




AVT

Exit report from the sandbox project AVT – Activity data for assessment and adaptation

Topics: Legal basis, data protection impact assessment (DPIA) and transparency.

February 2022

Contents

SUMMARY	3
ABOUT THE PROJECT	4
OBJECTIVE FOR THE SANDBOX PROCESS	5
LEGAL BASIS	6
INTRODUCTION	6
LEGAL BASIS REQUIREMENT	6
DISCUSSION OF LEGAL BASIS	8
THE RIGHT TO OBJECT	11
SUMMARY – LEGAL BASIS	12
IMPACT OM PUPIL PRIVACY	13
INTRODUCION	13
WORKSHOP WITH STAKEHOLDERS	13
OTHER IDENTIFIED RISKS	14
TRANSPARENCY	16
INTRODUCION	16
INFORMATION ADAPTED TO USERS	16
EXPLAINING THE SYSTEM’S UNDERLYING LOGIC	17
NEXT STEPS	18

Note

Technology and the law are constantly evolving, and adjustments and clarifications may have occurred after this report was written.

Summary

The AVT project – *Aktivitetsdata for vurdering og tilpassing* (Activity data for assessment and adaptation) – is a research and development project focused on digital learning analytics in schools. The project explores opportunities and challenges associated with the use of learning analytics and artificial intelligence (AI) to analyse pupil activity data from various digital teaching aids.

Through participation in the sandbox, the AVT project seeks to explore the legal framework, in addition to frameworks for responsibility and ethics, for the use of learning analytics in schools. The overall goal is to pave the way for adapted education for each individual pupil, while also protecting the pupil's privacy. The sandbox project has been split into three sub-objectives, where we have explored the legal basis, protection of the pupil's privacy through a data protection impact assessment, and transparency.

Conclusions

- **Legal basis:** The views of the Data Protection Authority and the AVT project differ in terms of the choice of legal basis and of whether the supplementary basis is adequate. The AVT project believe the processing is in response to a “legal obligation” and points to the supplementary basis in education legislation to support this conclusion. The Data Protection Authority finds that “in the public interest” is the most appropriate basis for the processing, and that the supplementary basis ought to be stronger than it currently is. The Authority, however, does not rule out that “legal obligation” may be used.
- **Impact on pupil privacy:** Through dialogue with stakeholders, all representing groups who, in various ways, would be affected by the learning analytics system, three main risks to pupil privacy were identified. These are the risk of altered behaviour (cool-down effect), the risk of incorrect personal data in the system, and the risk of the technology causing the pupils unwanted stress. In addition, we highlighted three other potential risks, namely risks associated with special categories of personal data, the processing of third-party data, and where to draw the line towards an automated decision-making system.
- **Transparency:** One challenge in the AVT project is the considerable range in ages and level of development in the group that will be using the learning analytics system. In the sandbox, we discussed how linguistic and visual tools can be used to provide information that is simple enough for the youngest pupils, while also meeting the information needs of older pupils and parents. Among other things, the AVT project has considered building a “dummy” version of the learning analytics system, which allows users to experiment with different variables. This type of visualisation is often quite effective at explaining advanced technology in a straightforward manner.

Going forward

The Ministry of Education is currently working on a new education act, which will not be ready until 2023, at the earliest. The education sector, therefore, currently has a unique opportunity to clarify how schools can contribute to the development of educational AI tools.

The sandbox project is completed, but the exploration of artificial intelligence in schools has only just begun. The AVT project is blazing the trail in a field that will only become more important in the future. The project's contributions in the sandbox have shed light on an area the Data Protection Authority does not have previous experience with. The opportunities are endless, and the importance of a sober, ethical approach could not be clearer. The AVT project has taken its responsibilities seriously and is now likely to build the foundation for a better education of our children in the time to come.

About the project

“Education and training in schools and training establishments must, in collaboration and agreement with the home, open doors to the world (...)”

(Objectives clause of the Education Act)

One of the values for Norway’s primary and secondary education is that it should “open doors to the world and the future”. In order to achieve this, it is essential that our children receive a good education. We have a little over half a million primary and lower secondary school pupils in Norway. All of these children are unique individuals, with unique prerequisites for learning. In order to ensure that the children benefit from education, school owners (primarily municipalities) have an obligation to assess the performance and progression of each pupil and adapt the education to each individual pupil’s needs, as far as it can be done (see Section 1-3 of the Education Act).

In practice, teachers are responsible for adapting the teaching. Today, the information base that teachers need in their work on assessment and adaptation is scattered across paper sources and various digital tools, as well as, in part, the teacher’s own memory. What if a system could gather this information in one place and systemize it to support teachers in their work? A system that analyses the progress made by each pupil, presents an overview of this progress and makes recommendations for what the pupil could focus on next. This is the goal for the so-called AVT project (*Aktivitetsdata for vurdering og tilpassing* – Activity data for assessment and adaptation).

The AVT project is a research and development project focused on digital learning analytics in schools. The project explores opportunities and challenges associated with the use of learning analytics and artificial intelligence (AI) to analyse pupil activity data from various digital teaching aids. Activity data is the term used for the data that is generated when a pupil completes activities in a teaching tool. The activity could be the answer the student gave to a question, information about which activity the pupil completed, how long the pupil spent working on it, and whether the pupil gave the correct answer or an incorrect one. Similar to how schools purchase textbooks from publishers, they also purchase digital teaching aids from different providers. The AVT project focuses on digital teaching aids, but we simply refer to them as “teaching aids” in this report.

The context for the project is the research field artificial intelligence (AI) in education, with a primary focus on intelligent tutoring systems and learning analytics. Specifically, the AVT project uses an open learner model, and analytics and recommendation algorithms for analysis of learning progress and recommendations for pupils. Results of the analysis are presented in an online portal (dashboard) customised for each user group – such as teachers, pupils and parents. Users will use Feide to log on to the dashboard.

The objective of the project is to explore whether it will be appropriate to develop a publicly managed solution that could support teachers in their work on adaptations for each individual pupil, give pupils greater insight into their own learning, and support teachers in their work on pupil assessment. The goal of adaptations is to ensure that the pupils achieve the best possible learning outcome from their education. On a more general level, the AVT project aims to drive the development of national solutions, guidelines, norms and infrastructure.

The first part of the project (AVT1) was initiated in 2017. During this phase, a prototype was developed, as well as a solution for collecting activity data from different providers. And a preliminary data protection impact assessment (DPIA) was also conducted. This assessment maps the potential impact of the system on the protection of personal data and develops measures to mitigate the risks. The second part of the project (AVT2) was initiated in 2020. During this phase, activity data on pupils at participating schools in Oslo was collected from various teaching aids, before the results were analysed and presented to teachers and pupils.

The project owner for the AVT2 project is the Norwegian Association of Local and Regional Authorities (KS). The project has been led by the University of Bergen (UiB) and its Centre for the Science of Learning & Technology (SLATE), and the City of Oslo’s Education Agency has been the main partner and driving force since the project’s initiation in 2017. Recently, the Municipality of Bærum and the regional intermunicipal collaboration Inn-Trøndelag have also joined, in minor roles. All project partners have been involved in the sandbox project.

So far, the system itself has only been tested with constructed test data and pseudonymised activity data related to mathematics from a group of pupils in Oslo schools. This data was collected from transition tests in mathematics after 4th and 7th grade, as well as from a couple of the participating suppliers, who have been able to share data in a suitable format.

Objective for the sandbox process

The objective of this sandbox project must be seen in relation to the Data Protection Authority's overall objective for the sandbox, which is to promote the development of artificial intelligence solutions that, from a data protection perspective, are both ethical and responsible. The Data Protection Authority plays a supervisory role in the sandbox, where we work with the AVT project to arrive at good data protection solutions for learning analytics.

In previous phases of the AVT project, project participants got the impression that the education sector is not quite prepared for the use of artificial intelligence – and that this applies both to the regulation of this technology and to its use. In addition, it is not quite clear what the boundaries for “responsible artificial intelligence” actually are. Through participation in the sandbox, the AVT project seeks to explore the legal framework, in addition to frameworks for responsibility and ethics. The overall goal is to pave the way for adapted education for each individual pupil, while also protecting the pupil's privacy in the best way possible.

The sandbox project has had three sub-objectives, where we have explored:

1. Legal basis. Whether the legal basis for the processing of pupil activity data meets the requirements of the General Data Protection Regulation (GDPR). If not: what will it take to ensure compliance?
2. Pupil privacy. What are the consequences of the learning analytics system on the protection of the pupils' personal data? What do the parties responsible have to keep in mind to ensure that the pupils' privacy is protected?
3. Transparency. How can information about the learning analytics system be given to its users (teachers, pupils and parents)? What do they **have to be** informed about, and what **should** they be informed about?

Legal basis

Processing of personal data: In short, personal data is any information that can be tied to a natural person, directly or indirectly. It does not matter which format the information has. Text, images, video and audio are all included. "Processing" includes everything that is done with the information: collection, structuring, changing, analysis, streaming, disclosure, transfer, storage, deletion, etc.

Introduction

Public authorities process a lot of personal data about us as citizens. In order for this processing to be lawful, they need a legal basis – a statutory authority. This statutory authority defines the framework of what the authority (such as a municipality) can use our personal data for, and serves to protect citizens from intervention in their lives.

Municipalities are responsible for Norway's primary and lower secondary school (see Section 13-1 of the Education Act). This means they are responsible for and have the authority to decide how this education is provided – within certain limits, of course. As an example, it would be difficult to achieve good and effective education if each parent had the power to decide which textbooks they wanted their child to have. In addition, the municipalities are "school owners" of the public primary and lower secondary schools. In other words, municipalities have the power to decide what happens in "their" public primary and lower secondary schools. Legally, the municipalities are then also the "data controller" for the processing of the pupils' personal data. The learning analytics system that is being developed in the AVT project is something municipalities will be able to offer schools, in line with other teaching aids and support systems teachers use to plan, teach and follow up on the pupils' learning. If the school chooses to implement the system, it will not be optional for the pupils.

Legal basis requirement

Municipalities need a legal basis in order to impose pupils to use the learning analytics system. And because the system uses artificial intelligence, they need a legal basis for both the application phase and the development phase (also called "post-learning") for the artificial intelligence (algorithm). The development of the algorithm is a process that often continues after the system has been implemented – as a perpetually repeated process – because the system learns from new data that is added and generated. The purpose of the two types of processing required by these phases must be defined in detail by the data controller, and we return to the AVT project's definitions later in this chapter.

It is natural to split the question of legal basis in two, based on the two main phases in an AI project; the development phase and the application phase. The two phases utilise personal information in different ways.

Article 6 of the General Data Protection Regulation defines several alternatives for lawful processing. Two of these are relevant for learning analytics:

Article 6 (1) (c) of the GDPR ("legal obligation), which says that

[...] processing is necessary for compliance with a legal obligation to which the controller is subject [...]

and Article 6 (1) (e) of the GDPR ("in the public interest"), which says that

[...] processing is necessary for the performance of a task carried out in the public interest or in the exercise of official authority vested in the controller [...]

Both of these legal bases require what we call supplementary legal basis. This means that the processing must have a legal basis in another law. The requirement for a supplementary legal basis follows from Article 6 (3) of the Regulation, which, among other things, provides that

*[...] The purpose of the processing shall be determined in that legal basis **or**, as regards the processing referred to in point (e) of paragraph 1, shall be necessary for the performance of a task carried out in the public interest or in the exercise of official authority vested in the controller. [...] [our emphasis]*

The supplementary legal basis does not necessarily need to regulate the processing in detail, even if the Regulation does allow it. However, we also need to consider the legal basis in light of other legal requirements. Among other things, the supplementary legal basis must be designed in such a way that the processing is predictable for the person whose data is being processed. This is part of the legal protection principle of predictability, which is a central principle in Norwegian law. In addition, the more invasive the processing is, the clearer the legal basis should be. Also, the principle of data minimisation (Article 5 of the GDPR) is relevant, because it limits the scope of personal data available for lawful processing. This principle states that processing must be limited to only what is necessary in relation to the purposes of processing. In order to make assessments and adapt the level of education to the needs of individual pupils, it is for example, necessary to process information about the pupil's learning progress in different subjects. But it is not necessary to process information about the pupil's leisure time activities or when the pupil did their homework, even if this information could be relevant for the assessment.

It is also worth noting that the conditions associated with "legal obligation" are more narrowly defined than those associated with "in the public interest". If based on the former, the controller must be legally **obligated** to perform the processing, and the purpose **must** follow from the supplementary legal basis. This means that the legislator in a democratic process must define the purpose of processing the controller is legally obligated to perform. Furthermore, the data subject has no right to object to the processing if the data controller is legally obligated to perform the processing. The limited scope of action that follows from "legal obligation" indicates that there are no real alternatives to the processing of personal data for achieving the stated purpose. A logical consequence of this would be that it would not serve any practical purpose to give the data subject any right to object.

However, if the processing is conducted for the purposes of performing a task carried out in the public interest or in the exercise of official authority vested in the controller, the purpose does not need to be specified in the legal basis (but it may). It is sufficient that the purpose in itself is necessary for performing the task or exercising the authority in question. As such, letter (e) gives the data controller room to define the purpose themselves. This means that the requirements for the supplementary legal basis are somewhat less expansive if the processing is based on "public interest" than if it is based on "legal obligation".

For processing based on "public interest", the data subject also has the right to object to the processing, see Article 21 of the Regulation. Such objections must be made on grounds relating to "that persons particular situation", but the provision does not specify further what this entails. The wording suggests it may concern individual circumstances of virtually any nature. The right to object, however, is not absolute. If the data controller can demonstrate "compelling legitimate grounds for the processing which override the interests", the processing may continue (see Article 21 (1) of the GDPR).

The Regulation also provides other legal bases for processing, which are not relevant for our project. Basing the processing on consent, for example, is not an option. Consent may be withdrawn at any time, and is not an appropriate solution where there is an imbalance of power between the parties. Basing the processing on contracts between the municipality and the parents is also not an option, because nobody can be forced to enter into a contract.

Discussion of legal basis

The application phase

Legal basis was the most frequently discussed topic in this sandbox project, and these discussions primarily focused on the application phase of the learning analytics system. The AVT project defined the purpose for this phase as follows:

Use learning analytics to support teachers in their assessment work, to provide better adaptations and to provide pupils with insight into their own learning

There is no doubt that the municipalities have a legal basis for processing the personal data of pupils, both to provide educational adaptations and to assess their performance. In 2021, a new provision was added to the Education Act (Section 15-10).

This provision grants school owners the general right to process personal data, including special categories of personal data, whenever necessary to perform tasks pursuant to the Education Act.

It follows from Section 1-3 of the Education Act that education must be adapted to the abilities and aptitudes of the individual pupil. Furthermore, it follows from the Education Regulations that pupils have a right to continuous assessments in subjects. The purpose of assessment is to promote both learning and the desire to learn, as well as to provide information about competence acquired at various stages of instruction in the subject.¹ The basis for assessment is the competence objectives in the subject curriculum.² The requirements for continuous assessment are defined in Section 3-10 of the regulations. The regulations specify that continuous assessment is an integrated part of education, and that this assessment shall be used to promote learning, adapt education and increase competence in the subject. The municipality is responsible for ensuring that the requirements of both act and regulations are met, which includes making available any and all resources necessary for ensuring compliance with these requirements.

The discussions in the sandbox has primarily focused on whether the legal obligation municipalities have to adapt education and assess pupil performance, also extends to the processing of the pupils' personal data using learning analytics and artificial intelligence. As mentioned above, this must be seen in light of, among other things, the supplementary legal basis, the principle of data minimisation, and the legal protection principle of predictability. Before we move on to the discussion, we would like to provide a quick summary of how adaptation and assessment work today.

There is no universal method for how municipalities assess pupil competence and provide adapted education. Methods vary between schools and inbetween teachers at the same school. Nevertheless, the AVT project has provided a general example: The teacher gathers assessment information manually by looking at completed activities in each teaching aid used by the pupils. Each teaching aid presents assessment data in different ways. The teacher must therefore try to connect the dots and build an overview of the pupil's competence and development, based on the pupil's answers and activities, in order to generate the best possible basis for assessment. This basis is then applied in the adaptation of education. This process is relatively unregulated, and the tools used spans from the teacher's own memory, via unstructured written notes, to more structured systems, such as digital learning platforms with various functionalities.

It could be relevant to compare the learning analytics system with teaching aids and support systems for the compilation of activity data from one major provider with a broad product portfolio. These providers often have similar functionality for compiling pupil performances from their own product portfolio. The learning analytics system of the AVT project will collect more of this essential functionality for teachers in a single system, which will also use artificial intelligence.

¹ Sections 3-2 and 3-3 of the Education Regulations

² Second paragraph of Section 3-3 of the Education Regulations

There are several benefits of implementing this type of centralized learning analytics system, which gathers activity data from different teaching aids and uses artificial intelligence to process it:

- The activity data for each individual pupil will be collected from different teaching aids and easily accessible from a single system.
- The presentation of the data is more structured, because they have a comparable format and are tied to elements in the curriculum. This will improve the quality of the analysis, performed by both the system and the teacher.
- Analyses of individual pupil learning will be more accurate, as the pupil's data can be compared with aggregated data about other pupils.
- New personal data is generated, and a pupil profile can indicate where the pupil is in their learning process and provide recommendations for what to do next. This would be difficult to extract without the compilation provided by the system's AI.

The AVT project has evaluated that continuous adaptation and assessment work, with or without the use of digital learning analytics, is the same processing, in principle. The legal obligation to provide assessment and adapted education is consistently implemented in the Education Act and the associated regulations. This obligation can be fulfilled in various ways, using various tools – both printed and digital. The scope of data processed about each individual pupil and the metadata (context) this is added to vary, and is not tool-dependent. The AVT project therefore maintains that there is no difference, in principle, in whether a teacher performs their own analyses of pupil competence and adaptation needs by, for example, studying the pupil's performance across various teaching aids, notebooks and verbal participation in class, and having a learning analytics system to help the teacher make these assessments by compiling and processing data from some of the same teaching aids in a visual interface.

Further on, the AVT project has considered that the choice of tool is not relevant in the consideration of whether or not a legal basis is sufficient. In support of this view, they refer to the Privacy Appeals Board's decision in the *Spekter* case, where the Board wrote:

*"While the data may be qualitatively different as a result of digital collection compared to collection through observations and communication with the pupils, this is also not something that would lead to the legal basis in this case being deemed insufficient. The Board does not believe it is relevant in this context to assess the appropriateness of *Spekter* versus other means of collecting the same information."*

As a further argument in support of the use of a learning analytics tool falling within the scope of the legal obligation, the AVT project emphasised that it is the school's responsibility to adapt education to ensure the best possible basis for assessment. This applies both to the breadth of the subject and to in-depth knowledge within certain areas of the subject. Furthermore, the AVT project points to the fact that specific competence objectives within the curriculum define the framework of what is necessary, and that the principle of necessity in the Personal Data Act must therefore be seen in light of this. Personal data that is not relevant for the pupils' attainment of the learning objectives for the subject will therefore be unnecessary, whereas personal data that is relevant for the pupils' attainment of the learning objectives will be necessary.

As a basis for the structuration of personal data in the learning analytics system, a domain structure (called *fagkart* – subject map – in the project) was developed. This subject map is a digital representation of the curriculum and more detailed subject topics and concepts within it. In addition, the project has defined which types of metadata to register, through an xAPI profile and examples. The AVT project claims this framework ensure that all personal data used are relevant, and therefore necessary. They claim that the elements of the metadata are what must be verified in light of the necessity principle, and not the scope of the data sources themselves. The AVT project argue that they have subjected these elements to a rigorous assessment. In the subject map, six of the ten elements, for example, are taken straight from the curriculum, whereas the last four are based on the curriculum.

The view of the Data Protection Authority is that the choice of tool may be relevant for the assessment of whether the processing is covered by the supplementary legal basis. Different tools process personal data differently, and processing using one tool may constitute a greater intervention than processing using another tool.

Tools are less relevant in the assessment if it involves two similar tools processing the same quantity of personal data in almost identical ways. The use of learning analytics and artificial intelligence will, among other things, generate new and different types of personal data about the pupils, which means that much more personal data about each individual pupil will be processed. The system will also be capable of learning, finding connections, conducting probability analyses and drawing conclusions far beyond what both humans and systems that do not use artificial intelligence are capable of. In our view, therefore, the use of learning analytics with the help of artificial intelligence represents a fundamental deviation from current methods of adaptation and pupil assessment.

Even so, whether or not there is a fundamental difference between the current method and the future method involving learning analytics is not the determining factor for which purpose of processing is most appropriate. Even if there is a fundamental difference, this does not necessarily entail that the processing cannot be based on a “legal obligation”. To consider this, it would be relevant to explore whether, and if so, how, the processing of personal data is affected by the implementation of the learning analytics tool.

The Data Protection Authority supports that, in connection with assessment of the pupil’s performance and adapting the education to meet the pupil’s needs, it may be necessary to process personal data that are relevant for the pupil’s attainment of the learning objectives defined by the curriculum. We do not, however, share the view that the scope of the data sources is irrelevant in a consideration of which types of data it would be necessary to process, to “comply with a legal obligation” or to “perform a task carried out in the public interest”, respectively. The necessity requirement defines the limits of what constitutes lawful processing of personal data within the defined purposes. The principle of necessity is a legal standard in the EU/EEA and must be interpreted in accordance with the Regulation’s purpose. If the processing is based on a “legal obligation”, the framework provided by the Education Act sets a firm boundary, in that the legislator lays down the purposes, and considerations of which types of data are considered necessary must be made in light of these. If the processing is based on “public interest”, the processing must also be performed within the scope of the Education Act, but the controller has more freedom to define what the purposes are. In the Data Protection Authority’s view, the learning analytics system appears to be processing more types of personal data than what is strictly necessary to fulfil any legal obligation. The Data Protection Authority, then, has some doubts about whether the municipalities can be said to have a legal obligation to use learning analytics and artificial intelligence to adapt education and assess pupil performance. Based on the Data Protection Authority’s limited understanding, therefore, it would seem more appropriate to base the processing on “public interest”, where the municipalities have more freedom to define what the purposes of the processing are.

It is important to point out that the sandbox project did not include a comprehensive analysis of how the obligation to provide assessments and adapted education to each individual pupil is currently handled, or of the degree to which current methods are suitable for fulfilling this obligation. The AVT project has highlighted that the quality of continuous assessment and adapted education is likely to increase considerably if the learning analytics system were implemented. The project has furthermore argued that one of the central values of the learning analytics system is that each pupil would have improved access to and use of their personal data, as a result of the data being available on a joint platform managed by public authorities.

The development phase

In considering the legal basis, it is natural to distinguish between the application phase, as described below, and the development phase. Among other things, the development phase entails that the learning analytics system uses activity data from the pupils to train the algorithm, making it more accurate in its predictions.

The AVT project concluded that the purposes for the two phases are not compatible with each other. One must therefore consider the legal basis for the algorithm’s development phase on its own, separate from the main purpose. In the sandbox, the partners agreed that “public interest” was the most appropriate legal basis for the development phase. The municipalities do not currently have any legal obligation to contribute to the development of AI tools for use in education.

The processing of personal data about individual pupils for the purpose of developing the algorithm in a learning analytics tool is a task carried out in the public interest, and not just in the interest of the individual pupil whose personal data is processed. There is reason to assume that processing for this purpose may appear less predictable for the pupils and their parents than does the processing of personal data in connection with the use of the learning analytics system.

By basing the processing on “public interest”, the pupils are granted the right to object. Read more about this later on in this report. The AVT project aims to design the learning analytics system in a way that protects this right, with an easily accessible function where pupils can request that their data should not be used to train the system. In addition, they want to limit the impact on the pupils’ privacy by pseudonymising personal data used to develop the algorithm. They

also want to look into the possibility of anonymising the pupils' personal data for this purpose. The latter is dependent on the policy adopted by the European Data Protection Board in its guidelines on anonymization, currently under preparation.

Like the processing of personal data necessary for compliance with a legal obligation, the processing of personal data necessary for the performance of a task carried out in the public interest also requires a supplementary legal basis.

The AVT project has primarily emphasised Section 13-3e, together with Section 15-10, of the Education Act as the supplementary legal basis for the processing of personal data in the development phase. Section 13-3e of the Education Act regulates the municipalities' obligation to promote quality development in education. The AVT project argues that this obligation entails making sure learning analytics are as accurate and effective as possible, which makes it necessary to develop the chosen algorithm. This includes the processing of personal data about pupils for post-learning purposes.

The topic was not discussed in detail in the sandbox, because most of the attention was directed at the legal basis for processing in the application phase. The Data Protection Authority finds it challenging to identify a supplementary legal basis for the development phase in the Education Act or associated regulations, but does not rule out that the basis for processing provided by the AVT project may be used. This challenge is not unique to the education sector; it is something most public bodies who wish to contribute to the development of AI tools risk encountering. The issue is also briefly discussed in the final report from the NAV sandbox project, which ran parallel to this project.

The Ministry of Education is currently working on a new education act, which will not be ready until 2023, at the earliest. The education sector, therefore, currently has a unique opportunity to clarify how schools can contribute to the development of educational AI tools.

The right to object

The processing of personal data on the basis of "public interest" means that the data subjects, with certain limitations, have the right to object to the processing (see Article 21 of the GDPR). In the sandbox, we discussed this right in light of the processing performed by the learning analytics system. As previously mentioned, an objection must be based on the data subject's "particular situation". As we have limited experience with the use of artificial intelligence in primary and lower secondary education, it can be difficult to envision the kinds of particular situations this may involve.

The right to object to processing by the learning analytics system can have both positive and negative implications. An objection would lead to an assessment by the municipality, which, in turn, may lead to improvements in the system. Also, there is an imbalance of power between the municipality and the pupil, and the right to object could contribute to levelling out this imbalance. The resources it would require to process potential objections can be a drawback, and it is difficult to estimate how many objections there will be. The use of artificial intelligence is innovative and unfamiliar territory for many people. Uncertainties may give rise to concerns in both pupils and parents, which may, in turn, lead to objections against processing. While objections on these grounds are not necessarily valid, and the controller does not necessarily need to grant them, the objections must of course be processed, and the burden of evidence is on the controller.

The number of objections could likely be reduced with transparency (which is also a requirement), such as about which types of data are being processed, how the data is used, what the pupils' rights are, and how they may exercise their rights. At the end of the day, the majority of people will likely appreciate the improvements in adaptations, with subsequent improved learning for the children. Provided that the system is developed responsibly.

Summary – legal basis

The Data Protection Authority and the AVT project dedicated two workshops to discussing the legal basis for processing of personal data by the learning analytics system, and our views differ in terms of the choice of legal basis and whether the supplementary basis is sufficient. The AVT project believe the processing is in response to a “legal obligation” and points to the supplementary basis in education legislation to support this conclusion. The Data Protection Authority’s argument is that “public interest” is the most appropriate basis for the processing, and that the supplementary basis ought to be stronger than it currently is.

In this context, it is important to point out that the partners in the AVT project know considerably more about the education sector, learning analytics and the Education Act than the Data Protection Authority. The project has been ongoing since 2017, and the partners have worked actively on these assessments and the system itself. The AVT project, therefore, has a much broader basis for assessment than the Data Protection Authority, and the project has been able to dive much deeper into this subject matter. The Data Protection Authority’s contributions are considered as a guidance and do not constitute a review of the legality of the learning analytics system. At the same time, the AVT project is conscious of the fact that there are some central questions concerning the processing of data in the application of the system they do not yet have any experience with, and that may entail a reconsideration of their view on the legal basis for processing at a later date.

We would also like to emphasise that the education sector has an opportunity to initiate the establishment of a stronger legal basis than it currently has. The new education act, for example, may include explicit provisions for the use of artificial intelligence and digital learning analytics, with a more clearly defined framework for their use. If so, this would ensure a stronger democratic foundation for the use of such tools and a higher degree of predictability, for pupils, parents, teachers, school owners and others who may be affected.

Impact on pupil privacy

Introduction

For some types of processing, data controllers are required to perform a data protection impact assessment (DPIA). More specifically: if the processing of personal data is likely to result in a high risk to the rights and freedoms of the data subject, especially as a result of new technology. But what is the threshold for “high risk”? This is not always easy to know. The Data Protection Authority has prepared guidelines for the performance of these assessments, as well as a list of processing activities that always require a DPIA. There are several items on this list that are relevant for the AVT project:

- Data collected via third parties in conjunction with at least one other criterion.
- Processing of personal data using innovative technology in conjunction with at least one other criterion.
- Processing of personal data for evaluating learning, coping and well-being in schools or kindergartens. This includes all levels of education from primary school to secondary school and higher education. (Vulnerable data subjects and systematic monitoring.)
- Processing of personal data to systematically monitor proficiency, skills, scores, mental health and development. (Sensitive data or data of highly personal nature and systematic monitoring.)

In the AVT project, they had already performed a DPIA before the project was accepted into the sandbox. One of the sub-objectives for the sandbox project was to further refine this assessment, with suggestions from the Data Protection Authority on issues related to artificial intelligence.

As the sandbox project progressed, we acknowledged that there would not be enough time to carry out all of the activities on the project plan. We therefore had to deprioritise activities related to this sub-objective. As a result, we were only able to carry out one activity related to DPIA in the sandbox: a workshop with stakeholders, all of whom represented groups who would be affected by the learning analytics system in various ways.

We have also been made aware of potential data protection risks through discussions in other parts of the sandbox process. We briefly go over these findings, which are also relevant for a data protection impact assessment, in the section headlined “Other identified risks”.

Workshop with stakeholders

The workshop included pupils, parents, teachers and data protection officers from municipalities working with the AVT project. The event began with a joint presentation of the learning analytics system, followed by discussions in smaller groups – one with a combination of children and adults and one with adults only. Most of the risks mentioned had already been identified by the participants in the sandbox project, but discussions with stakeholders contributed to further clarifying these risks. Findings follow below.

Risk of altered behaviour/cool-down effect

Cool-down effect: Unwanted change in one's own behaviour caused by uncertainties related to

- *who is processing personal data about us,*
- *which types of personal data about us are being processed,*
- *how personal data about us is being processed, and*
- *why personal data about us is being processed.*

(Read more about the cool-down effect in the Data Protection Authority's privacy survey 2019/2020.)

If pupils change their behaviour when they are working with digital learning tools because they feel the learning analytics system is monitoring them, there may be a cool-down effect. The pupils were especially worried about the system monitoring how long it took them to complete the task. They pointed out that if the time they spent working on a problem was recorded, it could lead to pressure of solving the problems as quickly as possible, at the cost of quality and learning from problem solving.

Another example of a cool-down effect could be that the pupils do not feel as free to “try and fail” in their problem-solving, because everything they do in the digital learning tools is recorded and may potentially affect the profile built by the learning analytics tool.

If the introduction of an AI-based learning analytics system in education leads to a cool-down effect, the AI tool may be counterproductive. Instead of the learning analytics system contributing to providing each individual pupil with an education adapted to their needs, the individual pupil adapts their education and behaviour to the system.

Risk of incorrect personal data in the system

A fundamental principle in the General Data Protection Regulation is that the personal data being processed must be correct.³ Incorrect and imprecise data in a learning analytics tool could have direct implications for the profile of each individual pupil. This could, in turn, affect the teacher’s assessment of the pupil’s competence and the learning resources recommended for the pupil.

The learning analytics system collects data on the pupils’ activities from the digital learning aids used by the school. One potential source of incorrect data, which was discussed by the adult participants in the workshop, is when a pupil solves problems on someone else’s behalf. This has likely always been a risk in education, and there is no reason to believe that a transition to digital activities has changed anything in this regard. However, the impact on the individual pupil may be far greater now, if the data from this problem solving is included in an AI-based profile of the pupil. As an example, the system may be tricked into believing that the pupil is performing at a higher level than they actually are, thus recommending problems the pupil does not yet have the skills to solve. This could be demotivating for the pupil, reinforcing the experience of not being able to master a subject or topic. A similar source of incorrect data is when a pupil deliberately gives the wrong answer, to manipulate the system into giving them easier or fewer activities. This, too, is a familiar strategy, used by children since long before the digitalisation of education. What both of these examples have in common is that the problems must be addressed both technologically and in raising awareness in general.

Risk of the technology causing the pupils unwanted stress

In using a learning analytics system, there is a risk that the distinction between a practice/learning setting and a test setting is blurred for the pupils. Teachers already use information from the pupils’ problem-solving and participation in class as a basis for assessment of what the pupils have learned. By using a learning analytics system, however, this assessment will be systematised and visualised in a different way, compared to the current situation. The pupils expressed concerns that there would be an expectation to show their “score” in the system to peers and parents, similar to how pupils currently feel pressure to share test results.

Other identified risks

Special categories of personal data

The AVT project does not plan to process special categories of personal data (sensitive data) in the learning analytics system. This type of processing requires a separate legal basis pursuant to Article 9. We discussed this issue in the sandbox, to promote dialogue and reflection around potential future needs, so that this could be addressed in the development phase to the extent possible.

Among other things, we discussed the possibility of the system being used to identify learning disabilities. Often, if the possibility to use a system for other useful purposes than what it originally was intended for exists, it is only a matter of time before it is.

³ Article 5 (1) (d) of the GDPR

It is also not unlikely that the education sector may be authorised to use artificial intelligence to uncover learning disabilities in the future. The earlier learning disabilities can be identified, the earlier the person with the learning disability can get the follow-up they need. The condition, of course, is that the municipality has a clear statutory authority and that the AI system has been developed responsibly, within the framework of relevant law.

The AVT project has been aware of this issue, and they plan to initiate measures to prevent the system from processing special categories of personal data. These measures include regular control of the algorithm and the data it generates. The solution does not currently have any mechanisms in place for blocking free-text responses from connected teaching aids, but the system does not process data from free-text fields in its learning analytics.

Processing of data about third parties

The processing of data about third parties was also discussed in the sandbox. This refers to the risk that municipalities process data about pupils' parents, siblings and friends in the system. The municipalities have no statutory authority to process personal data about anyone other than the pupils.

The AVT project has described that they do not believe the solution could include information about parents, siblings or other third parties. Many of the measures implemented to prevent processing of special categories of personal data would also be relevant in this context. Such as not processing data from free-text fields in the solution. It would, however, be necessary to review this regularly. Further development of the system and integration with new service providers will entail a risk that third-party personal data may be processed through the solution at some point in the future. If so, it will be essential to identify if, and if so, when, this happens.

Decision support

The AVT project assumes that teachers will be using this system as decision support, and not as an automated decision-making system. This is relevant for the assessment of whether the use is subject to Article 22 of the GDPR. This provision grants the right "not to be subject to a decision based solely on automated processing". There are several reasons why this was a topic of discussion in the sandbox. Among other things, users may want to use the system as a system for automated decision-making in the future. There had been several discussions within the AVT project previously, concerning the role of the teacher, and the importance of not marginalising the teacher with this learning analytics system.

Another risk is that the system is used as a de facto automated decision-making system, even if it is not designed to function that way. Some teachers may accept recommendations from the system without making their own assessments. This may be due to high workloads or a lack of knowledge about the algorithm, insight into how the system works, etc. We can also envision a situation where the recommendations from the system are so good that the teachers feel they cannot override the system.

In order to prevent a decision support system from slowly transitioning into a decision-making system, it is imperative to be aware of the risk and being conscious of this process. In addition, it will be important to implement measures to ensure that the system is actually used as decision support, such as the school owner developing procedures and training users on how to use the system, how it works and what the system's recommendations are based on. This dissemination of knowledge is important for enabling teachers to consider the system's recommendations in light of their own observations and knowledge of the pupils. Other relevant measures may include the implementation of procedures to reveal whether the decision support is real, or whether the system's recommendations, in practice, are entirely decisive for the teacher's decisions. To reduce this risk, the AVT project designed the learning analytics system so that the pupil or teacher has to actively choose whether to accept the system's recommendations.

Similar issues were also discussed in the NAV project. The report from this project is available on the Data Protection Authority website.

Transparency

Introduction

Transparency is a fundamental principle of GDPR. In addition to being a prerequisite for uncovering errors, discriminatory treatment or other problematic issues, it contributes to increased confidence and places the individual in a position where they are able to assert their rights and safeguard their interests.

The General Data Protection Regulation imposes specific requirements with which all those who process personal data must comply (see Articles 12, 13 and 14). These requirements include, among other things, that the data subject must be informed of the fact that personal data is being processed, how the personal data is processed and who the data controller is. This information must be easily accessible and must use clear and plain language. This especially applies to information provided to children.

The specific requirements for information that follow from the GDPR are considered minimum requirements. In some cases, it is necessary to provide additional information to comply with the fundamental principle of transparency, which was a topic for discussion in the AVT project.

Information adapted to users

It can sometimes be challenging to provide an intelligible explanation for how a system based on artificial intelligence processes personal data. For the AVT project, the age range of its users further complicates this situation. This system may potentially be used by children as young as six at one end of the range and by graduating pupils in upper secondary school at the other. The youngest children, especially, do not have the ability to understand how their privacy may be affected nor the ability to protect their own interests. Their parents must therefore also receive information, so that they can protect their children's interests.

As data controllers, the municipalities are responsible for complying with this information obligation. Norway has 356 municipalities, all with different prerequisites for implementing a learning analytics system. For that reason, the partners involved in the AVT project want to prepare an information packet, which includes some general information about learning analytics. In this context, they have prepared draft templates of written and graphic material that municipalities may use in their information activities. In the sandbox, we evaluated this information material and discussed potential solutions for how this information can be adapted to suit the needs of the different user groups.

For compliance with the information obligation, it is important to keep the intended recipients in mind. We cannot always know in advance who the recipients of the information are, e.g. because the controller does not always know who the users of a service may be. The AVT project has the benefit of having a clearly defined group of users for the learning analytics system. Their challenge is the considerable range of ages and levels of development within this group. One central discussion in the sandbox was how the AVT project can provide information that is simple enough for the youngest pupils, while also meeting the information needs of older pupils and parents.

These sandbox discussions can be summarised as follows:

- Use a language that takes into account the youngest pupils – adults also appreciate information that is simple and easy to understand.
- All of the information required by law must be included, but not necessarily in the same place and at the same time. Adults and children alike can lose heart if the document or online article is too long. One guiding principle may be to not only focus on what the pupils/parents need to know, but to also take into account when they need this information.
- It could be beneficial to provide information in layers, where the most basic information is presented first, while at the same time, giving the reader an opportunity to read more detailed information on the various topics. Take care to ensure that important information is not “hidden away” if this approach is used.
- Consider whether it would be appropriate to provide (or repeat) information when the pupils are in a setting where the information in question is relevant, e.g. with the use of pop-up windows.

- Use different approaches – what works for one group may not necessarily work for another. The AVT project has included text, video and images in their information material, and feedback from data subjects indicate that different user groups respond differently to different formats.
- Be patient and do not underestimate the complexity and demanding nature of trying to understand how the learning analytics system works, as well as the purpose and consequences of implementing this type of system. This applies to both children and adults.

The AVT project will continue its work on issues related to the obligation to provide and need for information after the sandbox project is completed.

Explaining the system’s underlying logic

For automated decisionmaking systems subject to Article 22 of the Regulation, it follows expressly from the GDPR that the data controller must provide relevant information about, among other things, the system’s underlying logic.⁴ Whether it is necessary to provide information about this logic if the processing does not entail automated decisionmaking or profiling, should be assessed on a case-by-case basis, based on whether it is necessary to ensure fair and transparent processing.

As previously mentioned, the learning analytics system in the AVT project is a decision support system, not an automated decisionmaking system subject to Article 22 of the GDPR. In the sandbox, we have not concluded whether the AVT project is legally obligated to provide information about the underlying logic of the learning analytics system. We did, however, discuss the issue in light of the objective of the sandbox, which is to promote the development of ethical and responsible artificial intelligence. In this context, we discussed how explanations that provide users with increased insight into how the system works, could increase trust in the system, promote proper use of the system and uncover potential defects.

But how detailed should an explanation of the system be? Is it sufficient to simply provide a general explanation for how the system processes personal data in order to produce a result, or should a justification for every single recommendation made by the system also be provided? And how does one provide the youngest pupils with a meaningful explanation? This sandbox project does not offer any final or exhaustive conclusions on these issues, but we did discuss benefits, drawbacks and various alternatives for solutions.

For the youngest pupils, creativity is a must when it comes to explanations, and these explanations do not necessarily need to involve text. For example, the AVT project created an information video, which was presented to various stakeholders. The video was well liked by the children, but garnered mixed reviews from the adults. The children thought it explained the system in a straightforward way, but the adults found it did not include enough information. This illustrates firstly, how different the needs of different people are, and secondly, how difficult it can be to find the right level and quantity of information. The AVT project has also considered building a “dummy” version of the learning analytics system, which allows users to experiment with different variables. This way, users can see how information fed into the system affects the recommendations in turn made by the system. Visualisation is often quite effective at explaining advanced technology in a straightforward manner. One could have different user interfaces for different target groups, such as one user interface for the youngest pupils and another aimed at older pupils and parents.

A privacy policy is useful for providing general information about the processing of personal data by the system. As previously mentioned, we have also discussed whether individual justifications for the system’s recommendations should be provided. This would be information about how the system has arrived at the specific recommendation and the data this recommendation is based on. Individual justifications could be easily accessible by users, but do not necessarily have to be presented alongside the recommendation. There are many benefits to providing justifications for the system’s recommendations. If pupils and teachers gain a broader understanding of how the system works, it could increase their trust in the system. The justifications could also make teachers better able to truly assess the recommendations made by the system, thus mitigating the risk of some using the system as an automated decision-making system. Finally, justifications may help uncover flaws in the system.

⁴ Article 13 (2) (f) and Article 14 (2) (g) of the GDPR

Going forward

In this sandbox project on “Activity data for assessment and adaptation”, we have explored the legal basis for processing, transparency and protection of pupil privacy through a data protection impact assessment.

The most intense debate has been focused on legal basis, where the participants have differing views. The Data Protection Authority has not concluded on what would be the most appropriate legal basis for those implementing the learning analytics system. That is neither our mandate for the sandbox, nor a goal in and of itself. The most important focus has been to explore challenges and opportunities within the regulatory framework, as well as to highlight relevant aspects to take into account when considering the use of artificial intelligence in education. The many good discussions have enhanced the competence on all those involved in this sandbox project. The fact that the legal basis for processing is such a controversial topic, in the NAV and Secure Practice projects as well as this one, is a finding in and of itself. Considerations of legal basis can be difficult, especially when artificial intelligence is involved.

The Ministry of Education is currently working on a new education act, which will not be ready until 2023, at the earliest. The education sector, therefore, currently has a unique opportunity to clarify how schools can contribute to the development of educational AI tools.

We have also documented the benefit of involving various stakeholders in assessments of risks and impacts to the pupils' privacy. The pupils contributed with scenarios the adults had not necessarily considered. Parents, data protection officers and others helped emphasise the need for information that artificial intelligence brings. We have identified several challenges inherent in this use, including in exploring how information can be adapted to different user groups. Needs vary considerably, even within groups. Younger children's needs differ from the needs of older children, and teachers' needs differ from the needs of parents. Discussions on transparency also highlighted resource issues. The preparation of information comes at a cost, and resources are not infinite. Going forward, the AVT project will have to find a happy medium between what the law requires and what the stakeholders need to know.

The sandbox project is completed, but the exploration of artificial intelligence in schools has only just begun. The AVT project is blazing the trail in a field that will only become more important in the future. The project's contributions in the sandbox has shed light on an area the Data Protection Authority does not have much experience with. The opportunities are endless, and the importance of a sober, ethical approach could not be more clear. The AVT project has taken its responsibilities seriously and is likely building the foundation for a better education for our children in the time to come.



**Norwegian Data Protection
Authority regulatory sandbox
for responsible artificial
intelligence**

Office address:
Trelastgata 3, Oslo
Norway

Postal address:
PO Box 458 Sentrum
0105 Oslo

sandkasse@datatilsynet.no
Telephone: +47 22 39 69
00

datatilsynet.no/sandkasse
personvernbloggen.no