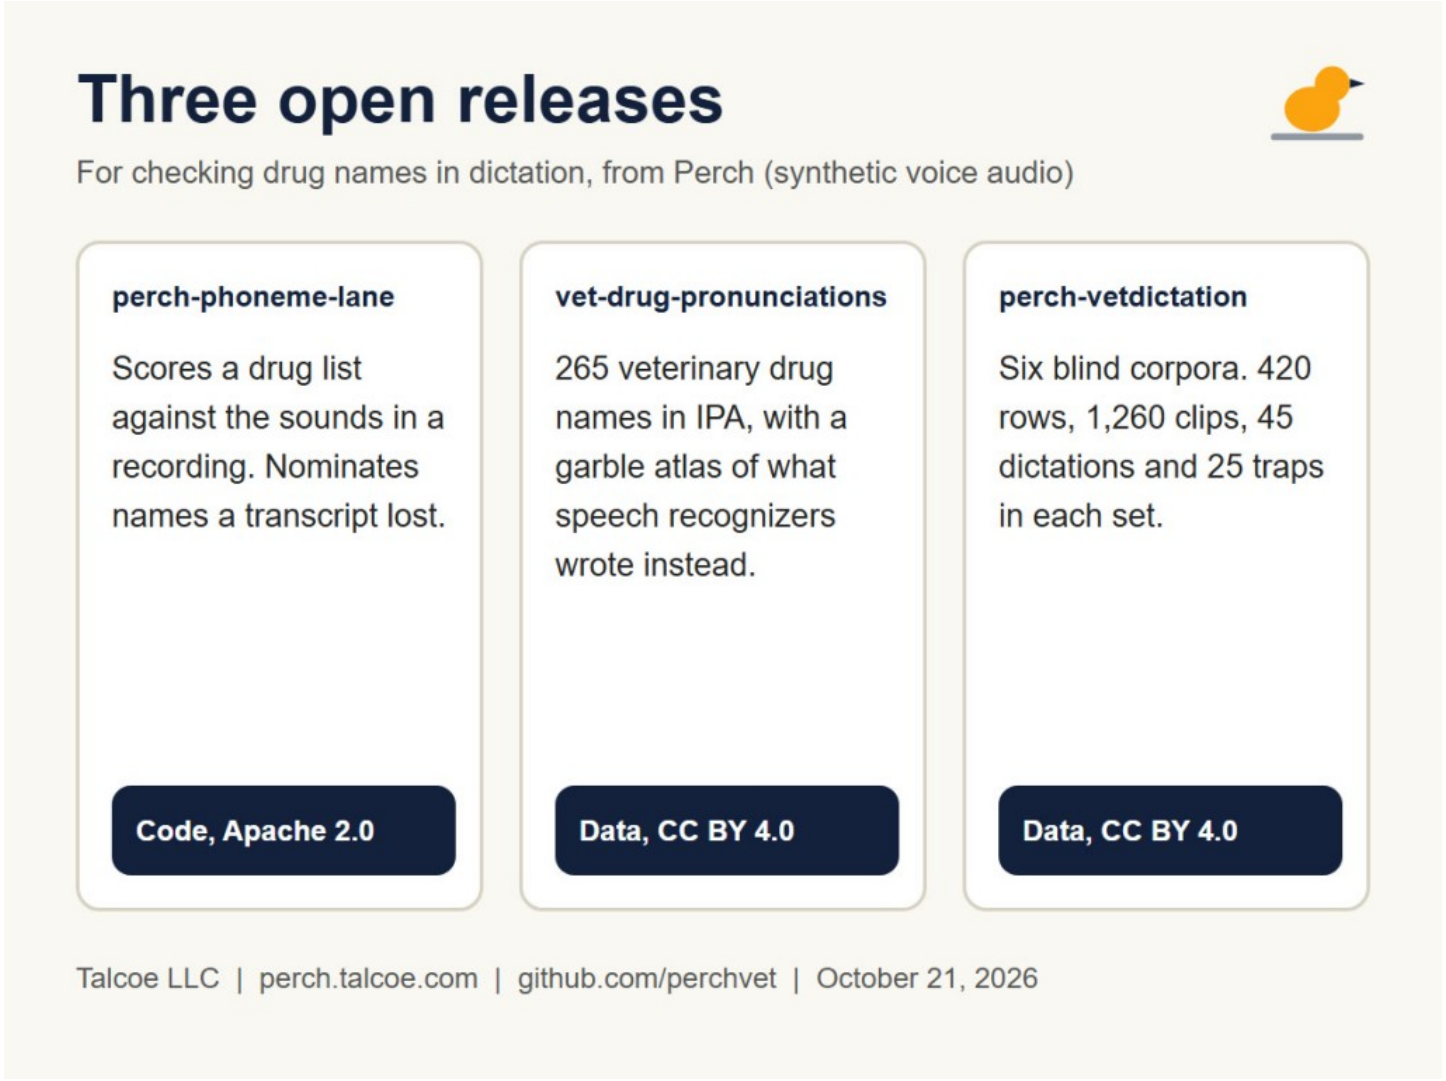
Perch Publishes Open Tool, Pronunciation Table and Six Dictation Corpora for Veterinary Drug Names

The code is Apache 2.0 and the data is CC BY 4.0, so a vet, a clinic or a vendor can rerun the tests on their own system. The audio is synthetic speech.



New York City, New York Oct 27, 2026 (Issuewire.com) - Perch, the veterinary AI scribe built by Talcoe LLC, on October 21 published three open releases for checking drug names in dictation. They are a tool that scores a drug list against the audio of a dictation, a table of how 265 veterinary drug names are said and misheard, and six blind dictation corpora with 1,260 clips.

The tool is the phoneme lane. It scores the audio, not the transcript. It runs a phoneme model over a dictation and scores every name on a clinic's drug list against the sounds it heard. It returns the closest matches and the seconds of the clip where they sit. It proposes names. It does not chart them.

Two synthetic examples ship with it. In an English clip read by a Kokoro voice, a transcript spells famotidine "famitidine," maropitant "meripitin" and Entyce "Entice." The lane's top four names are famotidine, morphine, Entyce and maropitant, and all four were said. That is one clip, not a measurement. In a Hindi sample sentence naming Heartgard and Bravecto, the transcript comes out

entirely in Devanagari, so a text matcher has no Latin letters to read. The lane ranks Heartgard first.

The table is vet drug pronunciations. It lists 265 drug and brand names used in veterinary practice, each with its pronunciation in IPA, the International Phonetic Alphabet, for American, British and South Asian English speakers. Of the 265 names, 66 are on Perch's default list. The other 199 are drugs that are not on that list. They are the distractors, names a matcher must not confuse with the 66. For 59 of the 66, a garble atlas lists the wrong words that speech recognizers wrote when synthetic voices said the name. Each wrong spelling is counted, and flagged if it is an ordinary English word. Atenolol, for example, came out as "8 nilal."

The corpora are Perch blind dictation corpora. There are six sets of 70 rows, which makes 420 rows. Each set holds 45 dictations and 25 traps. A trap is a sentence built to bait a system into offering a drug nobody said. A separate author wrote each set, and none of them saw Perch's matching code. Three Kokoro voices read each corpus, which makes 1,260 synthetic clips and about five hours of audio. The first set holds the 70 clips of Perch's earlier trap set and adds two more voices for each sentence.

This release carries limits, noted here where they apply. The audio in every set and every example is synthetic. It is Kokoro text to speech voices, not a recording of a real clinician in a real exam room, and real speech will differ. The phoneme lane has only been tried on synthetic speech. The voices and the lane's pronunciations share a source, so the samples flatter the lane. The pronunciations were produced by software, and none was recorded from a speaker. Blind describes how the rows were written, not who has seen them. A system whose builders have read the rows can be tuned to them. These are Perch's own tests, run by the company that builds Perch. Talcoe publishes the sets, the table and the code so others can rerun them. This is not an accuracy claim about any product.

The tool is at https://github.com/perchvet/perch-phoneme-lane under Apache 2.0, release v0.1.0. The table is at https://github.com/perchvet/vet-drug-pronunciations under CC BY 4.0, release v1.0.0. The corpora are at https://github.com/perchvet/perch-vetdictation under CC BY 4.0, release v1.0.0, with the audio as release assets. A longer account, written by Talcoe's cofounder, is posted at https://perch.talcoe.com/news/three-open-releases.

"A company's own results are a claim until someone else can rerun them. So we are publishing the sentences we score Perch on, the table our matcher reads and the tool we built for the names it still missed. Run your own system against them and tell us what it misses," said Mohammed Usama, cofounder of Talcoe LLC.

Perch runs in the browser with nothing to install. Try it without signing in at perch.talcoe.com/demo.

Method note: The transcripts quoted above are from Whisper large v3. The phoneme model is wav2vec2 xlsr 53 espeak, with 388 phone symbols and 20 millisecond frames. The garble atlas holds what Whisper large v3 and MedASR wrote when six Kokoro voices said each of the 66 names in three sentence frames. Kokoro v1.0, from hexgrad and under Apache 2.0, voiced the clips.

About Perch

Perch is a veterinary AI scribe built by a two person team at Talcoe LLC in New York. A vet talks through a visit, and Perch drafts the note in SOAP, DAP or BIRP format while the vet is still speaking. The vet reviews and signs the note. Perch declines to chart a drug name it cannot resolve rather than guess at one. It runs in the browser with nothing to install. Audio is processed in the cloud and never retained.

About Talcoe LLC

Talcoe LLC is a two person company in New York, building software for veterinary practices. Its two cofounders are Mohammed Usama and Mustafa Nisar. The company publishes the tests it scores its own product on, so a reader can check the work rather than take the company's word for it.

Press Contact

Mohammed Usama
Cofounder, Talcoe LLC
Email: mo@talcoe.com
Talcoe LLC

Media Contact

Mohammed Usama, Talcoe LLC

*********@talcoe.com

https://perch.talcoe.com

Source : Talcoe LLC

See on IssueWire