



## Privacy and ethical risks in learning technology development

*By David Barnard-Wills, Trilateral Research Ltd.*

The DEVELOP project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No 688127. This publication reflects only the author's view and the European Commission is not responsible for any use that may be made of the information it contains.

## Introduction

The DEVELOP project is building a personalised learning environment (PLE) for career development. The DEVELOP system will allow users to assess their transversal competencies and social capital to highlight learning opportunities for career development. The ambition is that personalised visualisations of potential career paths can inform and guide learners towards realistic and attainable careers. The DEVELOP project is investigating the ways in which data from workplace social networks can be used to provide such recommendations and support career development. These data are a potentially rich source of information on the way that workplace relationships and communication patterns can influence job opportunities and career progression, but they are also sensitive and private.

As part of the project, the consortium is conducting a privacy and ethical impact assessment exercise (PIA+). The aims of the PIA+ are better privacy protection, increased transparency and accountability of personal information processing technologies, as well as mitigation of the risks of data processing. The exercise is also intended to engage stakeholders and affect the direction of the project from its earliest stages. PIA+ acts as a foundational component for achieving meaningful "privacy-by-design", providing information to support privacy protecting design decisions. The form of the PIA+ is integrated with the Agile approach of DEVELOP.

As part of the PIA, the project team developed a set of vignettes to demonstrate potential privacy, ethical and legal impacts that might arise from a PLE such as DEVELOP, if careful thought was not given to these issues in the design process. These vignettes are useful for communicating potential impacts across a multidisciplinary project team, and with stakeholders and users. As a type of storytelling, the vignettes assist in convincing people of the validity of a design, or the reality of an ethical challenge, and developing buy-in among stakeholders. Because stories are engaging, memorable and easily passed on, they are powerful communication tools. They give developers something to work with and a concrete problem to solve. They also tie in with the use-case scenarios that make up a part of the project's more traditional requirements gathering. Furthermore, they encapsulate ideas drawn from the expanding academic and best practice literature on ethics in big data, privacy and data protection, responsible innovation, as well as science and technology studies.

These scenarios were developed at a point where the requirements for the DEVELOP system had been collected, and a broad scope of the project was defined, but specific design decisions were yet to be taken. The scenarios are therefore built around a potential, hypothetical system: one that can still be changed to reduce these risks. These challenges are not unique to DEVELOP, given the wide spread of big data analytics into many areas of work and life, and these vignettes could be adjusted for other contexts.

DEVELOP is responding to many of the questions created by these vignettes. Ways in which they might be addressed include technological solutions (such as the way the data storage is set up or access controls), as well as social and organisational solutions (such as training, guidance and consultancy on how organisations can best manage some of these issues). These solutions come back to a number of core principles:

- transparency and privacy visualisation;
- user control, voluntary use and consent;
- data minimisation and anonymisation;
- role-based access control, restrictions on the scope of data included;
- error-correction and revocability;
- specific solutions for AI-planning and social network analysis.

## Scenario 1: The trouble with biased role models

**Problem statement:** Lisa is a web-developer, relatively recently hired by a company. She has decided that this career is good for her and wishes to pursue it further. She attempts to use the DEVELOP system to determine what appear to be the needed social capital requirements to be a senior web developer in this company. She is happy with the learning interventions the PLE recommends to her, which seem sensible to her. However, looking at the social capital of the senior developer role she is discouraged. The social network described is extremely unlike the social network she has acquired in her time at the company, and it seems unobtainable. She decides that perhaps she needs to move to another organisation that is a better fit for her.

**DEVELOP solution:** Algorithms can give the impression that they are free from bias, but they are very dependent upon the data upon which they are built - and this data comes from the real social world, and as such can bring with them social biases. DEVELOP is working to make sure that the right training data-sets are used, so that patterns of bias and discrimination in the workplace are not replicated by the tool, and also to find ways to communicate the outputs of the tool to users in a way which informs, but does not restrict them. DEVELOP operationalises and uses measurements of social capital in a way which supports, rather than limits, career development. The way that the system communicates with its users is also critical - DEVELOP pays attention to the way that social capital is reported to the user. For example, it might provide prompts to users on how to expand and build their networks, rather than implying they are (and might always be) a poor fit for the organisation.

## Scenario 2: The collapse of consent

**Problem statement:** William is a development manager with substantial experience. He is responsible for a team of developers, and wishes to better understand their skills and capacities. Through DEVELOP, William has access to a graphical display of the different transversals skills of his team members in a number of categories. Each of his team members has consented to sharing their assessments and learning outcomes with William. In an attempt to motivate his team, William takes a screenshot of this screen, prints it out, and places it on the office notice board on a weekly basis, in the process breaking the team's assumptions about how confidentiality works in the tool.

**DEVELOP solution:** The DEVELOP project understands that technical approaches to system security must be supplemented by organisational measures and training, taking into account the whole context of use within an organisation. Its deployment will be supplemented with guidance and training materials developed by the project team, and informed by the PIA, which can then be made available to future customers to avoid such problems. DEVELOP is also implementing best-practice guidance on consent from EU data protection authorities.

### Scenario 3: Forced development

**Problem statement:** Marieke, a senior HR manager, is really enthusiastic about the new DEVELOP tool. She sends a memo to all departments in the company suggesting that they are obligated to go into their account in the DEVELOP tool and consent to their line managers and the entire HR department being able to view their information. Discussion about what use an employee has made of DEVELOP becomes a key part of annual performance reviews in the company. This is the risk of consent not being meaningful and freely given.

**DEVELOP solution:** DEVELOP is being designed with user-control at its centre, along with granular consent mechanisms, and role-based access controls, as well as with guidance to management on its privacy-protecting deployment. Training for HR during the tool implementation will highlight that employees should not be forced to participate.

### Scenario 4: Trapped by the machine

**Problem statement:** Lisa has been recommended a set of online courses by the DEVELOP tool. Lisa is not keen on one of these courses, as she has had bad experiences with the content provider in the past - finding their courses out-dated and of poor quality. She declines to take the course, but it keeps re-appearing on her suggested courses. Her manager believes that the course is objectively the best course for her, and makes her completion of it a performance target to be completed by her next review.

**DEVELOP solution:** DEVELOP responds to this by placing the individual user in control of her use of the tool. Recommendations for courses are first presented to the employee, with a subsequent option to pass these on to managers or career coaches. Users are given the option to reject the system's recommendations within the software itself, and this can be used as an input to further refine the recommendation process.

### Scenario 5: Invisible requirements

**Problem statement:** William has diligently followed all the suggested learning interventions from the DEVELOP tool. He is personally feeling ready for his next employment challenge. However, the DEVELOP system keeps informing him that he is not suitable for the available job roles because of something to do with his "social capital". William experiments with the tool but cannot identify the reason for this gap. He consults with the HR department, who cannot help him because they also do not understand the algorithm behind the system. This is the risk of non-transparent systems.

**DEVELOP solution:** DEVELOP is finding ways to make the system more transparent to its users, so they can make meaningful assessments of its value to them, and to help guide their career decisions. Ways to increase transparency include contextual information boxes pop ups and visualisations, all based on a very clear picture of how recommendations are derived and why particular data is used in that process. When the DEVELOP team understands the process, they are better able to communicate this to their users, both HR and employees.

## Scenario 6: The problem of the privileged path

**Problem statement:** Lisa and her colleague Liam are both web developers with similar training and experience. However, Liam's desk is located next to his boss, William, whilst Lisa's is in another room. As a result, Liam hardly ever communicates with William over the messaging tool Slack, instead he simply turns around to ask a question. Both Lisa and Liam use the DEVELOP tool which draws upon communications data to measure workplace social capital. Tool rates Liam's workplace social capital as significantly poorer than Lisa's, due to a lower level of measured communication with their senior team lead. This is the risk of assuming that what is measured is what matters.

**DEVELOP solution:** DEVELOP services are designed to communicate with users about the scope and limitations of this type of analysis. In practical terms, the project's testing and piloting phases will be particularly alert to this type of unintended consequence.

## Scenario 7: The games we play

**Problem statement:** Due to their interaction with the DEVELOP tool, Lisa and her colleagues become convinced that email and workplace social network communication with senior leadership is the best way to build social capital. As a result, senior leadership are overwhelmed by messages about minor or irrelevant issues from junior colleagues attempting to game the system.

**DEVELOP solution:** Transparency about the way the system works can help to reduce "gaming". It would be a mistake to assume that a hidden process is more secure against manipulation. Rather, if users can understand why and how the system uses information for their benefit, they will feel a reduced need to manipulate it. Guidance to users will communicate the reason for providing accurate data through real behaviour, whilst privacy protections increase their confidence in the system. If the employees remain in control of their use of the system, this reduces the risk that it can be used against them and therefore reduces the perception of a need to game the system.

## Scenario 8: Unintended consequences

**Problem statement:** Due to her role in both helping people dealing with HR problems, and managing disciplinary issues, having a high level of communication with Marieke becomes associated by the DEVELOP algorithm with having poor job performance. As this becomes known, Marieke is dropped from email lists, becomes a social pariah, and is left uninformed of problems until it is too late. As with Lisa and her colleagues, this represents the risk of unintended consequences and risk of measurement changing the thing measured.

**DEVELOP solution:** In a similar way to "The games we play", by presenting a realistic, rather than overstated account of the capabilities of the DEVELOP system, combined with a design focus upon data-informed guidance, rather than pseudoscientific objectivity, DEVELOP minimises the risks of both gaming and unintended consequences.

## Conclusion

Creating a set of evocative scenarios that help to visualise and make real the sometimes complex and abstract privacy and ethical risks is a powerful tool for ensuring that technological systems are designed in a privacy-friendly way, and for avoiding common design pitfalls. This set of vignettes contributes towards making the DEVELOP system safe and secure for its users.

Identifying the risks is not in itself sufficient. DEVELOP is responding to these ethical challenges through a bespoke set of interventions and methods. These include: increasing the transparency and visibility of data processing and informing the users about how their data is being processed; promoting user control through user-centric design, voluntary use and granular informed consent practices; decreasing the risk of function creep or over-reliance upon the system through role-based access controls, and careful selection of scope and data sources.

## References

- Ball, K. (2010). *Workplace Surveillance: an overview*. Labour History, Vol. 51, No. 1, pp. 87-106. doi: 10.1080/00236561003654776
- Bogdanovic, D., Crawford, C. & Coles-Kemp, L. (2009). *The need for enhanced privacy and consent dialogues*. Information Security Technical Report, Vol. 14, No. 3, pp. 167-172. doi: 10.1016/j.istr.2009.10.011
- Danezis, G., Domingo-Ferrer, J., Hansen, M., Hoepman, J., Le Métayer, Tirtea, R., Schiffer, S. (2014). *Privacy and Data Protection by Design - from Policy to Engineering*. European Network and Information Security Agency (ENISA). doi: 10.2824/38623
- Guillemin, M. & Gilliam, L. (2004). *Ethics, reflexivity and "ethically important" moments in research*. Qualitative Inquiry, Vol. 10, No. 2, pp. 261-280. doi: 10.1177/1077800403262360
- O'Neil, C. (2016). *Weapons of Math Destruction: How big data increases inequality and threatens democracy*. Crown, United States. ISBN: 978-0553418811
- Rip, A. & Schot, J. (2002). *Identifying loci for influencing the dynamics of technology development*. Soerensen, K.H. & Williams, R. (Eds.): *Shaping Technology, Guiding Policy: Concepts, Spaces and Tools*. Edward Elgar, Cheltenham, pp. 158-176. ISBN: 978-1840646498
- Wright, D. (2012). *The state of the art in privacy impact assessment*. Computer Law & Security Review, Vol. 28, No 1, pp. 54–61. doi: 10.1016/j.clsr.2011.11.007
- Wright, D. (2015). *Ethical Impact Assessment*. Holbrook, J. (Ed.): *Ethics, Science, Technology, and Engineering: A Global Resource*, 2nd edition. Macmillan Reference USA. ISBN: 9780028662138