1. Automatic Trailer Generation

Company: RTL

RTL Nederland is the largest commercial broadcaster in the Netherlands, with a yearly turnover of around €350m. RTL has the mission to tell unmissable stories that touch heart and mind. With a team of 10 data scientists, RTL works on challenges such as personalization, forecasting and automatic content generation.

Challenge

Every day, RTL produces many hours of video content that aims to touch the heart and mind of viewers. With advanced computer vision technologies, we believe that we can support our creative people, by improving the efficiency of video production and supporting the creative process of promotional content creation.

For ICT with Industry 2022, we pick up the challenge of automatically generating trailers for TV programs and movies. The task is to generate from a full length video a short teaser that increases interest in viewing the content, without spoiling the experience.

We created a dataset of episodes from our daily soap opera GTST with shot-level annotations of recurring video content. Each GTST episode starts with about a minute of recaps from previous episodes and ends with a short preview of the next episode. This recurring content makes for ideal training material in a weakly-supervised learning task for what content is interesting to include in promotional material.

We know from directors that using emotional salient material is an integral part of the production of promotional content, as emotion compels attention and elicits arousal. A better understanding of interesting content would on the one hand facilitate directors in finding better promotional material, and on the other hand unlock an important aspect of the automatic generation of promotional video content.

We are seeking to close the gap in vocabulary between our use cases and automatically generated content. Combining the content and applied expertise of RTL with the academic expertise of the two best computer vision groups in the Netherlands, will make for exactly the right atmosphere for tackling this challenge.

Input for the workshop

Before the workshop, we will invite interested participants for a preparatory afternoon at RTL, giving them a first introduction of the use case, a tour of the RTL and a studio, and an informal chat with some of the intended end users of the solution, our promotional video editors. Should COVID-19 still limit this, we will make this a virtual event with a similar program.

As working with video content is computationally intensive, we will prepare a first baseline approach before the workshop, if possible, with input from the participants. We can provide the participants with a selection of video material, subtitles and the added metadata as training material. We will also look at the aforementioned video summarization task and other publicly available datasets to ensure potential for reuse beyond our use case.

For our daily soap opera GTST we have a video dataset with shot-level annotations of recurring content. This dataset consists of around 1.000 annotated shots and will be available for participants.

Where needed and relevant we can provide participants with the output of video analysis and other general-purpose text, audio and visual content descriptors. For these pre-computed features, we consider ImageNet-trained deep neural net features per frame, face detection bounding boxes, and if there is a face, AffectNet output for the face, and OpenSmile sound features sampled regularly from the content, and NLP features derived from subtitles.

Team Leaders

Academic Team Leader

- Dr. Pascal Mettes, Video & Image Sense lab, Universiteit van Amsterdam

RTL Nederland

- Dr. Daan Odijk, Lead Data Scientist
- Dr. Hendrik Vincent Koops, Data Scientist

Intelligent Systems, Vision Lab, Delft University of Technology

- Dr. Seyran Khademi
- Ombretta Strafforello, MSc

2. The future of virtual racing

Company: Axini

Axini is a leader in the development and application of model-based testing with their Axini Modeling Platform. Axini is investigating how model-based testing can be applied to serious virtual racing. One of the challenges in gaming, and virtual racing in particular, is testing. A big and important part of testing is manual testing done by a human, often professional, driver. There are also automated tests and it is a big challenge to keep them running. To thoroughly test a game we need thousands of test-scripts and this is a big challenge to maintain. Wouldn't it be great to have automated testing without the need to program test-scripts and test-data. Model-based testing may provide a solution.

Challenge

What is the future of racing? Is the next generation of professional drivers still racing on tracks? How about our options to include sustainability aspects within racing? Just some questions we have to think over and some very challenging and ethical topics to discuss. Even today's top Formula One teams are considering their options and there seems to be no doubt that electrical racing will be their top priority. Shifting towards the future, innovations like on track rechargeable cars and pit stops with a full battery recharge in less than 10 seconds have become requirements instead of wishful thinking. Top F1 teams also rely on race simulator to correctly correlate with the real world properties of their cars. Or are we even shifting towards pure virtual racing only? The world of racing, racing car designs and our way of thinking about it will evolve in the upcoming years. New technological innovations will be introduced and cars will most likely look like they have come straight out of a sci-fi movie. Beside all the futuristic aspects, nowadays we already have to deal with an increasing amount of data generated. Data analysis, artificial intelligence and machine learning aspects become more and more important. Software development will be top priority whether or not in a virtual environment. Until we can really predict the future we can only think over all major aspects of a realistic and the virtual truth of racing.
Let's suppose racing becomes pure virtual. What kind of reality and physical feedback is necessary for a real racing experience? How can we cope with the driver's viewpoint while bringing it closer to the fans? Bringing racing experiences into our living rooms and allow the fans to strive with their idols? Therefore drivers, cars and tracks shall be transferred into models allowing artificial intelligence to virtualize all of these. The complexity and especially the interaction between all of the parameters is quite complex. How can we solve these interactions, keeping in mind the laws of physics, include the passion for the sport and still become very realistic in virtual racing?

Input for the workshop

Before starting the workshop all participants will be demonstrated how state of the art gaming development is done these days. They will be seated into a simulation to experience the impact of different racing setups while acting as a professional test driver.

Today racing games are developed using all kinds of configuration parameters to simulate the behavior of the car on a track. Final testing is always performed by a (professional) human test driver. The test driver's feedback is used to tweak the track and environmental conditions to get it as close as possible to the reality of a real racing car.

A first step to a pure virtual racing environment is the replacement of a test driver by a model. How can we transform the professional test diver into a model to execute the primary testing and reduce the testing effort of humans by 80%. During the workshop we would like to investigate all aspects and physical feedback parameters of a professional test driver and create a model out of it for testing purposes at first. Next steps are to take the physical aspects of the track, the environmental conditions and complex interactions between driver and environment into consideration.

Team Leaders

Academic Team Leader

- Ana Oprescu, Universiteit van Amsterdam

Axini

- Machiel van der Bijl, CEO

3. ARTificial Intelligence for Simplified Texts (ARTIST)

Company: KB, National Library of the Netherlands

The KB, national Library of the Netherlands collects and stores all written publications that appear in the Netherlands and makes them available to all Dutch citizens. This includes books, newspapers and periodicals, but also websites and e-books. With 450 fte the KB is one of the largest cultural heritage institutions in the Netherlands. The Research department (20 fte) focuses on topics ranging from providing insights into the development of the public library sector to applying AI techniques to open up our collections to a wide audience. The KB is an active player in the field of Artificial Intelligence: it is co-founder of the Cultural AI Lab and the working group Culture&Media of the Netherlands AI Coalition.

Challenge: Text simplification

In the Netherlands, about 2.5 million citizens between 16 and 65 years old find it hard to read. This means they face challenges to fully participate in today's society. Recently we have seen this problem when people with-level literacy received invitations for the COVID-19 vaccines that were too complicated for them. But also understanding the news by reading news articles in the newspaper or websites can be difficult making it hard to understand current issues.

The KB, national library of the Netherlands, aims to make all publications available to all Dutch citizens, including people who have reading disabilities. In this use case we propose to explore the possibilities to make news articles/books/publications more accessible to illiterate people by applying AI techniques to automatically rewrite publications. In the Netherlands, several initiatives have been undertaken to manually make books or news articles more accessible. However, this is very labour intensive and only makes a small selection of publications available for illiterates. During the ICT with Industry workshop we aim to explore several methods to automatically rewrite news articles/book, making them available for all Dutch citizens.

Background on text simplification

Text simplification (TS) represents a long-standing task in NLP, attracting considerable and continuous research interest from the community since the 90s [1,2]. In the last few decades, mainstream methods for both lexical and syntactic TS approaches have developed from rule-based to data-driven [3,4,5]. In contrast to rule-based methods that require manual input from trained human editors, data-driven methods extract text simplification rules automatically from training corpora. Recent development of data-driven methods has focused on neural methods that are capable of capturing the distributional semantics of words [6,7], a key requirement for meaning preservation and text coherence. Correspondingly, a visible trend in TS is the recent shift from statistical machine translation, e.g., the phrase-based approach [8] to neural machine translation [9,10].
Neural TS methods are at the early stage with only a few studies reported. They are still found to be unreliable, often capture spurious, biased correlations; and are hard to explain,

due to linear and non- linear data transformation incomprehensive to humans. The problem is more obvious when large training data do not exist, a common situation in resource-limited languages and domains. In high- stakes domains with safety, trust, or ethical requirements, the aforementioned limitations of neural models can lead to negative and sometimes damaging effects. In this workshop, participants will work together to address those challenges by investigating the new approaches from a wide range of angles such as data (augmentation), modeling (e.g., neural-symbolic), and humans in the loop for both evaluation, interpretation, and knowledge integration.

Input for the workshop

For this workshop, the KB will provide participants with relevant data and (possibly manual) labeled training data. We are currently exploring which dataset is most suitable for this use-case and will make sure this data will be prepared well ahead of the workshop. For this we can build on our large digital collection of both historical newspapers on https://www.delpher.nl/ as well as books on https://dbnl.nl/ that can be made available to academic researchers through KB's Data services.

Before the workshop we will invite all participants to the premises of the KB, giving them an introduction to the challenge, a tour of the building and a chat with our colleagues who are working on making collections available to end users.

During the workshop, we will provide participants with relevant data and (possibly manual) labeled training data. We are currently exploring which dataset is most suitable for this use-case and will make sure these data are available during the workshop.

We consider two open software libraries. The primary library is Huggingface (https://huggingface.co), which contains pre-trained, transformer-based models for various languages (including Dutch) and various tasks including text simplification. As a fallback option, we use OpenNMT (https://opennmt.net), an open-source library specifically for neural machine translation. Those libraries allow participants to investigate problems in existing neural text simplification techniques, as well as to further improve the performance by e.g., fine-tuning the neural models on our prepared data. Computational resources will be made available to the participants. We consider the GPU cluster provided by SURFsara and other solutions such as Google Colab and Amazon Satemaker.

Scientific findings of the workshop will be disseminated in the form of a conference paper. We further envision a manifesto written by the participants collaboratively to summarize the challenges and opportunities of text simplification discussed during the workshop. We expect this will to help bring forward both the research and possible application.

Team Leaders

Academic Team Leaders

- Dr. Jie Yang Assistant Professor, Delft University of Technology
- Dr. Marijn Koolen Researcher Huygens ING

KB, national Library of the Netherlands

- Dr. Martijn Kleppe, Head of Research

References

[1] Chandrasekar, R., Doran, C. and Bangalore, S., 1996. Motivations and methods for text simplification. In COLING 1996 Volume 2: The 16th International Conference on Computational Linguistics.

[2] Al-Thanyyan, S.S. and Azmi, A.M., 2021. Automated text simplification: A survey. ACM Computing Surveys (CSUR), 54(2), pp.1-36.

[3] Glavaš, G. and Štajner, S., 2015, July. Simplifying lexical simplification: Do we need simplified corpora?. In Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers) (pp. 63-68).

[4] Horn, C., Manduca, C. and Kauchak, D., 2014, June. Learning a lexical simplifier using Wikipedia. In Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers) (pp. 458-463).

[5] Specia, L., 2010, April. Translating from complex to simplified sentences. In International Conference on Computational Processing of the Portuguese Language (pp. 30-39). Springer, Berlin, Heidelberg.

[6] Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S. and Dean, J., 2013. Distributed representations of words and phrases and their compositionality. In Advances in Neural Information Processing Systems (pp. 3111-3119).

[7] Devlin, J., Chang, M.W., Lee, K. and Toutanova, K., 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

[8] Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R. and Dyer, C., 2007, June. Moses: Open source toolkit for statistical machine translation. In Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics companion volume proceedings of the demo and poster sessions (pp. 177-180).

[9] Wang, T., Chen, P., Amaral, K. and Qiang, J., 2016. An experimental study of LSTM encoder-decoder model for text simplification. arXiv preprint arXiv:1609.03663.

[10] Zhang, X. and Lapata, M., 2017. Sentence simplification with deep reinforcement learning. arXiv preprint arXiv:1703.10931.

4. Interactive Information Retrieval for Bug Resolution

Company: Philips

At Philips, our purpose is to improve people's health and well-being through meaningful innovation. We collaborate on designing innovations to seamlessly bring together people, data and technology to help our customers achieve their Quadruple Aim: improving the patient experience of care, improving the health of individuals and populations, improving the work life of health professionals, and reducing the per capita cost of healthcare.

Philips is the leader in image-guided therapy (IGT), providing integrated solutions that advance minimally invasive procedures. At Philips IGT we help healthcare providers to decide, guide, treat and confirm the right care in real time during the procedure to enable better outcomes for each patient and at population level. Working together with clinicians, we use our comprehensive portfolio of interventional imaging systems, smart devices, software and services to treat one patient every second worldwide.

Challenges

In Philips IGT, there is an attempt underway to accelerate the resolution of software bugs that are found during development and testing of interventional X-ray systems. In the figure below (from [1]) the steps of a typical bug resolution workflow are shown inside the grey box. Philips is developing automated support for the steps understanding, triaging, and fixing. For example, Machine Learning algorithms have recently been developed and are being used to save time during bug triaging.

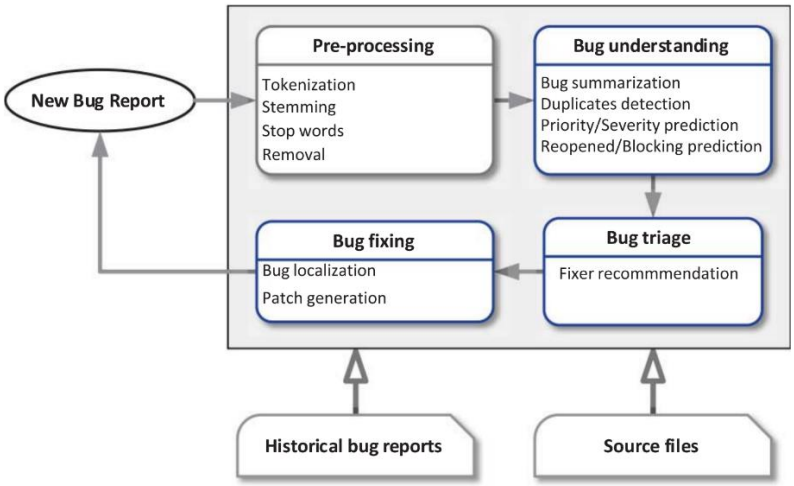

Automated solutions depend on data sources such as new and historical bug reports, source files, and other code artefacts like log files.

Challenge in bug triaging
The quality of a prediction algorithm for bug triaging depends heavily on the quality of the data that it is being trained on. A conclusion of a pilot project is that the quality of the training data used, historical bug reports, could be improved. Current shortcomings include missing data related to different aspects of a bug report (many fields in Philips' bug reporting tool are often not filled in), and a lack of structure in the free text fields.

It is expected that these shortcomings can be partly addressed by providing an alternative user interface for the task of submitting bug reports. In particular, making this task more interactive may help to elicit additional and more structured information that is crucial for training high-performing machine learning algorithms. Adding human-computer interaction solutions such as guided question- and-answer schemes, use of chat bots, and use of other modalities than text entry (e.g. voice) could be part of a solution.

Challenges in duplicates detection and bug localization
Beyond bug triaging, we are working on other steps in the bug resolution workflow, namely duplicates detection (part of bug understanding) and bug localization (part of bug fixing). Also these steps will provide challenges, and those challenges could also be cases for ICT with Industry. Since proof-of- concepts solutions are currently being developed by ourselves, it is not possible at this point to say what challenges we will face by the start of 2022. We mention those topics here because they may allow to broaden the opportunities for finding a match with an academic team.

Also duplicates detection and bug localization may benefit from approaches that add human-computer interaction to solutions that aim to (partially) automate these tasks. Duplicates may already be detected during submission of a new bug report, and therefore an improved user interface for the task of submitting bug reports may include features that help to quickly identify potential duplicates. Localizing the source files that need to be edited in order to fix a software bug is also a search process that may benefit from techniques that offer suggestions based on e.g. past fixes or formalized heuristics that could be applied to speed up the process of localizing buggy files.

Input for workshop

Philips uses several database systems for bug reports. The user interfaces of these will serve as the starting point ("baseline") for any proofs-of-concept to be developed during ICT with Industry, and therefore information about those user interfaces will be provided to the workshop participants.

The contents of the Philips bug databases are subject to non-disclosure. Instead, the workshop participants will be using a collection of bug reports from an open source project to develop and demonstrate their proofs-of-concept. It might be possible to prepare a separate database, populated with open source bug reports, that can be used during ICT with Industry. Otherwise, documentation of the user interfaces of these database systems will be made available.

Team Leaders

Academic Team Leader

- Guus Ramackers, Leiden Institute of Advanced Computer Science (LIACS), Leiden University

Philips

- Patrick Bronneberg, Department Manager

References

[1] T. Zhang, H. Jiang, X. Luo and A. T. S. Chan, "A Literature Review of Research in Bug Resolution: Tasks, Challenges and Future Directions," in The Computer Journal, vol. 59, no. 5, pp. 741-773, May 2016, doi: 10.1093/comjnl/bxv114.

5. Leveraging ecosystem-scale call graph datasets

Company: Software Improvement Group

The Software Improvement Group (SIG) helps business and technology leaders drive their organizational objectives by fundamentally improving the health and security of their software applications. SIG employs about 120 people, with offices in Amsterdam, Copenhagen, Antwerp, and New York.

SIG's Research team is active in scientifc R&D and innovation, often within joint academia-industry programs, such as H2020 FASTEN (https: //www.fasten-project.eu/), and NWO I-Cave (https://www.nwo.nl/en/ cases/i-cave-five-years-research-cooperative-and-autonomous-driving). In addition, SIG Research hosts internships for Master and PhD students, and supports various educational programs with guest lectures and board memberships.

Challenge

Modern code bases typically depend on dozens of third-party components originating from open source ecosystems. In running systems, third-party code may even be dominating locally developed code. This argues for a capability called Software Composition Analysis (SCA) that is able to analyze and remediate risks that exist within third-party components. Example risks are manifesting security vulnerabilities, poorly maintainable code that limits future quality, and license violations hidden away in library code.

To address these challenges, the H2020-FASTEN project (2019-2022) has developed capabilities to create fine-grained function/method-level call graphs for all artifacts and their versions within three open source ecosystems, Maven (Java), PyPI (Python), and Debian (C/C++). Provided with a local code base and a resolution of its transitive dependencies, these call graphs can then be merged together to provide a super-graph representing the precise usage of third-party code. In addition to call graph analysis, the project has created a plugin framework that enables source code analysis to extract meta-data and findings relevant for risk analysis.

The FASTEN tools have already been used to generate hundreds of thousands of call graphs, consisting of billions of nodes and edges, creating large databases of graphs together with various classes of meta-data, such as security vulnerabilities, code quality metrics, and licensing information.

Now that these databases are in place, the opportunity exists to attack a challenging research goal with a big potential impact in the software industry: **Providing developers with automated support to improve their decision making on third-party library usage. What library (and version) to implement, update, or remove, given its impact on architecture and overall quality?** In this direction, we see the following sub tasks:

- defining and validating metrics that express the quality of third-party library usage, involving notions of coupling, information hiding, ease of replacement, and update-ability.

- defining and validating metrics that express architectural quality concerns, eg. using centrality measures, graph topology, or time series analysis.
- designing a recommendation system to prioritize improvement opportunities, leveraging graph information to generate options and to provide cost heuristics.

In the workshop, you will be asked to collaborate on generating designs for the above goals, performing initial validation studies with the data available through the FASTEN project. A broad range of solutions and approaches is conceivable, ranging from traditional software metric (eg. Goal-Question-Metric) approaches, predictive analytics, to machine learning and recommendation systems.

The resulting proof-of-concepts should be coded in Java or Python, ideally wrapped in Docker containers, and made available on GitHub.

Input for the workshop

The participants will be able to work with the open source FASTEN repositories, and a large public data set of call graphs and meta-data hosted by the FASTEN project. The workshop organizers will further provide several open source projects to use as example code bases.

https://www.fasten-project.eu/ https://github.com/fasten-project
https://api.fasten-project.eu/api/mvn/packages

Team Leaders

Academic Team Leader

- Dr. Maurício Aniche, Assistant Professor in Software Engineering, Technical University Delft

Software Improvement Group

- Dr. Magiel Bruntink, Head of Research