A CONSEQUENCE ASSESSMENT FRAMEWORK FOR SITE-SPECIFIC ROAD AVALANCHE WARNING IN NORWAY

Markus Eckerstorfer[1*], Emil Solbakken[1], Halgeir Dahle[1], Tore Humstad[1], Njål Farestveit[1], Jens Øyen Tveit[1], Eivind Schnell Juvik[1], Dag Theodor Røsvik Andreassen[1]

[1] *The Norwegian Public Roads Administration (NPRA), Norway*

ABSTRACT: Site-specific avalanche warning for roads is a risk-reducing measure that combines an assessment of avalanche likelihood with an impact-based assessment of consequence for the road — whether a given avalanche reaches the road, and what happens to traffic if it does. This consequence dimension is rarely addressed as an explicit, repeatable methodology; site-specific warning has typically relied on qualitative, case-by-case judgment instead. We present a structured, path-specific consequence assessment framework, separating a static preseason consequence matrix per avalanche path from a small set of daily dynamic inputs, combined in an operational likelihood and risk matrix. The framework introduces specific runout length, relative runout length, and consequence classes as methodological contributions to site-specific road avalanche warning. Applied to Rv. 13 Vikafjellet in 2025/26 and two additional corridors since 2024/25, the framework governed daily alert-level decisions alongside active avalanche control, yielding a residual risk comparable to established control programs internationally. We discuss current limitations — particularly the reliance on expert judgment for terrain, traffic scoring, and design avalanches — and outline how the framework could be made more systematic and automated.

KEYWORDS: site-specific avalanche warning, consequence assessment, road avalanche risk, residual risk, impact-based forecasting

* Corresponding author address: Markus Eckerstorfer, The Norwegian Public Roads Administration, Norway; email: markus.eckerstorfer@vegvesen.no

1. INTRODUCTION

Site-specific avalanche warning, a risk-reducing measure for roads exposed to snow avalanches, is gaining traction in Norway. It gives the road owner a decision basis for keeping roads open when safety is sufficient, conducting preventive avalanche control when effective, closing roads preventively when risk is unacceptable, and reopening roads safely after an event.

Site-specific avalanche warning for roads draws on established hazard frameworks: the EAWS matrix combines snowpack stability, spatial distribution, and avalanche size into a danger level (Müller et al., 2025), while the Conceptual Model of Avalanche Hazard (CMAH) structures this assessment into a repeatable forecasting workflow (Statham et al., 2018). Both are, by design, independent of the element at risk and provide no methodology for evaluating consequence for a specific road or avalanche path.

We address this gap with a structured, path-specific consequence assessment framework for site-specific road avalanche warning. The framework introduces explicit classifications for avalanche runout and road impact, combined with terrain and traffic parameters, to produce repeatable consequence assessments for each avalanche path. We report on its design and first-season operational use on Rv. 13 Vikafjellet, including a residual-risk evaluation of the active avalanche-control program.

2. SITE-SPECIFIC ROAD AVALANCHE FORECASTING AT NPRA

2.1 <u>Remote forecasting and alert levels</u>

The Norwegian Public Roads Administration (NPRA) currently operates site-specific avalanche warning and active avalanche control using avalanche towers on three mountain passes in southern Norway: Rv. 15 Strynefjellet (since 2021/22), E134 Haukelifjell (since 2024/25), and Rv. 13 Vikafjellet (since 2025/26).

NPRA uses remote forecasting, basing their assessment on sensor and model data and field observations by local observers. Four alert levels (Figure 1), each tied to a predefined set of measures for road operation and traffic handling, are issued in a bulletin following an impact-based forecasting approach — translating hazard into expected consequences for road users (WMO, 2021). Use of active avalanche control is not uniquely assigned to an alert level; it may addi-

https://doi.org/10.15788/1790098548

tionally be recommended when suitable. Forecasts are issued for 6-hour timesteps within a validity period of one to seven days, updated whenever the current forecast expires or conditions change materially.

Alert level		Road operation	Traffic handling
Red	Critical	No stopping or presence at avalanche path	Road closed
Orange	Demanding	Stay in vehicle only at avalanche path	Convoy travel with risk assessment; traffic not stopped in avalanche path
Yellow	Moderate / uncertain	Limit time outside vehicle	Convoy travel through avalanche path; road remains open
Green	Normal	Normal operation	Open road

Figure 1: Alert levels and associated road operation and traffic measures (simplified from the internal alert level action table).

2.2 Study area: Rv. 13 Vikafjellet

The Rv. 13 Vikafjellet mountain pass in Vestland county, Norway, encompasses sixteen avalanche paths grouped into three sub-areas: Kvassdalen, Skjelingavatnet-Ygnisdalen (split up for the residual risk analysis), and Skredene.

Terrain and snowpack characteristics vary considerably between sub-areas. Low-lying Kvassdalen is dominated by wet slab and wet glide avalanches capable of D5 magnitude. Skjelingavatnet–Ygnisdalen and Skredene have small vertical drops (50–150 m), with the road running through the starting zones. The 2025/2026 season was distinctly milder than the 1991–2020 normal, with a pronounced dry winter followed by a wet spring.

3. A CONSEQUENCE ASSESSMENT FRAMEWORK FOR AVALANCHE PATHS

The framework is structured around two layers of assessment: a static layer, built once per avalanche path during preseason planning and revisited only when path characteristics change, and a dynamic layer, re-evaluated at every forecast during the operational season (Figure 2).

3.1 Static layer: Preseason assessment of consequence matrices

Design avalanches: For each avalanche path, every destructive size (EAWS, n.d) is assigned a specific runout length class describing whether and how it affects the road under a given relative runout length. Design avalanches are established using historic avalanche records and, to some extent, dynamic avalanche simulations.

Specific runout length is categorized into five classes, capturing if and how much the road is eventually affected by an avalanche (Figure 3). Relative runout length captures how snow conditions in the avalanche track shorten or lengthen the runout of a given destructive size (Figure 3). The latter is an essential part of the daily forecasting, allowing forecasters to account for varying specific runout lengths from avalanches of similar destructive size.

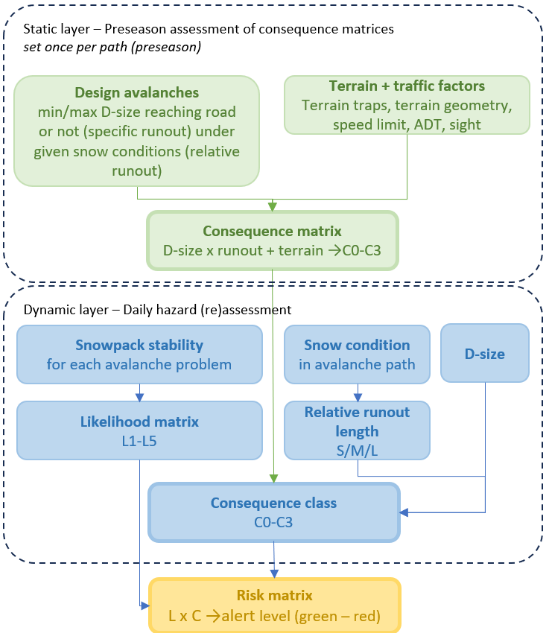

Figure 2: Logic of the consequence assessment framework for site-specific avalanche forecasting for roads.

Specific runout length		Description
CR	Cross road	Solid avalanche mass crosses the entire road
RR	Reach road	Solid avalanche mass reaches part of the road
SCR	Snow cloud	Powder cloud may reach the road; solid mass typically does not
NRR	Not reach road	Avalanche does not reach the road; sets a minimum size of concern for the path
EX	Excluded	Path is judged unable to produce an avalanche that reaches the road; risk excluded
Relative runout length		**Description**
L	Long	Large snow volume in the avalanche path; dry and soft/loose snow surface, or a hard snow surface
M	Medium	Normal snow volume in the avalanche path, mostly dry snow with varying firmness
S	Short	Little (or no) snow volume in the avalanche path and/or wet snow

Figure 3: Specific and relative runout length classes which are linked to destructive avalanche sizes to create path-specific consequence matrices.

Terrain and traffic factors: In addition to the size-runout relationship, each avalanche path is evaluated on a set of static terrain factors — path geometry (e.g., incline and confinement of the track and runout), exposure to terrain traps, and the position of the road within the avalanche path (starting zone, track, or runout zone). A set of static traffic factors is assessed alongside these: speed limit, annual average daily traffic (ADT), sight distance, and road curvature leading through the avalanche path. These terrain and traffic factors either increase or decrease the consequence score set by the size-runout relationship.

Consequence classes: Four consequence classes (C0-C3) describe the expected impact on vehicles and traffic if an avalanche reaches the road (Figure 4). The classes reflect both the physical impact of avalanche debris on vehicles and passengers and the influence of path-specific terrain and traffic exposure.

Consequence		Description	Extent of harm
C3	Large	Vehicles may be completely buried and/or displaced from the roadway. High risk, particularly where the road runs along a fjord/river or in steep terrain.	Serious injury or death
C2	Medium	Vehicles may be partially buried and/or displaced from the roadway, but avalanche speed is typically low and/or the surrounding terrain has no terrain traps.	Injury
C1	Minor	Vehicles may become stuck; low risk of serious collision with other traffic due to low speed.	No injury
C0	None	Vehicles are normally unaffected.	No damage

Figure 4: Consequence classes (C0-C3) for site-specific road avalanche warning, defined by the expected impact on vehicles and the extent of harm.

Consequence matrix: All outlined consequence factors are finally set up in a path-specific consequence matrix (Figure 5). For each destructive size class, the matrix records the expected specific runout length under representative snow conditions and whether terrain and traffic factors raise or lower the resulting consequence class.

This relationship is only partly rule-based: the framework connects size and runout to a base-consequence class, but the path-specific terrain and traffic adjustment relies on subjective judgment rather than defined scoring rules.

Some rules hold without exception: a specific runout class of NRR (Not reach road) or EX (Excluded) always carries consequence class C0. Consequence classes are always non-decreasing from short to long relative runout at a fixed destructive size. Class EX is mostly reserved for large D-sizes, because many of our starting zones simply cannot generate avalanches that large.

The NRR class follows a related but graded pattern: it accounts for roughly 75% of assignments at D1, falling to 38% at D2, 13% at D3 before disappearing at D4 and D5, due to uncertainty involving large to very large avalanches. Additionally, if D4 or D5 avalanches are possible but cannot reach the road, we would not have forecasted for that particular avalanche path. Finally, even when an avalanche does cross the road (CR), the resulting consequence class ranges freely from C1 to C3 across paths, since the actual severity depends on each path's terrain traps, road position, and traffic exposure rather than on size or runout in isolation.

Kvassdalen 4	D1	D2	D3	D4	D5
S	NRR C0	NRR C0	CR C2	CR C3	CR C3
M	NRR C0	NRR C0	CR C2	CR C3	CR C3
L	NRR C0	RR C1	CR C3	CR C3	CR C3

Skjelingavatnet 1	D1	D2	D3	D4	D5
S	RR C1	CR C2	CR C3	EX C0	EX C0
M	RR C1	CR C2	CR C3	EX C0	EX C0
L	RR C2	CR C3	CR C3	EX C0	EX C0

Figure 5: Examples of path-specific consequence matrices for Kvassdalen 4 and Skjelingavatnet 1 avalanche paths, combining destructive size (D1-D5) and relative runout distance (S/M/L) into specific runout class (CR, RR, NRR, EX) and consequence classes (C0-C3). Marked in yellow are historic avalanche events that function as design avalanches.

Consequence matrices for two avalanche paths exemplify the idea (Figure 5): Kvassdalen 4 can produce D4 and D5 avalanches under a variety of snow conditions in the avalanche path (relative runout length: S/M/L), all of which cross the road (specific runout length: CR), resulting in consequence class C3. Skjelingavatnet 1, on the other hand, cannot produce D4 or D5 avalanches (EX: excluded → C0); however, D1 avalanches with short relative runout length will reach the road (RR) and trigger class C1.

Avalanche examples across consequence classes: Photo examples from Rv. 13 Vikafjellet (Figure 6) illustrate how terrain and traffic exposure shape consequence beyond destructive size and runout. In (a), a D3-CR avalanche reaches only C2 — the flat, open downhill side disperses debris without a terrain trap. Panel (b) shows D2 avalanches reaching the road where a winding alignment and high snowbanks increase burial risk, giving C2. Panel (c) has the same D3-CR combination as (a), but a steep terrain trap on the downhill side pushed a snowplow off the road, illustrating why this path is assigned C3. Panels (d) and (e) both show D1-RR avalanches with short runout yet differ in consequence (C2 vs. C1): in (d), a curve combined with high vehicle speed accounts for the higher class. Panel (f) shows D2 avalanches that do not reach the road (NRR), resulting in consequence class C0.

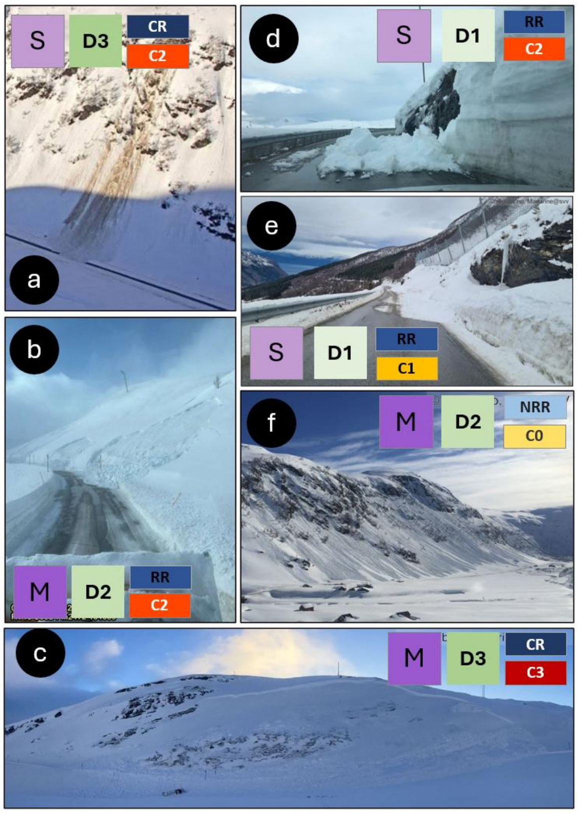

Figure 6: Photo examples of avalanches on Rv. 13 Vikafjellet spanning the range of consequence classes (C0–C3), illustrating how destructive size, runout, and path-specific terrain and traffic exposure combine to determine consequence.

3.2 Dynamic layer: Daily hazard (re)assessment

Daily hazard assessment workflow: Daily hazard assessments run through purpose-built forecasting software in the RESPONS map toolkit since season 2025/26 (Orset and Frekhaug, 2024). For each active avalanche problem, the forecaster enters snowpack stability, destructive avalanche size, and relative runout length. The software then issues the highest alert level per avalanche problem and avalanche path. The highest alert levels per sub-area are compiled into a bulletin for road owners describing weather and snowpack conditions, with a written hazard assessment and recommendations for preventive avalanche control.

The example in Figure 7 walks through the three steps behind this. First, for a wet slab problem on one of the Kvassdalen avalanche paths, the forecaster rates snowpack stability on a likelihood matrix that combines sensitivity to triggering with spatial distribution. The matrix is a variant of the EAWS matrix, refined from ADAM (Müller et al., 2016) with input from Thumlert et al. (2024). The combination of poor snowpack stability and "some" trigger points gives L3: Possible.

Second, the forecaster derives a consequence class from the path-specific consequence matrix, using destructive size and relative runout distance. Destructive size 3 with short relative runout, reflecting wet snow in the track, is expected to cross the road (CR). At Kvassdalen 4, the 80 km/h speed limit, good visibility from a distance, and no terrain trap on the downhill side result in consequence class C2.

Third, L3 likelihood and C2 consequence classes are combined in the risk matrix to obtain alert level "Yellow".

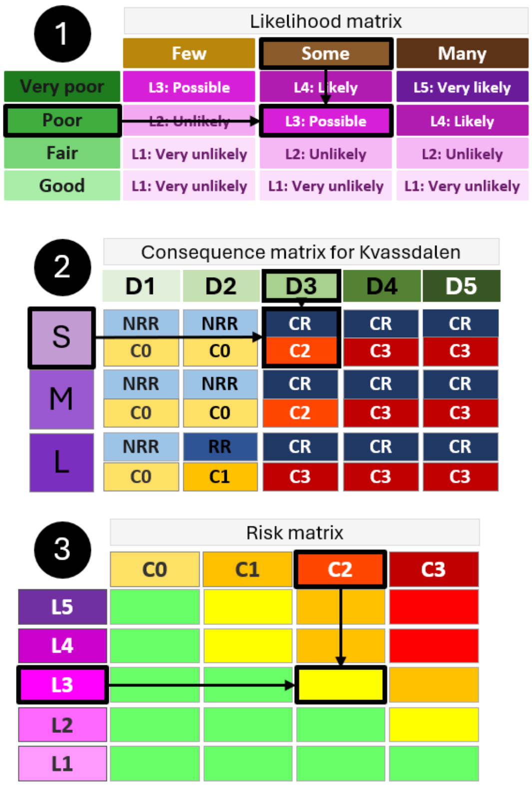

Figure 7: Worked example of the hazard assessment workflow for a wet slab problem on Kvassdalen 4, resulting in alert level "Yellow".

4. MATRIX USE ACROSS THE SEASON

Of 1,276 classified avalanche problem assessments, "Possible" likelihood (S3) was the most frequent rating (47%). S3×C0 was the single most used combination (20%), followed by S3×C1 (16%) and S3×C2 (11%). S5 (Very likely) occurred in fewer than 1.5% of assessments, always paired with low consequence, and the S5×C3 combination never occurred (Figure 8a).

In the consequence matrix (Figure 8b), destructive size was the primary driver of consequence class — C-class increased consistently from D1 to D3 regardless of relative runout length, and no assessment reached D4 or D5. Short relative runout combined with D1 was the most frequent entry, reflecting the corridor's predominantly small avalanche activity.

The dominance of S3 and near-absence of S4/S5 is consistent with a season driven by wet snow and persistent weak layer problems producing moderate rather than extreme triggering sensitivity (Figure 8c). The absence of S5×C3 indicates that the corridor never simultaneously experienced widespread instability and high-consequence conditions — the most critical scenario the framework is designed to detect.

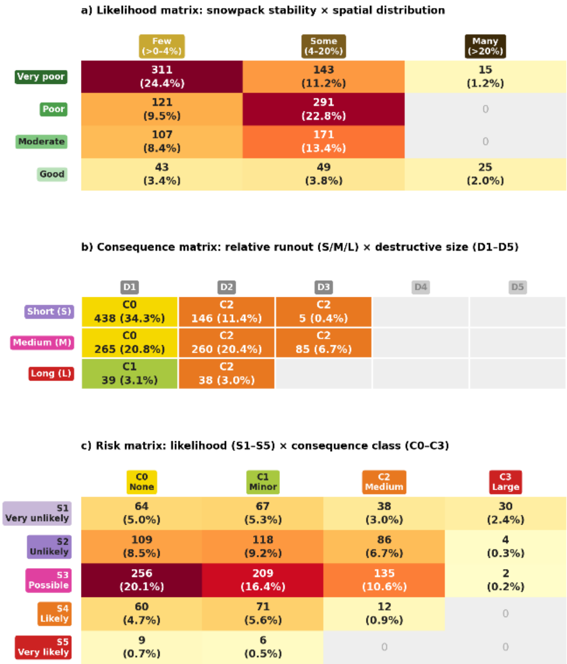

Figure 8: Matrix use across the 2025/26 season (n=1,276 classified avalanche problem assessments). (a) Likelihood matrix: assessed combinations of snowpack stability and spatial distribution of trigger points. (b) Consequence matrix: assessed combinations of relative runout length and destructive size; cell color indicates the dominant consequence class (C0–C3). (c) Risk matrix: assessed combinations of likelihood class (S1–S5) and consequence class (C0–C3). Numbers show count and percentage of total assessments per cell.

5. SEASON 2025/26 CHARACTERIZATION

Across the season, 64 avalanches were registered: 50 controlled, using the Wyssen avalanche towers, and 14 uncontrolled, comprising 13 natural releases and one avalanche incidentally triggered by a snowplow (Jan 1) (Figure 9). Alert levels varied substantially by path: Elvakroken stayed “Green” all season, Halsabakkane and Kvassdalen were elevated on only 6% and 8% of days, Skjelingavatnet and Skredene on 24% and 20%, and Ygnisdalshamrane — the most active path — on 44%, with all four “Orange” episodes concentrated in late February and March.

Wet snow (447 path-days — summed across all paths and days on which the problem was assessed as active) and persistent weak layer (378) were the season’s dominant avalanche problems, well ahead of new snow (245) and wind slab (120). A persistent weak layer of facets around a crust remained active through January–March before disappearing by April. Wet snow was on and off active from the season’s opening days through to closure in May, spiking sharply in March. The two problems co-occurred heavily through February and March, which explains why the season’s “Orange” episodes clustered in that window.

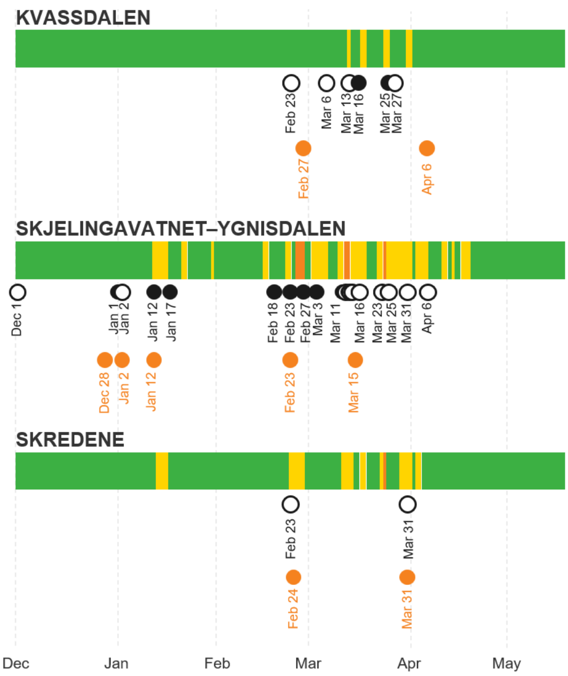

Figure 9: Daily alert level for the three sub-areas on Rv. 13 Vikafjellet in the 2025/26 season. Dots below each alert level band mark avalanche control (filled = release, hollow = no release) and natural avalanches (orange).

Alert levels tracked closely with weather-driven instabilities, peaking in late February and March during repeated rain-on-snow events. On 23 February, 9 avalanches and 14 control shots were recorded under "Yellow" — one of the busiest days of the season. Four days later, 13 mm of rain pushed Skjelingavatnet–Ygnisdalen to "Orange", with 6 avalanches recorded. A second cluster in mid-to-late March saw 11 shots triggering 7 avalanches on 16 March and 7 shots triggering 6 on 25 March, both under "Yellow" (Figure 9).

Not every elevated alert was matched by activity. On four days in early March — 6, 12, 14, and 27 March — control shots fired under "Yellow" or "Orange" conditions produced no avalanche. On 24 March, both Skjelingavatnet–Ygnisdalen and Skredene reached "Orange" under 19.9 mm of rain — the season's heaviest single day — yet no avalanche was recorded anywhere on the corridor. Similar mismatches occurred on 25–26 February. Together with 23 February's opposite pattern, this points to calibration issues in both directions.

Road closures were limited: only two were logged this season (short 10-15 min closures were not logged) — 28 minutes on 1 January and 186 minutes on 23 February — totaling 3.6 hours, well under the 30-hour corridor average since 2021 (14 events, 421 hours). This is consistent with active avalanche control containing disruption rather than the road waiting out the hazard.

6. QUANTITATIVE RISK CONTEXT: RESIDUAL RISK AND ACTIVE PROGRAM EFFECTIVENESS

The consequence framework presented governs day-to-day alert-level decisions. To place these decisions in a broader risk-management context, and to evaluate the effectiveness of the active avalanche-control program, we conducted a season-level quantitative risk assessment for Rv. 13 Vikafjellet.

Residual risk quantifies the risk remaining on an open road despite implemented measures, following Hamre et al. (2016): Residual risk = $N_{unmitigated}$ / N_{total}, where N_{total} is all avalanches that reached the road regardless of trigger or road status, and $N_{unmitigated}$ is the subset of natural avalanches that reached an open road (i.e., not captured by the control program).

Hamre et al. (2016) report residual risk of 0.03–0.19 for active avalanche-control programs internationally, versus 0.25–0.88 for passive programs (warning and closure only).

For season 2025/26, avalanche event data were compiled from two sources: Regobs.no, for avalanche observations, and the avalanche tower activation log, for all avalanche-control shots. Of the 84 avalanche-control shots fired, 50 triggered an avalanche (release efficiency 59.5%) (Table 1).

Of 40 avalanche events that reached the road, 7 were natural avalanches occurring on an open road (unmitigated), giving a season residual risk of 7/40 = 0.175 — within the active-program range and comparable to Hospental, Switzerland (0.18) and Pimentón mine, Chile (0.19) in the Hamre et al. (2016) dataset. Table 1 breaks this down by sub-area.

Table 1: Residual risk and avalanche control effectiveness metrics per sub-area, season 2025/26. *Skredene's residual risk of 1.00 is based on only 2 road-hit events and should be treated with caution.

Sub-area	*Shots*	*Release eff.*	*Road-hit ratio*	*Total to road*	*Unmitigated*	*Residual risk*
Ygnisdalen	30	63%	0.633	22	3	0.14
Skjelinga-vatnet	32	59%	0.375	15	2	0.13
Kvassdalen	18	56%	0.056	1	0	0.00
Skredene	4	0%	0.000	2	2	1.00*
TOTAL	84	59.5%	0.381	40	7	0.175

Ygnisdalen and Skjelingavatnet fall within the active-program range; Ygnisdalen's high road-hit ratio (0.633) reflects the fact that the road leads through the starting zone, rather than inadequate control. Controlled avalanches will nearly always reach the road in these sub-areas. Skredene's residual risk of 1.00, based on only two road-hit events and zero successful releases from four shots, is statistically unreliable given the small sample. Kvassdalen's residual risk of 0.00 rests on one road-hit event only.

The metric has a blind spot: on 27 February, four natural avalanches released on Kvassdalen 2–5 while the road was open, but none reached the road, so none counts toward N_{total} or $N_{unmitigated}$. A marginally longer runout under slightly different snow conditions would have turned an uneventful day into four unmitigated hits.

Residual risk captures realized misses, not near-misses. We therefore treat this season's results as a first operational baseline, not a settled verdict on the program's effectiveness.

7. DISCUSSION

7.1 Scalability and transferability

We present a consequence framework for site-specific avalanche forecasting currently applied to three road corridors in Norway, with three additional projects under development in Northern Norway. The framework's core innovations — impact-based alert levels, path-specific consequence matrices combining specific and relative runout length and scored consequence classes — represent a structured departure from qualitative, case-by-case judgment. Making these elements explicit is a prerequisite for the systematic improvements described below.

7.2 *Methodology and calibration*

Terrain and traffic factors are currently folded into each path's matrix through expert judgment rather than explicit scoring rules. A more systematic approach would derive a rule-based base-consequence class from the avalanche's physical impact potential and adjust it using scored terrain and traffic factors, keeping room for path-specific adaptation while making each matrix auditable across forecasters.

A related limitation concerns the reproducibility of static input parameters — particularly specific runout length and consequence class assignments. Both rely on expert judgement during the preseason planning phase and different forecasters may interpret the same terrain differently. A shift in either parameter for a high-frequency path can materially change the season's alert level distribution. Mid-season adjustments are possible but should be carefully documented, and systematic review between seasons is essential.

Dynamic avalanche simulations — run across release scenarios and friction parameters — should ground design avalanche calibration more systematically. How to translate simulated impact pressure and runout into specific runout length classes remains an open question.

7.3 *Operational performance and communication*

The likelihood matrix is applied at the scale of a single avalanche path, based on CMAH being applicable to all temporal and spatial scales (Statham et al., 2018). We have encountered situations where we did not have enough information to set the spatial distribution of trigger point stability, which we resolved by directly assigning the most conservative likelihood class.

Because the daily workflow is explicit and rule-based, it is in principle well suited to automation — for instance, by mapping avalanche problems and stability ratings from an independent regional bulletin directly onto the likelihood and consequence matrices.

A final consideration concerns communication. Bulletins are issued to road owners and contractors as recommendations, and feedback suggests they are read and valued. Yet there is a real risk that increasingly formalized forecasts become an intellectual exercise rather than a tool for dialogue — and the framework's value depends as much on maintaining that exchange as on methodological rigor.

8. CONCLUSION

Our consequence assessment framework for site-specific road avalanche warning combines a complete, static preseason consequence matrix with a small number of daily dynamic inputs, integrated with an operational likelihood and risk matrix. Applied to Rv. 13 Vikafjellet over one season, it produced a residual risk comparable to established active avalanche-control programs internationally, while surfacing concrete next steps in terrain and traffic scoring, the physical grounding of design avalanches, and the reliability of daily likelihood judgments.

ACKNOWLEDGEMENTS

The site-specific avalanche warning projects are part of the NPRA's internal project: Technology for Natural Hazard Management. Thank you to all external partners involved.

REFERENCES

European Avalanche Warning Services (EAWS), n.d. Avalanche size. Available at: www.avalanches.org (visited: 20.8.2026)

Hamre, D., Greene, E. and Margreth, S. (2016). Quantifying the effectiveness of active mitigation on transportation corridors. Proceedings, International Snow Science Workshop, Breckenridge, Colorado, 2016, pp. 435–441.

Müller, K., Mitterer, C., Engeset, R., Ekker, R. and Kosberg, S.Ø., 2016. Combining the conceptual model of avalanche hazard with the Bavarian matrix. Proceedings of the International Snow Science Workshop, Breckenridge, CO, 472–479.

Müller, K., Techel, F., and Mitterer, C. 2025. The EAWS matrix, a decision support tool to determine the regional avalanche danger level (Part A): conceptual development. Natural Hazards and Earth System Sciences, 25, 4503–4525.

Orset, K.I. and Frekhaug, M.H., 2024. Natural hazard preparedness in the Norwegian Public Roads Administration (NPRA). Proceedings of the International Snow Science Workshop, Tromsø, Norway, 252–256.

Statham, G., Haegeli, P., Greene, E., Birkeland, K., Israelson, C., Tremper, B., Stethem, C., McMahon, B., White, B. and Kelly, J., 2018. A conceptual model of avalanche hazard. Natural Hazards, 90(2), 663–691.

Thumlert, S., Stefan, M. and Langeland, S., 2024. Assessing and communicating likelihood and probability of snow avalanches. Proceedings of the International Snow Science Workshop, Tromsø, Norway, 1609–1617.

World Meteorological Organization, 2021. WMO Guidelines on Multi-hazard Impact-based Forecast and Warning Services, Part II: Putting Multi-hazard IBFWS into Practice. WMO-No. 1150. Geneva: WMO.