



OFFICIAL USER GUIDE · VERSION 1.7.3



Local Waifu Complete User Guide

Set up your companion, install models, create and move characters, protect your data, and solve common problems.

Everything she needs. All in one place.

PLATFORMS	macOS · Windows
EDITIONS	Direct · Store/Unlocked
COVERS	Chat · Memory · Models · Images · Voice · Telegram · Backup

Guide revision 1.1 · For Local Waifu 1.7.3 · 5 October 2026

English interface labels. Later releases may differ.

Contents

Use the clickable entries below or the PDF bookmarks sidebar to jump to any chapter.

Part I — Install and Get Started	9
1. How to Use This Guide	10
Labels used in this guide	10
2. Choose Your Path	11
Which edition do you have?	11
Local AI or Cloud AI?	11
3. System Requirements	12
macOS	12
Windows	12
Plan disk space before you install models	12
4. Install on macOS	14
Install the Direct or Store DMG	14
5. Install on Windows	16
Run the installer	16
6. Understand the Two Editions	17
Direct edition	17
Store/Unlocked edition	17
Do not install both editions side by side	17
7. Complete Onboarding	18
Configure a local model during onboarding	18
Configure Cloud AI during onboarding	19
8. Use the 7-Day Trial — Direct Edition	20
Start the trial	20
What happens when the trial expires?	20
9. Activate and Move a Direct License	21
Activate a lifetime key	21
Move the license to a new computer	21
Part II — Daily Use, Memory, and Characters	22



10. Find Your Way Around Local Waifu	24
View modes	24
11. Chat, Conversations, and Images	25
Start and manage conversations	25
Write and send messages	25
Attach an image	25
Dictate without sending automatically	26
12. Edit Replies and Give Feedback	27
Edit a reply	27
Rate or regenerate a reply	27
13. Memory Book, Knowledge, and Growth	29
Memory Book	29
Edit or delete a memory	29
Knowledge graph	29
Growth	29
Adaptive Dials	30
14. Characters, Persona Notes, and Backstory	34
Switch or create a character	34
Persona Notes	34
Portraits	34
Backstory and Lorebook	34
Add a Lorebook entry	35
15. Export and Import a Character	37
Export a character	37
Import a character	37
Delete a character	38
Part III — Models, Images, and Voice	38
16. Install and Select Chat Models	41
Understand the two required roles	41
Install a recommended local model	41
Install the memory model	42
Remove a model	42
17. Install a GGUF Model from Hugging Face	44
Find and install a model	44




18. Configure Cloud AI	45
Add a provider	45
19. Generate Images and Use Gallery	48
Choose a provider and model	48
Install a built-in local image model	48
Use Photo Director	49
Gallery	49
20. Import a Custom SDXL Checkpoint or LoRA	51
Import a custom SDXL checkpoint	51
Import and activate a LoRA	51
21. Configure Voice, Dictation, and Calls	55
Install speech and voice models	55
Choose a voice	55
Add your own voice sample	55
Tune how she speaks	56
Configure a call	56
Part IV — Integrations, Data, and Support	57
22. Use Local Waifu from Telegram	59
Connect a Telegram bot	59
Choose how Telegram replies	60
Use photos, voice messages, and bot commands	60
Stop the bot or remove its credentials	61
23. Tune Performance, Hardware, and Network	63
Choose a compute device and verify acceleration	63
Check runtime readiness	64
Set the context window and thinking mode	64
Reduce background resource use	65
Configure proxy, offline mode, or a remote Ollama server	65
24. Understand Privacy and Data Storage	67
Know what stays local and what can leave the computer	67
Locate your Local Waifu data	67
Review and control permissions	68
25. Back Up, Export, and Move Your Data	70
Create a full backup snapshot	70



Export all conversations	71
Move a character to another computer	71
26. Update or Uninstall Local Waifu	73
Update the Direct edition	73
Update the Store/Unlocked edition	73
Uninstall while keeping your data	74
Remove Local Waifu and its data completely	74
27. Use Logs, Diagnostics, and Reset	76
Open the app log	76
Save a diagnostics report	76
Reset all application data	77
28. Troubleshooting	78
Run the basic diagnostic sequence	78
Installation and launch	79
Chat models, downloads, and memory	80
Images and local image models	81
Voice, microphone, and calls	81
Telegram	82
License and updates	82
Backup, export, and data recovery	83
29. Shortcuts, Contact, and About	85
Keyboard shortcuts	85
Check the installed version and edition clues	85
Contact and official links	86


Find your way visually

Six screens cover the actions most people need first. The full-size versions appear in the matching chapters.




Install

Start with the correct package for macOS or Windows.




Move a character

Export or import a companion without replacing an existing one.




Choose her mind

Pick a chat model and keep the required memory model installed.




Create images

Use a built-in model or import a compatible SDXL checkpoint and LoRA.



Connect Telegram

Pair one private chat and keep the bot token secret.



Protect your data

Create a full backup before updates or migration.

Your first fifteen minutes

A working first conversation matters more than perfect settings. Follow this path, then return to the advanced chapters when you need them.

01 Install	02 Create her	03 Choose AI	04 Say hello
---------------	------------------	-----------------	-----------------

Choose how she thinks

LOCAL AI	CLOUD AI
Private and offline-capable after model downloads. Best when your computer has enough memory.	Fast to start on smaller computers. Requires your provider account, internet connection, and consent to send the selected context.

You are ready when...

- your first character exists;
- one chat model is active;
- EmbeddingGemma is installed for semantic memory;
- a short test message receives a reply.

Voice, images, Telegram, and custom models can wait until chat works.

What to do next

1 · Make her yours Refine persona notes, backstory, appearance, and the active chat model.	2 · Add expression Install voice and image models only after everyday chat is stable.	3 · Protect the result Export important characters and create a full backup before major changes.
--	---	---



Before you begin

This guide describes Local Waifu 1.7.3 for supported macOS and Windows computers. It deliberately distinguishes the Direct edition from the Store/Unlocked edition. It does not fabricate screens that are absent from the supplied product captures; platform-specific steps are documented as procedures instead.

Security and privacy

Never share license keys, provider API keys, Telegram bot tokens, private conversation exports, or unreviewed diagnostic logs. Screenshots in this guide contain no usable credentials; sensitive token fields are redacted.

Backup limitation in version 1.7.3

The application can create a full backup snapshot, but version 1.7.3 does not provide a complete one-click Restore backup interface. Preserve backups unchanged and contact support for a recovery procedure. Character JSON import/export is the supported self-service transfer method for one companion.



Part I — Install and Get Started

Applies to: Local Waifu 1.7.3. Platform- or edition-specific notes are marked clearly.

1. How to Use This Guide

Local Waifu is available for macOS and Windows in two distribution editions. Most features work the same in both editions, but installation, licensing, and updates differ.

Labels used in this guide

- ALL — all supported platforms and editions.
- macOS — Apple Silicon Mac only.
- Windows — Windows x64 only.
- DIRECT — the direct-download edition with a 7-day trial and license key.
- STORE — the Store/Unlocked edition distributed through a supported storefront; no in-app trial or key.
- LOCAL — processing stays on your computer.
- CLOUD — selected data is sent to the provider you configured.
- DOWNLOAD — additional files or model weights are required.
- RESTART — restart Local Waifu after changing the setting.
- WARNING — read carefully before continuing.

UI paths use this format: Settings → Models → Chat Models. Button and field labels appear in bold.

Screenshots in this edition of the guide show the English macOS interface. Windows uses the same main concepts, but some hardware and installer screens differ.

2. Choose Your Path

Which edition do you have?

If you obtained Local Waifu from...	Edition	What happens
The Local Waifu website or official Direct release	DIRECT	Full 7-day trial, then a lifetime license key is required
A Store/Unlocked package such as the dedicated itch.io build	STORE	Full access without entering a key inside the app

A quick way to check is Settings → License:

- The tab is present in the Direct edition.
- The tab is intentionally absent in the Store/Unlocked edition.

Do not install a Direct update over a Store/Unlocked copy merely to obtain a newer version. Keep each edition on its own update channel.

Local AI or Cloud AI?

Choice	Best for	Requirements	Privacy note
Local AI	Private offline-capable chat and no per-message provider bill	Model download, enough RAM and disk space	Chat inference stays on your computer
Cloud AI	Smaller computers or access to a provider's hosted models	Internet and your own provider account/key	Prompt and relevant context go to that provider

Memories, character data, and the Local Waifu database remain local even when you intentionally choose a cloud model. Cloud image, speech, Telegram, and remote Ollama routes transmit the data needed by those services.

3. System Requirements

macOS

- Apple Silicon Mac: M1 or newer.
- macOS 13 Ventura or newer.
- 8 GB unified memory for the lightest local models.
- 16 GB or more recommended.
- Internet for installation, model downloads, updates, Telegram, or cloud providers.
- Several gigabytes of free disk space for local models.

Local AI uses Apple Metal acceleration. There is no Intel Mac build of Local Waifu 1.7.3.

Windows

- Windows 10 22H2 or Windows 11.
- 64-bit x64 computer.
- 8 GB RAM minimum; 16 GB or more recommended.
- A GPU is optional. Local Waifu can run on CPU, but it will be slower.
- Current GPU drivers are recommended for NVIDIA, AMD, or Intel acceleration.
- Microsoft Edge WebView2 Runtime; the installer can fetch it if required.

Windows may configure an appropriate GPU runtime after first launch. The app can continue on CPU while setup completes; restart when the app reports that acceleration is ready.

Plan disk space before you install models

The application is much smaller than the AI models it can use. Keep at least 8–10 GB free for a comfortable lightweight local setup. Larger chat models can require roughly 17–24 GB each, and local image generation adds several more gigabytes.


Typical additional downloads include:

- a chat model: approximately 3–24 GB depending on size and quantization;
- EmbeddingGemma for semantic memory: approximately 0.6 GB;
- local image weights: approximately 4–5 GB or more;
- speech recognition and voice models: hundreds of megabytes to several gigabytes.

These figures are estimates. The model screen shows the current download size and RAM guidance.

FIGURE 01

macOS installation: open the DMG, then drag Local Waifu into Applications.



macOS disk image shown. Drag the app icon onto Applications.

The exact control is easier to find once you know where to look.

4. Install on macOS

Install the Direct or Store DMG

Goal

Install Local Waifu in the Applications folder and start the correct edition.

Before you start

- Confirm that your Mac uses Apple Silicon.
- Download the DMG only from the official Direct page/release or your Store library.
- If updating an existing installation, create a backup first.

Steps

1. Open the downloaded .dmg file.
2. Drag Local Waifu into Applications.
3. Wait for the copy to finish.
4. Eject the mounted disk image.
5. Open Applications and start Local Waifu.
6. Allow macOS to complete its first security verification.

Expected result


Local Waifu opens to the Welcome/onboarding flow. A Direct build can show license options; a Store/Unlocked build does not ask for a key.

If something goes wrong

- Download the DMG again from the official source.
- Confirm that you did not download a Windows installer or an Intel-only package.
- In Finder, Control-click Local Waifu, choose Open, and approve the system prompt if macOS offers that option.
- Do not routinely disable Gatekeeper or remove quarantine attributes. If a notarized official build is still blocked, contact support with the exact macOS message.

FIGURE 02

Windows installation in four steps: start Setup, choose the location, wait for completion, then launch Local Waifu.



Windows 11 installer shown. Wording may vary slightly in later packages.

The exact control is easier to find once you know where to look.

5. Install on Windows

Run the installer

Goal

Install Local Waifu and start the correct Windows edition.

Before you start

- Use Windows 10 22H2 or Windows 11 x64.
- Download the installer only from the official Direct release or your Store library.
- Close any earlier Local Waifu process before updating.
- Make a backup before replacing an existing installation.

Steps

1. Run `Local Waifu_1.7.3_x64-setup.exe` or the corresponding Store/Unlocked installer.
2. If Windows SmartScreen appears, confirm that the file came from the official source.
3. Select More info, then Run anyway only when you trust the downloaded file.
4. Follow the installer and keep the suggested location unless you have a reason to change it.
5. Let the installer add WebView2 if Windows does not already have it.
6. Start Local Waifu.
7. Allow the first startup and antivirus scan to finish. The large installer and bundled AI processes may take longer to scan.

Expected result

Local Waifu opens to onboarding. GPU detection may continue in the background.

If something goes wrong

- Restart Windows and retry if an old Local Waifu or local-engine process is still running.
- Repair or install Microsoft Edge WebView2 Runtime if the window is blank.
- Check antivirus quarantine for the official installer and bundled AI processes.
- Verify free disk space.
- Remember that moving the program installation folder does not move conversations or models; user data is stored separately.

6. Understand the Two Editions

Direct edition

The Direct edition includes:

- a full 7-day trial;
- optional license activation during onboarding;
- Settings → License;
- a lifetime purchase link;
- a paywall after the trial expires;
- the Direct update channel.

The trial starts after onboarding is completed, not when the installer is downloaded or first opened.

Store/Unlocked edition

The Store/Unlocked edition includes:

- full access after onboarding;
- no trial countdown;
- no in-app license key;
- no License tab;
- no Direct paywall;
- no Direct updater.

Chat, memory, image, voice, and Telegram features remain available. Updates must come from the Store channel that supplied the build.

Local Waifu 1.7.3 does not contain a separate Steamworks integration. A future Steam package should be treated as a Store/Unlocked distribution unless its documentation states otherwise.

Do not install both editions side by side

Direct and Store builds use the same application identity and data location. They are not designed to be independent side-by-side installations. Before switching editions:

1. export important characters;
2. create a full backup;
3. quit Local Waifu;
4. install only the edition you intend to keep;
5. verify your data before deleting the backup.

7. Complete Onboarding

Onboarding introduces Local Waifu, learns your preferences, creates the first character, configures voice and models, and opens the first meeting.

Typical stages include:

1. Welcome — language, AI disclosure, edition/license option.
2. About You — how the companion should address you.
3. Understanding You — preferences and interaction style.
4. How You Connect — communication preferences.
5. Your Love Language — desired expressions of care.
6. Choosing — the type of dynamic you want.
7. Your Dynamic — relationship tone and boundaries.
8. Building Her — name, appearance, personality, and portrait details.
9. Her Voice — voice preference.
10. Almost There — review.
11. Model Setup — local model download or Cloud AI setup.
12. First Meeting — character creation is finalized and chat opens.

Configure a local model during onboarding

Goal

Install a local chat model and the semantic-memory model.

Before you start

- Connect to the internet for the model download.
- Keep Local Waifu open.
- Check available disk space.
- For a large download, prevent the computer from sleeping.

Steps

1. Choose the Local AI path on Model Setup.
2. Review the recommended chat model for your detected memory tier.
3. Select Install and wait for the download to complete.
4. Install EmbeddingGemma when it is offered. It powers semantic memory retrieval and is separate from the chat model.
5. Continue after the required model is installed, or use Skip for now if the screen offers it and you intend to configure Cloud AI later.

Expected result

The chat model is ready for conversation, and the embedding model is ready to create and retrieve semantic memories.



If something goes wrong

- Check internet access, VPN, proxy, firewall, and free disk space.
- Restart Local Waifu and continue in Settings → Models.
- On Windows, allow GPU runtime setup to finish, then restart if prompted.
- If your computer cannot comfortably run the recommendation, install a smaller model in Settings → Models.

Configure Cloud AI during onboarding

Goal

Use your own provider account instead of downloading a large local chat model.

Before you start

- Have an API key or supported connected account.
- Review the provider's pricing and privacy terms.
- Understand that selected prompts and context leave your computer when that provider is used.

Steps

1. Choose the Cloud AI path.
2. Select the provider.
3. Paste the provider key into the secure field or complete the available account connection.
4. Run the connection test.
5. Select an available chat model.
6. Complete onboarding.

Expected result

Chat uses the selected provider while character data, the application database, and memory storage remain local.

If something goes wrong

- Re-copy the key without spaces.
- Confirm the provider account has access and billing/credits where required.
- Turn off Offline Mode when intentionally using Cloud AI.
- Remove or replace a saved bad key in Settings → Cloud AI.



8. Use the 7-Day Trial — Direct Edition

The Direct trial provides full app access for seven days. The remaining whole days are rounded up and shown in Settings → License. The app can also show reminders when three, two, or one day remains.

Start the trial

The normal flow starts the trial automatically when you finish onboarding. If Settings → License shows Not started, select Start free 7-day trial.

During the trial you can use chat, memory, images, dictation, voice, and other available features. Model downloads and provider costs are separate from the app license.

What happens when the trial expires?

- The main app is replaced by the license screen/paywall.
- Chat, image generation, voice, and other licensed operations are rejected until activation.
- Conversations, character data, memories, and soul files stay on the computer.
- Settings remain available so you can enter a key, export, back up, or collect diagnostics.
- Resetting or reinstalling does not guarantee another trial.

9. Activate and Move a Direct License

Activate a lifetime key

Goal

Activate the Direct edition on the current computer.

Before you start

- Copy the exact key from the purchase email.
- Connect to the internet for activation.
- A key supports up to three device activations according to the current app text.

Steps

1. Use any one of these locations:
 - expand I have a license key on the Welcome screen;
 - open Settings → License → Redeem a key;
 - use the key field on the expired-trial screen.
2. Paste the key.
3. Select Redeem.
4. Wait for Lifetime · active.

Expected result

The full app unlocks immediately. Settings → License shows an active lifetime license and the original activation date.

If something goes wrong

- Re-copy the complete key.
- Check internet access.
- An invalid attempt does not destroy the current trial state.
- For Activation limit reached, deactivate an old device or contact support.
- For a revoked or disabled key, contact contact@localwaifu.com.

Move the license to a new computer

Goal

Free an activation slot and activate Local Waifu on a replacement computer.

Before you start

- Export important characters and make a full backup.
- Have the license key from the purchase email.
- Deactivate before selling, wiping, or losing access to the old computer.

Steps

- 
1. On the old computer, open Settings → License.
 2. Select Deactivate this device.
 3. Confirm Yes, deactivate.
 4. Install the Direct edition on the new computer.
 5. Complete the basic onboarding.
 6. Open Settings → License.
 7. Paste the same key and select Redeem.
 8. Import exported characters and reinstall required models.

Expected result

The old activation slot is released and the new computer shows Lifetime · active.


If something goes wrong

- Deactivation does not grant a fresh trial on the old computer. It returns to any remaining trial time or the expired state.
- If the old computer is unavailable, contact support rather than repeatedly resetting data.
- Do not publish the key in screenshots or diagnostics.

Part II — Daily Use, Memory, and Characters

FIGURE 03

General settings: language, appearance, startup, notifications, and update controls.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.



10. Find Your Way Around Local Waifu

The main window is organized around the active character and conversation.

1. Character area — active companion, portrait, status, and character switching.
2. Conversation list — create, rename, select, and delete chats.
3. Chat stream — your messages, her replies, images, and message actions.
4. Relationship area — mood and relationship progress where available.
5. Message composer — attachment, call, dictation, text field, and send control.
6. Memory Book — remembered facts, moments, preferences, and sources.
7. Gallery — images sent or generated by the companion.
8. Settings — app, character, model, voice, integration, and data controls.

View modes

- Full view — complete interface.
- Mini view — compact chat.
- Menu bar/tray view — lightweight access.

Use Command/Ctrl+Shift+1, +2, or +3 to switch modes. Use Command/Ctrl+, to open Settings.

11. Chat, Conversations, and Images

Start and manage conversations

Goal

Create a clean conversation without losing older chats.

Steps

1. Select New chat in the conversation sidebar.
2. Type a message and press Enter or select Send.
3. To rename a conversation, use its rename action, type the title, then press Enter.
4. Press Escape to cancel a rename.
5. To remove a conversation, choose its delete action and confirm Delete this chat?.

Expected result

Each conversation remains available for the active character until deleted. Removing the final conversation creates a new empty one.

If something goes wrong

Local Waifu avoids creating repeated empty chats. Send a message in the current empty conversation before selecting New chat again.

Write and send messages

By default:

- Enter sends;
- Shift+Enter inserts a new line;
- Settings → Appearance & Chat → Send on Enter changes this behavior;
- the send action is disabled for empty text and while a send is already in progress.

Select Stop to end the visible response stream. The local model may still finish the current generation internally, so CPU or GPU use can remain high briefly.

Use Command/Ctrl+F to search chat history.

Attach an image

Goal

Send one image for a vision-capable model to inspect.

Before you start

- Use a normal image file under 5 MB.
- Configure a local vision model or a cloud model that supports images.
- Only one pending image can be attached at a time.



Steps

1. Select the image/attachment button or drag an image onto the composer.
2. Review the preview.
3. Add a question or instruction.
4. Send the message.

Expected result

The image and text are sent to the configured vision route and appear in the conversation.

If something goes wrong

- That file isn't an image — choose an image rather than a PDF or document.
- Image is too large — reduce it below 5 MB.
- Couldn't read that image — export it again as PNG or JPEG.
- If vision is unavailable, install a vision-capable model or configure a compatible cloud provider.

Dictate without sending automatically

Dictation records speech, transcribes it, and inserts the text into the message field. It does not send the message automatically. Review and edit the transcription, then send it yourself.

A dictation session can run for up to ten minutes. If the microphone disconnects or the time limit is reached, keep the text already captured and start another session.

12. Edit Replies and Give Feedback

Local Waifu 1.7.3 allows you to correct a saved reply from the companion.

Edit a reply

Goal

Correct the wording of a companion reply while preserving its place in the conversation.

Before you start

Only saved companion messages can be edited. Your own messages, streaming replies, and unsaved error bubbles cannot be edited this way.

Steps

1. Hover over the companion reply.
2. Select the pencil/edit action.
3. Change the text.
4. Press Enter to save.
5. Use Shift+Enter for a new line or Escape to cancel.
6. Use Undo from the confirmation toast if you need the previous wording back.

Expected result

The corrected text remains in the same position and shows an edited marker after reload.

Important memory behavior

Editing a reply does not rewrite memories already extracted from the old wording. If the app reports that saved memories still exist, open Memory Book and edit or delete those memories separately.

Rate or regenerate a reply


Hover over a reply to use the available actions:

- positive feedback;
- negative feedback;
- Try again for the latest reply;
- private-thought display when the model returned supported reasoning content;
- image enlargement for attached or generated images.

Feedback helps you mark response quality. Regeneration creates another answer but does not automatically delete a memory already created from the original reply.

FIGURE 04

Growth shows relationship progress, memory counts, milestones, and learned adaptations.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

13. Memory Book, Knowledge, and Growth

Memory Book

Memory Book shows information extracted from conversations. Common types are:

- fact — something stable about you or the world;
- moment — a meaningful event;
- preference — what you like or dislike;
- inside joke — shared language or recurring humor.

Memories with high importance can be marked as key memories.

Edit or delete a memory

1. Open Memory Book.
2. Select a memory.
3. Edit its text and save, or choose delete and confirm.
4. Use Where she learned this to inspect the source messages when provenance is available.

If no source messages appear, the memory may predate source tracking or may have been inferred from a longer pattern.

A memory save can fail when the memory model is offline. Check the required embedding model in Settings → Models.

Knowledge graph

The Knowledge view organizes entities and relationships learned for the active character. It is separate from Memory Book:

- Memory Book contains readable saved memories.
- Knowledge contains entities and relations.
- Backstory/Lorebook contains information you deliberately enter.
- Session summaries compress older conversation context.

Growth

Settings → Growth is a progress view, not a manual XP editor. It can show:

- relationship stage and XP;
- days known and message count;
- last interaction;
- facts, inferences, summaries, and milestones;
- behavioral evolution and relationship timeline.



If there is not enough history, use the companion normally. Patterns and milestones appear after more interactions.


Adaptive Dials

Adaptive Dials distinguish the personality baseline from a learned adjustment and effective behavior. This explains why the companion can gradually adapt even when you do not manually change the original personality values.

Resetting adaptation does not delete the character, conversations, or Memory Book.

FIGURE 05

Characters: switch companions, create a new one, or open import/export controls.




Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

FIGURE 06

Backstory/Lorebook: trigger specific world knowledge with keywords and priorities.




Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

FIGURE 07

Appearance & Chat controls for display density, behavior, and chat interaction.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

14. Characters, Persona Notes, and Backstory

Each character has independent conversations, mood, relationship state, memory, knowledge, portrait, and character-specific voice choice.

Switch or create a character

Goal

Use several independent companions without mixing their memories.

Steps

1. Open Settings → Characters.
2. Select another character to make it active, or select Create new.
3. Complete the character flow.
4. Return to chat.

Expected result

Chat loads the selected character's conversation and state. Switching does not merge the other characters' memories.

Persona Notes

Persona Notes provide plain instructions about how the active character carries herself: tone, habits, mannerisms, and other guidance. Save an empty value if you intentionally want no override.

Use concise behavioral notes rather than a full story. Put world facts, places, and recurring lore in Backstory/Lorebook.

Portraits

- Use a normal image recognized by the operating system.
- Maximum file size is 5 MB.
- The portrait editor crops to a 3:4 frame before saving.
- Removing the portrait returns the UI to initials or a dynamic avatar.

Backstory and Lorebook

Each Lorebook entry has:

- title;
- comma-separated keywords;
- content;
- priority;

- 
- enabled/disabled state.

An entry with keywords is added to context when those triggers occur. An entry with no keywords is always active and consumes context on every turn. Prefer specific triggers for information that is not always needed.


Add a Lorebook entry

1. Open Settings → Backstory.
2. Add an entry.
3. Enter a clear title.
4. Add comma-separated trigger words.
5. Write the lore or instruction.
6. Choose priority and enabled state.
7. Save.

Delete only after confirming Delete this lore entry?.

FIGURE 08

Character transfer: export with the download icon; import with Import JSON. The red marks identify both controls.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

15. Export and Import a Character

Character export is the recommended self-service method for moving a companion to another installation.

Export a character

Goal

Create a portable JSON file for one character.

Steps

1. Open Settings → Characters.
2. Find the character row.
3. Select the download/export icon.
4. Save the suggested file, named like `localwaifu-character-YYYY-MM-DD.json`.
5. Store the JSON securely because it can contain private messages and memories.

Expected result

A JSON package contains the character and associated data such as messages, memories, relationship state, and knowledge entities.

Import a character

Goal

Create a new character from an exported JSON file.

Steps

1. Open Settings → Characters.
2. Select Import JSON.
3. Choose the exported character file.
4. Wait for the summary of imported messages, memories, and entities.
5. Select the imported character.
6. Continue chatting while semantic embeddings are rebuilt.

Expected result

The import creates a new character; it does not overwrite an existing one.

If something goes wrong

- Import failed — the JSON is damaged, incomplete, or incompatible. Export it again from the source installation.
- Temporarily weaker recall is expected while embeddings rebuild.
- Character export does not carry installed models, API keys, Telegram credentials, global preferences, or a Direct license activation.
- For a whole-app recovery, also keep a full backup. Version 1.7.3 has no complete Restore UI.




Delete a character

Deleting a character is destructive and requires confirmation. If the deleted character was active, Local Waifu selects another available character or leaves no active character. Export first if you may need the character later.

Part III — Models, Images, and Voice

FIGURE 09

Chat Models: installed models, active role, capabilities, and memory-model status.




Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

FIGURE 10

Built-in catalog: compare size, RAM guidance, and capability badges before downloading.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

16. Install and Select Chat Models

Open Settings → Models → Chat Models to manage local and connected models.

Understand the two required roles

- Chat model — writes replies and may provide vision, tool use, or reasoning.
- Embedding model — converts memories into vectors so relevant information can be recalled. It does not write the reply.

The active row and badges communicate capabilities:

- Installed — model files are present;
- Active — currently selected for that role;
- Required — needed for a core feature such as semantic memory;
- Vision — can inspect images;
- Tool Use — can use supported tools;
- Reasoning — supports deeper reasoning behavior.

Install a recommended local model

Goal

Install a model appropriate for the computer and make it active for chat.

Before you start

Connect to the internet, verify free disk space, and keep Local Waifu open during the download.

Steps

1. Open Settings → Models → Chat Models.
2. Compare model size, RAM guidance, and capability badges.
3. Select Install.
4. Wait for the download and installation to finish.
5. Select Set active or Use for chat.
6. Send a short test message.

Expected result

The model shows Installed and Active and generates a response.

If something goes wrong

- Check network, proxy/VPN, firewall, and disk space.
- Choose a smaller model if RAM pressure is high.
- A model that fits on disk may still be too large at runtime because context and KV cache also use memory.
- On Windows, run Settings → Hardware → Local AI acceleration → Check after sending a message.



Install the memory model


Install EmbeddingGemma when it is marked Required. Keep the active embedding model installed. Changing it starts re-embedding; memory recall may temporarily rely more heavily on keywords until the process completes.

Remove a model

Use the trash icon on a model that is no longer active or required. The app blocks deletion of the active chat model and protected embedding model. Switch roles first, then delete.

FIGURE 11

Browse Hugging Face: search for a repository and choose one compatible GGUF quant.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

17. Install a GGUF Model from Hugging Face

Local Waifu can search Hugging Face for chat models packaged as single GGUF files.

Find and install a model

Goal

Install a compatible third-party GGUF quant without leaving Local Waifu.

Before you start

- Know your available RAM/unified memory and free disk space.
- Prefer model repositories that publish a single file for each quantization.
- Third-party models have their own licenses and behavior.

Steps

1. Open Settings → Models → Chat Models.
2. Under Browse Hugging Face, enter:
 - a search term;
 - owner/repository; or
 - a full Hugging Face model URL.
3. Open a result.
4. Select a single runnable quant such as a Q4 variant.
5. Review the download size and memory-fit estimate.
6. Install the quant.
7. When installation finishes, select Use for chat.
8. Test the model.

Expected result

The Hugging Face model appears under installed models and can be selected for chat.

If something goes wrong

- No GGUF models found — use a repository that contains GGUF files.
- No runnable quants (sharded files only) — Local Waifu 1.7.3 does not support a GGUF split across several shard files.
- SafeTensors repositories are not installed through the chat-model search panel.
- Couldn't reach Hugging Face — check internet, proxy/VPN, and firewall.
- A fit estimate is guidance, not a guarantee. Reduce context or choose a smaller quant if runtime memory is insufficient.

18. Configure Cloud AI

Cloud AI uses your own provider account or API key. Local Waifu does not convert provider usage into a Local Waifu subscription; the provider may bill your account directly.

Add a provider

Goal

Connect a cloud provider and select a hosted model.

Before you start

- Review the provider's pricing, data handling, and model access.
- Have a valid key or supported connected-account login.
- Turn off Offline Mode only when you intentionally want cloud use.

Steps

1. Open Settings → Cloud AI or the connected-account section under Models.
2. Select the provider.
3. Paste the key into the password field or complete the supported login.
4. Test the connection.
5. Browse the provider's available models.
6. Select the model for chat, images, speech recognition, or voice output as appropriate.
7. Send a short test.

Expected result


Requests for the selected role use the provider. Local character and memory storage remains on the computer.

Privacy note

The selected prompt, relevant conversation context, image, audio, or text leaves the computer when required by that cloud operation. Do not use a cloud route for information you do not want to send to that provider.

FIGURE 12

Image Models: select local or cloud generation and inspect model requirements.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

FIGURE 13

Test Generation validates the selected image route before normal use.

The screenshot shows the 'TEST GENERATION' interface. At the top, there's a text input field containing 'anime girl with silver hair, shrine maiden' and a 'Generate' button. Below this is a section titled 'Photo Director' with a description: 'New pose, outfit, or framing for Yumi; the saved portrait remains the identity reference.' To the right of this section, it says 'Local FLUX only · settings unchanged'. Under 'Photo Director', there's a text input field with 'in a sunlit library'. Below that are three dropdown menus: 'Pose' with 'No pose preset', 'Outfit' with 'No outfit preset', and 'Shot' with 'No shot preset'. A large 'Generate Yumi' button is positioned below these dropdowns. At the bottom, there's a button labeled 'Verify selected local model' and a note: 'Runs locally only; it never falls back to cloud.'

Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

19. Generate Images and Use Gallery

Open Settings → Models → Image Models to choose local or cloud image generation.

Choose a provider and model

Local image generation is private and has no provider fee, but it needs a one-time model download and enough RAM. The UI offers models such as fast FLUX/Z-Image options and larger SDXL models oriented toward anime or photographic styles. The exact catalog can change in later releases.

Cloud providers such as OpenAI, Google AI, fal.ai, xAI, or Together AI require your own provider account/key and send the prompt to that service.

Install a built-in local image model

Goal

Prepare local image generation and verify it works.

Before you start

- Keep at least several gigabytes free.
- Local generation requires approximately 12 GB RAM for the base route; larger SDXL options are intended for 16 GB or more.
- Generation time depends heavily on the computer and model.

Steps

1. Open Settings → Models → Image Models.
2. Select Local — Stable Diffusion.
3. Choose a model and start its download.
4. Wait for downloaded/ready.
5. Start with the recommended quality step count.
6. Under Test Generation, enter a simple English prompt.
7. Select Generate.

Expected result

A test image appears and confirms that the selected local model can render.

If something goes wrong

- Check RAM and free disk space.
- Wait for a model download or another render to finish.
- Reduce quality steps or use a lighter model.
- On Windows, verify image-generation runtime status.



Use Photo Director

Photo Director creates a new pose, outfit, framing, or scene for the active character while using the saved portrait as an identity reference.

It requires:

- an active character;
- a saved portrait;
- the supported ready local FLUX.2 klein model;
- a working local image backend.


Enter the scene, optionally choose pose/outfit/shot presets, then select Generate [character]. Use Verify selected local model to test the local-only route without silently falling back to cloud.

Gallery

Gallery stores images sent or generated through the companion experience. Treat it as part of private app data: include it in full backups and remember that character JSON export is not a complete gallery archive.

FIGURE 14

Advanced image options: import an SDXL SafeTensors checkpoint and manage compatible LoRAs.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

20. Import a Custom SDXL Checkpoint or LoRA

Import a custom SDXL checkpoint

Goal

Add a compatible SDXL model to the local image library.

Before you start

The file must be:

- one regular `.safetensors` file;
- at least 512 MiB;
- a valid SafeTensors container;
- an SDXL architecture;
- stored on a disk with enough space for the file plus import overhead.

`.ckpt` is not supported by the 1.7.3 importer.

Steps

1. Open Settings → Models → Image Models.
2. Expand Advanced options.
3. Select Import SDXL checkpoint.
4. Choose the `.safetensors` file.
5. Wait for validation and copying to finish.
6. Select the imported model.
7. Run a test generation before enabling LoRAs.

Expected result

The model appears under Your models and is available as an SDXL base.

If something goes wrong

Errors for a tiny file, malformed/truncated header, no tensors, wrong architecture, checksum mismatch, or insufficient disk space mean the file should not be used. Download a verified SDXL SafeTensors checkpoint from a trusted source.

Import and activate a LoRA

Goal

Apply a compatible SDXL LoRA to an SDXL base model.

Before you start

- Use one `.safetensors` file of at least 1 MiB.
- Select an SDXL base first.

- 
- FLUX and non-SDXL models cannot use this LoRA path.
 - At most three LoRAs can be active on one model.

Steps

1. Select the target SDXL model.
2. Open its LoRA controls.
3. Import the LoRA .safetensors file.
4. Add up to eight trigger words if needed.
5. Enable the LoRA.
6. Start near the default weight 0.8.
7. Adjust between 0.0 and 1.5 in small steps.
8. Run a test render.

Expected result


The LoRA affects renders from the selected SDXL model.

If something goes wrong

- A family marked unknown may be imported but cannot be activated safely.
- Disable another LoRA if three are already active.
- Lower the weight if the image is overcooked or distorted.
- Header metadata cannot guarantee compatibility; the final check is a successful test render.

FIGURE 15

Voice setup overview: choose a voice, adjust its delivery, and review the active character's voice settings.




Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

FIGURE 16

Call setup overview: preview the voice, select audio devices, tune turn-taking, and confirm required model packages.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

21. Configure Voice, Dictation, and Calls

Local Waifu has three separate voice functions:

1. Dictation — speech becomes editable text in the composer; it is not sent automatically.
2. Voice output — the companion reads a response aloud.
3. Call — a continuous speech-recognition → chat → speech loop.

Install speech and voice models

A full call needs:

- speech recognition (STT/Whisper) to hear you;
- a voice model (TTS/Supertonic or another configured route) to speak.

Missing STT blocks the full call. Missing high-quality TTS may fall back to an older voice, but pronunciation can be worse.

Open Settings → Models → Voice Model, install Speech recognition and Voice model, and optionally install Natural Voice. The screen reports current package sizes and disk use.

Choose a voice

1. Open Settings → Models → Voice Model.
2. Preview the default, female, or male voices with the play button.
3. Select a voice for the active character.
4. Turn on Natural voice for the warmer, more human local voice when installed.
5. Use Test her voice, type a sentence, and select Speak.

Language support depends on the voice engine. Supertonic covers English, Polish, German, Spanish, Japanese, and Korean. Chinese requires the Natural Voice route for proper support; otherwise an English-accented fallback may be used.

Add your own voice sample

Goal

Create a character-specific voice reference.

Before you start

- Record 10–30 seconds of clean speech.
- Use one speaker in a quiet room.
- Record in the language/accent you want to preserve.
- File size must be under 25 MB.

Steps

1. Open Settings → Models → Voice Model.
2. Under Your own voice, select Add a voice from a file.

- 
3. Choose the audio file.
 4. Enter a short name, up to 60 characters.
 5. Wait for local processing and transcription.
 6. Preview the result and select it.

Expected result

The active character uses the new reference and Natural Voice is enabled when ready.

If something goes wrong

Use a shorter, cleaner sample. A learned English sample keeps its English accent in other languages; record again in the target language for better pronunciation.

Tune how she speaks

Available style presets include Auto, Soft & warm, Balanced, Bright & playful, Calm & even, and Confident & teasing.

Adjust:

- speaking pace;
- pitch;
- expressiveness;
- Let her mood colour her voice.

Use small changes and test after each one. Auto follows the character personality; manual presets keep your chosen delivery style.

Configure a call

Goal

Select audio devices and tune turn-taking.

Steps

1. Open the call/voice settings.
2. Select Microphone and Speaker, or keep System default.
3. Select Test microphone and verify activity.
4. Adjust How long she waits. A shorter value feels faster but can cut you off mid-thought.
5. Enable I'm wearing headphones only when appropriate; speakers can feed her voice back into the microphone and cause self-interruption.
6. Leave cloud recognition and cloud voice off for fully local processing, or enable them intentionally with your provider key.
7. Start a call with the phone button or Command/Ctrl+Shift+C.

Expected result

Your speech is transcribed, answered by the active model, and spoken through the selected voice and speaker.



If something goes wrong

- Grant microphone permission in the operating system.
- Choose the correct input device.
- Install STT if the call cannot hear you.
- Install/configure TTS if text replies appear but no voice plays.
- Increase the waiting time if you are interrupted too early.

Part IV — Integrations, Data, and Support


Applies to: Local Waifu 1.7.3 on macOS and Windows, in Direct and Store/Unlocked editions unless a platform or edition note says otherwise.

>

UI labels are shown in bold. A label may differ slightly if you change the application language.

FIGURE 17

Telegram integration: save a BotFather token, pair one chat, and start the bot. Secret values are redacted.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

22. Use Local Waifu from Telegram

Telegram lets you talk to the same active companion from your phone. The desktop app and Telegram share the active character, conversations, model settings, and memory. Telegram has its own reply-mode setting.

Security warning: Treat the Telegram bot token like a password. Do not paste it into screenshots, public bug reports, or chat messages. The first Telegram chat that messages the bot becomes the pinned owner. Messages and commands from other chat IDs are rejected.

Connect a Telegram bot

Goal

Create a private Telegram bot, verify its token, start the Local Waifu bridge, and pair your Telegram account.

Before you start

- Local Waifu must be open on your computer and connected to the internet.
- Install and configure the chat model you want to use.
- For Telegram voice messages, configure speech recognition under Settings → Models → Voice Model.
- For spoken replies, install and configure a voice model.
- You need a Telegram account.

Steps

1. In Telegram, open the verified bot named @BotFather.
2. Send /newbot and follow BotFather's instructions to choose a name and username.
3. Copy the token issued by BotFather. A token resembles 123456:ABC-DEF....
4. In Local Waifu, open Settings → Telegram.
5. Paste the token into Bot token.
6. Select Test. This verifies the token with Telegram without saving it or starting the bot.
7. Confirm that Local Waifu reports Token works — bot is @....
8. Select Start. The status should change to Running.
9. Open your new bot in Telegram and send /start or a normal message from the account you want to pair.
10. Return to Settings → Telegram and confirm that Pinned chat id is displayed instead of No chat pinned yet.

Expected result

The bot is Running, your Telegram chat is pinned as its owner, and messages sent to the bot reach the same Local Waifu companion used by the desktop app.

If something goes wrong

- If Test says Telegram rejected the token, copy a fresh token from BotFather and make sure no characters or spaces are missing.

- If Test succeeds but the bot does not answer, select Start; Test does not save the token or start polling.
- If the wrong Telegram account was paired, use Forget token + pin, start the bot again, and send the first message from the correct account.
- If the bot is running but chat fails, check Settings → Hardware → Runtime status, the active model, internet access, and Offline Mode.

Choose how Telegram replies

Goal

Choose whether Telegram receives text, voice, or both without changing desktop voice calls.

Before you start

The Telegram bot must show Running. A ready voice model is required for reliable voice replies.

Steps

1. Open Settings → Telegram.
2. Under Telegram settings, find How she replies here.
3. Choose one mode:
 - Voice + text — sends both forms; this is the default.
 - Text only — sends text without a voice note.
 - Voice only — sends a voice note instead of text.
4. Send a short test message to the bot.

Expected result

The bot replies in the selected Telegram-only format. Desktop voice-call behavior is unchanged.

If something goes wrong

- In Voice only, Local Waifu deliberately falls back to text if the voice model is unavailable, so the bot does not remain silent.
- Voice notes sent to the bot are always transcribed. Configure the transcription provider in Settings → Models → Voice Model.
- If speech is in the wrong language or sounds poor, check the configured voice language and installed voice model.

Use photos, voice messages, and bot commands

Goal

Use the bot's supported media and discover the commands available in your installed build.

Before you start

For image understanding, a vision-capable route must be ready. For incoming voice notes, speech recognition must be ready.

Steps

- 
1. Send normal text, a photo containing text, or a Telegram voice message to the bot.
 2. In Local Waifu, open Settings → Telegram → Available commands.
 3. Use Telegram's / menu to choose a command. The list is generated by the app, so it is the authoritative list for your build.
 4. Common commands include /help, /ping, /status, /response, /think, /model, /voices, /lang, /imgmodel, /mood, /newchat, /reset, and /image <prompt>.
 5. For dynamic lists such as models or voices, select the numbered command presented by the bot.

Expected result

Text is handled as chat, text in photos can be read, voice messages are transcribed, and commands open the corresponding choices or actions.

If something goes wrong

- If a photo is ignored, verify that a vision-capable model or provider is available.
- If a voice note cannot be understood, check speech recognition under Settings → Models → Voice Model.
- Model changes made through Telegram also affect the desktop app because both interfaces share the active model.
- /reset asks for explicit confirmation. Read the bot's prompt before confirming.
- An unknown slash command may be treated as normal conversation text.

Stop the bot or remove its credentials

Goal

Pause Telegram access or completely remove the saved token and owner pairing.

Before you start

Understand the difference:

- Stop pauses the bridge but keeps the token and pinned owner.
- Forget token + pin stops the bridge and removes both the saved token and pairing.

Steps

1. Open Settings → Telegram.
2. To pause the bot, select Stop.
3. To remove access completely, select Forget token + pin.
4. Select Yes, forget to confirm.

Expected result


After Stop, the status is Stopped and the same configuration can be started again. After Forget token + pin, the token and owner pairing are removed, allowing a new token or owner to be configured.

If something goes wrong

Forget token + pin remains available even while the bot is stopped. Use it if a stale token or old owner pairing prevents a new setup.

FIGURE 18

Performance settings: context size, deeper thinking, model routing, and runtime behavior.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

23. Tune Performance, Hardware, and Network

Open Settings → Hardware to inspect storage, choose compute behavior, set the local context window, reduce resource use, and control network routing.

Choose a compute device and verify acceleration

Goal

Use the best available hardware and confirm whether the active local model is running on the GPU or CPU.

Before you start

- Send at least one message so the local model is loaded before checking acceleration.
- A GPU is optional on Windows; Local Waifu can fall back to CPU.
- macOS builds require Apple Silicon and use unified memory.

Steps — macOS

1. Open Settings → Hardware → Compute Performance.
2. Under Compute Device, choose:
 - Auto — recommended; lets the local engine select acceleration.
 - CPU — forces CPU-only inference.
 - Metal (GPU) — allows Metal acceleration on Apple Silicon.
3. Restart Local Waifu when prompted.

Steps — Windows

1. Open Settings → Hardware → Compute Performance.
2. Leave Compute Device on Auto-detect (Recommended) unless troubleshooting.
3. Send a message to load the model.
4. Under Local AI acceleration, select Check.
5. Read the result:
 - Running on GPU — the model is in VRAM.
 - Split CPU/GPU — only part of the model fits in VRAM.
 - Running on CPU — the GPU is not engaged.
 - No model loaded yet — send a message, then check again.
6. If NVIDIA or AMD hardware is detected but not fully used, select Retry GPU setup. The app may download a matching runtime and restart the engine.

Expected result

macOS uses Metal unless CPU-only mode was selected. Windows reports the real placement of the loaded model and, where available, the percentage held in VRAM.



If something goes wrong

- Restart the app after changing Compute Device.
- Try a smaller or more heavily quantized model if Windows reports Split CPU/GPU.
- Update GPU drivers, then use Retry GPU setup again.
- CPU fallback is functional but can be much slower.
- There is no separate Apple Neural Engine inference path in the macOS 1.7.3 build; use Auto or Metal (GPU).

Check runtime readiness

Goal

Confirm which local features can run now.

Before you start

This detailed panel is available in the Windows build. On macOS, inspect the corresponding model panels if the runtime readout is not shown.

Steps

1. Open Settings → Hardware.
2. Under Runtime status, select Check.
3. Review Local AI engine, Memory model, Chat model, Voice output, Sees images, and Image generation.
4. For an item marked Not available, open the relevant model section and install or select the missing component.

Expected result

Each feature you intend to use is marked Ready.

If something goes wrong

Restart Local Waifu after installing a runtime or model, then run Check again. If the local engine remains unavailable, verify Ollama Host and review App Logs as described in Chapter 27.

Set the context window and thinking mode

Goal

Balance conversation context, memory use, and response speed for local chat models.

Before you start

Larger context windows require more RAM or unified memory and may substantially slow replies. Cloud providers use their own limits.

Steps

1. Open Settings → Hardware → Context Window.
2. Start with 8K (Medium).

- 
3. Increase to 16K (High) or 32K (Ultra) only if your model and available memory can support it.
 4. Use Custom only when you know the model's context limit. Windows accepts 8,192–262,144; the macOS field can accept 512–262,144, although the active model may support less.
 5. Turn on Deeper thinking only when you prefer more considered but slower responses.

Expected result

Local replies use the requested context size within the active model's capabilities. Deeper thinking trades speed for more reasoning.

If something goes wrong

- Reduce Context Window if responses become slow, memory pressure rises, or the model fails to load.
- A high value is a request, not a guarantee; the active model may support less.
- Turn off Deeper thinking for faster, more responsive chat.

Reduce background resource use

Goal

Make Local Waifu release model memory sooner and perform fewer background wake-ups.

Before you start

Low Power Mode can make the next reply slower because the local model is unloaded sooner between chats.

Steps

1. Open Settings → Hardware.
2. Turn on Low Power Mode.
3. Restart Local Waifu.

Expected result

Background work occurs less frequently and the local model leaves memory sooner after inactivity.

If something goes wrong

If first-response latency becomes inconvenient, turn off Low Power Mode and restart the app.

Configure proxy, offline mode, or a remote Ollama server

Goal

Control network behavior without disabling local AI unnecessarily.

Before you start

- Use System Proxy is enabled by default and changes require a restart.
- Offline Mode blocks cloud providers but does not block local Ollama.
- A remote Ollama Host exposes prompts to that other computer and to the intervening network; use only a host you trust.

Steps

- 
1. Open Settings → Hardware → Network.
 2. Leave Use System Proxy enabled if your network depends on operating-system proxy settings. Restart after changing it.
 3. Turn on Offline Mode to prevent cloud chat providers from running.
 4. To use Ollama on another computer, enter its URL in Ollama Host, for example `http://192.168.1.50:11434`.
 5. To return to the bundled local engine, select Reset beside Ollama Host.

Expected result

Local Ollama continues to work in Offline Mode. Cloud models refuse to run. A valid custom Ollama Host routes local-model requests to the specified server.

If something goes wrong

- If cloud chat reports that it is blocked, turn off Offline Mode only if sending data to the configured provider is intentional.
- If local chat goes offline after editing Ollama Host, select Reset.
- If downloads fail behind a corporate proxy, enable Use System Proxy and restart.
- Check firewall access to the remote Ollama host and port; do not expose Ollama directly to the public internet.

Storage note: Local models can occupy tens of gigabytes. The Change Path control is disabled in the macOS 1.7.3 UI. Do not rely on it for moving models. The Windows build can support moving model storage, but perform a backup first and follow the controls shown by that build.

24. Understand Privacy and Data Storage

Local Waifu is local-first, not “always local.” Local models, memories, and embeddings can remain on your computer, but cloud AI, cloud speech services, Telegram, and remote Ollama send selected data outside the local app.

Know what stays local and what can leave the computer

Activity	Where data goes
Local chat through bundled Ollama	Processed on your computer
Memories and embeddings	Stored locally; embeddings remain local even when chat uses a cloud provider
Local image generation	Processed on your computer
Local speech recognition and voice	Processed on your computer
Cloud chat or image model	Prompt, relevant conversation context, and requested input go to the provider you selected
Cloud STT/TTS	Audio or text required by that service goes to the selected provider
Telegram	Messages and media pass through Telegram and the Local Waifu bot integration
Remote Ollama Host	Model requests go to the configured network computer
Model downloads and update checks	Network requests go to the relevant model host or update service

Important: Offline Mode blocks configured cloud model routes, but it does not make Telegram operate without Telegram’s servers, stop model downloads already requested, or block a custom network Ollama Host.

Locate your Local Waifu data

- macOS: `~/Library/Application Support/com.lumizone.localwaifu/`
- Windows: `%APPDATA%\com.lumizone.localwaifu\`

The app data directory contains the database, character data, soul files, portraits, generated images, voice references, settings, logs, and `secrets.json`. Model files may occupy separate model directories shown under Settings → Hardware → Storage.

Soul files containing personality, memory, and knowledge are encrypted with cocoon encryption. Their key material is tied to the original computer environment. This is one reason a copied full data directory is not the recommended self-service migration method between computers.

The `secrets.json` file can contain BYOK API keys, the Telegram bot token, and activation-related data. On Unix-like systems it is protected with owner-only file permissions, but anyone who can read your user



account's files may still read it. Protect the operating-system account and disk.

Review and control permissions

Goal

Enable only the operating-system permissions needed for the features you use.

Before you start

Permissions are requested when a related feature is first used, not simply because the application launches.

Steps

1. Enable microphone access only for dictation, voice calls, or voice-message workflows.
2. On macOS, open Settings → Advanced → Privacy and enable Proactive calendar reminders only if you want Local Waifu to check your calendar periodically.
3. Approve Calendar access when macOS asks. Reminder access may be requested when a reminder feature is used.
4. To revoke access later:
 - macOS: use System Settings → Privacy & Security and select the relevant category.
 - Windows: use Settings → Privacy & security and select the relevant permission.
5. Keep Blur on Background enabled under Settings → General if you do not want conversation content visible when the window loses focus.

Expected result


Only the features you intentionally use have operating-system access. Disabling a permission blocks that feature but does not delete existing data.

If something goes wrong

If a feature remains denied after you change the operating-system setting, quit and reopen Local Waifu. For microphone failures, also verify the input device under voice-call settings.

FIGURE 19

Advanced data controls: diagnostics and reset are separate from full backup and conversation export.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

25. Back Up, Export, and Move Your Data

Local Waifu offers three different data operations. They are not interchangeable.

Operation	Contains	Best use	Can it restore the whole app automatically?
Make a backup	Snapshot of the Local Waifu app data directory, including database, soul files, history, and settings	Recovery copy before an update or support-assisted restore	No. Version 1.7.3 has no complete Restore UI.
Export Conversations	Chat-message records in JSON	Reading, archiving, or processing messages	No; there is no general conversation-import control in this panel
Character export	One character plus related messages, memories, relationship data, and knowledge	Supported self-service transfer to another installation	Imports as a new character; it is not a full-app restore

Create a full backup snapshot

Goal

Create a consistent copy of the Local Waifu application data before a major update, reset, repair, or computer migration.

Before you start

- Choose a destination outside the Local Waifu data directory, preferably another drive or a synced backup location.
- Make sure the destination has enough free space.
- A full backup can contain private conversations, images, voice references, settings, and secrets. Store it securely.

Steps

1. Open Settings → Advanced.
2. Review Data health. Resolve a reported integrity problem before relying on a new backup.
3. Under Backup & safety, find Make a backup.
4. Select Choose folder.
5. Pick an existing destination folder outside Local Waifu's app data directory.
6. Wait until the row reports Saved, a file count, and a total size.
7. In the destination, confirm that a new folder named `local-waifu-backup-<timestamp>` exists.
8. Copy the backup to a second protected location if it is your only recovery copy.

Expected result



Local Waifu creates a timestamped folder and a consistent SQLite snapshot. Temporary SQLite WAL/SHM files are not required in the backup.

If something goes wrong

- If the app reports that the destination overlaps Local Waifu's data directory, choose another folder or drive.
- If the copy stops, check free space and destination permissions, then use Backup again.
- Do not judge success only by the folder's existence; confirm the in-app file count and size.

Restore warning: Local Waifu 1.7.3 does not provide a complete Restore backup button or automated full-backup import. Do not copy a backup over the live data directory while the app is running. For full restoration, quit the app and obtain a version-appropriate support procedure. A manual replacement may fail if encryption or machine-bound data no longer matches.

Export all conversations

Goal

Save chat messages as a JSON file for archive or analysis.

Before you start

This export is not a backup. It does not include the complete character, memories, soul files, models, credentials, or settings, and the same panel has no general import function for this file.

Steps

1. Open Settings → Advanced → Data Export.
2. Select Export Data beside Export Conversations.
3. Choose a new .json filename and location outside the Local Waifu data directory.
4. Wait for Saved!.
5. Keep the file private; message content is readable JSON.

Expected result

A file named like `localwaifu-export-YYYY-MM-DD.json` contains the exported message records.

If something goes wrong

- Choose an absolute path ending in .json.
- Do not select the Local Waifu data directory.
- Choose a new name; the exporter refuses to overwrite an existing export.

Move a character to another computer

Goal

Use the supported character export/import path for a self-service migration.

Before you start

- This transfers the exported character and related messages, memories, relationship data, and knowledge.

- It does not transfer installed models, provider keys, connected accounts, the Telegram token/pairing, app-level preferences, or the Direct license activation.
- The imported character is created as a new character; it does not overwrite an existing one.

Steps

1. On the old computer, open Settings → Characters.
2. Export the character and save the generated .json file.
3. Transfer that file securely to the new computer.
4. Install and complete onboarding in Local Waifu on the new computer.
5. Open Settings → Characters and import the character JSON.
6. Confirm the reported counts for messages, memories, and knowledge entities.
7. Reinstall the required chat, embedding, image, and voice models.
8. Re-enter cloud API keys or reconnect accounts.
9. Configure Telegram again with Settings → Telegram.
10. If using the Direct edition, deactivate the old device before disposal and activate the license on the new computer.
11. Continue using the imported character while the app rebuilds embeddings.

Expected result

A new character appears with the exported history and character-related data. Semantic recall improves as embeddings are rebuilt during continued use.

If something goes wrong

- Import failed means the JSON is damaged or incompatible; export it again from the source installation.
- If memory retrieval is temporarily weaker, allow time for embedding reconstruction and keep the required memory model installed.
- If you need complete app recovery rather than a character transfer, preserve the full backup and contact support. There is no complete Restore UI in 1.7.3.

26. Update or Uninstall Local Waifu

Update the Direct edition

Goal

Install an available Direct-edition update without risking known-bad local data.

Before you start

1. Create a backup with Settings → Advanced → Make a backup.
2. Check Data health.
3. Connect to the internet.
4. Note the platform difference: the macOS Direct updater is license-gated; the Windows 1.7.3 Direct updater does not apply the same lifetime-only gate.

Steps

1. Open Settings → General.
2. Find Check for Updates and select Check Now.
3. If a newer version appears, read the version number.
4. Select Install & Restart.
5. Wait for the download and allow the app to close and relaunch.
6. Open Settings → About and verify the new version.

Expected result

The update passes the pre-update database and encryption-metadata safety checks, installs only after your confirmation, and relaunches on the newer version.

If something goes wrong

- If Data health fails or installation is blocked, do not force the update. Make a backup and contact support.
- On macOS Direct, a trial license may be prevented from using the in-app update route; use the supported download route for your entitlement.
- If the updater cannot connect, check proxy, firewall, VPN, and Offline Mode, then use Download latest manually from GitHub if shown.
- Nothing installs merely because automatic update checks are enabled. Download and installation require user action.

Update the Store/Unlocked edition

Goal

Keep the Store/Unlocked build on its own distribution channel.

Before you start

The Store/Unlocked build has no Direct-edition updater, license tab, or Direct release-channel access.

Steps

1. Create a backup.
2. Check the store or client from which you obtained Local Waifu, such as itch.io or the itch app.
3. Download or install the update through that same channel.
4. Do not replace it with a Direct build merely to obtain the Direct updater.
5. After updating, open Settings → About and verify the version.

Expected result

The Store/Unlocked edition remains license-free and continues using its correct update channel.

If something goes wrong

Confirm that the downloaded installer is the Store/Unlocked build. If the update unexpectedly presents a trial, license field, or paywall, stop and obtain the correct store build before changing data.

Uninstall while keeping your data

Goal

Remove the application while preserving data for later reinstallation.

Before you start

Create a backup. If transferring or selling the computer and you use Direct, first open Settings → License and select Deactivate this device.

Steps — macOS

1. Quit Local Waifu.
2. Move Local Waifu from Applications to the Trash.
3. Do not delete `~/Library/Application Support/com.lumizone.localwaifu/`.

Steps — Windows

1. Quit Local Waifu and wait for its local engine processes to stop.
2. Open Settings → Apps → Installed apps → Local Waifu → Uninstall.
3. Leave Delete the application data unselected if that option is offered.

Expected result

The program is removed while the app data remains available to a later installation.

If something goes wrong

If reinstallation starts as a blank app, do not overwrite or reset anything. Confirm that the original data directory still exists, preserve it, and contact support.

Remove Local Waifu and its data completely

Goal

Remove the application, local content, and saved credentials from a computer you no longer intend to use.

Before you start



This is destructive. Export characters and create an external backup first. Deactivate a Direct license before wiping or selling the computer.

Steps

1. Remove cloud provider keys and connected-account credentials from their settings panels.
2. Open Settings → Telegram and use Forget token + pin.
3. Open Settings → Advanced, select Reset All, then Yes, delete.
4. Restart the app once if you need to confirm the reset completed.
5. Quit and uninstall Local Waifu using the platform procedure above.
6. Delete the remaining Local Waifu data directory only after verifying your backup:
 - macOS: ~/Library/Application Support/com.lumizone.localwaifu/
 - Windows: %APPDATA%\com.lumizone.localwaifu\
7. Remove separately stored model files if you also want to reclaim their disk space.

Expected result

The application, user data, secrets, and optional model files are removed from the computer.

If something goes wrong

Reset All does not remove every credential by itself. In particular, license/trial data, BYOK keys, and the Telegram token can remain in the secret store. That is why credentials must be removed explicitly and the remaining data directory deleted for a complete wipe.

27. Use Logs, Diagnostics, and Reset

Open the app log

Goal

Collect recent technical details after reproducing a failure.

Before you start

Logs may include filenames, device information, model names, error text, and other context. Do not post them publicly without review.

Steps

1. Reproduce the problem once.
2. Open Settings → Advanced → Developer Options.
3. Under App Logs, select Open Logs.
4. Note the approximate time of the failure and copy only the relevant section when possible.

Expected result

The detailed application log opens in the system's default text viewer.

If something goes wrong

If the log is empty or does not open, use Diagnostics report → Save report instead. The app log and audit log serve different purposes; the app log is normally more useful for chat, voice, and startup failures.

Save a diagnostics report

Goal

Create one text file describing the build, configuration, hardware, runtime state, and recent log activity for support.

Before you start

Privacy warning: Review the report before sending it. Local Waifu redacts license values, API keys, and Telegram tokens, but the report includes audio-device names and an app-log tail that is not fully content-filtered. Do not assume the UI phrase “No keys, no chat text” makes every line safe to publish.

Steps

1. Reproduce the issue.
2. Open Settings → Advanced → Developer Options.
3. Under Diagnostics report, select Save report.
4. Locate the revealed file named like `diagnostics-YYYYMMDD-HHMMSS.txt`.
5. Open it in a text editor.
6. Review every section, especially the recent app-log tail.
7. Remove or redact private text, usernames, paths, device names, prompts, or other details you do not want to share.

8. Attach the reviewed file to a private support message.

Expected result

A reviewed diagnostic text file is ready to send without exposing information you chose to remove.

If something goes wrong

If report creation fails, send the app version, platform, edition, exact error message, reproduction steps, and a reviewed excerpt from App Logs instead.

Reset all application data

Goal

Return the main application database and local content to a fresh state when repair is impossible or when intentionally starting over.

Before you start

Destructive action: Reset All is irreversible from within Local Waifu. Create a backup and export any characters you want to keep.

The reset deletes the database, characters, conversations, memories, soul files, portraits, generated images, voice references/clones, and local settings. It does not automatically remove the Direct license/trial record, BYOK keys, connected-account credentials, or Telegram token from the secret store.

Steps

1. Create a full backup and verify its folder, file count, and size.
2. Export important characters separately.
3. If you also want credentials removed, delete provider keys and use Settings → Telegram → Forget token + pin first.
4. Open Settings → Advanced → Developer Options.
5. Find Reset Database and select Reset All.
6. At Sure?, select Yes, delete.
7. Wait for All data deleted.
8. Quit and restart Local Waifu as instructed.

Expected result

The main app data is cleared, and the application starts in a fresh state after restart. License, trial, provider, or Telegram credentials may remain unless removed separately.

If something goes wrong

- Do not repeat the reset if the app appears partially initialized; quit and reopen it first.
- Do not attempt to copy a backup into the live data directory while Local Waifu is running.
- There is no complete Restore UI in version 1.7.3. Preserve the backup and contact support for a full restore procedure.

28. Troubleshooting

Run the basic diagnostic sequence

Goal

Identify the failing layer before changing or deleting data.

Before you start

Do not use Reset All as a first troubleshooting step. Preserve the exact error message and create a backup before destructive action.

Steps

1. Confirm the version under Settings → About.
2. Restart Local Waifu and reproduce the issue once.
3. On Windows, run Settings → Hardware → Runtime status → Check.
4. Check the active chat, memory, image, speech, and voice models relevant to the failure.
5. Check Settings → Hardware → Network, especially Offline Mode, Use System Proxy, and Ollama Host.
6. Check free disk space and available RAM/VRAM.
7. Open Settings → Advanced → App Logs.
8. If support is needed, save and review a Diagnostics report before sending it.

Expected result

The problem is narrowed to installation, model/runtime, resource pressure, network routing, permissions, Telegram pairing, license, update, or data health.

If something goes wrong

Stop before reset or manual file replacement. Preserve the data directory and backup, record the last successful action, and contact support.



Installation and launch

Symptom	Likely cause	What to do
macOS blocks the first launch	Gatekeeper verification or an unrecognized download	Confirm the installer came from the official source. In Finder, Control-click Local Waifu, choose Open, and approve the system prompt if macOS offers it. Do not routinely disable Gatekeeper or strip quarantine attributes.
Windows shows SmartScreen	The installer may not have paid Authenticode reputation	Confirm the file came from the official source, inspect the publisher/source, select More info, then Run anyway only if you trust the file.
Windows app does not open or shows a blank window	Missing/damaged WebView2, antivirus interference, or incomplete install	Reboot, install/repair Microsoft Edge WebView2 Runtime, allow the official app in security software, and reinstall.
Installer appears stuck	Antivirus scanning, low disk space, or running app processes	Wait for scanning, verify free space, quit Local Waifu and its local engine, then retry.

Chat models, downloads, and memory

Symptom	Likely cause	What to do
Model download will not start or stalls	Engine startup, network, proxy/VPN, registry access, or disk space	Check Use System Proxy, VPN/firewall, free space, and internet access. Restart the app, then retry in Settings → Models.
Hugging Face repo has no installable quant	It contains sharded GGUF files or no GGUF	Choose a repository with a single runnable .gguf quant. Sharded GGUF and SafeTensors are not installed through the chat-model search panel.
Chat shows offline or no reply	Local engine/model unavailable, bad Ollama Host, cloud blocked, or missing provider credentials	Run Runtime status, check the active chat model, reset Ollama Host, disable Offline Mode only for intentional cloud use, and verify the provider account/key.
Clicking Stop ends the visible stream but CPU/GPU stays busy	Stop releases the UI while the backend may finish the current generation	Wait for the current generation to settle before starting resource-intensive work.
Replies are very slow or the app runs out of memory	Model or context is too large; GPU placement is partial/CPU	Reduce Context Window, choose a smaller quantized model, turn off Deeper thinking, and check GPU placement on Windows.
Memory recall is weaker after changing the embedding model	Re-embedding is still in progress	Keep the required memory model installed and allow continued use to rebuild embeddings. Keyword-only recall may be used temporarily.

Images and local image models

Symptom	Likely cause	What to do
Image cannot be attached	File is not an image, exceeds 5 MB, another image is pending, or vision is unavailable	Use one valid image under 5 MB and verify Sees images or a vision-capable model/provider.
Custom checkpoint import fails	Wrong format, too small, invalid SafeTensors, or not SDXL	Use one regular .safetensors SDXL checkpoint of at least 512 MiB. .ckpt is not supported by the 1.7.3 importer.
Imported LoRA cannot be enabled	Active model is not SDXL, family is unknown, or three LoRAs are already active	Select an SDXL base model, use a verified SDXL LoRA, and keep at most three active. Confirm compatibility with a test render.
SDXL imports but will not run	Insufficient RAM or disk/runtime resources	Custom SDXL uses a 16 GiB RAM profile. Close other apps or use a lighter image model.
Image model cannot be deleted	Rendering or download is active	Wait for the operation to finish, then retry.

Voice, microphone, and calls

Symptom	Likely cause	What to do
A call cannot start	Speech recognition is missing or unavailable	Install/configure Whisper or a supported cloud STT route. STT is required for a full call.
No spoken reply	Voice model is missing or unavailable	Install Supertonic or configure a supported cloud TTS route. A fallback may be used, but quality can be lower.
Speech has the wrong accent or language	Wrong language, English fallback, or mismatched custom sample	Select the correct language and voice model. For a custom voice, use 10–30 seconds of clean speech in the target language.
Microphone is silent	OS permission denied or wrong input device	Grant microphone permission in the operating system and select/test the correct input device. Restart afterward.
Dictation stops unexpectedly	Ten-minute limit or disconnected microphone	Start another dictation session and reconnect or reselect the microphone.

Telegram

Symptom	Likely cause	What to do
Test succeeds but bot does not answer	Bot was not started	Select Start. Testing does not save or start the token.
Bot will not pair with a new account	It is pinned to the first chat	Use Forget token + pin, restart the bot, and message it first from the correct account.
Voice only returns text	Voice model is not ready	Install/configure the voice model. Text is the intentional fallback.
Telegram model change also changes desktop	Shared active-model setting	This is expected. Select the preferred model again from either interface.

License and updates

Symptom	Likely cause	What to do
Direct key is rejected	Typo, no network for activation, invalid key, or device limit	Paste the exact key, reconnect, and retry. Deactivate an old device or contact support if the device limit is reached.
Trial expired	Direct-edition trial period ended	Activate a valid lifetime key. User data is retained; do not reset in an attempt to obtain a new trial.
License tab is absent	Store/Unlocked edition	This is expected; that edition does not require a key.
Update control is absent	Store/Unlocked edition	Update through the store/client used to obtain the app. Do not mix it with the Direct updater.
Install & Restart is blocked	Data-health or encryption-metadata check failed	Create a backup and contact support. Do not force file replacement.




Backup, export, and data recovery

Symptom	Likely cause	What to do
Backup destination is rejected	It overlaps the live Local Waifu data directory	Choose an external folder or another location outside the app data tree.
Conversation JSON does not restore the app	It is an export, not a backup	Use character export/import for self-service transfer. Preserve the full backup for support-assisted recovery.
There is no Restore backup button	Version 1.7.3 has no complete Restore UI	Quit the app, preserve the backup unchanged, and request a support procedure. Never merge it into the live directory while the app is running.
Reset did not remove license, cloud keys, or Telegram token	Reset does not clear the secret store	Remove provider credentials explicitly, use Forget token + pin, and delete the remaining data directory only for a deliberate full wipe.

FIGURE 20

About shows the installed Local Waifu version, AI disclosure, open-source acknowledgements, and backup reminder.



Interface shown on macOS. Windows uses the same setting unless this guide marks a platform difference.

The exact control is easier to find once you know where to look.

29. Shortcuts, Contact, and About

Keyboard shortcuts

The following shortcuts are implemented in Local Waifu 1.7.3 even though About no longer contains a keyboard-shortcut reference section.

Action	macOS	Windows
Open Settings	Command+,	Ctrl+,
Search chats	Command+F	Ctrl+F
Start or end a voice call	Command+Shift+C	Ctrl+Shift+C
Switch to Full view	Command+Shift+1	Ctrl+Shift+1
Switch to Mini view	Command+Shift+2	Ctrl+Shift+2
Switch to menu bar/tray view	Command+Shift+3	Ctrl+Shift+3
Cancel or close the active transient UI where supported	Escape	Escape

Search note: both Command/Ctrl+F and Command/Ctrl+Shift+F open chat search.

Check the installed version and edition clues

Goal

Collect the information support needs and verify an update.

Before you start

Know whether you installed Local Waifu from the Direct download or a Store/Unlocked channel. The presence of Settings → License is a useful clue: Direct shows it; Store/Unlocked does not.

Steps

1. Open Settings → About.
2. Read the number shown after Version.
3. Note your operating system and whether Settings → License is present.
4. If reporting a model or voice issue, also note the active model names.
5. If requested, save a Diagnostics report, review it, redact private details, and attach it privately.

Expected result

You have the app version, platform, edition clue, reproduction steps, exact error text, and a reviewed diagnostic file if needed.

If something goes wrong

If About cannot load, include the installer filename or download channel and state that the version screen was unavailable. Do not reset the app merely to retrieve a version number.

Contact and official links

- Support email: contact@localwaifu.com
- Website: <https://localwaifu.com>
- Community: <https://discord.gg/RbaVVTADxz>
- Direct-edition source/releases: <https://github.com/lumizone/local-waifu>
- Privacy: <https://localwaifu.com/privacy/>
- Terms: <https://localwaifu.com/terms/>

The About screen also identifies Local Waifu as local-first and private by design, provides the AI disclosure, lists open-source components, and links to Privacy and Terms. The GitHub button is shown only where appropriate for the distribution edition.

When contacting support, include:

1. Local Waifu version.
2. macOS or Windows version.
3. Direct or Store/Unlocked edition.
4. What you expected and what happened.
5. Exact steps that reproduce the issue.
6. The exact error message.
7. Relevant model/provider names, without API keys or tokens.
8. A reviewed and redacted diagnostics report when useful.

Never send a license key, Telegram bot token, provider API key, unreviewed diagnostics report, or private conversation export in a public channel.

I want to...

Use this index when you know the task but not the Local Waifu term for it.

I WANT TO...	GO TO
Move one companion to another computer	Chapter 15 · Export and Import a Character
Move or protect the entire installation	Chapter 25 · Back Up, Export, and Move Your Data
Install my own chat model	Chapter 17 · Install a GGUF Model from Hugging Face
Import an image checkpoint or LoRA	Chapter 20 · Import a Custom SDXL Checkpoint or LoRA
Make memory recall work	Chapters 13 and 16 · Memory Book and EmbeddingGemma
Reduce RAM usage or speed up replies	Chapter 23 · Performance, Hardware, and Network
Use Local Waifu from my phone	Chapter 22 · Use Local Waifu from Telegram
Collect logs or a diagnostics report	Chapter 27 · Logs, Diagnostics, and Reset
Solve an error	Chapter 28 · Troubleshooting
Find shortcuts, version, or support	Chapter 29 · Shortcuts, Contact, and About



End of guide

Keep this PDF with your Local Waifu backup instructions. For current downloads and product information, use the official Local Waifu channels. For account, license, or recovery help, contact contact@localwaifu.com.

Local Waifu · Your companion, your computer, your data.