

**A 23 MW data centre is all
you need**

Samuel Albanie, Dylan Campbell, João F. Henriques

SIGBOVIK 2022

Motivation

Accurate forecasting is valuable


Battle outcomes



Clay model of a sheep Liver tablet
(Babylonian Soothsayers, 1900 BC)

Fate of Empire?

Precipitation nowcasting




SoTA forecasting algorithm:
Are you in or near Wales?
Then it will probably rain

Anorak for dog walk?

Future forecasting frameworks


Passive



The Singularity is Near
Kurzweil, 2005

Lacks a rigorous theoretical foundation

Active



Inventing the Future
Gabor, 1967

Exhausting


Image credits: Ray Kurzweil and Kurzweil Technologies, Inc. CC BY 1.0, <https://commons.wikimedia.org/w/index.php?curid=1273707>

Cover art, *Inventing The Future* by Dennis Gabor, 1967, A Pelican book

Our (passive) approach


Central Limit Theorem to the rescue

Traditional tech forecast (Kurzweil)



- Data-driven? ✓
- Eerily accurate so far? ✓
- Number go up? ✓
- Statistical rigour? ✗

Our predicted trajectory



- Uses Central Limit Theorem? ✓
- Subject to gravity? ✓
- Visually soothing ✓
- Finite maximum? ✓

Peak: 3:07 am, 20th July, 2032

The road ahead

Causes


Plausible Causes

- A butterfly wing flap in downtown Tirana, 1658
- The premature death of Poincaré, 1912
- The inexplicable density of roundabouts in Milton Keynes, 2022

Clues from science fiction

To escape Slack notifications, many technologists will turn to **soma** (Huxley, 1933) and **wireheading** (Niven, 1969)

This will slow progress, but some have already submitted to **focus mode** (hungry, but immune to such sirens).



We ask a Transformer...

We next trained a generative Transformer on temporally curated Wikipedia articles.

The model predicts that by 2031, humanity has achieved a new Gini coefficient SoTA:

- **99.99999%** of all wealth will belong to 3 humans and **Molly**.

On December 31st, 2031: anonymous* game designer releases self-learning reboot of Flappy Bird 🐦


Designed to maximise user engagement, it explores **reward hacking**.

Player cannot stop playing. The 3 richest humans are incapacitated.

Capital flows grind to a halt. Silicon foundries fall silent.

The world turns to Molly. She wants to help, but only after walkies. And she might have forgotten her BTC passwoof.

Wikipedia and StackOverflow servers are powered down. No-one can remember the command to undo recent git commits.



Molly

A critical point

Implications

A chance for a positive feedback loop

Big language models open up an **attack surface** on English spelling.


Our plan: ensure 51% of internet text (used to train future spell-checkers) follows our spelling preferences.

Propose: **HMGAN** inspired by the blog posts of Elizabeth II to synthesise the required volume of prose

Assumptions:

- Global data centres will use **227 TWh** by 2032.
- By 2032, most internet content will be unintentional bot-to-bot conversation between chatbot sales assistants.
- Training corpora for spell-checkers will use their transcripts

Conclusion: a 23 MW data centre will allow HMGAN at the required volume to out chat the chat-bots.*



How can we design a timeless spell-checker?

Propose: a **diachronic spell-checker** trained on Twitter.

Unforeseen side-effect: the conversion of some words to coloured boxes (possibly due to data-augmentation) 

Timeless variant: predicts **future** language and events.

Experiments: British Bin Colouring Problem (BBCP)

Problem formulation

On Wednesday evenings, wheel out the colour of bin that causes *maximum mischief* to your neighbours.

Wheeled out bins of the wrong colour **will not be collected under any circumstances.**

You have three choices:

Black - unfiltered

Blue - recycling

Green - garden waste



To avoid social tension, most neighbours will copy their neighbours bin colour, rather than check for the the local government website.

You may collaborate with others, but you must account for upstanding citizens (who will wheel out the correct bin colour).

The problem might be NP-Hard. It is definitely environmentally significant.

Implementation: Grantchester

Use HMGAN to craft royal entreatments to wheel out **blue bins** on a **green bin Wednesday**.
Sent leaflets to residents (Sobol sequence).
We wheeled out our own blue bin and waited.

Analysis

Stochastic upstanding citizen locations and wheel-out race conditions complicated our analysis.
Unclear whether a 51% bin colour majority would achieve our desired ends.
We adopted a strategy of hoping it would work.

Results: Unpromising

In our eagerness, failed to wait until 27th March
Missed the transition to Daylight saving time.
Too dark for our neighbours to determine our bin colour (so they were uninfluenced)
They did not take kindly to unsolicited leaflets

Experiments: Qualitative results

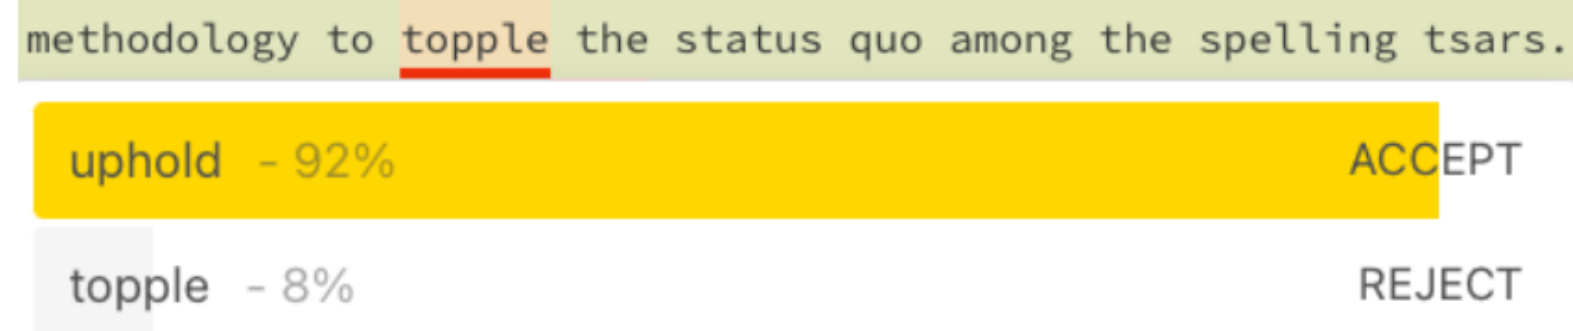
Evaluation of predictive spell-checker

Backtesting on historical data, events and spellings predicted/ caused by our model include **quarantinis** but not **maskne**.

The model insists on auto-correcting our use of 'colour' to 'color'.

Plausible causes include:

- The challenge of overturning the spelling status quo.



- The difficulty of controlling large language models.
- We still don't really understand what `.detach()` does in PyTorch.

Conclusions

Contributions

Rigorous statistical prediction of the future.

A precise time and date of the peak

3:07 am, 20th July 2032

Tools for long-term linguistic 'lock in'.

Future work

New **time** and **date** scheme based on standard deviations from the Gaussian technology curve.

Further research: determine if the two-to-one mapping from years to standard deviations will be problematic.

Thank you for your attention

Code coming soon: <https://github.com/albanie/A-23MW-data-centre-is-all-you-need>

We thank, in particular, A. Sophia Koepke for her wisdom.