

RUDYCT ACADEMIC SERIES 2026

Social Intelligence **in the Age of Automation**

Restoring the Human Touch

SOCIAL INTELLIGENCE • EMPATHY • TRUST • HUMAN-AI COLLABORATION

Prof. Ir. Rudy C. Tarumingkeng, PhD

Bogor, Indonesia | 29 July 2026

www.rudymt.com/cv.pdf

RUDYCT ACADEMIC SERIES 2026

SOCIAL INTELLIGENCE IN THE AGE OF AUTOMATION

Restoring the Human Touch



EMPATHY



TRUST



COMMUNICATION



HUMAN-AI
COLLABORATION

By
Prof. Ir. Rudy C. Tarumingkeng, PhD



www.rudycr.com/cv.pdf



Abstract

Automation is changing not only what work is performed but also how people interpret one another, exercise judgment, build trust, and experience dignity at work. This article argues that social intelligence is therefore not a decorative “soft skill.” It is a strategic capability that enables individuals, teams, and institutions to combine machine speed with human meaning. Drawing on research in social intelligence, emotional intelligence, human–AI collaboration, psychological safety, collective intelligence, job design, service management, and responsible innovation, the article develops an integrated framework for restoring the human touch without romanticizing either people or technology. It explains how social perception, empathic interpretation, self-regulation, communication, coordination, ethical judgment, and relationship repair create value in automated settings. It also examines risks: algorithmic management, de-skilling, surveillance, false empathy, automation bias, exclusion, emotional labor, and weakened accountability. Sector narratives, measurement principles, learning practices, governance mechanisms, and a phased implementation roadmap translate the argument into organizational action. The central conclusion is that successful automation is a sociotechnical achievement. Organizations gain durable value when they design technology to enlarge human agency, strengthen relationships, and reserve consequential judgment for accountable people.

Keywords: social intelligence; automation; artificial intelligence; human–AI collaboration; empathy; trust; psychological safety; job design; leadership; responsible innovation

Executive Summary

The central argument of this article is simple but demanding: automation does not remove the social nature of work. It redistributes it. Routine exchanges may disappear, yet the remaining human moments are often more ambiguous, emotional, consequential, and difficult to standardize. Employees inherit exceptions; managers must explain choices shaped by systems they did not build; customers seek a person when the automated path cannot recognize their circumstances. The organizations that prosper will not be those that merely preserve courtesy around digital machinery. They will design a coherent sociotechnical system in which human

attention, judgment, voice, care, and accountability are treated as productive infrastructure.

Social intelligence is presented here as an integrated capacity with eight linked layers: social perception, contextual interpretation, empathy, self-awareness and regulation, clear expression, coordination, ethical discernment, and relationship repair. This model avoids two mistakes. The first is sentimentalism—the belief that goodwill alone corrects unfair systems. The second is technological determinism—the belief that whatever can be automated should be automated and that social consequences will adjust by themselves. A mature organization instead makes explicit choices about task allocation, decision rights, escalation, explanation, and redress.

The practical agenda has five parts. Diagnose how automation changes relationships and emotional load. Redesign jobs so that people retain agency, comprehension, and a path to mastery. Develop social intelligence through deliberate practice in real situations. Govern consequential systems through meaningful human review and contestability. Finally, measure value with a balanced set of operational, human, customer, learning, and risk indicators. Restoring the human touch is therefore not a retreat from technology. It is the discipline of making technology serve human purposes under accountable institutional conditions.

1. Introduction: Automation Is Also a Social Transformation

Automation reorganizes attention, authority, time, and dependence among people. This matters because leaders often treat implementation as a technical installation and discover too late that altered relationships determine whether the system is trusted and used. The important analytical move is to stop separating technical performance from the social conditions that make performance interpretable and legitimate. A system can be accurate in aggregate and still fail a particular encounter if the affected person cannot understand the decision, express relevant context, or obtain a remedy. Conversely, a warm interaction cannot redeem a process that is systematically unfair. Social intelligence joins these levels by asking what people notice, how they interpret it, what authority they possess, and whose consequences count.

The governing principle is that social intelligence must be designed into workflows, roles, interfaces, and managerial routines. In practice, this means organizations

should map every automated decision together with the people who interpret, contest, explain, and live with it. The objective is neither to keep people inside every minor transaction nor to automate every activity that appears standardized. It is to create a deliberate distribution of cognition, emotion, and accountability. Machines can search, calculate, draft, detect patterns, and maintain consistency. People can frame ambiguous purposes, recognize moral salience, negotiate competing values, and accept responsibility. The boundary must remain revisable because tasks, populations, and system capabilities change.

Research on social intelligence and interpersonal skill emphasizes that effective conduct depends on both accurate social perception and adaptive expression (Riggio, 1986; Silvera et al., 2001). The research should be read as a design resource rather than a promise that a single intervention will generalize everywhere. Context changes the meaning of behavior. A question that signals curiosity in one setting may signal distrust in another; rapid response may demonstrate service in one context and disrespectful impatience in another. Social intelligence therefore requires local inquiry, feedback from affected groups, and the willingness to revise routines when observed consequences differ from intention.

A central danger is that efficiency can rise while voice, learning, and legitimacy quietly decline. The remedy is not another slogan about people first. It is evidence in the workflow: protected channels for dissent, explicit decision rights, sufficient time for conversation, meaningful human review, accessible explanations, and learning from overrides and failures. These arrangements convert human touch from an individual personality trait into an organizational capability. They also expose trade-offs. If a service promises empathy but budgets no time for listening, the operating model—not the employee's attitude—is the primary contradiction.

2. What Social Intelligence Means

Consider the proposition that the capacity to perceive social cues, understand motives and contexts, regulate responses, and act effectively and ethically with others. It changes the managerial question from “Can the tool perform the task?” to “What pattern of relationships will exist after the task is redesigned?” That second question reveals dependencies that a process map may hide: who teaches the system, who absorbs exceptions, who explains outcomes, who notices unusual harm, and who can stop execution. Because the term is sometimes reduced to charm,

networking, or extroversion, which obscures its cognitive, emotional, behavioral, and moral components, apparently minor interface and policy decisions can alter status, voice, and trust across an organization.

A workable response starts from this principle: competence is demonstrated through accurate interpretation and constructive action, not through social performance alone. Leaders can operationalize it when they evaluate both the quality of a person's reading of a situation and the consequences of the response. This is a form of sociotechnical design. It recognizes that output quality arises from interactions among data, models, incentives, expertise, language, identity, and power. Social intelligence supplies the interpretive capacity to see these interactions and the behavioral capacity to act on them. Ethical judgment supplies the boundary: effectiveness cannot be defined solely by the objective of the system owner.

Studies of collective intelligence show that social sensitivity and participation patterns help explain team performance beyond the intelligence of individual members (Woolley et al., 2010). These findings point toward complementarity, but complementarity is not automatic. It must be produced through role clarity, practice with edge cases, feedback about outcomes, and authority to intervene. A person who sees a problem but cannot alter the decision is not providing oversight. A person who clicks approval without time or information is not exercising judgment. The quality of collaboration is therefore visible in what happens when human and machine assessments disagree.

The failure mode deserves equal attention: a socially fluent person may manipulate others unless effectiveness is bounded by reciprocity and ethics. Organizations should use this possibility as a testable risk hypothesis. They can look for early signals in complaint narratives, override patterns, participation data, turnover, workarounds, and the gap between formal procedures and actual behavior. Such inquiry is socially intelligent because it treats resistance not merely as an adoption problem but potentially as information about bad fit, lost identity, unrecognized workload, or compromised values.

3. Related Concepts and Necessary Distinctions

Social intelligence overlaps with emotional intelligence, empathy, cultural intelligence, communication skill, and political skill without being identical to any of them. Seen from the employee's position, the issue is experienced through daily

moments: a recommendation that seems impossible to question, a customer whose need does not fit the form, a manager who quotes a score instead of holding a conversation, or an expert who discovers that the underlying skill is fading. Seen from the institution's position, conceptual confusion produces fragmented training and misleading assessment. The two views need to be held together; otherwise personal experience is dismissed as anecdote or aggregate data is dismissed as dehumanizing.

The core design commitment is that an integrated model should preserve the distinct contribution of emotion, context, culture, power, coordination, and moral responsibility. A concrete starting point is to define the behavioral domain before selecting instruments or learning programs. This creates an observable hypothesis about better work. For example, the organization can examine whether response quality, resolution time, safety reporting, understanding, and stakeholder confidence improve together. If one indicator improves while another deteriorates, leaders should resist declaring victory and investigate how benefit and burden are distributed.

Research on human–AI collaboration finds that value depends on task allocation, complementarity, and the conditions under which people can recognize and correct machine errors (Jarrahi, 2018; Fügner et al., 2021). The implication is that social intelligence is not located exclusively inside an individual. It is also scaffolded—or obstructed—by meeting design, interface defaults, span of control, employment security, data access, professional norms, and leadership example. An attentive employee in a punitive system will eventually learn silence. A considerate manager with no discretion cannot provide meaningful accommodation. Capability development and structural redesign must therefore proceed together.

The caution is straightforward: labels can become fashionable substitutes for observable capability. Prevention begins before deployment by involving those who perform and receive the work. It continues after deployment through review of actual consequences, including consequences that the business case did not predict. The aim is institutional humility: confidence sufficient to act, coupled with channels that make correction possible. Restoring the human touch means making attention and response reliable, especially when standardized pathways fail.

4. Why the Human Touch Becomes More Valuable

Machines make prediction and production cheaper while ambiguity, care, persuasion, reconciliation, and responsibility remain socially intensive. This matters because organizations may assume that human contribution shrinks wherever an algorithm becomes competent. The important analytical move is to stop separating technical performance from the social conditions that make performance interpretable and legitimate. A system can be accurate in aggregate and still fail a particular encounter if the affected person cannot understand the decision, express relevant context, or obtain a remedy. Conversely, a warm interaction cannot redeem a process that is systematically unfair. Social intelligence joins these levels by asking what people notice, how they interpret it, what authority they possess, and whose consequences count.

The governing principle is that value migrates toward framing, exception handling, relationship stewardship, and accountable judgment. In practice, this means organizations should redesign roles around complementary bundles of technical and relational work. The objective is neither to keep people inside every minor transaction nor to automate every activity that appears standardized. It is to create a deliberate distribution of cognition, emotion, and accountability. Machines can search, calculate, draft, detect patterns, and maintain consistency. People can frame ambiguous purposes, recognize moral salience, negotiate competing values, and accept responsibility. The boundary must remain revisable because tasks, populations, and system capabilities change.

The automation–augmentation paradox explains why the pursuit of substitution can weaken the human capabilities needed for effective augmentation (Raisch & Krakowski, 2021). The research should be read as a design resource rather than a promise that a single intervention will generalize everywhere. Context changes the meaning of behavior. A question that signals curiosity in one setting may signal distrust in another; rapid response may demonstrate service in one context and disrespectful impatience in another. Social intelligence therefore requires local inquiry, feedback from affected groups, and the willingness to revise routines when observed consequences differ from intention.

A central danger is that the phrase human touch can conceal poor processes or unpaid emotional labor unless backed by resources and authority. The remedy is not another slogan about people first. It is evidence in the workflow: protected channels for

dissent, explicit decision rights, sufficient time for conversation, meaningful human review, accessible explanations, and learning from overrides and failures. These arrangements convert human touch from an individual personality trait into an organizational capability. They also expose trade-offs. If a service promises empathy but budgets no time for listening, the operating model—not the employee's attitude—is the primary contradiction.

5. Human–AI Complementarity and the Automation–Augmentation Paradox

Consider the proposition that the same technology can substitute for human action, augment human capability, or do both at different moments. It changes the managerial question from “Can the tool perform the task?” to “What pattern of relationships will exist after the task is redesigned?” That second question reveals dependencies that a process map may hide: who teaches the system, who absorbs exceptions, who explains outcomes, who notices unusual harm, and who can stop execution. Because local automation decisions can erode the expertise required to supervise the system and manage rare exceptions, apparently minor interface and policy decisions can alter status, voice, and trust across an organization.

A workable response starts from this principle: the unit of design should be the whole task system rather than an isolated activity. Leaders can operationalize it when they specify when people lead, when machines advise, when machines execute, and how control returns to people. This is a form of sociotechnical design. It recognizes that output quality arises from interactions among data, models, incentives, expertise, language, identity, and power. Social intelligence supplies the interpretive capacity to see these interactions and the behavioral capacity to act on them. Ethical judgment supplies the boundary: effectiveness cannot be defined solely by the objective of the system owner.

Evidence on psychological safety demonstrates that learning and error reporting require an interpersonal climate in which candor is not punished (Edmondson, 1999; Edmondson & Lei, 2014). These findings point toward complementarity, but complementarity is not automatic. It must be produced through role clarity, practice with edge cases, feedback about outcomes, and authority to intervene. A person who sees a problem but cannot alter the decision is not providing oversight. A person who clicks approval without time or information is not exercising judgment. The quality of

collaboration is therefore visible in what happens when human and machine assessments disagree.

The failure mode deserves equal attention: augmentation rhetoric can disguise displacement, while human-in-the-loop language can disguise ceremonial oversight. Organizations should use this possibility as a testable risk hypothesis. They can look for early signals in complaint narratives, override patterns, participation data, turnover, workarounds, and the gap between formal procedures and actual behavior. Such inquiry is socially intelligent because it treats resistance not merely as an adoption problem but potentially as information about bad fit, lost identity, unrecognized workload, or compromised values.

6. A Multilayer Architecture of Social Intelligence

Effective social action links perception, interpretation, empathy, self-regulation, expression, coordination, ethics, and repair. Seen from the employee's position, the issue is experienced through daily moments: a recommendation that seems impossible to question, a customer whose need does not fit the form, a manager who quotes a score instead of holding a conversation, or an expert who discovers that the underlying skill is fading. Seen from the institution's position, training one visible behavior, such as active listening, rarely changes the underlying chain of attention and judgment. The two views need to be held together; otherwise personal experience is dismissed as anecdote or aggregate data is dismissed as dehumanizing.

The core design commitment is that development must connect internal sensemaking with external behavior and collective routines. A concrete starting point is to use observation, practice, feedback, reflection, and real-work application as one learning cycle. This creates an observable hypothesis about better work. For example, the organization can examine whether response quality, resolution time, safety reporting, understanding, and stakeholder confidence improve together. If one indicator improves while another deteriorates, leaders should resist declaring victory and investigate how benefit and burden are distributed.

Trust in automation is most useful when calibrated to capability and context rather than maximized without qualification (Lee & See, 2004; Glikson & Woolley, 2020). The implication is that social intelligence is not located exclusively inside an individual. It is also scaffolded—or obstructed—by meeting design, interface defaults, span of control, employment security, data access, professional norms, and

leadership example. An attentive employee in a punitive system will eventually learn silence. A considerate manager with no discretion cannot provide meaningful accommodation. Capability development and structural redesign must therefore proceed together.

The caution is straightforward: overconfidence at any layer can propagate errors through the entire interaction. Prevention begins before deployment by involving those who perform and receive the work. It continues after deployment through review of actual consequences, including consequences that the business case did not predict. The aim is institutional humility: confidence sufficient to act, coupled with channels that make correction possible. Restoring the human touch means making attention and response reliable, especially when standardized pathways fail.

7. Social Perception in Data-Rich Environments

Social perception involves noticing tone, silence, timing, participation patterns, status signals, and contextual constraints. This matters because dashboards privilege what is countable and can push weak signals outside managerial attention. The important analytical move is to stop separating technical performance from the social conditions that make performance interpretable and legitimate. A system can be accurate in aggregate and still fail a particular encounter if the affected person cannot understand the decision, express relevant context, or obtain a remedy. Conversely, a warm interaction cannot redeem a process that is systematically unfair. Social intelligence joins these levels by asking what people notice, how they interpret it, what authority they possess, and whose consequences count.

The governing principle is that quantitative indicators and situated human observation should correct one another. In practice, this means organizations should pair operational reviews with structured inquiry into stakeholder experience and anomalous cases. The objective is neither to keep people inside every minor transaction nor to automate every activity that appears standardized. It is to create a deliberate distribution of cognition, emotion, and accountability. Machines can search, calculate, draft, detect patterns, and maintain consistency. People can frame ambiguous purposes, recognize moral salience, negotiate competing values, and accept responsibility. The boundary must remain revisable because tasks, populations, and system capabilities change.

Research on social intelligence and interpersonal skill emphasizes that effective conduct depends on both accurate social perception and adaptive expression (Riggio, 1986; Silvera et al., 2001). The research should be read as a design resource rather than a promise that a single intervention will generalize everywhere. Context changes the meaning of behavior. A question that signals curiosity in one setting may signal distrust in another; rapid response may demonstrate service in one context and disrespectful impatience in another. Social intelligence therefore requires local inquiry, feedback from affected groups, and the willingness to revise routines when observed consequences differ from intention.

A central danger is that more data can create an illusion of social knowledge while suppressing curiosity. The remedy is not another slogan about people first. It is evidence in the workflow: protected channels for dissent, explicit decision rights, sufficient time for conversation, meaningful human review, accessible explanations, and learning from overrides and failures. These arrangements convert human touch from an individual personality trait into an organizational capability. They also expose trade-offs. If a service promises empathy but budgets no time for listening, the operating model—not the employee's attitude—is the primary contradiction.

8. Empathy Without Sentimentality

Consider the proposition that empathy combines perspective taking, affective resonance, and compassionate concern. It changes the managerial question from “Can the tool perform the task?” to “What pattern of relationships will exist after the task is redesigned?” That second question reveals dependencies that a process map may hide: who teaches the system, who absorbs exceptions, who explains outcomes, who notices unusual harm, and who can stop execution. Because organizations often celebrate empathy as warmth but fail to convert understanding into fair decisions, apparently minor interface and policy decisions can alter status, voice, and trust across an organization.

A workable response starts from this principle: empathy should inform boundaries, priorities, and action without eliminating evidence or accountability. Leaders can operationalize it when they ask what another person may be experiencing, what evidence supports that inference, and what responsible response is possible. This is a form of sociotechnical design. It recognizes that output quality arises from interactions among data, models, incentives, expertise, language, identity, and

power. Social intelligence supplies the interpretive capacity to see these interactions and the behavioral capacity to act on them. Ethical judgment supplies the boundary: effectiveness cannot be defined solely by the objective of the system owner.

Studies of collective intelligence show that social sensitivity and participation patterns help explain team performance beyond the intelligence of individual members (Woolley et al., 2010). These findings point toward complementarity, but complementarity is not automatic. It must be produced through role clarity, practice with edge cases, feedback about outcomes, and authority to intervene. A person who sees a problem but cannot alter the decision is not providing oversight. A person who clicks approval without time or information is not exercising judgment. The quality of collaboration is therefore visible in what happens when human and machine assessments disagree.

The failure mode deserves equal attention: unregulated empathy may create bias toward vivid or familiar individuals and neglect less visible groups. Organizations should use this possibility as a testable risk hypothesis. They can look for early signals in complaint narratives, override patterns, participation data, turnover, workarounds, and the gap between formal procedures and actual behavior. Such inquiry is socially intelligent because it treats resistance not merely as an adoption problem but potentially as information about bad fit, lost identity, unrecognized workload, or compromised values.

9. Self-Awareness and Regulation Under Acceleration

Automation compresses response time and amplifies the consequences of impulsive communication. Seen from the employee's position, the issue is experienced through daily moments: a recommendation that seems impossible to question, a customer whose need does not fit the form, a manager who quotes a score instead of holding a conversation, or an expert who discovers that the underlying skill is fading. Seen from the institution's position, people can outsource reflection to notifications, rankings, and machine recommendations. The two views need to be held together; otherwise personal experience is dismissed as anecdote or aggregate data is dismissed as dehumanizing.

The core design commitment is that socially intelligent action requires a pause between signal, interpretation, emotion, and response. A concrete starting point is to create escalation thresholds, deliberation windows, and reflective routines for

consequential decisions. This creates an observable hypothesis about better work. For example, the organization can examine whether response quality, resolution time, safety reporting, understanding, and stakeholder confidence improve together. If one indicator improves while another deteriorates, leaders should resist declaring victory and investigate how benefit and burden are distributed.

Research on human–AI collaboration finds that value depends on task allocation, complementarity, and the conditions under which people can recognize and correct machine errors (Jarrahi, 2018; Fügner et al., 2021). The implication is that social intelligence is not located exclusively inside an individual. It is also scaffolded—or obstructed—by meeting design, interface defaults, span of control, employment security, data access, professional norms, and leadership example. An attentive employee in a punitive system will eventually learn silence. A considerate manager with no discretion cannot provide meaningful accommodation. Capability development and structural redesign must therefore proceed together.

The caution is straightforward: calm behavior should not be confused with emotional suppression or passive compliance. Prevention begins before deployment by involving those who perform and receive the work. It continues after deployment through review of actual consequences, including consequences that the business case did not predict. The aim is institutional humility: confidence sufficient to act, coupled with channels that make correction possible. Restoring the human touch means making attention and response reliable, especially when standardized pathways fail.

10. Trust, Reliance, and Appropriate Skepticism

Trust in automated work has interpersonal, institutional, and technological objects. This matters because users may over-trust fluent systems or under-trust systems introduced through opaque and coercive processes. The important analytical move is to stop separating technical performance from the social conditions that make performance interpretable and legitimate. A system can be accurate in aggregate and still fail a particular encounter if the affected person cannot understand the decision, express relevant context, or obtain a remedy. Conversely, a warm interaction cannot redeem a process that is systematically unfair. Social intelligence joins these levels by asking what people notice, how they interpret it, what authority they possess, and whose consequences count.

The governing principle is that calibrated reliance grows from competence, transparency, recourse, predictable boundaries, and credible accountability. In practice, this means organizations should make uncertainty, provenance, failure modes, and appeal routes visible at the point of use. The objective is neither to keep people inside every minor transaction nor to automate every activity that appears standardized. It is to create a deliberate distribution of cognition, emotion, and accountability. Machines can search, calculate, draft, detect patterns, and maintain consistency. People can frame ambiguous purposes, recognize moral salience, negotiate competing values, and accept responsibility. The boundary must remain revisable because tasks, populations, and system capabilities change.

The automation–augmentation paradox explains why the pursuit of substitution can weaken the human capabilities needed for effective augmentation (Raisch & Krakowski, 2021). The research should be read as a design resource rather than a promise that a single intervention will generalize everywhere. Context changes the meaning of behavior. A question that signals curiosity in one setting may signal distrust in another; rapid response may demonstrate service in one context and disrespectful impatience in another. Social intelligence therefore requires local inquiry, feedback from affected groups, and the willingness to revise routines when observed consequences differ from intention.

A central danger is that trust campaigns cannot compensate for a system that repeatedly violates expectations. The remedy is not another slogan about people first. It is evidence in the workflow: protected channels for dissent, explicit decision rights, sufficient time for conversation, meaningful human review, accessible explanations, and learning from overrides and failures. These arrangements convert human touch from an individual personality trait into an organizational capability. They also expose trade-offs. If a service promises empathy but budgets no time for listening, the operating model—not the employee's attitude—is the primary contradiction.

11. Psychological Safety, Voice, and the Right to Question

Consider the proposition that teams learn when members can report mistakes, doubt an output, and challenge authority without humiliation. It changes the managerial question from “Can the tool perform the task?” to “What pattern of relationships will exist after the task is redesigned?” That second question reveals dependencies that a process map may hide: who teaches the system, who absorbs exceptions, who

explains outcomes, who notices unusual harm, and who can stop execution. Because algorithmic recommendations acquire social authority even when no formal rule requires obedience, apparently minor interface and policy decisions can alter status, voice, and trust across an organization.

A workable response starts from this principle: leaders must make dissent operationally useful and protect people who surface inconvenient evidence. Leaders can operationalize it when they record overrides, near misses, and contested recommendations as learning data rather than as signs of disloyalty. This is a form of sociotechnical design. It recognizes that output quality arises from interactions among data, models, incentives, expertise, language, identity, and power. Social intelligence supplies the interpretive capacity to see these interactions and the behavioral capacity to act on them. Ethical judgment supplies the boundary: effectiveness cannot be defined solely by the objective of the system owner.

Evidence on psychological safety demonstrates that learning and error reporting require an interpersonal climate in which candor is not punished (Edmondson, 1999; Edmondson & Lei, 2014). These findings point toward complementarity, but complementarity is not automatic. It must be produced through role clarity, practice with edge cases, feedback about outcomes, and authority to intervene. A person who sees a problem but cannot alter the decision is not providing oversight. A person who clicks approval without time or information is not exercising judgment. The quality of collaboration is therefore visible in what happens when human and machine assessments disagree.

The failure mode deserves equal attention: voice mechanisms become performative when raising a concern has no visible consequence. Organizations should use this possibility as a testable risk hypothesis. They can look for early signals in complaint narratives, override patterns, participation data, turnover, workarounds, and the gap between formal procedures and actual behavior. Such inquiry is socially intelligent because it treats resistance not merely as an adoption problem but potentially as information about bad fit, lost identity, unrecognized workload, or compromised values.

12. Communication in Algorithmically Mediated Work

Messages are increasingly drafted, filtered, summarized, translated, prioritized, and evaluated by machines. Seen from the employee's position, the issue is experienced

through daily moments: a recommendation that seems impossible to question, a customer whose need does not fit the form, a manager who quotes a score instead of holding a conversation, or an expert who discovers that the underlying skill is fading. Seen from the institution's position, communication can become faster yet less owned, less contextual, and harder to repair. The two views need to be held together; otherwise personal experience is dismissed as anecdote or aggregate data is dismissed as dehumanizing.

The core design commitment is that people remain responsible for intent, audience, disclosure, tone, and consequences. A concrete starting point is to establish norms for reviewing machine-generated communication and for switching to direct conversation when stakes or ambiguity rise. This creates an observable hypothesis about better work. For example, the organization can examine whether response quality, resolution time, safety reporting, understanding, and stakeholder confidence improve together. If one indicator improves while another deteriorates, leaders should resist declaring victory and investigate how benefit and burden are distributed.

Trust in automation is most useful when calibrated to capability and context rather than maximized without qualification (Lee & See, 2004; Glikson & Woolley, 2020). The implication is that social intelligence is not located exclusively inside an individual. It is also scaffolded—or obstructed—by meeting design, interface defaults, span of control, employment security, data access, professional norms, and leadership example. An attentive employee in a punitive system will eventually learn silence. A considerate manager with no discretion cannot provide meaningful accommodation. Capability development and structural redesign must therefore proceed together.

The caution is straightforward: synthetic fluency can create false confidence, false intimacy, and diffusion of authorship. Prevention begins before deployment by involving those who perform and receive the work. It continues after deployment through review of actual consequences, including consequences that the business case did not predict. The aim is institutional humility: confidence sufficient to act, coupled with channels that make correction possible. Restoring the human touch means making attention and response reliable, especially when standardized pathways fail.

13. Leadership as Attention, Interpretation, and Moral Presence

Leadership in automated organizations consists less of possessing information and more of directing collective attention and legitimizing responsible action. This matters because managers may hide behind scores or use technology to avoid difficult conversations. The important analytical move is to stop separating technical performance from the social conditions that make performance interpretable and legitimate. A system can be accurate in aggregate and still fail a particular encounter if the affected person cannot understand the decision, express relevant context, or obtain a remedy. Conversely, a warm interaction cannot redeem a process that is systematically unfair. Social intelligence joins these levels by asking what people notice, how they interpret it, what authority they possess, and whose consequences count.

The governing principle is that leaders must translate data into context, acknowledge trade-offs, and remain personally answerable. In practice, this means organizations should require leaders to explain consequential decisions in language that affected people can understand and contest. The objective is neither to keep people inside every minor transaction nor to automate every activity that appears standardized. It is to create a deliberate distribution of cognition, emotion, and accountability. Machines can search, calculate, draft, detect patterns, and maintain consistency. People can frame ambiguous purposes, recognize moral salience, negotiate competing values, and accept responsibility. The boundary must remain revisable because tasks, populations, and system capabilities change.

Research on social intelligence and interpersonal skill emphasizes that effective conduct depends on both accurate social perception and adaptive expression (Riggio, 1986; Silvera et al., 2001). The research should be read as a design resource rather than a promise that a single intervention will generalize everywhere. Context changes the meaning of behavior. A question that signals curiosity in one setting may signal distrust in another; rapid response may demonstrate service in one context and disrespectful impatience in another. Social intelligence therefore requires local inquiry, feedback from affected groups, and the willingness to revise routines when observed consequences differ from intention.

A central danger is that visible empathy without distributive fairness becomes symbolic management. The remedy is not another slogan about people first. It is evidence in the workflow: protected channels for dissent, explicit decision rights,

sufficient time for conversation, meaningful human review, accessible explanations, and learning from overrides and failures. These arrangements convert human touch from an individual personality trait into an organizational capability. They also expose trade-offs. If a service promises empathy but budgets no time for listening, the operating model—not the employee's attitude—is the primary contradiction.

14. Teams, Collective Intelligence, and Coordination

Consider the proposition that team performance depends on balanced participation, social sensitivity, shared representations, and reliable coordination. It changes the managerial question from “Can the tool perform the task?” to “What pattern of relationships will exist after the task is redesigned?” That second question reveals dependencies that a process map may hide: who teaches the system, who absorbs exceptions, who explains outcomes, who notices unusual harm, and who can stop execution. Because adding AI can improve information processing while disrupting shared understanding and participation, apparently minor interface and policy decisions can alter status, voice, and trust across an organization.

A workable response starts from this principle: the team should treat the system as a designed resource within a social process, not as an infallible teammate. Leaders can operationalize it when they clarify roles, interaction protocols, handoffs, and the moments when collective deliberation is mandatory. This is a form of sociotechnical design. It recognizes that output quality arises from interactions among data, models, incentives, expertise, language, identity, and power. Social intelligence supplies the interpretive capacity to see these interactions and the behavioral capacity to act on them. Ethical judgment supplies the boundary: effectiveness cannot be defined solely by the objective of the system owner.

Studies of collective intelligence show that social sensitivity and participation patterns help explain team performance beyond the intelligence of individual members (Woolley et al., 2010). These findings point toward complementarity, but complementarity is not automatic. It must be produced through role clarity, practice with edge cases, feedback about outcomes, and authority to intervene. A person who sees a problem but cannot alter the decision is not providing oversight. A person who clicks approval without time or information is not exercising judgment. The quality of collaboration is therefore visible in what happens when human and machine assessments disagree.

The failure mode deserves equal attention: strong individual performance can coexist with weak collective intelligence. Organizations should use this possibility as a testable risk hypothesis. They can look for early signals in complaint narratives, override patterns, participation data, turnover, workarounds, and the gap between formal procedures and actual behavior. Such inquiry is socially intelligent because it treats resistance not merely as an adoption problem but potentially as information about bad fit, lost identity, unrecognized workload, or compromised values.

15. Job Design, Agency, and Meaningful Work

Technology changes autonomy, feedback, task variety, identity, workload, and opportunities to learn. Seen from the employee's position, the issue is experienced through daily moments: a recommendation that seems impossible to question, a customer whose need does not fit the form, a manager who quotes a score instead of holding a conversation, or an expert who discovers that the underlying skill is fading. Seen from the institution's position, efficiency projects frequently remove the parts of work through which people develop mastery or understand the whole service. The two views need to be held together; otherwise personal experience is dismissed as anecdote or aggregate data is dismissed as dehumanizing.

The core design commitment is that good design distributes cognitive work so that people retain agency, comprehension, and growth. A concrete starting point is to assess the future job rather than only the automated task and involve incumbents in redesign. This creates an observable hypothesis about better work. For example, the organization can examine whether response quality, resolution time, safety reporting, understanding, and stakeholder confidence improve together. If one indicator improves while another deteriorates, leaders should resist declaring victory and investigate how benefit and burden are distributed.

Research on human–AI collaboration finds that value depends on task allocation, complementarity, and the conditions under which people can recognize and correct machine errors (Jarrahi, 2018; Fügener et al., 2021). The implication is that social intelligence is not located exclusively inside an individual. It is also scaffolded—or obstructed—by meeting design, interface defaults, span of control, employment security, data access, professional norms, and leadership example. An attentive employee in a punitive system will eventually learn silence. A considerate manager

with no discretion cannot provide meaningful accommodation. Capability development and structural redesign must therefore proceed together.

The caution is straightforward: nominal human oversight is harmful when responsibility remains but control and knowledge disappear. Prevention begins before deployment by involving those who perform and receive the work. It continues after deployment through review of actual consequences, including consequences that the business case did not predict. The aim is institutional humility: confidence sufficient to act, coupled with channels that make correction possible. Restoring the human touch means making attention and response reliable, especially when standardized pathways fail.

16. Service Encounters and Customer Experience

Automation can remove friction from routine service while making difficult moments feel more alienating. This matters because customers are often trapped in standardized pathways precisely when their situation is exceptional. The important analytical move is to stop separating technical performance from the social conditions that make performance interpretable and legitimate. A system can be accurate in aggregate and still fail a particular encounter if the affected person cannot understand the decision, express relevant context, or obtain a remedy. Conversely, a warm interaction cannot redeem a process that is systematically unfair. Social intelligence joins these levels by asking what people notice, how they interpret it, what authority they possess, and whose consequences count.

The governing principle is that human access should be treated as a risk-sensitive design choice rather than a premium decoration. In practice, this means organizations should define escalation based on vulnerability, uncertainty, emotion, financial consequence, and irreversibility. The objective is neither to keep people inside every minor transaction nor to automate every activity that appears standardized. It is to create a deliberate distribution of cognition, emotion, and accountability. Machines can search, calculate, draft, detect patterns, and maintain consistency. People can frame ambiguous purposes, recognize moral salience, negotiate competing values, and accept responsibility. The boundary must remain revisable because tasks, populations, and system capabilities change.

The automation–augmentation paradox explains why the pursuit of substitution can weaken the human capabilities needed for effective augmentation (Raisch &

Krakowski, 2021). The research should be read as a design resource rather than a promise that a single intervention will generalize everywhere. Context changes the meaning of behavior. A question that signals curiosity in one setting may signal distrust in another; rapid response may demonstrate service in one context and disrespectful impatience in another. Social intelligence therefore requires local inquiry, feedback from affected groups, and the willingness to revise routines when observed consequences differ from intention.

A central danger is that anthropomorphic interfaces may simulate care without possessing obligation or authority. The remedy is not another slogan about people first. It is evidence in the workflow: protected channels for dissent, explicit decision rights, sufficient time for conversation, meaningful human review, accessible explanations, and learning from overrides and failures. These arrangements convert human touch from an individual personality trait into an organizational capability. They also expose trade-offs. If a service promises empathy but budgets no time for listening, the operating model—not the employee's attitude—is the primary contradiction.

17. Health Care and the Ethics of Presence

Consider the proposition that clinical and administrative automation can increase capacity but cannot replace the relational work through which patients understand and consent. It changes the managerial question from “Can the tool perform the task?” to “What pattern of relationships will exist after the task is redesigned?” That second question reveals dependencies that a process map may hide: who teaches the system, who absorbs exceptions, who explains outcomes, who notices unusual harm, and who can stop execution. Because prediction may be accurate while explanations, priorities, and lived circumstances remain unresolved, apparently minor interface and policy decisions can alter status, voice, and trust across an organization.

A workable response starts from this principle: technology should release time for presence and shared decision making. Leaders can operationalize it when they measure whether clinicians gain usable time and whether patients experience comprehension, respect, and recourse. This is a form of sociotechnical design. It recognizes that output quality arises from interactions among data, models, incentives, expertise, language, identity, and power. Social intelligence supplies the

interpretive capacity to see these interactions and the behavioral capacity to act on them. Ethical judgment supplies the boundary: effectiveness cannot be defined solely by the objective of the system owner.

Evidence on psychological safety demonstrates that learning and error reporting require an interpersonal climate in which candor is not punished (Edmondson, 1999; Edmondson & Lei, 2014). These findings point toward complementarity, but complementarity is not automatic. It must be produced through role clarity, practice with edge cases, feedback about outcomes, and authority to intervene. A person who sees a problem but cannot alter the decision is not providing oversight. A person who clicks approval without time or information is not exercising judgment. The quality of collaboration is therefore visible in what happens when human and machine assessments disagree.

The failure mode deserves equal attention: automation can intensify documentation burden or shift hidden work onto patients and families. Organizations should use this possibility as a testable risk hypothesis. They can look for early signals in complaint narratives, override patterns, participation data, turnover, workarounds, and the gap between formal procedures and actual behavior. Such inquiry is socially intelligent because it treats resistance not merely as an adoption problem but potentially as information about bad fit, lost identity, unrecognized workload, or compromised values.

18. Education and the Formation of Social Judgment

Generative systems alter writing, feedback, tutoring, assessment, and access to explanations. Seen from the employee's position, the issue is experienced through daily moments: a recommendation that seems impossible to question, a customer whose need does not fit the form, a manager who quotes a score instead of holding a conversation, or an expert who discovers that the underlying skill is fading. Seen from the institution's position, education may optimize output while weakening dialogue, intellectual struggle, and the ability to encounter difference. The two views need to be held together; otherwise personal experience is dismissed as anecdote or aggregate data is dismissed as dehumanizing.

The core design commitment is that schools and universities should teach learners to collaborate with tools while retaining authorship and interpersonal responsibility. A concrete starting point is to design assessment around reasoning, oral defense, peer

dialogue, revision, and real-world application. This creates an observable hypothesis about better work. For example, the organization can examine whether response quality, resolution time, safety reporting, understanding, and stakeholder confidence improve together. If one indicator improves while another deteriorates, leaders should resist declaring victory and investigate how benefit and burden are distributed.

Trust in automation is most useful when calibrated to capability and context rather than maximized without qualification (Lee & See, 2004; Glikson & Woolley, 2020). The implication is that social intelligence is not located exclusively inside an individual. It is also scaffolded—or obstructed—by meeting design, interface defaults, span of control, employment security, data access, professional norms, and leadership example. An attentive employee in a punitive system will eventually learn silence. A considerate manager with no discretion cannot provide meaningful accommodation. Capability development and structural redesign must therefore proceed together.

The caution is straightforward: personalization can narrow exposure and turn learning into a private transaction. Prevention begins before deployment by involving those who perform and receive the work. It continues after deployment through review of actual consequences, including consequences that the business case did not predict. The aim is institutional humility: confidence sufficient to act, coupled with channels that make correction possible. Restoring the human touch means making attention and response reliable, especially when standardized pathways fail.

19. Human Resources, People Analytics, and Algorithmic Management

Employment systems increasingly screen, rank, schedule, monitor, recommend, and discipline. This matters because administrative consistency can be gained at the cost of opacity, contextual blindness, and constrained voice. The important analytical move is to stop separating technical performance from the social conditions that make performance interpretable and legitimate. A system can be accurate in aggregate and still fail a particular encounter if the affected person cannot understand the decision, express relevant context, or obtain a remedy. Conversely, a warm interaction cannot redeem a process that is systematically unfair. Social intelligence joins these levels by asking what people notice, how they interpret it, what authority they possess, and whose consequences count.

The governing principle is that high-impact people decisions require validation, human review, explanation, and appeal. In practice, this means organizations should govern models across their life cycle and examine disparate outcomes as well as average accuracy. The objective is neither to keep people inside every minor transaction nor to automate every activity that appears standardized. It is to create a deliberate distribution of cognition, emotion, and accountability. Machines can search, calculate, draft, detect patterns, and maintain consistency. People can frame ambiguous purposes, recognize moral salience, negotiate competing values, and accept responsibility. The boundary must remain revisable because tasks, populations, and system capabilities change.

Research on social intelligence and interpersonal skill emphasizes that effective conduct depends on both accurate social perception and adaptive expression (Riggio, 1986; Silvera et al., 2001). The research should be read as a design resource rather than a promise that a single intervention will generalize everywhere. Context changes the meaning of behavior. A question that signals curiosity in one setting may signal distrust in another; rapid response may demonstrate service in one context and disrespectful impatience in another. Social intelligence therefore requires local inquiry, feedback from affected groups, and the willingness to revise routines when observed consequences differ from intention.

A central danger is that a human approver does not ensure meaningful oversight if workload, incentives, or interface design discourage review. The remedy is not another slogan about people first. It is evidence in the workflow: protected channels for dissent, explicit decision rights, sufficient time for conversation, meaningful human review, accessible explanations, and learning from overrides and failures. These arrangements convert human touch from an individual personality trait into an organizational capability. They also expose trade-offs. If a service promises empathy but budgets no time for listening, the operating model—not the employee's attitude—is the primary contradiction.

20. Diversity, Culture, Power, and Inclusion

Consider the proposition that social cues and acceptable behavior are interpreted through culture, language, disability, gender, generation, profession, and status. It changes the managerial question from “Can the tool perform the task?” to “What pattern of relationships will exist after the task is redesigned?” That second question

reveals dependencies that a process map may hide: who teaches the system, who absorbs exceptions, who explains outcomes, who notices unusual harm, and who can stop execution. Because systems trained on dominant patterns can treat difference as error or low performance, apparently minor interface and policy decisions can alter status, voice, and trust across an organization.

A workable response starts from this principle: social intelligence includes humility about one's own interpretive frame and attention to unequal power. Leaders can operationalize it when they include affected communities in design, testing, monitoring, and redress. This is a form of sociotechnical design. It recognizes that output quality arises from interactions among data, models, incentives, expertise, language, identity, and power. Social intelligence supplies the interpretive capacity to see these interactions and the behavioral capacity to act on them. Ethical judgment supplies the boundary: effectiveness cannot be defined solely by the objective of the system owner.

Studies of collective intelligence show that social sensitivity and participation patterns help explain team performance beyond the intelligence of individual members (Woolley et al., 2010). These findings point toward complementarity, but complementarity is not automatic. It must be produced through role clarity, practice with edge cases, feedback about outcomes, and authority to intervene. A person who sees a problem but cannot alter the decision is not providing oversight. A person who clicks approval without time or information is not exercising judgment. The quality of collaboration is therefore visible in what happens when human and machine assessments disagree.

The failure mode deserves equal attention: surface-level personalization can coexist with structural exclusion. Organizations should use this possibility as a testable risk hypothesis. They can look for early signals in complaint narratives, override patterns, participation data, turnover, workarounds, and the gap between formal procedures and actual behavior. Such inquiry is socially intelligent because it treats resistance not merely as an adoption problem but potentially as information about bad fit, lost identity, unrecognized workload, or compromised values.

21. Emotional Labor, Well-Being, and Burnout

The human touch requires emotional effort that is often demanded most intensely in stressful automated systems. Seen from the employee's position, the issue is

experienced through daily moments: a recommendation that seems impossible to question, a customer whose need does not fit the form, a manager who quotes a score instead of holding a conversation, or an expert who discovers that the underlying skill is fading. Seen from the institution's position, organizations may automate routine work and leave people with a concentrated portfolio of conflict, grief, uncertainty, and exception handling. The two views need to be held together; otherwise personal experience is dismissed as anecdote or aggregate data is dismissed as dehumanizing.

The core design commitment is that relational quality depends on staffing, recovery time, supervisor support, and legitimate boundaries. A concrete starting point is to measure emotional load and redesign capacity alongside customer-facing behavior. This creates an observable hypothesis about better work. For example, the organization can examine whether response quality, resolution time, safety reporting, understanding, and stakeholder confidence improve together. If one indicator improves while another deteriorates, leaders should resist declaring victory and investigate how benefit and burden are distributed.

Research on human–AI collaboration finds that value depends on task allocation, complementarity, and the conditions under which people can recognize and correct machine errors (Jarrahi, 2018; Fügner et al., 2021). The implication is that social intelligence is not located exclusively inside an individual. It is also scaffolded—or obstructed—by meeting design, interface defaults, span of control, employment security, data access, professional norms, and leadership example. An attentive employee in a punitive system will eventually learn silence. A considerate manager with no discretion cannot provide meaningful accommodation. Capability development and structural redesign must therefore proceed together.

The caution is straightforward: training people to appear empathetic without changing conditions transfers system failure into individual strain. Prevention begins before deployment by involving those who perform and receive the work. It continues after deployment through review of actual consequences, including consequences that the business case did not predict. The aim is institutional humility: confidence sufficient to act, coupled with channels that make correction possible. Restoring the human touch means making attention and response reliable, especially when standardized pathways fail.

22. Ethics, Privacy, Accountability, and Responsible Innovation

Social intelligence at institutional scale means anticipating effects, reflecting on values, engaging stakeholders, and responding to evidence. This matters because ethical principles often remain detached from procurement, product targets, and operational decisions. The important analytical move is to stop separating technical performance from the social conditions that make performance interpretable and legitimate. A system can be accurate in aggregate and still fail a particular encounter if the affected person cannot understand the decision, express relevant context, or obtain a remedy. Conversely, a warm interaction cannot redeem a process that is systematically unfair. Social intelligence joins these levels by asking what people notice, how they interpret it, what authority they possess, and whose consequences count.

The governing principle is that responsibility must be assigned to roles with authority, resources, and enforceable review rights. In practice, this means organizations should embed impact assessment, documentation, contestability, incident response, and sunset criteria in governance. The objective is neither to keep people inside every minor transaction nor to automate every activity that appears standardized. It is to create a deliberate distribution of cognition, emotion, and accountability. Machines can search, calculate, draft, detect patterns, and maintain consistency. People can frame ambiguous purposes, recognize moral salience, negotiate competing values, and accept responsibility. The boundary must remain revisable because tasks, populations, and system capabilities change.

The automation–augmentation paradox explains why the pursuit of substitution can weaken the human capabilities needed for effective augmentation (Raisch & Krakowski, 2021). The research should be read as a design resource rather than a promise that a single intervention will generalize everywhere. Context changes the meaning of behavior. A question that signals curiosity in one setting may signal distrust in another; rapid response may demonstrate service in one context and disrespectful impatience in another. Social intelligence therefore requires local inquiry, feedback from affected groups, and the willingness to revise routines when observed consequences differ from intention.

A central danger is that ethics washing is likely when values are public but decisions and remedies remain opaque. The remedy is not another slogan about people first. It is evidence in the workflow: protected channels for dissent, explicit decision rights,

sufficient time for conversation, meaningful human review, accessible explanations, and learning from overrides and failures. These arrangements convert human touch from an individual personality trait into an organizational capability. They also expose trade-offs. If a service promises empathy but budgets no time for listening, the operating model—not the employee's attitude—is the primary contradiction.

23. Measuring Social Intelligence Responsibly

Consider the proposition that measurement can support development when it distinguishes constructs, contexts, behavior, and outcomes. It changes the managerial question from “Can the tool perform the task?” to “What pattern of relationships will exist after the task is redesigned?” That second question reveals dependencies that a process map may hide: who teaches the system, who absorbs exceptions, who explains outcomes, who notices unusual harm, and who can stop execution. Because self-report scores are vulnerable to impression management and may be treated as fixed traits, apparently minor interface and policy decisions can alter status, voice, and trust across an organization.

A workable response starts from this principle: assessment should combine validated instruments with observation, multisource feedback, scenarios, and outcome evidence. Leaders can operationalize it when they use measures for a declared purpose, protect data, and avoid high-stakes interpretation beyond validity. This is a form of sociotechnical design. It recognizes that output quality arises from interactions among data, models, incentives, expertise, language, identity, and power. Social intelligence supplies the interpretive capacity to see these interactions and the behavioral capacity to act on them. Ethical judgment supplies the boundary: effectiveness cannot be defined solely by the objective of the system owner.

Evidence on psychological safety demonstrates that learning and error reporting require an interpersonal climate in which candor is not punished (Edmondson, 1999; Edmondson & Lei, 2014). These findings point toward complementarity, but complementarity is not automatic. It must be produced through role clarity, practice with edge cases, feedback about outcomes, and authority to intervene. A person who sees a problem but cannot alter the decision is not providing oversight. A person who clicks approval without time or information is not exercising judgment. The quality of collaboration is therefore visible in what happens when human and machine assessments disagree.

The failure mode deserves equal attention: quantification can narrow the construct and reward performative conformity. Organizations should use this possibility as a testable risk hypothesis. They can look for early signals in complaint narratives, override patterns, participation data, turnover, workarounds, and the gap between formal procedures and actual behavior. Such inquiry is socially intelligent because it treats resistance not merely as an adoption problem but potentially as information about bad fit, lost identity, unrecognized workload, or compromised values.

24. Developing Social Intelligence as an Organizational Capability

Social intelligence grows through deliberate practice, feedback, reflection, role modeling, and reinforcement in real work. Seen from the employee's position, the issue is experienced through daily moments: a recommendation that seems impossible to question, a customer whose need does not fit the form, a manager who quotes a score instead of holding a conversation, or an expert who discovers that the underlying skill is fading. Seen from the institution's position, one-off workshops decay when incentives, workload, and leadership behavior contradict the lesson. The two views need to be held together; otherwise personal experience is dismissed as anecdote or aggregate data is dismissed as dehumanizing.

The core design commitment is that development should connect individual skill with team routines and system redesign. A concrete starting point is to use simulations, coaching, peer observation, after-action reviews, and transfer projects. This creates an observable hypothesis about better work. For example, the organization can examine whether response quality, resolution time, safety reporting, understanding, and stakeholder confidence improve together. If one indicator improves while another deteriorates, leaders should resist declaring victory and investigate how benefit and burden are distributed.

Trust in automation is most useful when calibrated to capability and context rather than maximized without qualification (Lee & See, 2004; Glikson & Woolley, 2020). The implication is that social intelligence is not located exclusively inside an individual. It is also scaffolded—or obstructed—by meeting design, interface defaults, span of control, employment security, data access, professional norms, and leadership example. An attentive employee in a punitive system will eventually learn silence. A considerate manager with no discretion cannot provide meaningful

accommodation. Capability development and structural redesign must therefore proceed together.

The caution is straightforward: development language can individualize problems that require structural change. Prevention begins before deployment by involving those who perform and receive the work. It continues after deployment through review of actual consequences, including consequences that the business case did not predict. The aim is institutional humility: confidence sufficient to act, coupled with channels that make correction possible. Restoring the human touch means making attention and response reliable, especially when standardized pathways fail.

25. Workforce Strategy, Reskilling, and Fair Transition

Automation changes task bundles unevenly and creates both opportunities and transition costs. This matters because reskilling promises fail when workers lack time, access, credible pathways, or demand for the new skill. The important analytical move is to stop separating technical performance from the social conditions that make performance interpretable and legitimate. A system can be accurate in aggregate and still fail a particular encounter if the affected person cannot understand the decision, express relevant context, or obtain a remedy. Conversely, a warm interaction cannot redeem a process that is systematically unfair. Social intelligence joins these levels by asking what people notice, how they interpret it, what authority they possess, and whose consequences count.

The governing principle is that social intelligence should be integrated with technical and domain expertise in transparent career architectures. In practice, this means organizations should fund learning time, recognize prior experience, and link training to real mobility. The objective is neither to keep people inside every minor transaction nor to automate every activity that appears standardized. It is to create a deliberate distribution of cognition, emotion, and accountability. Machines can search, calculate, draft, detect patterns, and maintain consistency. People can frame ambiguous purposes, recognize moral salience, negotiate competing values, and accept responsibility. The boundary must remain revisable because tasks, populations, and system capabilities change.

Research on social intelligence and interpersonal skill emphasizes that effective conduct depends on both accurate social perception and adaptive expression (Riggio, 1986; Silvera et al., 2001). The research should be read as a design resource rather

than a promise that a single intervention will generalize everywhere. Context changes the meaning of behavior. A question that signals curiosity in one setting may signal distrust in another; rapid response may demonstrate service in one context and disrespectful impatience in another. Social intelligence therefore requires local inquiry, feedback from affected groups, and the willingness to revise routines when observed consequences differ from intention.

A central danger is that celebrating adaptability can shift all transition risk onto individuals. The remedy is not another slogan about people first. It is evidence in the workflow: protected channels for dissent, explicit decision rights, sufficient time for conversation, meaningful human review, accessible explanations, and learning from overrides and failures. These arrangements convert human touch from an individual personality trait into an organizational capability. They also expose trade-offs. If a service promises empathy but budgets no time for listening, the operating model—not the employee's attitude—is the primary contradiction.

26. Strategic Value and a Balanced Scorecard

Consider the proposition that the value of social intelligence appears in adoption, learning, safety, retention, customer recovery, innovation, and legitimacy. It changes the managerial question from “Can the tool perform the task?” to “What pattern of relationships will exist after the task is redesigned?” That second question reveals dependencies that a process map may hide: who teaches the system, who absorbs exceptions, who explains outcomes, who notices unusual harm, and who can stop execution. Because organizations overemphasize short-term throughput because relational value is distributed and delayed, apparently minor interface and policy decisions can alter status, voice, and trust across an organization.

A workable response starts from this principle: a balanced case should connect human indicators to operational, financial, and risk outcomes. Leaders can operationalize it when they track leading behaviors, system conditions, stakeholder experience, and lagging performance together. This is a form of sociotechnical design. It recognizes that output quality arises from interactions among data, models, incentives, expertise, language, identity, and power. Social intelligence supplies the interpretive capacity to see these interactions and the behavioral capacity to act on them. Ethical judgment supplies the boundary: effectiveness cannot be defined solely by the objective of the system owner.

Studies of collective intelligence show that social sensitivity and participation patterns help explain team performance beyond the intelligence of individual members (Woolley et al., 2010). These findings point toward complementarity, but complementarity is not automatic. It must be produced through role clarity, practice with edge cases, feedback about outcomes, and authority to intervene. A person who sees a problem but cannot alter the decision is not providing oversight. A person who clicks approval without time or information is not exercising judgment. The quality of collaboration is therefore visible in what happens when human and machine assessments disagree.

The failure mode deserves equal attention: poorly chosen metrics can invite gaming and convert relationships into transactions. Organizations should use this possibility as a testable risk hypothesis. They can look for early signals in complaint narratives, override patterns, participation data, turnover, workarounds, and the gap between formal procedures and actual behavior. Such inquiry is socially intelligent because it treats resistance not merely as an adoption problem but potentially as information about bad fit, lost identity, unrecognized workload, or compromised values.

27. An Implementation Roadmap

Capability building must move from diagnosis through design, experimentation, governance, scaling, and renewal. Seen from the employee's position, the issue is experienced through daily moments: a recommendation that seems impossible to question, a customer whose need does not fit the form, a manager who quotes a score instead of holding a conversation, or an expert who discovers that the underlying skill is fading. Seen from the institution's position, large transformation programs frequently launch tools before agreeing on human outcomes or decision rights. The two views need to be held together; otherwise personal experience is dismissed as anecdote or aggregate data is dismissed as dehumanizing.

The core design commitment is that implementation should begin with a consequential work journey and expand through evidence. A concrete starting point is to create a cross-functional coalition with worker and stakeholder participation. This creates an observable hypothesis about better work. For example, the organization can examine whether response quality, resolution time, safety reporting, understanding, and stakeholder confidence improve together. If one

indicator improves while another deteriorates, leaders should resist declaring victory and investigate how benefit and burden are distributed.

Research on human–AI collaboration finds that value depends on task allocation, complementarity, and the conditions under which people can recognize and correct machine errors (Jarrahi, 2018; Fügener et al., 2021). The implication is that social intelligence is not located exclusively inside an individual. It is also scaffolded—or obstructed—by meeting design, interface defaults, span of control, employment security, data access, professional norms, and leadership example. An attentive employee in a punitive system will eventually learn silence. A considerate manager with no discretion cannot provide meaningful accommodation. Capability development and structural redesign must therefore proceed together.

The caution is straightforward: scaling a weak social design multiplies harm as efficiently as it multiplies value. Prevention begins before deployment by involving those who perform and receive the work. It continues after deployment through review of actual consequences, including consequences that the business case did not predict. The aim is institutional humility: confidence sufficient to act, coupled with channels that make correction possible. Restoring the human touch means making attention and response reliable, especially when standardized pathways fail.

28. Sector Narratives: Social Intelligence in Practice

The principles become visible when translated into concrete choices in banking, manufacturing, health care, education, public service, and small enterprise. This matters because abstract claims allow organizations to agree while imagining incompatible practices. The important analytical move is to stop separating technical performance from the social conditions that make performance interpretable and legitimate. A system can be accurate in aggregate and still fail a particular encounter if the affected person cannot understand the decision, express relevant context, or obtain a remedy. Conversely, a warm interaction cannot redeem a process that is systematically unfair. Social intelligence joins these levels by asking what people notice, how they interpret it, what authority they possess, and whose consequences count.

The governing principle is that cases should reveal tensions, decisions, stakeholder experiences, and feedback loops. In practice, this means organizations should use scenario rehearsal before deployment and after-action review after incidents. The

objective is neither to keep people inside every minor transaction nor to automate every activity that appears standardized. It is to create a deliberate distribution of cognition, emotion, and accountability. Machines can search, calculate, draft, detect patterns, and maintain consistency. People can frame ambiguous purposes, recognize moral salience, negotiate competing values, and accept responsibility. The boundary must remain revisable because tasks, populations, and system capabilities change.

The automation–augmentation paradox explains why the pursuit of substitution can weaken the human capabilities needed for effective augmentation (Raisch & Krakowski, 2021). The research should be read as a design resource rather than a promise that a single intervention will generalize everywhere. Context changes the meaning of behavior. A question that signals curiosity in one setting may signal distrust in another; rapid response may demonstrate service in one context and disrespectful impatience in another. Social intelligence therefore requires local inquiry, feedback from affected groups, and the willingness to revise routines when observed consequences differ from intention.

A central danger is that success stories become misleading when trade-offs and displaced work are omitted. The remedy is not another slogan about people first. It is evidence in the workflow: protected channels for dissent, explicit decision rights, sufficient time for conversation, meaningful human review, accessible explanations, and learning from overrides and failures. These arrangements convert human touch from an individual personality trait into an organizational capability. They also expose trade-offs. If a service promises empathy but budgets no time for listening, the operating model—not the employee's attitude—is the primary contradiction.

29. Research Agenda and Unresolved Questions

Consider the proposition that rapid diffusion of generative and agentic systems is changing interaction faster than institutions can accumulate evidence. It changes the managerial question from “Can the tool perform the task?” to “What pattern of relationships will exist after the task is redesigned?” That second question reveals dependencies that a process map may hide: who teaches the system, who absorbs exceptions, who explains outcomes, who notices unusual harm, and who can stop execution. Because short experiments cannot reveal long-term effects on expertise, relationships, identity, and inequality, apparently minor interface and policy decisions can alter status, voice, and trust across an organization.

A workable response starts from this principle: research should combine field experiments, longitudinal studies, qualitative inquiry, and participatory methods. Leaders can operationalize it when they publish failures, distributional effects, and boundary conditions as well as average gains. This is a form of sociotechnical design. It recognizes that output quality arises from interactions among data, models, incentives, expertise, language, identity, and power. Social intelligence supplies the interpretive capacity to see these interactions and the behavioral capacity to act on them. Ethical judgment supplies the boundary: effectiveness cannot be defined solely by the objective of the system owner.

Evidence on psychological safety demonstrates that learning and error reporting require an interpersonal climate in which candor is not punished (Edmondson, 1999; Edmondson & Lei, 2014). These findings point toward complementarity, but complementarity is not automatic. It must be produced through role clarity, practice with edge cases, feedback about outcomes, and authority to intervene. A person who sees a problem but cannot alter the decision is not providing oversight. A person who clicks approval without time or information is not exercising judgment. The quality of collaboration is therefore visible in what happens when human and machine assessments disagree.

The failure mode deserves equal attention: claims of inevitability can close inquiry precisely when plural futures remain possible. Organizations should use this possibility as a testable risk hypothesis. They can look for early signals in complaint narratives, override patterns, participation data, turnover, workarounds, and the gap between formal procedures and actual behavior. Such inquiry is socially intelligent because it treats resistance not merely as an adoption problem but potentially as information about bad fit, lost identity, unrecognized workload, or compromised values.

30. Conclusion: Human Touch as Designed Capacity

Humanity in automated work is produced through choices about attention, authority, time, care, and accountability. Seen from the employee's position, the issue is experienced through daily moments: a recommendation that seems impossible to question, a customer whose need does not fit the form, a manager who quotes a score instead of holding a conversation, or an expert who discovers that the underlying skill is fading. Seen from the institution's position, nostalgia for pre-digital

work is no more useful than technological determinism. The two views need to be held together; otherwise personal experience is dismissed as anecdote or aggregate data is dismissed as dehumanizing.

The core design commitment is that the goal is an institution in which machines expand capability while people retain dignity, agency, and responsibility. A concrete starting point is to treat every automation decision as a simultaneous decision about relationships and social learning. This creates an observable hypothesis about better work. For example, the organization can examine whether response quality, resolution time, safety reporting, understanding, and stakeholder confidence improve together. If one indicator improves while another deteriorates, leaders should resist declaring victory and investigate how benefit and burden are distributed.

Trust in automation is most useful when calibrated to capability and context rather than maximized without qualification (Lee & See, 2004; Glikson & Woolley, 2020). The implication is that social intelligence is not located exclusively inside an individual. It is also scaffolded—or obstructed—by meeting design, interface defaults, span of control, employment security, data access, professional norms, and leadership example. An attentive employee in a punitive system will eventually learn silence. A considerate manager with no discretion cannot provide meaningful accommodation. Capability development and structural redesign must therefore proceed together.

The caution is straightforward: the human touch disappears when everyone assumes somebody else is responsible for preserving it. Prevention begins before deployment by involving those who perform and receive the work. It continues after deployment through review of actual consequences, including consequences that the business case did not predict. The aim is institutional humility: confidence sufficient to act, coupled with channels that make correction possible. Restoring the human touch means making attention and response reliable, especially when standardized pathways fail.

Appendix A. Capability Profiles for Development and Observation

The profiles below are developmental descriptions, not a diagnostic license. They should be adapted to role and culture, observed across situations, and combined with direct feedback. None should be used as a single high-stakes score. Each profile links a human capacity to an automation-era risk and a concrete learning practice.

A.1 Attentive presence

Meaning: Attentive presence is sustaining undivided attention long enough to understand another person's purpose. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person summarizes before proposing; notices hesitation; reduces avoidable interruption. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Notification-heavy workflows reward fragmented attention. The development response is to practice short listening rounds and redesign meetings around explicit attention. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.2 Social cue recognition

Meaning: Social cue recognition is detecting verbal, paralinguistic, temporal, and participation signals. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person notices changes in tone, silence, pace, and who has stopped contributing. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Text-only and summarized communication removes contextual cues. The development response is to compare interpretations with the speaker and keep uncertainty visible. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.3 Perspective taking

Meaning: Perspective taking is constructing a disciplined account of how a situation may appear from another position. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person asks about constraints and incentives before assigning motive. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Personalized systems can reinforce the user's existing frame. The development response is to use stakeholder role reversal in scenario exercises. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.4 Empathic concern

Meaning: Empathic concern is allowing another person's welfare to matter in the choice of response. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person offers proportionate support while maintaining fair boundaries. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Synthetic empathy can simulate concern without responsibility. The development response is to connect listening to a concrete commitment, owner, and follow-up. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.5 Emotional self-awareness

Meaning: Emotional self-awareness is recognizing one's affective state and its influence on interpretation. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person names tension or urgency without projecting it onto others. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Accelerated workflows reduce time to notice emotion. The development response is to use a brief signal-interpretation-emotion-response journal. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.6 Self-regulation

Meaning: Self-regulation is choosing a constructive response under pressure rather than enacting the first impulse. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person pauses, checks evidence, and selects an appropriate channel. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Alerts and rankings encourage immediate reaction. The development response is to introduce deliberation windows for consequential communication. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.7 Epistemic humility

Meaning: Epistemic humility is recognizing the limits of personal and machine knowledge. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person states uncertainty and seeks disconfirming evidence. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Fluent output invites unwarranted confidence. The development response is to require confidence statements and alternative hypotheses. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.8 Contextual judgment

Meaning: Contextual judgment is integrating formal rules with relevant features of a particular case. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person explains why an exception is or is not warranted. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Standardization treats context as noise. The development response is to create governed exception pathways with learning reviews. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.9 Clear expression

Meaning: Clear expression is making intent, reasoning, uncertainty, and requested action understandable. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person adapts language without hiding material information. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Machine drafting can produce polished ambiguity. The development response is to review high-stakes messages for ownership, audience, and consequence. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.10 Dialogic inquiry

Meaning: Dialogic inquiry is using questions to generate shared understanding rather than to perform interrogation. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person asks open and specific questions and follows the answer. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Chat interfaces may encourage rapid answer retrieval. The development response is to train managers to hold inquiry before advocacy. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.11 Constructive disagreement

Meaning: Constructive disagreement is challenging ideas while protecting dignity and shared purpose. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person makes evidence and assumptions discussable. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Algorithmic authority can suppress dissent. The development response is to use pre-mortems, red teams, and protected challenge roles. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.12 Conflict navigation

Meaning: Conflict navigation is identifying interests, emotions, identity threats, and power in disagreement. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person separates the person from the contested problem. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Remote and automated exchange can escalate misinterpretation. The development response is to move complex conflict into facilitated synchronous conversation. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.13 Relationship repair

Meaning: Relationship repair is acknowledging harm, restoring understanding, and changing future behavior. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person offers a specific apology and verifies the remedy. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Automated channels make responsibility difficult to locate. The development response is to define a human owner for service and decision recovery. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.14 Trust calibration

Meaning: Trust calibration is matching reliance to evidence about competence, limits, and context. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person verifies more when stakes or novelty rise. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Automation bias and blanket distrust both distort judgment. The development response is to teach users to recognize system boundaries and failure modes. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.15 Coordination literacy

Meaning: Coordination literacy is understanding dependencies, handoffs, timing, and shared representations. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person anticipates how a local action affects downstream work. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Automation can hide displaced tasks. The development response is to map exception work and invisible coordination before scaling. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.16 Inclusive facilitation

Meaning: Inclusive facilitation is creating participation conditions for people with different status and communication styles. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person uses multiple channels and checks whose view is absent. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Dominant patterns are easier for systems to recognize. The development response is to audit participation and rotate agenda-setting roles. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.17 Cultural humility

Meaning: Cultural humility is approaching cultural interpretation as inquiry rather than mastery. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person tests assumptions and invites correction. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Models may universalize dominant linguistic norms. The development response is to include local reviewers and affected communities. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.18 Power awareness

Meaning: Power awareness is recognizing how authority and dependency affect what can safely be said. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person does not interpret silence as consent. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Surveillance and ranking increase perceived risk of voice. The development response is to separate developmental feedback from punitive monitoring. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.19 Ethical discernment

Meaning: Ethical discernment is recognizing when efficiency objectives conflict with rights, fairness, or dignity. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person articulates the affected values and escalation need. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Optimization can hide normative choices. The development response is to use impact assessment tied to enforceable decision gates. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.20 Moral courage

Meaning: Moral courage is acting on ethical judgment despite interpersonal or career pressure. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person raises evidence and seeks allies through legitimate channels. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Diffused responsibility makes inaction easy. The development response is to protect whistleblowing and reward well-founded challenge. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.21 Coaching orientation

Meaning: Coaching orientation is helping others construct capability rather than merely supplying answers. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person asks, demonstrates, observes, and gives specific feedback. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Generative tools can bypass productive struggle. The development response is to use AI output as material for critique and revision. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.22 Learning agility

Meaning: Learning agility is updating behavior from evidence across unfamiliar situations. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person runs small experiments and revises assumptions. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Rapid tool cycles can encourage novelty without learning. The development response is to institutionalize after-action review and transfer logs. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.23 Boundary management

Meaning: Boundary management is protecting attention, recovery, privacy, and professional limits. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person says no with reason and offers a viable alternative. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Always-on systems normalize availability. The development response is to define response windows and protected human recovery time. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

A.24 Accountable judgment

Meaning: Accountable judgment is owning a consequential decision and its explanation. It matters because automated work changes which signals remain available and how rapidly a response is expected. A practical indicator is that the person can name the decision maker, evidence, trade-off, and remedy. The indicator should be interpreted in context; quiet attention, direct challenge, or emotional restraint may carry different meanings across cultures, roles, and power relationships.

Automation tension: Human review can be reduced to rubber-stamping. The development response is to record material reasoning and maintain an appeal path. Practice should use a realistic work episode, followed by feedback on perception, interpretation, behavior, and consequence. Learners should identify what they inferred, which evidence they used, what alternative account remained possible, and how their action affected dignity, task progress, and future trust. Team leaders can reinforce transfer by making the behavior visible in reviews and by removing workflow conditions that punish it.

Appendix B. Twelve-Stage Implementation Roadmap

The stages form a learning cycle rather than a rigid waterfall. An organization may return to framing after a pilot reveals that the original problem was poorly defined. Each stage should have a named owner, participation from people who perform and receive the work, evidence for moving forward, and authority to pause.

B.1 Frame the human outcome

Purpose and action: Choose one work journey and state whose experience should improve, which relationship should be strengthened, and which harm must be prevented. The stage is complete only when the relevant people can describe what changed, what uncertainty remains, and how someone can challenge the current design. Leaders should document the trade-off being made, the groups likely to experience benefit or burden, and the signal that would trigger reconsideration. This discipline keeps implementation adaptive while preventing vague commitments from substituting for decision.

B.2 Map decisions and dependencies

Purpose and action: Trace data, recommendations, handoffs, exceptions, emotional load, authority, and recourse from beginning to end. The stage is complete only when the relevant people can describe what changed, what uncertainty remains, and how someone can challenge the current design. Leaders should document the trade-off being made, the groups likely to experience benefit or burden, and the signal that would trigger reconsideration. This discipline keeps implementation adaptive while preventing vague commitments from substituting for decision.

B.3 Listen to affected groups

Purpose and action: Use interviews, observation, complaint narratives, and participatory workshops to understand actual rather than imagined work. The stage is complete only when the relevant people can describe what changed, what uncertainty remains, and how someone can challenge the current design. Leaders should document the trade-off being made, the groups likely to experience benefit or burden, and the signal that would trigger reconsideration. This discipline keeps implementation adaptive while preventing vague commitments from substituting for decision.

B.4 Establish a baseline

Purpose and action: Measure current operational outcomes, stakeholder experience, voice, workload, trust calibration, and distributional differences. The stage is complete only when the relevant people can describe what changed, what uncertainty remains, and how someone can challenge the current design. Leaders should document the trade-off being made, the groups likely to experience benefit or burden, and the signal that would trigger reconsideration. This discipline keeps implementation adaptive while preventing vague commitments from substituting for decision.

B.5 Define decision rights

Purpose and action: Specify who recommends, decides, executes, explains, monitors, overrides, and repairs each consequential action. The stage is complete only when the relevant people can describe what changed, what uncertainty remains, and how someone can challenge the current design. Leaders should document the trade-off being made, the groups likely to experience benefit or burden, and the signal that would trigger reconsideration. This discipline keeps implementation adaptive while preventing vague commitments from substituting for decision.

B.6 Redesign roles

Purpose and action: Combine technical, domain, and relational responsibilities so that people retain comprehension, agency, and learning. The stage is complete only when the relevant people can describe what changed, what uncertainty remains, and how someone can challenge the current design. Leaders should document the trade-off being made, the groups likely to experience benefit or burden, and the signal that would trigger reconsideration. This discipline keeps implementation adaptive while preventing vague commitments from substituting for decision.

B.7 Build capability

Purpose and action: Use practice with realistic scenarios, feedback, coaching, peer reflection, and on-the-job transfer assignments. The stage is complete only when the relevant people can describe what changed, what uncertainty remains, and how someone can challenge the current design. Leaders should document the trade-off being made, the groups likely to experience benefit or burden, and the signal that would trigger reconsideration. This discipline keeps implementation adaptive while preventing vague commitments from substituting for decision.

B.8 Create governance gates

Purpose and action: Connect impact assessment, privacy, security, fairness, safety, and labor considerations to real approval authority. The stage is complete only when the relevant people can describe what changed, what uncertainty remains, and how someone can challenge the current design. Leaders should document the trade-off being made, the groups likely to experience benefit or burden, and the signal that would trigger reconsideration. This discipline keeps implementation adaptive while preventing vague commitments from substituting for decision.

B.9 Pilot with edge cases

Purpose and action: Test not only average flow but also ambiguity, vulnerability, disagreement, failure, and the need for human escalation. The stage is complete only when the relevant people can describe what changed, what uncertainty remains, and how someone can challenge the current design. Leaders should document the trade-off being made, the groups likely to experience benefit or burden, and the signal that would trigger reconsideration. This discipline keeps implementation adaptive while preventing vague commitments from substituting for decision.

B.10 Evaluate mixed evidence

Purpose and action: Combine performance data with qualitative experience and examine who receives benefit and who absorbs new work. The stage is complete only when the relevant people can describe what changed, what uncertainty remains, and how someone can challenge the current design. Leaders should document the trade-off being made, the groups likely to experience benefit or burden, and the signal that would trigger reconsideration. This discipline keeps implementation adaptive while preventing vague commitments from substituting for decision.

B.11 Scale through standards and communities

Purpose and action: Codify proven patterns while maintaining local discretion, practitioner learning, and reporting of failure. The stage is complete only when the relevant people can describe what changed, what uncertainty remains, and how someone can challenge the current design. Leaders should document the trade-off being made, the groups likely to experience benefit or burden, and the signal that would trigger reconsideration. This discipline keeps implementation adaptive while preventing vague commitments from substituting for decision.

B.12 Renew and retire

Purpose and action: Reassess models, roles, controls, and learning needs as technology and context change; retire systems whose value no longer exceeds their harm. The stage is complete only when the relevant people can describe what changed, what uncertainty remains, and how someone can challenge the current design. Leaders should document the trade-off being made, the groups likely to experience benefit or burden, and the signal that would trigger reconsideration. This discipline keeps implementation adaptive while preventing vague commitments from substituting for decision.

Appendix C. Applied Sector Scenarios

These short narratives are not universal prescriptions. Their purpose is to make design choices visible and to support scenario-based discussion among leaders, workers, technical teams, customers, regulators, and community representatives.

C.1 Retail banking

Situation: An automated fraud model blocks a transaction by a long-standing customer who is traveling for family care. **Social failure:** A service officer sees only a risk score and a script. **Socially intelligent redesign:** The bank displays uncertainty and evidence, authorizes a trained human to investigate context, provides rapid appeal, and records whether the case reveals a new fraud pattern or a model limitation. The critical distinction is between adding a courteous human at the end of a broken process and changing the process so that attention, discretion, explanation, and learning are available where they matter. Evaluation should include resolution quality, time, stakeholder understanding, worker load, distributional effects, and whether a similar case is handled better next time.

C.2 Manufacturing

Situation: Predictive maintenance reschedules work and implicitly devalues the craft knowledge of experienced technicians. **Social failure:** Experts comply publicly but withhold observations that the model cannot capture. **Socially intelligent redesign:** Leaders ask technicians to annotate anomalies, give them visible influence over maintenance decisions, and treat disagreement as joint diagnosis rather than resistance. The critical distinction is between adding a courteous human at the end of a broken process and changing the process so that attention, discretion, explanation, and learning are available where they matter. Evaluation should include resolution quality, time, stakeholder understanding, worker load, distributional effects, and whether a similar case is handled better next time.

C.3 Hospital care

Situation: A triage system prioritizes cases accurately on average but a patient cannot explain a complex social circumstance through the intake form. **Social failure:** The standardized pathway creates delay and fear. **Socially intelligent redesign:** A nurse can trigger contextual review, explain the process in plain language, and ensure that automation creates time for presence rather than another documentation layer. The critical distinction is between adding a courteous human at the end of a broken process and changing the process so that attention, discretion, explanation, and learning are available where they matter. Evaluation should include resolution quality, time, stakeholder understanding, worker load, distributional effects, and whether a similar case is handled better next time.

C.4 University education

Situation: Students use generative AI to produce fluent assignments while faculty struggle to evaluate learning. **Social failure:** The institution responds with surveillance and broad prohibition. **Socially intelligent redesign:** Courses redesign assessment around drafts, oral explanation, peer dialogue, documented tool use, and reflection on errors, preserving authorship while teaching responsible collaboration. The critical distinction is between adding a courteous human at the end of a broken process and changing the process so that attention, discretion, explanation, and learning are available where they matter. Evaluation should include resolution quality, time, stakeholder understanding, worker load, distributional effects, and whether a similar case is handled better next time.

C.5 Public benefits

Situation: An eligibility system requests additional evidence from a citizen with limited digital access. **Social failure:** No individual appears to own the consequence of non-completion. **Socially intelligent redesign:** The agency offers assisted channels, explains the requirement, logs recurring barriers, and uses case evidence to redesign policy rather than labeling applicants noncompliant. The critical distinction is between adding a courteous human at the end of a broken process and changing the process so that attention, discretion, explanation, and learning are available where they matter. Evaluation should include resolution quality, time, stakeholder understanding, worker load, distributional effects, and whether a similar case is handled better next time.

C.6 Logistics platform

Situation: Optimization assigns demanding routes and measures drivers against predicted times. **Social failure:** Traffic, disability access, and customer interaction are treated as deviations. **Socially intelligent redesign:** Drivers participate in model review, unsafe estimates can be challenged without penalty, and route quality includes safety and service recovery rather than speed alone. The critical distinction is between adding a courteous human at the end of a broken process and changing the process so that attention, discretion, explanation, and learning are available where they matter. Evaluation should include resolution quality, time, stakeholder understanding, worker load, distributional effects, and whether a similar case is handled better next time.

C.7 Small enterprise

Situation: A family-owned retailer adopts an AI assistant for marketing, inventory, and customer messages. **Social failure:** The owner saves time but the brand voice becomes generic and customers disclose sensitive information to a tool they assume is human.

Socially intelligent redesign: The business discloses automation, reviews consequential messages, protects data, and reserves personal contact for complaints and important relationships. The critical distinction is between adding a courteous human at the end of a broken process and changing the process so that attention, discretion, explanation, and learning are available where they matter. Evaluation should include resolution quality, time, stakeholder understanding, worker load, distributional effects, and whether a similar case is handled better next time.

C.8 Human resources

Situation: A recruiting model ranks applicants and reduces screening time. **Social failure:** Recruiters become anchored to the score and unconventional candidates disappear.

Socially intelligent redesign: The firm validates job relevance, withholds scores during an independent first review, audits group outcomes, documents overrides, and gives applicants a meaningful route to clarification. The critical distinction is between adding a courteous human at the end of a broken process and changing the process so that attention, discretion, explanation, and learning are available where they matter. Evaluation should include resolution quality, time, stakeholder understanding, worker load, distributional effects, and whether a similar case is handled better next time.

Appendix D. Measurement Framework

Measurement should begin with a theory of value. Social intelligence contributes through better interpretation, coordination, learning, trust calibration, inclusion, recovery, and ethical judgment. These mechanisms influence outcomes such as safety, adoption, service quality, innovation, retention, and legitimacy. A dashboard should therefore combine four levels. Capability indicators capture observable behaviors and scenario performance. System indicators capture workload, decision rights, access to escalation, and response to voice. Experience indicators capture understanding, fairness, dignity, and confidence. Outcome indicators capture operational, financial, learning, and risk performance.

No single metric should become a target detached from its purpose. A rising satisfaction score may hide customers who abandoned the process; a falling override rate may indicate better accuracy or suppressed dissent; faster handling may move work to another unit or to the customer. Teams should review indicator pairs and narratives. Recommended pairs include speed with first-contact resolution, adoption with appropriate override, customer effort with employee emotional load, model accuracy with subgroup error, voice frequency with response quality, and productivity with learning opportunity.

Evaluation should use a baseline, a comparison where feasible, repeated measurement, and explicit attention to distribution. Qualitative inquiry is not merely illustrative; it identifies mechanisms and unintended effects that were not encoded in the original measure. Privacy and proportionality are essential. Social-interaction data can be intimate and easily misinterpreted. Organizations should minimize collection, separate development from discipline, restrict access, set retention limits, and explain what is measured and why.

Glossary

Affective empathy: The capacity to resonate emotionally with another person's state; distinct from perspective taking and compassionate action.

Agentic AI: AI systems designed to pursue goals through sequences of actions with a degree of operational autonomy.

Algorithmic management: The use of data and algorithms to allocate, monitor, evaluate, or discipline work.

Automation bias: The tendency to favor automated recommendations even when contrary evidence is available.

Augmentation: The use of technology to enlarge human capability rather than simply substitute for human action.

Calibrated trust: Reliance matched to a system's demonstrated competence, limitations, and operating context.

Collective intelligence: A group's general capacity to perform across diverse tasks, shaped partly by interaction quality and social sensitivity.

Compassionate concern: Motivation to support another person's welfare in response to perceived need.

Contestability: The practical ability to question, appeal, or seek revision of a decision.

Cultural intelligence: The capability to function effectively across culturally diverse contexts.

De-skilling: Erosion of expertise or judgment when meaningful parts of a task are removed or rarely practiced.

Emotional intelligence: Abilities related to perceiving, using, understanding, and managing emotion.

Emotional labor: Effort involved in regulating emotional display to meet occupational expectations.

Empathy: A family of processes involving understanding, sharing, or responding to another person's experience.

Explainability: The degree to which relevant people can understand the basis, limits, and implications of a system output.

Human-in-the-loop: A configuration in which a person participates in some stage of an automated process; meaningfulness depends on authority, time, information, and competence.

Human touch: Reliable human attention, understanding, care, judgment, and accountability in an interaction or institution.

Job crafting: Employee-initiated changes to tasks, relationships, or interpretations that make work more workable or meaningful.

Moral agency: The capacity and responsibility to make choices on the basis of ethical considerations.

Perspective taking: The cognitive effort to understand how a situation may appear from another person's position.

Psychological safety: A shared belief that interpersonal risk taking, such as asking questions or admitting mistakes, is safe in a team.

Responsible innovation: An approach that anticipates effects, reflects on purposes and values, includes stakeholders, and responds to evidence.

Social intelligence: The integrated capacity to perceive social situations, interpret people and context, regulate oneself, and act effectively and ethically with others.

Social sensitivity: Accuracy and responsiveness in perceiving other people's states and interpersonal cues.

Sociotechnical system: A work system whose outcomes emerge from the joint design of people, technology, structure, tasks, and environment.

Task allocation: The deliberate distribution of activities and decision rights among people and technological systems.

Voice: The expression of ideas, concerns, or challenge intended to improve an organization or prevent harm.

References

- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
- Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3–30. <https://doi.org/10.1257/jep.29.3.3>
- Avolio, B. J., Sosik, J. J., Kahai, S. S., & Baker, B. (2014). E-leadership: Re-examining transformations in leadership source and transmission. *The Leadership Quarterly*, 25(1), 105–131. <https://doi.org/10.1016/j.leaqua.2013.11.003>
- Brown, M. E., Treviño, L. K., & Harrison, D. A. (2005). Ethical leadership: A social learning perspective for construct development and testing. *Organizational Behavior and Human Decision Processes*, 97(2), 117–134. <https://doi.org/10.1016/j.obhdp.2005.03.002>
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work (NBER Working Paper No. 31161). National Bureau of Economic Research. <https://doi.org/10.3386/w31161>
- Cascio, W. F., & Montealegre, R. (2016). How technology is changing work and organizations. *Annual Review of Organizational Psychology and Organizational Behavior*, 3, 349–375. <https://doi.org/10.1146/annurev-orgpsych-041015-062352>
- Davis, M. H. (1983). Measuring individual differences in empathy: Evidence for a multidimensional approach. *Journal of Personality and Social Psychology*, 44(1), 113–126. <https://doi.org/10.1037/0022-3514.44.1.113>
- Deci, E. L., Olafsen, A. H., & Ryan, R. M. (2017). Self-determination theory in work organizations: The state of a science. *Annual Review of Organizational Psychology and Organizational Behavior*, 4, 19–43. <https://doi.org/10.1146/annurev-orgpsych-032516-113108>
- Dellermann, D., Ebel, P., Söllner, M., & Leimeister, J. M. (2019). Hybrid intelligence. *Business & Information Systems Engineering*, 61, 637–643. <https://doi.org/10.1007/s12599-019-00595-2>
- Demerouti, E., Bakker, A. B., Nachreiner, F., & Schaufeli, W. B. (2001). The job demands-resources model of burnout. *Journal of Applied Psychology*, 86(3), 499–512. <https://doi.org/10.1037/0021-9010.86.3.499>
- Dignum, V. (2019). *Responsible artificial intelligence: How to develop and use AI in a responsible way*. Springer. <https://doi.org/10.1007/978-3-030-30371-6>
- Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28), eadn5290. <https://doi.org/10.1126/sciadv.adn5290>
- Edmondson, A. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2), 350–383. <https://doi.org/10.2307/2666999>
- Edmondson, A. C., & Lei, Z. (2014). Psychological safety: The history, renaissance, and future of an interpersonal construct. *Annual Review of Organizational Psychology and Organizational Behavior*, 1, 23–43. <https://doi.org/10.1146/annurev-orgpsych-031413-091305>
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280. <https://doi.org/10.1016/j.techfore.2016.08.019>

- Fügener, A., Grahl, J., Gupta, A., & Ketter, W. (2021). Will humans-in-the-loop become borgs? Merits and pitfalls of working with AI. *Organizational Behavior and Human Decision Processes*, 167, 24–38. <https://doi.org/10.1016/j.obhdp.2021.02.005>
- Gittel, J. H. (2000). Organizing work to support relational co-ordination. *Organization Science*, 11(5), 517–539. <https://doi.org/10.1287/orsc.11.5.517.15217>
- Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627–660. <https://doi.org/10.5465/annals.2018.0057>
- Gmyrek, P., Berg, J., Bescond, D., Gubernot, D., & Horne, R. (2025). *Generative AI and jobs: A refined global index of occupational exposure* (ILO Working Paper 140). International Labour Organization. <https://doi.org/10.54394/HETP0387>
- Grandey, A. A. (2000). Emotion regulation in the workplace: A new way to conceptualize emotional labor. *Journal of Occupational Health Psychology*, 5(1), 95–110. <https://doi.org/10.1037/1076-8998.5.1.95>
- Green, A. (2024). Artificial intelligence and the changing demand for skills in the labour market. *OECD Artificial Intelligence Papers*, No. 14. OECD Publishing. <https://doi.org/10.1787/88684e36-en>
- Hancock, J. T., Naaman, M., & Levy, K. (2020). AI-mediated communication: Definition, research agenda, and ethical considerations. *Journal of Computer-Mediated Communication*, 25(1), 89–100. <https://doi.org/10.1093/jcmc/zmz022>
- Huang, M.-H., & Rust, R. T. (2018). Artificial intelligence in service. *Journal of Service Research*, 21(2), 155–172. <https://doi.org/10.1177/1094670517752459>
- Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human–AI symbiosis in organizational decision making. *Business Horizons*, 61(4), 577–586. <https://doi.org/10.1016/j.bushor.2018.03.007>
- Joseph, D. L., & Newman, D. A. (2010). Emotional intelligence: An integrative meta-analysis and cascading model. *Journal of Vocational Behavior*, 76(1), 54–78. <https://doi.org/10.1016/j.jvb.2009.10.007>
- Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410. <https://doi.org/10.5465/annals.2018.0174>
- Lane, M., Williams, M., & Broecke, S. (2023). The impact of AI on the workplace: Main findings from the OECD AI surveys of employers and workers. *OECD Social, Employment and Migration Working Papers*, No. 288. OECD Publishing. <https://doi.org/10.1787/ea0a0fe1-en>
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392
- Lewis, K. (2003). Measuring transactive memory systems in the field: Scale development and validation. *Journal of Applied Psychology*, 88(4), 587–604. <https://doi.org/10.1037/0021-9010.88.4.587>
- MacCann, C., Jiang, Y., Brown, L. E. R., Double, K. S., Bucich, M., & Minbashian, A. (2020). Emotional intelligence predicts academic performance: A meta-analysis. *Psychological Bulletin*, 146(2), 150–186. <https://doi.org/10.1037/bul0000219>
- Maslach, C., & Leiter, M. P. (2016). Understanding the burnout experience: Recent research and its implications for psychiatry. *World Psychiatry*, 15(2), 103–111. <https://doi.org/10.1002/wps.20311>

- Mayer, J. D., Salovey, P., & Caruso, D. R. (2004). Emotional intelligence: Theory, findings, and implications. *Psychological Inquiry*, 15(3), 197–215.
https://doi.org/10.1207/S15327965PLI1503_01
- Mehler, M., Balint, E., Gralla, M., Stiefel, J., Tüscher, O., & Zech, A. (2024). Training emotional competencies at the workplace: A systematic review and meta-analysis. *BMC Psychology*, 12, 718. <https://doi.org/10.1186/s40359-024-02198-3>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192.
<https://doi.org/10.1126/science.adh2586>
- O’Boyle, E. H., Jr., Humphrey, R. H., Pollack, J. M., Hawver, T. H., & Story, P. A. (2011). The relation between emotional intelligence and job performance: A meta-analysis. *Journal of Organizational Behavior*, 32(5), 788–818. <https://doi.org/10.1002/job.714>
- Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics—Part A*, 30(3), 286–297. <https://doi.org/10.1109/3468.844354>
- Parker, S. K., & Grote, G. (2022). Automation, algorithms, and beyond: Why work design matters more than ever in a digital world. *Applied Psychology*, 71(4), 1171–1204.
<https://doi.org/10.1111/apps.12297>
- Raisch, S., & Krakowski, S. (2021). Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review*, 46(1), 192–210.
<https://doi.org/10.5465/amr.2018.0072>
- Riggio, R. E. (1986). Assessment of basic social skills. *Journal of Personality and Social Psychology*, 51(3), 649–660. <https://doi.org/10.1037/0022-3514.51.3.649>
- Salovey, P., & Mayer, J. D. (1990). Emotional intelligence. *Imagination, Cognition and Personality*, 9(3), 185–211. <https://doi.org/10.2190/DUGG-P24E-52WK-6CDG>
- Seeber, I., Bittner, E., Briggs, R. O., de Vreede, T., de Vreede, G.-J., Elkins, A., Maier, R., Merz, A. B., Oeste-Reiß, S., Randrup, N., Schwabe, G., & Söllner, M. (2020). Machines as teammates: A research agenda on AI in team collaboration. *Information & Management*, 57(2), 103174. <https://doi.org/10.1016/j.im.2019.103174>
- Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human–Computer Interaction*, 36(6), 495–504.
<https://doi.org/10.1080/10447318.2020.1741118>
- Silvera, D. H., Martinussen, M., & Dahl, T. I. (2001). The Tromsø Social Intelligence Scale, a self-report measure of social intelligence. *Scandinavian Journal of Psychology*, 42(4), 313–319.
<https://doi.org/10.1111/1467-9450.00242>
- Stilgoe, J., Owen, R., & Macnaghten, P. (2013). Developing a framework for responsible innovation. *Research Policy*, 42(9), 1568–1580.
<https://doi.org/10.1016/j.respol.2013.05.008>
- Sundar, S. S. (2020). Rise of machine agency: A framework for studying the psychology of human–AI interaction (HAI). *Journal of Computer-Mediated Communication*, 25(1), 74–88.
<https://doi.org/10.1093/jcmc/zmz026>
- Vanneste, B. S., Puranam, P., & Kretschmer, T. (2024). Trust over time in exchange relationships: Meta-analysis and theory. *Academy of Management Review*, 49(4), 881–908.
<https://doi.org/10.5465/amr.2022.0041>

- Wirtz, J., Patterson, P. G., Kunz, W. H., Gruber, T., Lu, V. N., Paluch, S., & Martins, A. (2018). Brave new world: Service robots in the frontline. *Journal of Service Management*, 29(5), 907–931. <https://doi.org/10.1108/JOSM-04-2018-0119>
- Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004), 686–688. <https://doi.org/10.1126/science.1193147>
- Wrzesniewski, A., & Dutton, J. E. (2001). Crafting a job: Revisioning employees as active crafters of their work. *Academy of Management Review*, 26(2), 179–201. <https://doi.org/10.5465/amr.2001.4378011>