NOVARA
CONSULTING GROUP

This document was translated from English by machine and may contain mistakes. The English original is at https://www.novaracg.com/2026/11/25/signplay-tells-hearing-parents-when-their-auslan-is-right-it-has-not-said-how-often-the-tick-is-wrong/.

SignPlay Tells Hearing Parents When Their Auslan Is Right. It Has Not Said How Often the Tick Is Wrong.

Heather Grizzle · November 25, 2026

SignPlay grades hearing parents' Auslan with a phone camera. Its one accuracy figure has no method, and its 250,000-video dataset has no consent terms.

目录

1. A green tick on a kitchen bench
2. What SignPlay is, as of November 2026
3. What SignPlay gets right
4. A grader is not a dictionary
5. Eighty-one percent of what?
6. Hands are not the whole sign
7. A quarter of a million videos, described by size alone
8. Input is not authority
9. Voting for an app nobody can download
10. What a school, a funder or a family should ask
11. The green tick has to be earned
12. 参考文献
13. Tips

简要说明

- SignPlay, built by the University of Sydney's AID-LAB with Deaf Connect, uses a phone camera to tell hearing parents of deaf children whether their Auslan signs are correct. It is a finalist for a Shaping Australia Award, and Deaf Connect is asking the public to vote for it.
- The need is real and well documented. Deaf Connect reports 159 families in its Auslan courses, 115 more waiting, and waits of up to three months, longest outside the cities.
- The only performance number published is 81.1 percent accuracy across 2,000 signs, with no method, no test signers and no split between the two kinds of error. For a tool that grades parents, the error that matters most is the one that calls a wrong sign right, because the parent then teaches it to the child.
- Roughly a quarter of a million signing videos were collected from the Deaf community. Nothing published says how many people gave them, whether they were paid, what they agreed to, or whether the data can travel to partners, one of which is Google.org.
- Before a school, a funder or a family relies on SignPlay's green light, they should be able to read how often it is wrong, in which direction, on whose signing, and who decided what counts as correct.

A green tick on a kitchen bench

Picture a mother in Dubbo, in central west New South Wales. Her son is fourteen months old and was identified as deaf through newborn hearing screening. The family has decided to learn Auslan, and she has put her name down for a course with Deaf Connect. The course is full. She is on a waiting list that Deaf Connect itself says can run to three months, and longer for families outside the major cities [3]. In the meantime, someone has told her about an app.

She props the phone against the fruit bowl, opens it, and is asked to sign MILK. She makes the sign she half remembers from a video. The app watches her hands through the camera and, a second later, shows a green tick. She signs it again that night at the

high chair, and the next morning, and the morning after that. Her son is at the age when every sign he sees at home becomes his language.

The family is invented, and so is the tick. The waiting list, the phone camera and the verdict on whether a sign is correct all come from how SignPlay and Deaf Connect describe the app. The question this piece asks is a simple one, and it is the question that mother has no way to ask: when the app showed her a green tick, how often is that tick wrong, and who checked?

What SignPlay is, as of November 2026

SignPlay is a research project of AID-LAB, a human-computer interaction group in the School of Computer Science at the University of Sydney, led by Anusha Withana and Zhanna Sarsenbayeva [1]. The lab calls it an "AI for Good" initiative to help hearing people, especially the parents and families of deaf children, learn Australian Sign Language. It is a set of short, game-based activities on an ordinary smartphone. What sets it apart from the Auslan dictionaries and video courses that already exist is sign recognition: the phone camera watches the learner sign and the app returns a judgment on whether the sign was right. The University's awards entry puts it plainly. Unlike videos or flashcards that only demonstrate signs, SignPlay can tell users whether they are signing correctly [3].

The project page gives a start date of August 1, 2025, and names three partners: Deaf Connect, Georgia Tech, linked to its PopSign sign language learning game, and Google.org, the philanthropic arm of Google [1]. The awards entry adds Georgia Tech researchers including Thad Starner to the team, and describes Deaf Connect as having helped collect vocabulary data and as giving ongoing guidance on the app's design [3]. PopSign's partner page describes Deaf Connect as Australia's largest whole-of-life service provider for Deaf, deafblind and hard of hearing people [6], and the awards entry says it supports around 30,000 of them [3].

SignPlay is a finalist in the Community Champion category of the 2026 Shaping Australia Awards, run by Universities Australia. Public voting for the People's Choice award runs until January 19, 2027, and the winners are announced on February 24, 2027, at Parliament House in Canberra [4]. This week Deaf Connect posted on

LinkedIn asking its followers to vote, describing SignPlay as an app that "offers" short, game-based activities so families can learn Auslan at their own pace [5].

What SignPlay gets right

It is worth saying clearly what is good here, because a lot of it is good.

The problem statement is specific and sourced rather than borrowed. Plenty of sign language AI pitches open with a global statistic and never mention a real waiting list. SignPlay's entry cites Deaf Connect's own enrolment figures: 159 families in its family Auslan courses, 115 waiting, waits of up to three months, and the longest waits outside the major cities [3]. That is a measured gap in a real service, described by the organisation that runs it.

The use case, as the awards entry frames it, is modest and sensible. SignPlay is presented as a way for parents to practise between lessons, not as a replacement for being taught by Deaf people. That framing matters, and it is a better instinct than most of the market shows.

The reason given for building a new dataset is also sound. According to the entry, Deaf Connect pointed out that existing Auslan datasets were mostly recorded in studios and did not reflect everyday signing, and the project set out to collect signing in ordinary conditions instead [3]. A recognition model that has only seen studio lighting, studio cameras and expert signers will struggle in a kitchen with a parent who is still learning. Wanting real-world data is the right call.

And the project is built with a real community organisation, named publicly, rather than with an anonymous "Deaf advisory panel" of the kind NCG has found on vendor websites with no members listed. None of what follows disputes the need or questions anyone's motives. It concerns what an outsider can verify about the system that has been built to meet that need, and what a family is entitled to know before trusting it.

A grader is not a dictionary

Most of the sign language AI that NCG reviews produces signing: an avatar or a synthetic signer turns English into a signed language and a Deaf person watches the result. When that output is wrong, the Deaf person can usually tell. SignPlay runs the other way. Its AI watches someone sign and passes judgment on them. The person

being judged is a hearing adult who is learning, which means the one person in the room who cannot tell whether the judgment is right is the person relying on it.

The person most affected never touches the app at all. It is the deaf child. NCG made the same point last week about Thatsign, a British prototype aimed at the same families: a family tool is not low stakes, because the parent learns from the tool and the child learns from the parent [10]. SignPlay concentrates that risk further, because it does not just show the parent a sign. It certifies the parent's version of the sign as correct.

A grading app can be wrong in two different ways, and they do not cost the same.

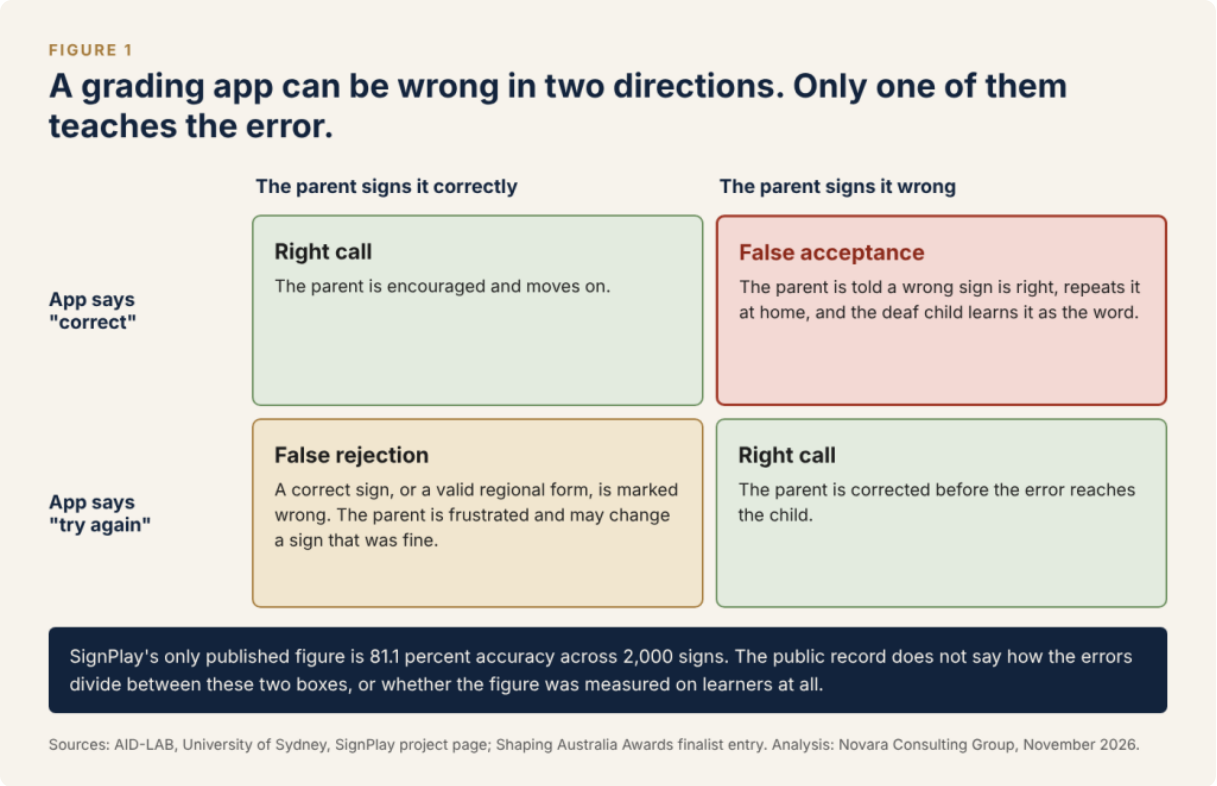

Figure 1. A grading app can be wrong in two directions. A false acceptance teaches the error; a false rejection frustrates the learner. SignPlay's published accuracy figure does not say how its errors divide.
Sources: AID-LAB project page; Shaping Australia Awards finalist entry. Analysis: Novara Consulting Group, November 2026.

A false rejection, where the app tells a parent a correct sign is wrong, is annoying. It might push a parent to change a sign that was fine, or to give up on a sign altogether. A false acceptance, where the app tells a parent a wrong sign is right, is worse. The

parent now has a machine's confirmation that the error is correct. They will use it at the high chair with confidence, and their child, who has no other model at home, will learn it as the word.

Any honest account of a grading tool's accuracy has to say how its errors split between those two boxes. A teacher who marks every answer correct has a perfect record on correct answers and is useless. A system tuned to be encouraging, which is a reasonable design choice for a game aimed at nervous beginners, will tend to accept more borderline signs. That might be the right trade for motivation. It might be the wrong trade for a child's first language. Either way, it is a trade a family deserves to know about.

Eighty-one percent of what?

SignPlay has published one performance number. The project page reports that the team collected over 246,000 data points and that its current model reaches 81.1 percent accuracy across 2,000 distinct signs [1].

That is a real achievement as a research result. Two thousand signs is a large vocabulary for isolated sign recognition, and a model that picks the right sign four times out of five across that range is not trivial to build. But as evidence that the app grades parents correctly, the number is missing nearly everything a reader would need.

It does not say who was in the test set. If the model was tested on fluent Deaf signers, the figure describes how well it recognises fluent signing. The app's actual job is to judge beginners, whose signing is by definition off in ways a fluent signer's is not, and nothing published describes performance on learners. It does not say how many people were tested, or whether any of them were also in the training data, which flatters a score. It does not break the figure down by sign, so a parent cannot know whether the signs they practise most are among the ones the model handles well. And it does not separate the two kinds of error in Figure 1.

There is a further problem in the shape of the number itself. "81.1 percent accuracy across 2,000 signs" reads like a recognition score: show the model a clip, and it names the sign correctly about four times in five. Grading is a different task. When a parent is asked to sign MILK, the app does not need to pick MILK out of 2,000 options. It needs to decide whether this attempt at MILK is good enough. A model can be decent

at the first task and poorly calibrated at the second, and a recognition score cannot tell you which. Under the SLAT Index, NCG's evaluation standard for sign language AI, a figure without its denominator and test conditions is not treated as evidence of performance, however precise it looks [13]. Precision to one decimal place is not the same thing as information.

Hands are not the whole sign

The awards entry shows SignPlay's practice screen and describes it as giving live hand-tracking feedback [3]. That phrase deserves a careful read.

Auslan, like every signed language, does not live only on the hands. Raised or lowered eyebrows mark the difference between a statement and a question. Mouth patterns sharpen or change meaning. Eye gaze points to people and places set up in the signing space. A headshake can carry negation on its own. When NCG looked at Thatsign last week, its team had already paused production after its own Deaf reviewers flagged missing facial grammar [10]. Missing faces are among the most common failures in machine signing, and the same blind spot applies when the machine is the one grading.

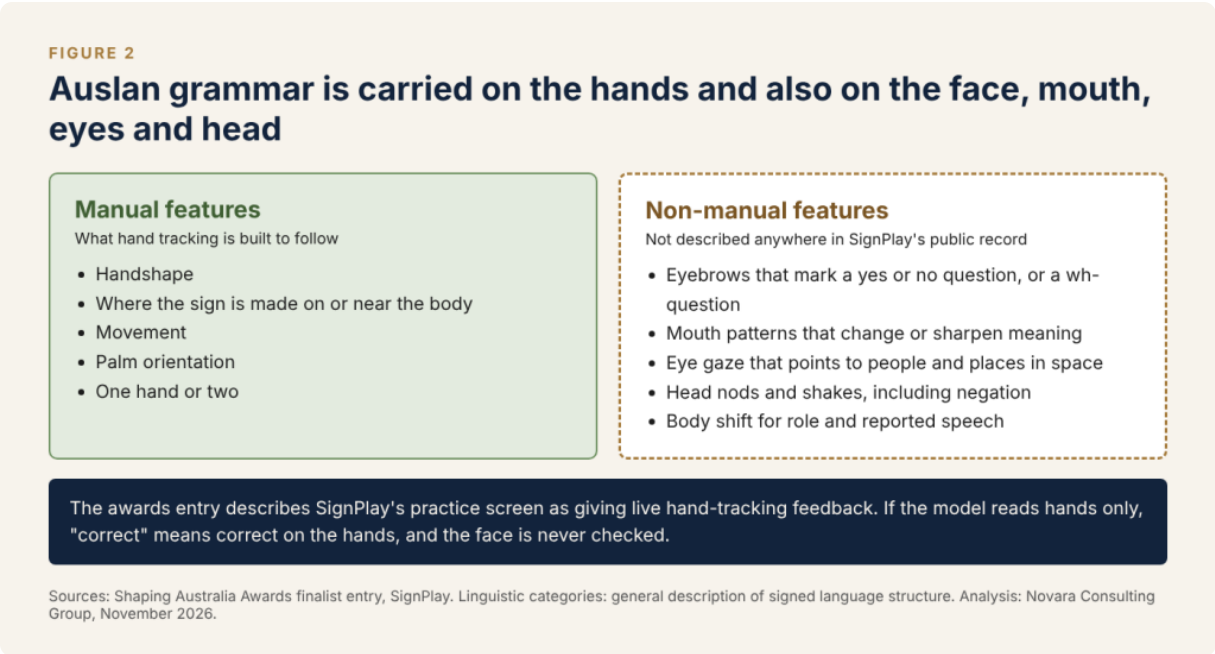

Figure 2. If a model reads hands only, a sign can be marked correct with its facial grammar missing or wrong. SignPlay's public record does not say whether the model observes the face.
Source: Shaping Australia Awards finalist entry. Analysis: Novara Consulting Group, November 2026.

Nothing published says whether SignPlay's model observes the face at all. If it reads hands only, then "correct" in SignPlay means correct on the hands. A parent could produce the right handshape, location and movement for a question while signing it with a statement face, and the app would show the same green tick. For single vocabulary items in a beginner game this may matter less. As SignPlay grows toward phrases, and toward the everyday communication families actually need, it matters a great deal.

Regional variation raises a second question. Auslan has recognised regional differences, commonly described as Northern and Southern dialects, and established learning resources teach both [9]. A family in Brisbane and a family in Melbourne may be taught different, equally correct signs for the same thing. The public record does not say which variety SignPlay accepts as correct, or what happens when a parent signs a valid regional form the model was not trained to expect. If the answer is that it marks it wrong, the app is quietly standardising a language its Deaf community has not asked to have standardised, one green tick at a time.

A quarter of a million videos, described by size alone

The dataset behind SignPlay is the project's proudest claim. The awards entry describes it as more than 250,000 videos covering 2,000 Auslan signs, collected with members of the Deaf community, and calls it "the largest real-world Auslan dataset" assembled for learning and recognition [3]. The project page describes the same collection as "over 246,000 data points" [1]. Those are different units and different counts, and neither source explains the difference.

The "largest" claim holds only with its qualifier attached. MM-WLAuslan, a dataset from another Australian team published at NeurIPS in 2024, holds over 282,000 sign videos across 3,215 Auslan glosses from 73 signers [7]. It was recorded in a studio, so SignPlay's claim rests on the words "real-world". That is a fair distinction to draw, but it is worth noticing that the published comparison point tells you how many signers were filmed and SignPlay's does not.

That is the larger gap. Every one of those videos shows a person's face, hands and body. In a community the size of Australia's signing Deaf community, a face in a signing video is close to a name. NCG has argued for months that the first questions to ask about any sign language AI are whose signing built it, what they agreed to,

whether they were paid, and whether they can take their faces back out [11]. For SignPlay, the public record answers none of them. It does not say how many people contributed, how many were Deaf, how they were recruited, whether consent was explained in Auslan, how long the video is kept, whether a contributor can withdraw, or whether the trained model forgets them if they do.

Those questions are answerable, because other projects answer them. A QUT team collecting Auslan for Sydney Trains this year said in its public call for participants that it would pay signers for their time [8]. PopSign, SignPlay's own partner at Georgia Tech, publishes that its American Sign Language dataset came from 47 Deaf signers who use ASL as their first language, recruited through the Deaf Professional Arts Network [6]. SignPlay could publish the same facts about its own collection tomorrow.

The partner list makes one further question unavoidable. Google.org is named as a partner, and Google has its own sign language recognition work, including the SL2T technology NCG reviewed in August [12]. NCG has found nothing to suggest SignPlay's data has gone anywhere it should not, and this is not an accusation. It is a question the project should be able to answer in one sentence: is any SignPlay video, derived data or trained model shared with Google.org, Georgia Tech or anyone else, and can it be used commercially? A family filming themselves at the kitchen bench, and a Deaf contributor who lent their signing to a university, both deserve that sentence in writing.

The same applies to the app itself. Whether learner video is processed on the phone or sent to a server is not stated anywhere NCG could find. In a family home, a phone pointed at a signing parent will often catch a child in the frame.

Input is not authority

The awards entry says SignPlay was co-designed with families and with Deaf community input, and that Deaf Connect gave ongoing guidance on design [3]. That is more than many projects can say. But input and guidance describe Deaf people being consulted. They do not describe Deaf people deciding.

In a grading app, the decision that matters most is a linguistic one: what counts as a correct sign. That decision sets the model's thresholds, decides which regional forms are accepted, and determines when a parent sees a green tick. The public record does



not say who makes it. The AID-LAB project page lists eight researchers and does not identify any of them as Deaf [1]. Deaf people appear on the public record as data contributors, as a testimonial writer, as a research student, and as the partner organisation's staff. Nowhere does it say that a Deaf person holds authority over what the app calls correct, or that any design decision changed because a Deaf person objected.

That pattern, Deaf participation arranged around a project rather than placed inside its decisions, is the most common governance shape NCG finds across sign language AI. It may well be that Deaf Connect's role is larger than the public record shows. If so, the fix is simple: say who decides, and show one decision their input changed.

Voting for an app nobody can download

There is one more gap, and it is the oddest. NCG could not find SignPlay.

There is no listing on the Apple App Store or Google Play under that name. There is no privacy policy and no terms of use on the AID-LAB site, no published paper on SignPlay, and no code for it on the lab's GitHub page, which lists the lab's other projects [2]. The awards entry gives no number of families using the app, no learning outcomes, and no evaluation, and its impact language is forward-looking: SignPlay has "the potential to improve outcomes for families across Australia" [3].

Deaf Connect's own words point both ways. Its LinkedIn post says SignPlay "offers" game-based activities, in the present tense, and asks people to vote [5]. Its supporting statement in the awards entry says families "will be able to" practise at home [3]. The awards page and the post also frame the app slightly differently. One describes practice between lessons. The other describes families learning Auslan at their own pace, which reads, to a parent in Dubbo on a three-month wait, very much like a substitute for lessons.

None of this is unusual for a university project in its second year, and an awards nomination is not a product launch. But a public vote is being asked for, on the strength of a tool the public cannot inspect, and a parent reading that post could reasonably go looking for an app to download tonight. The time to publish how it is tested, what data it keeps, and what it cannot judge is before families start relying on it, not after.



What a school, a funder or a family should ask

If SignPlay reaches families in 2027, it will reach them through organisations: early intervention services, schools, hospitals' newborn hearing programmes, and charities paying for licences. Those organisations are the ones with the standing to ask, and NCG's SLAT Index exists to give them the questions [13]. Checked against the public record today, most of them have no answer yet.

FIGURE 3

Ten questions, checked against SignPlay's public record

Question	Public record
Who is it for, and what problem does it address?	Stated
How accurate is it, measured how, on whom?	Partly: one figure, no method
Split between false acceptance and false rejection	Not stated
Does it read the face, mouth, eyes and head?	Not stated
Which regional varieties count as correct?	Not stated
How many signers made the dataset, and how many are Deaf?	Partly: "Deaf community", no count
Consent, payment, retention and withdrawal for contributors	Not stated
Is data shared with partners, or usable commercially?	Not stated
Is learner video kept on the phone or sent to a server?	Not stated
Can families download it today, and where are its terms?	Not stated

Not stated means absent from the public sources listed, not proof that no answer exists.

Sources: AID-LAB project page; Shaping Australia Awards finalist entry; PopSign partner page; AID-LAB GitHub; app store searches. All viewed November 25, 2026.

Figure 3. Ten questions checked against SignPlay's public record. Not stated means absent from the public sources listed, not proof that no answer exists.
Sources: AID-LAB project page; Shaping Australia Awards finalist entry; PopSign partner page; AID-LAB GitHub; app store searches. All viewed November 25, 2026.

Before relying on SignPlay's verdicts, a buyer should ask for written answers to these:

1. How was the 81.1 percent figure produced: how many test signers, were they Deaf or hearing, fluent or learners, and were any of them also in the training data?
2. Of the errors, what share are false acceptances, where a wrong sign is marked correct, and what share are false rejections?

3. Does the model observe the face, mouth, eye gaze and head, or the hands only? If hands only, how is that disclosed to families inside the app?
4. Which regional varieties of Auslan does the app accept, and what happens when a parent signs a valid form the model was not trained on?
5. Is the dataset 246,000 data points or more than 250,000 videos, how many individual signers contributed, and how many of them are Deaf?
6. What did contributors consent to, was the consent explained in Auslan, were they paid, and can they withdraw their video and have it removed from trained models?
7. Is any SignPlay video, derived data or model shared with Google.org, Georgia Tech, PopSign or anyone else, and may it be used commercially?
8. Is learner video processed on the phone or sent to a server, and what happens to footage that includes children?
9. Who holds the authority to decide what counts as a correct sign, and do Deaf people hold it?
10. Is the app available to families today? If so, where are its privacy policy and terms; if not, when, and how many families are in any pilot?
11. Is an independent evaluation planned, rated by fluent Deaf signers, comparing children's home language input from families who used SignPlay with families taught by people alone?

The green tick has to be earned

The problem SignPlay is trying to solve is real, and it is urgent in the way only early childhood problems are. A deaf child's first years do not wait for a course place to open up. A tool that helps a parent practise between lessons, built with a Deaf community organisation and trained on real signing in real homes, could do genuine good.

But the thing SignPlay adds to the existing Auslan apps, the thing that makes it AI, is a judgment. It tells a hearing parent that their signing is right. That judgment will be trusted most by the people least able to check it, and its mistakes will be inherited by



children who cannot check it at all. A video course that shows a sign badly is a weak teacher. An app that tells a parent their wrong sign is correct is a confident one.

Before that mother in Dubbo trusts the green tick on her kitchen bench, someone should be able to tell her how often it is wrong, in which direction, on whose signing it was tested, and which Deaf people decided what the tick means. Right now, the public record cannot.

参考文献

1. AID-LAB, University of Sydney. (2026). SignPlay: Learning Australian Sign Language through interactive mobile games. Viewed November 25, 2026. https://www.aid-lab.org/research/signplay
2. AID-LAB. (2026). aid-lab-org repositories on GitHub. Viewed November 25, 2026. https://github.com/aid-lab-org
3. Universities Australia. (2026). SignPlay: Learning Auslan through an app (University of Sydney). Shaping Australia Awards 2026 finalist, Community Champion Award. Viewed November 25, 2026. https://www.shapingaustraliaawards.com.au/finalists/signplay-learning-auslan-through-an-app-university-of-sydney
4. Universities Australia. (2026). Shaping Australia Awards: The finalists. Viewed November 25, 2026. https://www.shapingaustraliaawards.com.au/the-finalists
5. Deaf Connect. (2026, November). LinkedIn post on SignPlay′s nomination in the Shaping Australia Awards. https://www.linkedin.com/company/deaf-connect-of-australia/
6. PopSign, Georgia Institute of Technology. (2026). PopSign: American Sign Language learning app, partners. Viewed November 25, 2026. https://www.popsign.org/
7. Shen, X., Du, H., Sheng, H., Wang, S., Chen, H., Chen, H., Wu, Z., Du, X., Ying, J., Lu, R., Xu, Q., and Yu, X. (2024). MM-WLAuslan: Multi-view multi-modal word-level Australian Sign Language recognition dataset. Advances in Neural Information Processing Systems 37, Datasets and Benchmarks Track. https://arxiv.org/abs/2410.19488

8. Deaf Australia. (2026). Sydney Trains Auslan data collection, April 13 to 14, 2026. https://www.deafaustralia.org.au/events/sydney-trains-auslan-data-collection
9. NextSense. (2021). Auslan Tutor (app listing, including Northern and Southern dialects). Apple App Store. https://apps.apple.com/app/id307217026
10. Grizzle, H. M. (2026, November 20). Thirteen signs at the dinner table: What Thatsign gets right, and what a family sign language tool owes a deaf child. Novara Consulting Group. https://www.novaracg.com/2026/11/20/thirteen-signs-at-the-dinner-table-what-thatsign-gets-right-and-what-a-family-sign-language-tool-owes-a-deaf-child/
11. Grizzle, H. M. (2026, November 5). The Corpus Question, part two: When the signer is the supplier. Novara Consulting Group. https://www.novaracg.com/2026/11/05/the-corpus-question-part-two-when-the-signer-is-the-supplier/
12. Grizzle, H. M. (2026, August 23). The coordinates of trust: A governance reading of Google DeepMind's SL2T. Novara Consulting Group. https://www.novaracg.com/2026/08/23/the-coordinates-of-trust/
13. Grizzle, H. M. (2026). SLAT Index Evaluation Standard, version 2.0. Novara Consulting Group. https://doi.org/10.5281/zenodo.21537271

> **回应权**
>
> NCG did not contact the University of Sydney, AID-LAB, Deaf Connect, Georgia Tech or Google.org before publication. This analysis relies on the public sources listed above, all viewed November 25, 2026. NCG will put the questions in this piece to AID-LAB and Deaf Connect and will publish any response in full.

Tips

NCG welcomes confidential tips about sign language AI deployments, procurement decisions and vendor practices, including how training data is collected, how signers are paid and how products are tested, at tips@novaracg.com. For readers in the United Kingdom: writing to NCG is not, by itself, a protected disclosure under the

Public Interest Disclosure Act 1998. If you are considering whistleblowing about your employer, independent advice is available from Protect at protect-advice.org.uk.

> **SUPPORT NCG**
>
> NCG does this work without vendor money, so our findings stay ours. If it has been useful to you, you can [support NCG's independent work here](https://buy.stripe.com/8x2bIV5qH7PP580aeu1wY00). NCG is a for-profit company, not a charity. Contributions are not tax-deductible and buy no service or influence.

引用方式

Grizzle, H. M. (2026, November 25). SignPlay Tells Hearing Parents When Their Auslan Is Right. It Has Not Said How Often the Tick Is Wrong.. Novara Consulting Group. https://www.novaracg.com/zh/2026/11/25/signplay-tells-hearing-parents-when-their-auslan-is-right-it-has-not-said-how-often-the-tick-is-wrong/