

On MacOS v26 (on MacBook Pro M5)

```
curl -fsSL https://ollama.com/install.sh | sh
```

In use

ollama pull qwen3.8:27b-mix

Fast: MoE daily driver (~23 GB)

ollama pull qwen3.6:35b-a3b

the official qwen3.6:35b-a3b tag is a 23GB Q4_K_M. That's tight on 32 GB, so if memory gets squeezed, pull a smaller community IQ4_XS instead (~17.7GB):

```
ollama pull hf.co/bartowski/Qwen_Qwen3.6-35B-A3B-GGUF:IQ4_XS
```

Small model: for autocomplete

```
ollama pull qwen3.5:9b
```

```
anton-w@ant1mbp7 ~ % ollama pull qwen3.5:9b
pulling manifest
pulling dec52a44569a: 100%
pulling 9371364b27a5: 100%
verifying sha256 digest
writing manifest
success
anton-w@ant1mbp7 ~ %
```

```
6.6 GB/6.6 GB 57 MB/s 0s
65 B
```

Tested

- good, but too big, calculations are long (40'--4h), almost can not predict workflow pace (timing). Sometimes takes too much time for smaller tasks.

```
ollama pull qwen3-coder:30b
```

Revision #2

Created 25 September 2026 07:22:15 by Anton

Updated 25 September 2026 08:07:54 by Anton