

Generative Models and the Digital Humanities Towards Synthetic Literature

Algoliterary Lectures
Maison du livre, Bruxelles
10 November 2017

mike.kestemont@uantwerp.be
University of Antwerp, Department of Literature
www.mike-kestemont.org



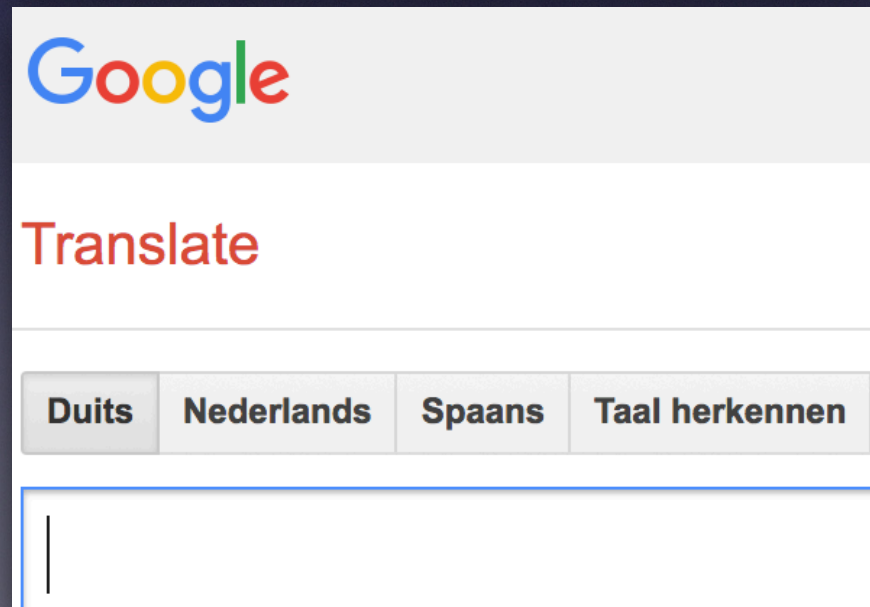

“In literary criticism and analysis, probably 40 or 50 years ago, literary critics were a kind of cultural hero. Now they’re kind of a national joke.”

Steven Pinker (2003)




Yves Saint Laurent presents


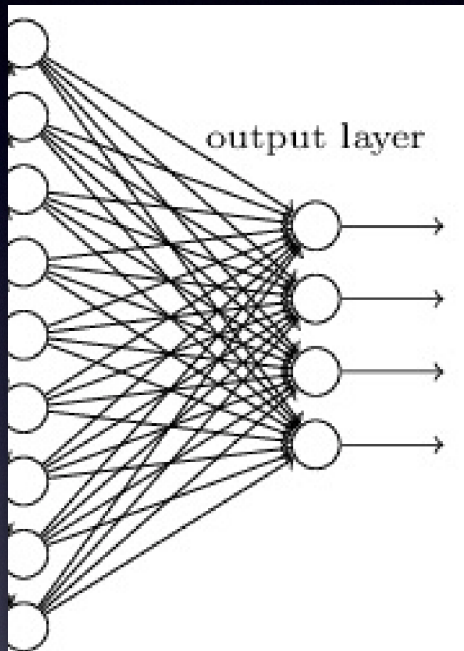
The Hype of Deep Learning



Neural Networks




“Deep” Learning?



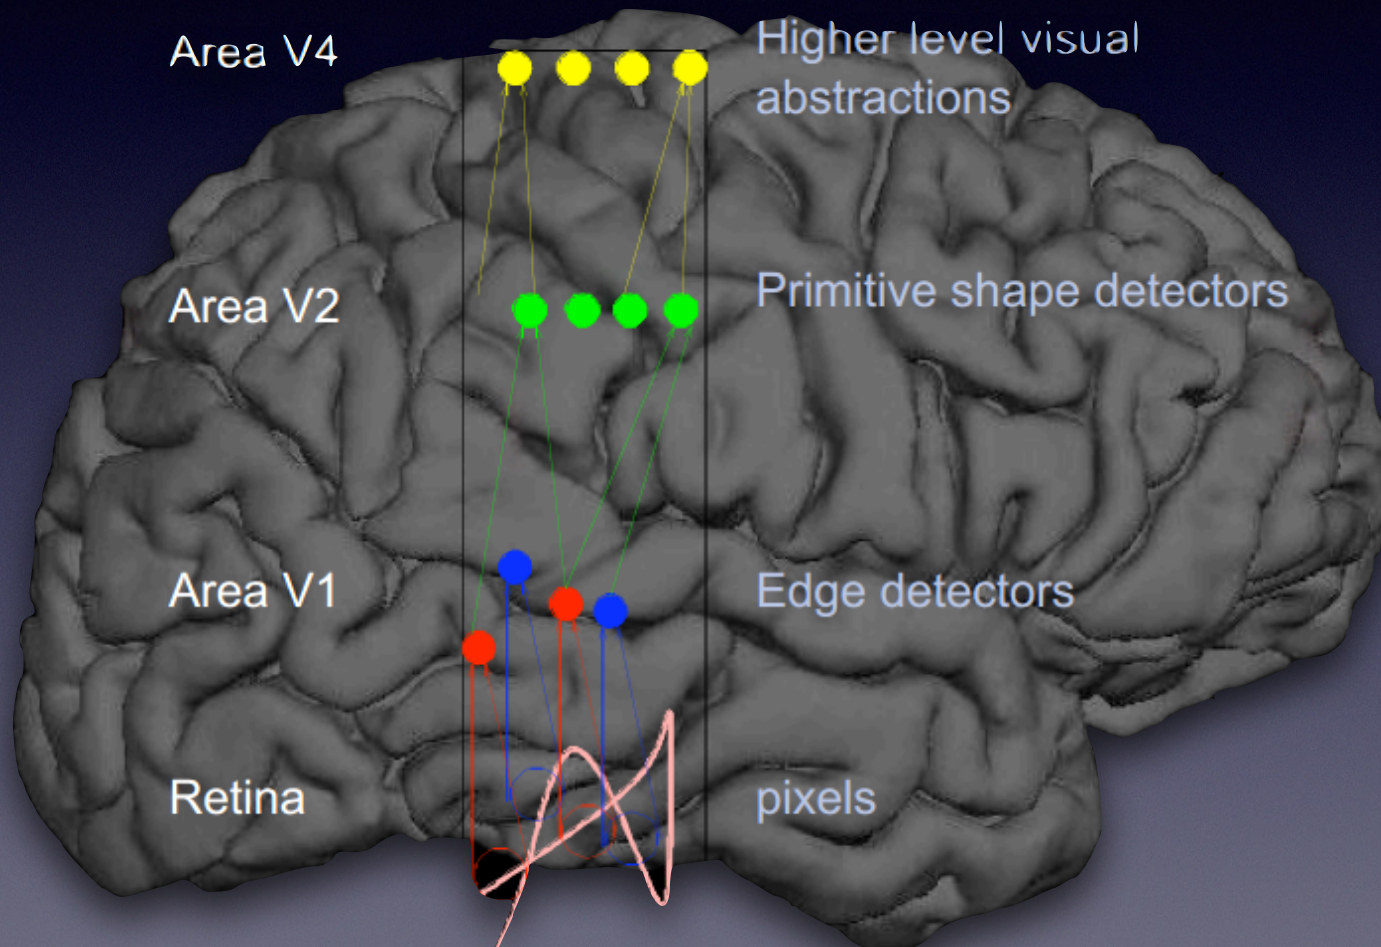
Stack 'hidden' layers between input and output layer

Computer Vision

Importance of layers



Analogies human brain



e.g. [Cahieu et al. 2014]

This repository Search

karpathy / neuraltalk2

Code Issues 1 Pull requests 0 Wiki Pulse Graphs

Watch 11 Unstar 185 Fork 19

```
th /Users/kyle/Documents/Learning/neuraltalk2 — luajit -e package.path="/Users/ky...
a man in a hat and glasses is holding a cell phone
a man in a hat and glasses is holding a cell phone
a man in a hat and glasses is holding a cell phone
a man in a hat and glasses is holding a cell phone
```

Efficient Image Captioning code in Torch, runs on GPU

15 commits 1 branch 0 releases 2 contributors

Branch: master

binoculars coco

Search

- binoculars
- biomf-gpu
- caffe
- cgt
- chainer
- chainer-char-rnn
- char-rnn
- cnn-vis
- cudaarray
- deeppose
- deepdream
- DeepDream.torch
- DeepLearningTutorials
- deeppose
- deeppy
- EmbeddingScripts
- GRUV
- hebel
- Kayak
- keras
- CameraImage
- coco
- coco-caption
- cv
- data
- description.txt
- eval.lua
- luaosc
- misc
- models
- prepro.py
- README.md
- test_language_model.lua
- train.lua
- vis

Application - 17.2 MB

Created Today, 14:39

Modified Today, 15:57

Last opened Today, 16:04

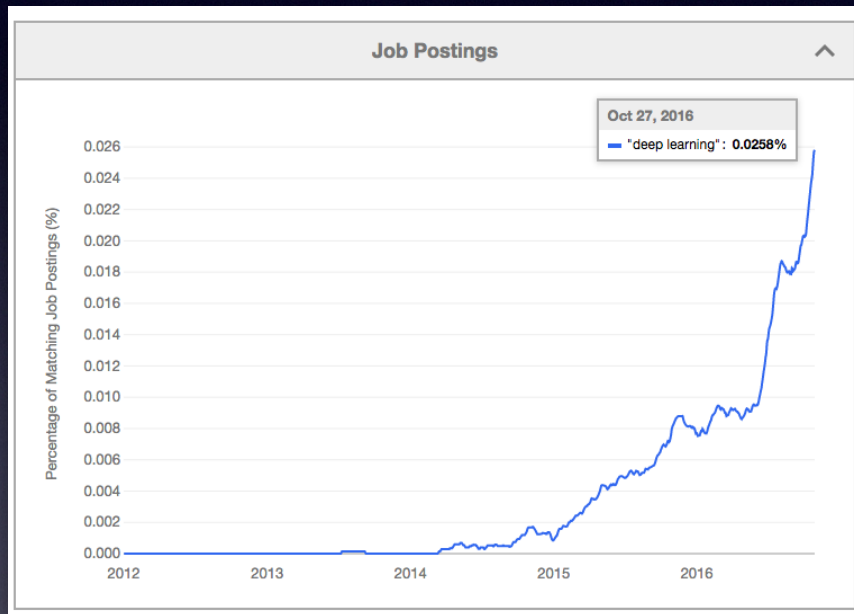
Version --

[Add Tags...](#)

train.lua cleaning up the main directory

README.md

The field



- Almost dead in 90s...
- Now massive growth
- Academia and industry
- Insane economic potential



Geoffrey Hinton
(Toronto, Google)

UK



Yann LeCun
(New York, Facebook)

FR



Yoshua Bengio
(Montreal)

FR

- Three influential tenors
- Advocates of open data, open science, open software
- Indirect European impact via Enlightenment ideas?

French Enlightenment Ideas



Yann LeCun

November 14, 2015 · 🌐



Dear Friends, thank you for all the prayers and and all the #PrayForParis. We know you mean well.

But please be aware that many French people are pretty mad at everything religious right now.

France is one of the least religious countries in the world. A relatively small number of French people believe in god, even fewer believe in the power of prayers.

Many see religion as the most powerful means to indoctrinate people to commit atrocities of the kind we just witnessed.


France has a long history of freeing itself and its citizens from the dangers of religious indoctrination, from the enlightenments, to the French revolution, to the principle of "laïcité" and the separation of church and state, to the ban of overt signs of religion in public schools.

Throughout its history, France has suffered through religious inquisition, religious wars and massacres, and the persecution of religious minorities. This is why the French are so attached to "la liberté" and so intolerant of indoctrination, particularly religious indoctrination.

The battle between rational thought and religious obscurantism, between humanism and religious "morals", between freedom of thought and indoctrination is one that France is very familiar with.

France is used to fight this battle and win.

The French will never enter a battle that pits one religion against another. But they will unite to fight against medieval obscurantism.



“We have incredible people in FAIR who are making significant progress in AI, but to really move the bar it’s equally as important to be outward focused. To push the envelope, push the science and technology forward, we must be actively engaged with the research community. We publish a lot of things we do, distribute a lot of code on open-source, and engage deeply with academia to drive the progress.”

— Yann LeCun, Director of AI Research

Fundamental idea that the accessibility of open data to any citizen is crucial to progress through participation. Anyone should be able to participate.
Active engagement and participation as key terms.



Democracy

- How should AI research and its applications, at the institutional level, be controlled?
- In what areas would this be most pertinent?
- Who should decide, and according to which modalities, the norms and moral values determining this control?
- Who should establish ethical guidelines for self-driving cars?
- Should ethical labeling that respects certain standards be developed for AI, websites and businesses?

Proposed principle:




The development of AI should promote informed participation in public life, cooperation and democratic debate.

Labeled data

- Importance of labeled data
- Supervised learning
- $Y = \text{target}$ (cf. YSL!)
- E.g. ImageNet (competition):
 - ~1,000 per “concept”
 - hand-labeled
- Expensive, difficult to sustain


IMAGENET 14,197,122 images, 21841 synsets indexed

Your query "dog" matches 61 synsets.

	<p>Synset: dog food Definition: food prepared for dogs. <i>Popularity percentile:</i> 91% <i>Depth in WordNet:</i> 6</p>
	<p>Synset: hotdog, hot dog, red hot Definition: a frankfurter served hot on a <i>Popularity percentile:</i> 91% <i>Depth in WordNet:</i> 8</p>
	<p>Synset: water dog Definition: a dog accustomed to water <i>Popularity percentile:</i> 90% <i>Depth in WordNet:</i> 10</p>

Autoencoder

- Need for unsupervised learning
- “Diabolo” principle
- Have network reconstruct data
- Learn “lossy” compression
- Limited capacity enforces generalisation
- Importance: unsupervised (“free” like in “free beer”)



“Cat Paper” [Quoc et al. 2012]




Face detector



Human body detector




Cat detector




10 Million 200x200 images
from YouTube (1 week)
Visualize what individual neurons
“see”

Sampling

- Learn what typical data looks like (cf. cats)
- Through sampling from the bottleneck layer
- Variational Autoencoding
- Create “new” images...
- Generate “**synthetic**” data






Synthetic digits (MNIST)

[http://
blog.fastforwardlabs.com/
2016/08/22/under-the-
hood-of-the-variational-
autoencoder-in.html](http://blog.fastforwardlabs.com/2016/08/22/under-the-hood-of-the-variational-autoencoder-in.html)

Encourage network to generate data that looks and feel “real”, genuine?

“Generative Adversarial Nets” (2014), Goodfellow et al.



"Generative Adversarial Networks is the **most interesting idea in the last ten years** in machine learning."


Yann LeCun, Director, Facebook AI



Two competing networks



D: Detective




R: Real Data



G: Generator (Forger)



I: Input for Generator



PROGRESSIVE GROWING OF GANs FOR IMPROVED QUALITY, STABILITY, AND VARIATION

Submitted to ICLR 2018

Karras et al. [2017]

That's really cool!

But...

What should we do with these models?

How should we evaluate this?

“Do not ask what GANs can do for you,
ask what you can do for GANs” [?]

What I cannot create,
I do not understand.

Know how to solve every
problem that has been solved

Why cost \times \$0.01 .PC

TO LEARN:

Bethe Ansatz Probs.

Kondo \rightarrow

2-D Hall

accel. Temp


Non linear Classical Hydro



$$\textcircled{A} f = U(r, a)$$

$$g = 4(r - \bar{r}) u(r, \bar{r})$$

$$\textcircled{B} f = 2|K \cdot a| (u \cdot a)$$



$$\textcircled{D} f = \sqrt{|K \cdot a|} (u \cdot a)$$

“What I cannot create, I do not understand.”
[R. Feynman]

Can I claim to understand literature,
if I am unable to generate it?

Idea of generative models is gaining traction elsewhere but what about
literature?

GANs don't work (yet?) for text...

Medium



Yoav Goldberg [Follow](#)

Senior Lecturer at Bar Ilan University. Working on NLP. Recently with Neural Nets. Published a book ...

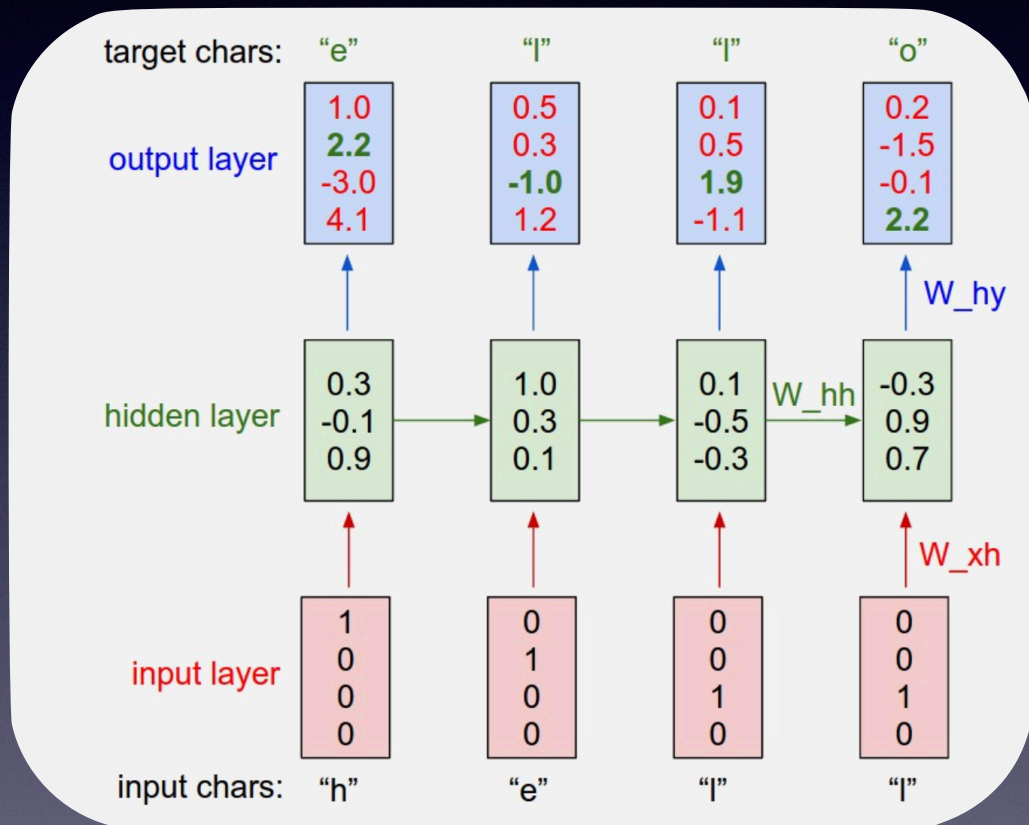
Jun 9 · 14 min read

An Adversarial Review of “Adversarial Generation of Natural Language”

Or, for fucks sake, DL people, leave language alone and stop saying you solve it.

[edit: some people commented that they don't like the us-vs-them tone and that “deep learning people” can—and some indeed do—do good NLP work. To be clear: I fully agree. #NotAllDeepLearners]

Text Generation



"The Unreasonable Effectiveness of Recurrent Neural Networks"

- Language model
- Predict next letter, on the basis of previous 256 letters
- Recurrent neural network
- Long-Short Term Memory (LSTM)
- Capture long-range dependencies

PANDARUS:

Alas, I think he shall be come approached and the day
When little strain would be attain'd into being never fed,
And who is but a chain and subjects of his death,
I should not sleep.

Second Senator:

They are away this miseries, produced upon my soul,
Breaking and strongly should be buried, when I perish
The earth and thoughts of many states.

DUKE VINCENTIO:

Well, your wit is in the care of side and that.

Second Lord:

They would be ruled after this chamber, and
my fair nues begun out of the fact, to be conveyed,
Whose noble souls I'll have the heart of the wars.

Clown:

Come, sir, I will make did behold your worship.

VIOLA:

I'll drink it.

Replies:

- [Yoav Goldberg](#) compared these RNN results to [n-gram maximum likelihood \(counting\) baseline](#)
- [@nylk](#) trained char-rnn on [cooking recipes](#). They look great!
- [@MrChrisJohnson](#) trained char-rnn on Eminem lyrics and then synthesized a rap song with robotic voice reading it out. Hilarious :)
- [@samim](#) trained char-rnn on [Obama Speeches](#). They look fun!
- [João Felipe](#) trained char-rnn irish folk music and [sampled music](#)
- [Bob Sturm](#) also trained char-rnn on [music in ABC notation](#)
- [RNN Bible bot](#) by [Maximilien](#)
- [Learning Holiness](#) learning the Bible
- [Terminal.com snapshot](#) that has char-rnn set up and ready to go in a browser-based virtual machine (thanks [@samim](#))

Cell that turns on inside quotes:

"You mean to imply that I have nothing to eat out of.... On the contrary, I can supply you with everything even if you want to give dinner parties," warmly replied Chichagov, who tried by every word he spoke to prove his own rectitude and therefore imagined Kutuzov to be animated by the same desire.

Kutuzov, shrugging his shoulders, replied with his subtle penetrating smile: "I meant merely to say what I said."

Cell that turns on inside comments and quotes:

```
/* Duplicate LSM field information. The lsm_rule is opaque, so
 * re-initialized. */
static inline int audit_dupe_lsm_field(struct audit_field *df,
                                     struct audit_field *sf)
{
    int ret = 0;
    char *lsm_str;
    /* our own copy of lsm_str */
    lsm_str = kstrdup(sf->lsm_str, GFP_KERNEL);
    if (unlikely(!lsm_str))
        return -ENOMEM;
    df->lsm_str = lsm_str;
    /* our own (refreshed) copy of lsm_rule */
    ret = security_audit_rule_init(df->type, df->op, df->lsm_str,
                                  (void **)&df->lsm_rule);
    /* Keep currently invalid fields around in case they
     * become valid after a policy reload. */
    if (ret == -EINVAL) {
        pr_warn("audit rule for LSM '%s' is invalid\n",
              df->lsm_str);
        ret = 0;
    }
    return ret;
}
```

One example application of generative models: work with



Folgert Karsdorp



Enrique Manjavacas



Ben Burtenshaw

CPNB

- Collective for the Promotion of the Dutch Book (CPNB)
- Dutch Book Week
- The theme of robots, based on the novel “I, Robot” by Isaac Asimov (in translation)
- Asked us to make a robot write an entire novel...




Ronald Giphart

Co-create 10th story to *Ik, Robot*



Method

- Train Recurrent Neural Net on
 - 4,392 novels by 1,600 authors
- Literary “autocompletion” (Roemmele 2015)
- Different “voices”: mixing e.g. Kristien Hemmerechts with Asimov
- Co-creative system



Example output for different temperatures

τ \ seed	Mijn vrouw is dood en al begraven.	My wife is dead and already buried.
0.1	<i>De volgende dag werd er een bericht van de politie gevonden. De verklaring van de politie was een van de belangrijkste verklaringen van de man die het verhaal van de politie had gehoord.</i>	<i>The next day, a police message was found. The police's statement was one of the most important statements of the man who heard the story by the police.</i>
0.3	<i>De dokter keek op zijn horloge. Ze was er zeker van dat hij die dag niet meer zou komen. Hij was niet gek. Hij was niet van plan om zijn dochter te vertellen dat hij haar een keer had gezien.</i>	<i>The doctor looked at his watch. She was sure he wouldn't come anymore that day. He wasn't crazy. He did not intend to tell his daughter he'd seen her once.</i>
0.5	<i>De tranen weerklonken over haar gezicht en voelde ze in haar keel. Ze was er niet. Ze was er zeker van dat ze alleen maar een affaire had gehad met de meisjes.</i>	<i>The tears resounded over her face and she felt them in her throat. She wasn't there. She was sure she had only had an affair with the girls.</i>
0.7	<i>Ze zegt dat het weer in San Sebastian naar Californië haar deel van de wereld is. Ze moet van de wereld naar de hemel hebben gelopen.</i>	<i>She says that the weather in San Sebastian to California is her part of the world. She must have walked from the world to heaven.</i>
0.9	<i>Lucas kon niet met zichzelf trouwen en toen hij achteloos met zijn gebaren zat te praten maakte hij deel uit van de lessen van de waarheid.</i>	<i>Lucas couldn't marry himself, and when he spoke painlessly with his gestures, he was part of the lessons of truth.</i>

Table 1: Example of our current NLG system with translation seeded by "Mijn vrouw is dood en al begraven" (My wife is dead and already buried) for different temperature τ values.

AsiBot Interface

The screenshot shows the AsiBot web interface. At the top, there is a browser window with the address bar showing 'localhost:5000/index'. Below the browser, the interface has a header with 'AsiBot' on the left, a document dropdown menu labeled 'My document' with file and trash icons, and an 'admin' button with a right-pointing arrow on the right.

Below the header, there are two sliders: 'Creativiteit' set to 0.45 and 'Lengte' set to 100. To the right of these sliders are four colored buttons: 'A' (green), 'IA&RG' (orange), 'IA' (blue), and 'IA&KH' (pink).


On the left side, there is a vertical list of three buttons: 'B', 'I', and 'U'. The main area contains a large text box with the Dutch text: 'Mijn vrouw is dood en al begraven. Ik ben alleen in huis, alleen met de twee meiden. Dus ben ik weer vrij; maar wat baat me nu die vrijheid?'. A cursor is visible at the end of the text.

At the bottom, there is a panel with a dropdown menu on the left, a 'Clear All' button, and a refresh icon on the right.

Genetic edition/analysis

(*ongoing*)

- With Dirk van Hulle, Wouter Haverals and Vincent Neyt (UA)
- Turn into a “genetic story”
 - Effect of editorial revisions
 - Tenses changed in one go
 - Use of model (gets more “boring”)
 - Interaction with different “voices”
- Authorship attribution



Number of edits per minute

Evaluation?

- Hard to evaluate using numbers: is predictable good?
- Output is linguistically “perfect” (long-term memory)
- Semantics are hard to keep in check over longer sequences (+256 characters): very local consciousness
- Longer narratives remain elusive: mainly useful as a source of inspiration (serendipity)

Writing competition

With *Algemeen Dagblad* (NL)

AsiBot

Schrijf samen met een robot

Doe mee aan de [schrijfwedstrijd](#)



Ik ga akkoord met de [gebruiksvoorwaarden](#).

start

Begin hier met schrijven...

Ik, Robot

I. Nescimov

G. Revimov

A. Dante

R. Giphart

I. Asimov

K. Hembot

I. Robhart

Creativiteit

0.35

Aantal letters

100

genereer suggesties

Het was donker buiten en de sterren waren lichter geworden. Er was geen enkel teken van leven te zien in de stad. Jane keek op haar horloge en zag dat het al tien uur was. De geur van haar zachte parfum was verdwenen en ze voelde een koude rilling over haar rug lopen.

'Heb je al gezegd waarom je niet mee mocht naar de stad?'

'Nee.' Zei ze kort. Ze keek op naar de gestalte in de deuropening.

'Hoe lang sta je daar al?'

De gestalte stapte uit de deuropening naar Jane. Ze stapte naar voren en greep de hand van de robot beet.

By Lauren Fonteyn

Conclusion

AI's like AsiBot offer *tools that we don't know how to use yet*. The electric guitar was there before Jimi Hendrickx but he's was among the first to realise its potential.



Selected References

- Manjavacas e.a., Synthetic Literature: Writing Science Fiction in a Co-Creative Process. *Computational Creativity in Natural Language Generation (CC-NLG 2017)*, pp. 29-37
- Cahieu et al., Deep Neural Networks Rival the Representation of Primate IT Cortex [...]. *Plos Comp Biol* (2014).
- LeCun Y.; Bengio Y. & Hinton G., Deep Learning. *Nature* (2015).
- Quoc Le et al. Building high-level features using large scale unsupervised learning. *ICML* (2012).
- Bod, *A New History of the Humanities* [...] (2013).
- Kingma ea., Auto-Encoding Variational Bayes. *ICLR* (2014).
- Roemmele ea. 2015. Creative help: a story writing assistant. *International Conference on Interactive Digital Storytelling* (2015).