

DINQ User Manual

The talent search engine built for the AI era. Not a resume. Real work, made discoverable.

About this guide

This guide takes you from zero to confident with DINQ: how to find the right person in a single sentence, how to read a candidate's *real work*, and how to shortlist and reach out in a way that actually gets replies. Everything is organized around what *you* want to accomplish, with screenshots to follow along.

Whether you're a hiring lead, an engineering manager, an investor, or scouting a co-founder, think of DINQ as your AI sourcing partner.

1. What is DINQ

DINQ is an AI-powered talent search engine. You describe the person you're looking for in plain language, and DINQ searches 500+ data sources at once — LinkedIn, GitHub, Google Scholar, personal sites, papers, open-source projects, technical communities — then assembles the scattered evidence into a candidate profile you can trust, evaluate, and act on.

What makes it different from traditional sourcing tools:

- **It finds *work*, not *resumes*.** In the AI era, real ability lives in code, papers, and projects — not in a self-written CV. DINQ makes those real signals discoverable by AI.
- **Conversation replaces filters.** No Boolean strings. Just describe what you need, the way you'd tell a colleague.
- **The process is transparent and every result is traceable.** Each conclusion comes with source links, so you can see *why* DINQ reached it.




Figure 1: The DING home — describe the person you need in a single sentence.

2. Getting started

2.1 Create an account & sign in

1. Go to ding.me and open sign-in / sign-up from the top right.
2. Create an account with **email verification**, **Google**, or **GitHub**.
3. If you have an invite code, enter it during sign-up to unlock credits.
4. You'll land in your workspace.

💡 Tip: Signing in with GitHub means that when you later build your own DING Page, your open-source footprint can be aggregated in one click.

2.2 Know your workspace

After signing in, the left sidebar is your main navigation. It groups everything DING can do:

Entry	What it does
Search	The core. Launch a talent search in natural language.
Talent Radar	Automated sourcing: scheduled scans, new matches delivered to your Inbox.
Shortlist	Save, organize, and export candidates by folder.
API Playground	Plug DING search into your own systems via API.
Integration	Connect email for outreach; access the API.
Personal	Build your own DING Page (a living profile).
Organization	Create a team / lab / community page.
Inbox	In-app messages and notifications.

Your remaining **Credits** are always shown in the top right. The interface is available in English, Simplified Chinese, and Traditional Chinese, and you can switch the language at any time.

2.3 What are Credits?

Deep searches, profile enrichment, and revealing contact details consume credits. Your balance has three parts:

- **Subscription credits** — reset monthly with your plan, no roll-over;
- **Reward credits** — earned through referrals, campaigns, and similar;
- **Purchased credits** — bought separately.

The free plan is enough to experience the full flow. You can review every charge under **Settings** → **Subscriptions**.

 **Tip:** A deep search typically costs tens to a few hundred credits, depending on depth and how many candidates you request. Start with a small candidate count to test direction, then scale up.

3. Precision search: find the right person

Search is the heart of DING. The whole thing feels like talking to a seasoned sourcing consultant, in three steps: **state the need** → **refine together** → **review results**.

3.1 Step 1 — Say what you need, in one sentence

Type the person you're looking for right into the search box, for example:

- Find video generation AI researchers

- B2B SaaS growth lead, 3+ years
- Infrastructure engineer who can own vLLM inference optimization

You can also **upload a job description** — text, PDF, or image — and DINQ will distill the key search criteria from it.

Press Enter (or the search button) and DINQ first restates how it understood your need, then asks whether you'd like to refine the criteria together for more precise results.




Figure 2: DINQ confirms its understanding of your need before deciding whether to refine.

3.2 Step 2 — Progressive refinement (optional)

Choose "OK, help me refine my search" and DINQ asks one dimension at a time, fill-in-the-blank style, with clickable suggestions:

1. **Location** — SF Bay Area / New York·Boston / Remote (global) / China (Beijing·Shanghai·Shenzhen)...
2. **Experience level** — 3—5 years (mid) / 5—8 years (senior)...
3. **Education** — PhD required / Master's+ / Bachelor's+ / no strict requirement...
4. **Compensation range** — pick to match your budget;
5. **How many candidates** — around 5 / 10 / 20 / 30+.

For each question you can pick a suggestion or type your own under "Or type your own."




Figure 3: Progressive refinement — pick one dimension at a time as DING turns a fuzzy need into a precise profile.

When refinement is done, DING shows a summary of the brief. If it looks right, click **Start search**.

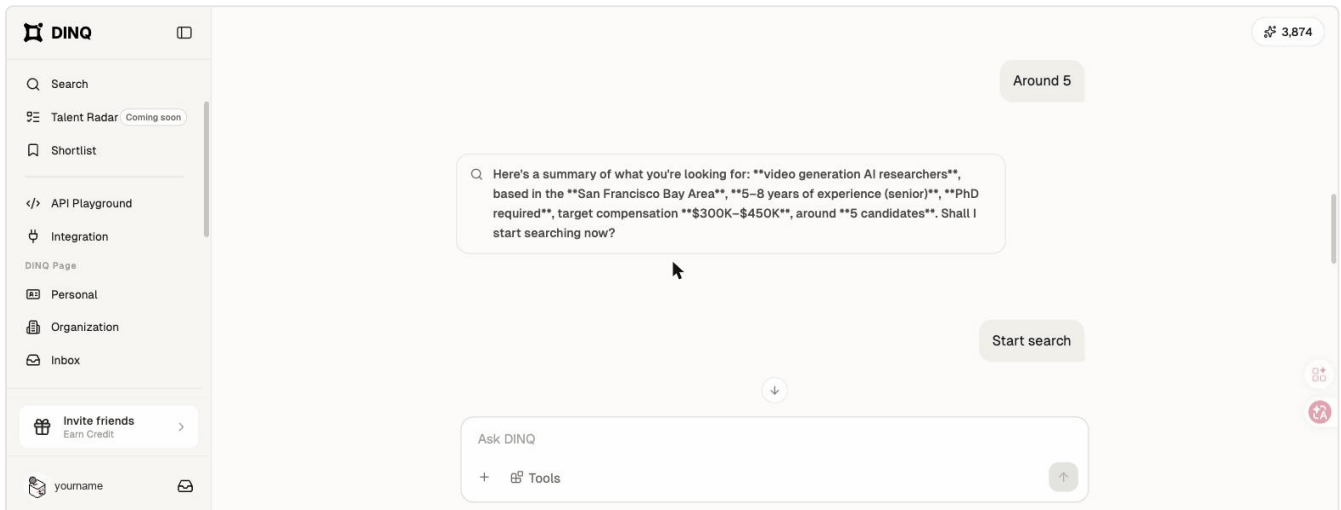


Figure 4: The pre-search brief summary — click "Start search" to begin sourcing.

💡 Tip: Don't want the back-and-forth? Choose "Skip, search directly" at step 1 to search immediately. The more specific your brief, the better the results.

3.3 Step 3 — Watch DING search the web (a transparent process)

DING isn't keyword matching — it genuinely

researches the open web

. As it works, you'll see its sourcing trail in real time: launching parallel searches, cross-checking candidates' LinkedIn / Scholar / GitHub, filling in missing details, broadening the net...

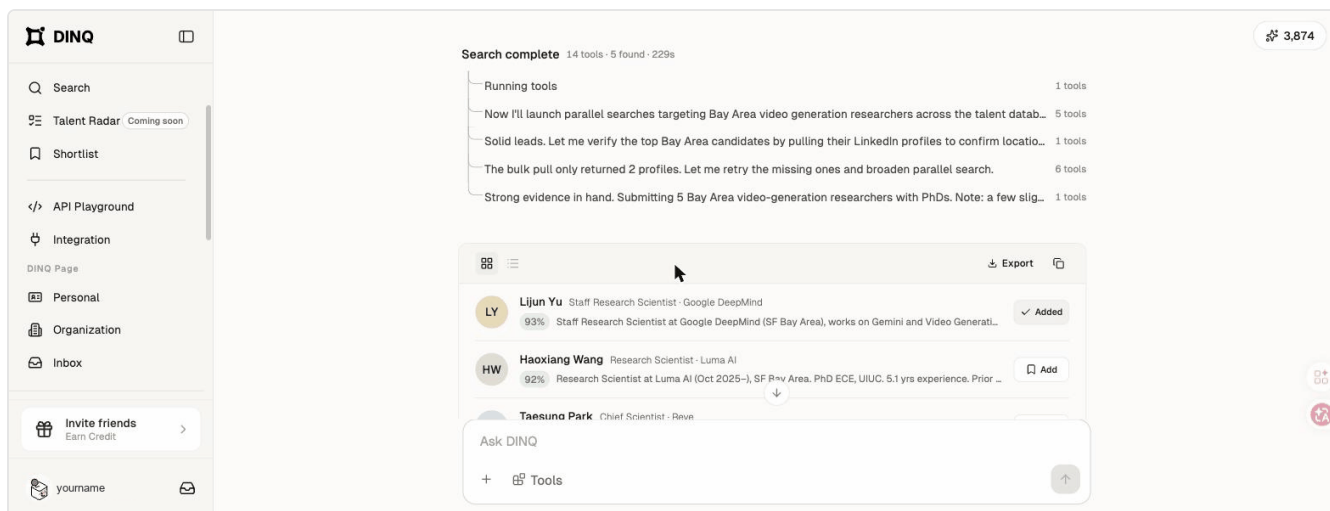


Figure 5: A transparent sourcing process — see which sources DING is checking and why.

💡 Why does a search take a moment? Because DING runs multi-source deep research rather than querying a static table. That's exactly why the results are more trustworthy.

3.4 Reading your results

When the search finishes, candidates are ranked **by match score, high to low**. Each row shows:

- Avatar, name, current title • company;
- A **match percentage** — overall fit across many signals;
- A one-line highlight;
- An **Add** button — save to a shortlist in one click.

Toggle **grid / list view** at the top right of the results, or **Export** them. Below the list, an AI-written summary explains why these people were surfaced, with each one's strengths and trade-offs.




Figure 6: Candidates ranked by match score, with an AI summary and one-click shortlisting.

💡 How to read the match score? It reflects how well a candidate's *public signals*

fit your need. The richer the signals (active GitHub, cited papers, a personal site), the more confident the match.

4. Candidate profiles: see the real work

Click any candidate and DING enriches their profile on demand. Seconds later it opens a *living profile* (the DING Card):

- **Header:** photo, title, location, and social links — LinkedIn / Google Scholar / Twitter / Homepage;
- **Match context:** a short paragraph on *why* they match, with key capability tags (e.g. multimodal foundation models, video generation, reinforcement learning);
- **Education:** schools, degrees, fields;
- **Work experience;**
- **Publications / notable work;**
- **Recent activity:** every item carries a **source link** you can click to verify.




Figure 7: A candidate's living profile — match context, education, work, publications, and recent activity, all traceable to their sources.

This is the fundamental difference between DING and a resume database: you're not seeing what a candidate

claims

about themselves — you're seeing objective evidence aggregated from their public work, with provenance.

From the profile you can directly:

- Click **Shortlist / Add** to save them;

- Click **Get email** (consumes a few credits) to reveal contact details.

💡 Tip: DINQ is an assistant to

expand your search and organize clues

— not a final hiring decision. Treat results as a high-quality starting point and verify with your own judgment.

5. Shortlists: keep candidates organized

Once you **Add** candidates, manage them in **Shortlist**.

- **Folders:** organize by role or project; click **Create Folder** to add one;
- **Table view:** see everyone's name, status, company, title, and domain tags at a glance;
- **Status:** track outreach progress, e.g. "Not obtained / Obtained" (contact);
- **Fuzzy search + status filter:** locate people fast in a large list;
- **Export:** send the list out for sharing across your team or archiving.




Figure 8: The shortlist — organized by folder, with outreach status, search, and export.

💡 Tip: Create a folder per open role and Add promising people as you go. Your shortlist grows with every search.

6. Contacting candidates: built on context

Finding someone is only the start — **outreach** is what converts. DINQ's philosophy: **reach out from context**, never from a mass template.

1. In a candidate's profile, click **Get email** to reveal contact details (a few credits);

2. DING drafts a **personalized** opener based on their public work — referencing something they actually did (an open-source project, a paper, a direction that closely matches the role);
3. Before sending, **edit** the tone, role details, and the way you frame the invitation;
4. Send via the in-app **Inbox** or your connected email.

A note that says "I saw your recent work on vLLM inference optimization — it lines up with where my team is headed, and I'd love to talk" gets replies that a generic recruiting blast never will.

💡 Outreach etiquette: use your real identity and communicate professionally; no harassment or spam; disclose AI-generated content honestly.

7. Inbox

Inbox is your communication hub inside DING, with two tabs:

- **Inbox:** conversations with candidates live here — a conversation list on the left, the chat window on the right;
- **Notifications:** system alerts and activity.




Figure 9: The Inbox — in-app messages and notifications in one place.

8. Talent Radar: put recurring roles on autopilot

⚠️ **Coming soon:** Talent Radar is built and on the way — it currently appears as "Coming Soon" in the product. The section below previews how it will work once it's available.

When a role needs **continuous** sourcing, **Talent Radar** is your automated recruiting assistant — "tell us who you need, we'll handle the rest." Setting up a radar takes about 30 seconds:

1. **Describe the need** in natural language (e.g. "Senior ML Engineer, 5+ years");
2. **Refine across six dimensions** — field, role & company type, location, experience level, education, salary range — guided by AI;
3. **Set execution preferences** — scan frequency (hourly / daily / weekly / monthly) and candidates per run (10 / 20 / 30 / 50);
4. **Launch** — confirm and go; the first results appear **within minutes**.




Figure 10: Talent Radar — describe your need in natural language; the system scans 24/7 and delivers new matches to your Inbox.

Two radar types

- **Talent Search** — continuously match new people around a profile;
- **Candidate** — continuously track specific individuals.

Run states

A radar has five states: **Tracking** / **Completed** / **Paused** / **Failed** / **Cancelled**.

 **Tip:** "Daily new additions" counts only deduplicated, previously undiscovered matches — no repeat noise. Each automated search consumes credits just like a manual one, and the radar pauses automatically when credits run low. Search conditions can't be edited after creation — create a new radar to change them.

9. Build your own DING identity

DING isn't only for finding people — you can have your own DING identity too.

9.1 DING Page (your personal page)

In **Personal**, build your page — a living profile that aggregates your real work, with its own link `dinq.me/yourname` :

- Fill in basics: position, degree, email, location;
- Write a short bio;
- Add **Analysis** blocks: visualizations such as career trajectory and network graph;
- Add **image** / **video** media.




Figure 11: The DING Page editor — assemble your real work into one shareable living profile.

Note: A DING Page is **public** and may be accessed, indexed, cached, or archived by users, visitors, and search engines. Only include what you're comfortable making public.

9.2 Organization

In **Organization**, create a page for a team, lab, or community: set its type (Community / Open source / Lab), members, and description, and get a custom link `dinq.me/orgname` . You can also browse and join recommended organizations.




Figure 12: Organization pages — a DING identity for your team, lab, or community.

9.3 Verification

Complete identity verification under **Settings** → **Verification** to make your page more trustworthy.

10. Advanced: API & integrations

10.1 API Playground

To plug DING's talent search into your own systems, the **API Playground** lets you try it directly:

- Create and manage your API key under **Settings** → **API Keys**;
- In the Playground, set query, number of results, active time, and sources, then run;
- Ready-to-use code samples in **Python** / **TypeScript** / **Curl** are on the right — copy and go.

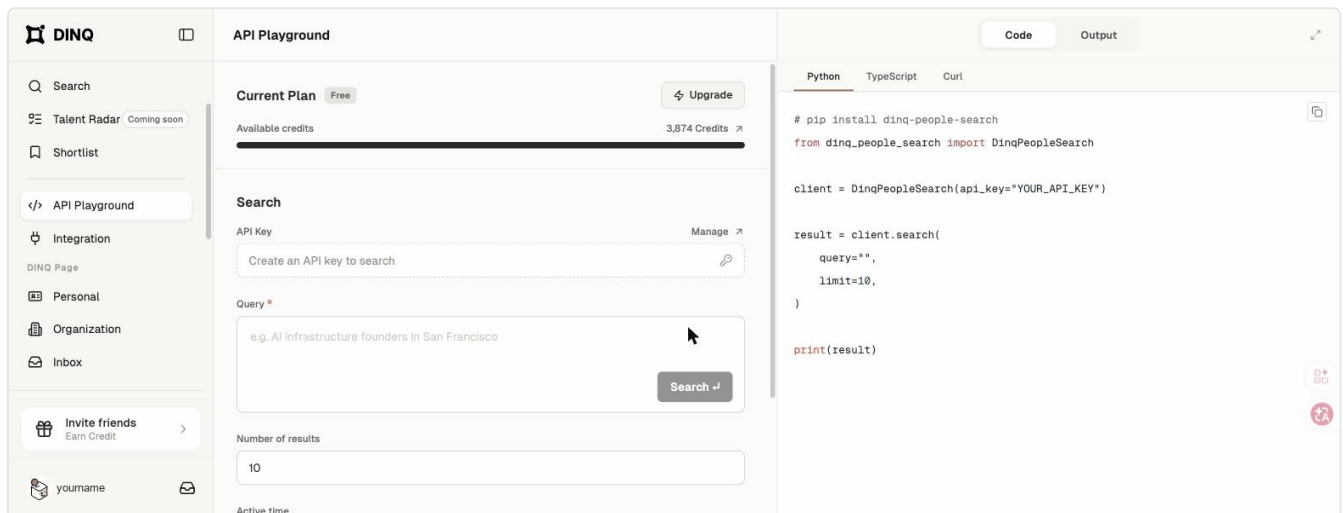


Figure 13: The API Playground — try the talent search API visually, with multi-language code samples.

10.2 Integration

In **Integration**:

- **Email integration:** connect an email account (e.g. Gmail), set your sender name and signature, and send automated outreach sequences;
- **API access:** jump to the API Playground and bring DING into your workflow.




Figure 14: Integration — connect email for outreach, or plug in via API.

11. Account, credits & settings

Click your account name in the bottom left for **My DING / Verification / Settings / Upgrade / Sign out**. In the **Settings** dialog:

- **Profile:** complete details like birthday, gender, and status to get smarter matches (your info stays private);
- **Account:** set a password, bind / unbind email, connect GitHub, and account deletion;
- **Verification;**
- **Invite Friends:** invite others to earn credits;
- **DING Card:** your card settings;
- **Subscriptions:** view your plan, credit breakdown, line-item usage, and billing;
- **API Keys:** manage API keys.




Figure 15: The settings hub — profile, account, verification, invites, subscriptions, and API keys in one place.

Plans & upgrading

DINQ offers a free plan that's enough to experience the full flow. When you need more credits or higher-frequency use, click **Upgrade** (monthly / yearly, with a discount for annual).




Figure 16: Subscriptions & credits — see your balance breakdown and every charge.

Earn credits by inviting friends

"Invite friends · Earn Credit" in the bottom left generates your invite; when a friend signs up, you both earn credits.

Switching language

You can switch the interface language at any time; DINQ natively supports English, Simplified Chinese, and Traditional Chinese.

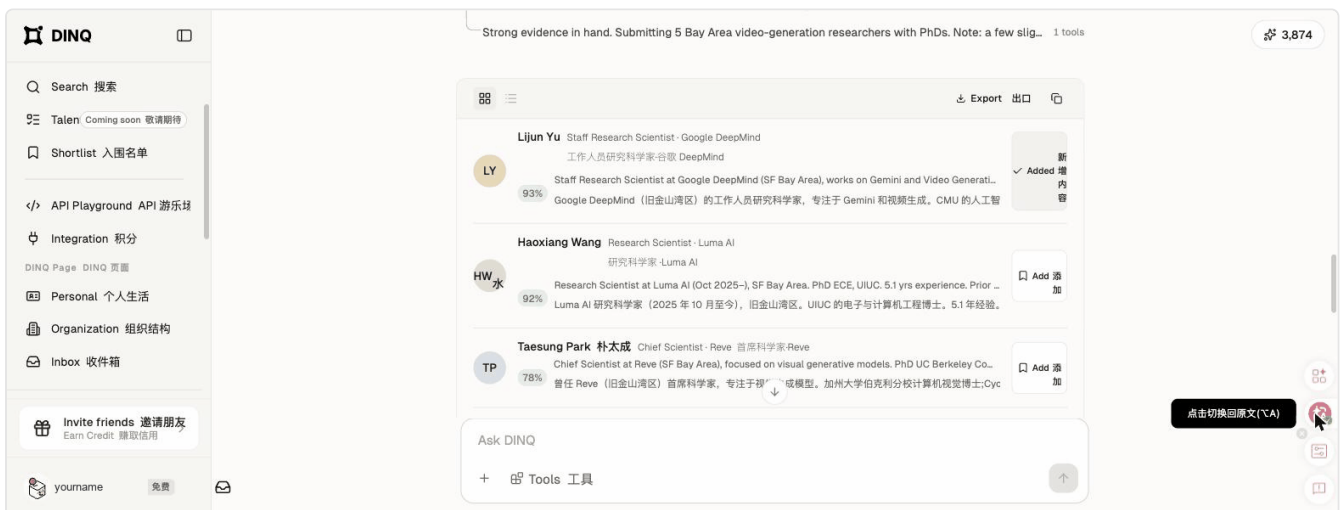


Figure 17: Language switcher — switch the interface between English, Simplified Chinese, and Traditional Chinese.

12. Data & privacy

- **Sources:** DINQ's candidate data comes from public web information and structured talent data — professional profiles, GitHub, papers, projects, personal sites, and other public sources.
 - **Public signals first:** DINQ focuses on public evidence beyond resumes, not self-description.
 - **Your data isn't used to train models:** per the current privacy policy, DINQ does not use user data to train or improve its underlying AI models.
 - **Delete / correct:** you can correct or delete personal data in account settings; after an account is closed, data typically stops displaying within 24 hours and is usually deleted within about 30 days (legal / security exceptions apply).
 - **DINQ Pages are public:** they may be accessed, indexed, cached, or archived by visitors and search engines.
-

13. FAQ

Q: How is DINQ different from LinkedIn? A: LinkedIn is a self-reported professional profile; DINQ focuses on public evidence beyond the resume — GitHub, papers, projects, personal sites, technical-community activity — and brings it into one sourcing workflow.

Q: I already have LinkedIn + GitHub — why DINQ? A: DINQ consolidates information scattered across the web into one discover → evaluate → contact flow, sparing you the manual stitching and tab-switching.

Q: Are search results always accurate? A: No. Treat results as an assistant to expand your search and organize clues — not a final hiring decision.

Q: Do I have to write the search query myself? A: No. DINQ guides you with multiple-choice questions to build a complete, precise brief step by step.

Q: Can I upload a job description? A: Yes — text, PDF, or image.

Q: Why does a search take time? A: Because DINQ runs multi-source deep research rather than simple keyword matching, and shows the process transparently.

Q: Which roles is DINQ best for? A: AI/ML, infrastructure, quantitative finance, security, robotics, biotech, research, and medical technology — roles where a resume alone can't capture real ability.

Q: Is a DINQ Page public? A: Yes. It may be accessed, indexed, cached, or archived by users, visitors, search engines, and third parties.

14. Tips for getting the most out of DINQ

1. **The more specific the brief, the better the results.** Be clear on direction, location, level, and key skills.
2. **Lean on progressive refinement.** Not sure how to phrase it? Let DINQ ask you step by step.
3. **Start small, then scale.** Try "around 5" first; once the direction is right, raise the count.
4. **Verify via sources.** Every activity item links to its source — use it as your due-diligence starting point.
5. **One folder per role.** Let your shortlist grow with your sourcing.
6. **Reach out with context.** Reference what a candidate actually did — it transforms reply rates.
7. **Mind your credits.** Deep searches and revealing contacts consume credits; usage is always visible.

DINQ · Find with precision. Make real work visible to AI.