

[Children's cancer charity receives nuclear cash boost](#)

Children's cancer charity Danica's Dream Fund has received a cheque for £2,671 from LLW Repository Ltd (LLWR).

The LLWR workforce chose the Egremont based organisation as its charity of the year and have been busy initiating fundraisers, including a sponsored walk up Scafell Pike, England's highest peak, a sponsored shave and cake sales.

Danica Maxwell passed away in 2010 aged just 15, following three brain tumours, and Dad John, who works on the LLWR Site, went on to launch the charity in her name.

John said: "This donation is excellent and will help families going through cancer. We use donations like this to buy gifts for children with cancer, such as a laptop or we fund a family break, to help them get back to normality."

Amanda Raynor, who also works on the LLWR Site, took part in the Scafell Pike walk and said: "I'm from Egremont and I know what a great charity this is. I'm close to John and I know what he's been through. What he does for children is great.

"Scafell Pike was hard work but thoroughly enjoyable. Team spirit got us through."

The LLWR workforce is continuing to raise funds for the charity which nationally has now raised over £200,000 for those aged up to 23 living with cancer, and their families.

[Graham wins £8 million LLW Repository Ltd contract](#)

An £8 million Enabling Works contract for LLWR's Repository Development Programme (RDP) has been awarded to Graham.

The two year contract will deliver key infrastructure elements on the Repository Site, such as construction compounds, haul roads, material stockpile areas and drainage lagoons, that will allow the final cap over waste disposal vaults and legacy trenches to be constructed under RDP, LLWR's biggest ever programme, spanning more than 30 years.

The Enabling Works contract, which commences in January 2020, will see LLWR continuing its relationship with Graham following completion of a range of infrastructure and security projects delivered under their Engineering Design and Construction Framework.

Jonathan Evans, Head of Programmes and Projects at LLWR, congratulated Graham's on the contract award and said construction activities were expected to start in April 2020.

He added: "It is worth recognising that whilst we will be continuing to work with Graham, the Enabling Works contract is new and will operate entirely separately from the existing Framework with a defined scope and a different commercial delivery model."

Activities are well underway on the Repository Site to provide the space needed to import the large volumes of materials that will be required to construct the cap.

Steve Beattie, Contracts Director, Civil Engineering, at Graham, said: "We are delighted to have been selected by LLWR for the Enabling Works contract, which is testament to the successful delivery of the joint objectives and project deliverables that shaped our Engineering Design and Construction Framework portfolio over the last five years.

"It continues our established relationship with LLWR and is a positive reflection of the strength and professionalism of our site team. We place a strong emphasis on repeat business, and this contract demonstrates the impact that our industry leading processes have made in the development of a mutually beneficial long term partnership."

Graham is a privately owned company with a 200 year history that specialises in building, civil engineering, interior fit out, facilities management and investments, across the UK and Ireland.

[Women in Nuclear UK appoints David Peattie as new Patron](#)



David Peattie, CEO of the Nuclear Decommissioning Authority

WiN UK's mission is to address the industry's gender balance, improving the representation of women in leadership and engaging with the public on nuclear issues.

The role of Patron involves helping to gain government and industry support, and raising the profile of WiN UK and its objectives.

David Peattie said:

I'm honoured to take on this role with WiN UK because I am passionate about the positive impact of diverse teams in business. I believe WiN UK has an important part to play in making change happen.

I'd like to thank Adrienne Kelbie, CEO of the Office for Nuclear Regulation, who has achieved so much during her time as WiN UK Patron – driving many businesses and individuals to make a difference across the industry.

Jack Gritt, President of WiN UK, said:

David Peattie's appointment as our Patron is great news for us and reflects, yet again, the importance of gender balance in our industry... We're energised by his positive view of the future.

Published 20 December 2019

Plating up for Christmas

With the countdown to Christmas Day looming, DVLA is today reminding people still looking for a gift for their loved one that it's never too late to pick up the perfect personalised registration online.

Last year over 600 people bought a personalised registration from DVLA between Christmas Eve and Christmas Day, with 90 of those purchasing on Christmas Day before 9am.

DVLA sold more than 32,000 personalised registrations between 1 and 25 December 2018.

DVLA Chief Executive Julie Lennard said:

With millions of registrations available to suit all tastes, and with prices starting at £250, there's still time to find that perfect present – even if it is on Christmas morning.

A personalised registration with the person's name may be the ideal gift as 3 out of 4 motorists said they would prefer a registration that bears their name.

The top 3 search terms on the DVLA Personalised Registrations website are currently 'Sam', 'Dan' and 'Tom', but no matter what a person is called they can [search for their ideal personalised registration online](#).

Notes to editors:

90 purchases were made on Christmas Day 2018 between midnight and 8:59am and 544 purchases were made on Christmas Eve 2018 between midnight and 11:59pm.

The 3 in 4 motorists saying they would prefer a personalised registration bearing their name is taken from responses to a DVLA survey.

How GDS used data to improve content and user journeys on GOV.UK

Summary

Project timing – May – August 2019

Team – Data science team and developers from the Government Digital Service (GDS)

Tools – machine learning, network analysis and creating semantic vectors using [Google's universal sentence encoder](#)

Objective

The GDS team working on this project wanted to improve navigation on GOV.UK so it was easier for users to find what they needed. The team aimed to:

- reduce duplicate or overlapping pieces of guidance
- group related content with links

To do this the team used data to:

- identify similar or duplicate GOV.UK content
- use an automated process to generate related content links for most GOV.UK pages

Background

[GOV.UK](#) is the main website for government. It enables users to interact with government and find services, guidance and news from all government departments. There are over 400,000 unique pieces of content on GOV.UK and statistics showed that over 6 months, 700 pieces of content were published per week.

This amount of content, and high publishing rate, leads to duplicate content and makes it hard for users to find what they need.

GOV.UK provides different types of navigational help, including:

- a search function
- [breadcrumb trails](#) that show a list of nested web pages
- related links, which are shown to the right-hand side of content on a page

Before this project, only around 2% of GOV.UK's content (about 8,000 pages) had related links.

Investigating and trying solutions

To improve users' experience, we investigated:

- semantic vectors to group content and find similar or duplicate entries
- network analysis to align how GOV.UK pages are connected with hyperlinks (the structural network) with the way users travel between pages (the user generated functional network)
- machine learning to auto generate links on GOV.UK pages

Semantic vectors

To reduce the amount of similar or duplicate content on GOV.UK we first needed to find this content. We wrote a [blog post on how we tested and created semantic vectors](#) to plot the content with similarities and duplications.

After testing various tools we chose Google's [universal sentence encoder](#) to turn content into semantic vectors. By using the universal sentence encoder we could tell publishers where their content overlapped with pre-existing content. Publishers could use the information at the publication stage to prevent duplicate content going live, or retrospectively to clean up existing duplicate content.

We could also use semantic vectors to make decisions on changes to GOV.UK's taxonomy. The taxonomy is structured into topics and publishers tag content with the most appropriate topics.

Network analysis

We wrote [a blog post on how we used network analysis](#) to discover and assess:

- whether all content areas of GOV.UK are accessible and have links
- individual pages' position and significance within the overall network

The team identified hub pages which have more connections than others. To investigate these, we calculated network properties, including:

- network density
- connectedness
- link distribution
- centrality measures

The team created the [GOV.UK Network Data Pipeline](#) using Python to automatically extract user journeys from BigQuery – the database that stores our Google Analytics data. We aggregated these journeys for a specific time period to learn how users interact with the site. The result of these journeys was a functional network, or map of connected content.

We compared the functional network to how GOV.UK was structured, and used the results to inform changes to that structure and improve navigation for users.

Machine learning and A/B testing

When publishers create a new piece of content they add links to other related pages. Navigational links that facilitate browsing are automatically linked to the new content item. This creates GOV.UK's structural network which consists of approximately 350,000 links. We wrote [a blog post about using machine learning and A/B testing](#) to check the structure of GOV.UK content and automatically generate related content links.

After considering different algorithms, we decided to implement the [node2vec](#)

[algorithm](#) to create links and display them on the right-hand side of GOV.UK pages. These would help people find the content they needed more easily.

We tested the user journeys through the links using 2 [null hypotheses](#):

1. There is no significant difference in the proportion of journeys using at least one related link. If this is true, then people are not clicking our links, probably because they do not find them useful or interesting.
2. There is no significant difference in the proportion of journeys that use internal search. If this is true, then we did not provide people with relevant links to click on so they still use internal search to navigate.

We used an [A/B test](#) to randomly assign users to one of two possible versions of the site. Half of the users were directed to the control site (A) where only a small percentage of the pages had links which were added by the publishers. The other half of the users were directed to the version of the site (B) with the algorithmically generated related links. To speed up our analysis [we wrote a software package](#) that allowed us to run routine analysis, as soon as each experiment completed.

The results showed that both of our null hypotheses were false. The first implied that users found the related links interesting and/or relevant so they clicked them. User journeys showed an improvement for more than 10,000 users per day. The second result implied that there's a potential reduction in internal search by about 20% to 40%.

Applying our solutions

We wrote [a blog post on how we put our findings into production](#). We used [Jupyter](#) notebooks for rapid development and prototyping which allowed us to code, document and [publish data in a single place](#).

To automate the process of generating related links, we decided to modularise the code from the Jupyter notebooks and write it in object-oriented Python. By using a combination of user journey information and the structure of GOV.UK we could keep adapting the generated links.

To maintain a high quality of related links on GOV.UK, we added a number of steps to our process, including:

- excluding a number of pages that should not have links coming from or going to them
- applying a confidence threshold to the links generated by node2vec which only allowed through the most relevant links
- generating a spreadsheet of the top 200 most popular pages so that our content designers could check the generated links and make changes where necessary

Results

We've been running the process to generate links every 3 weeks. These links are automatically displayed on GOV.UK when publishers have not set their own links, which ensures that users have a way to find the content they're looking for.

We're continuously iterating and refining the related links process, and monitoring results to make sure users' experiences are improving. In the future, we aim to bring the link generation online and use it as part of the publishing process. All of our code is available to [view on GitHub](#).

Automatically generated links will not replace hand-curated links as they will not have the same considered context as those created by a subject matter expert.

Prior to this work, the vast majority of GOV.UK pages did not have curated recommended links. We automated that process, improving thousands of user journeys per day.