Railway Gazette INTERNATIONAL

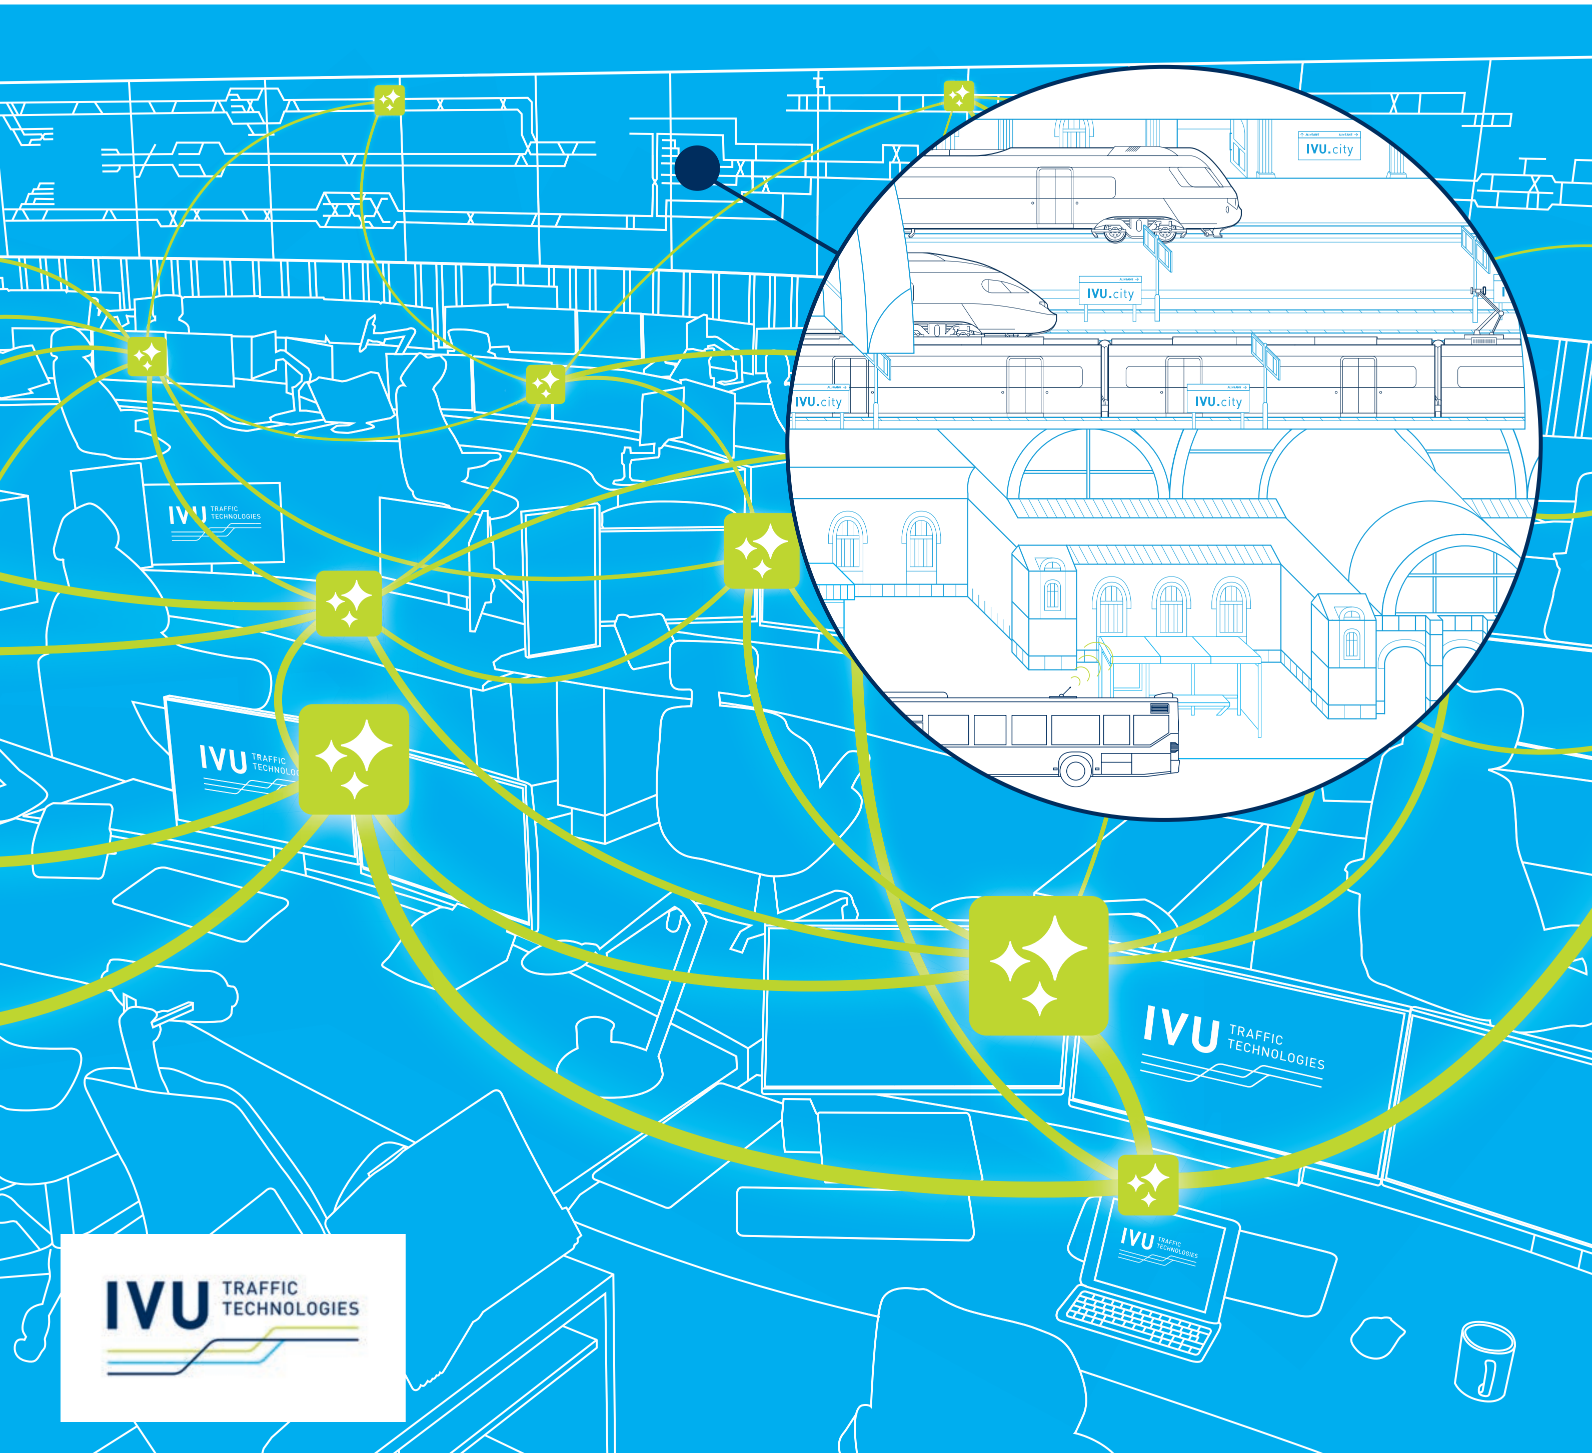

DACH REGION
Germany, Austria and Switzerland drive ATO and RTO research
Page 24

SKILLS
Realising the limits of AI in the control centre
Page 54

IN FOCUS
InnoTrans 2026 showcasing mobility on a global scale
Page 84

Realising the limits of AI in the control centre

Photos: DB AG

Despite the hype around AI, the rail sector faces a delicate challenge in balancing clear opportunities for digitalisation of control centre functions while recognising that 'a better algorithm' is no silver bullet for delivering more reliable or efficient services.

Dr Mareike Massow
Head of Product Management, IVU Traffic Technologies

Prof Dr Birgit Milius
Professor of Railway Operations & Management, School of Land & Sea Transport Studies, Technische Universität Berlin

The idea of a 'digital dispatcher' has been circulating in railway IT for the better part of a decade. But the shift in vocabulary from 'decision support' to 'autonomous operations' has occurred with a speed that the reality in control centres has not matched. Now the advance of AI is raising expectations even further.

This article is an attempt to map the terrain honestly — for the people who run those control centres, procure systems, and are accountable when the tools fall short. The central argument is simple: the digital dispatcher is not a magical tool to acquire once, but a direction to pursue one step at a time. Understanding what that means requires clarity about products, operating models, and the nature of the work itself.

A note on scope and terminology

This article focuses on staff managing train operations and resources from railway control centres — not on roles that directly set signals or manage shunting moves. This distinction matters, because control centre decisions are, in the main, not directly safety-critical in the technical sense. Boundary cases exist — authorising a crew to work marginally beyond regulated working time, for instance — but these remain exceptions. The automation of directly safety-relevant decisions lies outside the scope of this article.

Across European railways, the control centre role described in this article goes by many different names: disponent, régulateur, trafikledare — each reflecting its own operational tradition and regulatory context. Even within the English-speaking railway world, no single term has achieved consensus: 'dispatcher' dominates in North America, while 'controller' is the established usage in UK rail; in Britain, 'dispatcher' typically refers to staff on platforms at major stations authorised to allow trains to depart. Throughout this article, the authors use 'dispatcher' as the broader common denominator — and they hope British readers will forgive the compromise.

No easy automation

Anyone who has spent time in an operations control centre during a major disruption knows what is at stake. A dispatcher managing a busy corridor on a bad afternoon is simultaneously tracking dozens of delayed trains, fielding calls from drivers, station staff and other control centres, protecting connections, monitoring crew compliance, updating passenger information systems, and documenting decisions — all in real time, with incomplete information, under great pressure. This is not a workflow that yields easily to automation. It is a cognitive environment of considerable intensity, and the people who do it well have typically spent years developing the situational awareness and judgment it requires.

That expertise is increasingly difficult to replace. A significant cohort of the most experienced dispatchers across European railways is approaching retirement, taking with them tacit knowledge — for example, about how a specific junction behaves under perturbation, which connections to protect and which to sacrifice — that is nearly impossible to transfer at scale. Networks, meanwhile, are operating closer to capacity, meaning disruption propagates faster and recovery margins are tighter. This is why the digital dispatcher is not merely a technology ambition. It is a response to an operational and organisational problem that is getting harder.

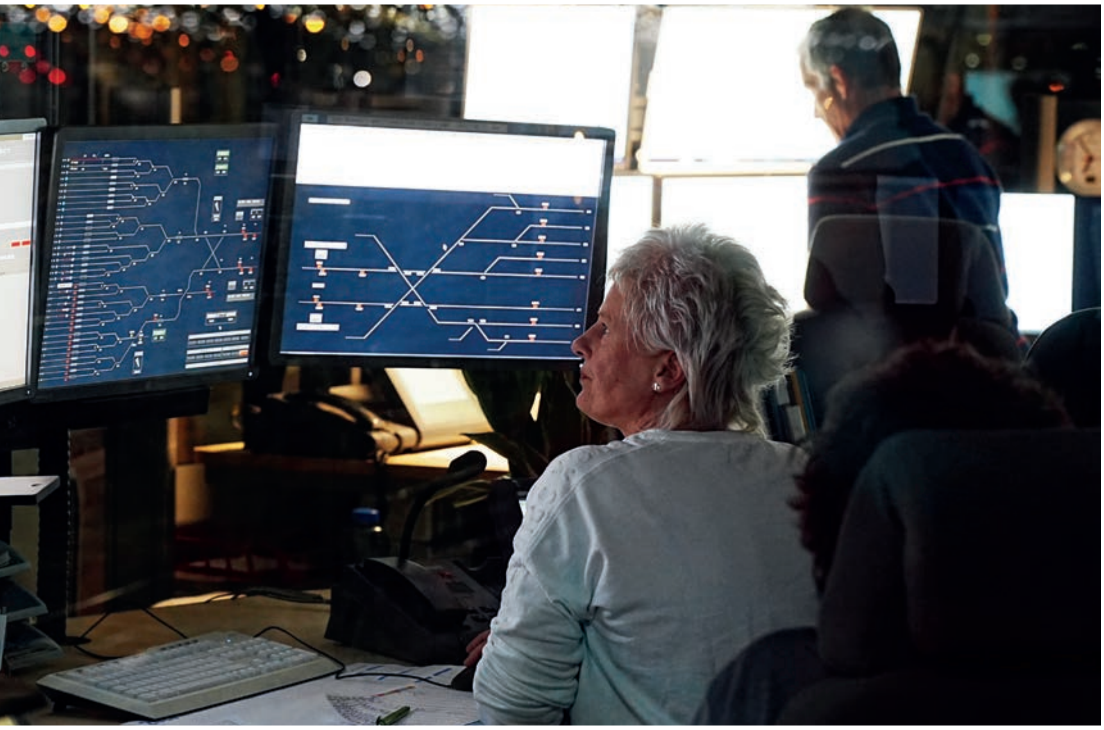

Table I. Who benefits from AI capability?

Capability	Value to infrastructure manager	Value to train operator
Network-wide delay prediction	High	Medium
Information consolidation and incident management	High	High
Conflict detection & re-scheduling	High	Low (constrained by IM primacy)
Rolling stock propagation analysis	Low	High
Crew optimisation	Low	High
Passenger information automation	Medium	Medium

The 'digital baseline'

Before assessing what comes next, it is worth being precise about what is already operational. The picture is more advanced than sceptics acknowledge — and less advanced than technology suppliers sometimes claim.

Delay prediction models are running at some infrastructure managers and larger railway operators. These systems provide probabilistic forecasts of how disruption will propagate through the network. Their accuracy varies depending on data quality and network complexity, but the best implementations are operationally useful.

Automated conflict detection is an option. Given a real-time picture of train positions and a timetable, it is straightforward to identify where two trains will need the same piece of infrastructure. The harder problem — what to do about it — is where automation becomes more limited.

Optimisation-based re-scheduling tools exist and are in use, primarily at infrastructure manager level, for resolving conflicts on defined sub-networks. Rolling stock and crew propagation analysis — identifying which downstream resource rotations are broken by a delay — is increasingly automated. Passenger information, once dependent on manual dispatcher input, is now largely driven by real-time data feeds.

Dispatchers must juggle a complex flow of data from disparate sources in order to manage train operations; this task becomes even harder at times of significant disruption.

The honest summary: the toolkit for control centres of 2026 is substantially better equipped than the one of 2016. But to bring these tools into widespread adoption is still a large undertaking in an industry that is complex, often safety critical, and relies on routines that have been proven to work in practice for decades.

So today, human dispatchers can be augmented by digital tools that reduce their workload in routine conditions. But the step from augmentation to autonomy — from a system that supports a decision to one that makes it — is still further ahead. Understanding why this is the case requires looking at two things: the nature of the tools themselves, and the organisational context in which they operate.

Mathematical optimisation

In discussions, AI and mathematical optimisation are often plotted against each other, but they both have their areas of usefulness. The two are related but distinct — and understanding which method solves which problem best matters for setting realistic expectations.

Mathematical optimisation — including Mixed Integer Linear Programming and Constraint Programming — finds the best available solution within a problem that has been fully defined by human specialists. The objective function, the constraints, the variables: all are specified in advance. The algorithm searches the solution space with precision and, in the case of exact methods, with provable guarantees. The result is auditable. This is why optimisation engines have found their way into railway applications: they do exactly and only what they are designed to do. Duty scheduling optimisation is a prime example: systems such as IVU.rail apply these methods to generate and replan crew duties across large rosters in minutes — a task that would take human planners many hours — while fully respecting all relevant contractual, regulatory and operational constraints.

Between classic optimisation and machine learning sits a family of nature-inspired metaheuristic methods that are legitimately both. Ant Colony Optimisation, modelled on the pheromone trail behaviour of ant colonies, is one such method — and it is already in operational use for predicting how delays will propagate through a network.

Classic AI approaches like machine learning operate differently: rather than solving a pre-specified problem, they discover structure from data. It generalises across situations which it has not explicitly seen before. This is powerful, but it introduces uncertainty: the model may behave unexpectedly in conditions that differ from its training data, and its reasoning is not directly interpretable in the way an optimisation solution is.

The distinction has direct practical implications for real-world application. For example, classic optimisation is broadly compatible with existing safety certification frameworks. Machine learning components in the operational loop face a significantly harder path through EN 50128 and related CENELEC standards.

Proposing an alternative train path during disruption is generally much more suitable for an optimisation algorithm. Predicting how a delay will propagate through the network is a problem well suited for metaheuristics like Ant Colony Optimisation, which can navigate the combinatorial complexity involved. Estimating how long an infrastructure fault is likely to last — based on fault type, location, time of day, season, and the historical resolution times of comparable incidents — is a machine learning challenge: there is no formula for it, but there are patterns in years of operational data that a trained model can learn to recognise. Reading a free-text disruption notification and extracting the operationally relevant information is a role for natural language AI.

These are not competing tools. They are complementary layers, each doing what the other cannot.

> ‘Predicting how a delay will propagate through the network is a problem well suited for metaheuristics like Ant Colony Optimisation’

Operating model

The most important variable in assessing the realistic scope of AI-assisted dispatching is one that is rarely analysed in AI discussions: the operating model.

In the predominant European model, infrastructure management and train operations are separated. An infrastructure manager owns and operates the network, train operators (or Railway Undertakings in European policy vernacular) run trains on it. During disruption, the IM typically holds operational primacy — it decides how the network is managed, which trains are prioritised, and how paths are reallocated. AI-based disruption impact prediction and the rapid evaluation of alternatives give the infrastructure manager tools that can extend what a human team can do under pressure.

Operators must adapt to decisions made by an organisation over which they have no direct control. Often there are general prioritisation rules, but RUs have limited real-time visibility into the IM's reasoning.

AI-driven dispatch support can help to deliver network-wide transparency for IMs and operators alike — a single operational picture integrating all train movements. To achieve that at the highest quality, the track authority and the operator need to co-operate. For example, infrastructure managers have the best visibility of what is happening on the lineside: they can predict arrival delays with reasonable confidence. Departure delays are a different matter. These are driven by vehicle conditions, crew availability and the critical handover points in operational planning — information that sits with the operator, not the infrastructure manager. The best delay predictions are therefore not produced by either party alone, but emerge from continuous, ideally automated, data exchange between both.

For operators alone, the primary AI lever is more constrained: optimising their own resources — rolling stock rotations, crew plans, connection protection decisions — within the operational reality set by the IM. Expanding that scope requires not just better algorithms, but better data agreements. This distinction matters for procurement. A train operator evaluating a dispatch support system may be looking at capabilities it can neither technically exploit nor organisationally act upon. The fit between system architecture and operating model is more determinative of value than algorithmic capability.

In vertically integrated models — where infrastructure and operations sit within the same organisation — the scope of possible AI-based applications broadens substantially. Real-time simulation, integrated scenario comparison across timetables, resources and network state, and end-to-end adjustment of operational plans become genuinely achievable. The digital dispatcher in this context is not more intelligent in principle. But it is embedded in an organisational structure where its recommendations can be acted upon without cross-boundary negotiation — which in practice makes an enormous difference to what automation can deliver.

Two trains are competing for access to the same section of infrastructure: managing competing priorities is a key human skill for control centre staff.

The overlooked bottleneck

Ask a dispatcher what consumes most of their cognitive effort during a major disruption, and the answer is rarely ‘deciding what to do’. More often, it is keeping track of what is happening.

Disruption notifications arrive by telephone. Updates come via chat messages, emails, and free-text entries in operational systems. A train driver reports an issue. Information about maintenance work arrives via a different channel. The dispatcher integrates all of this mentally, in real time, while simultaneously managing communications with multiple parties and updating an operational picture that is never quite complete.

This fragmented information environment is the most significant and maybe least discussed bottleneck in railway dispatching. It is also the area where AI — specifically natural language processing and information structuring tools — offers the most credible near-term value. Not by making autonomous decisions, but by doing what current systems cannot: reading unstructured inputs, extracting operationally relevant content, consolidating it into a coherent situational picture, and ideally suggesting the response options that past disruptions of a similar type advise.

The prerequisite for this is systematic data capture. Learning from past disruptions requires that those disruptions were recorded in structured,

queryable form — which in most operational environments today they are not. Railway operators are increasingly recognising that they need a centralised Incident Management System: a single, authoritative record of every disruption event, response action and operational decision, serving as the nucleus from which all further communication, reporting and analysis flows — and that can later be queried, analysed and fed back into prediction models.

Infrastructure managers have generally moved faster on this than train operators, driven by their network-wide co-ordination role and the volume of cross-operator incident data they handle. For many RUs, systematic incident recording remains a gap rather than a capability, but software providers are responding. IVU Traffic Technologies, for instance, is developing an Incident Management System as an extension of its integrated resource planning and dispatching suite, recognising that a system which already holds the operational plan, the crew assignments and the vehicle rotations is the logical place to capture and contextualise disruption data.

Looking further ahead, the way infrastructure managers and train operators exchange operational information may itself be transformed. Today, cross-organisational data sharing depends on agreed structured interfaces — standardised formats, defined fields, negotiated protocols. In a longer-term vision, AI agents on either side of the IM-operator interface could communicate directly with each other: querying, interpreting and acting on information without requiring every data exchange to be pre-specified in a formal interface contract.

Today, the investment that will most directly accelerate the digital dispatcher is not a more powerful algorithm. It is the instrumentation of the control centre itself: ensuring that what happens today can be learned from tomorrow.

Limits, realism and the future

There are things AI will not resolve in the foreseeable future, and it is worth naming them plainly. Legal accountability for operational decisions will remain with a human. Safety certification of learned AI components in the operational loop is a long road under current standards.

Truly novel disruption scenarios — the ones that have never appeared in training data — will continue to require human judgment. The cross-operator crisis co-ordination that experienced dispatchers conduct through personal relationships and phone calls has no algorithmic equivalent.

The roadmap, for organisations making investment decisions today, should look like this. In the near term, the priority is information consolidation and structured decision support: tools that reduce cognitive load and make the operational picture more legible under pressure. In the medium term, conditional automation of well-defined, low-risk standard actions becomes viable. Broader operational autonomy will likely emerge first in vertically integrated environments with mature data infrastructure and evolving regulatory frameworks.

The common thread across all of this is that technology is not the binding constraint. Data quality, interface standardisation, operating model fit, organisational readiness, and the genuine involvement of working dispatchers in system design: these are what will determine whether a digital dispatcher capability actually delivers in a control centre, or remains an idea that never quite makes it into production.

The direction is clear, the tools are maturing, and the operational need is real. What the industry needs now is less ambition in the marketing and more rigour in the implementation. The destination remains, appropriately, in human hands.

Note: The authors acknowledge that AI tools were used in the preparation of the initial draft of this feature. The editing by *Railway Gazette International* was undertaken manually.

‘The algorithm is the easy part’

Danish national operator DSB’s Head of Current Operations Dr Thomas Gordon on how AI is changing life in the control room.

The term AI is everywhere in railway discussions right now. What does it mean to you?

In the public realm, AI is used in such a broad sense that it can simultaneously mean everything and nothing. I prefer to look at specific examples.

Take phone calls. In a control centre, a driver calls in to report a problem with the train. Today, that call depends on the experience of whoever picks up. But if you analyse that conversation in real time, AI could suggest the most probable causes and offer the dispatcher a set of ready responses. We are already experimenting with exactly this kind of call analysis in customer support. There is no fundamental reason it could not be applied in a control centre.

Where are you seeing the most tangible operational results from automation today?

In re-planning. If we have a problem in operations today that will continue into tomorrow, we need to replan all the duties for the following day. Done manually, that is a task of many hours — in practice, it is nearly impossible to complete properly the night before. With the adjustment optimiser we use at DSB S-Tog [in København], it takes around 20 min. That turns an almost impossible situation into a manageable one. It turns something humans genuinely cannot solve in the available time into something solvable. Technically, this is not AI in the strict sense — but we might as well describe it that way.

The next step is to use the same tool in the morning to replan the afternoon. And gradually, you move that window closer and closer to the present moment. That is the trajectory.

You mentioned rolling stock management as another area. Can you give an example?

Sometimes you need smaller, simpler tools that deliver a tiny solution in seconds rather than a perfect solution to a global problem in hours. That is actually what dispatchers do all day. A train needs to be stopped for 10 min to clean insects off the windscreen — trivial, but those 10 min need to be planned. The vehicle also needs to be in the depot tomorrow. The system should be able to absorb those small constraints and quietly replan the related tasks around them. We have built exactly this kind of tool for drivers: assign a task, and the system automatically replans the break. Small, fast, and immediately useful.

What is the hardest part of this change?

For other professions, this is completely normal — think of a taxi driver. In rail, we have historically treated the duty as something close to a contract. Changing that mindset is much more a matter of culture than of what is technically possible — and it has to be done gently, over years, not overnight.

What do we have to start with to make any of this work?

Trust. Between the company and the trades union, between management and the people in the control centre, between the system and the dispatcher who is supposed to rely on it. If that trust is not there, no algorithm will compensate for it.

What we have found is that people's actual preferences vary considerably. Perhaps 20% genuinely need to know their duty for tomorrow in advance. Many have no strong view either way. And a further 20% actively prefer flexibility. The system has to make room for all of them. Drivers, infrastructure controllers, dispatchers — everyone has to be able to see themselves in the technology. That is the base of it all.

Questions posed by Dr Mareike Massow.