LinkedIn AI-Assisted Search & Projects

Data Security & Compliance

LinkedIn

Introduction

Our core value of putting our members first powers all the decisions we make, including how we manage and protect the data of our members and customers. We remain committed to this value as we leverage AI in the development of our services.

As part of its Information Security program, LinkedIn follows industry standards as well as its own best practices to stay ahead of the increasing number of threats facing all Internet services and infrastructure. Our platform incorporates multi-layered security controls which have been independently validated against frameworks such as ISO 27001 and SOC 2. LinkedIn's services are designed to be used in compliance with GDPR, and our Data Processing Agreement sets out our obligations to comply with applicable data protection laws around the world.

Through the provision of transparent and accessible information on our security, privacy and compliance practices, we aim to empower our customers to use our services with confidence.

LinkedIn Trust Center

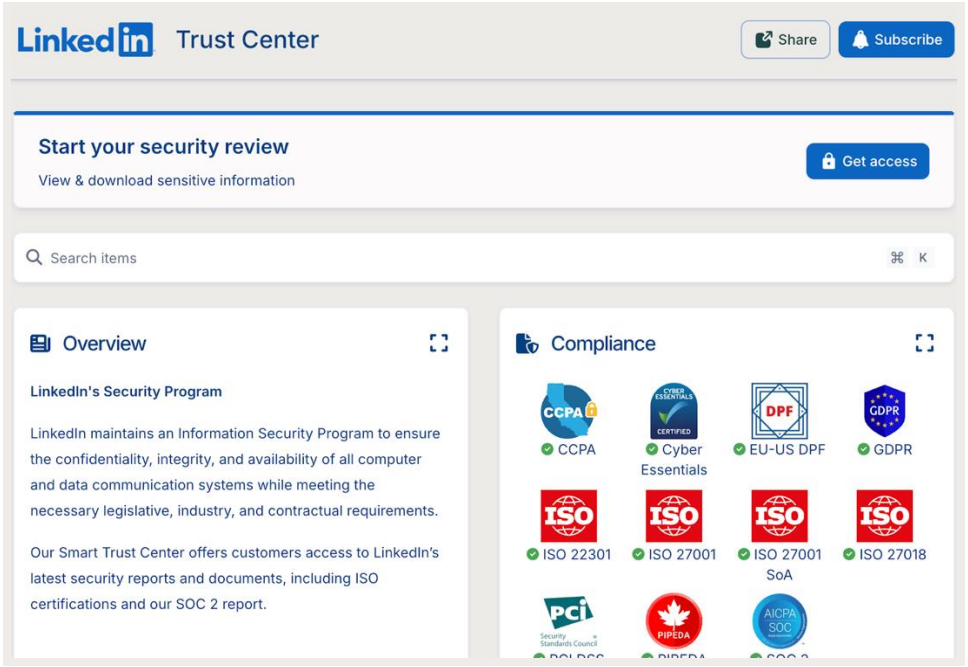

Visit LinkedIn's Trust Center for a suite of documentation on data governance and compliance assessment.

Responsible AI principles

LinkedIn was founded with a clear vision to create economic opportunity for every member of the global workforce. In 2025, we are seeing transformative advances in AI that have the potential to help us accelerate our progress toward that vision.

AI is not new to LinkedIn. LinkedIn has long used AI to enhance our members' professional experiences. By leveraging the power of AI, we help our members connect, increase productivity and achieve success in their careers.

While AI has enormous potential to expand access to opportunity and ultimately transform the world of work in positive ways, the stakes are high. The use of AI comes with risks and potential for harm. That's why, consistent with our commitment to build a trustworthy platform, we must continue to use AI responsibly. Inspired by, and aligned with, Microsoft's leadership in Responsible AI, we are sharing the Responsible AI Principles that we use at LinkedIn to guide our work:

- **Advance Economic Opportunity:** People are at the center of what we do. AI is a tool to further our vision, empowering our members and augmenting their success and productivity.
- **Uphold Trust:** Our commitments to privacy, security and safety guide our use of AI. We take meaningful steps to reduce the potential risks of AI.
- **Promote Fairness and Inclusion:** We work to ensure that our use of AI benefits all members fairly, without causing or amplifying unfair bias.
- **Provide Transparency:** Understanding of AI starts with transparency. We seek to explain in clear and simple ways how our use of AI impacts people.
- **Embrace Accountability:** We deploy robust AI governance, including assessing and addressing potential harms and fitness for purpose, and ensuring human oversight and accountability. We are committed to learning from, and helping, others as AI best practices, norms and laws evolve.

Underlying these principles is our commitment to listen and learn about how AI can continue to be a tool to accelerate progress towards economic opportunity for all.

What is AI-Assisted Search & Projects

LinkedIn's AI-Assisted Search & Projects in LinkedIn Recruiter is a generative AI (GAI) feature that can help you save time by enabling you to quickly search for candidates, create projects, and get recommendations just by simply inputting your hiring needs in your own words.

Development of AI at LinkedIn

In building and approving AI features, we have a design review process to ensure that product design aligns with LinkedIn's Responsible AI Principles and privacy and security standards. This includes steps such as: evaluation of data privacy controls, review of AI disclosures, measurement of AI output quality, integration of output moderation, and implementation of quality and trust measurement systems.

LinkedIn is committed to ensuring all of the products we provide can be used in compliance with applicable law, our privacy policy, and customer agreements as applicable. We also conduct risk assessments and data protection impact assessments for LinkedIn products that incorporate generative AI (GAI) models to power features that create or suggest content, in accordance with applicable law.

A thorough security standard has been set up at LinkedIn to lessen the risks linked with the application of Generative AI technology. This standard encompasses requirements for the complete lifecycle of Generative AI models, ranging from their creation and training to their implementation and supervision. Generative AI models that are created and upheld by third parties for LinkedIn are required to comply with secure software development controls. This involves consistent auditing and penetration testing to guarantee the security of the models and their adherence to LinkedIn's policies. LinkedIn performs an internal penetration test before releasing the functionality to the public. External penetration tests are conducted annually thereafter.

Addressing bias

LinkedIn strives to ensure that its products and services are fair, including by measuring and mitigating algorithmic bias. Our objective is to ensure equally qualified members receive equal treatment by our models. As with all our products, our teams continuously assess our systems, and if harmful biases are identified, we will work to address them.

LinkedIn conducts fairness and bias reviews on a per-model level, including models used in talent solutions products. When potential biases are identified, the team investigates and, as appropriate, retrains the model to mitigate the biases (or deramps the model).

Performance monitoring

We have ongoing monitoring of stability, latency, and output quality signals. We perform regular evaluation to ensure consistency of AI model quality. We have processes to monitor fulfillment failure rates and user quality feedback signals.

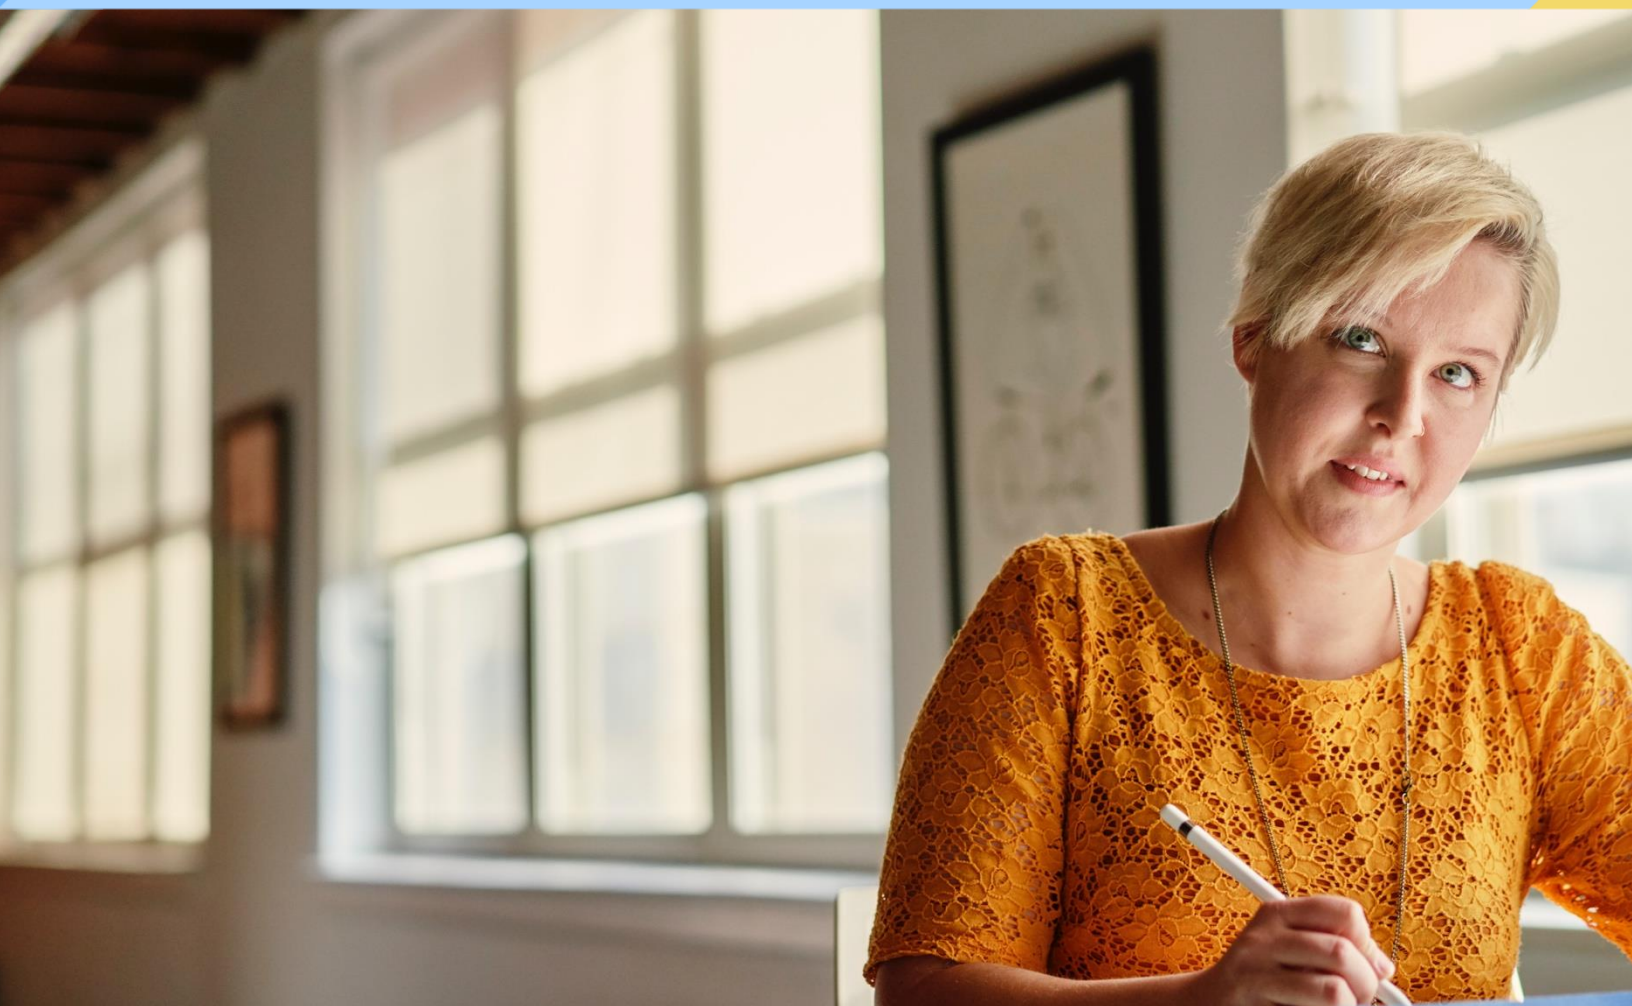

Data Storage and Processing

Processing of input and output data

As a user of LinkedIn enterprise products and services, the information that you choose to provide as inputs, as well as the responses to your inputs and your use of this feature, will be associated with your enterprise seat and will be considered "customer personal data" under our Data Processing Agreement. While your prior "sessions" with these features may not be viewable by you after the close of a session, we will retain this information beyond the session. Like other personal data on LinkedIn, the purchaser of your enterprise seat can request to see what we have retained and request to take action on your personal data. Enterprise users must submit their deletion and export requests to their master administrator.

Input data may be used to improve the AI in two ways:

1) we look at searches, particularly ones with feedback, to help us understand where we can improve our translation of input into a search and to identify areas to prioritize for product improvements. This is similar to what we've always done with search data.

2) As with other search data, we constantly run experiments to improve search ranking and retrieval and to understand how changes to our models improve outcomes for search. For instance, we might make changes to the way we create queries from user input and we will run an experiment to see if this is leading to more successful conversations between recruiters and candidates. We may also use this data in the future to train in-house models, for example to classify whether a request is a search creation request or a project creation request. LinkedIn data will not be stored or used to train third-party AI models, including OpenAI.

Processing of sensitive information

AI-Assisted Search & Projects in Recruiter does not require the processing of sensitive personal data provided by customers. Any data customers provide is subject to the LinkedIn Data Processing Agreement, including with respect to confidentiality, security, and privacy.

Information that members share on Linkedin.com like their names, job title, and additional information on their profile are used to match their information with recruiters' search results. However, this data is not considered sensitive customer data and is provided by members on Linkedin. If users input an ask to the AI feature for sensitive personal information, they will get an error that it's not supported.

Data Training

Data collected for the AI tools in LinkedIn Recruiter come from LinkedIn's economic graph -- information that members and companies share and update on LinkedIn.

The generative AI (GAI) models currently used to power LinkedIn AI-Assisted Search & Projects in Recruiter were developed by OpenAI and LinkedIn uses these models through Microsoft's Azure AI services. If you wish to learn more about how these models were trained, please see OpenAI and Azure AI services.

Third-party technology

Third-party technology is not permitted to use customer data to train their models. LinkedIn's AI-powered features, such as those in LinkedIn Marketing Solutions, ensure that customer inputs into LinkedIn's Generative AI (GenAI) features are not sent to OpenAI. Furthermore, Microsoft has committed not to use these inputs for further training of AI models. This is supported by the standard Azure OpenAI Services terms, which specify that LinkedIn data will not be stored or used to train third-party AI models, including those by OpenAI.

Personal data

We have implemented several security controls to ensure that customer user personal data is not used or accessible for training LinkedIn's Generative AI (GenAI) Language Learning Model (LLM). These controls include:

- **Data Processing Agreement Compliance:** LinkedIn's use of customer data, including for AI features, adheres strictly to the terms outlined in the Data Processing Agreement (DPA). This agreement specifies how customer data can be used, ensuring it aligns with privacy and data protection standards.
- **Microsoft Assurances:** Microsoft processes data in accordance with standard Azure OpenAI Services terms. This ensures that LinkedIn data, including customer inputs into LinkedIn's GenAI features, will not be stored or used to train third-party AI models, including OpenAI.
- **Commitment to Data Privacy and Security:** Microsoft has committed not to use customer inputs for further training of AI models. This commitment is part of the broader data privacy and security measures outlined for Azure OpenAI Services, ensuring that customer data is protected and used in compliance with applicable regulations.
- These measures collectively ensure that customer user personal data is protected and not used to train LinkedIn's GenAI LLM model, maintaining the confidentiality and integrity of customer data in line with LinkedIn's privacy and security commitments.

Keeping our members in control of their data

In regions where LinkedIn or its affiliates use member data to train generative AI models for content creation, members can choose to opt-out of having your personal data and content you create on LinkedIn used for training (including fine-tuning). To opt out, use the Data for Generative AI Improvement member setting. Opting out means that LinkedIn and its affiliates won't use the user's personal data or content on LinkedIn to train models going forward, but does not affect training that has already taken place. We are initially making this setting available to members whose profile location is outside of Canada, the EU, EEA, UK, Switzerland, Hong Kong, or Mainland China. If the members live in the stated regions above, LinkedIn will not use personal data or content on LinkedIn to train or fine-tune generative AI models for content creation without further notice. In addition to the setting set forth above, members can also object to the use of their personal data for training non-content generating GAI models using the LinkedIn Data Processing Objection Form, accessible here.

The opt out setting on LinkedIn.com allows individual LinkedIn members to control if their LinkedIn data can be used to train generative AI models for content creation - this includes any personal data, including personal data in LinkedIn Recruiter. If the member opts-out using that setting, we won't use their Recruiter personal data to train GenAI models for content creation. The controls for data training of AI models will remain at the member level to keep our individual members in control of their own data.

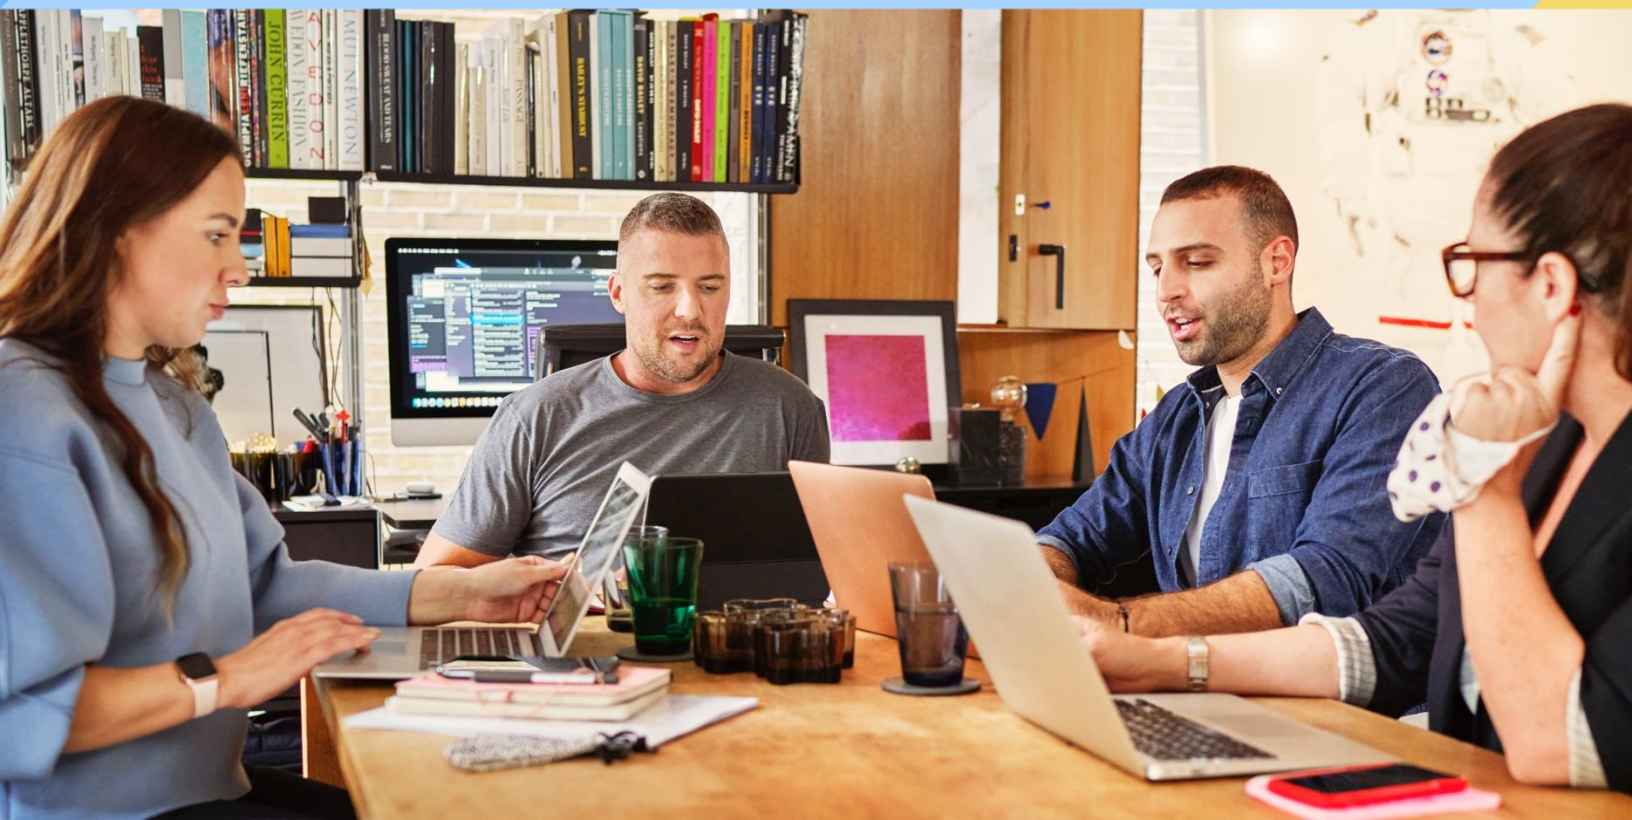

AI-Assisted Search & Project's AI explainability

AI-Assisted Search enables recruiters searching for potential job candidates to provide a plain language input into the search tool in LinkedIn Recruiter, which then gets translated via AI into a structured search to leverage Recruiter's existing search functionality and filters (e.g., title, industry, schools, companies, and location) to obtain search results. For example, a prompt like "Marketing Managers in Dallas with sales enablement experience" would create the following search filters: job title (Marketing Manager), location (Dallas, Texas), and skills (sales enablement).

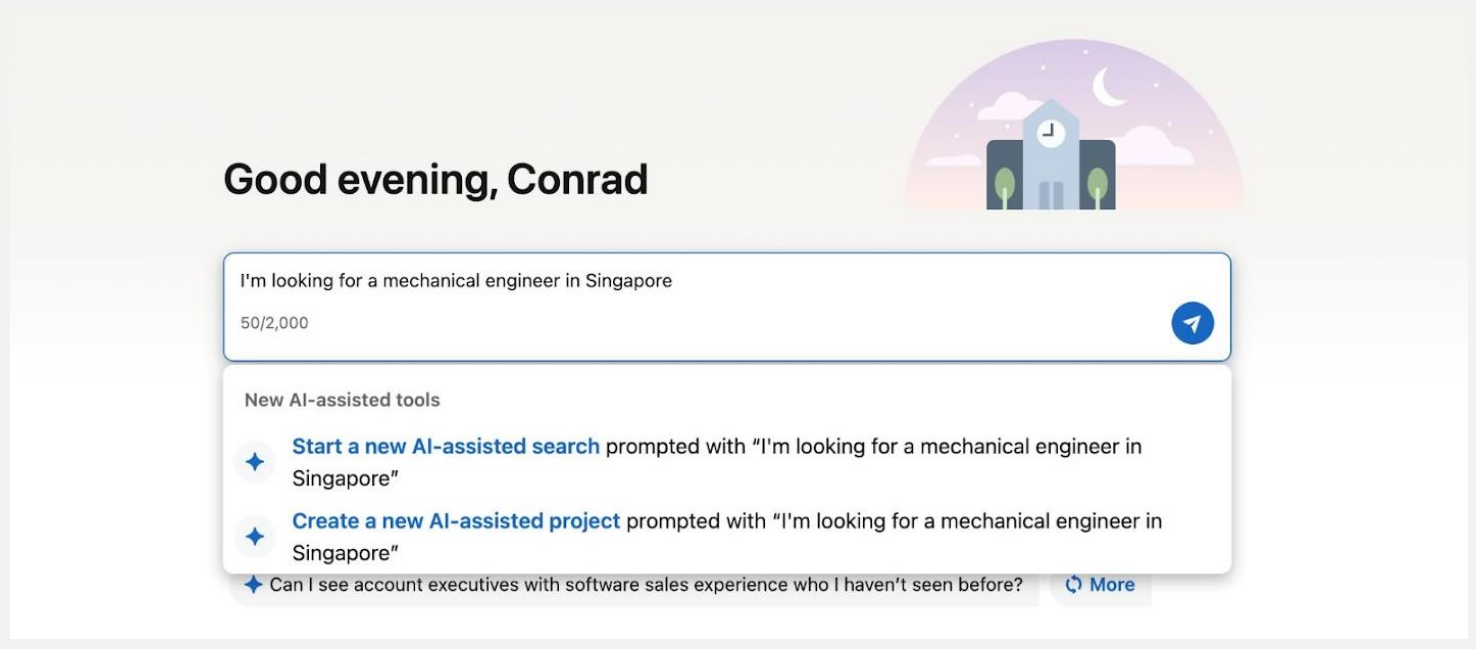

Recruiters can start a search with AI-Assisted Search from the Recruiter homepage or the Talent Pool tab of a project. AI-Assisted Search will analyze what a user wrote and instantly populate the right search filters to match the provided criteria. After initial search results, AI-Assisted Search will also provide the user with additional suggestions.

Recruiters can manually adjust the search filters or try different inputs to reach the desired outcomes. Users also have the ability to provide in-product feedback, via a thumbs up/down, when the results do not meet their expectations so the product can learn and improve over time.

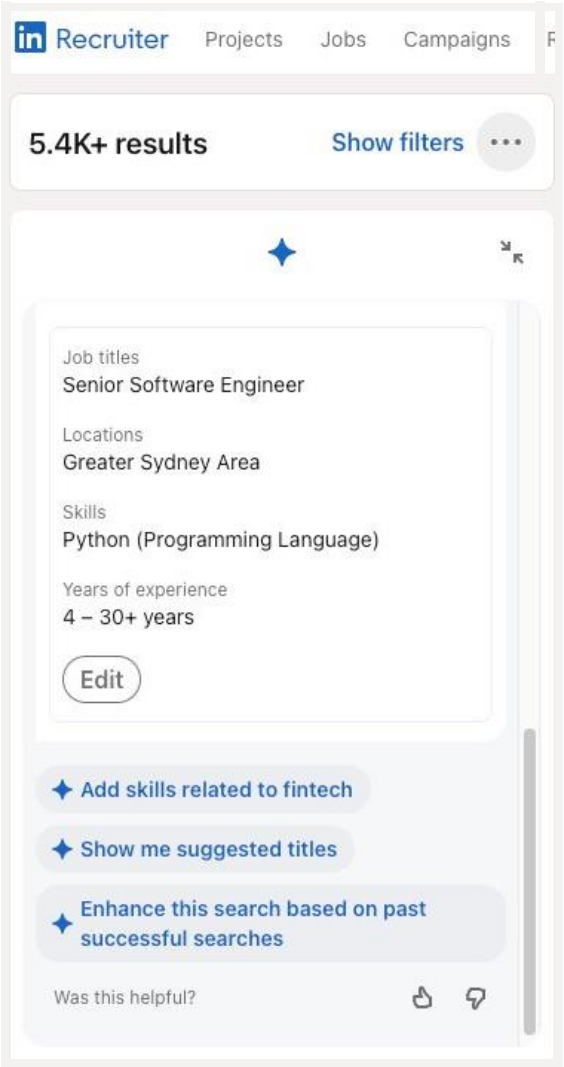

The LinkedIn Recruiter product provides a ranked list of candidates/members corresponding to a search query. Given a search request, candidates/members that match the requested keywords and facets are selected and then ranked based on various factors.

This includes signals such as the similarity of their work experience/skills with the search criteria, and the likelihood of a response from an interested candidate/member, which are weighted using machine-learning models.

Customers choose what qualifications are relevant for open roles and members choose how much information they wish to share with LinkedIn. Members can highlight certain experiences and skills, making them more likely to be responsive to recruiter search efforts.

AI features in Recruiter are primarily recommendations and do not make hiring or sourcing decisions for the users. Users will still own the decision making of engaging or not engaging candidates on Linkedin Recruiter.

Using AI-Assisted Search is not mandatory for users. They still have the ability to run manual keyword or Boolean searches if preferred. To give users more control, they can also manage their default AI settings under AI Preferences within their Account Settings.

AI-Assisted Search has been subject to our review process to ensure it can be used in compliance with applicable laws, our privacy policy and our contracts with customers. Our security commitments with respect to customer data are set forth in our agreement with customers. Like all of our services, AI products are subject to security reviews designed to prevent harm.

AI technology and models used

The generative AI (GAI) models currently used to power AI-Assisted Search & Projects in Recruiter are developed by OpenAI and LinkedIn uses these models through Microsoft's Azure AI services. If you wish to learn more about how these models were trained, please see OpenAI and Azure AI services.

Compliance with applicable privacy, data protection, and AI/ADT laws, regulations, executive orders, and guidance.

LinkedIn complies with applicable law in the jurisdictions where our products are offered and we always seek to enable our customers to comply with their legal obligations when using our products. We are committed to keeping our platform trusted and professional, and we respect the laws that apply to us in the countries in which we operate.

Fair Hiring Practices: LinkedIn's hiring product offerings focus on widening the hiring funnel through candidate sourcing and discovery. Specifically, AI-Assisted Search & Projects in Recruiter helps recruiters build a pipeline of individuals in an effort to encourage those individuals to apply for roles and more efficiently manage candidates to make sure they are engaging with the right people at the right time. Recruiters decide who to reach out to and who to advance in the pipeline once a person has applied. While AI-Assisted Search & Projects delivers insights about an applicant's qualifications against recruiters' hiring criteria, it will not screen out applicants or make hiring decisions. Recruiters will still need to take action and make final decisions on applicants whether they advance or not.

Additional Resources

- LinkedIn's Responsible AI Principles
- Microsoft: Governing AI: a Blueprint to the Future
- Microsoft: Advancing the Future Through Responsible AI
- Data, Privacy, and Security for Azure OpenAI Service
- LinkedIn's Smart Trust Center
- LinkedIn's Public Trust and Compliance page
- LinkedIn Subscription Agreement
- LinkedIn Data Processing Agreement
- LinkedIn Privacy Policy
- Regional privacy web page