

Manufacturing misinformation

Dr Norman Lewis

The EU-funded propaganda war against free speech

May 2025



Manufacturing misinformation:
The EU-funded propaganda war
against free speech

Dr Norman Lewis

Contents

Executive summary	7
Introduction	9
1 Regulating language and political legitimacy	12
1.1 The Ministry for Narrative Control	15
2 EU Newspeak from the Digital Services Act to 349 projects	21
2.1 FAST LISA: Fighting hate Speech Through a Legal, ICT and Sociolinguistic approach	24
2.2 REASON: REAct in the Struggle against ONline hate speech	27
2.3 RECO-DAR: Right-wing extremist ecosystems driving hate speech: dissemination and recruitment strategies	30
2.4 VIGILANT: Vital IntelliGence to Investigate ILlegAl DisiNformaTion	32
2.5 VERA.AI: VERification Assisted by Artificial Intelligence	36
2.6 TITAN: AI for Citizen Intelligent Coaching against Disinformation	38
2.7 ORBIS: Augmenting participation, co-creation, trust and transparency in Deliberative Democracy at all scales	41
3 Conclusion: the emperor has no clothes	45
Appendix, 'NEUspeak' into English	49
End notes	53

Executive summary

This report exposes a covert campaign conducted by the European Commission to regulate the boundaries of legitimate public debate in Europe.

It uncovers how the Commission has funded hundreds of unaccountable non-governmental organisations (NGOs) and universities to carry out 349 projects related to countering 'hate speech' and 'disinformation' to the tune of almost €650million.

Taxpayers' money has been consciously used to fund an Orwellian disinformation complex to dictate and control the language of public debate.

This narrative-driven crusade hides how the Commission is engaged in a systematic assault on free speech in Europe.

The EU is engaged in a silent war to regulate language and, through this, the de-legitimisation of alternative narratives, like the rising tide of populist opposition. This is a battle over language and the legitimacy to dictate the terms of public communications. It is a top-down, authoritarian, curated consensus, where expression is free only when it speaks the language of compliance established by the Commission.

These projects – while framed in the language of digital safety, empowerment and democratic participation – are designed to construct an ideological infrastructure for controlling political narratives and shaping public opinion across the continent.

At the heart of this regime is an unofficial yet systematically constructed network dedicated to shaping European thought to delegitimise populist

opposition. Controlling the language provides an infrastructure of political control, a semi-invisible censorship apparatus.

The Commission uses ambiguous, euphemistic terms like ‘deliberative democracy’, ‘co-creation’, ‘capacity building’, ‘awareness raising’ and ‘fact-checker-in-the-loop’ to disguise its top-down narrative enforcement. Far from engaging in open dialogue, these terms sanitise coercive speech control as technocratic efficiency. They justify the increasing deployment of algorithms rather than human judgement to dictate what is truth or lies, information or disinformation, and who can speak in public. What is presented as neutral ‘research’ is often the ritual confirmation of preordained political assumptions, with academia, NGOs and media partners financially incentivised to legitimise the narrative of the Commission’s institutionalisation of speech regulation.

This report is not merely an exposé of financial misuse but a democratic intervention. It challenges the moral and intellectual legitimacy of the EU’s assault on speech, showing how language has become the software infrastructure of control. When the Commission defines what may be said, who may say it and how it must be framed, they do not protect democracy but undermine it. Once the cornerstone of European identity, free expression is buried beneath the soothing language of trust, safety and cohesion, while those who speak against this process are silently erased.

While exposing how taxpayer money is being used without any public accountability is essential, exposing the language war is even more so. This report is a necessary act of democratic vigilance. When language is narrowed, softened, obfuscated or stripped of meaning, so is the possibility of resistance and the development of alternatives. This report aims to sound the alarm about how the Commission fatally undermines free speech and democracy in the name of democracy and free expression.

Introduction

A spectre is haunting Europe, but it is not the spectre of disinformation or hate speech. It is the spectre of linguistic control and censorship to curb free speech. And it's not coming from China or Russia but the heart of the EU itself: from the European Commission.

Since 2016 – after the Brexit vote and the first election of Donald Trump as US president – the EU Commission, spooked by these developments, has been on a crusade to control Europe's political narrative. One form of this crusade has been the 'hate speech' and 'disinformation' narrative, which the EU Commission has argued is a growing threat to social stability and democracy in Europe. Its content, however, is far from the benign act of responsible government the EU Commission would have us believe. It is an authoritarian assault on free speech and the European demos, whom the Commission considers lack the moral independence to think and act in their own best interest. It also requires the paternal hand of unelected experts and technocrats for guidance.

This report focuses on the much-neglected *means* through which the EU Commission realises its narrative objectives. Our research has uncovered the staggering fact that the Commission has funded hundreds of unaccountable non-governmental organisations and universities to carry out 349 projects related to countering 'hate speech' and 'disinformation' to the tune of €648,890,016. (A breakdown is provided in the next section.) Taxpayers' money has been consciously used to fund an Orwellian disinformation complex to dictate and control the language of public debate. This crusade hides in public the systematic assault on free speech in Europe.

In a previous report published in 2024, *Controlling the Narrative: The EU's Attack on Online Speech*, we outlined the EU's Censorship Operating System and how the EU's Ministry of Truth manages this.¹ In this report, we uncover what might be called the EU Ministry for Narrative Control – a shady network funded to the tune of nearly €649million – to sustain the most extensive quest to regulate the language of political speech in European history.²

The EU is engaged in a silent war to regulate language and, through this, the de-legitimisation of alternative narratives like the rising tide of populist opposition. This is a battle over language and the legitimacy to dictate the terms of public communications. It is a top-down, authoritarian, curated consensus, where expression is free only when it speaks the language of compliance established by the Commission. Our report catalogues this deeply dishonest Orwellian crusade, which, unlike historical attempts to outlaw free speech, does not burn books or squash dissent with jackboots. Instead, it is a silent and focused war conducted in public to control the language of conversation. The Commission rightly understands that controlling the language of communications means it can dictate what is information and disinformation, truth or lies, what is legitimate or illegitimate speech and who can speak or not. And if it can control information and the truth, it controls history, the past and the future.

The motivation for this report is more than exposing how taxpayer money is being used without any public accountability. That is very important. But this report is a necessary act of democratic vigilance. Because when language is narrowed, softened, obfuscated or stripped of meaning, so is the possibility of resistance and the development of alternatives.

The report is structured as follows:

In chapter 1, we explain the EU Commission's quest to regulate the language of public communication. The overriding objective of the Commission is to categorise certain speech as 'hateful' or 'disinformation', so that it can be policed. But, as we show, the EU is at the same time (some-what paradoxically) dependent on the existence of such speech. Solving the 'problem' of so-called hate speech and disinformation is now a core part of the *raison d'être* of the EU. Thus, 'discovering' more and more such speech becomes a priority. The EU is therefore dependent on inventing the problem it then must solve.

Chapter 2 examines several projects and their objectives to demonstrate the dishonest subterfuge the Commission is engaged in, which it funds using taxpayers' money without public scrutiny. These examples highlight how the Commission erodes not only civil society but the integrity of academic research, too.

Chapter 3 concludes the report by examining the inconvenient truth and paradox of the EU Commission's 'hate speech' and 'disinformation' narrative: namely that the EU, in a quest to legitimise itself and its authority, depends upon the institutionalisation of the very thing it purports to want to extinguish. The EU Commission is addicted to hate speech and disinformation not because it wants to foster these but because it fears the energy and unpredictability of free speech. The swaggering arrogance the Commission projects in public hides the fact that it has no clothes beneath the hot air of its Orwellian bombast. It is weakness, not strength, illegitimacy and the absence of authority that forces the Commission to sustain its narrative artificially. Calling this out is an essential step towards democratic renewal in Europe.

1 Regulating language and political legitimacy

It is shocking that the taxpayers' money spent to sustain the Ministry for Narrative Control is almost 31 per cent higher than the money allocated for transnational research and innovation projects addressing various cancers-related objectives (€494million).³ This is not an accident. The EU Commission regards stemming the cancer of free speech as more of a priority than the estimated 4.5million new cancer cases and almost two million cancer deaths in Europe in 2022, for example.⁴ This may be callous, but the Commission is in a life-and-death struggle to maintain its legitimacy and authority. All is fair in love and war.

The €649million spent, however, needs to be seriously caveated. This underestimates the amount the Commission is spending on upholding its narrative. The EU Commission hides information in plain sight in numerous disjointed databases containing details of the programmes and projects it funds. The lack of consistency in naming and labelling organisations and activities means that those without a PhD in forensic accounting and data science find it almost impossible to calculate real spending and budgets for projects or organisations.

The figure we have calculated is drawn from 'Hate Speech' and 'Disinformation' projects listed in the EU Funding and Tenders Portal and the Horizon Programme database, Cordis. When we calculated this at the beginning of April 2025, the figures broke down as follows:

'Hate Speech' Projects	Number	Amount EU budgeted €
The EU Funding and Tenders Portal	60	21,113,539
Horizon Cordis Data Bank	75	190,560,320
'Disinformation' Projects		
The EU Funding and Tenders Portal	72	60,612,142
Horizon Cordis Data Bank	142	376,604,015
TOTAL	349	648,890,016

However, many projects are not designated 'hate speech' or 'disinformation', but promote the same narrative. For example, the EU Funding & Tenders portal reveals that the EU has funded 227 'mental health' projects over the past decade.⁵ While many may be genuine mental-health projects, many deal with 'hate speech' and 'disinformation'. For example, the project 'Social Media: Measuring Effects and Mitigating Downsides',⁶ fully funded under the Horizon programme to the tune of €1,494,625, masquerades as a project concerned about social media's negative mental-health impact on children. In substance, however, it is an ideologically behavioural conditioning project aimed at promoting the idea that social media requires domestication, and thus that it projects the same message that this needs to be steered by unelected and unaccountable EU experts for the sake of mental health and democracy, of course.

Another, the 'e-Intervention Enhancing Mental Health in Adolescents' – funded under the Horizon Europe programme for €6,998,198 – contains a project titled '(Em)poweryouth: a co-created model towards a non-discrimination culture'.⁷ It describes itself as 'an innovative model for combating Islamophobia based on changing narrative frameworks' and focuses on Muslims aged 18 to 30. The goal, however, is to 'co-create new narrative

frameworks to combat Islamophobia by seizing the public debate'. And how will that be accomplished? Well, by training 'young people to help them to become opinion leaders'. To this end, they propose 'resources and tools for the empowerment of these target groups to identify intra-group discriminatory practices and generate internal transformation processes and progress in equality'. What this Orwellian newspeak means by 'empowerment' has nothing to do with autonomy; it is about embedding EU institutional norms within this group. 'Internal transformation processes' and 'progress in equality' might sound universal and positive but means conformity with the EU's prescribed ideology of inclusion. Those who question or resist these transformations may be recast as discriminatory – even within their communities. What this has to do with mental health is anyone's guess, but the 'opinion leaders' this aims to create will be nothing but mouthpieces for the EU narrative.

These examples and many more are not included in the overall funding we've uncovered so far. However, there is another level of funding similarly excluded, and that is the funding for many minority groups who feature in almost every official EU definition of hate speech and are regarded as the victims of hate speech and disinformation. Each group is covered by an EU strategy position, which involves numerous initiatives and funding in multiple areas, making it almost impossible to calculate the extent of the EU's real spending.⁸

In short, €649million is the tip of a funding iceberg. No doubt, more research will uncover the full extent of this funding. However, this only demonstrates how important the control over the narrative is for the Commission.

1.1 The Ministry for Narrative Control

In 2022, the Commission funded a pilot project for €1,190,500 to establish a ‘European Observatory on the Narrative Distribution’.⁹ This was officially known as the ‘Narratives Observatory combatting Disinformation in Europe Systemically’ (NODES).¹⁰ Its primary aim was to analyse and monitor how narratives – particularly those contributing to disinformation – emerge and spread within the European public sphere. The project focused on key topics such as climate change, migration and Covid-19, operating in four languages: English, French, Spanish and Polish.

Led by the think tank Re-Imagine Europa,¹¹ the NODES consortium included partners like Agence France-Presse (AFP), the French National Centre for Scientific Research (CNRS), PlusValue, Sotrender, Science Feedback and Università Ca’ Foscari Venezia. The project concluded its initial phase with a notable exhibition at the European Parliament in December 2024 titled ‘Bridging the Divide: New Narratives for Climate Action’.¹² This event showcased innovative methods to address polarisation and foster a constructive dialogue on climate issues.

Building upon the foundations laid by NODES, a second initiative was launched under the European Commission’s call for proposals. This new project, ‘Predictive Research on Misinformation and Narratives Propagation Trajectories’ (PROMPT), is led by the French organisation OpSci.¹³ PROMPT aims to analyse emerging narratives related to the war in Ukraine, gender-based disinformation, and electoral misinformation and had a particular focus on the 2024 European elections. The project employs advanced AI methodologies, including large language models (LLMs), to detect and understand the spread of disinformation across various media platforms.

This project should be considered the vanguard foray into establishing the EU's Ministry for Narrative Control. When one analyses the objectives of the initial project, what sounds like a scientific or cultural initiative is, on closer inspection, a programme designed to map, influence and ultimately manage public perceptions of narratives at the level of emotion, identity and ideological affiliation.

The language used to promote this initiative is itself a masterclass in newspeak. The project begins by declaring that the Covid-19 pandemic has become 'a major challenge for the fragile web of trust that connects the citizens with their states and the EU'. The daemon is social media because it 'provides a crucial boost for the international disinformation network'. This claim is immediately followed by the observation that social media also provides 'a great weapon against it'. Social media is regarded as both disease and cure, enemy and tool. This sets the ground for greater institutional intervention in digital communications under the comforting guise of protecting the public. 'Arming against disinformation' is a cute way of establishing the need for censorship.

However, the most revealing phrase comes with the assertion that 'narratives play an important role in framing facts and information in a package that can be easily transmitted across society'. This seemingly harmless line is, in effect, a quiet concession that facts alone are no longer the battleground – the narrative wrapper counts. The goal is no longer to correct falsehoods but to outcompete them at the emotional and cultural storytelling level. Truth becomes less a matter of evidence and more about which 'package' gains traction.

When the project states its aims to 'decipher the emotional values that guide successful narratives' and 'develop recommendations for effective communication and policy', this transforms political expression into

something that can be psychologically reverse-engineered and reprogrammed. In this model, emotion becomes a metric of threat, not a sign of democratic engagement. The target is no longer misinformation, but the wrong feeling triggered by the wrong story.

To this end, the project aims to uncover the ‘dominant narrative patterns’ by ‘analysing the most qualitatively representative examples’. This is a polite way of saying the project will create a system to classify and rank stories that deviate from desired ideological norms the Commission wants to uphold. A series of seemingly neutral questions follows, such as, ‘What are the values and narratives that unite and separate us?’ However, these are not posed to citizens for debate. They are addressed to a consortium of researchers, algorithm designers and EU-funded institutions tasked with monitoring and categorising how Europeans think and feel.

By embedding phrases like ‘dominant narrative models’, ‘identity and community’ and ‘weaponised information ecosystems’ in policy language, the project disguises its true objective: constructing a system of narrative legitimacy, a hierarchy of permissible meaning. It is a war not on lies, but on language itself. The problem is not what people say, but how and why they say it – and whether it aligns with what the Commission considers trust-building, inclusive or democratic.

In this context, language is no longer descriptive. It is performative and strategic. Terms like ‘deliberative democracy’, ‘resilience’ and ‘values’ are emptied of their traditional meanings and refilled with technocratic purpose. Debate is not welcomed, but ‘facilitated’; disagreement is not resolved, but ‘monitored’; identity is not expressed, but ‘mapped’. Under the guise of inclusivity and cohesion, the citizen becomes a subject of interpretive language management – not a participant in democracy, but a data point in narrative and semantic engineering.

This is not censorship in the blunt, authoritarian sense. It is the creation of a discursive architecture that manages dissent through language control – by redefining what counts as acceptable discourse and who gets to shape it. It does not silence voices; it submerges them under waves of euphemism, analytics and policy. At heart, the battle over narratives is a struggle over meaning, fought with the soft weapons of framing, emotion and ambiguity. And in this war, whoever controls the language controls the limits of the political imagination.

In this regard, the use of language, particularly the acronyms used to designate these projects, is vital to note. Acronyms like NODES and PROMPT are, as can be seen, deliberately engineered and created to disguise the real purpose of each project. Another example is FAST LISA (made from the project title ‘Fighting hAte Speech Through a Legal, ICT and Sociolinguistic Approach’).¹⁴ This project, as we will demonstrate in detail in the following chapter, is an indoctrination programme for young people about behaviour change, coaching them ‘what not to say’ rather than equipping them to challenge hate speech critically.

These chirpy acronyms don’t just sound like digital voice assistants or wellness apps that might schedule your meetings or check your vitamins; they are deliberate, dishonest strategic terms chosen to disguise a real authoritarian purpose. They communicate safety while infantilising the public by treating them as anxious children in need of comfort, not truth. They hide from view that these projects are about power, the automation of speech control through algorithmic semantic engineering, and the complexities of democracy and censorship vs free speech. Search the EU’s databases, and you will not find CONFORM, the acronym for a project titled ‘Automated CONTENT Control InFrastructure for Narrative COMpliance.’ You will not find PANOPTIC for a project titled ‘Programme for Algorithmic Narrative

Oversight and Perception Tuning through Informational Control’ – even though such putative project titles and acronyms would be rather more honest.

This disingenuous use of language highlights another essential dimension of the Ministry for Narrative Control. And everyone who participates or aspires to participate plays a well-scripted and templated game. The verbal contortions they go through to create project names that yield feminised acronyms that sound more like toothpaste brands indicate that participation demands performative compliance. But it is more than just compliance. It is collaboration in a dishonest charade that masquerades as academic ‘research’ but obscures reality, perpetuating the assertion that hate speech and disinformation are real problems in need of urgent attention.

The battle over language is often overlooked or regarded as of secondary importance. But, from what we see above and will be seen in more detail below, language is not just a technique for communication. It is how we think, imagine, and decide what is real and meaningful. Every society, whether democratic or authoritarian, depends on language to shape its values and meaning, its conflicts and its limits. The words we are given determine what we can see, what we can name and what we can challenge. When language is controlled – by states, institutions or NGOs – so is the range of thought and dissent. A society that redefines surveillance as ‘safety’ or censorship as ‘content moderation’ does not need to silence citizens outright; it simply changes the meaning of their silence.

Language is the EU Ministry for Narrative Control’s software infrastructure of control. When the EU Commission defines hate speech, disinformation or extremism, it is not identifying problems – it is drawing the lines around what can be said, by whom, and with what consequences. These definitions are not neutral. They carry ideological weight, especially

when delivered in the neutral, obscurantist tones of policy language. Entire categories of political speech are being re-coded as illegitimate. Most notably, populist political movements – particularly those critical of EU integration, immigration or the Green Deal – are increasingly framed not as political ideas to be debated but as algorithmic vectors of hate, extremism or misinformation. Populist language is scrutinised not on ideological or democratic grounds, but through technical and moral frameworks that pose rhetorical questions. The answers to these questions are known beforehand, such as whether populist language promotes harmful stereotypes. Or does it erode trust in institutions? Or does it violate community guidelines?

Populism is not outlawed directly (yet). But it is systematically linguistically degraded, rendered suspect by default, always placed on the edge of unacceptability, a quiet form of de-legitimisation, silently enforced through the language of civility and tolerance. And when this becomes the norm, the terrain of democratic contest shrinks. Once populist dissent is pathologised as hate or treated as a cybersecurity threat, it no longer needs to be engaged with. It can be monitored, fact-checked, defunded, quarantined and removed. Through the vocabulary of public safety and moderation, public debate is increasingly managed like a public-health crisis, a hygiene regime that cleanses speech and purges the toxins to promote ‘healthy voices’.

But who decides what is toxic? Who defines what counts as healthy speech? This question is not allowed nor posed. There is no open democratic deliberation. Instead, we have a narrative reproduced in hundreds of EU-funded projects that produce toolkits, algorithmic classifiers and behavioural ‘nudge’ strategies developed by unaccountable experts whose legitimacy is assumed rather than questioned. In the following chapter, we provide examples illustrating the systemised obfuscation of the EU Commission’s well-funded ‘hate speech’ and ‘disinformation’ narrative.

2 EU Newspeak from the Digital Services Act to 349 projects

The EU's success in determining the language of public communications is remarkable. It is a lexicon that is never questioned, whether in European Parliament debates, online or in the media. Yet it pervades everything. The Digital Services Act (DSA) is its crowning glory. Through this, the Commission has not only established its right to determine what can or cannot be said online, but also codified the Orwellian newspeak at the heart of the EU language – or what we will call NEUspeak from now on.

The DSA is presented as a milestone legislation aimed at creating a safer digital space where users' fundamental rights are protected and social-media platforms are obliged to act responsibly. The language is deliberately sterile, technocratic and reassuring: it speaks of transparency, systemic risk mitigation, trust and safety mechanisms in a world of harmonised obligations. But nothing is further from the truth.

The term 'service' is one of the most subtly insidious pieces of bureaucratic NEUspeak. On the surface, it sounds benign, even benevolent. But when used in the language of digital governance, 'service' is a rhetorical Trojan horse. This term masks authoritarian regulation, surveillance and control behind the comforting suggestion of convenience, neutrality and the protection of the public good.

At face value, calling platforms like Facebook, X, YouTube or TikTok 'services' makes them sound like utilities – passive infrastructures that deliver content like electricity or water – technical and procedural entities rather than inherently political and ideological. This depoliticises the public square of the twenty-first century, obscuring the fact that these corporations are

private, unaccountable, profit-driven entities with enormous powers to shape public debate through algorithmically enforced speech boundaries. It hides from view that what's at stake is whether private, unaccountable corporations or the equally unaccountable Commission gets to determine the truth and who can mutter it.

'Service' shifts the relationship between citizens and social-media platforms. In a true democratic public service, the user has rights, protections and some democratic oversight. However, under the DSA, the user is not an active citizen, but a data-producing endpoint within a tightly regulated commercial framework increasingly dictated by state-aligned priorities. Speech is presented as a commodity delivered conditionally rather than something inalienable. The citizen is reduced to an 'end-user'; dissent and anything deemed dangerous becomes a risk factor to be managed. Regulating speech censorship is not a political tool, but a technical means of service optimisation.

The word 'service' is part of a conscious misdirection. It enables the EU Commission to describe the policing of speech and unofficial political narratives as ensuring the fine-tuning of a technical delivery system. Policy replaces politics, while resistance is isolated as a breach of contract. Terms like 'illegal content' and 'systemic risks' are never defined but are deliberately vague. This creates an environment of constant doubt where self-censorship by everyone involved becomes the default. The Commission, in turn, avoids the appearance of direct censorship while effectively outsourcing enforcement to private actors, thus excusing itself from responsibility. This is censorship by unaccountable outsourced proxy.

The technical enforcement through designated 'trusted flaggers' – entities empowered to report content for expedited removal – is another layer of obfuscation. These are not independent, non-aligned, neutral

organisations sworn to enforce objectivity. Often, they are unelected NGOs or state-aligned organisations who agree with the Commission's federalist ideological agenda. The technocratic language suggests neutral expertise, but deploying these proxies establishes a hierarchy of speech where certain voices are given institutional priority in shaping the information landscape. Meanwhile, the 'Transparency Data Base' established as part of the DSA is opaque, hiding the more profound shift towards centralised narrative management in public.

In these ways, the DSA doesn't openly censor but rebrands the regime it enforces as the infrastructure of neutral content moderation. It doesn't silence voices directly – it builds systems in which silence becomes the safest option for social-media platforms and users. And it does all this while speaking the smooth, managerial dialect of EU policy-speak: a language engineered not to alarm, but to comfort, confuse and control.

The language embedded in the DSA sets the legal terms of speech controls. It frames the public conversation and thus the terms within which the debate about free speech is conceived and acted upon across Europe. The key point, however, is that this is backed by a curated narrative that reinforces the idea that this is the only acceptable way public debate can be conducted. Setting the language sets the terms of the conversation.

What is not fully understood is what strengthens this power: the strict way in which the Commission has curated the parameters of the 'hate speech' and 'disinformation' narrative through its funding of projects that involve civil-society groups, universities and research companies. These organisations and institutions legitimise, sustain and advance the assumptions this subterfuge is built upon.

As we have stated, at least 349 identifiable projects are being marshalled by hundreds of NGOs, universities and profit-making research organisations

for this purpose. Each of these projects ought to be scrutinised in detail to expose the extent of the Commission's deception. Space precludes this. Instead, we examine some to bring out the fabricated underhand and unscrupulous manipulation the Commission is engaged in. We also do so to show how the Commission hides its crookedness in public while behaving like a dealer doling out financial incentives behind closed doors to feed a network of collaborators who sustain its narrative.

2.1 FAST LISA: Fighting hate Speech Through a Legal, ICT and Sociolinguistic approach¹⁵

The FAST LISA project is framed as a cross-border initiative to tackle the growing challenge of online hate speech through digital tools, legal expertise and youth engagement. This was partially funded by the Commission for €568,165 and was coordinated by the Alma Mater Studiorum, University of Bologna. It involved three more universities, three local government bodies and one private limited company.

The FAST LISA project presents itself as a progressive and innovative response to the challenge of online hate speech in Europe. However, behind the polished language of 'protocols', 'tools' and 'capacity building' lies a coordinated attempt to implement a centralised system of speech monitoring and behavioural influence wrapped in the soft language of civic 'empowerment' and digital literacy. This project is not merely about countering harmful content, but embedding an institutional infrastructure for narrative control – shaping how people speak, think and engage online through surveillance, algorithmic classification and targeted re-education programmes.

The objective of the project states:

In a European scenario where online hate speech represents an increasing challenge for national governments, due to its elusiveness

*and to the lack of efficient reporting and counteracting,
FAST LISA will produce a consolidated and cross-border
standard protocol designed for local authorities...*

Without explaining why governments are struggling to manage and define hate speech online, which might have something to do with its subjective and contextual subtlety, they assert this content is problematic, inferring this is the result of deliberate behaviour. Instead, this project aims to create a one-size-fits-all EU-wide system to help local governments deal with it. But this hides the creation of a uniform speech-governance framework – a set of top-down rules that overrides national nuances, streamlines regulation, and enables state-aligned actors to determine acceptable speech boundaries across borders.

The first stated methodological step is:

*Design, dedicated to develop the FAST LISA dashboard,
a specific tool designed on a state-of-the-art analysis covering
legal, linguistic and sociological aspects of hate speech online...*

Put simply, this is a digital platform that will scrape and analyse content from major social-media platforms using AI to detect and map what it will determine hate speech to be. In practice, this amounts to the automated surveillance of public expression, with machine-learning tools trained to flag sentiment, tone and phrasing that deviates from approved norms. This dashboard isn't about understanding context. It's about creating a data-driven classification system for pre-censorship in order to reinforce dominant ideologies under the guise of deploying neutral technology.

The second methodological step is:

*Capacity Building, dedicated to the delivery of the project's
training programme addressing groups of young adults in the
3 FAST LISA target cities/areas*

Empowering young people in civic responsibility sounds very positive. However, the ‘capacity building’ they want to build has nothing to do with arming young people with critical skills to increase their autonomy and counter hate speech when they come across it. This is about the indoctrination of youth in the EU’s narrative. It is soft behavioural conditioning to adopt the project’s worldview as normative, not as a point of debate.

The final methodological step draws this out more succinctly:

Engage, leading youngsters, experts, cities and local authorities, to become engaged FAST LISA agents by developing pilot actions at local level, both on a community and on an institutional level.

The ‘capacity building’ is, in fact, the indoctrination of young people to behave and act as speech police and, through this, to embed this ideological model at both community and governmental levels. By transforming citizens, particularly youngsters, into ‘agents,’ this project outsources enforcement and legitimacy to supposedly grassroots actors, who then feed back into institutional policymaking, ensuring that what appears to be bottom-up reform is, in fact, a pre-scripted system of narrative compliance, rolled out in stages and dishonestly cloaked in the language of ‘policy co-design’.

In essence, FAST LISA is less about protecting vulnerable communities from hate and more about embedding a system of narrative regulation into the local, legal and digital fabric of European society. By transforming young people into ‘agents,’ authorities into collaborators, and data into a justification for behavioural engineering, the project normalises a future in which speech is constantly measured, flagged and shaped by invisible hands – all in the name of progress and protection. At its core, what appears to be a participatory programme is a centralised mechanism for the ideological management of online expression.

2.2 REASON: REAct in the Struggle against ONline hate speech¹⁶

The REASON project, established in 2021 and provided with €250,000 by the Commission, was a project presented as a national initiative to combat hate crimes and hate speech through institutional coordination, targeted training and the use of advanced technology. The Italian government's presidency of the Council of Ministers coordinated it.¹⁷ It involved one university, one NGO, the Associazione Carta di Roma¹⁸ and the *Istituto per la Ricerca Sociale* (IRS), a non-profit independent research institute.¹⁹

The REASON project hides its goal of norm enforcement, ideological filtering and narrative control under highly technocratic language. The aim is to 'improve the responses of the Italian authorities on hate crime and hate speech by envisaging several actions addressed to specific target groups', which are never defined.

To accomplish this, it aims to

establish formal public authorities networks...with a view to designing and implementing a specific National Action Plan against Hate Crime and Hate Speech...a preliminary study...for the elaboration of a specific training module for each target group...[and] a national Observatory for the identification of online hate speech...[using] automatised text mining...for objective representations.

In plain English, the state proactively shapes what kinds of speech are socially and legally permissible. It aims to coordinate a nationwide bureaucratic structure to develop and enforce a top-down definition of hate. It reduces nuance, erases legitimate dissent, and allows broad and subjective interpretations to be acted upon by police and other authorities. The term 'observatory' is especially disingenuous. It's a cute way of establishing a technological arm of enforcement. This tool gives the illusion of neutrality while hiding that the algorithms trained will be based on institutional

definitions of ‘hate’. What is described as ‘objective’ is a subjective filtration system automated at scale, placing public discourse under continuous algorithmic surveillance.

A final goal includes an ‘awareness-raising action’ involving national and local stakeholders to ‘facilitate mutual learning and exchange of good practices’. In essence, this is about embedding the ideological framework into all levels of governance, ensuring that what begins as a centralised policy becomes a standardised national norm. It is a project that builds the infrastructure for institutionalised speech correction, where compliance with dominant norms is enforced through training, technology and bureaucratic alignment.

The final report by the authors demonstrates this. They proudly report that a set of training workshops has been carried out with young people from potentially targeted groups as a proposal for participation and digital activism: the research action path consisted in a mini-conference on forms of hate speech and their validation with activists, finding and analysing cases, the design of counter-narrative and alternative narration, the creation of videos for the social web. This activity included and raised awareness among groups of people directly involved in the spread of hate speech. These latter took part in the project in the mini-conference with a ‘learning by doing’ approach, through which they experienced acting as content producers and consumers. These young people consolidated their knowledge by developing their ‘activist’ consciousness.

The phraseology in this passage contains some of the most important recurring NEUspeak, which needs to be understood.

The ‘training workshops’ and ‘participation and digital activism’ are not about engaging young people in critical thinking or radical activism. It is about recruiting and shaping targeted participants into compliant

disseminators of officially approved narratives. Similarly, the ‘research action path’ is not about open inquiry. It is a structured programme to validate pre-defined definitions and views on hate speech. The bastardisation of the concept of ‘research’ is discussed in more detail below.

‘Activism’ means the creation of officially acceptable speech gatekeepers. Teaching participants to ‘design counter-narratives and alternative narration’ means training them to reframe contested issues using institutionally sanctioned perspectives that delegitimise alternatives. The inclusion of those ‘directly involved in the spread of hate speech’ under the guise of ‘awareness raising’ is a Big Brother re-education programme dressed as a health intervention. Its aim is behavioural correction, not through debate and contestation, but enforced transformative compliance. This culminates in participants developing what the report proudly calls their ‘activist consciousness’ – a euphemism for accepting ideological compliance – not critical independence. What is presented as grassroots empowerment is, in fact, a form of top-down ideological grooming, the training of a new generation of censors and digital-speech police dressed in the language of participation, care and social justice.

Even more disingenuous is how ‘research’ is manipulated to give the impression of objectivity while reducing it to a confirmation ritual.

The ritual of ‘research’ as conformation

The reference above to ‘a ‘research action path’ is one form in which the Commission redefines ‘research’ to mean something that has nothing to do with honest inquiry. The Oxford English Dictionary defines ‘research’ as ‘the systematic investigation into and study of materials and sources to establish facts and reach new conclusions’. In any field, proper research upholds an open-mindedness where discovery reveals new and uncertain

things and assumptions are tested, weighed, discarded or updated according to new insights and discoveries. At its core, research is meant to generate new knowledge, not manufacture pre-defined agreements. In every Commission-funded project related to ‘hate speech’ and ‘disinformation’ with some research component, ‘research’ reinforces the hate speech/disinformation hypothesis rather than testing it. When ‘research’ becomes an obligatory tick-box exercise in funding applications, the Commission rebrands inquiry as a confirmation ritual rather than any honest pursuit of truth, knowledge or change.

The 349 projects we’ve uncovered are all geared towards confirming what the EU Commission knows before one euro of funding has been spent: namely, that hate speech and disinformation is a rising problem and represents a threat to democracy. ‘Research’ that systematically ‘proves’ this assumption is not research; it is the manufacturing of propaganda used to legitimise the narrative, pre-empt criticism and thus delegitimise any ideas or narratives that do not conform to this apparent ‘objective’ perspective. Based on this ‘research’, the exclusion of alternatives is not a position gained through rigorous, challenging and honest sweat; it is a manufactured deceit laboured to present the Commission’s narrative as the only game in town.

The following example demonstrates this.

2.3 RECO-DAR: Right-wing extremist ecosystems driving hate speech: dissemination and recruitment strategies²⁰

Unlike many projects, this title is more honest and open about its bias and intentions despite its Orwellian acronym ‘RECO-DAR’. This involves SCENOR,²¹ an NGO based in Vienna that coordinated the project and Modus|Zad, a not-for-profit research organisation based in Berlin. The Commission granted this project €269,310.

The project objectives start with what any actual research-based project would seek to prove instead of assuming that ‘hate speech has been on a constant rise in the German-speaking online community over the past years’. The rise in ‘right-wing extremist’ hate speech is presented as an established, measurable fact. There is no definition of what qualifies as this hate speech, no data, no clarification of who is doing the counting or under what criteria. Presented as a mapping exercise through automated data capture, this ‘research’ is a surveillance and categorisation operation focused not on illegal content – which is assumed rather than proven – but on the ‘ecosystems spreading violent hate speech’. Under the guise of a research-based project, the project blurs the line between actual incitement to violence and politically undesirable speech. This is a deliberate attempt to delegitimise ‘alternative narratives’ as ‘right-wing extremism’.

The project’s real goal, however, is to have a ‘direct and tangible impact on the work of public authorities, social-media companies, academia and practitioners engaged in preventing and countering hate speech’. To this end, the project provides the following:

increased knowledge on hate speech online and indicators to identify and report hate speech and to create effective counter-narratives; contribute to strengthened cooperation among the stakeholders and increased awareness of hate speech online.

This ‘research’ justifies the arbitrary identification of hate speech without defining it precisely, being presented as objective ‘knowledge’. This can thus be used to impose norms to be acted upon without the need for any public or democratic debate. The ‘indicators to identify and report hate speech’ are automated or institutional surveillance mechanisms that operationalise subjectivity, enabling complex speech acts to be reduced to binary flagging

systems. This deputises individuals or platforms into speech police without any accountability.

Creating ‘effective counter-narratives’ is the production of propaganda by another name, which is measured not by the democratic debate it produces but by the level of conformity and behavioural change it engenders.

‘Strengthened cooperation among the stakeholders and increased awareness of hate speech online’ is a polite way of enforcing policy alignment and gatekeeping by the Commission, NGOs, tech firms and researchers outside the scrutiny of the public. ‘Increased awareness’ is a vague term that is about increasing exposure to the EU’s hate speech and anti-populist crusade, not prompting any critical engagement with ‘hate speech’.

In short, this ‘research’ aims to prove why the Commission’s crusade against ‘hate speech’ and ‘disinformation’ needs to exist and that its narrative is the only acceptable one. It develops and disseminates tools that institutional actors can use to define, detect and delegitimise politically undesirable speech online. The goal is to equip platforms, governments and NGOs with the language and systems necessary to classify ‘far-right’ or populist dissent as harmful, to justify intervention, and to promote official narratives as a public good. Under the banner of ‘cooperation’ and ‘awareness’, it advances a quiet but powerful regime of ideological conformity in Europe’s digital public sphere.

2.4 VIGILANT: Vital IntelliGence to Investigate ILlegAl DisiNformaTion²²

This project aims to develop a new platform to help police fight online hate crimes. It is funded under the Horizon Europe programme with €3,376,604. Coordinated by Trinity College Dublin, it involved one think tank, three

cybersecurity technology firms, two not-for-profit NGOs, four universities and five police authorities across 10 EU member states.

The VIGILANT project presents itself as a solution to the complex problem of online disinformation. In true EU Commission fashion, social media and speech are assumed to be a problem that requires intervention and should be led by law-enforcement agencies – rather than the public, through democratic oversight. This starting assumption justifies a project presented in technical language, which represents a profound shift in how speech is policed across Europe.

At its core, VIGILANT builds a powerful AI system to automate the surveillance and classification of online content, blurring the line between harmful speech and political dissent. It frames the inability of police to track digital conversations as a technological failure, not a legal or ethical boundary. By treating ideologically charged but legal expressions like ‘separatism’ or ‘nationalist’ rhetoric as potential security threats, it extends counterterrorism logic into the realm of digital communications. Furthermore, by aligning public authorities, tech companies and NGOs in a coordinated detection and response system, the project institutionalises a soft but expansive architecture of narrative control, while framing itself as ethical, inclusive and user-centric.

Overleaf is a table which breaks down the project’s objectives expressed through Commission NEUspeak.

VIGILANT project: NEUspeak decoded

Project in NEUspeak	in plain English	what it means
'Identifying, tracking and investigating online disinformation and other problematic content is an extremely complex problem'	Because disinformation and hate speech online depend on context, subjectivity and nuance, it is difficult to define and identify.	Complexity downplays the role of human judgement, which justifies algorithmic intervention. Also, framing speech as a security problem justifies deploying these systems to intervene behind the scenes in public debate.
'Many Police Authorities (PAs) do not have access to specialised tools...'	Police lack good software for tracking hate speech and disinformation	The police and authorities, not the public, are the rightful actors to regulate speech.
'Off-the-shelf products... are unsuitable'	Commercial software wasn't made for this.	Civilian tools don't give law enforcement enough power – so we'll build them a dedicated surveillance platform.
'The VIGILANT project solves this problem'	We've made a new solution.	VIGILANT is an AI surveillance suite aimed at monitoring, classifying, and profiling speech, users and networks, which takes the complexity out of controlling freedom of expression.
'Advanced disinformation identification and analysis tools...state-of-the-art AI methods...'	It uses the latest AI technology.	The automation of language policing and ideological classification at scale.

'Ethical-by-design and user-centric approach'	It's ethical and designed with users in mind.	PR subterfuge language without explaining what or how ethics are ensured, nor who 'the user' is.
'Social and behavioral aspects are also taken into account'	It looks at how people think and act online.	Includes psychological profiling and behavioural analysis to model online 'threats'.
'Covers disinformation from all major sources ...in all modalities... in multiple languages'	It works on all content types and platforms.	An expansive mandate for surveillance with no boundaries defined.
'Suitable for investigating hate speech, violent nationalist or separatist movements, radicalisation, extremist groups, incels, lone wolves...'	It can be used for hate speech and security threats.	Blurs lines between terrorism and ideology, speech and actions, allowing the state to treat nonviolent speech as threats.
'Leverage the knowledge and experience of stakeholder organisations and (social) media companies'	Work with unaccountable NGOs, think tanks, universities and social media firms	A public-private partnership in speech control – platforms and NGOs share data and align methods to advance the EU's Censorship Operating System.
'Includes training for PAs...as part of a long-term sustainable training network.'	Police will be trained to use the system.	Embeds these powers and tools into standard law-enforcement operations.
'The interdisciplinary consortium includes expertise from...social sciences, humanities, ethics, computer science, and four European Pas.'	Experts from different fields are involved.	Academia-washing: gives the project the appearance of objectivity and an aura of legitimacy.

In summary, VIGILANT is building a sophisticated AI-powered surveillance platform to enable European police authorities to police speech and political dissent online. Framed as a necessary tool for countering extremism and enhancing security, its real function is automating and detecting counter-narratives across all digital media to consolidate state control over democratic debate. Shrouded in the language of ethics and innovation, VIGILANT is the infrastructure for real-time ideological policing deployed by the police, who are trained to interpret and act on algorithmically flagged content.

2.5 VERA.AI: VERification Assisted by Artificial Intelligence²³

The vera.ai project positions itself as a technical response to the increasing sophistication of online disinformation, aiming to build AI-powered tools to detect fake content to support media professionals. Under Horizon Europe, it received €5,691,875 from the Commission. Coordinated by the Centre for Research and Technology Hellas (founded and supervised by the Greek Ministry of Development), it involves one technology research organisation, one not-for-profit research institute, six universities, two large media organisations and one prominent EU-supported NGO, the Brussels-based EU-DISINFOLAB. The project involves entities from eight member states.

Once again, the project's starting point is the unequivocal acceptance of the EU's 'hate speech' and 'disinformation' narrative. Its objectives begin with the statement that 'online disinformation and fake media content have emerged as a serious threat to democracy, economy and society'. This justifies its goal of centralising an algorithmically determined narrative detection system that can be used to verify online speech and information. And yet again, complexity is cited to belittle human judgement, which naturally requires 'trustworthy' AI algorithmic monitoring systems to compensate.

The constant reference to ‘fact-checking’ and integrating these AI tools simply means that ‘truth’ will be determined algorithmically based on training data sets defined by unaccountable entities hidden from the public. Vera.ai is a platform for standardising permissible speech and suppressing alternative narratives under the guise of safety, innovation and media ‘empowerment’.

In passing, it is striking to note how elastic the Commission’s idea of ‘safety’ becomes when it serves the interests of their narrative control. AI systems deemed too risky for general deployment under the EU’s own Artificial Intelligence Act are suddenly embraced when repurposed to police speech and information online. When defending the Commission’s messaging, the only precautionary principle on display is that against the dangers of free speech.

Again, we break down the NEUspeak to demonstrate how the Commission consciously manipulates language to accomplish its narrative promotion:

vera.ai project: NEUspeak analysis

Project in NEUspeak	in plain English	what it means
‘Disinformation and fake media content have emerged as a serious threat to democracy...’	Disinformation is dangerous and destabilising.	Democracy needs to be safeguarded through curbing free speech.
‘Highly realistic synthetic content...AI-powered bot networks...’	Deepfakes and spam bots are spreading misinformation.	This is technocratic language to amplify fear and create urgency for invasive AI monitoring.
‘Challenging for researchers and media professionals...’	Truth is relative.	Delegitimises human editorial judgement to justify algorithmic surveillance.

'Trustworthy AI solutions... co-created with and for media professionals'	We'll build tools with trusted journalists.	Embeds AI inside media workflows, outsourcing judgement to software.
'Fairness, transparency, robustness...explainability ... continuous adaptation ...'		Vacuous, ethical-sounding buzzwords will legitimise the work without clear accountability.
'Fact-checker-in-the-loop approach'	Humans will episodically be used to train and verify the AI.	The truth will be algorithmically determined, but not open to debate.
'Tools for deepfake detection in all formats'	It will scan videos, audio, images and text.	The extension of surveillance to all forms of expression.
'Grounding of the AI models on continuously collected fact-checking data...'	The AI learns from verified content.	Centralises truth by bypassing public scrutiny and deliberation.
'Expose disinformation campaigns and measure their impact.'	Find and track bad actors and measure their influence.	Policing narratives to shut down voices too influential to go unchecked.

This project has delivered three reports, one pilot prototype, a website, one data management plan and 12 peer-reviewed articles. It has been cited in 33 conference proceedings and related publications. The centralisation of narrative control under the language of 'trustworthy AI' and 'public safety' establishes dissenting or alternative perspectives as threats to be detected and neutralised, not debated.

2.6 TITAN: AI for Citizen Intelligent Coaching against Disinformation²⁴

The TITAN project is presented as an empowering tool to help citizens assess the truth of online information through AI-guided investigation and

fact-checking. It promises to improve critical thinking and media literacy via ‘intelligent coaching’ tailored to a user’s profile, skills and cognitive biases. TITAN is funded through the HORIZON.2.4 – Digital, Industry and Space programme for €5,734,395. A prominent private Italian company specialising in information-technology services, ENG, coordinates the project. It involves five universities, several not-for-profit research organisations and one media company, located across seven EU member states.

However, TITAN is constructing a surveillance-based, ideologically aligned, content-filtering system that uses AI to ‘assist’ users toward institutionally pre-approved interpretations of what is true or false. Through the language of ‘empowerment’ and ‘co-creation’, it aims to shape citizen reasoning along pre-determined lines, embedding a top-down fact-checking orthodoxy into personalised, algorithmic nudging. This project is about algorithmic obedience-training dressed in the soft language of digital citizenship. Far from promoting independent thought, TITAN aims to standardise it.

Again, a cute acronym disguises a malevolent intent. TITAN evokes strength, mythological grandeur and heroic protection, designed to conjure images of mighty guardians battling chaos. But far from empowering citizen’s agency, TITAN cloaks a paternalistic, AI-mediated guidance system aimed at enforcing the Commission’s narrative dressed in the language of heroism and autonomy. The project’s true nature is not to enable a force for independent investigation, but to create a soft-power structure designed to shape perceptions, manage behaviour, and standardise critical thinking according to the Commission’s narrative. It projects myth-making and awe-inspiring power while obscuring the fact that the ‘TITAN’ it upholds is the EU Commission.

It is a masterclass in NEUspeak. It claims to create an ‘open, distributed and citizen-engaging ecosystem’ where citizens are empowered to investigate the truthfulness of statements. This is achieved through ‘AI-driven, intuitive and personalised coaching’ that guides individuals in the ‘logical interpretation and critical assessment’ of claims, using fact-checking and media-literacy tools. The system adapts to the user’s profile, including their digital and media literacy skills, thinking difficulties and linguistic characteristics. The result, it claims, is the development of enhanced critical-thinking skills and the ability to ‘detect disinformation...at scale’.

Translating this into English, TITAN means the citizens will not fact-check themselves, but the platform will walk them through the ‘correct’ method, using a guided, AI-mediated interface that leads to the desired conclusion. This is an automated behavioural intervention masquerading as education. The AI coaching (or, more accurately, coaxing) is not neutral. It is trained on pre-determined definitions of truth, credibility and disinformation. ‘Critical thinking’ channels thought along approved lines, reducing citizens to profiles to be corrected, not free-thinking participants. The phrases ‘human-centred approach’ and ‘co-creation’ suggest inclusiveness, but these sessions are implementation-stage rituals used to create the appearance of public input. Citizens do not shape the system; it shapes them. As we have seen elsewhere, using the term ‘ecosystem’ deliberately gives the illusion of an organic, decentralised and collaborative dynamic. But what is described is a tightly controlled behavioural and educational pipeline that relies on a pre-determined algorithmic structure, coaching scripts and ideological assumptions about what needs to be indoctrinated.

TITAN is about imposing the Commission’s values and narrative. It is not censorship in the traditional sense, but a form of algorithmic

obedience-training dressed in the soft language of digital citizenship. The gap between its presentation and objective reveals the depth of the Commission's chicanery.

2.7 ORBIS: Augmenting participation, co-creation, trust and transparency in Deliberative Democracy at all scales²⁵

This project does not appear to have anything to do with hate speech or disinformation. It does not appear under these categories in the Horizon or Funding and Tenders databases. Yet, it is a project that promotes the same narrative.

ORBIS is a fully funded Commission project under the Horizon Programme for €2,498,882, scheduled to be completed in 2026. The Politecnico di Milano coordinates it. It involves seven universities, one NGO and seven independent research organisations across seven EU member states (and the UK Open University, which does not receive any funds).

ORBIS presents itself as a project to improve democracy using technology. Upon examination, it proposes a system for managing public participation, not expanding it. Wrapped in the language of democratic renewal, it is a technocratic project to legitimise the Commission's right to shape public opinion – the same assumption that underpins its 'hate speech' and 'disinformation' narrative. In both cases, the demos are regarded as data points robbed of agency where democratic political contestation becomes an institutionalised form of 'deliberation'. Democracy is engineered to reduce human judgement and political accountability. Its vision of 'inclusive, trustful democracy' is, in practice, an institutional apparatus that absorbs dissent, channels opinion and re-legitimises the status quo.

The table, overleaf, of the project's objectives demonstrates this clearly:

Project in NEUspeak	in plain English	what it means
'Addresses the disconnects between ambitious ideas and collective actions at a large socio-technical scale'	An abstract placeholder to justify intervention in how citizens behave collectively	We aim to engineer coherence between public expression and institutional priorities.
'A theoretically sound and highly pragmatic socio-technical solution'	We've designed a credible, practical system combining social and tech elements.	We're using academic legitimacy to obscure that the solution is top-down and expert-led, driven by the belief that AI technology can solve the problems of democracy.
'Enable the transition to a more inclusive, transparent and trustful Deliberative Democracy in Europe'	We want people to feel included and trust the discussion-based process rather than a politically contested democracy.	Impose 'trust' in institutions through engineered participation rather than political accountability.
'Deliberative democracy is not a theory...but a process through which we can collectively imagine and realise new practices'	We're not prescribing democracy, but facilitating its evolution through managed dialogue.	We aim to depoliticise democracy by turning it into a soft, non-confrontational workshop exercise, avoiding political demands.
'AI-enhanced tools for deliberative participation across diverse settings'	We're using AI technology to manage participation in debates and discussions.	Algorithms will provide agency rather than the demos, shaping speech and participation through data-driven filtering.
'Demonstrated measurable impact of such innovations in real-world settings'	We will test our tools and show results from real use-cases.	Tightly controlled use-cases will simulate performative democratic interaction, not give citizens political power or decision-making authority.

The language of ‘deliberative democracy’ and ‘citizen participation’ is a framework for engineering and controlling political engagement to align with the European Commission’s anti-democratic narrative. It redefines what a legitimate public debate looks like. This has little to do with political contestation but with the steering of collective opinion within boundaries set by the Commission. The technological solutions offered filter dissent to institutionalise a top-down consensus. The connection between this and the ‘hate speech’ and ‘disinformation’ narrative is that it reframes free speech not as a fundamental tenet of democracy – as a right – but as a managed process best supervised by experts to ensure ideological compliance.

It should be apparent from the above examples that the Commission’s funded projects are far more than efforts to tackle ‘hate speech’ or ‘disinformation’. (To help the reader and provide a reference point, Appendix I contains two tables that codify the key common phrases consistently used in many of these projects, which are decoded from NEUspeak into English.)

The narrative codifies a coordinated attempt to reshape the boundaries of acceptable speech, thought and political engagement across Europe. Funding this is a conscious effort to build and sustain a network of collaborators whose compliance legitimises the de-legitimisation of any alternatives or political challenges. The sneaky language of ‘safety’, ‘trust’, ‘inclusion’, ‘empowerment’ and ‘capacity building’ is a ruse to establish a semantic architecture determining the language of the public conversation. In the name of democracy and transparency, the Commission has built a narrative regulated and enforced by a network of speech police – supplemented by algorithmic tools, behavioural conditioning and centralised oversight. The resulting surveillance infrastructure aims to embed ideological compliance into the digital life of Europe’s demos, particularly among younger people. In the name of countering ‘hate speech’ and ‘disinformation’, the Commission

is crusading to re-engineer the information environment and the language of public life – from the language used to determine who is allowed to speak – to ensure that only state-sanctioned narratives remain visible, legitimate, and thus, repeatable.

This deliberately engineered network controls the European political narrative and embeds freedom *from* speech as the EU's pre-eminent value for the twenty-first century.

3 Conclusion: the emperor has no clothes

The above analysis reveals the paradox at the heart of the EU Commission's crusade against 'hate speech' and 'disinformation'. Without the existence of these spectres that must be exorcised, the institution would lose a key plank in its quest for legitimacy. The level of spending reveals an inconvenient truth: the EU relies upon the institutionalisation of the very thing it purports to want to extinguish. And it is willing to spend over half a billion euros to sustain it.

This should not be understood in any conspiratorial sense. The EU does not promote hate speech or disinformation; its target is free speech. They fear free speech because of its unpredictable energy, because it enables alternative narratives to be voiced and considered, and – horror of horrors – because it suggests that European citizens still retain moral independence and thus the ability to tell truth from lies and information from disinformation without the need to defer to experts or unelected technocrats who allegedly know what's best for them. Half a billion euros is a small sum if it enables the politically unaccountable EU technocrats to dictate what can or can't be said or thought in Europe. It is an authoritarian graft machine deployed to manage Europe's 'democracy' in the interests of maintaining the status quo.

The projects funded to sustain the narrative and the Commission's need to manage this constantly reveal that public money has been deployed to create one of the world's largest ideologically compliant echo chambers in history.

Policing the narrative is not an abstract or theoretical endeavour; it is creating an ideological and material force in the fabric of Europe's political life. The network of surveillance created sustains the narrative not simply by echoing and 'proving' the assumptions underpinning the narrative. They sustain themselves and the intuitions and people they employ by aligning their material interests with the Commission, whose pragmatism has no honourable boundaries apart from self-service.

The Commission pays lip service to evidence-based policymaking. However, our research exposes that what they consciously practice is, in fact, 'policy-based evidence', presented as 'research' to give it a veneer of academic legitimacy. It is the outsourcing of the manufacturing of propaganda masquerading as evidence-based policy. The fact that so many independent think tanks and universities are willing to collaborate in this charade speaks to a far more serious problem: the corruption of the academy and the end of the very idea of objective inquiry. The damage this does to democracy – indeed, to the future of expertise – is incalculable and is beyond the scope of this paper.

This is no accident. It is not a funding flaw or oversight, but structural and embedded in the Commission's DNA. It is doubtful that any organisation that has bid for project funding that disputed the asserted premise – that hate speech and disinformation are a rising threat to democracy – would see a penny of backing. Any legitimacy offered to a counter-narrative would disrupt the self-serving recursive reinforcement loop these projects are designed to provide. The dishonesty at the heart of this crusade is monumental: pre-determined results are presented as if they were the impartial outcome of 'independent' research and practice.

However, the fact that the Commission must resort to such machinations and is willing to spend 31 per cent more on this rather than on cancer research

exposes the Commission's weakness, not strength. The need to constantly manufacture an artificial consensus (maintained, as we have seen, through the liberal deployment of AI algorithms) reveals this narrative has no organic connection to the social reality it purports to reflect. It is a top-down conceit, sustained to legitimise the status quo that millions of Europeans are now questioning and speaking out against. €649million is a lot of money to hide the fact that the Commission can only rule through negative authority and manipulation; indeed, it has no clothes.

We have demonstrated that a fevered battle over language is underway in Europe. It is hidden and pernicious. It is about the legitimacy of the EU and its authority to dictate the terms of public communications. As we have demonstrated, this is a top-down, authoritarian, curated consensus, where expression is free only when it speaks the language of compliance established by the Commission.

We have catalogued this deeply dishonest Orwellian crusade that, unlike historical attempts to outlaw free speech, does not burn books or squash dissent with jackboots. Instead, it relies on a carefully created machinery of speech policing, supported by a vast system of public financing. The aim of the enormous financing operation is to present the thought police as being on the 'right side of history'. Imbued with the authority of hundreds of NGOs and academics, it is even able to claim that it polices speech in order to 'protect democracy'.

As we stated in the introduction, the Commission rightly understands that controlling the language of communications means it can dictate what is information and disinformation, truth or lies, what is legitimate or illegitimate speech, and who can speak or not. And if it can control information and the truth, it controls history, the past and the future.

CONCLUSION: THE EMPEROR HAS NO CLOTHES

Exposing how taxpayers' money is being used without any public accountability is an important step. However, exposing the language war is an even more necessary act of democratic vigilance. When language is narrowed, softened, obfuscated or stripped of meaning, so is the possibility of resistance and the development of alternatives. To borrow the much-abused language used by the Commission, this report is an essential step to 'raise awareness' about how the Commission is fatally undermining free speech and democracy in the name of democracy and free expression.

Appendix, 'NEUspeak' into English

Below is a table of key phrases found in numerous Project Objectives Statements. For the sake of those who cannot speak Orwellian 'newspeak', we have translated each phrase into English and then explain, in some instances where it is not apparent, what this means in reality.

Project objectives in NEUspeak	Translated into English
To promote resilience against disinformation and hate speech	We want to manage how people interpret information, define what counts as harmful or fake, and promote institutional narratives over alternative or dissenting ones. Meaning: we are recruiting you to become part of the EU's self-appointed speech police
To foster media literacy and critical thinking	We want people to be trained to recognise and reject <i>certain types of content</i> while implicitly trusting EU institutional sources Meaning: help us criminalise speech we fear
To develop innovative tools for the detection and counteraction of online hate speech and disinformation	We want automated systems to flag and suppress content we consider harmful, with limited human oversight and questionable accuracy. Meaning: outlaw nuance, context, and free speech
To enhance cross-sectoral cooperation among stakeholders (for example, civil society, tech platforms, journalists, fact-checkers)	We'll bring together NGOs and research organisations we trust to co-create narratives, strategies and solutions. Meaning: we set the terms that can sustain the Censorship Operating System

To support inclusive and rights-based approaches in tackling harmful content	We will restrict or regulate speech as inclusive and rights-based – even if the affected communities had little say in how harm was defined or addressed Meaning: identity politics means that freedom <i>from</i> speech is the new defence of 'free speech'
To empower citizens and communities to respond to digital threats	We want to use the foil of the 'community' to help in the monitoring, reporting and reinforcement of approved narratives. Framing it as empowerment gives us the power to determine legitimacy.
To raise awareness and build capacity in identifying and countering disinformation	We want to train people to recognise and reject information that challenges the institutional consensus – while presenting this as capacity-building and civic education. Meaning: build a network of approved unaccountable fact-checkers and informers
To contribute to a safer, more inclusive digital environment	We want to regulate speech and shape digital behaviour in ways that align with the Commission's values while framing the process as inclusive, safety-driven and responsible
To address the root causes and societal impacts of hate speech	We aim to re-educate the public and reinforce the need for deference to experts
To monitor and assess the effectiveness of countermeasures	We'll evaluate our interventions based on metrics we define, which will prove why the problem remains hidden and thus needs more research Meaning: this is an ongoing quest which must be sustained and constantly proved to be necessary

The same can be seen in the themes and policy-linked terms repeatedly referred to:

'Newspeak' themes	Translated into English
Democratic resilience	We will define and defend democracy by controlling speech, marginalising dissent, and shaping acceptable narratives – primarily online. Meaning: help redefine democracy
Civic engagement	We want people – especially youth – to participate in ways we recognise and can manage Meaning: exclude the general public
Fundamental rights and freedoms	We're going to protect freedom of speech – unless we believe it threatens a victimised minority's rights, at which point we'll restrict it in the name of free speech Meaning: help redefine free speech as freedom from speech
Transparency and accountability	We'll hide all our spending in plain sight in numerous databases, publish some high-level reports, issue vague commitments, and say we're listening – while keeping real decision-making and power structures opaque.
Multi-stakeholder governance	We let a curated group of unelected and unaccountable 'stakeholders' sit at the table – but power and decision-making still sit with us Meaning: exclude the general public

APPENDIX, 'NEUSPEAK' INTO ENGLISH

Online platforms and algorithmic accountability	We'll require and publish reports and compliance steps from platforms, but users still won't have control or insight into the systems governing their digital lives. Meaning: use AI to police speech in ways that are opaque to the public
Fact-checking and debunking	We decide which narratives are legitimate, and we partner with institutions or approved NGOs to signal what the public should trust Meaning: prop up our right to decide what is a fact, hate or disinformation, or you don't get a piece of the funding action
Narrative building and counter-narratives	We craft state-aligned messaging that looks grassroots – then promote it as an antidote to oppositional or radical content. Meaning: criminalise populist or far-right political organisations and ideas
Intercultural dialogue	We invite people to engage in polite, depoliticised conversations celebrating diversity but not power. Meaning: indoctrination of people from different cultural, ethnic, and religious backgrounds
Safe and trusted online environment	We will define what is harmful and what is trustworthy – and regulate online spaces accordingly, even if it means censorship and centralised control Meaning: make us look benign, disinterested, but responsible governors.


End notes

- 1 Norman Lewis, *Controlling the Narrative: The EU's attack on online speech*, MCC Brussels, May 2024 brussels.mcc.hu
- 2 This is supplemented by the unaccountable funding of the media, which serves the same purpose and which is covered in a forthcoming MCC Brussels report by Thomas Fazi.
- 3 Made up of €387million for 2021–23 and €116million for 2024. See 'Working Together', EU Commission Research and Innovation ec.europa.eu and '2024 Horizon Europe call under the Mission Cancer is out!', European Institute for Biomedical Imaging Research eibir.org
- 4 See Tadeusz Dyba, et al, 'The European cancer burden in 2020: Incidence and mortality estimates for 40 countries and 25 major cancers', *European Journal of Cancer*, November 2021 pmc.ncbi.nlm.nih.gov
- 5 EU funded projects: mental health, EU Funding & Tenders Portal ec.europa.eu
- 6 'Social Media: Measuring Effects and Mitigating Downsides', EU Funding and Tenders Portal ec.europa.eu
- 7 (EM)POWERYOU(TH): a co-created model towards a non-discrimination culture (empoweryouth), EU Funding & Tenders Portal ec.europa.eu
- 8 See 'The LGBTIQ Equality Strategy 2020–2025' commission.europa.eu 'The EU Anti-Racism Action Plan 2020–2025' commission.europa.eu, 'The new EU Roma strategic framework for equality, inclusion and participation (full package)' commission.europa.eu, 'Gender Equality Strategy 2020–2025' commission.europa.eu, 'Union of equality: Strategy for the rights of persons with disabilities 2021–2030' commission.europa.eu and 'EU Strategy on Combating Antisemitism and Fostering Jewish Life, 2021–2030' commission.europa.eu
- 9 'Financing of pilot projects and preparatory actions in the field of "Communications Networks, Content and Technology" and on the adoption of the work programme for 2022', 22 February 2022. See the 'Annex to the Commission decision': tinyurl.com
- 10 NODES: Narratives Observatory combatting Disinformation in Europe Systemically nodes.eu
- 11 Reimagine Europa re-imagine.eu
- 12 'The European Narratives Observatory NODES presents its work on New Narratives for Climate Action' Media Connect, 9 December 2024 mediaconnect.com
- 13 'The consortium led by Opsci will carry out the work of the European Narrative Observatory to combat disinformation', Opsci, January 2024 opsai.ai
- 14 FAST LISA: Fighting hAte Speech Through a Legal, ICT and Sociolinguistic approach, EU Funding & Tenders Portal ec.europa.eu
- 15 The project's objectives and the organisations involved can be seen here: FAST LISA, *ibid.*
- 16 'REASON: REAct in the Struggle against ONLINE hate speech', EU Funding & Tenders Portal ec.europa.eu
- 17 Istituto per La Ricerca Sociale irsonline.it
- 18 L'Associazione Carta di Roma cartadiroma.org
- 19 Istituto per La Ricerca Sociale irsonline.it
- 20 'RECO-DAR: Right-wing extremist eco-systems driving hate speech: dissemination and recruitment strategies', EU Funding & Tenders Portal ec.europa.eu
- 21 SCENOR – VEREIN ZUR ERFORSCHUNG AKTUELLER GESELLSCHAFTLICHER HERAUSFORDERUNGEN, EU Funding & Tenders Portal ec.europa.eu
- 22 'VIGILANT: Vital Intelligence to Investigate ILlegAl DisiNformaTion', Cordis cordis.europa.eu
- 23 'vera.ai: VERification Assisted by Artificial Intelligence', Cordis cordis.europa.eu
- 24 'TITAN: AI for Citizen Intelligent Coaching against Disinformation', Cordis cordis.europa.eu
- 25 'ORBIS: Augmenting participation, co-creation, trust and transparency in Deliberative Democracy at all scales', EU Funding & Tenders Portal ec.europa.eu

About the author

Dr Norman Lewis

Visiting Research Fellow



Dr Norman Lewis is a writer, speaker and consultant on innovation and technology. He is recognised worldwide as an expert on future trends and user behaviours regarding technology innovation and adoption. Norman was formerly a Director at PwC, responsible for running its crowdsourced innovation programme. Prior to this, he was the director of technology research at Orange. He is an advisory board member of Bubbletone Blockchain in Telecom – the world's

first decentralised mobile roaming service.

He was also an executive board member of the MIT Communications Futures Programme and a former chairman of the ITU TELECOM Forum Programme Committee. He is a co-author of *Big Potatoes: the London manifesto for innovation*.

About MCC Brussels

At a time of unprecedented political polarisation, MCC Brussels is committed to providing a home for genuine policy deliberation and an in-depth exploration of the issues of our time.

MCC Brussels is committed to asking the hard questions and working with people of goodwill from all persuasions to find solutions to our most pressing problems. An initiative of MCC (Mathias Corvinus Collegium), the leading Hungarian educational forum, MCC Brussels was founded in the autumn of 2022 to make a case for celebrating true diversity of thought, diversity of views, and the diversity of European cultures and their values.

Unmasking the EU's €650million war on free speech

A spectre is haunting Europe – not of disinformation, but of linguistic control and censorship, coming from the heart of the EU Commission. This report exposes a covert campaign conducted by the European Commission to regulate the boundaries of legitimate public debate in Europe. It has used hundreds of millions of euros in taxpayer money to fund an Orwellian disinformation complex and what might be called the EU Ministry for Narrative Control.

Through hundreds of unaccountable non-governmental organisations and universities carrying out 349 projects, masked by ambiguous 'NEUspeak' and euphemistic terms like 'deliberative democracy' or 'capacity building', the Commission is waging a silent war to regulate language. The objective is the de-legitimisation of alternative narratives, like the rising tide of populist opposition.

What's presented as combating 'hate speech' and 'disinformation' is, in fact, a systematic assault on free speech in Europe, designed to construct an ideological infrastructure for controlling political narratives and shaping public opinion. This is a top-down, authoritarian, curated consensus, where expression is free only when it speaks the language of compliance established by the Commission.