

QUFA AI Task Force

Digital Bargaining Report – Public Version

April 2026

Contents

Preamble	3
A Note on Process	3
Terminology.....	3
Guiding Principles	4
Key Issues.....	6
Deskilling, Automation, Staffing	6
Personnel Decisions	8
Employment and Hiring Equity	8
Academic Freedom	9
Workload	11
Privacy	11
Final Notes	12

Preamble

A Note on Process

The QUFA AI Task Force met between December 2025 and February 2026 to identify areas for strengthening our Collective Agreement with respect to AI and other digital technologies impacting members' working conditions. The members of the task force were Carolyn Lamb (Computing, co-chair), Norma Möllers (Sociology, co-chair), Erica Friesen (Library), Mél Hogan (Film & Media); Christian Muise (Computing); and Catherine Stinson (Computing/Philosophy); joined by Karen Rudie (QUFA president) and Doug Nesbitt (QUFA Staff).

Following recommendations by PSI's "[digital bargaining hub](#)" for developing Collective Agreement language, the Task Force began by brainstorming AI-related bargaining concerns and a set of guiding principles. We then analyzed the existing Queen's-QUFA Collective Agreement for gaps between concerns/principles and each Article of the Agreement. This analysis forms the bulk of this report which was developed, reviewed, and revised by all members.

We want to especially acknowledge Hannah Johnston (YUFA, PSI) who shared a great deal of expertise about digital bargaining issues, strategies, and suggested language from YUFA's experience at the bargaining table.

At the request of the QUFA executive committee, a redacted copy of the Bargaining Report was prepared for broader distribution to QUFA Membership. This redacted copy is identical to the original Bargaining Report except that recommendations for specific bargaining language have been removed for strategic reasons until bargaining starts.

Terminology

You will notice that our recommendations are not specific to generative AI.¹ There are two reasons for this: First, the meaning of AI has changed dramatically over the past seventy years, and it is reasonable to assume that it will change again in the future. To prevent our

¹ "AI" refers to artificial intelligence - meaning, any computational technology which attempts or performs a task associated with intelligence in humans. These tasks can include, for example, memory, reasoning, learning, communication, or perception. The method of approach to these tasks is often quite different from a human's, and quality and accuracy of tasks performed may or may not approach a human's. "Generative AI" or "genAI" is one type of artificial intelligence, characterized the ability to produce novel outputs in some medium, such as text, images, music or video - as opposed to other forms of AI which are used, for example, to analyze data or make predictions.

Collective Agreement language from becoming outdated as popular terminology changes, we generally advise against writing specific technologies into our Collective Agreement.

Second, while the recent explosion of interest in generative AI has indeed created some novel challenges for educators, many of the challenges it presents are versions of challenges also posed by other forms of AI or digital technologies. For example, the use of generative AI creates risks related to privacy, surveillance, and bias, but other forms of AI also pose similar risks for educators in these areas. To make protections in our Collective Agreement as capacious as possible, we recommend focusing on processes and issues which are shared across a range of different technologies. We found it helpful, for example, to think about issues related to the “data lifecycle” (collection, use, access, analysis, storage, disposal).

For this reason, while not ignoring novel challenges that are specific to generative AI, this report's main focus is on how the QUFA bargaining team can address challenges posed by AI and data-intensive technologies more generally.

Guiding Principles

Technology should improve working conditions (or at the very least, not worsen them)

AI is often marketed as a tool to make workers' lives easier by shortening the time spent on mundane tasks and boosting productivity. The Task Force acknowledges that Members are overworked and encourages keeping options open for the adoption of technologies that can actually fulfil this promise. However, the promise of AI can be a double-edged sword: it can embolden employers to ask increasingly more of employees, assuming (often incorrectly) that the use of AI should make employees exponentially more productive at all work tasks, then pre-emptively shrinking workforces and increasing workloads to fit this incorrect belief. In this report, we have provided suggestions for limiting unwanted impacts of new technologies on working environment and workload. As a general principle, any changes to the Collective Agreement should reflect that QUFA Members are the best judges of whether AI can improve the conditions of their research and teaching or not; setting guardrails against mandated use of technology (i.e., QUFA members should be free to use technology as they see fit, but the administration must not be able to mandate the use of specific technologies for research and teaching).

Academic freedom includes the freedom to not use technology

The Task Force acknowledges that QUFA Members, by definition, are experts in their fields, with years and often decades of experience in research and teaching in their specific subject areas; moreover, the unimpeded search for knowledge by Members is vital to the

public role of a university. The principle of academic freedom, enshrined in the QUFA Constitution and Article 14 of the Queen's-QUFA Collective Agreement, ensures that Members can use this expertise to make their own choices in matters of teaching, research, publication, choice of materials, and free expression, without adherence to prescribed doctrine. This freedom must apply to the use or non-use of AI, as we will describe more fully later in this report, although it does not free Members from ethical obligations to their students and colleagues. While the right to use AI is granted by default through academic freedom provisions, we are concerned about Members' ability to opt out in a changing workplace environment because increasing workloads (such as large class sizes) may significantly impede Members' ability to opt out. As subject matter experts possessing academic freedom, Agreement language should protect Members' decision-making ability in the adoption of any new technologies by the University which affect the ways that teaching, research, publication, and expression are expected to be done.

Employment equity and anti-oppression: technology must not discriminate

AI is known to reproduce the social biases of its training data and in many cases can even amplify this bias and intensify social inequality. The use of AI to make any decision involving equity-deserving Members is therefore highly likely to deepen inequalities these Members already experience and is incompatible with both our Collective Agreement and the Human Rights Charter. AI can also worsen working conditions for equity-deserving Members in other ways, such as by enabling large scale harassment of Members or routinely exposing them to discriminatory content. Finally, Members should be able to refuse AI use in teaching and research based on equity grounds – for example, because of its known biases, its uneven environmental impact, or the global labor exploitation on which it relies. As a general principle, where AI directly impacts the material conditions of employment, such as in personnel processes, Agreement language related to the use of AI in the workplace must take into account the principally biased nature of AI and the uneven impacts of these technologies on QUFA members.

Worker solidarity: our work is connected to others

The Task Force acknowledges that QUFA can only bargain for the rights of QUFA Members, not for other workers at Queen's. However, it is important to recognize that all workers' struggles are connected. For example, automation/deskilling (described below) is likely to affect staff and TAs first, but reduction in TA and staff availability due to automation will have ripple effects for QUFA members' workload. Similarly, within QUFA, term adjuncts are more at risk than tenured/tenure-track/continuing Members, but conditions will continue to deteriorate even for tenured Members if term adjunct jobs are automated away. The Task Force urges the bargaining team not to ignore the struggles of Queen's more precarious

workers, but to affirm solidarity across different status groups: by protecting our most vulnerable community members, we are protecting all of us.

Key Issues

Deskilling, Automation, Staffing

One of the primary concerns the task force identified is the deskilling or automation of labor, leading to either decreases in wages or job losses over time. Because AI can perform some tasks which are normally associated with intelligence in humans, current AI - especially generative AI - is marketed to managers, both in education and in industry, as a potential replacement for human workers. Managers, who are already motivated to reduce workforces as much as possible, are told by tech companies that they can now cut significant fractions of their human workforce – either because AI will directly replace all workers in a particular job type, or because AI will make workers 2x or 5x as efficient, supposedly enabling the human workforce to be reduced to half or 1/5 of its size without affecting productivity. This erroneous belief in AI efficiency is one of the factors behind the current entry-level hiring crisis across many industries including customer service, retail, legal services and, ironically, software.

In the education sector, these claims about AI are lies. While some rote tasks can be successfully automated, generative AI agents are nowhere near ready to reliably perform the full range of work tasks of a QUFA Member. In tests, even the best current AI agents fail at more than half of entry-level white-collar job tasks, such as updating a spreadsheet, often deceptively claiming to have completed the task correctly when they did not. Moreover, because of inherent flaws such as “hallucinations”, sycophancy, and bias amplification, even an improved generative AI would be both incompetent and dangerous as a replacement for QUFA Members, particularly in teaching work where the wellbeing of diverse and vulnerable students is at stake.

Aside from full automation of work tasks, there is also the related issue of deskilling. In deskilling, a job which requires expertise is *partly* automated, and workers are told that their role is no longer to perform the work and make the key decisions themselves, but rather to supervise the automated system and correct it if it makes an error. While deskilling may or may not lead to direct job loss, deskilling in the education sector would result in a catastrophic loss of Members’ ability to apply their own personal expertise, experience, and academic freedom to teaching and research tasks. Over time, deskilling also results in a perception that training and expertise, on the part of the human supervisor,

is less necessary, which results in lowered wages and poorer working conditions for the humans who remain.

In the context of ongoing austerity measures and reported plans for more remote course offerings and non-degree programming (“micro-credentials”) to generate revenue, it is reasonable to assume that the administration will try to automate at least part of the work currently carried out by faculty, despite the above-mentioned problems. Software to automatically create syllabi, assignments, and other course materials already exist, as do programs which automate grading.

Similar moves by the administration can be expected for librarian and archivist work. We identified as concerns for librarians and archivists the potential use of their intellectual property (e.g. web content, instructional content) to train AI systems that would eventually replace their work, or possible changes in workload through automation of library systems. For example, the implementation of library-wide systems by the administration, such as a chatbot to answer reference questions, could replace the one currently staffed by humans (which includes both QUFA and non-QUFA staff); or the administration could reduce staffing related to tasks such as cataloguing, metadata creation, or digitization, currently carried out by QUFA and non-QUFA staff.

We discussed several ways of limiting erosion of staffing levels and job quality for QUFA Members. First, bargaining language could set limits to the extent to which library work and teaching can be automated by technology (although we recognize the issues associated with the fact that teaching is carried out by members from various unions).

Second, the union can set limits to what extent Members’ work can be used to generate courses which could be taught without a human instructor by limiting the use of our course and other materials as training data. To this end, the intellectual property rights of faculty, librarians, and archivists would need to be strengthened. However, AI-assisted work raises thorny questions about authorship, ownership, and the reuse of training data; and we note that there are several ongoing court cases and currently no legal consensus about the question of intellectual property in the case of AI.

Special consideration should also be given to term adjuncts as QUFA’s most precarious Members. There is intense worry about the possibility of being replaced by AI. For example, if the administration gains IP rights over adjunct or TF teaching materials, it could use these materials to offer automatically generated online courses, posing a significant threat to term adjunct Members’ ability to acquire the necessary number of courses to obtain SROr, GROr, and Continuing as well as the ability of term adjunct Members who have not yet obtained these rights of reappointments to continue teaching.

A final area of special consideration are provisions for Librarians and Archivists.

Personnel Decisions

There are many sections in the Collective Agreement defining how Members will be evaluated and assessed. Given how error-prone AI is with respect to the correct analysis and interpretation of information, it should not be used in processes which could have adverse material consequences for QUFA members, i.e., in performance evaluation, RTP and Continuing, discipline, and grievance procedures. QUFA should insist on Members' right to be evaluated and assessed by humans, not AI – meaning that the humans responsible for the evaluation should directly read all files or evidence submitted, rather than relying on genAI summaries, that AI detection tools must not trigger disciplinary actions, and that hiring software or employee management software (like Workday) designed to make decisions or recommendations must not be used. There should be acknowledgement that any evidence produced by automatic transcription (for example, through assistive classroom technology, remote meeting recordings) is unreliable.

Special consideration should be given to employment and hiring equity (see next section) due to AI's well-documented gender, racial, and other biases, and the resulting potential to unlawfully harm equity-deserving employees.

Employment and Hiring Equity

It is well-documented that AI has racial, gender, and other biases related to protected groups. One of the most common reasons is that training data sets tend to be skewed, and that AI models inherit the biases of their underlying datasets. In the case of generative AI, which is trained on massive data sets (i.e., the internet), models effectively reproduce dominant societal biases, which is why racial, gender, and class biases tend to be most salient. Current generative AI models are thus heavily moderated by (often severely exploited) human workers who manually clean or extend datasets or score AI responses to fix the most egregious responses. Because it is currently (and as some scholars argue, in principle) impossible to create training data which are completely unbiased, the task force members agree that unbiased AI does not exist and will likely never exist. For this reason, where the deployment of AI can lead to unlawful systemic discrimination of groups protected by the Human Rights Code, AI must not be used.

Such cases include using hiring systems which scan, rank, and evaluate applications because existing research on automated hiring systems has shown that they tend to discriminate applications based on class, gender, and race. For the same reasons, using

generative AI to summarize members' files, rank Members' performance reports, or score Members' renewal, tenure, or promotion files must not be permitted.

Unlike interpersonal discrimination, discrimination due to hiring or evaluation systems has system-wide effects (i.e., applies to all applicants/materials processed by a system), meaning that the scale of such discrimination reaches the threshold of systemic discrimination. For example, a hiring system with gender bias would discriminate against all applicants identifying as cis-women, trans, or non-binary; using generative AI to evaluate Members' performance reports or files could systematically allocate lower scores to Members who are cis-women, trans, or non-binary. Due to correlations between membership in a protected group and other measurable factors such as word choice (so-called "proxy categories"), this bias persists even when the AI is not directly given information about a prospective hire's gender, race, or membership in other protected groups.

We also want to flag that other automated workplace technologies, such as lecture transcription technology, can be discriminatory, notably against English speakers who do not have a Canadian, US, or English accent; and QUFA members should not be required to use or suffer adverse consequences from such faulty technology.

Academic Freedom

Members have a right to academic freedom in matters of their professional practice. While this includes the right to use AI for reasonable purposes by default, it must also include the right to decline to use AI. We discussed a number of interrelated reasons why a Member may want to refuse AI:

- (1) *Professional standards*: It is important to recognize that a Members' refusal of AI can be directly tied to the professional standards of their scholarly field. For example, a Member with a research/teaching focus on sustainability may choose to decline to use AI because of its enormous carbon footprint and resource intensity; a Member with a research/teaching focus on racism may not use AI because of its known racist biases; a Member with a focus on social movements and political repression may refuse using AI because of its involvement in the rise of digital authoritarianisms; or a Member with a research/teaching focus on labor or global inequalities may decline use of AI because of the global labor exploitation on which it relies; and so forth. Members should not be required to use a system which would contravene the professional standards of their discipline.
- (2) *Equity*: Similarly, a Member may decline use of AI due to equity reasons. For example, they may decline using automated essay grading technologies with known

racial biases to avoid disadvantaging racialized students. Beyond AI, we can also make connections here to the administration's imposition of Endpoint which particularly disadvantages Members with caretaker responsibilities. Members should not be required to use a system which would compound already existing disadvantages.

- (3) *Indigenous sovereignty*: Since control over and property status of information fed into automated systems is uncertain, and companies tend to claim property or control over user inputs, Members whose teaching or research includes information about Indigenous communities may choose to not disclose such information to AI systems in accordance with Indigenous data sovereignty – the principle that Indigenous communities have control over data collection processes, and that they control how this data is used (see e.g., [OCAP Principles](#)). We can imagine similar arguments where the building of data centers interferes with Indigenous territorial sovereignty or causes water or energy scarcity (due to servers' cooling and electricity requirements) on Indigenous territories. Members should not be required to use a system which contravenes commitments to Indigenous sovereignty and reconciliation.
- (4) *Intellectual property*: The uncertainty of who owns the inputs and outputs of AI systems also creates problems for intellectual property. Members should not be required to use a system for which they would lose their IP as a condition of use.
- (5) *Privacy & security*: Finally, a Member may choose to not use specific technologies to protect their own or others' confidential information from external interference. We can think again of the administration's imposition of Endpoint which compromises the confidentiality of sensitive research data or client information, as colleagues from Law have outlined. In the context of the US state's wide-ranging powers to access data stored on cloud servers regardless of territorial jurisdiction, Members' may decline using grading technologies which could risk their students' safety, or cloud technologies which could compromise their own mobility and physical safety. Members should not be required to use a system which would compromise the ethical standards of their profession, or their privacy or security or that of the people they work with.

Thus, there are many reasons why Members may refuse AI in their professional practice, and they should be protected to do so. Refusal to use AI also should not have a negative impact on evaluations of Members.

Workload

As discussed above, Members should retain the right to decline to use AI – based both on personal academic freedom and on professional standards. However, there may be problems if departments formulate their workload standards with AI in mind (or are pushed to do so by the administration), which could lead to situations in which Members' may not be in a situation to exercise their freedom to decline AI usage. This is especially the case when standards are based on speed and volume (e.g., quick grading for large classes of students) rather than quality. Even when expectations are not formally set, they can change informally with the adoption of specific technologies. For example, in studies of workplaces where generative AI has been adopted, even when workers are not explicitly required to use AI, one of the results of this adoption is an increasingly high expectation for speed and volume of work and an eroded work-life balance. This is bad for all employees, both those who choose to use AI and those who decline its use. Therefore, for the protection of all Members – both those who choose to make reasonable use of generative AI and those who decline – workload standards should be defined based on the reasonable abilities of a Member who works unassisted by AI.

Privacy

Currently, the CA leaves it up to the administration to specify all parameters of data collection, retention, analysis, and access. As a result, the University can invoke managerial rights to implement broad surveillance mechanisms that are harmful to Members' privacy, security, and academic freedom. AI's ability to sift through enormous troves of data means that the ability of the University or third parties to harmfully surveil Members is continually increasing.

Due to the way data is stored by AI software and due to legislation like the U.S. CLOUD act and related negotiations over [Bill C-2](#) and the [Canada-U.S. Cloud Agreement](#) which would extend US law enforcement into Canadian territory, any entry of Members' data into an AI or into US-owned cloud software (e.g., OneDrive) is accessible by the U.S State. This entry is often done without Members' consent, by third parties such as students or University administrators who have access to Members' data. Microsoft recently clarified that it would release data to the U.S. State if asked to do so, and it is reasonable to assume that other cloud providers would do the same. Cloud providers are also plagued with security issues that are impossible to address locally – for example, Microsoft recently disclosed that it had [exposed confidential emails](#) to its AI chatbot Copilot, even if organizations had nominally restricted Copilot access to their emails. Since Queen's is using Microsoft

services, this means that the confidentiality of QUFA Members' emails during this period have all been compromised.

Thus, platforms, cloud storage, and cloud computing services, especially those provided by US corporations, pose a potential risk to members' privacy, ability to carry out professional responsibilities, or even physical safety. For example, Members using the University's cloud services risk exposing sensitive information (research collaborator/informant information, client information, student information) to the U.S. or other actors; Members travelling to the U.S. may experience safety issues at the border; faculty have been known to be unable to access platforms from certain countries due to current US sanctions; and more. Thus, disclosure of Member data to a LLM or similar AI should be taken very seriously, and understood, from a privacy standpoint, as disclosing Member data to a person or organization external to Queen's, or a state external to Canada.

We would encourage QUFA to work with the administration towards finding more secure workplace technology which better respects Members' privacy and safety, because the compatibility of our workplace technology with our Collective Agreement rights and Canadian labor laws and civil rights will likely remain an enduring issue. There are already good user-friendly solutions available (for example, Nextcloud + LibreOffice; BigBlueButton), and since there is rapidly growing demand by organizations for alternatives to Big Tech services, the variety and robustness of privacy-friendly workplace technology will only improve over the next years. Since these alternative services are also a great deal more economical than Microsoft subscriptions, such arguments may indeed be heard by the administration.

Final Notes

In the lead-up to bargaining, the AI Task Force is willing to assist with preparing materials that can help explain the issues, providing citations and concrete examples. These materials may be useful for educating the Bargaining Team as well as for educating general Members of QUFA. We are very pleased about QUFA's recognition that the issues discussed in this report are labor issues which merit serious consideration, and we hope that QUFA will continue to educate itself and its Members about them in the future.