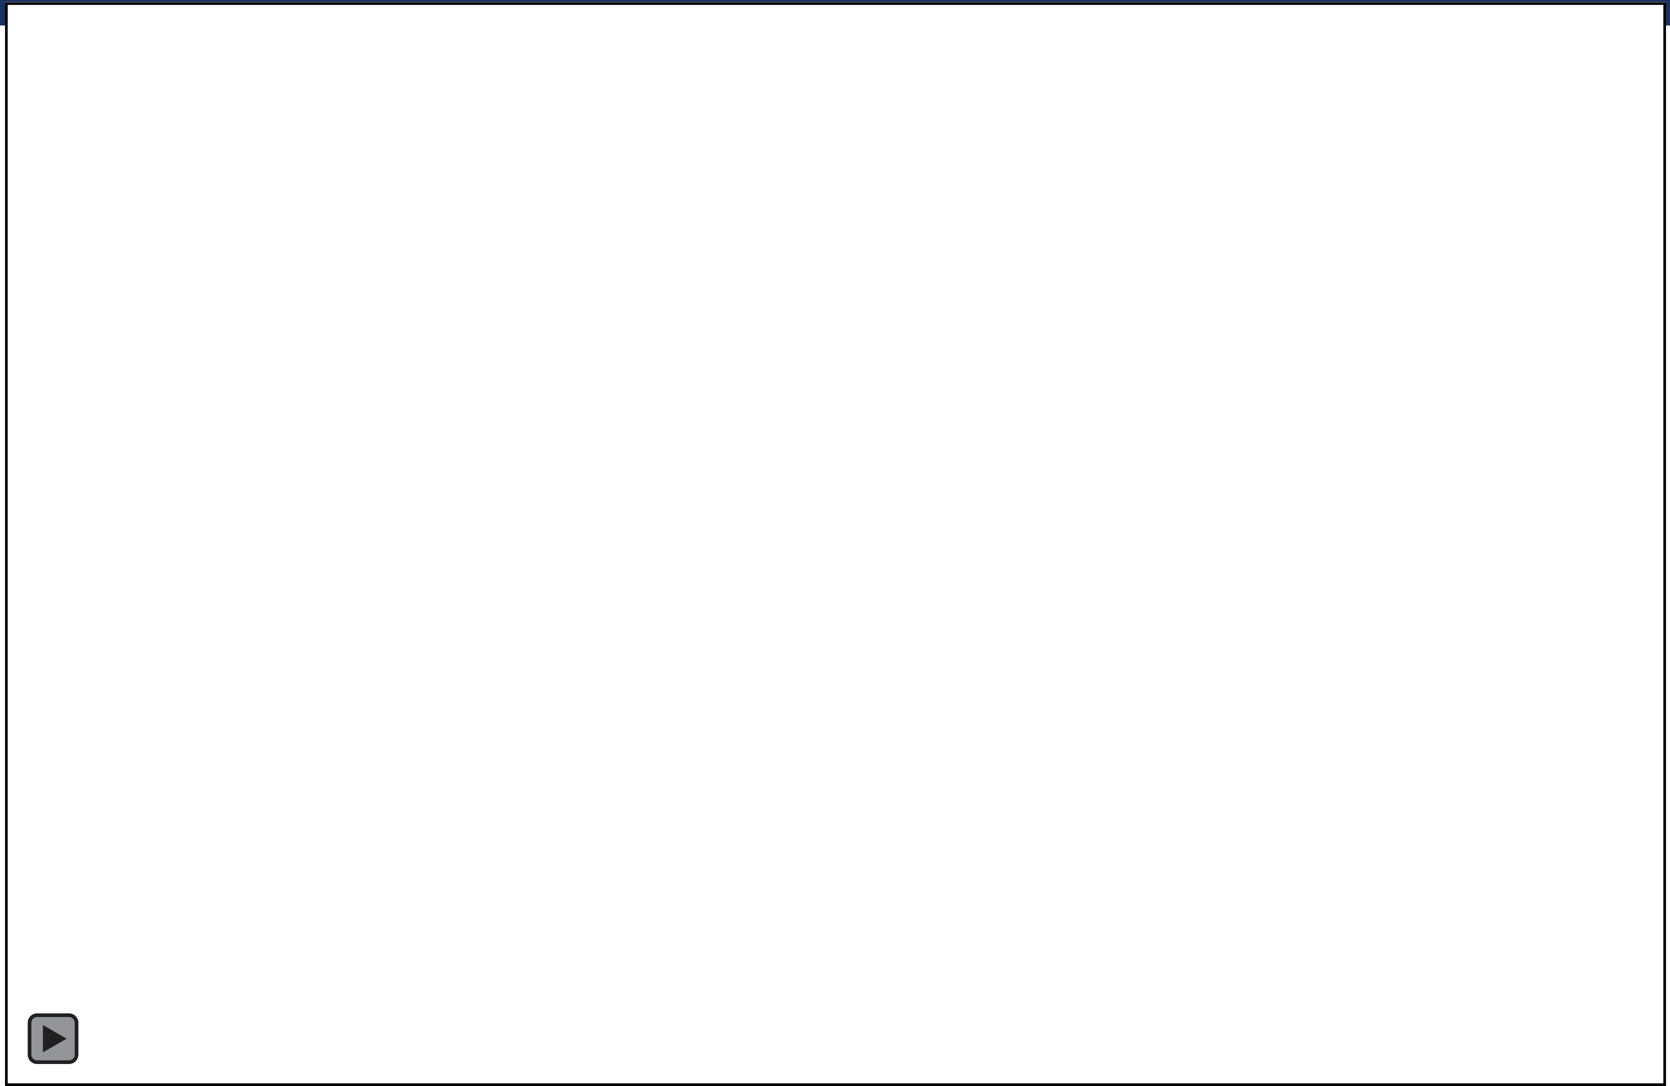

Example #2: ADR extraction

Shen,Z., & Spruit,M. (2021). Automatic Extraction of Adverse Drug Reactions from Summary of Product Characteristics. *Applied Sciences, 11*(6), Applications of Artificial Intelligence in Pharmaceutics, 2663. [JIF: 2.679] [pdf] [online]

- *"Automatic Extraction of Adverse Drug Reactions from Summary of Product Characteristics"*
- European Medicines Agency's…
  - The Electronic Medicines Compendium (EMC @UK)
    - >14,000 documents
    - ~ Structured Product Labels (US Food & Drug Admin. (FDA))
- *Aim:* Document how to safely use medicines, incl. ADRs
  - In Section 4.8 of the SoPC →
  - But, SoPCs still have a heterogeneous nature…
  - NLP for webscraping! → 647 medicines, in tablet form
  - **MedDRA:** Medical Dictionary for Regulatory Activities

- What *is* a Summary of Product Characteristics (SmPC) ?

https://www.medicines.org.uk/emc

4.8 Undesirable effects

Summary of the safety profile

Headache, abdominal pain, diarrhoea and nausea are among those adverse reactions that have been most commonly reported in clinical trials (and also from post-marketing use). In addition, the safety profile is similar for different formulations, treatment indications, age groups and patient populations. No dose-related adverse reactions have been identified.

Tabulated list of adverse reactions

The following adverse drug reactions have been identified or suspected in the clinical trials programme for esomeprazole and post-marketing. None was found to be dose-related. The reactions are classified according to frequency (very common > 1/10; common ≥1/100 to <1/10; uncommon ≥1/1000 to <1/100; rare ≥1/10000 to <1/1000; very rare <1/10000); not known (cannot be estimated from the available data).

Blood and lymphatic system disorders

Rare: Leukopenia, thrombocytopenia

Very rare: Agranulocytosis, pancytopenia

Immune system disorders

Rare: Hypersensitivity reactions e.g. fever, angioedema and anaphylactic reaction/shock

Metabolism and nutrition disorders

Uncommon: Peripheral oedema

Rare: Hyponatraemia

Not known: Hypomagnesaemia (see section 4.4); severe hypomagnesaemia can correlate with hypocalcaemia. Hypomagnesaemia may also be associated with hypokalaemia

Psychiatric disorders

Uncommon: Insomnia

Rare: Agitation, confusion, depression

Very rare: Aggression, hallucinations

Example #2: ADR Knowledge Base: MedDRA

We'll meet MedDRA again later!

https://www.meddra.org/browsers

- **MedDRA:** Medical Dictionary for Regulatory Activities
  - Hierarchical, multilingual, standardised taxonomy of Adverse events

	Class	Term
1	**System Organ Class**	Cardiac disorders
2	High Level Group Term	Heart failures
3	High Level Term	Left ventricular failures
4	Preferred Term	Chronic left ventricular failure
5	Lowest Level Term	Chronic diastolic heart failure

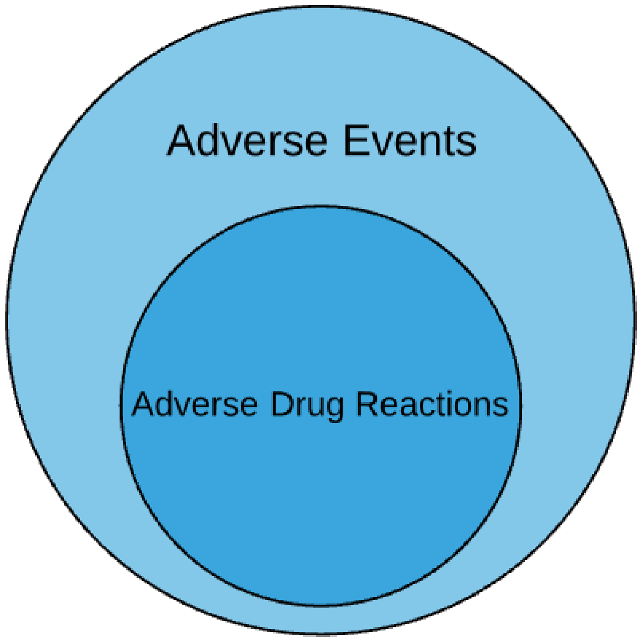

EXAMPLE #2: ADR EXTRACTION - METHOD

Shen,Z., & Spruit,M. (2021). Automatic Extraction of Adverse Drug Reactions from Summary of Product Characteristics. *Applied Sciences, 11*(6), Applications of Artificial Intelligence in Pharmaceutics, 2663. [JIF: 2.679] [pdf] [online]

- Adverse Drug Reactions (ADRs) extraction dev pipeline

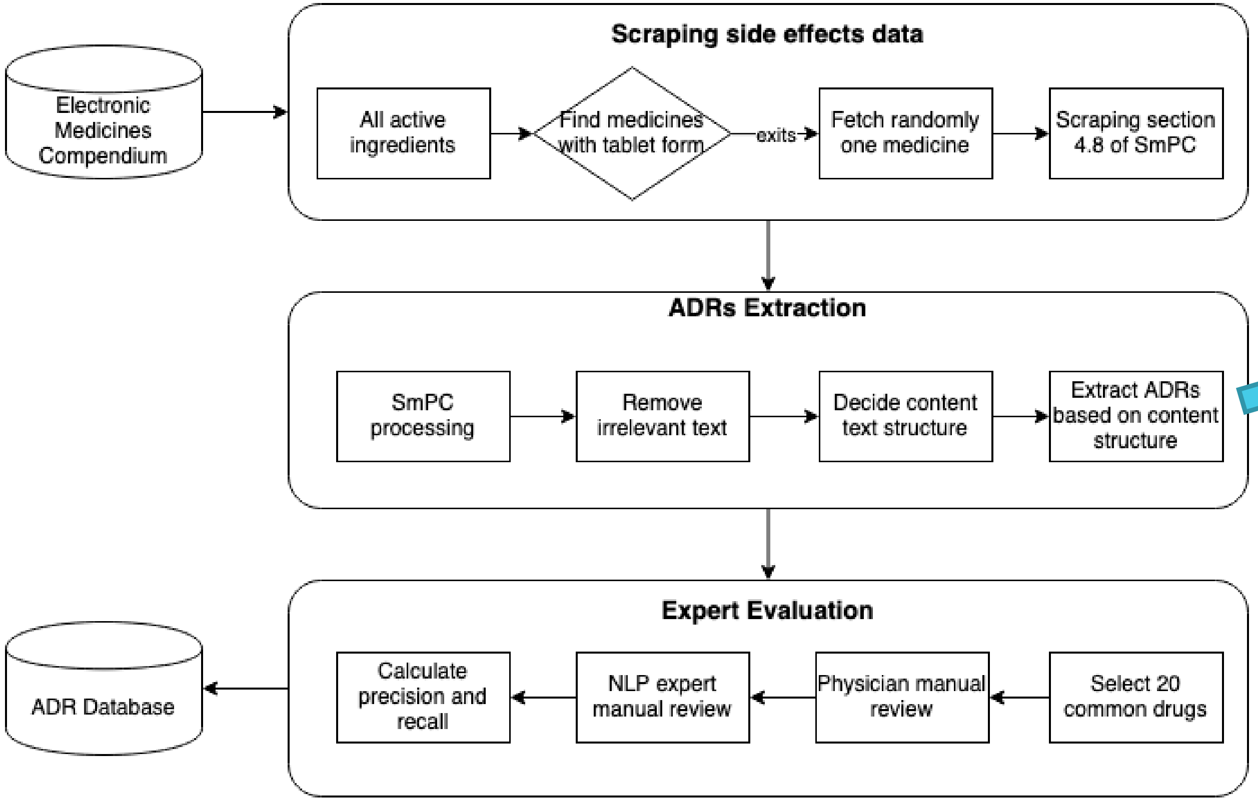

- 3 flows within the ADR Extraction Pipeline

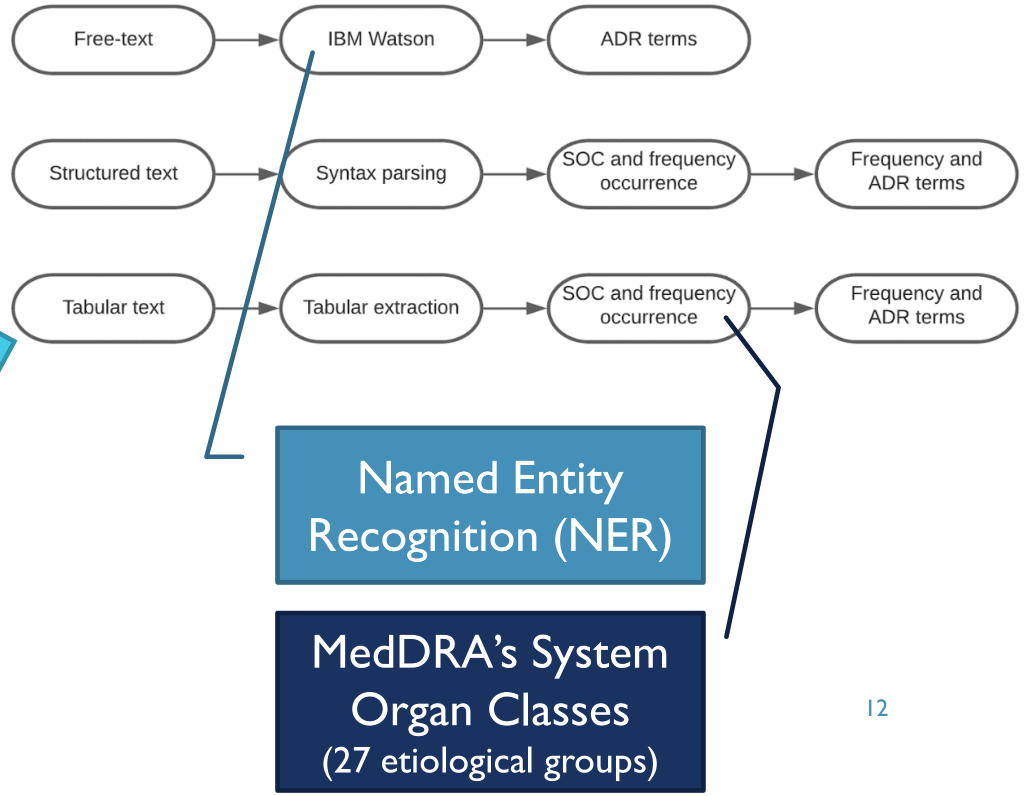

EXAMPLE #2: ADR EXTRACTION - PARSING

Shen,Z., & Spruit,M. (2021). Automatic Extraction of Adverse Drug Reactions from Summary of Product Characteristics. *Applied Sciences, 11*(6), Applications of Artificial Intelligence in Pharmaceutics, 2663. [JIF: 2.679] [pdf] [online]

- *Parsing:* An Adverse Drug Reactions (ADRs) extraction example for the structured text:
  https://www.medicines.org.uk/emc/product/2729#UNDESIRABLE_EFFECTS

Immune system disorders:

Very rare: Hypersensitivity reactions including urticaria, angio-oedema or anaphylactic reactions.

Psychiatric disorders:

Common: Decreased libido.

Uncommon: Increased libido.

Gastrointestinal disorders:

Very common: Diarrhoea.

Common: Abdominal pain, nausea, vomiting, flatulence.

Skin and subcutaneous tissue disorders:

Common: Pruritus, maculo-papular rash.

Not known: Vesiculo-bullous eruptions.

Reproductive system and breast disorders:

Common: Frigidity or impotence.

ADRs Extraction →


```
"atc_code": "N07BB03",
"adrs": {
  "very rare": [
    "hypersensitivity reactions including urticaria",
    "angio-oedema or anaphylactic reactions."
  ],
  "uncommon": [
    "increased libido."
  ],
  "common": [
    "decreased libido.",
    "abdominal pain",
    "nausea",
    "vomiting",
    "flatulence.",
    "pruritus",
    "maculo-papular rash.",
    "frigidity or impotence."
  ],
  "very common": [
    "diarrhoea."
  ],
  "unknown": [
    "vesiculo-bullous eruptions."
  ]
}
```

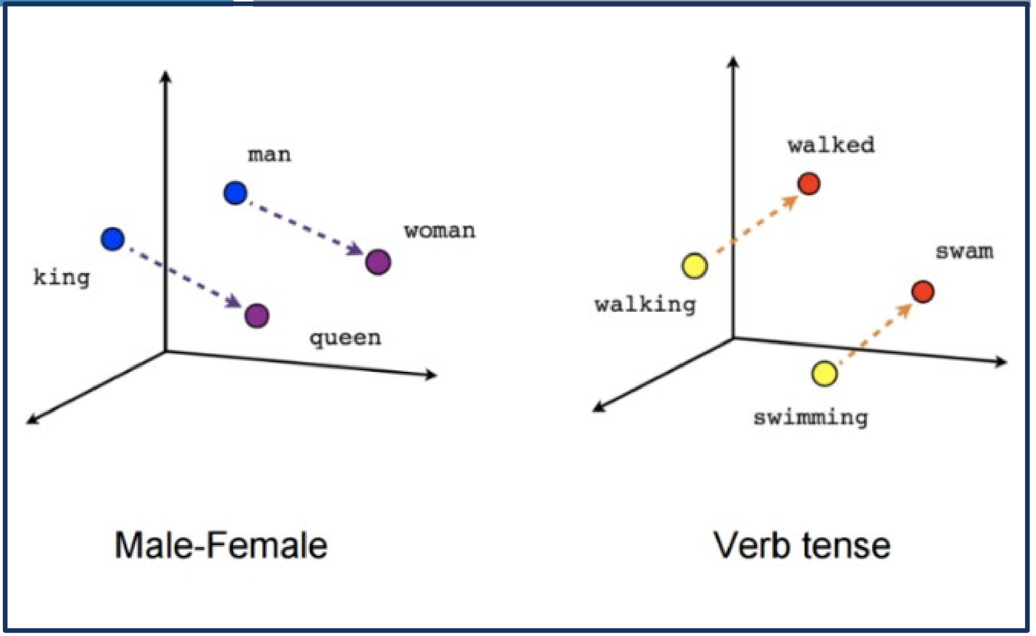

“Modern” NLP → Probabilistic NLP

EMBEDDINGS FOR CLASSIFICATION (VRA)

WORD EMBEDDINGS AS TEXT REPRESENTATIONS

- A **word embedding** is one of the most popular representations of *document vocabulary.*
- It is capable of capturing context of a word in a document, semantic and syntactic similarity, relation with other words, etc.
- Word embedding are simply : **vector representations** of a particular word.
  - Each word is mapped to one vector, and
  - the vector values are learned in a way that resembles a neural network
  - *Objective:* to have words with similar context occupy close spatial positions.

- **word2vec** is a "predictive" model using
  - *Continous Bag Of Words (CBOW):* takes the context of each word as the input and tries to predict the word corresponding to the context
    - captures co-occurrence one window at a time
    - *Skip-gram* is the inverse of CBOW (is better for rare words)

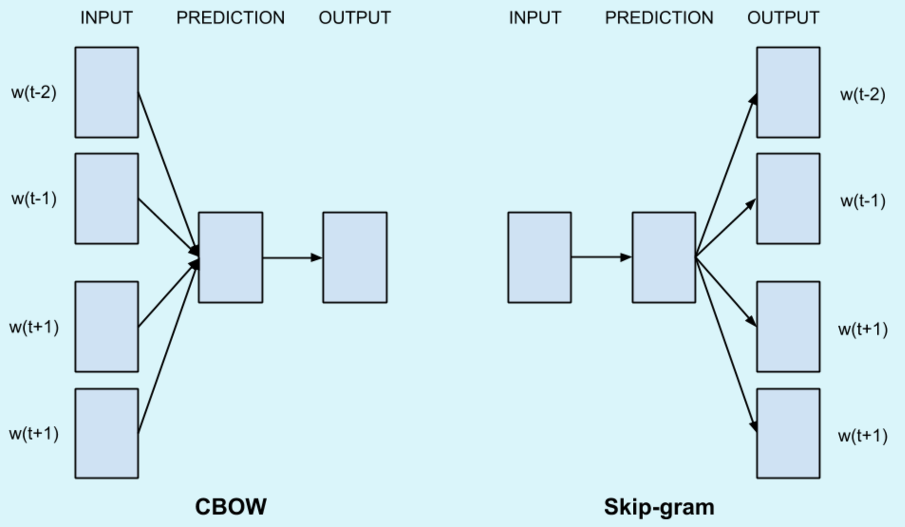

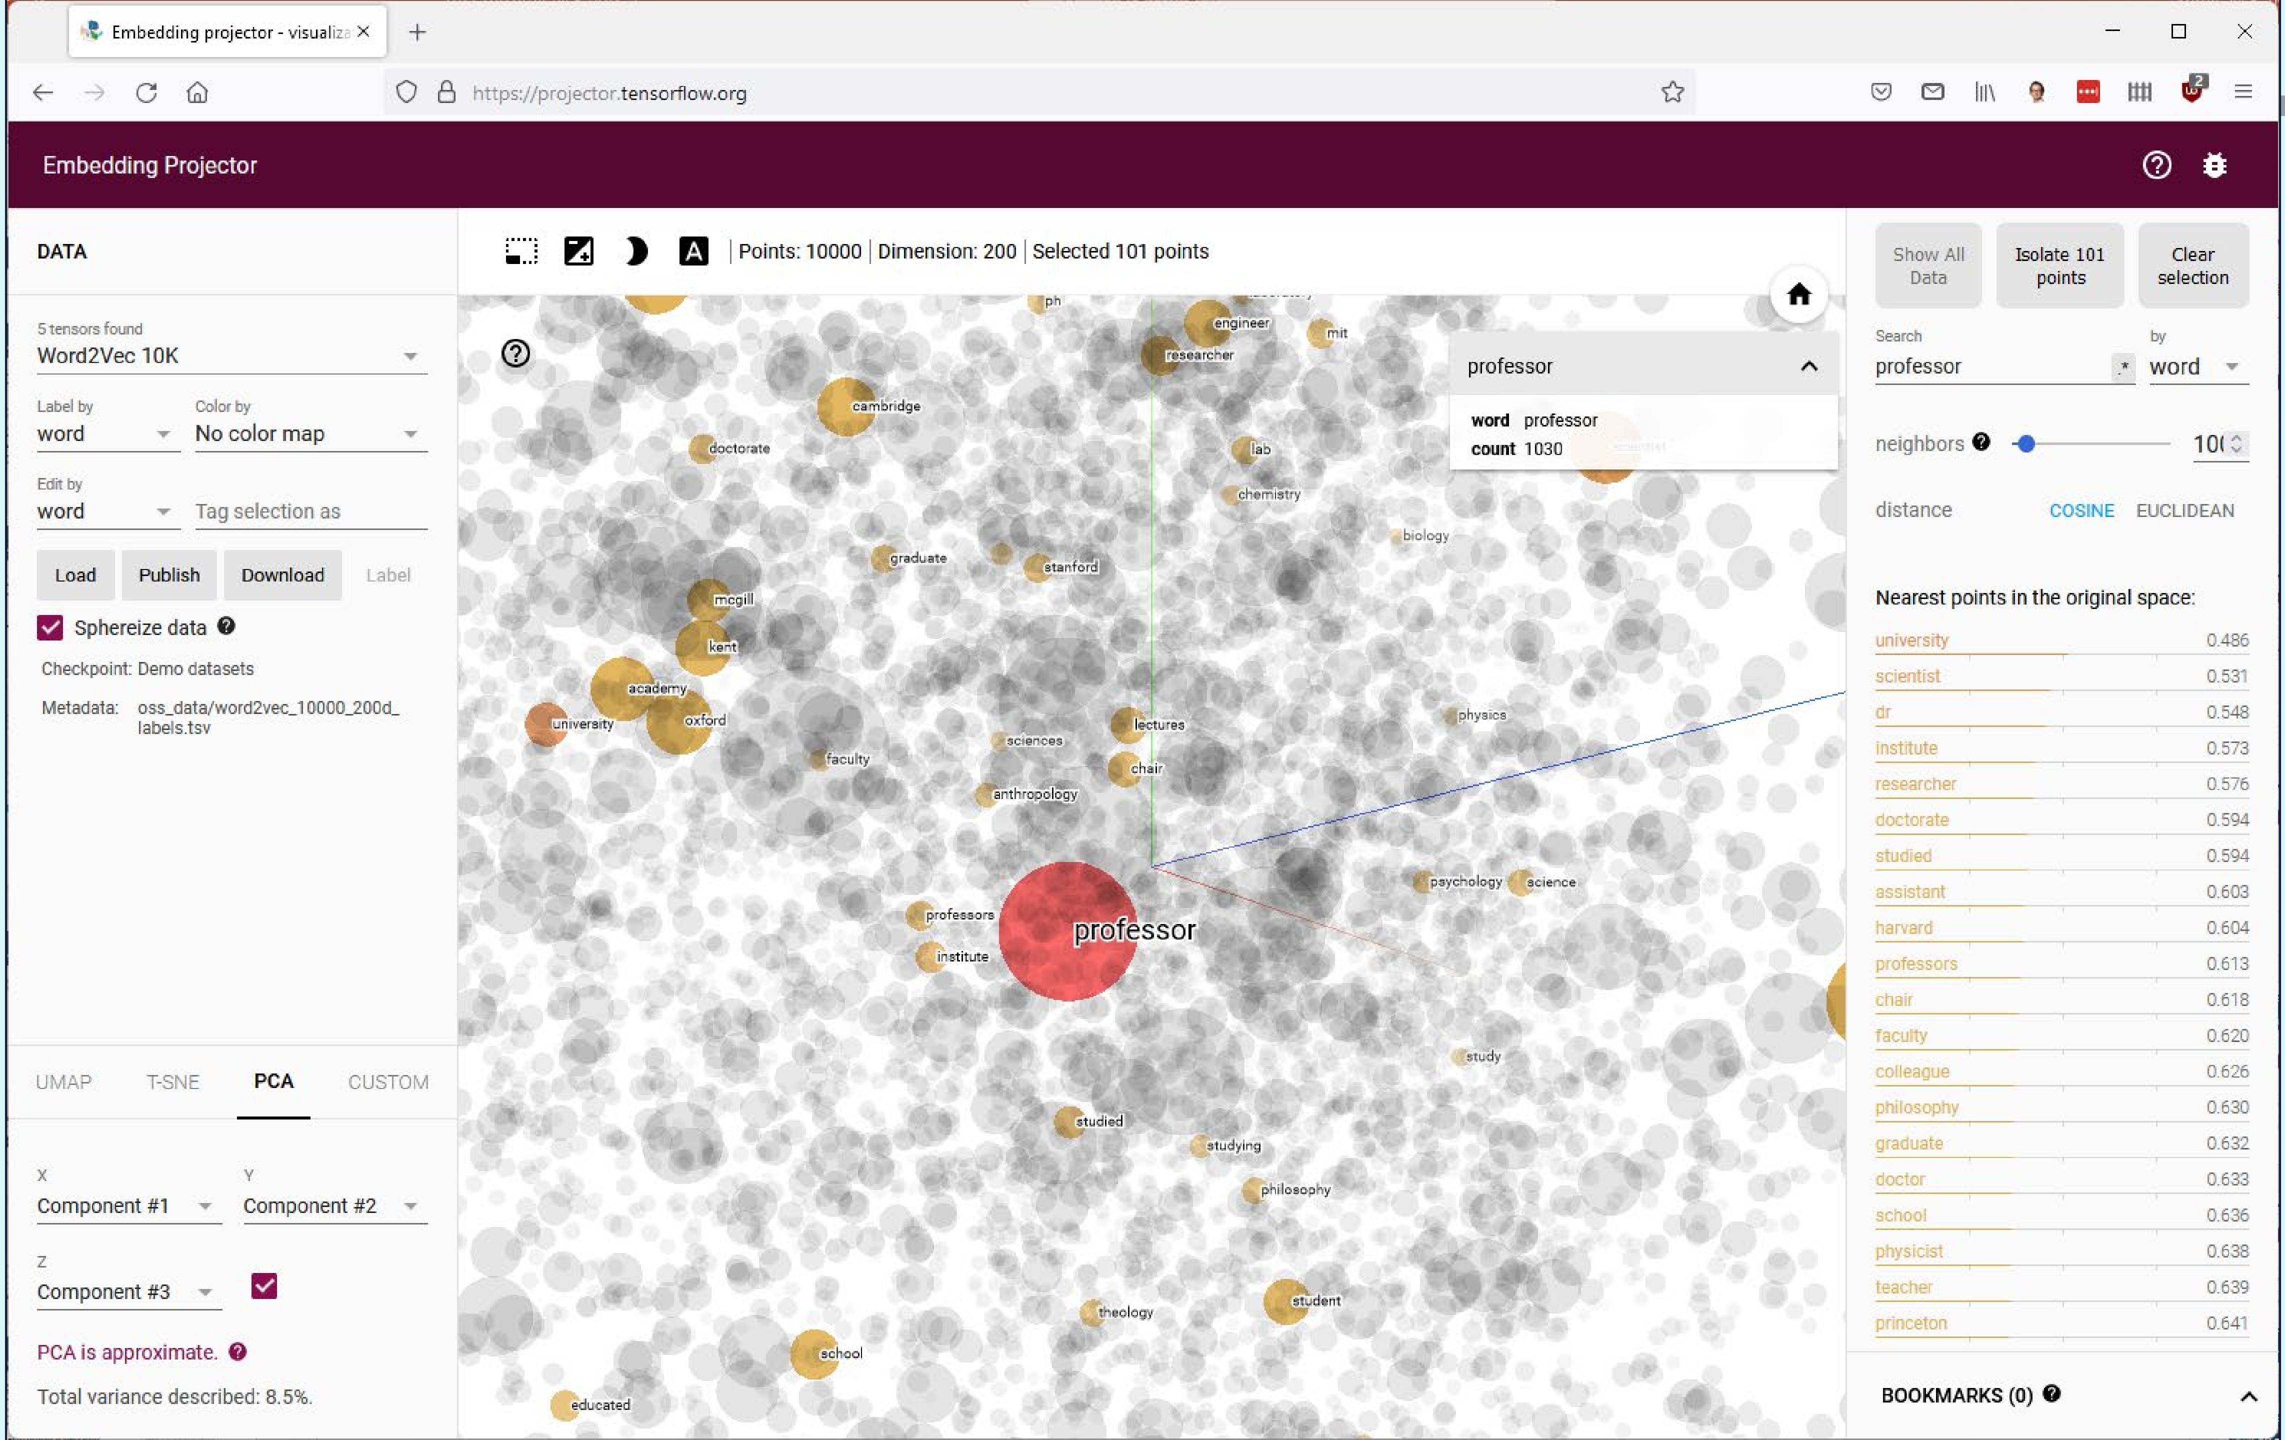

https://projector.tensorflow.org/

… Many state-of-the-art Text Representation Techniques!

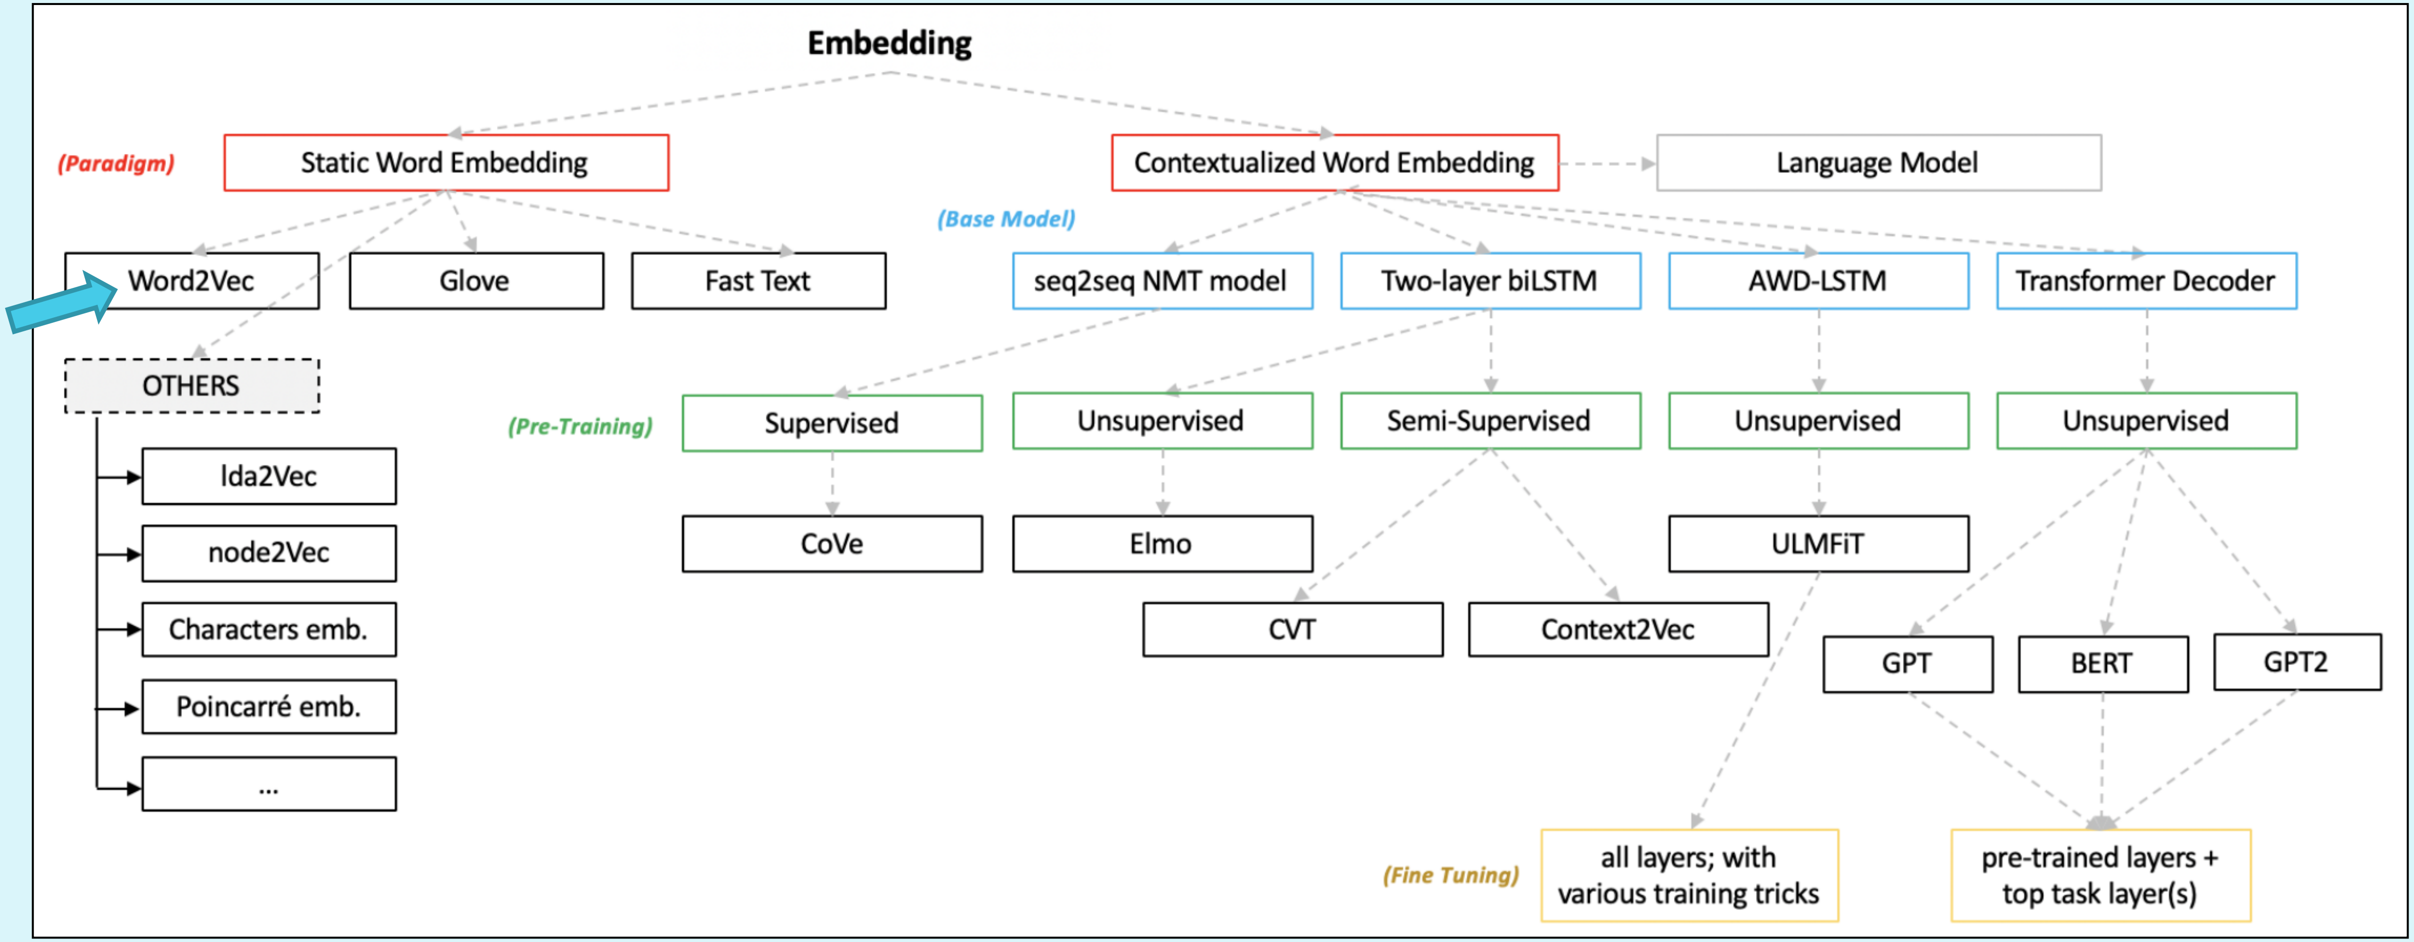

EXAMPLE #3: VIOLENCE RISK PREDICTION

Menger, V., Spruit, M., Est, R. van, Nap, E., & Scheepers, F. (2019). Machine Learning Approach to Inpatient Violence Risk Assessment Using Routinely Collected Clinical Notes in Electronic Health Records. *JAMA Network Open*, 2(7), e196709. [JIF: 8.483] [pdf] [online]

- "Predict for which admissions a violence incident will occur in the first 30 days, **based on clinical texts** that are written up to and including the first day of admission"
  - 2*3200 admissions, 2000 words/note, 950+650 incidents
  - Prediction task excludes incidents on Day 1 of admission
  - 30 days interval chosen for sufficient specificity
- Internal and external validation (UMCU, Antes R'dam)
  - Area Under Curve (AUC) to report performance

(2012-07-29)
"Mw heeft matig geslapen, sliep van 1.00 uur tot 4.00 uur. Kwam toen uit bed, at koekjes en dronk thee. Nog geadviseerd medicatie te nemen en mijn zorgen geuit over evt. doorschieten in een manie. Mw was er niet gevoelig voor en reageerde geagiteerd. Mw had spreekdrang maar gaf aan dat wanneer zij zich goed voelt ook veel praat. Mw gaat vandaag naar <PERSOON-1> met haar zoon, ziet daar nu niet meer tegenop omdat de klachten die zij gisteren aan haar voeten ervaarde verdwenen zijn. Mw ging na 4.00 uur weer naar bed en kwam niet meer uit haar kamer tot de ochtend."

?

EXAMPLE #3: VIOLENCE RISK – DATA SAMPLE

Menger,V., Spruit,M., Est,R. van, Nap,E., & Scheepers,F. (2019). Machine Learning Approach to Inpatient Violence Risk Assessment Using Routinely Collected Clinical Notes in Electronic Health Records. *JAMA Network Open*, 2(7), e196709. [JIF: 8.483] [pdf] [online]

Text representation

- Represent all clinical notes related to 1 admission as 1 vector (*i.e.* not as words)

- *Representation:* paragraph2vec
- *Classification:* SVM

EXAMPLE #3: VIOLENCE RISK – PERFORMANCE

Menger,V., Spruit,M., Est,R. van, Nap,E., & Scheepers,F. (2019). Machine Learning Approach to Inpatient Violence Risk Assessment Using Routinely Collected Clinical Notes in Electronic Health Records. *JAMA Network Open*, 2(7), e196709. [JIF: 8.483] [pdf] [online]

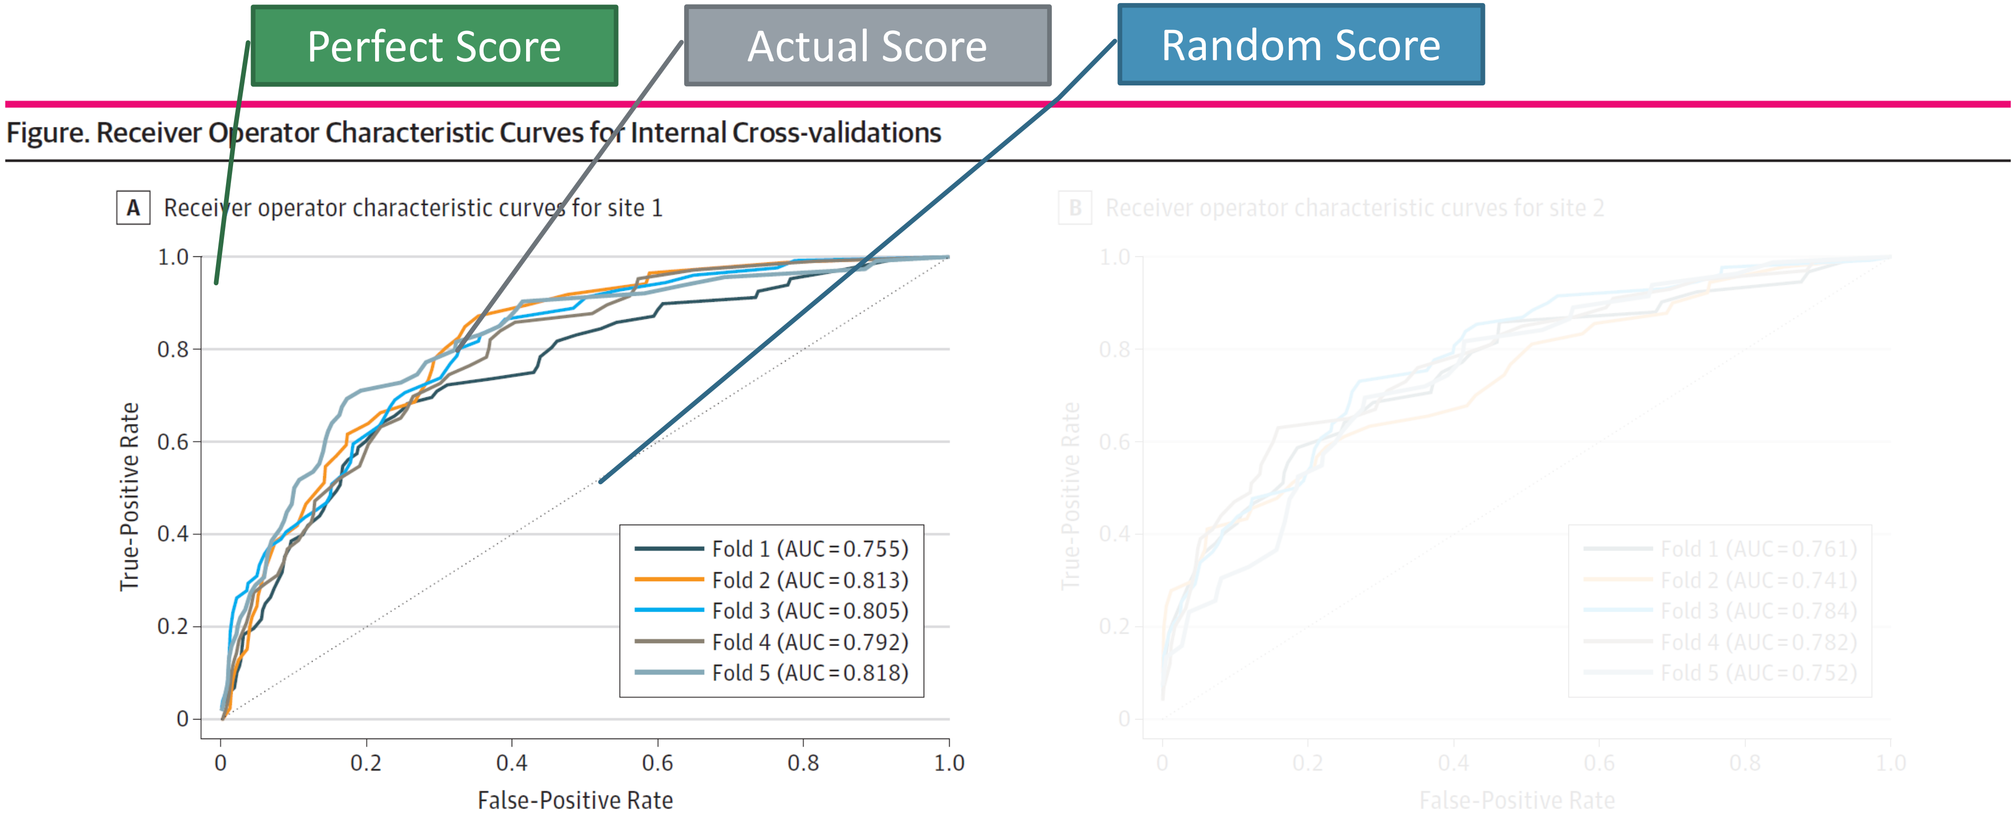

Receiver operator characteristic curves are shown for each fold, according to internal cross-validation in site 1 (A) and site 2 (B). Dashed diagonal lines denote an area under the curve (AUC) of 0.5, ie, predictive validity equivalent to chance. AUC indicates area under the curve.

Example #3: Violence Risk – Interpretation

Menger, V., Spruit, M., Est, R. van, Nap, E., & Scheepers, F. (2019). Machine Learning Approach to Inpatient Violence Risk Assessment Using Routinely Collected Clinical Notes in Electronic Health Records. *JAMA Network Open*, 2(7), e196709. [JIF: 8.483] [pdf] [online]

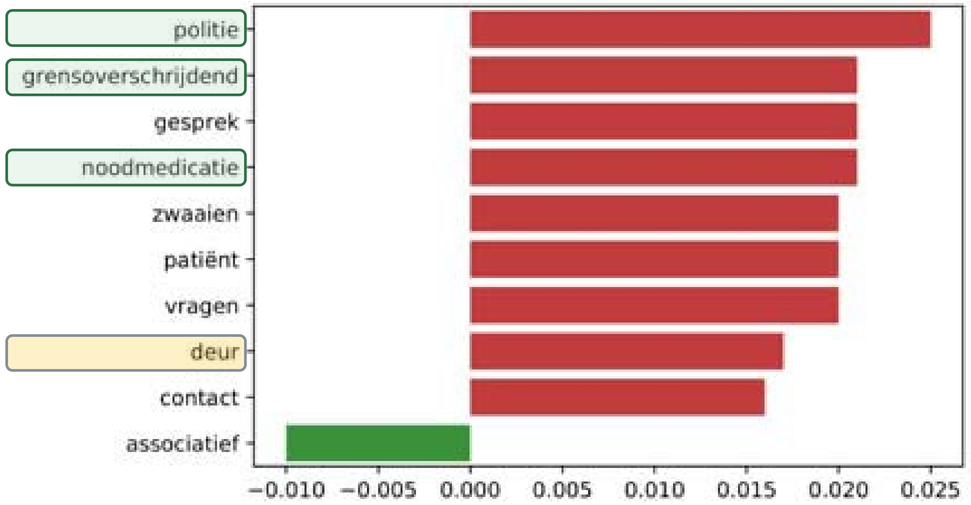

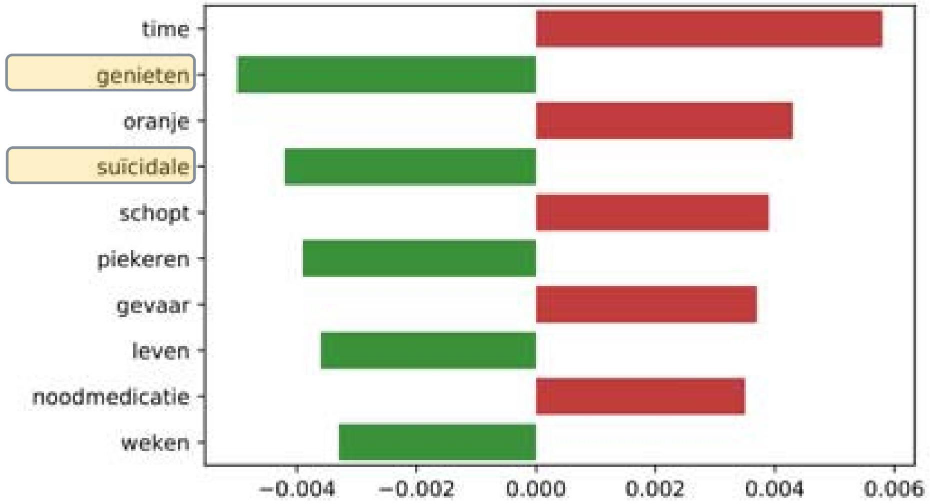

- Sample of Local Explanation predicting <u>high</u> risk of aggression

The *"Linear Model-Agnostic Explanations"* (LIME) method

- Sample of Local Explanation predicing <u>low</u> risk of aggression

EXAMPLE #3: ALTERNATIVE MODELS

Menger,V., Scheepers,F., & Spruit,M. (2018). Comparing Deep Learning and Classical Machine Learning Approaches for Predicting Inpatient Violence Incidents from Clinical Text. *Applied Sciences, 8*(6), Data Analytics in Smart Healthcare, 981. [JIF: 2.679]pdf] [online]

- In previous work, we determined SVM as an appropriate classifier for VRA, based on literature and experiments

Table 4. The performance for optimal hyperparameter values for each of the representations combined with the models, based on a 5-fold stratified cross validation. The performance is measured in AUC, along with its standard deviation. The best performance over different models is marked with an [a], the best performance over representations with a [b].

Model	Bag-of-Words Binary	Bag-of-Words tf-idf	Word Embeddings	Document Embeddings
RNN [1]	0.771 ± 0.018 [b]	0.753 ± 0.031	0.654 ± 0.043	0.788 ± 0.018 [a,b]
CNN [2]	0.729 ± 0.030	0.716 ± 0.038	0.684 ± 0.038	0.763 ± 0.024 [a]
NN [3]	0.727 ± 0.033	0.717 ± 0.038	0.751 ± 0.036 [a]	0.745 ± 0.022
NB [4]	0.686 ± 0.026	0.704 ± 0.034 [a]	0.700 ± 0.051	0.692 ± 0.046
SVM [5]	0.759 ± 0.040	0.756 ± 0.036 [b]	0.764 ± 0.024 [b]	0.770 ± 0.029 [a]
DT [6]	0.727 ± 0.018 [a]	0.719 ± 0.041	0.685 ± 0.041	0.665 ± 0.035

[1] Recurrent Neural Network; [2] Convolutional Neural Network; [3] Neural Network; [4] Naive Bayes; [5] Support Vector Machine; [6] Decision Tree.

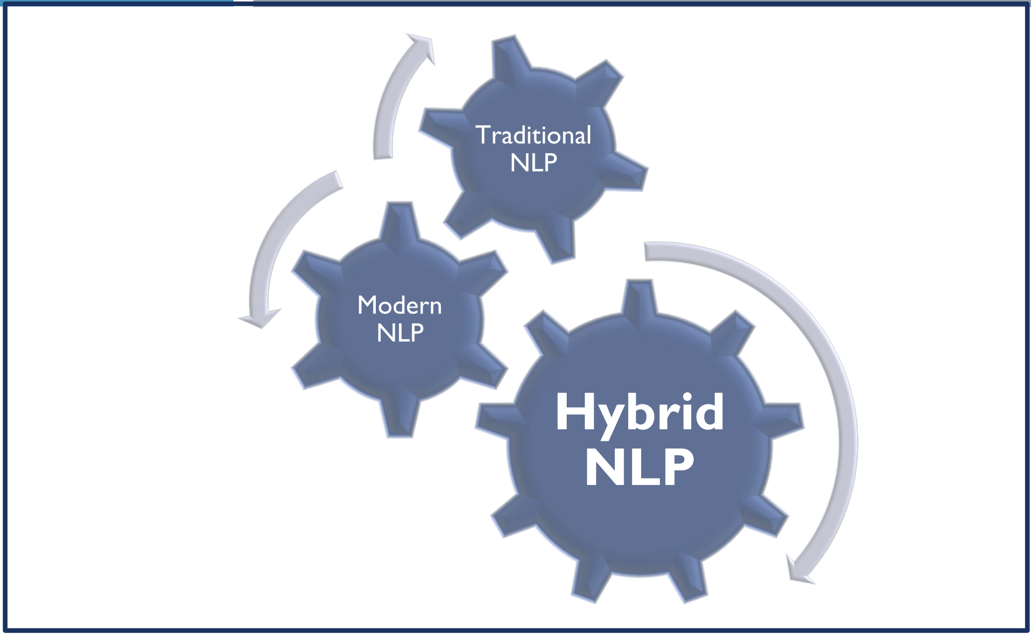

“Effective” NLP → Hybrid NLP

Combining traditional and modern approaches (*e.g.* ADRIN)

EXAMPLE #4: ADR IDENTIFICATION, REVISITED

Siegersma,K., Evers,M., Bots,S., Groepenhoff,F., Appelman,Y., Hofstra,L., Tulevski,I., Somsen,A., Den Ruijter,H., Spruit,M., & Onland-Moret,C. (2022). Adverse Drug Reactions Identification in clinical Notes (ADRIN): Word embedding models and string matching. *JMIR Medical Informatics, 10*(1), e31063. [JIF: 2.96] [pdf] [online]

- *Hypothesis:* Information on ADRs is present in clinical notes, but is underreported in EHRs
- *Goal:* Method for recognising ADRs in Dutch clinical notes
- *Method:* Case study, using clinical notes from the UMCU for developing a prototype that automatically identifies ADRs in clinical notes
- *Dataset:* Cardiology Centre Netherlands (CCN)
  - 109.151 patients between 2007-2018
  - 277.389 unique clinical notes
  - In 36.533 clinical notes, why medication is stopped
    - 1.556 notes where the doctor noted an ADR
- Manual labelling
  - 3.156 clinical notes: validation; 1.000 notes: test

- What *is* a Dutch CCN clinical note ?

Hartfrequentie over het algemeen te hoog met 90 gemiddeld. Wat nu ook opvalt is dat de nierfunctie in korte tijd is verslechterd naar een klaring van 20 ml.\nDerhalve wordt de medicatie aangepast.
Met de patient gaat het goed. Verdraagt Monocedodard niet. Heeft afgelopen maand een keer druk op de borst gehad, in rust. Inspanning (lopen) gaat goed, 2 kg afgevallen.\n\nBeleid:\nContinueren\nControle over 6 maanden
Bloeddruk blijft goed na halveren Olmetec. Echter zeer forse spierkrampen met name van de kuiten en fors haaruitval. Mogelijk bijwerking van de medicatie?\nB/ stop metoprolol plus controle 2 weken.

Example #4: ADR identification - Method

Adverse Drug Reactions Identification in clinical Notes (ADRIN)

Siegersma,K., Evers,M., Bots,S., Groepenhoff,F., Appelman,Y., Hofstra,L., Tulevski,I., Somsen,A., Den Ruijter,H., Spruit,M., & Onland-Moret,C. (2022). Adverse Drug Reactions Identification in clinical Notes (ADRIN): Word embedding models and string matching. *JMIR Medical Informatics, 10*(1), e31063. [JIF: 2.96] [pdf] [online]

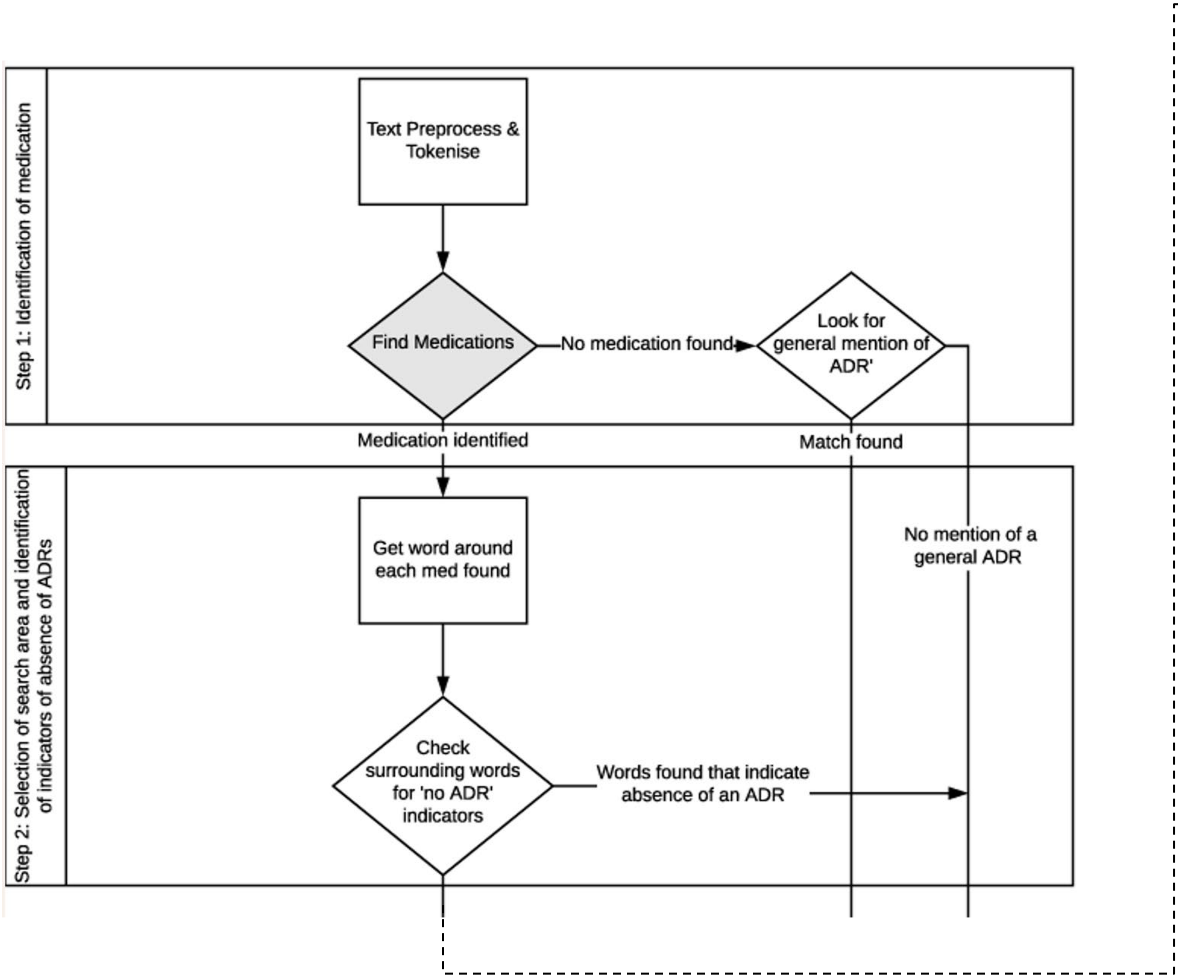

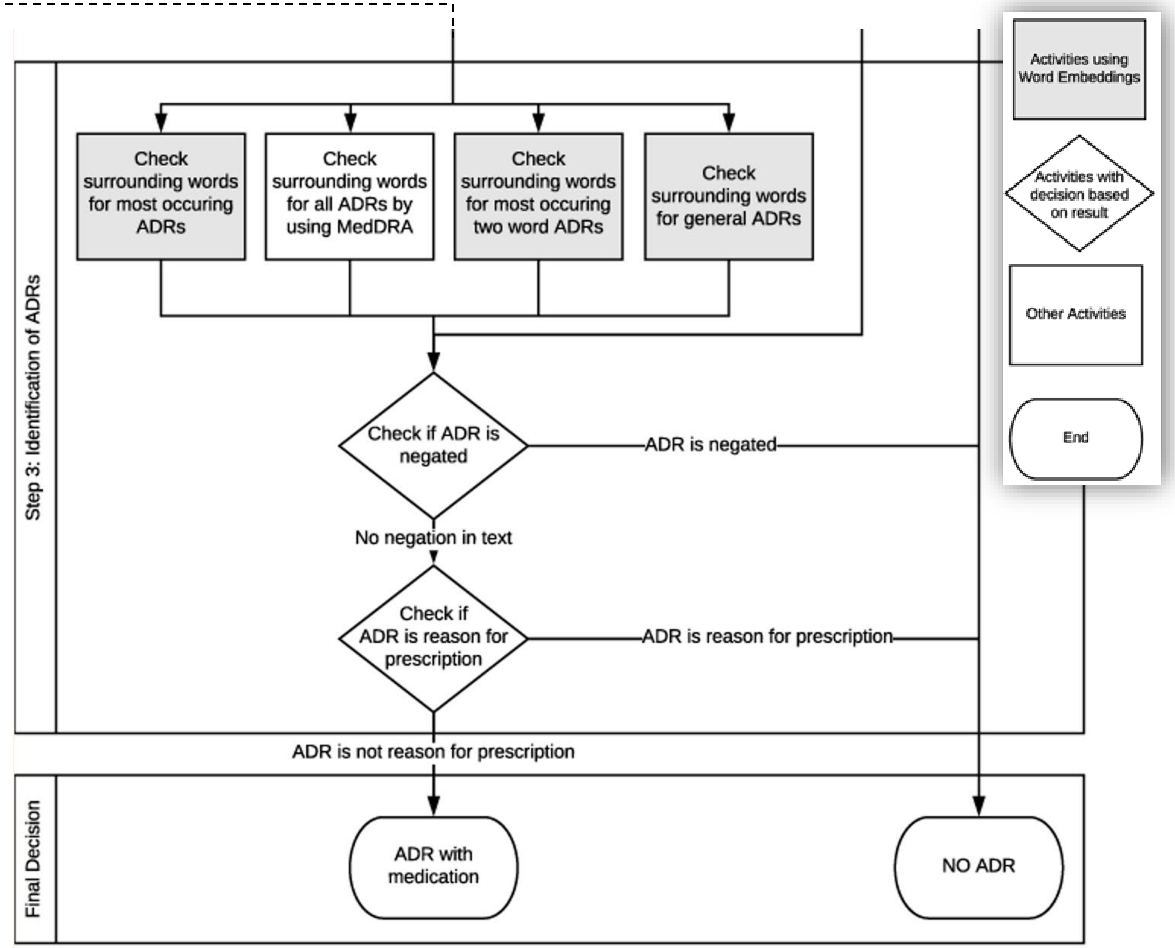

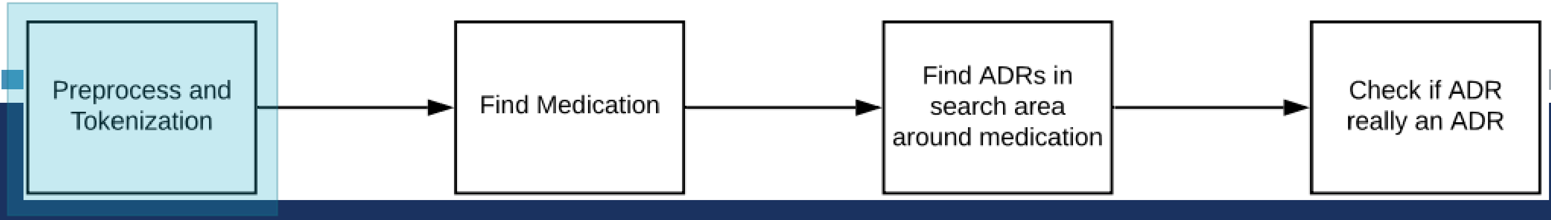

Example #4: ADR identification – Method – Step 1

1. Lowercase
2. Unidecode
3. Tokenize
4. Change numbers
5. Remove unknown words

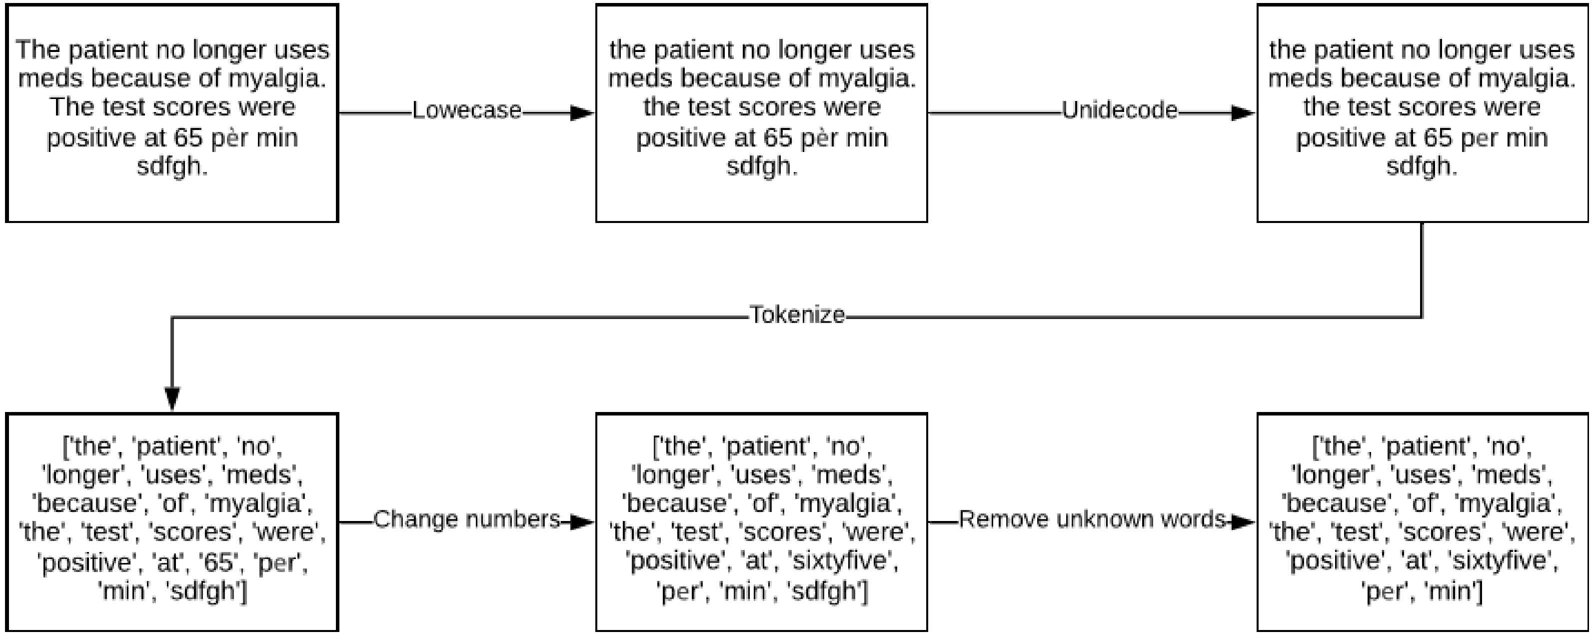

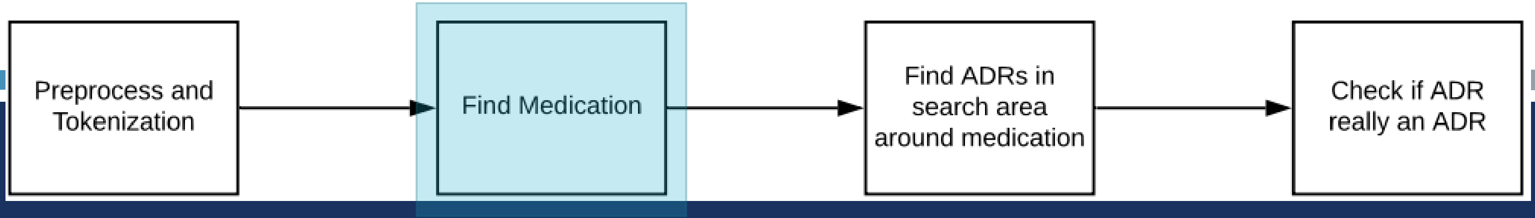

EXAMPLE #4: ADR IDENTIFICATION - METHOD – STEP 2

- *Idea:* Words that have similar neighbouring words are similarly shaped
  - Every word is represented in a numerical vector
  - Trained on all 277.389 clinical notes
- Because models are trained on domain specific text, domain specific results →
- The **Word2Vec** approach is used, *i.e.* vectors are shaped based upon their neighbouring words
  - "King – Man + Woman = Queen" →
  - **"Patient – Man + Woman = Patiente"**

```
In [11]: model_all.wv.most_similar(['rood'])

Out[11]: [('jeukend', 0.7554248571395874),
          ('opgezwollen', 0.7433047890663147),
          ('jeukende', 0.742124915122985 8),
          ('gezwollen', 0.7398363351821899),
          ('geirriteerd', 0.7380100488662720),
          ('verkleuringen', 0.7336500287055969),
          ('verkleuring', 0.7277984023094177),
          ('zere', 0.7272820472717285),
          ('paars', 0.7250189185142517),
          ('verkleurd', 0.7249985933303833)]
```

```
In [79]: model_all.wv.most_similar(positive=['patient', 'vrouw'], negative=['man'], topn = 1)

Out[79]: [('patiente', 0.8561874032020569)]
```

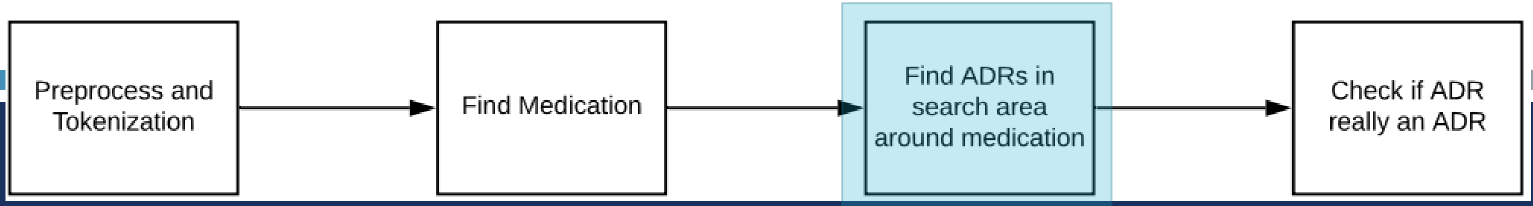

Example #4: ADR identification - Method – Step 3

- For ADR recognition with the word embedding models, for every word in the search area the similarity with predefined search words is computed.
- The predefined search words consist of the most occurring ADRs
- If the similarity is above a certain threshold, the word is identified as an ADR
- For the **MedDRA** search, a *Regular Expression* for each term in MedDRA is executed to see if it occurs in the search area

- Then, Check if ADRs *are* indeed ADRs
- When the prototype has found one or more [medication, ADR] combinations, two things are checked:
  1. Is the ADR not negated
  2. Is the ADR not the reason for prescription of medication

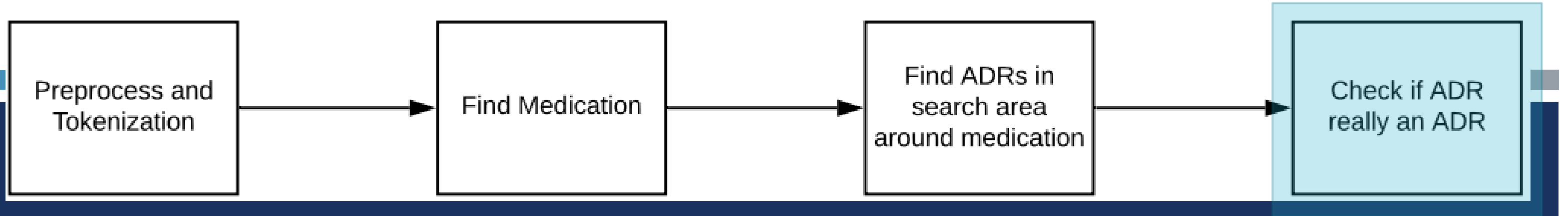

Example #4: ADR identification - Method – Step 4, with Data Sample

- *"The betablocker causes headaches and muscle pain. Because of these complaints the decision is made to stop bisoprolol and metoprolol."*

1. ['the', 'betablocker', 'causes', 'headaches', 'and', 'muscle', 'pain', 'because', 'of', 'these', 'complaints', 'the', 'decision', 'is', 'made', 'to', 'stop', 'bisoprolol', 'and', 'metoprolol']
2. ['the', '**betablocker**', 'causes', 'headaches', 'and', 'muscle', 'pain', 'because', 'of', 'these', 'complaints', 'the', 'decision', 'is', 'made', 'to', 'stop', '**bisoprolol**', 'and', '**metoprolol**']
3. Search area of 5 words before and 5 words after medication

   betablocker, ['the', 'betablocker', 'causes', 'headaches', 'and', 'muscle', 'pain',]

   bisoprolol, ['decision', 'is', 'made', 'to', 'stop', 'bisoprolol', 'and', 'metoprolol']

   metoprolol, ['made', 'to', 'stop', 'bisoprolol', 'and', 'metoprolol']

   → [betablocker, headaches]

   → [betablocker, muscle pain]
4. None of the ADRs are negated or the reason for prescription

Example #4: ADR identification - Evaluation

Siegersma,K., Evers,M., Bots,S., Groepenhoff,F., Appelman,Y., Hofstra,L., Tulevski,I., Somsen,A., Den Ruijter,H., Spruit,M., & Onland-Moret,C. (2022). Adverse Drug Reactions Identification in clinical Notes (ADRIN): Word embedding models and string matching. *JMIR Medical Informatics, 10*(1), e31063. [JIF: 2.96] [pdf] [online]

- *Computational experiment:* Evaluate different versions of the prototype by varying the search area size that is used for identifying ADRs

Version	Search area	Sentences?	MedDRA?
1	All	No	Yes
2	10	No	Yes
3	5	No	Yes
4	All	Yes	Yes
5	10	Yes	Yes
5b	10	Yes	No
6	5	Yes	Yes

No MedDRA…

- *Evaluation tasks:* Four tasks for prototype evaluation:

1. Predict if a text contains one or more ADRs or none
   - 5b: F-score = 0.71
2. Find all present **[medication, ADR]** combinations
   - 5b: F-score = 0.59
3. Find all present ADRs, regardless of medication
   - 2: F-score = 0.67 (5b: 0.66)
4. Find all medication that triggers an ADR, regardless of which ADR
   - 5b: F-score = 0.69

```
>>> import stanza
>>> stanza.download('en')       # This downloads the English models for the neural pipeline
>>> nlp = stanza.Pipeline('en') # This sets up a default neural pipeline in English
>>> doc = nlp("Barack Obama was born in Hawaii.  He was elected president in 2008.")
>>> doc.sentences[0].print_dependencies()
```

THANK YOU FOR LISTENING 👍

m.r.spruit@lumc.nl

https://www.universiteitleiden.nl/en/staffmembers/marco-spruit

Leiden University Medical Center

liacs Leiden Institute of Advanced Computer Science

Universiteit Leiden

Leiden University Campus The Hague

EXAMPLE #5: SIMPLE INTEGRATION

Tawfik,N., & Spruit,M. (2019). Towards Recognition of Textual Entailment in the Biomedical Domain. In Métais, E. et al. (Eds.), *Lecture Notes in Computer Science 11608, NLDB 2019: International Conference on Applications of Natural Language to Information Systems* (pp. 368–375). NLDB 2019, University of Salford, MediaCityUK Campus, United Kingdom, 26–28 June 2019: Springer. [pdf] [online]

- *Computational experiment:* A *feature-assisted* neural network architecture for a Natural Language Inference (NLI) task.

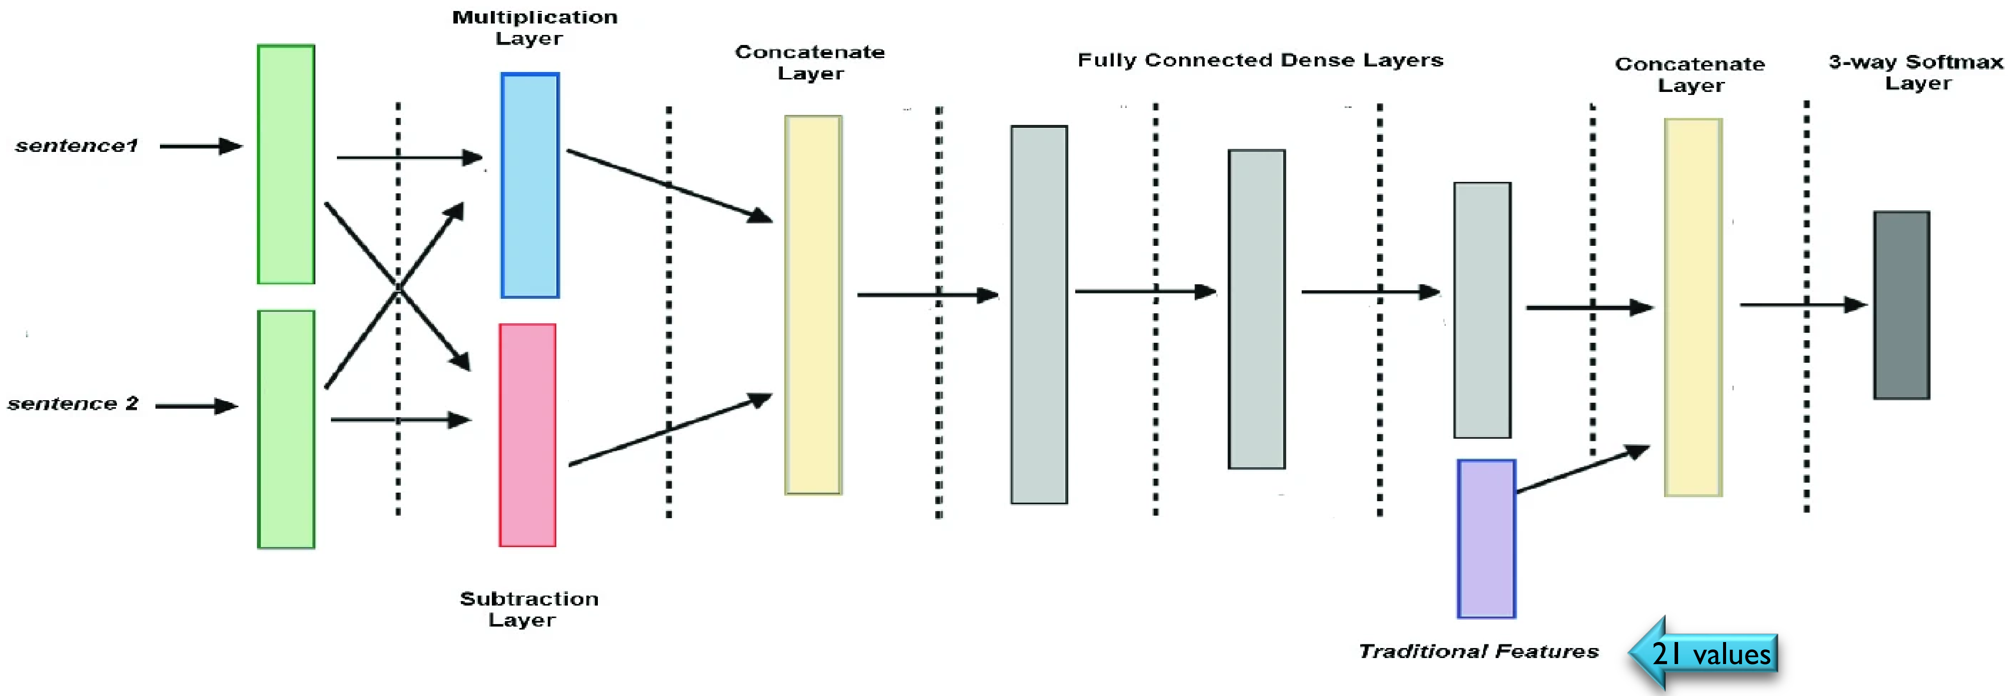

String-Based Features (e.g. editDist)
Contradiction-Based Features (e.g. Negation)
Context-Based Features (e.g. embedSim)

Symbolic NLP	Probabilistic NLP	Hybrid NLP: Combined	Hybrid NLP: Integrated
Menger,V., Scheepers,F., Wijk,L. van, & Spruit,M. (2018). DEDUCE: A pattern matching method for automatic de-identification of Dutch medical text. *Telematics and Informatics, 35*(4), Patient Centric Healthcare, 727–736. [JIF: 6.182] [pdf] [online]	Menger,V., Spruit,M., Est,R. van, Nap,E., & Scheepers,F. (2019). Machine Learning Approach to Inpatient Violence Risk Assessment Using Routinely Collected Clinical Notes in Electronic Health Records. *JAMA Network Open, 2*(7), e196709. [JIF: 8.483] [pdf] [online]	Siegersma,K., Evers,M., Bots,S., Groepenhoff,F., Appelman,Y., Hofstra,L., Tulevski,I., Somsen,A., Den Ruijter,H., Spruit,M., & Onland-Moret,C. (2022). Adverse Drug Reactions Identification in clinical Notes (ADRIN): Word embedding models and string matching. *JMIR Medical Informatics, 10*(1), e31063. [JIF: 2.96] [pdf] [online]	Tawfik,N., & Spruit,M. (2019). Towards Recognition of Textual Entailment in the Biomedical Domain. In Métais, E. et al. (Eds.), *Lecture Notes in Computer Science 11608, NLDB 2019: International Conference on Applications of Natural Language to Information Systems* (pp. 368–375). NLDB 2019: Springer. [pdf] [online]
Shen,Z., & Spruit,M. (2021). Automatic Extraction of Adverse Drug Reactions from Summary of Product Characteristics. *Applied Sciences, 11*(6), Applications of Artificial Intelligence in Pharmaceutics, 2663. [JIF: 2.679] [pdf] [online]	Rijcken,E., Kaymak,U., Scheepers,F., Mosteiro,P., Zervanou,K., & Spruit,M. (2022). Topic Modeling for Interpretable Text Classification from EHRs. *Frontiers in Big Data, 5*, Section Data Mining and Management, 846930. [pdf] [online]	Spruit,M., Verkleij,S., Schepper,C. de, & Scheepers,F. (2022). Exploring Language Markers of Mental Health in Psychiatric Stories. *Applied Sciences, 12*(4), Current Approaches and Applications in Natural Language Processing, 2179. [JIF: 2.679] [pdf] [online]	Zhou,J., Zhang,Z., Zhao,H., and Zhang,S. (2020). LIMIT-BERT : Linguistics Informed Multi-Task BERT. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4450–4461, ACL. <? DEDUCE v2 (in progress) ?>