LinkedIn Hiring Assistant

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 new products and 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

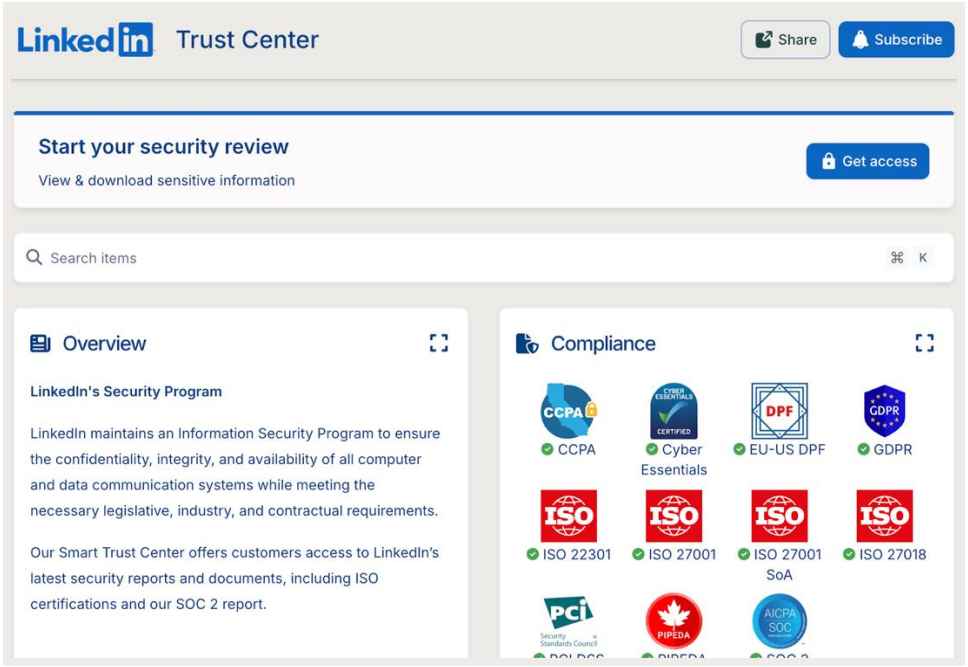

Visit LinkedIn's Trust Center for a suite of documentation on data security, governance and compliance.

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 Hiring Assistant

LinkedIn's Hiring Assistant product is an AI agent embedded within the existing LinkedIn Recruiter & Jobs experience.

It streamlines the recruiting process by tackling time-consuming tasks on users' behalf like multiple searches for the same role and composing initial outreach. Please visit this page for more details.

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.

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.

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.

In LinkedIn Recruiter, the candidates surfaced from the search filters will always reflect representative binary gender results.

This means, if there are 6500 potential engineering candidates (40% women, 60% men), we'll display 40% of women on each page of the search results to more fairly represent the available pool. To prevent LinkedIn members from being excluded due to their demographics (e.g., race, gender, disability), LinkedIn does not provide any search filters to target specific demographics in Recruiter search. We attempt to prevent customers from targeting specific demographics (e.g., find black engineers) with the use of error messages, and the search would not be executed.

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).

Data Storage and Processing

Processing of input and output data

As users of LinkedIn enterprise products and services, the information that customers choose to provide as inputs, as well as the responses to those inputs and use of the feature, will be associated with their enterprise seat and will be considered "Customer Personal Data" under our Data Processing Agreement. While prior "sessions" with these features may not be viewable after the close of a session, we will retain this information beyond the session. Like other personal data on LinkedIn, the purchaser of the enterprise seat can request to see what we have retained and request to take action on their 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

Hiring Assistant 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 LinkedIn Hiring Assistant utilizes data collected from LinkedIn's economic graph, which comprises information that LinkedIn members and companies share and update on the platform. This data includes but is not limited to LinkedIn profiles, work experiences, skills, education, certifications, location, and job postings.

Runtime or Inference Data

During the runtime or inference phase, Hiring Assistant uses real-time data inputs provided by recruiters and ongoing updates from LinkedIn profiles.

- **Suitability:** This data includes specific information on job roles, qualifications needed, skills, and candidate profiles, courtesy of recruiter inputs.
- **Sourcing:** Real-time data is sourced from updated LinkedIn profiles and job descriptions provided by users in the hiring workflow.
- **Completeness:** The data is enriched through consistent updates by members and companies, ensuring current information is always used.
- **Representativeness:** Runtime data captures live member activities and job market trends, providing an up-to-date representative dataset.
- **Accuracy:** Automated validation processes and member feedback contribute to maintaining precise and relevant data during the runtime phase.

Ensuring consistency between the training and runtime datasets is crucial:

- **Alignment:** Both datasets leverage the same data sources and fields of information.
- **Continuous Learning:** Feedback loops from recruiter interactions and outcomes feed back into the training data to continuously refine and enhance LIHA's performance.
- **Validation:** Data consistency and quality checks are employed to evaluate alignment and accuracy, ensuring seamless interoperability.

The generative AI (GAI) models currently used to power Hiring Assistant 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.

Customer data is not used to train third-party AI models

Customer data will not be used to train third-party AI models, including those developed by OpenAI and Microsoft. LinkedIn ensures that customer inputs into LinkedIn's AI-powered features are not sent to third-parties through several controls, including:

- **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.

Hiring Assistant's AI explainability

The new Hiring Assistant streamlines a recruiter's search for potential job candidates. A recruiter provides their hiring goals and role details (and optionally link to existing LinkedIn job post), and Hiring Assistant creates searches to source candidates, initially crafts outreach InMails, and helps Recruiters build a shortlist of applicants for the role. The Hiring Assistant will only do tasks the recruiter wants them to.

As part of the intake and sourcing workflow, recruiters can upload notes from an intake meeting, job description, link to a member profile, and job posting into the Hiring Assistant intake tool within LinkedIn Recruiter. The Hiring Assistant then parses through that information and translates that information into an outline of the role details, which are used for sourcing and in candidate outreach messages (e.g. job titles, locations, workplace type, employment type, seniority level, company hiring for). For example, meeting notes from a discussion with a Hiring Manager that state *"looking for Marketing Managers in Dallas, that has Sales Enablement experience"* could create the following role details: job title (Marketing Manager), location (Dallas, Texas), and skills (Sales Enablement). Recruiters can manually adjust the role details to ensure accuracy.

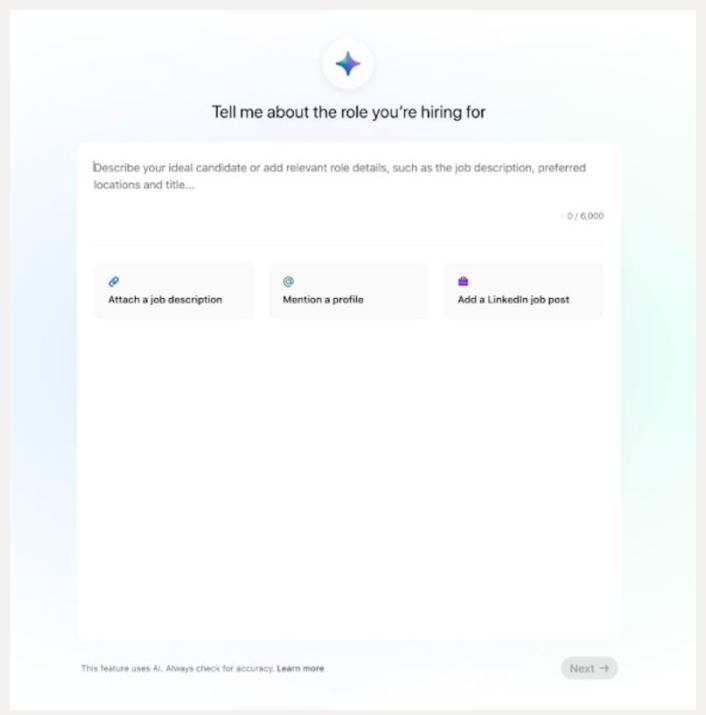

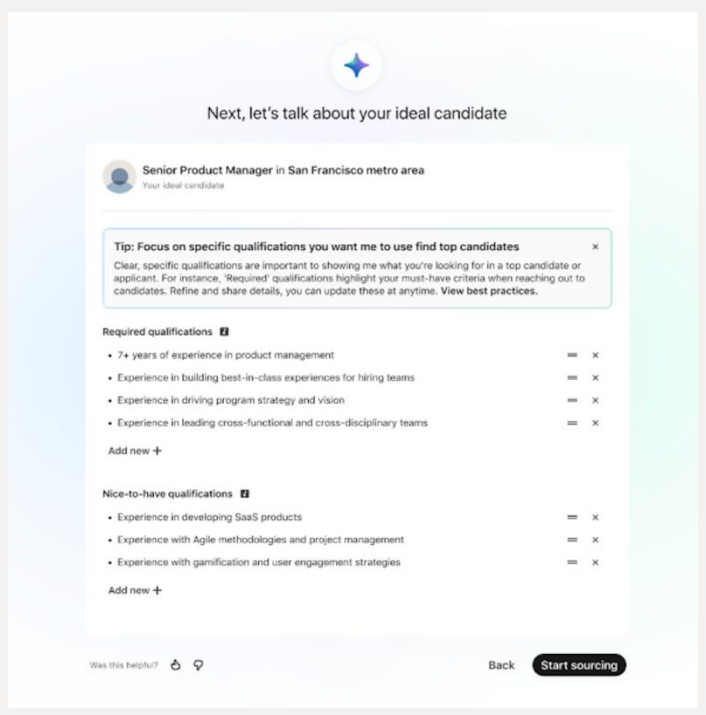

Next, LinkedIn Hiring Assistant suggests recommended required & desired qualifications for the ideal candidate that will be used in the back end to find candidates (e.g. education, years of experience, skills). Users can edit and refine this directly.

Once the role's details and qualifications are confirmed, a hiring plan is served to the user, suggesting an InMail goal and target number of candidates to pipeline; this can be edited. When candidates (both job applicants and those sourced by the Hiring Assistant) are ready for review, the user is notified. The Hiring Assistant will not move applicants into the pipeline but will evaluate and categorize them based on user-defined thresholds of requirements met.

Each candidate profile includes a summary of why the Hiring Agent thinks they are a good fit and how many qualifications they meet, based on a combination of the following sources:

- The candidate's LinkedIn profile (the following sections, if complete: work experience, skills, education, location, certifications, summary, volunteer experience, publications, and patents)
- Resume (if available)
- Responses to screening questions (if available)
- Recruiting notes (if available)

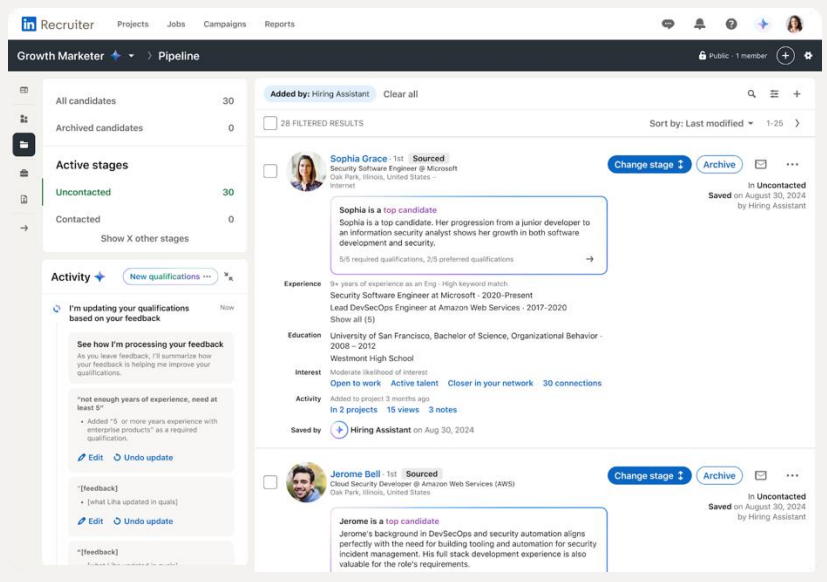

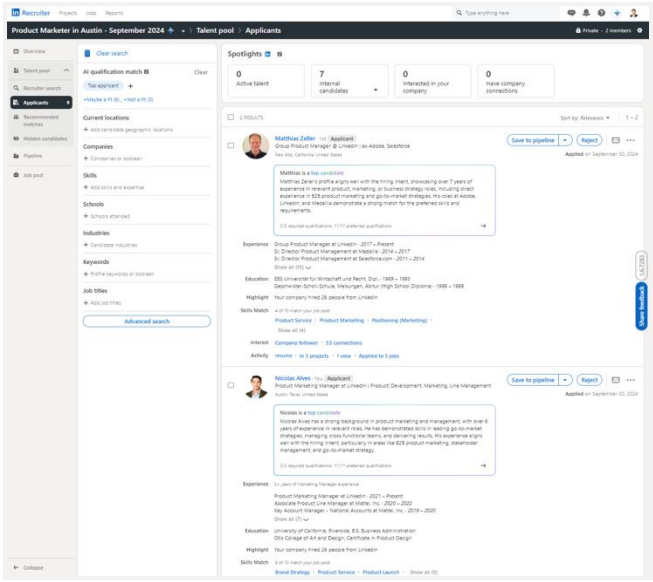

The Hiring Assistant will also prompt the recruiter to reach out to the candidates. When the recruiter wants to take action on this, the Hiring Assistant will bring the recruiter to the messaging tool. Using our existing AI-assisted messaging feature, the Hiring Assistant will draft outreach messages using data that our customers and members have chosen to share with us that is already input into our platform. The AI helps format and create all the elements of an outreach message. Certain fields include:

- **The opportunity:** job title, responsibilities, location, salary, workplace type, employment type, seniority, etc. for messages with job context
- **Good fit:** prioritized skills/skill match with job title, experience, past companies/roles, past interest in company, past applications, open to work, current/past schools, etc.
- **Recruiter/Company/Team background:** recruiter info, company info, culture, values, etc. (share language here from AI-assisted messaging).

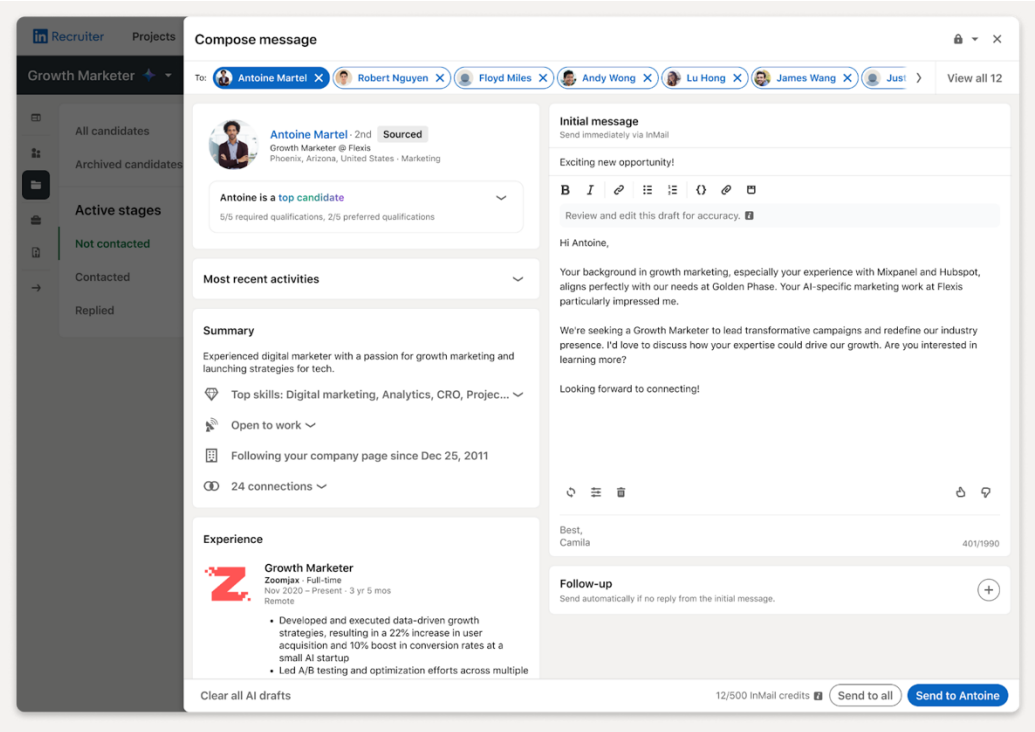

LinkedIn Hiring Assistant learns from the user's actions (e.g. if a candidate is archived, or contacted, a prompt surfaces to ask why this action was taken so it can feed learnings back into the sourcing model). LinkedIn Hiring Assistant presents what it has learned based on these user signals so users can review & improve our ability to deliver on recruiters' desired pipeline.

LinkedIn Hiring Assistant 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 Hiring Assistant in Recruiter are:

- **Intake and search:** 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.
- **Messaging:** a (text) Generative AI model – We use a combination of in-house language model for English and GPT for non-English languages. The in-house models are trained using machine learning to adapt to recruiter messaging use cases and generate draft messages that recruiters can edit and send. The models are able to generate content unique to a recruiter-candidate pair based on available parameters provided to it, for example: member profile info (position, company, open to work preferences, etc.), recruiter profile info (name, job title, etc.), recruiter company info (company summary), and job post info (job title, location, job type, etc.).
- **Summarization** of candidates: an in-house LLM which takes in hiring intent (list of required and preferred qualifications), candidate LinkedIn profile, candidate resume, screening questions, recruiting notes as input and generates a response about the qualifications the candidate matches, partially matches, or does not match. The output includes the source, summary of the qualification match, and an overall candidate summary.

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.

Privacy and Data Protection: LinkedIn takes a global approach to privacy and data protection, applying GDPR as our baseline compliance standard for all jurisdictions. LinkedIn's contractual Data Processing Agreement (published at https://legal.linkedin.com/dpa), addresses LinkedIn's obligations to comply with data protection laws around the world, including but not limited to the GDPR, LGPD, and CCPA.

AI Regulation: 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, including EU legislation like the AI Act. As the legal and public policy landscape relating to AI continues to develop, LinkedIn is closely monitoring emerging requirements which may apply to our services.

Fair Hiring Practices: LinkedIn's hiring product offerings focus on widening the hiring funnel through candidate sourcing and discovery. Specifically, Hiring Assistant 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 Hiring Assistant 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
- Regional privacy web page
- 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