SCRIPTORIUM
INITIATIVE

THE *[_SYNTHETIC]* CEO IS READY TO SPEAK.

IS YOUR ORGANIZATION READY TO TAKE THE RESPONSIBILITY?

SCRIPTORIUM-INITIATIVE.AI

MAKING SENSE OF THE *[_UNKNOWN]*

EXECUTIVE SUMMARY

Organizations can now create a recognizable synthetic representation of an executive, generate video from approved text, reproduce a voice and adapt the communication across languages and markets. This can make executive communication more accessible, consistent and scalable. It does not transfer the judgement or accountability of the person being represented.

The Scriptorium Initiative studied how experienced communications professionals responded to human and synthetic CEO communication. Respondents did not reject the synthetic CEO simply because it was artificial. Their willingness to consider it depended mainly on whether it solved a useful problem and whether the organization could manage the associated risks.

Context was important. Routine and approved communication was seen as the most appropriate starting point. There was greater concern about internal culture communication and little support for using a synthetic CEO in crises, layoffs or major organizational change. The higher the emotional, financial or moral stakes, the stronger the preference for the real executive to be present.

The research also suggests that hesitation is not primarily a matter of age, seniority or familiarity with AI. Respondents focused on the use itself. They wanted clear disclosure, human approval and explicit boundaries. They also expected the organization to remain answerable for every message and to provide access to a responsible person when necessary.

This whitepaper translates those findings into three practical frameworks. The Authority Test defines what a synthetic CEO may say and do. The Human Minimum Standard identifies the situations in which a real executive must appear. The Synthetic CEO Protocol sets out the evidence and controls required before deployment, including purpose, consent, authority, disclosure, approval, provenance, escalation and exit.

The Chief Communications Officer should lead or formally co-lead this decision. The CCO is best placed to assess what audiences may infer from the synthetic CEO and whether its use supports or weakens trust in the organization. Legal, privacy, security, HR, investor relations, compliance, brand and, where appropriate, the board should contribute to the decision and validate their respective controls.

The recommended approach is limited and deliberate. Begin with a disclosed, low-stakes use case that addresses a measurable stakeholder need. Maintain an approval record, use a verified channel and provide a clear route to a person. Reconsider approval whenever the system gains new capabilities, such as memory, live interaction, personalization or access to internal systems.

Synthetic communication can extend the reach of an executive. Judgement, vulnerability and responsibility must remain human.

Amsterdam, September 2026

Nanne Bos
Co-Founder of the Scriptorium Initiative

AT 8:07 AM THE CEO APPEARED ON EVERY SCREEN

Imagine a company announcing a difficult restructuring. At 8:07 a.m., a calm video from the CEO reaches every employee at once, translated into twelve languages and adapted for each market. The CEO never recorded it. The face is synthetic, the voice is cloned and the performance comes from an approved script. A disclosure appears beneath the video. Nobody has been technically deceived, yet the company has created the experience of leadership without requiring the leader to appear.

AI can now create a credible representation of an executive, and several companies are already experimenting with it. This raises a practical question: how should organizations use synthetic executive presence when accountability still rests with real people?

1. THE CEO IS MORE THAN A MESSENGER

Executive communication is often treated as a distribution problem. The CEO has limited time, while employees, investors and customers expect consistent access across markets. Synthetic representation offers an obvious answer: capture the leader once and reproduce that presence everywhere.

A CEO is more than a high-value content creator. After an accident, scandal or strategic failure, audiences look for more than information. They ask whether the leader is willing to be seen, questioned and held to the next decision.

A synthetic CEO therefore differs from a training avatar or product explainer. It does not only reproduce a face and voice. It borrows institutional authority. Klarna and Zoom illustrated the attraction in 2025 when AI avatars of their CEOs appeared in earnings-related communication. These were controlled, scripted experiments that also performed the companies' claims to be AI-forward.

Consequential steps

The next step is more consequential. The Guardian reported in 2026 on a proposed AI version of Mark Zuckerberg for employee communication. A representative that responds rather than only speaks is no longer a media asset. It becomes a leadership interface. Companies will find that interface difficult to resist. Executive attention is scarce, demand for executive presence is not, and every channel rewards speed, consistency and personalization.

The business case is immediate. A synthetic CEO is always available, consistent across markets and able to produce a localized message before the real CEO leaves the previous meeting. For communications teams, the promise addresses four persistent problems.

- **Scale**: one executive presence can reach thousands of audiences and formats.
- **Consistency**: approved language can be reproduced without improvisation or drift.
- **Localization**: messages can be translated and culturally adapted at speed.
- **Control**: every word, pause and emotional cue can be reviewed before publication.

Each advantage also changes the relationship with the audience. Scale can create distance, consistency can become sterility, and localization can simulate intimacy. Tight control may remove the vulnerability through which leaders demonstrate commitment.

The same features that strengthen the business case can therefore create the reputation risk: optimizing the signals of leadership while reducing the human exposure that makes them credible.

2. THE TECHNOLOGY IS MOVING FASTER THAN GOVERNANCE

Organizations can now create a recognizable synthetic representative from recorded material, generate video from approved text, clone a voice and localize the result. Costs and production effort have declined, although capabilities and governance safeguards still vary by product and use case.

Scripted corporate video has moved beyond demonstration. Interactive representatives remain less mature. The phrase synthetic CEO covers both, even though a fixed video and a live agent create different risks.

A reusable likeness

The most established capability is the scripted synthetic executive. A company trains a digital representation of an executive and generates new videos from supplied text. Some platforms also reproduce aspects of expression, movement and delivery style.

Vendor documentation illustrates the pace of development. HeyGen describes persistent digital representatives created from short reference video. Synthesia supports enterprise production and translation. Tavus distinguishes scripted generation from real-time conversation. These product claims describe technical capability, not organizational readiness.

For routine and approved messages, this layer is technically deployable. It can perform words supplied by the organization. It cannot show that the executive personally delivered, reviewed or saw every version.

A voice that can say anything

Voice cloning raises the stakes because the system can generate speech the executive never recorded. ElevenLabs distinguishes instant cloning from professional cloning trained on more material. Its documentation requires verification for professional clones. Product safeguards matter, but the organization must still control who can generate speech, for which topics and for how long. Consent must therefore be managed throughout the lifecycle. The organization needs rules for access, permitted topics, approval, storage, revocation and executive departure.

Multilingual presence

Localization is one of the strongest use cases. Major platforms offer multilingual generation or translation, allowing headquarters to reach markets faster and more accessibly. The benefit should not be confused with personal interaction. An employee may see the CEO apparently speaking their language even when the CEO does not speak it and never delivered that version.

A representative that answers back

The largest shift is from generated video to real-time interaction. Products such as HeyGen LiveAvatar and the Tavus Conversational Video Interface combine listening, language-model responses, synthetic speech and a visual replica.

At this point the synthetic CEO becomes a corporate actor. It must interpret questions, retrieve information and choose how confidently to answer. Connections to memory, company data or tools can also personalize a response or trigger action.

The executive's face and voice can make those machine decisions appear more authoritative than they are. A fluent response may feel like executive guidance even when it is an inference from a language model.

What is mature and what is not

Scripted synthetic presence is technically deployable for controlled, approved communication. Interactive presence is also deployable, but its reliability depends on sources, configuration, testing and human escalation. Neither form can exercise executive judgement or become answerable for consequences.

Technical capability tells an organization what it can build. It does not establish when a use is legitimate. That permission decision belongs with the Chief Communications Officer, working with the CEO and the functions that own legal, technical and operational risk.

3. WHO GIVES THE SYNTHETIC CEO PERMISSION TO SPEAK

Whether synthetic presence is appropriate is a communications question before it is a production question. The CCO must judge what audiences may reasonably infer, when disclosure is sufficient and when only the real leader will do. The Scriptorium Initiative studied this permission problem because synthetic representation changes the relationship between message, messenger and institutional responsibility.

Working with Damian Phan, a student at the University of Amsterdam, Scriptorium investigated the conditions under which experienced communication professionals would consider using a synthetic CEO. Respondents watched one of two short CEO videos, one human-delivered and one synthetic. They assessed adoption, usefulness, risk, authenticity, ease of use, suitable contexts and governance conditions.

A sample close to the decision

The study included 78 valid respondents with relevant experience. Of these, 67.9% were executive-level and 25.6% senior-level. More than 85% had at least thirteen years in corporate communication, and 53.8% worked in organizations with more than 5,000 employees.

The largest national groups came from the Netherlands, the United States, Germany and Switzerland. The sample is relevant to executive communication and governance, but it is not globally representative.

Sample characteristic	Result
Valid respondents	N = 78
Seniority	67.9% executive-level; 25.6% senior-level
Experience	85.9% had 13+ years in corporate communication
Organization size	53.8% worked in organizations with 5,000+ employees
Largest country groups	Netherlands 30.8%; United States 16.7%; Germany 15.4%; Switzerland 11.5%

Respondents noticed the difference

Respondents in the synthetic condition recognized the video as AI-delivered, while those in the human condition rated their video as more human-delivered. The synthetic CEO was not avoiding detection. Some suspicion remained even in the human condition. As synthetic media spreads, genuine executive communication may also need stronger proof of origin.

Finding 1. Synthetic was not automatically rejected

Respondents who saw the synthetic CEO were no less willing to consider using it than those who saw the human CEO. However, willingness was modest in both groups. The synthetic CEO did not prompt additional rejection, but neither did it attract broad acceptance.

Finding 2. Risk created the penalty

Respondents perceived the synthetic CEO as riskier than the human CEO. At the same time, they did not judge it as meaningfully less authentic, useful or easy to use. Risk, rather than the quality of the representation, was the main point of difference.

The management question is therefore not simply whether a synthetic CEO appears credible. It is whether the benefits justify the risks, including inaccurate messages, misuse of an executive's likeness, uncertainty about authority, employee overreliance and reputational damage.

Improving the representation will not address these concerns. Organizations need clear approval processes, visible disclosure, auditable records, defined authority boundaries and a reliable route to an accountable person.

Finding 3. Usefulness pulls and risk pushes

Respondents were more willing to consider a synthetic CEO when they saw a clear practical benefit, and less willing when they perceived greater risk. Usefulness and risk mattered more than authenticity or ease of use when these factors were considered together. These findings describe relationships within this sample and should not be treated as universal conclusions.

The decision therefore rests on a practical balance: does the synthetic CEO solve a meaningful stakeholder problem, and can the organization manage the associated risks?

Decision force	Evidence	Management implication
Usefulness	$r = .695$; $\beta = .470$; $p < .001$	Specify the stakeholder benefit and why a synthetic representative is needed
Perceived risk	$r = -.620$; $\beta = -.292$; $p = .003$	Show the controls, boundaries, owners and failure response
Authenticity	$r = .413$; not independently significant in the model	Treat realism as an entry condition, not the business case

Decision force	Evidence	Management implication
Ease of use	r = .316; not independently significant in the model	Operational convenience cannot compensate for weak governance

Finding 4. AI confidence did not divide the sample

Most respondents were confident in their understanding and use of AI, but this did not make them more willing to adopt a synthetic CEO.

Willingness also did not vary meaningfully by professional background, seniority, experience, industry, organization size or education. Within this sample, hesitation was not primarily a matter of age, status or AI literacy. It reflected concerns about the use itself.

Finding 5. Context created a professional boundary

Respondents drew a clear distinction between different forms of executive communication. Routine organizational announcements were seen as the most acceptable starting point, although support remained cautious.

Internal updates and culture-related communication raised greater concern. There was little support for using a synthetic CEO in crises, layoffs or major organizational change.This pattern appeared across roles, levels of seniority and organization sizes. The higher the emotional, financial or moral stakes, the stronger the preference for the real leader to be present.

Finding 6. Control creates permission

Clear disclosure was the most important condition for considering adoption. Respondents also wanted human review before publication and explicit guidelines defining how the synthetic CEO could be used.

These safeguards work together. Disclosure helps audiences understand what they are seeing. Human approval establishes responsibility for the message. Usage guidelines prevent a limited experiment from gradually moving into more sensitive areas.

Before launch, the organization should define the value, assess the risks, assign responsibility and establish clear boundaries.

What the study establishes and what it does not

The sample is experienced but modest, predominantly senior, international and Western. It should not be treated as a universal estimate of employee or public attitudes. Apart from the experimental comparison between video conditions, the analyses show relationships rather than causality. The findings are directional evidence for pilot design, governance and further research.

The research moves the debate from detection to permission. Participants assessed whether the usefulness justified the risk and whether the organization could control the expectations created by putting a human face on a machine-generated message.

Those expectations can exceed the system beneath them. The face implies presence, the voice implies identity, and a direct answer implies knowledge and authority. The CEO may never have seen the message, and the dialogue may never reach anyone able to change a decision.

4. THE ACCOUNTABILITY GAP BEHIND THE HUMAN INTERFACE

The familiar uncanny valley describes a face that looks almost human. A synthetic CEO creates a different kind of discomfort: an organization may appear present and responsive without providing the human involvement that audiences expect.

A synthetic representative can speak in the executive's voice, answer questions and express concern. Yet the executive may not have reviewed the message, heard the question or become involved in the response. The interface and the organization can therefore communicate different things. The interface suggests access, attention and authority. The organization may only be providing information generated from approved material.

This gap matters because audiences respond not only to what an executive says, but also to what the executive's presence represents. A CEO who appears after an accident, a restructuring or a strategic failure is signaling a willingness to be associated with the consequences. A synthetic CEO can reproduce the message, but it cannot assume that responsibility.

Five risks behind the human interface

1. Accountability laundering

A synthetic CEO can create the appearance of leadership presence without exposing the leader to questions or challenge. The message may be accurate and approved, while the organization appears more accountable than it is.

2. Simulated access

An interactive representative may give employees or other stakeholders the impression that they are engaging with leadership. If their questions do not reach responsible people or influence decisions, the organization is offering an information service rather than genuine dialogue.

This distinction should be made clear. A system that answers questions is not necessarily a system that listens.

3. Authority confusion

Audiences may not know whether a response is an approved statement, a summary of existing information, an interpretation or a commitment. Using the executive's face and voice can give a generated answer more authority than the underlying system deserves.

The risk increases when the synthetic representative moves beyond delivering a fixed script and begins responding to questions. A fluent answer can sound like executive guidance even when it is generated from incomplete information or general instructions.

4. Automated empathy

A synthetic representative can reproduce the language, tone and appearance of concern. This may be useful in routine communication, but it becomes more sensitive when the subject involves harm, grief, misconduct or organizational responsibility.

In such situations, audiences may reasonably expect the real leader to appear. A synthetic expression of concern can otherwise create distance at the moment when personal involvement matters most.

5. Authenticity contamination

As synthetic executive communication becomes more convincing, genuine communication may also become easier to question. Employees, journalists and other stakeholders may become uncertain about whether a message was recorded by the executive or generated on the executive's behalf.

Organizations therefore need to protect both forms of communication. Synthetic messages must be clearly identified, while genuine executive appearances may require stronger evidence of origin.

Trust requires more than realism

Improving the quality of a synthetic CEO may make the representation more convincing, but it does not resolve questions of authority, responsibility or accountability. A realistic face does not show who approved the message. A familiar voice does not establish whether the information is current. A confident answer does not prove that the executive agrees with it.

Trust must therefore be based on information that audiences can verify. They should be able to understand that the representative is synthetic, know what it is authorized to say, identify the sources behind its answers and reach an accountable person when necessary.

Provenance tools and content credentials can help establish where a piece of media came from and whether it has been altered. Disclosure requirements can also make the synthetic nature of an interaction clear. These measures are useful, but they do not answer every question. They cannot determine whether a response is appropriate, whether it can be relied upon or whether the real executive would defend it under challenge.

The central governance question is therefore not whether the synthetic CEO looks or sounds real. It is whether the organization can stand behind what it says and meet the expectations its presence creates. A synthetic CEO may extend the reach of executive communication. It should not create the impression of authority, access or accountability that the organization cannot provide.

5. DEFINING THE BOUNDARY

Organizations considering a synthetic CEO need to make two decisions. They must define the authority given to the synthetic representative and identify the situations in which only a real executive will be appropriate.

The Authority Test and the Human Minimum Standard provide a framework for making these decisions. The first governs what the synthetic CEO may say and do. The second identifies when executive presence is itself part of the organization's responsibility.

The Authority Test

The Authority Test assesses a synthetic CEO by its permitted actions, rather than by how realistic it appears.

The synthetic CEO may...	The synthetic CEO must not...
Render a fixed, approved statement.	Imply that a generated statement is spontaneous.
Retrieve current, approved source material.	Invent policy, facts, commitments or executive views.
Explain routine information within a defined scope.	Handle ambiguity as if it possesses executive judgement.
Translate and personalize without changing meaning.	Turn localization into undisclosed simulated intimacy.
Escalate questions to accountable humans.	Simulate listening when no human is prepared to act.

The important boundary is between communicating authority and exercising it. Delivering an approved statement resembles publication. Interpreting strategy, recommending action, expressing a new executive view or making a commitment involves executive judgement.

This boundary should be defined before the synthetic CEO is introduced. It should specify the subjects it may address, the sources it may use, the audiences it may engage and the circumstances in which it must stop and refer the matter to a person.

Authority should also be reconsidered whenever the system changes. Adding memory, live interaction, personalization or access to internal systems alters what the synthetic representative can know and do. Each additional capability therefore requires fresh approval.

Without this review, a system introduced to deliver approved communication can gradually begin to interpret information and influence decisions. At that point, it is no longer functioning only as a communications channel.

The Human Minimum Standard

Defining the authority of the synthetic CEO is not sufficient. Organizations must also identify situations in which the real executive needs to be present.

The research suggests that routine communication offers the most acceptable starting point, although support remained cautious. Respondents expressed greater concern about internal culture communication and little support for using a synthetic CEO in crises, layoffs or major organizational change.

The higher the emotional, financial or moral stakes, the stronger the expectation of human involvement.

Start here	Proceed carefully	Keep human
Training, onboarding, approved operational updates, product explanations	Scripted executive updates, multilingual messages, internal culture communication	Layoffs, fatalities, misconduct, crisis apologies, trauma, moral judgement
Low stakes; repeatable; easily corrected	Leadership symbolism; public interpretation; reputational exposure	Irreversible consequences; vulnerability; need for challenge and accountability

The Human Minimum Standard identifies situations in which the presence of the real leader is part of the organization's response. In these moments, audiences are looking for more than accurate information. They want to see who accepts responsibility, who is prepared to listen and who can make or change a decision.

AI may still support the communication. It can help prepare the message, anticipate stakeholder reactions, translate materials and coordinate distribution. It should not replace the person whose presence demonstrates that the organization understands the significance of the situation.

The standard is therefore not based only on the subject of the communication. Context also matters. A routine operational update may become sensitive if it affects jobs, safety, personal rights or public trust. Communications leaders should assess both the content and the expectations surrounding it.

The two frameworks work together. The Authority Test defines what the synthetic CEO is allowed to do. The Human Minimum Standard defines when it should not be used at all. A synthetic CEO may communicate approved information within a defined scope. When stakeholders need judgement, vulnerability or proof that someone is answerable, the real executive must appear.

6. GOVERNING THE SYNTHETIC CEO

A synthetic CEO may first appear as a technology purchase or a production experiment. But the decision to use an executive's likeness, voice and authority is ultimately a matter of communications governance.

The Chief Communications Officer should own, or formally co-own, that decision. The CCO is best placed to assess what audiences may infer from the synthetic CEO, whether its use is appropriate and how it may affect trust in the executive and the organization.

This responsibility cannot sit with communications alone. The CCO should bring together the CEO, legal, privacy, security, HR, investor relations, compliance and brand. The board may also need to be involved when the proposed use carries significant reputational, financial or governance implications.

Each function remains responsible for its own controls. The CCO's role is to bring those controls together and protect the integrity of the organization's leadership communication.

Four decisions before approval

Before approving a synthetic CEO, the organization should answer four questions in plain language:

1. Which stakeholder problem does the synthetic representative solve, and why would a neutral or clearly non-human interface not work as well?
2. Which messages, audiences and situations are permitted, and which remain outside its scope?
3. What may the representative repeat, explain or interpret, and when must it stop and refer the matter to a person?
4. Who approves each capability, how can audiences verify the communication, and how will the organization correct, suspend or withdraw the system if necessary?

These questions prevent the discussion from beginning with the technology. The first consideration should be whether the proposed use creates sufficient value for stakeholders to justify the additional complexity and risk.

Begin with a limited pilot

The first use should be scheduled, low stakes and based on an approved message. It should address a specific and measurable need, such as accessibility, translation or consistent delivery across markets.

The synthetic origin should be disclosed clearly. The organization should use a verified channel, retain the approval record and provide an accessible route to a person. The pilot should measure comprehension, usefulness, perceived risk and trust.

A successful technical demonstration is not enough. The organization must also establish that the use is appropriate, understood by its audience and supported by effective controls.

The Synthetic CEO Protocol

The Synthetic CEO Protocol turns these requirements into a deployment standard. Approval depends on evidence across eight areas, from purpose and consent to escalation and exit.

Control	The evidence it should leave behind
Purpose	A specific stakeholder need and measurable benefit
Consent	Documented rights for likeness, voice, scope, duration and exit
Authority	An explicit list of permitted and prohibited behaviours
Disclosure	A clear, accessible notice before exposure or interaction
Approval	A named human accountable for content and capability
Provenance	Verifiable origin, channel and available Content Credentials
Escalation	A tested route to a qualified human
Incident and exit	Revocation, correction and executive-departure procedures

The CCO should maintain the decision record and confirm that every control has an accountable owner. Approval should be reconsidered whenever the system's scope or capability changes.

A representative that gains memory, live interaction, personalization or access to internal systems is not simply an improved version of the original tool. These additions change what it can know, infer and do. They therefore require a new decision.

Preserve the role of real presence

Synthetic communication can make executive messages easier to produce, translate and distribute. It can extend the reach of leadership communication, but it cannot replace every function of leadership presence.

A real executive can listen to an unexpected question, acknowledge uncertainty, change course and make a commitment. These are not production features. They are expressions of judgement and responsibility.

Organizations should therefore use AI to remove unnecessary friction from communication while preserving the moments in which human involvement matters. The aim is not to maximize the use of synthetic presence. It is to use it selectively, where it creates value without weakening accountability.

The organizations that manage this well will distinguish between communication that can be scaled and responsibility that must remain personal. Scale information synthetically where it helps. Keep judgement, vulnerability and responsibility human.

APPENDIX: RESEARCH METHOD AND LIMITATIONS

The field study was conducted by Damian Phan in collaboration with the Scriptorium Initiative as part of Communication Science ’26 at the University of Amsterdam. Respondents were randomly exposed to one of two short CEO communication videos: a human-delivered condition and an AI/synthetic condition. They then evaluated adoption willingness, usefulness, perceived risk, leadership authenticity, ease of use, contextual suitability and governance conditions. AI self-efficacy and professional background variables were also assessed.

The study included 78 valid respondents. It provides unusually relevant evidence from experienced communication decision-makers, but the sample is modest and predominantly senior, international and Western. Subgroup tests may have limited statistical power. Findings about associations and predictors should not be interpreted as proof of causality beyond the experimental comparison between the two stimulus conditions.

The results are therefore best used as directional evidence for governance design and future investigation. Further research should include larger, preregistered samples and examine whether responses vary by region, gender, generation and cultural context. It should test multiple CEOs and message contexts, repeated exposure and interactive rather than video-only representatives.

Future work should also move beyond stated willingness. Self-reported attitudes are vulnerable to recall, social-desirability and hypothetical-response bias. Behavioural measures, whether participants verify a message, challenge an answer, escalate to a human or actually choose the synthetic option, would provide stronger evidence. Where appropriate, experimental studies could also combine self-report with attention and physiological measures such as eye-tracking or EEG to examine responses participants may not be able to articulate directly.

SELECTED REFERENCES

European Union. Regulation (EU) 2024/1689, Article 50.
https://eur-lex.europa.eu/eli/reg/2024/1689/

European Commission. Transparency obligations under Article 50 of the AI Act, 2026.
https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act

Chandra, B. et al. Reducing Risks Posed by Synthetic Content. NIST AI 100-4, 2024; updated 2026.
https://www.nist.gov/publications/reducing-risks-posed-synthetic-content-overview-technical-approaches-digital-content

Barrington, S., Cooper, E. A. & Farid, H. "People are poorly equipped to detect AI-powered voice clones." Scientific Reports, 2025.
https://www.nature.com/articles/s41598-025-94170-3

Coalition for Content Provenance and Authenticity. C2PA Technical Specification 2.2, 2025.
https://spec.c2pa.org/specifications/specifications/2.2/specs/C2PA_Specification

Partnership on AI. Responsible Practices for Synthetic Media.
https://syntheticmedia.partnershiponai.org/

Synthesia. Enterprise video translation and multilingual video player documentation, 2026.
https://help.synthesia.io/en/articles/8423908-how-do-i-translate-a-video

HeyGen. Digital Twin, Avatar V and LiveAvatar product documentation, 2026.
https://help.heygen.com/en/articles/14602974-avatar-v-is-now-available-on-heygen

Tavus. Video Generation and Conversational Video Interface documentation, 2026.
https://docs.tavus.io/sections/conversational-video-interface/overview-cvi

ElevenLabs. Voice cloning concepts and Professional Voice Cloning documentation, 2026.
https://elevenlabs.io/docs/eleven-api/concepts/voice-cloning

Scriptorium Initiative / Damian Phan. Would You Trust a Synthetic CEO? Unpublished field research presentation, 2026.

TechCrunch. Reporting on Klarna and Zoom CEO avatars in earnings-related communication, 2025.
https://techcrunch.com/2025/05/21/klarna-used-an-ai-avatar-of-its-ceo-to-deliver-earnings-it-said/

The Guardian. Reporting on Meta's proposed AI version of Mark Zuckerberg for employee communication, 2026.
https://www.theguardian.com/technology/2026/apr/13/meta-ai-mark-zuckerberg-staff-talk-to-the-boss


MAKING SENSE
OF THE [_UNKNOWN]