Copilot through the lens of data protection

Final report from the sandbox project with NTNU

Topics: Generative AI, information management and data protection impact assessments (DPIAs)

Datatilsynet

Content

Summary

Generative artificial intelligence (AI) is no longer just a fun tool on the side, it is now being integrated into the digital solutions we already use. Microsoft launched its Copilot for Office Suite in November 2023, offering the potential to significantly simplify working life. Some people have started using it to varying degrees, while others are still sitting on the fence. For what actually happens when you turn on Microsoft 365 Copilot?

The Norwegian Data Protection Authority and the Norwegian University of Science and Technology (NTNU) have looked at what data protection requirements apply and what assessments NTNU should carry out before they use Microsoft's AI assistant. NTNU conducted a parallel pilot project to examine whether it is ready to introduce M365 Copilot, as well as to propose a framework for management, operation, maintenance and development. NTNU has published its own findings report. It provides an overview of how M365 Copilot functions and a great deal of insight for others who are considering turning on Copilot.

Read NTNU's own findings report (in Norwegian only)

The Norwegian Data Protection Authority supports NTNU's report but recommends a more specific approach when conducting a data protection impact assessment (DPIA). Each organisation must carry out its own DPIAs based on what data they have and what tasks they wish to use M365 Copilot for.

M365 Copilot is an active component that retrieves and recreates information in new and unfamiliar ways. It is a challenge that this new technology's ability to formulate language well – including Norwegian – can make it seem human, as though it is capable of assessment and logical reasoning.

It is also important to emphasise that this is pioneering work. To the best of our knowledge, no other supervisory authority has looked at the use of M365 Copilot in relation to data protection regulations. This report should be seen as a first step in understanding and assessing whether such tools can be adopted in a (cautious and step-by-step) manner that is compliant with data protection regulations.

Main points

1. **M365 Copilot requires that the organisation's data are already stored in Microsoft's cloud solution.** M365 Copilot sits at the top of Microsoft's M365 cloud solution. Before introducing M365 Copilot, it is a prerequisite that you have conducted all necessary security and data protection assessments related to the M365 platform itself. You also need to have the necessary resources and expertise to manage service providers and the cloud solution in a responsible manner over time, particularly because of frequent changes from the supplier side. Responsibility for the data used in Copilot rests with the organisations that use the tool.

   The Norwegian Agency for Public and Financial Management (DFØ) has created a guide for public enterprises on the procurement of cloud services (in Norwegian only), which may be of help. See also section 4 of NTNU's findings report for further information and recommendations from NTNU on the management of the system.

2. **Get your own house in order.** Copilot will have access to the same information that the user of the tool has. That means that challenges and weaknesses in the 'digital foundation', such as poor access management and control of personal data, will be made visible and significantly amplified by M365 Copilot. It is important to emphasise that Microsoft, as a service provider, largely assumes that 'everything is in order' with the management of the underlying M365 platform, in order for Copilot to be used responsibly and in a lawful manner. It is therefore important to get your own house in order first, and any project to introduce it will probably require thorough (re)assessments of your own information management. This requires a certain amount of effort and resources but is a critical and necessary step when introducing new technology.

   The Norwegian Digitalisation Agency has prepared a guide to information management, (in Norwegian only), and information about what a record of processing activities must contain can be found on the Norwegian Data Protection Authority's website (in Norwegian only).

   NTNU has concluded that they are not yet ready to introduce M365 Copilot throughout their organisation, partly because they consider that their own house is not yet in order.

3. **Identify and limit what M365 Copilot will be used for.** Consider which tasks and associated processing of personal data M365 Copilot should and should not be used for. Some tasks are poorly suited for the use of generative AI, for example when it is important that the responses are correct and the user does not have the time or expertise to check what is generated. In addition, using M365 Copilot in e.g. HR and personnel management poses a particularly high risk to data protection. This is because access to personal data is difficult to manage and control, and because the consequences for individuals can be very severe. Tasks involving special category (sensitive) personal data should also be carefully assessed or avoided in connection with the use of M365 Copilot.

   Map and describe the processing operations that will occur if M365 Copilot is used for a specific purpose, i.e. from when a prompt is given to Copilot until it provides an answer. The record of processing activities is an appropriate place to start, where one can review and assess each instance of processing per purpose. That will provide a good starting point for assessing which tasks you can and would like to use Copilot for.

   If M365 Copilot has access to information that contains (sensitive) personal data, the information must be classified, identified and labelled, at least at the document level. We emphasise that Microsoft acknowledges that this necessary in order for M365 Copilot to be used responsibly.

4. **Assess the legal basis.** When tasks and associated processing of personal data are considered "M365 Copilot candidates", the legal basis for processing must be checked. For existing processing activities, you must assess whether using M365 Copilot would result in any changes to the processing, such as which or whose personal data are being processed. If there are changes, you must assess whether the existing legal basis for processing could still be used, including whether the processing remains 'necessary'. If this is not the case, M365 Copilot cannot be used for this processing.

   Processing personal data for new purposes requires the identification of an appropriate legal basis for processing. In cases involving the re-use of personal data for new purposes, as will often be the case, you must assess whether the new processing is compatible with the original purpose.

5. **Assess the impact to data protection.** As a general rule, a data protection impact assessment (DPIA) will be required when using generative AI that processes personal data. That is because the law emphasises 'using new technologies' as a particularly important factor, and because the understanding of risks associated with generative AI is still immature. A DPIA must be carried out for each processing operation or set of similar processing operations. Tasks that do not in themselves require the processing of personal data may nonetheless do so when using M365 Copilot, because Copilot uses all of the information accessible to the user and can thus link it to personal data.

   The DPIA process should identify technical and organisational measures that can reduce the risk to an acceptable level, and these must be in place before M365 Copilot is taken into use. Testing can be used as a measure to minimise risk. If the risk is too high even after measures are implemented, it is probably best not to use M365 Copilot for the processing in question. Alternatively, contact the Norwegian Data Protection Authority for a prior consultation.

6. **Will using it violate the Norwegian e-mail regulation?** M365 Copilot logs all user interactions. The log is stored in the user's own digital area and, in NTNU's case, is accessible to the M365 administrators. Overall, we consider it likely that this log of user interactions could fall under the prohibition against monitoring employees' use of electronic equipment. However, we understand that the primary purpose of such a log is to ensure that the quality of the service is as it should be. That purpose may fall under the prohibition's exemption that applies to managing the organisation's computer network. Whether the other exemption, to 'detect or resolve security breaches in the network', may be applicable must be specifically assessed in relation to the purpose of such log.

7. **The use of large language models requires expertise and awareness.** Large language models provide a new user experience for many, with both opportunities and limitations that remain unclear. It can be difficult to understand what information is used to form the basis of statements. Expertise is required to formulate prompts that produce relevant and good answers. It is the responsibility of the organisation deploying M365 Copilot that users of the solution have sufficient knowledge, awareness and training in its use. This competence not only ensures that what is generated is of a high quality, but also that the solution is used in a way that safeguards data protection.

8. **Consider alternative solutions**. M365 Copilot can be used for a great many things, so it is no small task to ensure that the system is used in a responsible and lawful manner. Some of Copilot's features may challenge the principles of purpose limitation and data minimisation. Measures that could in theory reduce risks and consequences may in practice be very difficult to implement. It is therefore important to consider whether other

AI solutions that present lower risks to data protection are able to meet your specific needs. These might include solutions that transcribe audio recordings, customised chatbots or support tools tailored to specific purposes and limited to carefully selected and quality-assured internal information sources.

9. **Introduce Copilot in small and controlled steps.** It is possible for Norwegian organisations to start using M365 Copilot, but not by everyone and not for everything. We strongly recommend that these types of solutions are introduced in small and controlled steps, using selected roles carrying out suitable processing operations in the organisation first. Structured plans must also be created for post-implementation monitoring and following up the quality of what the solution produces, through both organisational and technical measures.

The road ahead

NTNU has done an impressive, socially beneficial and extensive job of acquiring knowledge and awareness of the use of large language models in general, and integrated AI solutions such as M365 Copilot in particular. They have chosen not to introduce M365 Copilot throughout the entire organisation, but instead introduce the tool in small and controlled steps, limited initially to selected roles.

M365 Copilot is still in the early stages of development and does not provide control at a granular level, such as the ability to make local and flexible adaptations (e.g. disabling access to users' mailboxes or specific deletion policies). Microsoft probably considers unlimited access to the user's mailbox as an important and central feature, although it is perhaps one of the features that creates the most uncertainty for many organisations.

The Norwegian Data Protection Authority expects the issues that customers, organisations, authorities and wider society identify in the product are taken seriously by the product supplier. At the same time, there are clear requirements for organisations that wish to benefit from using the tool. The prerequisite of having an extremely well-functioning information management system may make it difficult to succeed with such solutions, but obviously has a positive upside that goes far beyond the implementation of one specific solution.

Note:

An assessment of cloud services in general, the transfer of personal data to third countries and Microsoft's role(s) under the GDPR was outside the scope of this project. We would however like to mention that the European Data Protection Supervisor (EDPS) recently made a decision that partly encompasses Microsoft's role in the provision of cloud services to a number of EU bodies. The decision has been appealed by both Microsoft and the European Commission. The outcome of the case may have an impact on how the use of cloud solutions must be organised in the future to be in line with the GDPR.

About the project

NTNU is an international university based in Trondheim with campuses in Gjøvik and Ålesund. The university's main profile is in science and technology, and offers a variety of programmes of professional study. Its academic breadth includes the humanities, social sciences, economics, medicine, health sciences, educational science, architecture, entrepreneurship and the arts. NTNU has 9,000 staff members and 43,000 students.

The main purpose of the project was to investigate whether and how a large public organisation like NTNU could use M365 Copilot. It is important to note that Microsoft uses the term 'Copilot' in various ways and several services operate under that name. This project is specifically about M365 Copilot, which integrates AI into existing Microsoft 365 services.

NTNU examined a number of different problems relating to the use of AI tools in the public sector. A key challenge was whether M365 Copilot can be used without personal data being processed in conflict with the GDPR. Another question was whether people would accept that their data could be used in contexts other than those for which they were originally collected. In addition, there are several ethical and organisational challenges related to the use of generative AI tools in general. NTNU also wanted to investigate risks associated with incorrect decisions as a result of, for example, discrimination and so-called 'hallucination'.

NTNU also wanted to develop a toolbox with guidelines, frameworks and DPIAs that could be used by other public and private organisations. The goal was to make it easier to assess whether and, if relevant, how generative AI tools such as M365 Copilot can be implemented in the public sector in a responsible manner. NTNU also wanted to look at how suppliers could be influenced to consider data protection by design and by default early in the development process, with the aim of preventing privacy issues being an afterthought towards and the end of the procurement process.

Objective of the sandbox project

The scope of NTNU's project was broad and covered more topics than just data protection. It was therefore important for the Norwegian Data Protection Authority to narrow down the scope of our involvement and assistance. The main objective has been to explore and clarify **what data protection regulations require** for NTNU and other public organisations to use tools such as M365 Copilot in a responsible and legal manner.

To do this, it has been necessary to look at:

1. What M365 Copilot actually is and how it works, as well as generative large language models in general.
2. How M365 Copilot can be understood in light of the data protection regulations at a high level.
3. What prerequisites must be in place, including 'getting your own house in order'.
4. Whether one or more DPIAs are required, and what is particularly relevant to consider in light of M365 Copilot.
5. Application of the Norwegian e-mail regulation.

Processing of special categories of personal data, cloud services in general, the transfer of personal data to third countries, and Microsoft's role under the GDPR were outside the scope of the project.

Relationship to NTNU's own findings report

NTNU published its findings report on 17 June 2024 in order to share its experience of M365 Copilot with other organisations. The report presents eight main findings that address not only data protection, but also ethical, legal, technical, and organisational issues.

The findings report may support and inspire both public and private organisations in their planning and assessment of generative AI tools, as well as contribute to the development of risk mitigation measures. In particular, we emphasise:

- NTNU's toolbox, which provides information about what generative AI is and how to use it in a smart, safe and secure way.
- NTNU's AI journey, with suggestions for an AI strategy, assessments per service and tips on procurement (NTNU's findings report, pages 29–36).
- NTNU's proposal for guidelines for generative AI.

We recommend that NTNU's findings report be read in addition to this final report, which is intended to supplement the NTNU report in selected areas.

NTNU's toolbox contains a data protection impact assessment (DPIA). NTNU chose to make an 'overall' assessment of M365 Copilot in the operational phase, where they looked at the technology as a whole. NTNU has not considered specific processing operations in light of specific purposes. This means that the work does not meet the requirements for what a DPIA must contain pursuant to Article 35 GDPR. However, the work provides information about NTNU's experience of M365 Copilot and how the tool works in general, which can be useful if a DPIA is to be prepared.

Limitation

M365 Copilot is a multi-functional tool that can be used for a variety of tasks. Given its wide range of functions, it can also be used to process personal data as part of its operations. It is not possible to conclude on a general basis that the tool can be used in accordance with data protection regulations. Although certain tasks performed by M365 Copilot do not in principle require the processing of personal data, such processing will almost always occur due to the tool's inherent characteristics.

In this final report, we have chosen to focus on a few basic topics. First, we explain what M365 Copilot is and how it works. We then provide a general description of how we understand M365 Copilot in light of data protection regulations, together with a review of key terms and factors to be aware of. We also look at the basic prerequisites for using M365 Copilot, including the need for 'getting your house in order'. These assessments are also relevant for other AI tools. Finally, we highlight the importance of data protection impact assessments (DPIAs) and what to be aware of if you are considering using M365 Copilot.

NTNU tested M365 Copilot on three use cases: 'beginning an official study', 'minutes function' and 'case management via e-mail'.[1] These use cases were chosen because they may also be relevant for other public sector organisations. In this final report, we have taken NTNU's use cases as a starting point, but we have made use case C a little more specific in order to have a clear and distinct purpose. This example does not necessarily reflect NTNU's actual work process, but is used for the sake of illustration. It is important to note that each use case may involve several types of processing of personal data.

Use case A:	A researcher uses M365 Copilot to collect information (data collection) before the official study can start. The researcher has access to information from the internet, previous documents they have written themselves or documents they have access to (but written by others). M365 Copilot can help the researcher to get an overview of relevant data material in order to make necessary assessments in line with the instructions for the study, get help with the actual writing (drafting) and proofreading/improving the language.[2]
Use case B:	An employee is assigned responsibility for ensuring that agreement on something is reached in an internal meeting between two or more parties. They call a digital Teams meeting, or a physical meeting where Teams actively listens to the meeting. The meeting is recorded and transcribed. M365 Copilot uses the transcript, the calendar invitation information, and 'nearby documents' to summarise the meeting.[3]
Use case C:	An employee is going to assess whether an application for admission to a master's programme submitted by email is complete (i.e. contains all the required information) and respond to the application by either confirming that the application is complete or requesting more information.

[1] NTNU's findings report, pp. 43–51.

[2] Ibid p. 43.

[3] Ibid p. 46.

What M365 Copilot is and how it works

Microsoft 365 Copilot differs from the chatbots you may already know in that large language model technology is integrated into the Microsoft 365 applications many people already use every day, such as Word, Excel, PowerPoint, Outlook and Teams. It thus represents an important trend referred to as 'integrated AI', i.e. AI solutions and functionality that appear in products you already use. Sometimes this is a functionality you can choose separately (M365 Copilot requires its own additional license). Other times, you get the technology 'thrown in' whether you want it or not. An example of the latter is the formerly named 'Bing chat', which is now also called 'Copilot', but which is a separate chatbot solution similar to, for example, ChatGPT. This type of solution can answer general questions and does not by default have access to the user's and the organisation's information.

M365 Copilot, on the other hand, acts as a personal assistant and combines the functionality of large language models with an organisation's internal data and information structures, made available via knowledge graphs and indexed semantic searches (RAG – see Appendix 1). The tool can also be configured to retrieve external information (from the Internet) to 'enrich' internal searches. It is worth noting that one of the potential challenges associated with this is that prompts that are initially perceived as internal or even confidential are suddenly exposed externally as searches in open search engines. M365 Copilot is a product that requires the organisation to have extremely good control of the solution's abilities and limitations in practice.

It is worth noting that the term 'Copilot' is used by Microsoft in many different contexts and does not refer to a specific product. Rather, it is a support tool implemented in different ways in different parts of the Microsoft 365 platform. The promise from Microsoft is that M365 Copilot can help you with everything from generating content and analysing data to improving communication and collaboration in your organisation – in return for giving it access to all the information you have access to.

Knowledge graphs (Microsoft Graph) and semantic searches constitute a data platform that connects data and services across Microsoft 365. It collects information from emails, calendars, documents, meetings, chats and more. Through Microsoft Graph, Copilot can access and understand the context of the information a user has access to. In addition, the solution builds an experience-based profile of the user over time, which is intended to provide more relevant and personalised assistance.

Because M365 Copilot is tightly integrated into the Microsoft 365 applications, the solution is adapted to the interface and work tasks of the various tools and it can therefore assist directly in the application without having to move information between different applications or interfaces.

By automatically retrieving information from meetings, emails, documents and chats (within given access rights) in combination with user profiling over time, the solution has the potential to provide more tailored assistance. This rather extensive access to the individual user's and the organisation's information will make it possible to automate complex tasks:

- **Outlook**: Summarise long email threads. Generate email response suggestions based on context, tone, and previous communication. Identify and suggest calendar appointments or tasks based on the content of your emails.
- **Word**: Prepare draft documents based on prompts or key points. Improvement of text, wording, grammar and style, and adaptation of tone to the target audience. Summary of documents and reports.
- **Excel**: Interpret datasets, identify trends and patterns, and present layouts in a simple manner. Create complex formulas by describing what you want to achieve in natural language. Suggest appropriate charts and graphs to represent data visually.
- **PowerPoint**: Generate entire presentations based on a document or idea, including suggested text, images and design. Provide recommendations for layout, colours and graphics. Convert a document or report to a presentation by extracting the key points.
- **Teams**: Summarise discussions in real time, write minutes, and identify action items. Generate a list of next steps and assign tasks to team members.
- **SharePoint**: Help to create and edit content, including text suggestions and structuring. Semantic searches to find relevant content based on meaning, not just keywords. Analyse content across SharePoint to identify knowledge gaps or overlapping information.

In other words, a lot of functionality is promised, and it is all made possible by two things: the power of large language models and more or less free access to all the information the individual user in the organisation already has access to.

How can M365 Copilot be understood in light of data protection regulations?

We begin by highlighting some key concepts and terms from data protection legislation that we believe are important to keep in mind when considering implementing and using M365 Copilot in your organisation. We hope this will help to avoid misunderstandings from the start that could lead to consequential errors.

Key concepts and terms

Personal data	This is defined in Article 4(1) GDPR as 'any information relating to an identified or identifiable natural person'. Even if a piece of information about someone is incorrect, for example when a large language model has generated a fact about an individual that is wrong, it counts as (incorrect) personal data. The same applies to predictions and assumptions about a person.
The data subject	An identified or identifiable natural person (Article 4(1) GDPR). In other words, it is the individual to whom information can be linked.
Erasure and accuracy	Every reasonable step must be taken to ensure that personal data that are incorrect, having regard to the purposes for which they are processed, are erased or rectified without delay, in accordance with the accuracy principle (Article 5(1)(d) GDPR). This means that users must be adequately trained, and NTNU must have procedures in place to reduce the risk of M365 Copilot generating incorrect personal data. If this nevertheless happens, the personal data must be erased or rectified without delay.
Processing	Article 4 (2) GDPR defines processing as 'any operation or set of operations which is performed on personal data or on sets of personal data, whether or not by automated means, such as (...)'. The legislator has deliberately given the term 'processing' a wide scope. This is evident both from the term 'any operation' and from the non-exhaustive nature of the definition, made clear by the use of 'such as'.[4] Any processing of data must have a legal basis pursuant to Article 6 GDPR. To determine the correct legal basis, the purpose of the processing and which personal data are to be processed must first be clarified. Which specific processing operations will take place must also be identified before the legal basis can be assessed and selected. In addition, if special categories of personal data are to be processed, a valid exception to the prohibition in Article 9 must be identified. 'M365 Copilot in the operational phase' or 'introduction of M365 Copilot' is not a specific processing operation. [5]
Purpose of the processing and the purpose limitation principle	'Purpose' is the very cornerstone of the GDPR. **The purpose** is the reason why a processing operation takes place, and it is the purpose that sets the limits for which personal data are to be processed and how. The GDPR states that personal data shall be collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes (Article 5(1)(b)). In the context of M365 Copilot, the terms 'specified' and 'explicit' are especially important. The purpose must be determined at the time of

[4] Judgment of 24 February 2022 [C5], *Valsts ieņēmumu dienests*, C-175/20, EU:C:2022:124, para. 35.

[5] NTNU's findings report, p. 102.

	collection of the personal data at the latest,[6] unless a new purpose is compatible with the original purpose pursuant to Article 6(4) GDPR. It is necessary to look at the purpose in order to comply with, among other things, the data minimisation principle, where personal data must be adequate, relevant and limited to what is strictly necessary **to achieve the purpose**.
The data minimisation principle	The data minimisation principle means that personal data must be adequate, relevant and limited to what is **necessary for the purposes** for which they are processed. It is therefore necessary to consider the purpose (which should already be identified in line with the purpose limitation principle) when assessing which and whose personal data are adequate, relevant and necessary **to achieve the purpose**.
Recipient	The concept is defined in Article 4(9) GDPR as 'a natural or legal person, public authority, agency or another body, to which the personal data are disclosed, whether a third party or not.' M365 Copilot can easily be personified, because it poses as a natural person through the way it answers. This can lead to an incorrect way of thinking when the tool is assessed from a data protection perspective and can lead to consequential errors in subsequent assessments. For example, M365 Copilot is not a 'recipient' within the meaning of the GDPR.
Records of processing activities	Each data controller is obliged to maintain a record of its processing activities, including the purposes of the processing, as well as which and whose personal data are processed to achieve the purpose (Article 30 GDPR). The information in the record of processing activities largely coincides with what the data subjects must be informed about. Some new processing activities will inevitably follow from the introduction of M365 Copilot, such as logging interactions (the content of interactions log). New processing activities must be recorded in the record of processing activities and the data subject must be informed in accordance with Article 13 GDPR.

Map and describe the processing

The provisions of the GDPR relate to 'processing' of personal data, as defined in Article 4(2) GDPR (see above). M365 Copilot is not in itself a processing operation, but a tool or set of functions – i.e. the 'means' – that can be used to process personal data in many different ways and for different purposes. However, the predefined purpose, and what is necessary to achieve this purpose, limits which personal data can be processed and in what way. The first thing to do is therefore to map and describe the processing that will take place if M365 Copilot is used in connection with a specific purpose, i.e. from when a prompt is entered into M365 Copilot until it generates an answer. Often one will want to use M365 Copilot in connection with processing that is already taking place, and by mapping what is new when taking M365 Copilot into use, you will be able to compare the 'old' processing and the 'new' processing, and then identify the new processing operations ('set of operations') that may or will occur.

A systematic description of the 'new' processing offers several advantages.

- **Choice of legal basis:** It will be possible to determine which legal basis is most suitable for the processing.
- **Assessment of necessity and proportionality:** A comparison between the old and new processing helps to assess the need for and the proportionality of the new processing.
- **Risk mitigation:** It will be possible to identify what technical or organisational measures should be put in place to mitigate risk, for example by having specific guidelines or procedures for effective prompt design, by

[6] Judgment of 24 February 2022 [C5], *Valsts ieņēmumu dienests*, C-175/20, EU:C:2022:124, para. 64.

changing what kind of access a particular user role should have, or by turning available settings on or off in M365 Copilot.

NTNU looked at three selected use cases in the sandbox project (NTNU's findings report, pp. 43–51). However, they deliberately chose not to consider them in their DPIA. Instead, they look at the product at higher level. NTNU should specify and describe each new processing operation that will take place when M365 Copilot is used.

The record of processing activities may be an appropriate place to start. More information about systematic descriptions of processing can be found in the Norwegian Data Protection Authority's DPIA checklist (in Norwegian only).

Often, several processing operations are carried out with personal data for one specific purpose. When a case officer is processing an application for admission to a study programme and must respond to an enquiry, they can first look for relevant information in the databases they have access to. This may include previous correspondence with other applicants about similar enquiries, previous decisions or internal guidelines. Such searches involve the processing of personal data, the results of which may contain both relevant and irrelevant information and personal data. It is the case officer who decides what is relevant and what they want to use going forward. One of the innovations with M365 Copilot is that such processing operations are done automatically, and the content is summarised and made available in a different format than before. It is not a given that this entails a *new* processing operation, but it must be considered. For example, if more personal data are processed or personal data are collated in a different way when M365 Copilot is used in connection with the task, it is important to identify and describe these new processing operations.

Example: In **use case A**, the task does not necessarily involve the processing of personal data without the use of M365 Copilot.

By using M365 Copilot, information about the user and possibly others may be processed because M365 Copilot will look for and use information it finds in 'nearby documents' (emails, chats, calendar invitations etc.) to enrich prompts and create outputs that are more relevant to the user. The scope of personal data that can be processed will depend on the individual user's access control, and may also be affected by the M365 Copilot settings (for example, by turning off 'Graph-grounded chat') or by the use of 'prompt engineering' (i.e. how the prompt is designed). The purpose of the processing can be described as helping the user to write report more effectively. This can be considered a completely new processing operation, for a purpose that did not exist before M365 Copilot was used in connection with the task.

Example: In use case B, the task involves the processing of personal data without the use of M365 Copilot, and should already be described in the record of processing activities. This may be described as follows: The purpose of keeping minutes of internal meetings is to document internal decisions made in the organisation. The meeting participants' names, roles and (a summary of) what was said in the meeting are recorded in writing and stored in a place where those who have an objective need for it have access.

When using M365 Copilot, new processing operations will take place in the form of recording and transcribing the meeting that will involve the processing of more personal data than before, such as voice, tone of voice, form of expression, gender (assumption from voice), as well as (personal) data that M365 Copilot finds when looking for 'nearby documents'. In addition, the optional settings can influence what other processing operations take place (e.g. an overview of who is talking, when and for how long).

Example: In **use case C,** the task involves the processing of personal data also without the use of M365 Copilot, and the processing should already be described in the records of processing activities. The purpose of the processing is to process applications for admission to a study programme. The case officer must assess whether all the necessary information is included in the application and respond to the enquiry as part of the case processing. Personal data received in the email are used to assess whether all the required information has been received.

When using M365 Copilot, new processing operations may take place, but this must be considered in relation to how applications are currently processed, including any existing processing activities related to, for example, searches. It is important to investigate whether

the scope of personal data processed will be expanded. Information about the applicant, the user and possibly others may be processed because M365 Copilot will look for and use information it finds in 'nearby documents' (emails, chats, calendar invitations etc.) to enrich prompts and create outputs that are more relevant to the user, such as a tailored answer directed at the applicant but similar to answers given to previous applicants. In addition, the optional settings can affect what other processing operations take place (e.g. by turning off 'Graph-grounded chat'), or by using 'prompt engrineering' (i.e. how the prompt is designed).

It is also important to consider new processing operations that will occur regardless of the task for which M365 Copilot is used, such as the content of interactions log and, if relevant, profiling of the user, which may occur for other, new purposes.

Assess the legal basis

According to Article 6 GDPR, the processing of personal data is only lawful if one of the conditions set out in (1) (a) to (f) are met. In its DPIA, NTNU has described that it is difficult to define one or more clear and distinct purposes for using M365 Copilot. NTNU concludes that the legal basis for 'M365 Copilot in operational phase' is legitimate interest pursuant to Article 6(1)(f) GDPR. It is important for us to point out that this is not in line with the GDPR because 'M365 Copilot in the operational phase' does not constitute 'processing operation'. M365 Copilot can be used as a means of performing many different processing operations for different purposes with different legal bases.

It is important to know what legal basis can be applied for each planned processing operation preferably before the DPIA stage, and at least before using M365 Copilot. If special categories of personal data are going to be processed, such processing must be based on one of the exemptions listed in Article 9. When the same personal data are processed for different purposes, the processing for each of those purposes must have a legal basis.[7]

If it concerns an existing processing operation, the conditions of the original legal basis must be reassessed, based on the description of the new processing operations.

All alternatives in Article 6(1) (b) to (f) GDPR contain a condition that 'processing [of personal data] is **necessary**' (our emphasis). The necessity condition will be met if the *purpose* of the processing cannot *reasonably* be achieved *as effectively* by other means that are *less restrictive* of the rights and freedoms of data subjects.[8]

The necessity condition must be interpreted restrictively, as it allows the processing of personal data without the data subject's consent.[9] Necessity must also be considered in the context of the data minimisation principle set out in Article 5 (1)(c) GDPR, which requires personal data to be adequate, relevant and limited to what is necessary for the purposes for which they are processed.[10] The data minimisation principle is an expression of the proportionality principle.[11] Among other things, proportionality requires that the benefits of restricting a right are not outweighed by the disadvantages of exercising that right.

We discuss the importance of the necessity condition in the context of Article 6(1) (e) and (f) GDPR below.

In some Norwegian legal sources, it has been argued that effective case management in public administration can be regarded as a 'substantial public interest' under Article 9(2)(g) GDPR.[12] In our view, effectiveness may in some cases be

---

[7] EDPB Guidelines 1/2024 on processing of personal data based on Article 6(1)(f) GDPR, p.6, via https://www.edpb.europa.eu/our-work-tools/documents/public-consultations/2024/guidelines-12024-processing-personal-data-based_en.

[8] Judgment of 4 July 2023 [GC], *Bundeskartellamt*, C-252/21, EU:C:2023:537, para. 108 and judgment of 22 June 2021 [GC], *Latvijas Republikas Saeima*, C-439/19, EU:C: 202i:504, para. 110 and para. 113.

[9] Judgment of 4 October 2024 [C5], *Koninklijke Nederlandse Lawn Tennisbond*, C-621/22, EU:C:2024:857, para. 31.

[10] Judgment of 4 July 2023 [GC], *Bundeskartellamt*, C-252/21, EU:C:2023:537, para. 109.

[11] Judgment of 22 June 2021 [GC], *Latvijas Republikas Saeima*, C-439/19, EU:C:2021 504, para. 98.

[12] Proposition No 135 (Bill) to the Storting 2019–2020 section 5.3.3.

interpreted as being a part of a purpose based on Article 6(1)(e) GDPR. This may be relevant for NTNU as a public institution.

NTNU's assessment of the necessity condition under Article 6(1)(e) should, among other things, take the following into account:

- Is M365 Copilot suitable to achieve NTNU's purpose in a better way?
- To what extent will NTNU be better placed to achieve the purpose of the processing if M365 Copilot is used?
- Are there any other ways that NTNU can reasonably achieve the purpose just as well?
- How much more restrictive are the new processing operations for the data subjects' rights and freedoms?
- Can NTNU implement any measures to make processing with Copilot less invasive?

If it is not possible to assess whether the necessity condition is met at this stage, it may be considered at the DPIA stage where the following are considered: 'the necessity and proportionality of the processing operations in relation to the purposes', 'an assessment of the risks to the rights and freedoms of data subjects' and which measures can be implemented to manage the risks and ensure the protection of personal data (Article 35(7) GDPR).

If the necessity condition cannot be met, even after measures identified in the DPIA are implemented, M365 Copilot cannot be used for the applicable processing.

For processing currently performed by NTNU for a purpose it has decided itself and which is based on the pursuit of a legitimate interest pursuant to Article 6(1)(f) GDPR, NTNU may consider including efficiency as a legitimate interest it wishes to pursue. The Court of Justice of the European Union has stated that making a service more efficient cannot be ruled out as a legitimate interest.[13] However, this will often involve adjusting the purpose of the processing and expanding which processing operations are necessary to achieve the legitimate interests. This is, however, contingent on:

- the processing not being carried out by public authorities in the performance of their tasks (Article 6(1) second subparagraph GDPR);
- the new purpose being compatible with the original purpose if, as will often be the case, the personal data to be processed were collected for another purpose (Article 6(4) GDPR);
- NTNU conducts a new, updated balancing of interests' assessment [14] that falls in NTNU's favour; and
- NTNU complies with all the other obligations in the GDPR.

If one of the above conditions cannot be met, even after measures identified in the DPIA are implemented, M365 Copilot cannot be used for the applicable processing.

For processing for a new purpose, NTNU must be able to identify a legal basis for the processing in the usual manner.

In order for NTNU to use consent as a legal basis for processing, it must be voluntary, specific, informed and unambiguous. In the context of M365 Copilot, this means, among other things, that NTNU must be able to clearly explain to the data subject how personal data are to be processed when M365 Copilot is used, thus ensuring foreseeability for the data subject. This may be difficult, especially if the data subject has little knowledge of generative AI and how M365 Copilot works. The power imbalance between NTNU and the individual must also be assessed. For example, public authorities or employers will normally not be able to use consent as a legal basis for processing since the data subject is in a dependent relationship. This does not mean that the use of consent as a legal basis for processing can be completely ruled out in relation to Microsoft 365 Copilot, but whether the consent conditions can be met must be considered specifically per use case and its associated processing operations. There may also be cases in employee-employer relationships where the employer can demonstrate that consent is voluntary, with no disadvantage to the employee if they do not consent to the processing.[15] It is also important to remember that a data subject can only

---

[13] Judgment of 4 July 2023 [GC], *Bundeskartellamt*, C-252/21, EU:C:2023:537, para. 122.

[14] The European Data Protection Board's guidelines 1/2024 on the processing of personal data based on Article 6(1) (f) GDPR is on public consultation and provides guidance on how to balance interests. They are available at https://www.edpb.europa.eu/our-work-tools/documents/public-consultations/2024/guidelines-12024-processing-personal-data-based_en.

[15] Guidelines 05/2020 on consent under GDPR, Section 22.

consent to the processing of their own personal data, while the use of Copilot may often involve the processing of several people's data, even if the input or output only concerns one person.

Careful consideration must therefore be given to whether consent is an appropriate legal basis in light of the individual processing operation in question, and whether the conditions for consent can be met.

Get your own house in order

Getting your own house in order is a fundamental prerequisite for complying with the Norwegian Personal Data Act and the GDPR. This is especially important when using M365 Copilot, because the tool acts as an 'accelerator' and can surface all the information it has access to in seconds.

M365 Copilot can be thought of as a 'clone' of its user. M365 Copilot has the same access and rights as its user. This means that all documents, emails, chats and other things the user has access to are available to M365 Copilot. Although M365 Copilot does not give the user access to new information, the tool makes it possible to quickly retrieve information that has previously been difficult to access. This could be information the user should not have access to and probably did not realise they had access to. This increases the risk of unintended or unauthorised use of data. Access control must therefore be closely linked to the user's role and needs in the organisation.

The Norwegian Digitalisation Agency has prepared a guide for getting your own house in order (in Norwegian only), including a maturity model that helps public organisations map and improve information management, focusing on establishing an overview of their own data sets. This is a resource we recommend all public entities familiarise themselves with.

NTNU must first get an overview of and control over

1. the system of agreements and settings for the actual cloud service that M365 Copilot sits on top of
2. information management in general, including classification, categorisation and access control
3. processing of personal data, including an updated and exhaustive record of processing activities

This is a challenging task, both for large organisations with a lot of data and many different systems, such as NTNU, but also for smaller organisations that may not have the expertise required.

Information management

Good information management helps to achieve several goals:

- **Information quality:** Ensures that the information is accurate, up to date and reliable.
- **Security:** Protects sensitive information against unauthorised access and security breaches.
- **Compliance:** Helps the organisation comply with legal requirements such as the GDPR and the Freedom of Information Act.
- **Efficiency:** Improves workflow and decision-making by making information readily available and understandable.
- **Reduced risk:** Minimises the risk of data loss, legal sanctions and reputational damage.

Good information management is contingent on basic guidelines being established for how information is to be handled within the organisation. Information mapping to identify what information is available, where it is stored, and who has access is an important step and also lays the groundwork for access control. Relevant training of employees is an important responsibility for the organisation.

The success of this is contingent on robust procedures that continuously classify information based on sensitivity and legal requirements, including the GDPR. Access to the different categories of information must be limited based on objective needs. Lifecycle management from creation to archiving or deletion is also part of this process.

Modern information management requires automation and the use of tools that ensure efficient, simple and consistent processes. They must also be able to handle the need to regularly evaluate and update practices to adapt to changing needs and regulations, as well as support the need to conduct internal and external audits to ensure compliance with laws and internal guidelines.

Modern information management requirements can also trigger the need for organisational changes, if roles have not been established with clear mandates to support these processes.

As emphasised in the findings from the NTNU report, a tool such as M365 Copilot can affect the organisation and should primarily be considered an organisational change project and an information management project rather than an IT project.

Organisations should consider whether any processes should be adapted or changed to be able to integrate M365 Copilot effectively, as well as adapting the product to the organisation's existing processes. This type of tool can be configured to a certain extent to meet specific needs, but it is important to understand the tool's limitations and strengths. This will require awareness and knowledge of which measures best support the need to generate gains, as well as ensuring compliance with the requirements for accountability and legality.

Records of processing activities

All entities that process personal data must keep a record of their processing activities (Article 30). The record of processing activities must show the purpose of each processing activity, a description of whose and which personal data are processed, recipients of personal data (if applicable) and whether personal data are transferred to countries outside the EEA. They should also, where possible, include a general description of the technical and organisational security measures referred to in Article 32(1).

Before considering whether M365 Copilot can be implemented, and in order to assess whether and how it could be implemented, an exhaustive and up-to-date record of processing activities must be in place, as already mentioned above. Then the organisation can, as a first step, assess whether M365 Copilot is suitable to be taken into use for each processing operation and, if so, how.

Using M365 Copilot will inevitably give rise to several new processing operations. This includes, as mentioned, storing each user's log of interactions with M365 Copilot (content of interactions log). This is a new processing operation, and particular consideration must be given to its purpose, when the log should be deleted and who in the organisation should have access to it. We will also discuss the application of the Norwegian e-mail regulation below. Administrators have access to the user's content of interactions log and have the ability to search it using eDiscovery. [16] Other new processing operations may arise from M365 Copilot generating information when it responds to prompts, which could be personal data. It is also unclear to NTNU whether user profiling occurs, but it considers it highly likely that profiling occurs.[17] Which new processing operations are triggered by using M365 Copilot will vary somewhat based on which settings are on or off, and these new processing operations must also be included in the records of processing activities.

Access control

In the following, we discuss access control as a security measure under Article 32 GDPR.

The general requirement pursuant to Article 32(1) GDPR is that the controller implements 'appropriate technical and organisational measures to achieve a level of security appropriate to the risk' in the processing of personal data. The purpose is to ensure the security measures are appropriate and proportionate to the specific risk linked to the processing operation.

The GDPR does not set out specific requirements for the content of the security measures. However, public authorities are obliged to use established standards when procuring, developing, setting up, operating and using IT solutions (Section 14 of Regulation No 959 of 5 April 2013 on IT standards in public administration). There are a number of such standards for personal data security, which all require that measures such as access control, logging and log control are in place, see, for example, ISO/IEC 27002:2022 chapters 5 and 8. Access control is a necessary element in the measures required under Article 32.

---

[16] eDiscovery, or 'electronic discovery', is the process of identifying, collecting, and analysing electronically stored information that can be used as evidence in legal cases or internal investigations. Microsoft Purview eDiscovery is a solution in Microsoft Purview that helps organisations to manage the eDiscovery process.

[17] NTNU's findings report, p. 119.

What is actually meant by access?

We have noted that the term 'access' is used in slightly different ways in practice when people talk about M365 Copilot. We will therefore explain in more detail what we mean by access control as a security measure.

In its findings report, NTNU emphasises the importance of 'actively deciding which data M365 Copilot should have access to'.[18] In order to prevent the tool from being *used* on incorrect data, measures to control end-users' *access* to data will be useful. Access control can support efforts to ensure that personal data are only used within the framework of a defined legal basis pursuant to Articles 6 and 9 GDPR. However, this does not form the core of access control as a *security measure*. As a security measure, the goal of access control is primarily to ensure that personal data have an appropriate level of confidentiality (Article 5(1)(f) and Article 32 (1)(b) GDPR).

Example from UiO

The significance of this can be illustrated by an example from real life. On 27 June 2024, Khrono referred to a personal data breach at the University of Oslo (UiO) where job applicants' CVs and the appointment committee's assessments were openly available to all employees at the university.[19] As a mitigating factor, it was highlighted that the information was difficult to access. In the breach report that was sent to the Norwegian Data Protection Authority, reproduced in Khrono's article, UiO wrote:

> *'To find [the information], employees must either actively search for it, or come across it by accident. This reduces the likelihood of the information actually being exposed to unauthorised persons, and the risk of the incident causing harm to those affected, but UiO cannot rule out that personal data have unintentionally been exposed to unauthorised persons.'*

Such arguments usually have some validity with respect to the assessment of whether a security breach is likely to have affected the data subject. However, this will be different for organisations using M365 Copilot. M365 Copilot retrieves information from obscure sources to which the user may not realise they have access to. This increases the likelihood of personal data being exposed unlawfully.

The Norwegian Data Protection Authority's experience is that this type of breach – where personal data are stored in places where they are available to unauthorised persons – is very common. Access control and classification of information should therefore be a priority security measure for organisations considering introducing M365 Copilot.

---

[18] NTNU's findings report page 2.

[19] https://www.khrono.no/alle-uio-ansatte-hadde-tilgang-til-informasjon-om-jobbsokere/885123 (last reviewed 4 Sept. 2024).

Data protection impact assessment

The obligation to carry out a data protection impact assessment (DPIA) follows from Article 35 GDPR. DPIAs shall, as a minimum, contain the four elements set out in Article 35(7) (a) to (d):

a. A systematic description of the envisaged processing operations and the purposes of the processing, including, where applicable, the legitimate interest pursued by the controller;
b. An assessment of the necessity and proportionality of the processing operations in relation to the purposes;
c. An assessment of the risks to the rights and freedoms of data subjects; and
d. The measures envisaged to address the risks, including safeguards, security measures and mechanisms to ensure the protection of personal data and to demonstrate compliance with [the GDPR] taking into account the rights and legitimate interests of data subjects and other persons concerned.

The obligation to carry out a DPIA arises if it is likely that 'a type of processing, in particular using new technologies (…), is likely to result in a high risk to the rights and freedoms of natural persons' (Article 35(1)). It is also worth noting that the wording 'rights and freedoms of natural persons' must be understood as a reference to the EU Charter of Fundamental Rights, which is largely in line with the European Convention on Human Rights. In other words, it is not only the impact on data protection that is to be analysed, but also the impact on rights such as freedom of expression and information and non-discrimination.[20] The Charter is not part of the EEA Agreement, but has an indirect effect through the GDPR.

For the sake of context, we would also mention that public entities that use what is classified as 'high risk' AI systems[21] pursuant to the AI Act will, in most cases, be subject to a similar impact assessment obligation through Article 27 of the AI Act as under Article 35 GDPR.

New technology, new consequences?

Whether a DPIA is required when using M365 Copilot, i.e. whether the processing of personal data entails a high risk to the rights and freedoms of natural persons, depends on several factors. It is relevant to look at the specific tasks the tool will perform and for what purposes, the context in which the tool is used, as well as the nature and scope of the processing of personal data. What can prove difficult, especially when using new technology such as generative AI, is that the user is not familiar with how the technology or product works, which makes it difficult to identify the potential risks involved, not to mention the likelihood of such risks.

We consider that, as a general rule, a DPIA will be required when using generative AI tools such as M365 Copilot in connection with the processing of personal data, as the 'use of new technologies' is highlighted as a particularly important factor, and the understanding of the risks associated with generative AI is still immature. Carrying out one or more DPIAs, regardless of whether it is required or not, will help NTNU to assess specific risks and the likelihood of such risks arising in a given context. It will also help clarify what NTNU does not know about either the product, the technology or the prerequisites for the technology in the context of a specific processing operation (e.g. to be able to assess whether the organisation's 'own house is in order'). It will also help to demonstrate compliance with the principle of responsibility under Article 5(2).

Several DPIAs?

It may seem time-consuming to have to carry out several DPIAs, especially if the organisation has several hundred different processing activities, but not doing so will probably make it difficult to comply with the principles of purpose limitation, data minimisation and lawfulness. This is why it will rarely be responsible or lawful to give the entire organisation and all its roles access to M365 Copilot. We find that a step-by-step approach to the introduction of M365 Copilot is most appropriate, where introduction is first considered for a limited area, e.g. a role and the associated processing that this role performs.

---

[20] Charter of Fundamental Rights of the European Union Articles 11 and 21.

[21] We have not considered whether M365 Copilot will be classified as a 'high risk' AI system.

It follows from Article 35(1) final sentence that a single DPIA may address a set of similar processing activities that present similar high risks. It is therefore possible to assess several processing operations under consideration for use with M365 Copilot in the same DPIA, as long as they are 'similar', and here the purpose, scope and what is done with personal data will be relevant factors here. With respect to the scope of the personal data that can be processed, it is important to look at what role will perform the task and what their access is. For example, a case officer will not have the same access as an HR employee or member of the management.

In addition, information or assessments in one DPIA will transferable to another DPIA.

Assess before consequences can arise

It is important that assessments are carried out *before* processing operations begin. If M365 Copilot is to be used in connection with an existing processing operation, the assessment must be carried out before the tool is used. But it does not stop there. As NTNU points out in its findings report, M365 Copilot is at an early stage of the development process and it is challenging to manage due to frequent changes that affect the risk situation.[22] DPIAs must therefore be carried out continuously, see also Article 35(11). One of NTNU's clear recommendations is to develop an exit strategy in the event of changes that lead to e.g. its use being deemed unlawful.[23] Further guidance on DPIAs can be found on the Norwegian Data Protection Authority's website[24] and in publications from the Article 29 Working Party and EDPB, respectively.[25]

Below we review some selected topics from a DPIA that we have given particular consideration in the sandbox project. However, it is important that NTNU also considers the other topics required when carrying out DPIAs for specific processing operations.

A systematic description of the processing

It is important to note that a DPIA pursuant to Article 35 requires that the envisaged processing operations and the purposes are specified, including the legitimate interest to be pursued, if relevant. This coincides with 'map and describe the processing' that we refer to above and includes all the processing operations that are covered. However, much of the information that NTNU has obtained in connection with preparation of its overall assessment will be transferable and will make carrying out specific DPIAs much easier.

Necessity and proportionality of the processing

NTNU has said that it will not be feasible for it to justify the necessity and relevance of the purpose of each of the variables in the data sets contained in a user's access to the Microsoft 365 platform, without a more thorough review and systematic follow-up.[26] However, it should be possible to consider necessity and relevance if the relevant processing operation(s) are first described systematically in light of a specific use case.

Purpose limitation

The purpose limitation principle states that personal data shall be collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with these purposes. When assessing the purpose

---

[22] NTNU's findings 4 and 5, see NTNU's findings report pp. 12–15.

[23] NTNU's findings report pp. 12 and 31.

[24] See https://www.datatilsynet.no/rettigheter-og-plikter/virksomhetenes-plikter/vurdering-av-personvernkonsekvenser/ (In Norwegian only – last reviewed 16 July 2024).

[25] Article 29 Working Party: *Guidelines on Data Protection Impact Assessment (DPIA) and determining whether processing is 'likely to result in a high risk' for the purposes of Regulation 2016/679* (WP 248), available at https://ec.europa.eu/newsroom/article29/items/611236/en.

EDPB: *Recommendation 01/2019 on the draft list of the European Data Protection Supervisor regarding the processing operations subject to the requirement of a data protection impact assessment*, available at https://www.edpb.europa.eu/our-work-tools/our-documents/recommendations/recommendation-012019-draft-list-european-data_en.

[26] NTNU's findings report, p. 104.

of processing, it is important to remember that M365 Copilot is a tool or function – a means – to achieve the purpose of the processing. Using M365 Copilot is not a purpose in and of itself.

'Helping the user to perform their tasks' is too vague and general, but the purpose limitation principle may be met if it is further specified, for example by specifying in detail what kind of task it is and why it is being performed.

As mentioned, a particular challenge of M365 Copilot is that the purpose, and what personal data are used to achieve the purpose, is in practice defined (controlled) by the individual user in each prompt. Using M365 Copilot allows personal data to be processed in a different context – and for a different purpose – than originally intended. This happens because M365 Copilot uses all the information that is available to the user via Microsoft Graph. NTNU identified features of M365 Copilot that allow personal data collected for one purpose to be further processed for new or other purposes.[27] It is therefore important that NTNU sets a clear framework for its users, for example in the form of guidelines, procedures and training, to ensure e.g. as much purpose limitation as possible when using M365 Copilot. This should preferably be seen in the context of the user's role' in the organisation, which will also correspond to this role's access.

Personal data may be further processed for new purposes, as long as the new purpose is compatible with the original purpose. Article 6(4) GDPR sets out a non-exhaustive list of what should be emphasised in this assessment. There are currently no guidelines or court decisions on how this provision should be applied or understood, but such a compatibility assessment must be performed per processing operation and not for M365 Copilot as a whole.

This is obviously difficult in practice, but may be easier when using M365 Copilot for some selected roles that perform a limited range of operations. A particular difficulty associated with the current version of M365 Copilot is that it is not possible to disable access to a user's electronic mailbox. This means that emails and the personal data they contain can easily be used for purposes other than that originally intended. The Norwegian Data Protection Authority does not have a set answer to how a data controller can ensure that the purpose of processing personal data in the context of M365 Copilot is compatible with the original purpose for which they were collected.

One possible measure may be to train users to delimit the search area through the prompt by using prompt engineering. In this context, it will be a prerequisite that the organisation has good 'order in its own house' and guidelines in place. In addition, NTNU's findings report[28] sets out that M365 Copilot can be set not to use information from certain areas, e.g. Teams chat. We believe that Microsoft should develop settings that also make it possible to block access to information from emails, as it is virtually impossible to have any control over what they include.

Data minimisation

Data minimisation is an absolute requirement under Article 5(1)(c) GDPR: '[personal data shall be] adequate, relevant and limited to what is necessary in relation to the purposes for which they are processed.' Data minimisation must therefore be assessed in light of the purpose of the processing in question, to find out what personal data are needed to achieve the purpose. As stated above, NTNU has yet to identify a specific processing operation with a specific purpose ('M365 Copilot in the operational phase' is not a purpose), which NTNU has to do itself.

The data minimisation principle is perhaps one of the most difficult obligations to fulfil when using M365 Copilot, because the tool is built to have access to everything a user has access to and therefore has the ability to process the data it 'considers' relevant based on the prompt after the enrichment process. It is not possible to determine exactly how M365 Copilot 'chooses' what is relevant based on the prompt, due to both the black box issue and the fact that this is, regardless, proprietary information.

Prompt engineering may be a measure that could be used to minimise data access, but it is unclear whether that would lead to full compliance with the data minimisation principle. Another potential measure is the activation of Double Key

[27] NTNU's findings report, p. 68.

[28] NTNU's findings report, p. 18.

Encryption (DKE) to block files that should not to be accesed by M365 Copilot.[29] However, granular settings in M365 Copilot would be preferable in terms of data access, especially in terms of access to a user's electronic mailbox.

Accuracy

The accuracy principle in the GDPR entails that 'every reasonable step must be taken to ensure that personal data that are inaccurate, having regard to the purposes for which they are processed, are erased or rectified without delay'. As mentioned above, the answers generated by M365 Copilot may be incorrect. They may, however, appear convincing and correct. The risk of incorrect answers will never be eliminated. It is therefore important that users are sceptical to AI-generated answers.

An AI-generated answer is based on probability and depends on the model's training data and weighting. NTNU writes that it is likely that M365 Copilot will come up with things that are both false and incorrect[30] (also known as 'hallucination'), and the nature of the tool means that it can produce incorrect information.[31] M365 Copilot can make errors even when it seemingly has access to information of sufficient quality.[32] The likelihood of errors increases if the user gives imprecise prompts to the large language model. This may become an even bigger problem if the user is supposed to check whether the answer is incorrect or not, but does not have enough time or information to do so.

This is even more relevant when personal data are processed. The accuracy principle entails a statutory obligation to ensure that personal data are accurate, and that every reasonable step must be taken to ensure that personal data that are inaccurate, with respect to the purposes for which they are processed, are erased or rectified without delay. If M365 Copilot generates incorrect personal data about someone, it may, firstly, be difficult for the user to check whether the answer contains errors, and, secondly, pose a high risk to the rights of the data subject.

NTNU would like users to be able to actively assess the information provided by the solution and be fundamentally critical of the information provided by a large language model to counteract the risk of incorrect information being perceived as correct. At the same time, NTNU acknowledges that special consideration should be given to whether M365 Copilot should be used in processes where the principle of contradiction is important. It is therefore a good idea to determine which areas or tasks are not suitable for the use of generative AI tools. This could include certain tasks relating to HR or the exercise of public authority, which require a high degree of precision and accuracy and where the impact of errors can be serious. NTNU points this out in its findings report (finding 2).

In the context of generative AI, we believe that this indicates that measures should be implemented to reduce the risk of incorrect personal data being generated (e.g. by prompt engineering or rules about what M365 Copilot should not be used for) and measures to rectify or erase incorrect personal data without delay (e.g. effective post-monitoring of what is generated). If M365 Copilot is to be used as a decision support tool, measures such as rules, guidelines, training and selection of users with the right expertise must be considered and introduced. It is important that those who use M365 Copilot are critical of the answers and have both the time and expertise to detect and correct incorrect personal data that may occur in the output.

This may limit what M365 Copilot can be used for. For example, it will not generally be advisable to use M365 Copilot where people or personal data are at the core of the task to be performed. We also agree with NTNU in that the threshold should be set very high for using M365 Copilot in the exercise of public authority, where accuracy is key.[33]

---

[29] NTNU's findings report, p. 117.

[30] NTNU's findings report, p. 71.

[31] NTNU's findings report, p. 105.

[32] NTNU's findings report, p. 9.

[33] Ibid.

The rights and freedoms of data subjects

NTNU has identified numerous risks related to the rights and freedoms of data subjects in NTNU's use of M365 Copilot at a general level. These are described in NTNU's findings report pp. 106–117. Below we focus in particular the right to information, the right of access, the right to object and the prohibition on automated decision-making.

The right to information

The cornerstone of data subjects' rights is the right to information. The data controller must explain, in a clear and simple way, how personal data are processed. This is a prerequisite for the data subjects' ability to exercise their rights. Articles 12 to 14 GDPR as well as the transparency principle in Article 5 require controllers to provide data subjects with information on how their personal data are to be processed, and these obligations are most often fulfilled by using both external and internal privacy policies.

When new technology is introduced, it is important that the impact the technology may have on the individual is reflected in the relevant information and relevant privacy policies and that it is clear when, how and in what context such technology is to be used on the data subject's personal data. In order for it to be understandable to the data subjects, it will often also be necessary to concisely explain how the technology itself works.

NTNU has found that their privacy policy needs to be updated with information about the use of M365 Copilot, based on an examination of what actually happens in practice.[34] This should include information about whether M365 Copilot is used in connection with the processing and for what purpose, as well as what new processing operations and processing arise from the use of M365 Copilot and what legal basis is used for each processing operation. How the tool works should also be explained in as simple a way as possible to ensure transparency.

At the point of implementation of M365 Copilot, data subjects should be sent information directly, where possible, about the processing NTNU is going to use M365 Copilot in connection with and what it means for the them. For example, this can be done in an email sent directly to the data subject. If it is not possible to contact the data subjects directly, the information should be clearly stated on an appropriate interface between the data subjects and NTNU, e.g. on NTNU's website.

Information to data subjects must be reviewed and assessed continuously or at regular intervals in line with technological developments and whether and how M365 Copilot is taken into use in other areas of NTNU.

It should also be clear what content is created using generative AI. This is especially important when the content includes personal data, and will make it easier for the data subjects to exercise their rights and have control over their personal data.

The right of access

Outputs generated by M365 Copilot may contain personal data and will then be subject to the right of access of the person concerned. NTNU has said that it must be considered whether it is possible to fully comply with an access request, as it will be challenging for NTNU to identify all the places where personal data may be processed in M365 Copilot.[35] NTNU has also identified this as a problem in relation to the use of the Microsoft 365 platform generally, even without the use of M365 Copilot.

As stated above, content generated using generative AI should be labelled. Where the generated content is stored will depend on the specific use case, but it should match the storage location where other documents for the task/use case are stored before using M365 Copilot. One particularly new feature of using M365 Copilot is storage of the content of interactions log. This log will be stored in accordance with NTNU's applicable retention policy (see the section on storage limitation above) and can be searched by administrators. However, the data subject will only have right of access

---

[34] NTNU's findings report, pp. 106–7.

[35] NTNU's findings report, p. 107.

to their own personal data stored in the content of interactions log, and not the entire log in general. The right of access shall not adversely affect the rights and freedoms of others (Article 15(4) GDPR).

The right to object

Data subjects may object to NTNU's processing of their personal data when the legal basis for processing is Article 6(1)(e) or (f) GDPR (see Article 21(1) GDPR). If the objection is granted, it may be difficult for NTNU to comply with the same because of the inherent properties of M365 Copilot, where its access reflects the user's access. NTNU has proposed a potential solution to this, which is to use Double Key Encryption (DKE) which can be activated for files containing the personal data in question.[36]

Prohibition on automated individual decision-making

When considering the legal basis for processing, the prohibition on automated individual decision-making in Article 22 GDPR should also be considered, which sets some limits on what M365 Copilot can be used for. Article 22 GDPR contains a prohibition on automated individual decision-making consisting of three cumulative conditions: (1) there must be a 'decision', which is a term to be interpreted broadly,[37] (2) the decision must be based solely on automated processing, and (3) it must produce legal effects or similarly significantly affect the person concerned. There is a high threshold for what is covered by the prohibition. As a rule, most things are not affected, but assessments may be affected if the decision-maker in reality relies solely on M365 Copilot's assessment.

Example: The assessment of the application for adaptation in use case C will have 'legal effect' for or 'significantly affect' the applicant, and therefore cannot be handled by M365 Copilot alone. However, a person can use M365 Copilot as a decision support tool, as long as they do not rely solely on M365 Copilot's assessment.

NTNU identified functions where a user can ask M365 Copilot to assess a colleague's behaviour and work performance. Although this action is not generally affected by the prohibition in Article 22 GDPR, it will also entail the processing of personal data that is unlikely to have a valid legal basis.

Risk mitigation measures

If a legal basis for processing is identified and the DPIA gives the processing the go-ahead, testing is a recommended risk mitigation measure. Testing must take place within the framework of the identified legal basis/bases for processing. Depending on the purpose of the testing and what is to be tested, Article 32 GDPR may also constitute what is known as a supplementary legal basis.

NTNU has identified many different risks to the rights and freedoms of data subjects as well as possible risk mitigation and damage limitation measures, which are discussed in their findings report, pp. 121–126. A total of 41 measures are listed. These risks and measures have been identified based on a general review of M365 Copilot. Nevertheless, it will be relevant to consider many of the risk mitigation measures in a more specific data protection impact assessment in light of a specific processing operation or a set of similar processing operations.

A particular challenge in the case of M365 Copilot is that the purpose of a processing operation is in practice defined (controlled) by the individual user in each prompt. It is therefore important that the data controller sets a clear framework for its employees, including in the form of procedures and training, in order to ensure as far as possible that the processing is carried out lawfully.

However, it will not be realistic for the controller to achieve complete control. A subsequent audit of the actual use will therefore be necessary. Such an audit must itself have a legal basis. Here it is worth noting that the provisions of the GDPR may themselves constitute a supplementary legal basis pursuant to Article 6(3). The larger the organisation, the

[36] NTNU's findings report, p. 117.

[37] Judgment of 7 December 2023 [GC], *Bundeskartellamt*, C-634/21, EU:C:2023:957, para. 60.

more difficult it will be to achieve control, and the greater the risk of undesirable incidents. NTNU, with its 70,000 users, refers to it as 'utopian' for all users to comply with routines.

Involvement of the data protection officer

When an organisation is considering using a new AI tool such as M365 Copilot, it is important to involve the data protection officer at an early stage. The data protection officer should be considered a key resource in the assessment, introduction and post-monitoring of the AI tool. Article 39 GDPR describes the tasks of the data protection officer, which include advising on data protection obligations and data protection impact assessments, as well as monitoring the performance of such assessments. The data protection officer shall perform their tasks independently.

The data protection officer must have an understanding of the entire lifecycle of the AI system that the company is considering acquiring and how it works. This means that the data protection officer must, among other things, receive information about when, why and how such a system processes personal data, how the data flow works (input and output) and decision-making processes in the model.

The ban on monitoring in the Norwegian e-mail regulation

Norway has a separate regulation that regulates an employer's access to employees' electronic mailboxes and other electronically stored material (the e-mail regulation).

When a user interacts with M365 Copilot, data about these interactions are stored containing the user's prompts and M365 Copilot's answers, including references to source material (the interaction log).[38] The interaction log is stored in a hidden folder in the user's 'mailbox'. Such hidden folders are not designed to be directly accessible to users or administrators, but can be searched by compliance administrators using 'eDiscovery tools'.[39]

Deletion of the log

Whether and when the interaction log is permanently deleted depends on the organisation's storage policy. A user may have the opportunity to delete the interaction log themselves as an option in the M365 Copilot settings, but it is unclear to NTNU whether a user deleting their own log also entails deleting the log that the administrator can see.[40] Microsoft itself states that *'Messages visible in Copilot are not an accurate reflection of whether they are retained or permanently deleted for compliance requirements'*.[41] The interaction log will first be moved to the SubstrateHolds folder. The interaction log will not be permanently deleted until the storage period set by the organisation expires.

The e-mail regulation[42] encompasses electronic mailboxes, personal areas in the organisation's computer network and other electronic equipment that an employer has placed at the employee's disposal for use in their work for the organisation. The regulation also applies to data that have been deleted, if they are found on backup copies or similar. It is clear that the interaction log falls within the scope of the regulation.

The e-mail regulation contains, firstly, conditions as to when the employer may, in an individual case, access information stored in the above-mentioned area. Secondly, Section 2 of the e-mail regulation contains a ban on monitoring the employee's use of electronic equipment, unless the purpose of monitoring is

a. to manage the organisation's computer network or
b. to detect or resolve security breaches in the network

The Norwegian Data Protection Authority's guidelines[43] state that the ban applies if: (1) the measure concerns monitoring, (2) the monitoring is aimed at employees' use of electronic equipment and (3) the employer has access to the information.

Is the interaction log monitoring?

The interaction log will show the history of an employee's use of electronic equipment. The question is therefore whether the interaction log is to be regarded as monitoring pursuant to the e-mail regulation. The interaction log can be considered monitoring, even if this is not the intention of the employer. Relevant factors in the assessment include the type and amount of information in the log, how long the information is stored and how much of the workday can be traced.

---

[38] Data, privacy and security for Microsoft Copilot for Microsoft 365 | Microsoft Learn (opened on 26 August 2024)

[39] Learn about retention for Microsoft Copilot for Microsoft 365 | Microsoft Learn (opened on 26 August 2024)

[40] NTNU's findings report, p. 79.

[41] Learn about retention for Microsoft Copilot for Microsoft 365 | Microsoft Learn (opened on 26 August 2024)

[42] Regulation on employer's access to electronic mail boxes and other electronically stored material of 2 July 2018

[43] Monitoring of employees' use of electronic equipment | Data Protection Authority (in Norwegian only)

The interaction log can reveal a lot about someone's behaviour and workday. M365 Copilot is built into all the tools the employees use on a daily basis, such as Word, Excel, PowerPoint and Teams. The extensive mapping that takes place when employees use M365 Copilot indicates that the interaction log should be regarded as monitoring.

Log access

If only the employee had access to the log, the ban would not apply. As described above, the ban applies if the employer has access to the data.

NTNU found that the employer has access to the interaction log using eDiscovery and Purview. There is also a function that causes a user's prompts to be sent for control if certain criteria are met – in other words, an 'alarm' is triggered. NTNU can change these criteria itself. However, it is unclear how much of the interaction log is sent for control, to whom and for what purpose.

Overall, we consider it likely that the interaction log could fall under the prohibition on monitoring employees' use of electronic equipment as set out in Section 2 second paragraph of the e-mail regulation. In order for the interaction log to be created lawfully, the purposes of the interaction log must fall under one of the exemptions.

Exemptions from the prohibition

One exemption applies if the purpose is to 'manage the organisation's computer network'. This is to be understood as all practical and technical measures necessary for the functioning of systems, networks, equipment and software.[44] We understand that the main purpose of the interaction log is to ensure that the quality of the service is as it should be. By reviewing the user's prompts and M365 Copilot's answers, NTNU can identify weaknesses and improvement potential in the system. Potential systematic wrong answers can be detected and corrected. That purpose may fall under the exemption that applies to managing the organisation's computer network.

The second exemption may apply if the purpose of the interaction log is to 'detect or resolve security breaches in the network'. We consider that what is meant by security breach is a breach of information security in general. This is generally said to be about ensuring that information does not become known to unauthorised persons (confidentiality), is not changed unintentionally or by unauthorised persons (integrity) and is available when necessary (availability). Whether this exemption applies must be assessed specifically in light of the purpose of the interaction log.

For both exemptions, NTNU must be aware of the principles of data minimisation and purpose limitation.[45]

---

[44] You can read more about the purposes in the Norwegian Data Protection Authority's guidelines on Monitoring of employees' use of electronic equipment | Data Protection Authority (in Norwegian only)

[45] A more detailed description and examples of data minimisation and purpose limitation for both exemptions is provided in Monitoring of employees' use of electronic equipment | Norwegian Data Protection Authority (in Norwegian only).

Conclusion

M365 Copilot has considerable potential, but will require that one takes small and controlled steps. For Microsoft's part, the tool is contingent on an organisation having extremely good control over its own information management, something that many organisations probably will not have.

It is unlikely that it will be possible to use M365 Copilot in a responsible and lawful manner without considerable preparation in advance. This includes getting one's own house in order and carrying out thorough data protection impact assessments for its planned applications. The technology also makes high demands of training of employees, and of awareness and knowledge among the organisation's users.

The positive aspect of this work is that a strong focus on information management can generate major benefits way beyond the actual use of this tool. Efficient, well-functioning information management is the very foundation for succeeding with digitalisation, socially beneficial data sharing and cost-effective compliance with laws and regulations, including the GDPR. Adopting new and advanced technology without thorough preparation, understanding its possibilities and limitations, and securing the necessary expertise, will not be responsible.

It is possible for Norwegian organisations to use M365 Copilot, but use cases should be chosen carefully to ensure compliance with, among other things, data protection requirements. At the same time, it is both right and important to test new technology and gain practical experience of the opportunities it presents, contingent on the necessary assessments being made in advance, and the organisation having good processes and establishing measures to reduce identified risks.

The same large language model technology that underpins M365 Copilot can also be used in other and more targeted ways than purely as a general office support tool. Such approaches can reduce requirements for organisational changes, ensure faster investment recovery and, not least, improve control of quality and compliance with regulations. It is therefore important to assess whether other AI solutions exist that can meet the organisation's specific needs, but which entail a lower data protection risk. It could also be the case that more focused solutions could be a good starting point for later use of integrated AI solutions such as M365 Copilot. By first getting the organisation's own house in order and establishing support mechanisms for compliance with requirements, organisations will be better equipped to use advanced solutions when products such as M365 Copilot have had time to mature and adapt.

The road ahead

NTNU has chosen not to introduce M365 Copilot throughout the entire organisation, but instead introduce the tool in small and controlled steps, limited initially to selected roles. As M365 Copilot requires its own licenses, increased costs must be justifiable through actual, realisable gains, and it is important that both direct and indirect costs are included in the overall assessment. M365 Copilot is still in the early stages of development and does not provide control at a granular level, such as the ability to make local and flexible adaptations (e.g. disabling access to users' mailboxes or specific deletion policies). Microsoft probably considers unlimited access to the user's mailbox as an important and central feature, but it is perhaps one of the features that gives rise to the most uncertainty among many organisations.

The Norwegian Data Protection Authority expects the issues that customers, organisations, authorities and wider society identify in the product are taken seriously by the product supplier. At the same time, there are clear requirements for organisations that wish to benefit from using the tool. The prerequisite of having an extremely well-functioning information management system may make it difficult to succeed with such solutions, but obviously has a positive upside that goes far beyond the implementation of one specific solution.

NTNU has done an impressive, socially beneficial and extensive job of acquiring knowledge and awareness of the use of large language models in general and integrated AI solutions such as M365 Copilot in particular. However, if NTNU wishes to expand its use of M365 Copilot, it is important that the necessary data protection impact assessments are carried out for specific processing operations in light of given use cases.

Appendix

Large language models

Large Language Models (LLMs), such as the Generative Pre-trained Transformer (GPT), are machine learning models that have been trained on very large quantities of text data. These models process and generate text by using contexts in the training data to predict the next word in a text or when it answers a question. They are used in a variety of applications, such as chatbots, text generation and language analysis.

LLMs are based on neural networks. They do not store languages and words as text, but as numerical representations called vectors, which effectively describe very complex relationships between language elements. An LLM does not 'understand' language in a human sense, but models the relationship between language elements based on how language is used by people.

Since LLMs are so well articulated, it is easy to perceive them as models of knowledge, but they are not. They are models of language itself and of how it is used in practice. In other words, today's LLMs do not have built-in knowledge of disciplines such as law, chemistry, physics, philosophy and mathematics. At the same time, language, by nature, reflects information about the world around us. LLMs are trained on large amounts of text that contain (random) information about different topics, and different types of information are thus often reflected in the language in which the model is trained. The technology is undergoing rapid development, and LLMs that are combined with knowledge models are both being tested and will increasingly also be available for general use, with the potential for more accurate and reliable answers.

An important challenge associated with LLMs is the phenomenon 'hallucination'. This means that the model generates text that is linguistically correct, but that contains incorrect or fictitious information. LLMs function as advanced, statistical models without a built-in understanding of the facts. They have a certain randomness built in (in order to be able to provide varied and/or alternative formulations of answers), but lack mechanisms for assessing the truth of the content, which can lead to generated text appearing to be credible, but actually being incorrect.

In addition to hallucination, errors may occur due to misconceptions in the training data. If a widespread error or misconception is present in the data, the model may repeat or reinforce this error. For example, if many sources in the training material incorrectly claim something, the model is likely to reflect this as if it were correct. This can be problematic when models are used in situations that require a high degree of precision or professional accuracy.

In practice, however, many answers will be good and relevant, given that the vast majority of the texts on which they are trained contain the most correct information, and because the linguistic contexts in many cases also contain the relevant facts.

Most major suppliers' LLMs are primarily trained on English-language data, which often generate better answers in English. Efforts are under way to adapt language models to national languages. But will they also be able to adapt them to national cultures? It is important to be aware of this when using LLMs, because they also reflect cultural context. A predominance of English and American text sources in the training material will mean that the text generated is also influenced by and reflects British and American culture.

The largest and most dominant LLMs, such as OpenAI's models, are created, trained and operated by large, private American companies. Microsoft uses OpenAI models to provide LLM services on its Azure platform, with adjustments and adaptations to its own products such as M365 Copilot.

Unlike 'classic' AI/machine learning systems that are mainly modelled and trained for specific purposes, LLMs can be used for many and unspecified tasks. They are therefore also referred to as foundation models. This makes them very useful, but at the same time challenging in terms of ensuring accuracy, relevance and responsible use.

Adaptation of LLMs

M365 Copilot uses several techniques to customise the product, including Knowledge Graphs and Retrieval-Augmented Generation (RAG). The purpose of RAG is to control the quality of text-based answers by retrieving selected and updated information from internal information sources before generating an answer. This new additional information is vectorised and indexed in the same numerical format as the original foundation model.

RAG comprises three main components:

1. **Retrieval:** The model searches for information in a database or external sources. This is similar to how traditional search engines work, but RAG uses what are known as semantic search methods to find relevant information based on linguistic context instead of keywords.
2. **Augmented:** The information collected is used to enrich the LLM's answers. This makes the answers more precise and fact-based, compared with the foundation models that are primarily based only on pre-trained data.
3. **Generation:** After the relevant information has been collected, answers containing additional information from steps 1 and 2 are generated using the foundation model itself.

RAG's semantic searches are based on the linguistic context of a question (prompts). This allows the system to retrieve information that is relevant even if the exact words do not match and makes searches more flexible and relevant to the user.

RAG has other features:

- **Updated knowledge:** In contrast to LLMs that can only draw on their own static training data, RAG can retrieve in-house information and information from new, self-checked sources, providing more accurate answers.
- **Flexibility:** The system can perceive complex intentions and contexts in questions, even if all keywords are not present.
- **Accuracy:** By specifying and controlling what information is used as the basis for answers, RAG reduces the risk of errors or 'hallucinations' compared to answers produced by the foundation model alone.

Applications for RAG

The purpose of RAG is to be able to control to a greater extent what information to include in answers from LLMs. It also enables some adaptation to specific subject domains to increase the likelihood that the information presented is relevant and correct. However, this requires very good control of what information is included in the RAG model. General LLM solutions such as M365 Copilot will be more able to tie answers to the organisation's own information, but will still depend on the quality of this information. Unclassified, old, outdated or incorrect information in internal sources will negatively affect quality.

Although RAG can improve LLMs' ability to provide answers based on the organisation's own information, there are also challenges associated with implementation, including monitoring the quality of answers. In addition, there are technical challenges related to scaling and performance when these models are used on a large scale, including that many additional operations can lead to longer response times.

LLMs such as GPT represent an important technological innovation in text generation, but have limitations when it comes to updated and fact-based information. However, the RAG system, if implemented correctly, can increase the quality of answers by using in-house information.

The Norwegian Data Protection Authority's regulatory sandbox for responsible artificial intelligence

Office address:
Trelastgata 3, Oslo

Postal address:
P.O. Box 458 Sentrum
NO-0105 Oslo

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

datatilsynet.no/sandkasse
personvernbloggen.no
twitter.com/datatilsynet