Universal Library of Business and Economics Research Article

ISSN: 3064-9951 | Volume 3, Issue 4

Open Access | PP: 41-46

DOI: https://doi.org/10.70315/uloap.ulbec.2026.0304005

A Generative Artificial Intelligence Pipeline for Large-Scale Product Content Updates on Online Marketplaces with Human Quality Control

Aleksei Konoplev

Founder and CEO, Copterdrone, Saint Petersburg, Russia.
ORCID: 0009-0004-4463-394X

Abstract

Retailers that sell through several online marketplaces accumulate catalogs in which thousands of items stay on sale with outdated images, thin descriptions, and incomplete attribute tables. Manual repair of such catalogs can stretch over years of content work and result in uneven quality, since content managers change and each brings personal habits. This review examines how generative artificial intelligence can renew product content at catalog scale when generative models work as one component of an end-to-end pipeline with checkpoints planned in advance and a mandatory visual review by a person before publication. The study applies the antecedents, decisions, and outcomes framework to academic literature published from 2021 to 2026 and to first-hand experience of designing and running such a pipeline at a Russian online retailer of radio-controlled models that sells through Ozon, Wildberries, Yandex Market, and Avito. The case is therefore described from inside the company, with design reasons drawn from direct practice. In the case pipeline, a person prepares links to original supplier or manufacturer content. The system then collects reference images, expands the description, fills the attribute template according to each marketplace's rules, writes prompts, and generates a standard set of five card images with infographics for each item. After that, it matches the new cards to internal stock-keeping units, assembles an HTML preview, and, once a person approves the preview, uploads the content through marketplace interfaces. The cost of a card with five images fell from 1,500–2,500 to 80 rubles, which is 19 to 31 times lower, preparation time went from one to one and a half weeks to ten minutes, and 2,500 of the 5,000 catalog items were updated within one and a half months. Spending on content shifted from wages to a model subscription and API tokens, which together cost about 1,100 US dollars per month. The review names three conditions for transfer to another catalog, namely access to original manufacturer content, an internal article numbering system, and marketplace application programming interfaces.

Keywords: *Content Automation, E-Commerce, Generative Artificial Intelligence, Human-in-the-Loop, Online Marketplaces.*

INTRODUCTION

A product card on a marketplace is a small production unit. It combines a set of images, a description, and an attribute table, and each platform sets its own rules for all three parts. Buyers judge an item through this card before they can hold it, and empirical work links listing content to demand. Seller-created multimedia richness shows a significant association with product sales on Amazon (Cai et al., 2023), and interpretable image features account for differences in demand among comparable Airbnb listings (Zhang et al., 2022).

The case behind this review comes from first-hand experience. The pipeline was designed, built, and run inside the company that the review describes, and this direct position gives the review access to the reasons behind each design decision and to operating records that an outside analysis would lack. The case shows the scale of catalog debt, the stock of cards that keep selling long after their images and text have aged. A Russian online retailer of radio-controlled models, which has operated in e-commerce since 2014 and runs stores and warehouses in Moscow and Saint Petersburg, had 5,000 items already on sale with outdated designs, weak descriptions, unfinished attribute tables, and missing hashtags. The company buys from Russian suppliers and imports from Chinese factories, and it sells through Ozon, Wildberries, Yandex Market, and Avito, so each item needs a card adapted to the rules of several platforms. The company estimated that content managers would need about two years to update the catalog by hand.

The reviewed research on generative artificial intelligence examines tools for this task one by one. Marketing research maps the uses of generative models for personalization,

Citation: Aleksei Konoplev, "A Generative Artificial Intelligence Pipeline for Large-Scale Product Content Updates on Online Marketplaces with Human Quality Control", Universal Library of Business and Economics, 2026; 3(4): 41-46. DOI: https://doi.org/10.70315/uloap.ulbec.2026.0304005.

content creation, and sales lead generation, links content creation with higher efficiency, and sets out a research agenda for the field (Kshetri et al., 2024). Work on product data shows that large language models can extract attribute values from unstructured product text and normalize them (Brinkmann et al., 2024). Experiments with professional writing tasks report shorter completion times and higher output quality with a generative assistant (Noy & Zhang, 2023), and field data from customer support show productivity gains that are largest for less-experienced workers (Brynjolfsson et al., 2025). In experimental and field settings, the model assists a writer or support agent, who remains responsible for the result.

The same literature describes the risks that grow once generated content reaches a storefront. Language models produce statements that are unsupported by their sources (Ji et al., 2023), and in a marketplace, such statements can become incorrect specifications on a product card. In hospitality services, generated images lower the intention to use and recommend a service compared with real photographs, and the effect is stronger for hedonic services and for customers with high involvement (Belanche et al., 2025). The finding comes from services, yet radio-controlled models are a high-involvement hobby purchase, which makes the faithfulness of a generated card to the physical product a commercial requirement in this case. Human review is the usual safeguard, yet people tend to over-rely on machine output, and a review of 35 experimental studies identifies active user verification as the primary lever against this automation bias (Romeo & Conti, 2026). Research on human-in-the-loop machine learning classifies the points at which people can enter a learning cycle and the roles they can play there (Mosqueira-Rey et al., 2023).

No description of an end-to-end pipeline that starts from a link to original manufacturer content, ends with upload through marketplace interfaces, places checkpoints along the way, keeps a person at a single review stage, and reports the unit economics of a product card was found in the reviewed literature. The review addresses three questions. The first asks what drives a seller away from manual content production. The second asks which design decisions keep quality under control when generation runs at catalog scale. The third asks what such a pipeline delivers and under which conditions another catalog can reproduce it.

MATERIALS AND METHODS

Review Design

The review follows the antecedents, decisions, and outcomes framework, which belongs to the family of framework-based reviews that include explicit procedures for assembling, organizing, and assessing sources (Paul et al., 2021). The framework suits a practical question that unfolds as a chain of conditions, choices, and results. Antecedents cover the pressures that made manual content production unsustainable in the case company. Decisions cover the pipeline stages, the checkpoints within it, and the risks identified in the literature that these checkpoints address. Outcomes cover the unit economics of a card, the organizational effects, and the commercial indicators recorded after the update. The comparison with alternative approaches and the conditions for transfer draw on all three categories together.

Literature Sources

The literature base consists of journal publications and conference papers published from 2021 to 2026 that have digital object identifiers. All of them are peer-reviewed research papers except Kshetri et al. (2024), which the publisher classifies as an editorial and which serves the review as an agenda-setting overview. Sources were selected along four streams, which are the effect of product content on demand on digital platforms, applications of generative models in marketing and product data management, the reliability and human oversight of model output, and the productivity effects of generative tools. A source entered the review when it informed at least one pipeline stage or one outcome indicator. Keyword searches of publisher, repository, and indexing records identified the candidates, Crossref metadata confirmed their bibliographic details and identifiers, and the final search was completed on 5 November 2026. Ten substantive sources met this criterion. Antecedents draw on Brynjolfsson et al. (2025) and Cai et al. (2023), decisions draw on Belanche et al. (2025), Brinkmann et al. (2024), Ji et al. (2023), Kshetri et al. (2024), Mosqueira-Rey et al. (2023), and Romeo and Conti (2026), and outcomes draw on Cai et al. (2023), Noy and Zhang (2023), and Zhang et al. (2022). One further methodological source defines the review procedure.

Case Setting

The case company is a Russian online retailer of radio-controlled models for children and adults, for amateurs and for professionals. It has operated since 2014, runs two stores with warehouses in Moscow and Saint Petersburg, serves customers across Russia and in part of the Commonwealth of Independent States, and sells through Ozon, Wildberries, Yandex Market, and Avito. The company relies on in-house software, which includes its own customer relationship management system for orders, a content management system for the web store, and a partial enterprise resource planning module for the warehouse with built-in analytics. The review rests on direct involvement in the design of these systems and of the content pipeline, which gives the analysis access to the reasons behind each design choice.

At the start of the update, 5,000 catalog items needed new content, and 2,500 of them passed through the pipeline within 1.5 months. The main generation tool is ChatGPT by OpenAI, powered by the GPT-6 Astra model, accessed via a subscription of about 100 US dollars per month and paid

API tokens of about 1,000 US dollars per month. Claude, DeepSeek, Grok, and Gemini serve as additional checks.

Data and Constraints

All case figures come from the company's own operating records. Raw data exports were not available for the review, and the update did not run as a controlled before-and-after study, so the figures describe operating practice. The changes in click-through rate, conversion, and orders are observations recorded after the update, and the review does not treat them as pipeline effects because season, advertising, assortment, and marketplace algorithms also shape these indicators. Preparation times appear as recorded values, one to one and a half weeks for a manual card and ten minutes for a pipeline card, without a speed multiplier, since the records do not state whether they measure working time or the interval from task to finished card. The records also leave unspecified what the additional models check. The literature section of the study is a narrative review of a focused set of sources, so its statements about gaps refer to the reviewed literature.

RESULTS

Antecedents of the Shift

The first antecedent is the price of a manual card. A card with five images cost the company 1,500 to 2,500 rubles and took a remote content manager 1 to 1.5 weeks. On this basis, the company set a manual update of the catalog at about two years, and the cost grows with every marketplace that demands its own card format.

The second antecedent is the instability of manual work. The content manager position saw constant turnover, and each newcomer brought personal habits in layout, wording, and attribute filling, so cards made in different months looked like products of different shops. This argument carries as much weight as price. Field evidence shows that generative assistance narrows the performance gap between experienced and novice workers (Brynjolfsson et al., 2025), and a pipeline with fixed rules carries this effect from a team of assisted workers to the output itself.

The third antecedent lies in the platforms' requirements. Each marketplace defines its own attribute set and image rules, and the case company had to meet them separately for every platform. Trimming the visual part of the card to save money carries a commercial risk because seller-created multimedia richness is the one seller factor with a significant association with sales in large-sample Amazon data (Cai et al., 2023).

Pipeline Design Decisions

The central decision of the case is to treat generative models as one component of a production chain. Fig. 1 shows this chain as eleven consecutive steps grouped into three phases, with two human steps and an additional cross-model check.

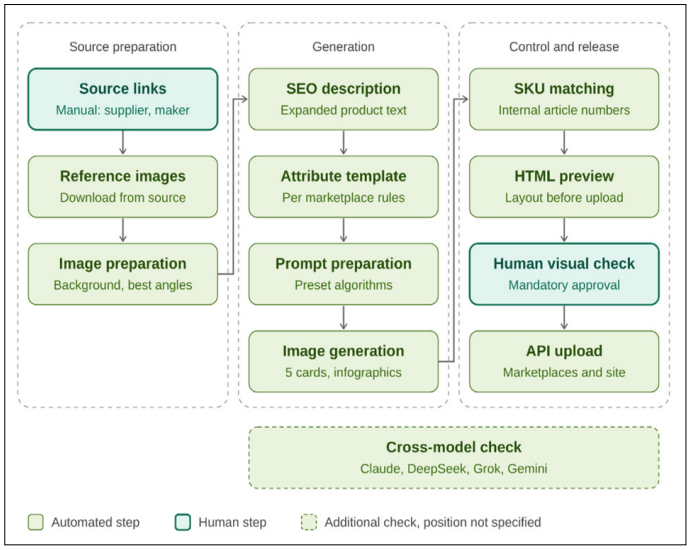


Fig. 1. Product content pipeline with human steps and checkpoints

Source preparation opens with the only manual step before generation. A person collects links to original content on the supplier or manufacturer's site, and from there the system downloads reference images, removes the background as needed, and selects the most attractive angles. Generation starts from text. The system expands the product description into full text for marketplace upload, with search optimization in mind, and populates the attribute table according to the rules of each marketplace and the company web store. The attribute table comes before image generation because the infographics on the cards display product properties, and these properties need to be fixed and verified first. Large language models can extract and normalize attribute values from product text (Brinkmann et al., 2024), which shows that this step lends itself to automation. Text expansion and visual card production are among the content-creation uses of generative models identified in marketing research (Kshetri et al., 2024). After that, the system generates prompts using preset algorithms, connects to image-generation providers via their APIs, and produces a standard set of five bright cards per item, each with infographics and other graphic elements, along with rich content for the listing.

Preset prompt algorithms encapsulate the design knowledge that a content manager once kept in personal practice. Composition, the placement of infographics, and the choice of product properties for the cover become rules applied to every item, so the visual standard of the catalog lives in the algorithm and survives any staff changes.

Control and release begin with matching. The new cards receive article numbers from the internal system, which ties every generated stock-keeping unit to a real stock position. The system then assembles an HTML preview of the listing, a person checks it visually, and approved content is sent to the marketplaces and the web store through their interfaces.

Most checkpoints answer a risk named in the literature.

Reliance on original manufacturer content narrows the space for unsupported statements, a failure mode documented across natural language generation tasks (Ji et al., 2023). The attribute template under marketplace rules converts a free-generation task into a fill-in task with a fixed structure and is designed to prevent missing or misplaced fields. Matching with internal article numbers keeps a generated card from attaching to the wrong stock position. The visual review closes the gap between a generated image and the physical product, which matters because, in services, generated images lowered the intention to use and recommend the offer, and customers described such images as less credible and as a potential source of deception (Belanche et al., 2025). Every card in the case starts with real reference images of the product, and the visual review is the point at which the generated card meets these references. Cross-model checking with Claude, DeepSeek, Grok, and Gemini adds a further layer beyond the main model.

Human-in-the-loop research on machine learning systems treats the point at which a person enters a process and the control that person holds at that point as design variables (Mosqueira-Rey et al., 2023). The pipeline involves no model training and applies the same logic to content production, with a person at the entry, where the link to source content is chosen, and at the exit, where the visual review releases the card. The late position of the review has a practical rationale. The preview gathers the text, attributes, and images of a single item on a single screen, so the reviewer can judge a finished card and have a concrete object to compare with the product. Such a setting supports the active verification that a review of experimental studies proposes as the main lever against automation bias (Romeo & Conti, 2026). The in-house order, content, and warehouse systems enable the closing steps, since matching and uploading depend on direct access to the company's own article numbers.

Outcomes

The outcomes of the pipeline are clearest at the level of a single card, a unit that any marketplace seller can set against their own figures. Fig. 2 compares the recorded cost and preparation time of a manual card with those of a pipeline card.

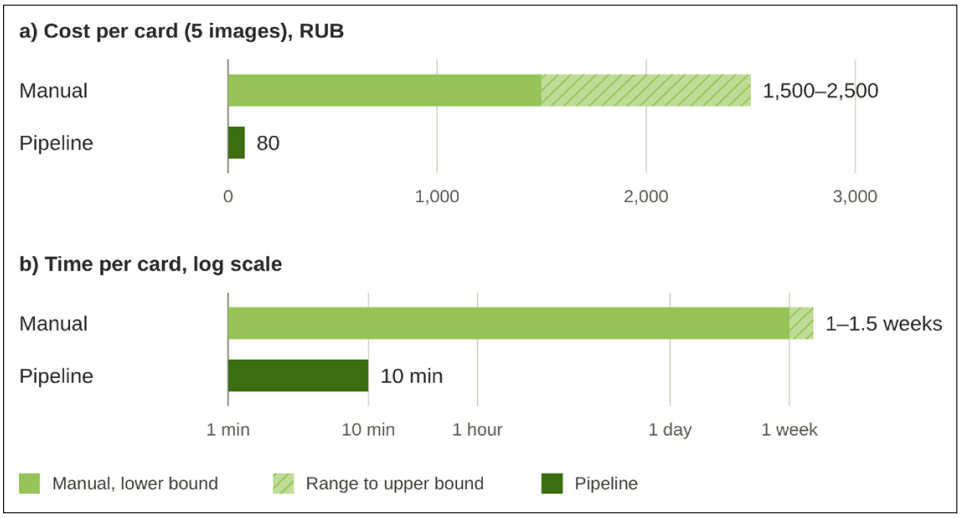


Fig. 2. Cost and preparation time of one product card with five images

The cost of a card fell to 80 rubles, which is 19 to 31 times lower than the manual range, and preparation time dropped to 10 minutes, compared with 1 to 1.5 weeks for a manual card. Over 1.5 months, the pipeline updated 2,500 items. The cost structure changed as well. Wages for a staff of content managers gave way to a model subscription and API tokens with a combined monthly bill of about 1,100 US dollars, and the two human steps remained with company employees.

Organizational outcomes followed from the same shift. The company no longer needed a staff of content managers, and the load on marketplace managers decreased. Every card now passes the same sequence of rules, which removes the dependence of layout, wording, and attribute filling on the individual executor.

After the update to 2,500 items, the company recorded a doubling of the click-through rate, an increase of about 50 percent in conversions, and a threefold increase in monthly orders. The direction of these observations aligns with evidence from image features (Zhang et al., 2022) and the multimedia richness of listings (Cai et al., 2023).

DISCUSSION

Comparison with Alternative Approaches

Table 1 sets the pipeline against other ways of preparing product content and compares manual work by a content manager, template automation, generation without review, and the pipeline with review in terms of cost, time, visual content, attribute completeness, quality stability, main risk, and the role of a person.

Table 1. Comparison of approaches to product content preparation

Criterion	Manual content manager	Template automation	Generation without review	Pipeline with review
Cost per card, 5 images	1,500–2,500 RUB	Low, no new images created	Low, tokens only	80 RUB
Time per card	1–1.5 weeks	Minutes	Minutes	10 minutes
Visual content	Made by hand	Reuses existing assets	Generated, not checked	Generated, checked by a person
Attribute completeness	Depends on the executor	Limited to source fields	Prone to errors and fabricated values	Template per marketplace rules
Quality stability	Low, executors change	High but rigid	Unpredictable	Same rules for every card
Main risk	Cost, delays, staff turnover	Generic content	Errors reach the marketplace, card rejection	Dependence on source content and APIs
Role of a person	Does all the work	Sets up templates	None	Prepares links, approves the preview

Note. Cost and time values for manual work and the pipeline come from the operating records of the case company. Other cells give analytical assessments based on the case and the reviewed literature.

Template automation keeps quality stable, yet it reuses existing assets and cannot produce new visual content, so the visual debt of a catalog stays where it was. Generation without review comes close to the pipeline in speed and price per card and sends unsupported attributes to the storefront (Ji et al., 2023), along with images that no one has compared with the product. Manual work keeps a person responsible for each card and comes at the cost of time, money, and exposure to staff turnover. The pipeline with review maintains generation speed and moves the review to the point where a single look covers the whole card. Experimental gains from generative assistance in writing tasks (Noy & Zhang, 2023) show what the model adds at the level of an individual task, and the case shows how this gain scales once the surrounding steps are automated as well.

The comparison also reveals a change in human work. A person in the pipeline makes two decisions: which source to trust and whether the finished card matches the product, and both decisions require knowledge of the assortment.

Conditions for Transfer

Another catalog can reproduce the pipeline when three conditions hold. The first condition is access to the original content from suppliers and manufacturers, since the text, attributes, and images all ground themselves in this material. The second condition is an internal article numbering system that can receive generated stock-keeping units and tie them to stock positions. The third condition is the availability of marketplace APIs for automated upload and update. Without these three elements, the pipeline cannot be assembled in the described form. Catalogs of unbranded goods with poor manufacturer content would need another source of reference material, and sellers who depend on third-party software for orders and stock would need an integration layer before the closing steps can run.

Limitations and Future Research

The case figures come from the company's own records, raw exports were not available, and no controlled before-and-after design was used, which keeps the review within the bounds of operating evidence. The literature part covers a focused set of sources, and its gap statements apply to that set. The recorded growth in click-through rate, conversion, and orders could also be explained by seasonality, advertising, assortment, and marketplace algorithms. Preparation times are given as recorded values without a multiplier, because the records do not define what they measure. The task of the additional models in the cross-model check is not specified. Future research could run controlled tests of generated and manual cards on the same marketplace, measure the share of cards that human reviewers reject, along with the reasons for rejection, track card quality over time, and examine how disclosure rules for generated images affect such pipelines.

CONCLUSION

Generative artificial intelligence can renew a marketplace catalog of thousands of items when embedded in the full production chain. In the reviewed case, the chain runs from a manually prepared link to original manufacturer content up to upload through marketplace interfaces, with generative models as one component.

The shift away from manual content production follows from three antecedents: the price and duration of a manual card, uneven quality due to constant staff turnover, and the distinct requirements of each marketplace. Quality under generation at scale rests on checkpoints planned in advance, including reliance on original manufacturer content, attribute templates per each platform's rules, matching to internal article numbers, and a mandatory visual review of an assembled preview by a person, with a cross-model check as an additional layer. The outcomes are a card cost of 80 rubles against 1,500–2,500 rubles, a preparation time of ten minutes against one to one and a half weeks, 2,500 updated items within one and a half months, and cards produced under the same rules regardless of who launches the process.

For sellers with large catalogs, the case offers a reproducible design that depends on access to original manufacturer content, an internal article numbering system, and marketplace APIs. Where these conditions hold, a person remains responsible for everything that reaches the storefront, and the size of the staff limits the review step alone.

REFERENCES

1. Belanche, D., Ibáñez-Sánchez, S., Jordán, P., & Matas, S. (2025). Customer reactions to generative AI vs. real images in high-involvement and hedonic services. *International Journal of Information Management, 85*, Article 102954. https://doi.org/10.1016/j.ijinfomgt.2025.102954
2. Brinkmann, A., Baumann, N., & Bizer, C. (2024). Using LLMs for the extraction and normalization of product attribute values. In *Advances in Databases and Information Systems* (Lecture Notes in Computer Science, Vol. 14918, pp. 217–230). Springer. https://doi.org/10.1007/978-3-031-70626-4_15
3. Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics, 140*(2), 889–942. https://doi.org/10.1093/qje/qjae044
4. Cai, X., Cebollada, J., & Cortiñas, M. (2023). Impact of seller- and buyer-created content on product sales in the electronic commerce platform: The role of informativeness, readability, multimedia richness, and extreme valence. *Journal of Retailing and Consumer Services, 70*, Article 103141. https://doi.org/10.1016/j.jretconser.2022.103141
5. Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys, 55*(12), Article 248. https://doi.org/10.1145/3571730
6. Kshetri, N., Dwivedi, Y. K., Davenport, T. H., & Panteli, N. (2024). Generative artificial intelligence in marketing: Applications, opportunities, challenges, and research agenda. *International Journal of Information Management, 75*, Article 102716. https://doi.org/10.1016/j.ijinfomgt.2023.102716
7. Mosqueira-Rey, E., Hernández-Pereira, E., Alonso-Ríos, D., Bobes-Bascarán, J., & Fernández-Leal, Á. (2023). Human-in-the-loop machine learning: A state of the art. *Artificial Intelligence Review, 56*(4), 3005–3054. https://doi.org/10.1007/s10462-022-10246-w
8. Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science, 381*(6654), 187–192. https://doi.org/10.1126/science.adh2586
9. Paul, J., Lim, W. M., O'Cass, A., Hao, A. W., & Bresciani, S. (2021). Scientific procedures and rationales for systematic literature reviews (SPAR-4-SLR). *International Journal of Consumer Studies, 45*(4), O1–O16. https://doi.org/10.1111/ijcs.12695
10. Romeo, G., & Conti, D. (2026). Exploring automation bias in human–AI collaboration: A review and implications for explainable AI. *AI & Society, 41*, 259–278. https://doi.org/10.1007/s00146-025-02422-7
11. Zhang, S., Lee, D., Singh, P. V., & Srinivasan, K. (2022). What makes a good image? Airbnb demand analytics leveraging interpretable image features. *Management Science, 68*(8), 5644–5666. https://doi.org/10.1287/mnsc.2021.4175

Copyright: © 2026 The Author(s). This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.