

# ICT with Industry 2019

21-25 January @Oort

Under the umbrella of ICT-research Platform Netherlands (IPN) the workshop ICT with Industry was organized by NWO in collaboration with the research schools ASCI, IPA and SIKS. At the same time as part of the economic priority areas (topsectorenbeleid) set by the Dutch government the workshop ICT with Industry provides an excellent platform where scientists and the industrial partners have the opportunity to explore possible collaboration.

The workshop strives for direct and rapid interaction between ICT researchers and industrial partners with the following objectives:

- To stimulate contact between ICT research and industrial R&D, by identifying and analyzing challenging problems and emerging technologies in multidisciplinary teams.
- To obtain creative solutions for industrial problems and to find new approaches that could lead to such solutions.
- To give insight into the wide range of possibilities ICT research offers (primarily to industry, but also to the general public), and thereby enable accelerated innovation.
- To enrich the PhD students' and postdocs' experience in collaborating with industry.

In this edition five case proposals were assessed and assigned to one academic team leader. the format of the workshop programme followed the previous editions and slightly adapted to the wishes of academic team leader and industrial partner resulting from the intakes. The communication with the Lorentz Center went excellent, although more flexibility in communication tools, e.g. website, invitation letter, would be recommended. One of lessons learnt is to have the posters distributed along with the publication of Call for Participation in June.

A total of over 60 participants from industry and academia were working on these cases, of which more than 20% female participants. Participants were affiliated to different Dutch universities, mostly Amsterdam, and a few from US and Hungary.

The cases for this edition were media-oriented and all of them resulted in promising methods and concrete directions to further develop. For future collaborations NWO provides several funding instruments, e.g. KIEM, Take Off, etc. The teams were invited to present a demo and a poster at ICT.OPEN2019. Exposure of the workshop ICT with Industry will be done at ICT.OPEN by a participant's journal of the week made by one participant, Myrthe Reuver (UL).

Besides this short report the Steering Committee together with the teams develop a full report on the cases and the outcome. Please find below the abstracts of all five cases:

KB (National Library of the Netherlands): (Semi-) Automatic Cataloguing of Textual Cultural Heritage Objects

Participants: Martijn Kleppe (company representative), Iris Hendrickx (academic leader), Sara Veldhoen (company representative), Alex Brandsen, Hugo de Vos, Karen Goes, Lin Huang, Hugo Huurdeman, Areumbyeol Kim, Sepideh Mesbah, Myrthe Reuver, Shenghui Wang & Richard Zijdeman

As a National Library, the KB collects and stores all publications of and about the Netherlands. These publications are described by manually assigning keywords by KB's annotators. While for centuries, the KB stored physical objects, recently digital publications are also being collected. Since these digital publications are available as full text, the KB now wants to explore the possibilities to automate the labelling of publications. During the ICT with Industry Workshop 2019, the main research question

was: “*Can we automatically label scientific texts with relevant key words?*”.

For this case study we used a set of dissertations manually tagged with keywords from the Brinkman thesaurus (‘Brinkeys’), and a set of metadata scraped from university websites, combined with abstracts and full text where available. The data is sparse and heterogeneous, and partly English and partly Dutch, which complicates any classification attempt as the Brinkman keywords list is completely in Dutch.

We investigated a range of approaches, for example testing existing tools, matching (translated) keywords in titles, and applying multilingual semantic embeddings. These initial experiments show promising results and possibilities for future support tools of the KB annotators. A list of suggested keywords can then be used by KB employees to quickly select the relevant keywords, which could provide consistency and a significant time saving in the day to day work of the annotators.

**OrangeNXT (ICTgroup): Detecting and classifying damage to traffic signs from images**

*Participants:* Jeroen Delcour (company representative), Bernardo Marques (company representative), Efstratios Gavves (academic leader), Sadaf Gulshad, Masoud Mazloom, David Ruhe, Zenglin Shi, Ivan Sosnovik, Gjorgji Strezoski, Tom Viering, Mingshu Wang

In this study, we worked for feasible solution for a real life problem of classification of damages on road traffic signs acquired by a moving panoramic cameras. As this work was the first attempt to create a dataset specific to Dutch traffic signs, there were some major challenges to solve in data preparation and annotations. Since there were not sufficient properly annotated signs, we augmented some samples with artificial graffiti and stickers for training. Some annotations made per stacked signs were manually classified before to the training stage.

Traffic sign detection from multi view panoramic images were successfully done via Faster r cnn model. Detection of type of damage and its severity for replacement were not defined properly. Therefore we focused on a more obvious damages like sticker and graffiti. There were dominantly more undamaged images in the dataset which leads overfitting. To overcome this imbalanced dataset problem, we customized our model with additional methods like cut-out and focal loss.

An efficient model for detection of damages was also important for accuracy. After some experiments with plain convolutional networks, we implemented customized version of Resnet18.

Finally we defined proper metrics for the evaluation of the results and developed a future roadmap for improving our results.

**FD Mediagroep: Opening the black box of user profiles in content-based recommender systems**

*Participants:* David Graus (company representative), Feng Lu (company representative), Dung Manh Chu (company representative), Anca Dumitrache (company representative), Nava Tintarev (academic leader), Dimitrios Bountouridis, Jaron Harambam, Domokos Kelen, Felicia Loecherbach, Mykola Makhortykh, Shabnam Najafian, Emily Sullivan, Marieke Vinkenoog, Daricia Wilkinson.

Personalized experiences powered by recommender systems have, after years of being a mostly academic endeavor, finally permeated our daily lives. Whether it is through personalized recommendations in web shops (e.g., Amazon), personalized media consumption (e.g., Spotify, Netflix), search engines, or virtual assistants (e.g., Apple Siri, Amazon Echo). However, driven by data breach incidents and ad-driven business models, we’ve recently seen a rise in distrust and skepticism around the collection of personal data—a requirement for recommender systems. In addition, the GDPR has generated an increased interest in aspects of explainability and transparency of black-box machine learning algorithms and models.

In this project, we aimed to explore methods of explaining one aspect of how our content-based recommender system works: the user profile. More specifically, we aimed to automatically summarize and visualize the recommender system's high dimensional internal representations of users. These profiles are automatically constructed from their reading behavior, by leveraging attributes of items, e.g., topics, entities, and tags.

A multi-disciplinary team of 9 participants tackled this problem of generating explanations for news recommendations. The team consisted of PhD students and postdocs, and even attracted participants who travelled from Hungary and the USA for the event. We developed a conceptual framework where we identified different levels of explanations, and different values that users and FD media may share. Following exploratory data analysis, we developed mockups for the interface for different explanatory values. This led to the development of an interactive explanation interface, as well as a static version for which we planned a user-study.

#### Persgroep: Urgent, or can it wait? Personalising push for Algemeen Dagblad

*Participants:* Marlies van der Wees (company representative), Menno van Zaanen (TU, academic leader), Daniel Berecz, Anne Dirkson, Chris Emmerly, Lieke Gelderloos, Chris Kamphuis, Mozghan Karimi, Radhesh Krishnan Konoth, Tudor Paisa, Gustavo Penha, Susan Vermeer.

The current online publishing of news articles also enables notifying newspaper readers of interesting articles through the so called process of pushing. At the moment, online editors decide which articles to push to readers and it is unclear exactly on which basis this decision happens. During the project, we investigated patterns in millions of news articles, whether they were pushed or not, and their click stream (each click corresponds to someone reading an article). It turns out that predicting whether a new news article should be pushed or not is difficult to model computationally, but the data does contain useful information that can be used to predict the expected number of clicks.

#### RTL Nederland: Captioning News Footage

*Participants:* Daan Odijk (company representative), Maurits van der Goes (company representative), Pascal Mettes (academic leader), Thomas Mensink (academic leader), Arthur Câmara, Yunlu Chen, Shuo Chen, Jiaojiao Zhao, William Thong, Emiel van Miltenburg, Devanshu Arya, Sarah Ibrahim

The RTL Nieuws team produces hours of video content every day. All this raw footage and edited clips are manually archived for future reuse. RTL Nieuws' documentation department spends over 50 hours every week on adding annotations and metadata to all this footage. On the other end, editors search the archive for relevant footage around current news events. We believe that with the right ICT solution we can make this a less labor-intensive process and potentially even increase the retrievability of the content.

We have developed a pipeline to match videos with user queries. This first involved extracting semantics from videos, where we leveraged the state-of-the-art in deep learning and optical character recognition to detect concepts within the videos. For each shot of each video in the database, we extracted information about 1,080 objects, 365 scenes, 102 scene attributes, and 400 action using pre-trained deep networks. We furthermore employed optical character recognition to extract textual information on screen, e.g. from subtitles or information boxes. The video detections were matched with the user queries either using elastic search or through word embeddings. As a final step, we developed an interactive front-end where users can input queries and find to retrieved videos. To enable a quick exploration of the video database, we have added the option to weight the importance of specific detections to the ranking of the videos. This helps with finding a better match between query and video, while also dealing with semantic ambiguity.