T

Sergio Sulmicelli

Algorithmic Discrimination and AI Governance: A Comparative Perspective

1. *Why AI puts equality to the test?*

Traditional equality analysis presupposes discrete decisions taken by identifiable actors, assessed against stable categories and legible reasons. AI disrupts each presupposition[1]. First, high-capacity models generalize from historical regularities, so the legal "baseline" to which equality compares is itself co-produced by prior stratifications. Second, models routinely infer protected traits from innocuous inputs, so antidiscrimination strategies focused solely on excluding sensitive attributes are systematically outflanked by proxies. Third, scale and pervasiveness matter: automation now conditions access to core goods, employment, credit, housing, education, and health, across both public and private domains. Under these conditions, insisting on formal neutrality risks misrecognizing structured harm as mere statistical noise[2].

As a matter of fact, AI systems are social artifacts shaped by design, organizational routines, and incentives; once deployed, they in turn reshape those very structures[3]. Moreover, design choices about data, labeling, objectives, and oversight distribute opportunities and burdens: in a sense, even AI has politics[4]. Seen more candidly, AI is both the product and producer of social order.

[1] J. Kleinberg *et al.*, *Discrimination in the Age of Algorithms*, in «Journal of Legal Analysis», 10, 2018, pp. 113-174.

[2] D. Morondo Taramundi, *Discrimination by Machine-Based Decisions: Inputs and Limits of Anti-discrimination Law*, in B. Custers, E. Fosch-Villaronga (eds), *Law and Artificial Intelligence*, Asser Press, The Hague, 2022, 73 ff.; R.J. Whelchel, *Is Technology Neutral?*, in «IEEE Technology and Society Magazine», 5, 1986, pp. 3-8.

[3] W.J. Orlikowski, *The duality of technology: Rethinking the concept of technology in organizations*, in «Organization science», 3, 1992, pp. 398-427.

[4] L. Winner, *Do artifacts have politics?*, in «Computer Ethics», 2017, pp. 177-192.

 DOI: 10.4454/0k7bac72

Design choices about data, labeling, objectives, and oversight are made within institutions and incentives; once deployed, those systems redistribute power. This duality and politics of technology suggest a different governance posture: treat design and deployment as legally salient moments where equality must already operate (*ex-ante*), rather than as a pre-legal background that only becomes visible at the remedial stage (*ex-post*).

Against that backdrop, the article proceeds in three moves. It first maps how harms arise along the socio-technical pipeline; it then revisits equality doctrine to show why classic tools under-perform unless they are fed with earlier-stage records and reasons; and it proposes a practical principle of "technological equality" that embeds *ex ante* duties across the AI life-cycle. A discursive comparison of leading instruments follows.

2. *From bias to discrimination: where harms arise*

Bias can infect machine-learning systems at multiple, mutually reinforcing entry points. Two influential mappings in the literature, the one associated with Kleinberg[5], and the taxonomy by Barocas and Selbst[6], cover largely those entry points. At the model level, discrimination can emerge along several well-understood pathways. Target selection is the first: what a system is asked to optimize often embeds a normative judgment that is not neutral. Optimizing "likelihood of reoffending" rather than "risk of failing to appear", or "predicted cost" rather than "clinical need", quietly reshapes who is prioritized and why. Targets can also be mis-specified (poor construct validity) or label-biased when ground-truths reflect prior unequal treatment (e.g., arrest counts standing in for criminal activity).

Second, training data encode histories that are unevenly observed. Sampling frames may under-represent certain communities; labels can reflect institutional bias. Even with scrupulous curation, dataset shift means that distributions seen in development differ from those in deployment, with disparate effects on minority subgroups.

Third, feature choice determines what the model treats as relevant. Variables that look innocuous can carry social meaning (ZIP code as a proxy

[5] J. Kleinberg *et al.*, *Discrimination in the Age of Algorithms*, in «Journal of Legal Analysis», 10, 2018, pp. 113-174.

[6] S. Barocas, A.D. Selbst, *Big data's disparate impact*, in «California Law Review», 2016, pp. 671-732.

for race; employment gaps as a proxy for caregiving and gender). Feature engineering choices, normalization, aggregation, dimensionality reduction, can dilute or amplify group differences.

Fourth, proxying formalizes the point: even when protected attributes are excluded, the model reconstructs them from correlated signals. Attempts to "drop the attribute" rarely defeat this and can backfire by making disparate impact harder to detect.

Finally, masking describes neutral-seeming pipeline choices used to conceal discriminatory intent or to create a façade of compliance.

Those pathways are nested within broader layers that law cannot ignore.

Before a line of code is written, someone frames the problem: which harms count, which outcomes matter, how categories such as sex/gender or race are operationalized, who annotates the data, and which communities are visible to sensors, surveys, and administrative registers. Foundational choices fix the epistemic lens: categories can be flattened (erasing non-binary identities), histories can be sanitized (removing context that explains behavior), and harms can be unmeasured (what is not measured will not be optimized)[7].

Model class, objective functions, regularization, and the selection of fairness metrics all carry normative trade-offs. Calibration, demographic parity, predictive parity, these cannot generally be satisfied together. Optimizing for headline accuracy can raise false-negative rates for smaller or noisier subgroups, a classic case of aggregation bias and fairness gerrymandering[8].

Systems do not act alone: humans interpret scores, organizations set thresholds, and institutions respond to outputs. Automation bias can lead reviewers to over-trust model recommendations; feedback loops can re-target enforcement and data collection towards already-over-policed communities, making future predictions appear self-confirming. Seemingly technical choices about presentation alter behavior and therefore outcomes[9].

Taken together, these layers show why discrimination in AI is rarely an episodic "bug" or a single actor's fault. It is the predictable crystallization of social structure in code, sustained by institutional routines.

[7] See among others A.E. Waldman, *Gender data in the automated administrative state*, in «Columbia Law Review», 123.8, 2023, pp. 2249-2320.

[8] Among others M. Kearns *et al.*, *Preventing fairness gerrymandering: Auditing and learning for subgroup fairness*, in *International conference on machine learning*, PMLR, 2018.

[9] See L. Harbarth *et al.*, *(Over) trusting AI recommendations: How system and person variables affect dimensions of complacency*, in «International Journal of Human-Computer Interaction», 41, 2025, pp. 391-410.

For legal analysis, it is helpful to distinguish foundational, technical, and implementation/interpretation biases. The taxonomy's function is to locate *ex ante* duties where *ex post* proof of intent, causation, or explicit classification is systematically elusive. Foundational bias calls for participatory problem-framing, inclusive category design, and data governance that records sampling frames and labeling rationales. Technical bias calls for justified target selection, representative evaluation, metric transparency. Implementation/interpretation bias calls for human-in-the-loop protocols calibrated to risk, user-facing notices and explanations, audits, and post-deployment monitoring that watches for drift and feedback loops. Put differently: the taxonomy is a map for where law should bite, not only at the moment of harm but throughout the life-cycle, so that equality is a design constraint and a deployment duty, not merely a courtroom argument after the damage is done[10].

3. *Rethinking equality: formal, substantive, transformative (and why it points beyond remedies)*

Equality law that enshrines formal neutrality, "like cases alike", often locks in unjust baselines; proxy-rich prediction then launders old exclusions through new correlations. A substantive turn improves on this by examining effects: disparate impacts, structural headwinds, and the ways institutional routines channel model outputs. Yet as categories proliferate and systems grow opaque, proof burdens rise and remedies arrive late. A transformative approach therefore asks institutions to redesign practices that generate exclusion in the first place: open participation channels, build documentation and reasons into systems, and adopt metrics that capture not only allocation (who gets what) but also representation and meaning (who is seen, how, and with what consequences). Because private deployers now mediate access to essential goods, these equality concerns cannot be cabined to "the state". Public-law values migrate horizontally, imposing positive duties of prevention, explanation, and redress on private actors through procurement, conformity assessment, and sectoral oversight. This does not displace classic doctrines of direct and indirect discrimination, but it reframes their use:

[10] On the taxonomy see also S. Sulmicelli, *Queer-responsive regulation for artificial intelligence in healthcare: a comparative study*, in «University of New South Wales Law Journal», 48, 2025, forthcoming.

intent and categorical visibility are often elusive in high-dimensional inference, and proxy webs frustrate *post-hoc* inquiries.

The principle of equality is one of the cornerstones of modern democratic constitutions and has traditionally been articulated in several theoretical models: formal, substantive, and transformative equality. Each has shaped anti-discrimination techniques and oriented the relationship between citizens and public power. However, the advent of artificial intelligence (AI) has called into question the capacity of these traditional categories to protect vulnerable subjects[11] effectively, making a critical reconsideration of equality in the digital age urgent.

In its formal conception, equality corresponds to the prohibition of arbitrary differential treatment among similarly situated individuals, on the premise that public authorities, constrained by constitutions, act neutrally. This principle, which provides uniform rules for those in equivalent conditions but also allows reasonably differentiated treatment for different categories, marks the first step in the constitutional path toward full equality. According to Peter Westen's well-known critique, however, in its formalist-Aristotelian sense ("treat like cases alike"), the principle collapses into an empty tautology, incapable of guiding substantive choices absent an external criterion of relevance[12]. While formal neutrality undoubtedly underwrote important early advances in civil rights, it has often proved inadequate to capture the complexity of social inequalities and today risks perpetuating pre-existing power structures, a point widely made in scholarship that shows how merely formal equality can level downward and reinforce the status quo.

From this awareness emerges the notion of substantive equality, which supplements abstract, formal neutrality by recognizing that equal treatment can, in some instances, be profoundly unjust when it ignores material starting conditions. Substantive equality can be articulated across four fundamental dimensions: a redistributive dimension, which justifies corrective measures to compensate for historical disadvantages; a recognition dimension, which aims to remove stigma and stereotypes; a participatory dimension, which requires the effective involvement of minorities in decision-making processes; and, finally, a transformative dimension, which seeks to change the social structures that generate exclusion[13].

[11] M. Tomasi, L. Busatta, M. Fasan, C. Nardocci, S. Penasa, S. Sulmicelli, *Vulnerabilità e Intelligenza Artificiale*, in «BioLaw Journal-Rivista di BioDiritto», 1S, 2024, pp. 1-4.

[12] P. Westen, *The Empty Idea of Equality*, in «Harvard Law Review», 95, 1982.

[13] S. Fredman, *Substantive Equality Revisited*, in «International Journal of Constitutional Law», 14, 2016, pp. 712-738.

From a constitutional perspective, the protection of equality then unfolds along vertical and horizontal lines. Traditionally, equality binds the state in the exercise of its powers (verticality), obliging it to ensure that legislation and administrative action do not discriminate arbitrarily. In systems such as South Africa's and Canada's, however, the equality principle also extends to relations between private parties (horizontality), imposing limits on freedom of contract and positive obligations beyond direct state action. The 1996 South African Constitution is paradigmatic: Article 8(2) provides that fundamental rights apply *inter privatos*, imposing duties of non-discrimination on all social actors. Similarly, Canadian case law under Section 15 of the Charter of Rights and Freedoms has developed a substantive conception of equality, moving beyond a formalist approach based on a neutral comparator[14].

This vertical/horizontal distinction becomes decisive in the AI context. Many algorithmic decisions that affect fundamental rights are now taken by private entities, employers, financial institutions, digital platforms. In the absence of genuine horizontal application of equality, or of anti-discrimination rules specifically capable of protecting citizens-consumers in cases of algorithmic discrimination, such actors would remain free to perpetuate or amplify systemic discrimination through technological tools that are opaque by their very nature.

As for the legal instruments that promote equality, anti-discrimination law has historically revolved around direct and indirect discrimination. Direct discrimination arises when a person is treated less favorably than another on the basis of a prohibited ground, such as sex, ethnicity, religion, disability, age, or sexual orientation, in a situation that is comparable. The decisive element is the causal link between the unfavorable treatment and the discriminatory ground, which must be evident and direct, even if proof of subjective discriminatory intent is not required. The link with the prohibited ground is thus explicit both formally (because it is directly invoked) and substantively, since the treatment adversely affects the entire reference group protected by anti-discrimination law. Indirect discrimination occurs when an apparently neutral provision, criterion, or practice places persons belonging to a protected group at a particular disadvantage, by reason of a prohibited ground, compared with others who do not share that characteristic. Unlike direct discrimination, it turns on the effect of the measure. Such disadvantage can be justified only if the challenged measure pursues

14 M. Tushnet, *The issue of state action/horizontal effect in comparative constitutional law*, in «International Journal of Constitutional Law», 1, 2003, p. 79.

a legitimate aim and the means used are appropriate and necessary, under a proportionality test. Assessing disadvantage requires identifying a reference group against which to conduct the comparative analysis, rather than comparing against a single individual. In this sense, the category of indirect discrimination is essential to capture and counter structural or systemic discrimination, even when not immediately visible[15].

In short, direct discrimination concerns adverse treatment explicitly based on a prohibited characteristic (e.g., ethnicity, sex, religion), whereas indirect discrimination concerns facially neutral practices that nonetheless produce disproportionately adverse effects for certain groups. If the move from direct to indirect discrimination marked an important evolution toward substantive equality, the emergence of AI has revealed structural limits in both models[16].

Direct discrimination, as noted, presupposes intentionality or at least a clear causal relationship between the discriminatory action and a protected category, elements often lacking in algorithmic processes: algorithms learn from historical data and operate on the basis of statistical correlations. As Barocas and Selbst observe, the reproduction of bias in AI systems is often the product of stratified social structures embedded in training data rather than the conscious choices of programmers. In this scenario, the direct-discrimination paradigm proves substantially ineffective[17].

Indirect discrimination also faces practical hurdles. In theory, it is better suited to algorithmic phenomena because it focuses on effects. In practice, however, technical opacity, barriers to data access, the complexity of predictive models, and the capacity of AI to generate new, intersectional categorizations make proof of discriminatory impact extremely difficult, draining substantive protection of effectiveness. Moreover, traditional criteria for comparing "like situations" and the focus on fixed protected categories are anachronistic when faced with forms of exclusion that manifest along multiple, intersectional lines, often invisible to existing legal frameworks, and produced by a technology capable of multiplying a subject's membership categories[18].

[15] S. Fredman, *Discrimination Law*, Oxford University Press, Oxford 2022.

[16] R. Xenidis, *Tuning EU Equality Law to Algorithmic Discrimination: Three Pathways to Resilience*, in «Maastricht Journal of European and Comparative Law», 27, 2020.

[17] S. Barocas, A.D. Selbst, *Big data's disparate impact*, in «California Law Review», 104, 2016, p. 671.

[18] S. Wachter, B. Mittelstadt, C. Russell, *Why fairness cannot be automated: Bridging the gap between EU non-discrimination law and AI*, in «Computer Law & Security Review», 41, 2021, p. 10.

The crisis of traditional principles emerges vividly in litigation involving predictive algorithms. The U.S. COMPAS case, concerning the assessment of recidivism risk, showed how an apparently neutral system can systematically overestimate risk for Black defendants compared with white defendants, even though race was not explicitly considered as an input variable. This phenomenon underscores the inadequacy of a conception of equality grounded solely in explicit data and instead requires an approach capable of interrogating the social structures underlying data and algorithms.

The complexity of algorithmic phenomena therefore demands a profound rethinking of equality. It is not enough to update existing categories: a transformative and preventive approach is needed, incorporating equality protections into the design, training, and implementation of AI systems. Protection against algorithmic discrimination requires ex ante tools, procedural transparency obligations, independent auditing, and accountability mechanisms that can prevent traditional ex post litigation and complement its usual instruments. In short, equality cannot remain anchored to a purely remedial conception. In this framework, constitutional law must reaffirm its role not only as a negative limit on power but also as a source of positive obligations to promote substantive justice in a technologically complex society.

The need for a principle of equality responsive to the technological challenge translates, first, into redefining the normative parameters that govern the design, training, and implementation of AI systems. Second, it implies structuring positive obligations for both public and private actors involved in automated decision-making, imposing standards of fairness, transparency, participation, and accountability that reflect the multidimensionality of the principle. Third, it entails recognizing that algorithmic discrimination, precisely because of its structural nature, must be countered through *ex ante* governance mechanisms that stand alongside the ex post remedial approach that has traditionally characterized anti-discrimination law.

Against that backdrop, technological equality names a governance requirement to operationalize equality values throughout the AI life-cycle by means of enforceable, reviewable duties binding designers, deployers, procurers, and regulators. *Ex ante* obligations include robust data governance; documentation artifacts (data sheets, model cards, change logs); risk and fundamental-rights impact assessments where appropriate; fairness-relevant testing with representative evaluation; human oversight scaled to risk; and auditability that survives organizational turnover. Structural awareness treats algorithmic harm as systemic, not episodic: it demands attention to representation and labeling practices, and active management of feedback loops that can magnify disparities over time. Accountability and participation require traceable decisions, intelligible reasons, user-facing notices in high-stakes contexts, accessible routes to challenge and escalation, and the meaningful involvement of affected communities in design and review. Life-cycle duties matter for two pragmatic reasons. First, they rebalance information and proof by creating artifacts intelligible outside the development team, enabling regulators, courts, and affected people to see and contest what would otherwise remain opaque. Second, they deter "compliance theater" by anchoring obligations to concrete processes that can actually be audited, sampling frames, justifications for target and metric selection, post-deployment monitoring for drift and subgroup error. Finally, the politics of measurement must be explicit. Fairness metrics are not neutral; group-based parity, individual consistency, counterfactual fairness, and calibration reflect different normative commitments and often trade off. Governance should therefore require domain-justified metric choices, reporting of distributional effects (including error at relevant thresholds and by subgroup), and alignment between evaluation practice and the social meaning of categories such as sex/gender, race, or disability. In short, technological equality turns equality from a remedial slogan into a design-time constraint and a deployment-time duty, so that by the time a dispute reaches a court, the reasons and records needed to vindicate rights already exist.

4. *A comparative perspective*

How do regulatory instruments frame equality? In other words: what legal models of AI regulation operationalize the equality principle, and how? The analytical structure adopted here evaluates present and future strategies on three planes at once: horizontally, by comparing jurisdictions; vertically, by locating rules within multilevel legal orders; and across the regulatory life-cycle, by examining phases and modalities of governance (from design to deployment to oversight). In this perspective, comparison is not only an analytical method but a critical device: it tests the adequacy of legal responses to algorithmic inequality and points toward a proactive, transformative, and technically substantive reformulation of the equality principle, one that can travel with AI systems across contexts and institutions.

Rather than slotting every instrument into a rigid template, this section reads them in context with an eye to how each can be mobilized for equality in practice.

4.1. *Council of Europe Convention on AI, Human Rights, Democracy and the Rule of Law (CAI)*

The CAI is an international law instrument that links equality to concrete, *ex ante* duties[19]. Doctrinally, it combines a substantive and transformative conception of equality: Article 10 (equality and non-discrimination) emphasizes positive obligations aimed at actively overcoming inequality, with a recognizably redistributive and participatory orientation. Article 16 requires impact assessments, a preventive hook that forces equality considerations upstream in the life-cycle. Article 14 secures effective remedies, ensuring that preventive governance is paired with avenues for redress. Two features matter for practice. First, the CAI speaks across public and private spheres, acknowledging the socio-technical nature of algorithmic risk and the need for safeguards that travel with systems across sectors and borders. Second, its effectiveness will turn on domestic implementation. Properly embedded, the CAI's vocabulary of rights, participation, and accountability gives legal traction to technological equality by legitimating *ex ante* safeguards and contestability beyond the state/market divide.

4.2. *EU AI Act*

The EU AI Act[20] is rules-based and risk-oriented, translating equality concerns into enforceable technical and organizational requirements. Its operational core includes: Article 10 on data governance (quality, relevance, representativeness); Article 15 on accuracy and robustness; Article 27 on fundamental-rights impact assessments (FRIA) for certain public-sector deployments or assimilated uses; Post-market monitoring, documentation, and conformity assessment scattered throughout the high-risk regime; Sanctions with strong deterrent effect (Articles 99-101). As an equality frame, the Act is substantive but not fully transformative. It excels at making harms visible and actionable (logs, model cards, conformity files, records) and at disciplining process; but it treats discrimination chiefly as a technical risk to be prevented, rather than as a structural phenomenon to be transformed. Notably absent are socio-technical levers such as mandated inclusive team composition, intersectionality-aware evaluation, or sustained participatory

[19] Council of Europe, Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law, Vilnius, 5.IX.2024, CETS 225.

[20] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence.

design duties. Equality gains will depend on how regulators, notified bodies, and courts use the Act's artefacts – reconstructing sampling frames from data logs, reading error reports for distributional impact, and auditing drift that re-introduces disparate effects over time.

4.3. *Brazil (PL 2338/2023)*

Brazil's PL 2338/2023[21] reads as an integrated governance charter with social inclusion at its core: Article 2 anchors the law in social justice, signaling a transformative, socio-technical view of equality. The bill imposes operational duties, transparency and traceability, algorithmic impact assessments, and stakeholder involvement both at design time and during evaluations. Distinctively, for high-risk AI it requires inclusive composition of development teams, tying technical performance to representation and participation. This hybrid-combining rules-based obligations with principles-based commitments, frames equality as capacity-building as much as risk control: deployers must show how systems meet contextual needs and how affected communities can understand, challenge, and shape deployment.

4.4. *Italy Law n. 132/2025*

Italy's framework after the enactment of L. 132/2025[22] is principles-forward but operational. Art. 3 anchors deployment in fundamental rights, non-discrimination and gender equality; sector-specific rails include healthcare (ban on discriminatory access, notice of AI use, periodic reliability checks) and public administration (traceability of AI use and preserved human responsibility). Governance is split between AgID and ACN, in coordination with sectoral regulators, with alignment to the EU AI Act for conformity assessment and sanctions. Ex post levers include deepfake and AI-aggravated offenses and updated copyright/TDM rules. Equality gains will hinge on implementing decrees (templates, registers, evaluation suites); meanwhile, procurement and PA duties can drive day-to-day accountability.

[21] Projeto de Lei No. 2338/2023 (Dispõe sobre o uso da Inteligência Artificial).

[22] Legge 23 settembre 2025, n. 132 (Disposizioni e deleghe al Governo in materia di intelligenza artificiale).

5. *Convergence, but mostly divergences. A conclusion*

Across the compared regimes, a shared spine is visible: *ex ante* documentation and testing, oversight by identifiable authorities, and user-facing transparency where stakes are high. However, read side by side, the instruments part ways along several fault lines. On the equality model, both the CAI (Art. 10) and Brazil's PL (Art. 2) take an explicitly substantive/transformative tack, coupling positive duties with participation and social inclusion. The EU AI Act remains substantive but risk-mitigating, its equality work is channeled through technical requirements (Arts. 10, 15, 27) rather than structural redesign. Italy's Law 132/2025 is principles-forward and sectoral, embedding equality chiefly in healthcare and public-administration duties and in national governance arrangements. When it comes to socio-technical levers, only Brazil's PL hard-codes inclusive team composition for high-risk AI and structured stakeholder involvement. The CAI gestures strongly toward participation but leaves national implementation to do the heavy lifting, while the EU Act and Italy lean more on technical/process controls, with participation largely a matter of policy rather than mandate. The picture is similar for impact assessments. The CAI requires rights-impact assessments; the EU Act's FRIA currently concentrates on specific public-sector uses; Brazil's PL envisages broader algorithmic impact assessments; and Italy relies on sectoral rails and institutional oversight rather than a single, universal FRIA obligation. On contestability and remedies, the CAI (Art. 14) pairs prevention with effective redress. The EU Act places weight on conformity files and muscular ex post sanctions (Arts. 99-101). Italy supplements administrative enforcement with targeted criminal and copyright updates, while Brazil's PL combines audits and impact assessments with both judicial and administrative remedies. Finally, enforcement capacity and integration diverge. The EU Act offers the most mature hard-law machinery; the CAI depends on national transposition; Brazil's PL still awaits full enactment and implementation; and Italy hinges on implementing decrees and coordination among AgID/ACN and sectoral regulators.

Abstract

This article examines how contemporary AI systems unsettle traditional equality law and why remedies limited to ex post *adjudication are insufficient. It maps where bias crystallizes into discrimination across the socio-technical flow of problem framing, data and labeling, model targets and features,* *deployment choices, and feedback loops—and argues for a principle of "technological equality" that embeds* ex ante*, lifecycle duties of documentation, participation, testing, oversight, and contestability. A comparative analysis of the Council of Europe's CAI, the EU AI Act, Brazil's PL 2338/2023, and Italy's Law 132/2025 highlights convergences on documentation and oversight, but diverging ambitions on participation, impact assessments, and structural change.*

Keywords: Artificial Intelligence; bias; equality; discrimination; comparative law.

Sergio Sulmicelli
University of Trento
sergio.sulmicelli@unitn.it