



Extending Knowledge Graphs with Subjective Influence Networks for personalized fashion

Kurt Bollacker, Natalia Díaz-Rodríguez, Xian Li

► To cite this version:

Kurt Bollacker, Natalia Díaz-Rodríguez, Xian Li. Extending Knowledge Graphs with Subjective Influence Networks for personalized fashion. Designing Cognitive Cities, 2018. hal-01952205

HAL Id: hal-01952205

<https://hal.science/hal-01952205>

Submitted on 12 Dec 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Extending Knowledge Graphs with Subjective Influence Networks for personalized fashion

Kurt Bollacker, Natalia Díaz-Rodríguez, Xian Li

Stitch Fix Inc, San Francisco, California, USA
{kbollacker,xli}@stitchfix.com, ndiaz@decsai.ugr.es,
<http://multithreaded.stitchfix.com/>, <http://www.stitchfix.com>

Abstract. This chapter shows Stitch Fix’s industry case as an applied *fashion* application in cognitive cities. Fashion goes hand in hand with the economic development of better methods in smart and cognitive cities, leisure activities and consumption. However, extracting knowledge and actionable insights from fashion data still presents challenges due to the intrinsic subjectivity needed to effectively model the domain. Fashion ontologies help address this, but most existing such ontologies are “clothing” ontologies, which consider only the physical attributes of garments or people and often model subjective judgements only as opaque categorizations of entities. We address this by proposing a supplementary ontological approach in the fashion domain based on subjective influence networks. We enumerate a set of use cases this approach is intended to address and discuss possible classes of prediction questions and machine learning experiments that could be executed to validate or refute the model. We also present a case study on business models and monetization strategies for digital fashion, a domain that is fast-changing and gaining the battle in the digital domain.

Key words: Ontology, Folksonomy, Knowledge Graph, Fashion, Subjectivity, Temporal Networks, Social Networks, Influence, Natural Language Processing, Recommendation Systems, Personalization, Business Development

1 Introduction

1.1 Cognitive cities and the Fashion domain

Cognitive computing aims at improving the quality of life in cities, especially aiding in decision-making, handling linguistic information - which is usually imprecise and developing applications towards achieving Smart Cities [8].

This chapter shows a use case beyond the cognitive cities, through a case on cognitive-cultural capitalism emanating from the new urbanism [18]. Because creativity is a concept whose time has come in economic and urban geography, and because much existing research on creative cities fails adequately to recognize that interdependent processes of learning, creativity and innovation are situated within concrete fields of social relationships [18], we show a use case where data



Fig. 1: Stitch Fix service

science and cognitive computing are applied to a successful and real business case. In spite of such modern urban mechanisms, which *tend to offer a flawed representation of urban dynamics* and lead sometimes to *essentially regressive policy advocacies*, cognitive-cultural capitalism is a robust theoretical framework through which contemporary urbanization processes can be described, as it is well motivated in [18]. This framework, and our use case concretely, show to have larger impacts on urban outcomes leading to higher degree of automation and job transformations.

Because of the AI transformation we are living in, we demonstrate how it is important to combine analytics, NLP, Web knowledge, fuzzy (cognitive maps), etc. with the fashion industry to develop new experiences in smart/cognitive cities that adapt to the citizens' new lifestyle.

1.2 Stitch Fix's case

Stitch fix (SF) is a personalization company that delivers fashion to your door and it is purely driven by data science (Figure 1). Stitch Fix is listed as number 2 in the top 15 companies to watch in 2017¹, and has been popular in the US, now running for 5 years, since CEO Katrina Lake tested the viability of the idea in her apartment in Massachusetts. The company counts with a peculiar

¹ Inc: Top 15 companies to watch in 2017 <http://www.inc.com/guadalupe-gonzalez/ss/top-companies-to-watch-2017.html>

title of Chief Algorithms Officer Eric Colson (and a subset of the team) that comes from the heavy recommendation system-based company Netflix. Stitch Fix is thus sometimes known as the "Netflix of fashion", "Pandora for clothing", or the "clothing company that blends AI and Human Expertise"². In the San Francisco Bay Area, Stitch Fix is popular as well for having a thorough curated blog on how they tackle different problems with data science (see Multithreaded series blog³).

Competitive asset: Why is Stitch Different to the rest of online fashion retailers? Big fashion retailers in US such as Macy's, Nordstrom and Gap continue to lay off employees and shutter stores in the face of increased competition from online players. Stitch Fix, is not characterized for having a fast delivery of fashion at home, nor having very popular nor high price point luxury brands; they have more exclusive vendors, as well as their own unique brands, to provide a more personalized experience to the client's profile, budget, requests, and lifestyle. As Colson says, "We're not going to be better priced, or faster shipping (their delivery is not supposed to be really fast), or a better brand, So we have to be really good at relevancy."

At Stitch Fix there is no online catalog, and the customer does not choose what is getting in a box. The aim is to have effortless shopping taken care of by a personal stylists that hand-picks a look for you. You can try in the comfort of home, keep what you like and return what you do not, for free, and giving feedback on how the stylist did for better improvement in next delivery. The personalization and styling experience is supported with social networks information such as the information that users add to their Pinterest boards, in order to give the stylist a visual glimpse of liked garments.

SF is not distinguished by high prices, its main products are within an affordable budget versus other services such as *Trunk Club*⁴, which targets larger budget customers. SF focuses on partnership, and cross collaboration, "together we are better" is a motto, and customer experience is the priority: if a feature does not help the client, they do not go for it. Stylists are recognized, and can work flexibly part time from home. The objective is enhancing the quality of their products and get standard procedures in place.

2 Beyond Clothing Ontologies: Modeling Fashion with Subjective Influence Networks

If Stitch Fix's personal stylists were to search for the perfect outfit for one of our customers who is asking advice for an outfit for his first Burning Man attendant or for a 50 year wedding anniversary party, it would be ideal if we could search

² Harvard Business Review <https://hbr.org/2016/11/how-one-clothing-company-blends-ai-and-human-expertise>

³ Multithreaded series blog <http://multithreaded.stitchfix.com/>

⁴ <https://www.trunkclub.com>



Fig. 2: Stitch Fix

in our database for styles such as rockabilly, retro, boho or “best pieces from the 70’s”. However, despite having ample of physical attribute annotations from different vendors, as well as algorithmically generated garments, we still do not have a sufficient purchase on the abstract attributes of a consumer’s perception, such as the aforementioned categories.

We address this problem by proposing a supplementary ontological approach in the fashion domain based on subjective influence networks. They measure novelty, impact and represent influence mechanisms that can validate or refute hypotheses that contain subjective or aesthetic components. The components of an influence vector instance consist of the elapsed time between the beginning of styles, the magnitude of the influence, the mechanism of influence, and the agent of influence.

We enumerate a set of use cases this approach is intended to address and discuss possible classes of prediction questions, hypothesis testing and machine learning experiments that could be executed to validate or refute the model. For instance, using network mapping, we could find answers to: *Can influences in other cultural domains such as music be used to predict fashion influences?* Other application is quantifying subjective attributes, for instance, subjective influence networks could characterise the differences between retro and classic glasses and their mechanisms of influence.

Example hypotheses to confirm/reject could be: *Are retro glasses those worn by 2 generations back? Are classic glasses those that never go out of fashion?* Other example of use case is predicting fashion cyclicalness, and so, we could postulate: Can influence networks tell apart fashion cycles periodicity to predict when bell trousers will be fashionable again? Representing fashion evolution on social media (e.g. hashtags on *Instagram*) as influence mechanism is an example to measure influence within time and space scales to evaluate the “viral” nature of rapid style changes. Through the use of subjective influence networks, we plan to augment the Knowledge Graph of fashion information, a search engine that allows advanced faceted search to infer features even if they are not explicitly tagged in the merchandise. The final aim is understanding and translating into machine-consumable manners the way humans perceive and transmit aesthetics

and style in fashion, and more generally, subjectivity, something that computers do not (yet) do better than us.

3 Background: Ontologies and Knowledge Graphs

As on-line fashion retail industry has been growing rapidly against traditional physical shopping, there has been a corresponding shift to a much more data-driven paradigm for business operations including manufacturing, merchandising, and marketing. In particular, future-focused data analysis has become a particularly important activity, such as predicting fashion trends, price forecasting, construction of recommender systems, and identification of consumer influencers. Often, these activities are approached using statistical, machine learning or other data-driven techniques. However, much of the data in the fashion domain comes from deep, diverse, cultural entities and phenomena. While fashion in itself is part of and can define culture, it also borrows from other cultural domains, such as music, language, film, religion, mythology, local folklore and many others. In most cultural domains, it is important to understand the narrative of history and contemporary subjective judgements and opinions. For example, in music, Italian words are used to contextualize abstract musical concepts (e.g., *allegro*, *largo*, *presto*). However the meaning of these words in the context of music has evolved and diverged from their original, common definitions. Knowing the history as well as the current interpretation of these words by the composers who use them is required to fully understand their musical meaning. Similarly, fashion is an inherently subjective, cultural notion. It is defined not by quantitative, testable measures, but by its history and the perceptions of people who care enough to form opinions about it. Therefore, in order to understand fashion in any rigorous way, this subjectivity must be an intrinsic part of the model.

One of the techniques for addressing the subjective, cultural parts of a knowledge domain is to use ontologies. Schemas, ontologies and its data population through knowledge graphs (KGs) are formal tools for expressing organized meaning and provide sense or context to a domain. More concretely, ontologies often integrate common-sense and human expert knowledge as well other external knowledge sources into machine readable computational models. Unfortunately however, most existing ontological work in fashion partially avoids subjectivity by simply focusing on “clothing ontologies” rather than fashion as a whole. Clothing ontologies primarily model the structure of physical feature values (e.g., sleeve length, colors, fabric). A particular garment can be represented in a multidimensional feature space chosen from such an ontology. Usually each garment class (e.g., top, bottom, shoe, hat) is considered to have a distinct feature space from other classes. When they do include subjective elements, clothing ontologies often do this through the inclusion of non-objective features (e.g., expected occasion, style category), but these features are usually opaque categorizations of entities, with no explicit semantics. Despite the limitations, these clothing features spaces are still useful because they provide semantic structure to data

that can be used when applying analytic/prediction techniques (e.g., similarity measures, classifiers, function estimators).

We believe that deeper, richer representations of the subjective features of fashion data is possible and would help in many important use cases. In this paper, we propose an architectural augmentation to traditional clothing ontologies that includes the notion of a subjective influence network in a way that may be able to capture subjective semantics that simple categorical features do not. We enumerate a set of potential use cases, and propose types of measurements and applications that can be carried out to measure the usefulness of our approach.

The rest of the paper is organized as follows. Section 4 exhaustively summarizes the state of the art on existing fashion ontologies and frameworks and Section 4.4 describes machine learning applications as motivating use cases for our fashion ontological modelling approach. Section 5 proposes the theoretical foundations of the subjective model of influence, entities, relations and the mechanisms to quantify influence and subjectivity. Section 6 discusses evaluation approaches and utility of the model once populated with empirical data. Section 9 concludes with further insights.

4 Related Work

4.1 Related Fashion Ontologies and Schemas

Ontologies have been used to represent knowledge in a large set of real-life problems, from genetics⁵ to decision support systems, optimization, matchmaking and human activity recognition [7]. In the fashion world, ontologies have sporadically been used for recommendation systems. For example, ontologies have been combined with fuzzy logic for personalized garment design, where fuzzy decision trees serve in learning a set of representative samples. Fuzzy cognitive maps model complex relations between sensory descriptors and fashion themes given by consumers to provide more fine grained recommendations as well as the evaluate how much a specific body shape is relevant to a desired emotional fashion theme [25].

An important existing ontology is the Garment Style Advice Ontology SERVIVE (SERVice Oriented Intelligent Value Adding nEtwork for Clothing-SMEs embarking in Mass- Customisation)⁶[22]. The Servive Fashion Ontology (SFO) includes relations among different categories of entities such as *colors*, *companies*, *garment features*, *materials*, etc. and provides a similarly structured and unified vocabulary to represent human, fashion and manufacturing concepts. The project includes the design of a Virtual Customer Advisor (VCA) which expresses preferences for a given garment that is evaluated via SWRL rules and Pellet reasoner. Fig. 3 shows the most abstract or top layer classes as well as the highest hierarchical layer of object properties modelled in SERVIVE ontology.

⁵ <http://geneontology.org/>

⁶ SERVIVE EU Project <http://www.servive.eu/>

Despite being the most complete ontology publicly available to the best of our knowledge, except for the subjective season labels (*hasHumanStyleColour*) and suitability classifications (*isForOccasion*), the ontology consists only of physical object hierarchies.

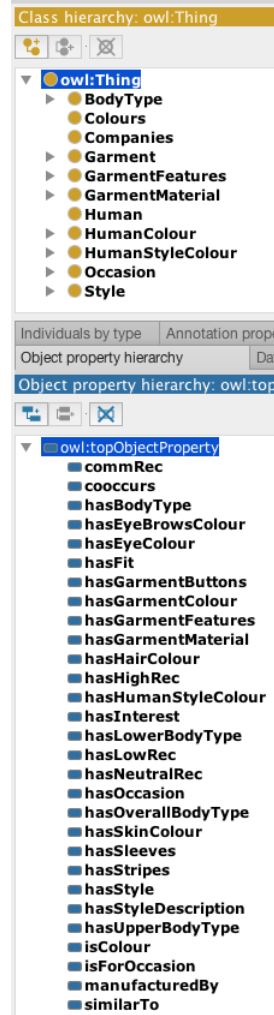


Fig. 3: SERVIVE Ontology main entity classes (above) and main object properties (below) [22]

Ontologies per se act primarily as a modelling tool, and for them to be useful, they are to be integrated into some kind of application (be it search, recommendation, classification or decision making applications). For instance, ontologies have also been integrated into probabilistic and media-rich approaches for per-

sonalized garment recommendation systems. Expert subjective knowledge from public online media is used to compute compatibility among products and user profiles according to context and probabilistic reasoning. [2] concretely focuses on dresses (*sarees*) and its evaluation of several individuals’ fashion preferences and celebrities’ actual choices compared with automated recommendations. The format of the ontology is MOWL, that *enables the analysis of visual properties of garments with respect to fashion concepts*, but it is not publicly available.

Another ontology, which considers designers, models, trends, seasons and celebrities is in [14], which exploits lexico-syntactic patterns as NLP tools for ontology learning, relation extraction and curation through domain experts. Table 1 summarizes the main ontologies’ concepts and relations modelled.

Table 1: Existing Clothing Ontologies

Ontology/Model and Language	Main Entities	Relations
SERVIVE [22], OWL	Body type, colors, companies, garments (features, material), human colour categories, seasonal human style color, occasion, style	Co-occurs, hasInterest, hasHigh/Low/NeutralRecommendation, hasBody-Type/Fit/EyeColour, hasGarmentButtons/Colour/Feature/Material/HumanStyleColour, hasOccasion, hasSleeves/Stripes/Style/styleDescription, isForOccasion, manufacturedBy, similarTo, isColour
Fashion ontology [14], RDF	Celebrity, designer, model, clothing term, trend, season	
Indian garment ontology [2], MOWL	Craft (stitch, print, embroidery), material, textile categories,	Celebrity validation
Fashion cognitive model [12]	Garment parts (silhouette, waist, length, collar, sleeve, ornaments, symmetry)	
Fashion cognitive model [25]	Body shape, desired emotional theme	Effectiveness, acceptability, realizability

Considering work that is more general than the fashion domain, open data portals such as Dbpedia and Freebase [4] contain 1K topics and 3K facts around fashion, clothing and textiles⁷. Despite the richness and structure found in these

⁷ <https://developers.google.com/freebase/>

formal base resources, the creative and subjective, contextual part of fashion is missing from these knowledge bases.

4.2 Cognitive Models for Fashion Modelling

In the literature there are non-ontological models which frame similar problems. They blend human and machine models for evaluating specific body shapes' relevance to a desired emotional fashion theme or intention to be transmitted. For instance, in [25], *effectiveness* evaluates whether recommended styles are relevant to the design objective or desired fashion theme, *acceptability* refers to whether the best recommended style is accepted by the expert, and *realizability* assesses if the proposed recommender system can be applied to the fashion [25].

An example of a cognitive model for fashion style decision making is in [12], where Genetic Algorithms enhanced with Multi-alternative Decision Field Theory (MDFT) tackle the context and choice set problem in decision making by using *psychological distance* between alternatives. The latter is based on the Euclidean distance among positions in a multi-attribute-dimensional subjective evaluation space.

4.3 Subjectivity in other domains

We identify a lack of a *subjective style* schema in the related work that goes beyond the biology or mechanics of clothing, and that expresses a more wholistic personal approach than the existing *inventory* clothing ontologies. By *inventory ontologies*, we mean those based on static attribute-based or physical feature spaces.

Other subjective and hard to describe domains such as music also benefit from having taxonomical classifications in form of ontologies. For instance, projects such as MusicBrainz⁸ collects music metadata, and the Music Ontology⁹[16] is a formal framework to deal with music-related information on the Semantic Web including editorial, cultural and acoustic information. Just like in music, a fashion ontology can integrate fashion-related data across multiple sources, or enrich search-engine results around decades, styles or influencers. Because of this, musicians might be useful allies for the fashion industry, (e.g., *thanks to their status as bohemian individuals*) and music industry might need fashion [13], e.g., to model music taste or predict fashion cliques.

Another similar natural phenomenon is language, where influence networks, among many other factors in time, model organically the evolution of its spread, its vocabulary, grammar rules, tonality, etc. In all, music, fashion and languages, influence and subjectivity are inherent to the domain and for them to fully be considered into machine learning systems, they need to be modelled quantitatively.

⁸ The Open Music Encyclopedia <https://musicbrainz.org/>

⁹ <http://www.musicontology.com>

4.4 Fashion Ontology Use Cases

In fashion, the human component of algorithm evaluation is necessary [19, 23]. Guided by this, we identify candidate applications where a fashion ontology enhanced with a better subjective data representation would likely be helpful.

1. Defining stylistic rule guides and recommendations or predicting specific trends. For instance, to answer questions on: how to be edgy and ahead of the fashion trend without being too far off, or how to predict the Oscars’ ceremony outfits?¹⁰.
2. Predicting mass production trends. For example, the problems of cost-efficient budget and resource allocation as well as market demand optimization.
3. Providing organizing structure, e.g., taxonomy or folksonomy, for fashion annotation systems that leverage crowd-sourced online data (e.g., [27, 26]).

In next section we specify our augmentation for clothing ontologies, including a description of modeling obligations needed to make it useful and examples of first order measurements of represented data.

5 Modeling Subjective Influence

5.1 Styles as Regions in a Feature Space

So how can subjectivity semantics be modeled as an influence network? Let us first consider a somewhat traditional interpretation of features of garments, based on physical properties from a clothing ontology. A particular garment g can be represented as a point in a clothing feature space G (see Figure 4). Let there be a theoretical set of all clothing styles Φ such that $\forall g \in G$ a subjective judge function $s()$ assigns a classification $s(g)$ such that $s(g) \in \Phi$. We define a distinct “style” x to be a region $S_x \subset G$ such that $\forall g \in S, s(g) = x$.

Because the S_x depends only on a single subjective function $s()$, it does not consider the fact that for any x , there may be multiple subjective functions that are contradictory. However, we believe this reflects the actual messiness of the real world.

5.2 Styles in a Network

Styles represented as a collection of points in a physical clothing feature space do little to capture the (subjective) semantics of fashion beyond the opaque categorical $s(g)$ features. To capture richer semantics, first consider a style as the human perception of the physical features of a single garment (or entire style region). Each style can then be described as a coherent aesthetic entity in the mind

¹⁰ <http://www.usatoday.com/story/life/entertainthis/2016/02/23/oscar-fashion-predicting-what-stars-wear-red-carpet/80747356/>




Fig. 4: Traditional clothing feature space

of an observer. While traditionally this style may be quantitatively described by its physical features, consider the alternative aspect of its subjective qualities shared with other cultural entities. These entities could be other clothing styles, or could be from other cultural domains external to fashion (e.g., music, sports, film, art, literature). We model this subjectivity as a network of influence.

We treat each style x as a node in an acyclic graph/network N (see Figure 5) such that there is a temporally directional edge function $e(x, y)$ that specifies the influence between nodes. Moving backward in time (y to x), an edge between styles describes the stylistic borrowing that occurs. Moving forward in time (x to y), the edge represents the influence from older to newer styles. This influence is not a single measure, but rather a collection of influences of different mechanisms. The strength of each mechanism can be represented as a single positive number. More formally:

$$\forall x, y \in \Phi, \exists e(x, y) \quad (1)$$

such that $e(x, y) = \overrightarrow{\mu_{xy}}$ where $\overrightarrow{\mu_{xy}}$ is the influence vector from x to y ,

Each element of $\overrightarrow{\mu}$ can be treated as a quad (t, i, m, a) where t is the amount of elapsed time between the influencing and influenced style, i is the intensity or strength of influence, m is the mechanism of influence, and a is the agent of influence. While t and i can both be represented as positive reals, m is a class that exists in a (likely) vast space of possible mechanisms M . Some categories of $m \in M$ might be:

- *Explicit*: The creator of a style explicitly declares previous styles that have been influential in the current creative process.




Fig. 5: Influence Network

- *Calculated*: Algorithmic or other mechanical means may estimate influence mechanism and strength based on garment features or causal cultural models.
- *Extrinsic*: The influence may be caused by cultural influences in one or more external parallel cultural influence networks (e.g music, religion, sports) For example, a musician that borrows musical style from a revered earlier musician, may also borrow fashion elements for their own public image.

There are many different types of possible agents of influence a , including:

- *Well known individual persons or small groups*: These extrinsic influencers may be fashion designers, well known artists/performers, cultural icons, or celebrities who are admired for the artistic or political talents.
- *Organizations*: Corporations whose business is in the fashion create styles and attempt to maximize the desirability of the products they sell.
- *Emergent Social Networks*: In the age of almost-instant, wide information dissemination, feedback loops of influence among highly fashion-conscious groups of people may result in rapid evolution and exposure of styles.

5.3 Modeling Obligations

The subjective influence network model simply lays a framework for building an ontology that is capable of representing some aspects of subjectivity in fashion. In order for this model to be practically useful, a full ontology would need to be constructed, including:

- Enumerating (at least some of) the members in G , Φ , N , and M .
- Characterizing a relevant set of subjective functions $s()$.
- Calculating, estimating, or assuming values for the quads (t, i, m, a) for the edges between the nodes $x \in N$.
- Consideration of cycles. For example, 70's Disco fashion has come back in multiple times in past decades. The approach described here would model this return as a new style that is heavily influenced by the original. However, explicit modeling of this dynamic would be important.

5.4 First Order Interpretations of the Network

Interpreting an existing fashion network might allow us to make useful, testable judgements, including identification of important styles properties, including:

- *Novelty*: This is the subjective notion of a style that is different from previous styles in a pleasantly surprising way. Using our influence network model, one naive first order measure of a style's novelty is that the sum of intensity of influencing styles is low; i.e. that it is influenced only weakly by the combination of all previous styles. The novelty ν of y could be defined as:

$$\nu_y = e^{-\sum_{x \neq y \in N} i_{xy}} \quad (2)$$

where N is our influence network, and i_{xy} is the intensity element of $\overrightarrow{\mu_{xy}}$. In this case, when the sum of i values is high, $\nu_y \approx 0$, and when i is zero, $\nu_y = 1$. Other, more sophisticated measures of novelty could include deeper network analysis approaches or more nuances summing of intensity based on mechanism m and or agent a . This measure of novelty requires there to be no missing nodes or links in the influence network. A more sophisticated variant that tolerates missing information and noise would likely be needed in a practical application.

- *Impact*: This is a measure of how much a particular style has influenced all other styles as a whole. A simple (and very naive) measure of impact ι of a style x on the network N could be:

$$\iota_x = \sum_{y \neq x \in N} i_{xy} \quad (3)$$

If x has little impact then $\iota_x \approx 0$ and if x is heavily influential, then ι_x would be large. There is much previous, mature work on the topic of measuring influence in networks such the concepts of centrality, node influence metrics, page rank, etc. As such is beyond the scope here, and a likely important direction for future research.

6 Evaluation Approaches

In this section, we propose quantitative evaluation strategies to assess the practical usefulness of representing knowledge in the fashion domain using the influence network model presented here. In particular, we suggest measurements on values of such representation and potential applications.

6.1 Quality Measurements

In order to assess practical values of the proposed approach, we describe a number of evaluation strategies to measure its quality. Specifically, we focus on data-driven and task-driven evaluations which have been applied to ontologies in other domains [17]. For the former, we aim to measure how well the ontology represents empirical data related to fashion. For the latter, we examine information retrieval and recommender systems which could be consumers of the ontology and data.

Domain data approximation This is a data-driven approach to quantify how well the proposed ontology approximates empirical data in the fashion domain. Since fashion is a highly non-static, subjective and high-dimensional domain, we propose a few metrics which may capture expressiveness, both in terms of topics and temporal evolution, including:

- **Categorical precision**: Count how many styles encoded in the influence network are real-world recognizable styles in empirical domain data.
- **Temporal bias**: If we repeat the above categorical measurements on datasets from different time spans, the resulting metrics might stay stationary if this property of influence network’s is time-invariant; otherwise, a network which fails to represent future datasets could indicate variable predictive power. The length of time span before the divergence is the representativeness timescale of the network, and the scope indicates its robustness.

- **Semantic similarity:** This measurement provides a distance metric between a traditional ontology augmented with an influence network and the text data in fashion domain in terms of “meanings” they express. We project both labels (class names O_c and property names O_p) in the ontology and the tokens in the text corpus D in the same vector space (e.g., using word2vec). Then we compute the overall similarity based on all labels’ distances weighted by their importance within the network:

$$\sum_{c_i \in O_c} \sum_{t \in c_i} Sim(t, D) * \theta(c_i) + \sum_{p_i \in O_p} \sum_{t \in p_i} Sim(t, D) \quad (4)$$

The importance score $\theta(c_i)$ is a normalized score $[0, 1]$ and can be defined depending on the context and usage. For example, $\theta(c_i)$ refers to how knowledgeable a network is w.r.t. class c_i , which can be approximated by the cumulative distribution function of its number of attributes u_i

$$\theta(c_i) = \sum_{u_m < u_i} P(U = u_m) \quad (5)$$

Task-specific expressiveness This is a task-driven measurements to quantify how expressible the ontology is compared to user’s mental representation in the context of a task. Here we use information retrieval in the fashion domain as an example task, and developed statistical measures of “expressiveness”.

- **Query concept recall:** Given a query stream like “natural fabric button down from banana republic”, we derive a mapping between concepts that appeared in the query and concepts encoded in the ontology. Specifically, we ask expert judges to identify important classes $\{c_i\}$ in the query. For each class being detected, we ask them to map it to the most similar label in the ontology. Then we compute the recall of concepts in the query as $\frac{\# \text{ concepts mapped to ontology}}{\# \text{ concepts detected}}$
- **Search result ranking:** We use rankings of search results for a given query as a proxy of “golden standards”. Then the Normalized Discounted Cumulative Gain (NDCG) metric can be adapted to measure the ontology’s relevancy to search quality. Specifically, the “gain” is quantified by the ontology’s recall of concepts in search results. We examine the top K returned search results. For each of them D_p at position p , we compute the ontology’s document concept recall, which is then discounted by its logarithmic rank $\log(p)$.

$$\sum_{p=1}^K \frac{Recall(O, D_p)}{\log(p)} \quad (6)$$

7 Fashion Knowledge Graph applications using Subjective Influence Networks

7.1 Web Data Markup through schema.org

There are massive data sources on the Web (including mobile applications). However, they are mostly unstructured, and there is no common vocabulary which facilitates collective curation of domain knowledge. Schema.org markup has been the major adoption for web data (about 31.3% of all pages by Dec 2015[10]) and is used by a variety of high traffic applications like search engines and news portal. Although it contains vertical specific schemas such as movies, music, medical and products, schema that can represent fashion content is absent.




Fig. 6: Integration with schema.org

In Figure 6 we illustrate that the proposed ontology framework can easily adapt to a lightweight ontology and integrated as an external extension of the core schema.org vocabulary, while also linking to other relevant common vocabularies such as the *GoodRelations* for E-commerce[11], *SIOC* for influence mechanisms on the social Web [5].

Consumed by machine learning systems The data represented in a subjective influence network as proposed here could be used for a variety of different data analysis and processing efforts, including the following types:

- **General machine learning problems:** The knowledge base represented by a populated influence network would contain instances associated with high-quality categorical types, which provide labeled data to train models for entity recognition. Also, the edges on their own in the network contain both numeric and categorical features which can be used in whole-network modeling experiments.

- **Fashion data retrieval:** The integration with `schema.org` enables community content publishers to explicitly annotate their posts with their perceived subjectivity of fashion contents, which are basic building blocks of a crowd-sourcing system. As a result, the marked up Web data in return allows for information organizers such as search engines to index rich contents and answer queries which contain both entities and subjective projections.
- **Recommender systems:** The taxonomy defined in the ontology provides a perfect complement to recommendations learnt in a bottom-up fashion. Therefore, it could be a very useful approach to deal with data sparsity situations such as cold start problems.

We now propose some concrete different use cases where the proposed work [3], can be applied.

Example 1: Virality in Fashion We may want to predict what new garments just seen on Milan Fashion week are going to become viral, to better supply our warehouses in the next 5 years (see Fig. 7).

- **Question 1:** Does fashion spread in the same way as viral music/art/cinema influence mechanisms function?
- **Question 2:** Can we use game virality mechanisms (e.g. Pokemon Go) to predict fashion virality?
- **Question 3:** Can a musical influence network be used to construct a predictive fashion influence network and find closely influential styles? Several hypothesis will have different graph instantiations in reality with different vectors of influence.
- **Hypothesis 1:** Social media mechanisms such as Facebook and Instagram are influential agents’ tools of spread.
- **Hypothesis 2:** Counterfeiting, as another explicit mechanism of influence (copy) from the Western world.
- **Experiment:** Collect social media popular mobile games (Pokemon Go), and popular fashion items (bare shoulder blouse) and see time elapsed since first players/wearers have it until the game/garment style becomes viral. If in Asian countries such as China, there are samples of the game/style in similar proportions relative to other games/garments within the same percentage rates to what it occurs in the same time span in the Western world, we can confirm our hypothesis 1. On the contrary, if there are not similar percentages of presence of “viral games” / “viral garment styles” in the Asian countries, we can contradict our initial hypothesis of Instagram being the agent. In the latter case, there must have been another agent conducting the influence (since Facebook and Instagram are banned in China), and further copyright experiments can be done to confirm hypothesis 2.

Example 2: Quantifying subjective attributes such as *retro* glasses. We want to build a quantitative model of what “retro” means, e.g., to distinguish it from “classic”, or for this term to be computable, retrievable or indexable in our fashion Knowledge Graph advanced search engine (see Fig. 8).




Fig. 7: Use case 1: **Mapping Networks**: Can influences in other cultural domains be used to predict fashion influences?

- **Question 1:** How to characterize quantitatively the attributes of what retro glasses mean?
- **Question 2:** Do retro and classic glasses belong to the same kind of persona/ style/ physical attributes?
- **Hypothesis 1:** Retro glasses are those that have been worn since mid-century, they stop being in fashion, and they come back.
- **Hypothesis 2:** Retro glasses may or not have been carried by an icon. Retro glasses can be those that have been worn by a second degree family member (i.e., grandparent generation), in which case, the influence mechanism is based on appreciation/ homage.

- **Hypothesis 3:** Items that never really go out of fashion (i.e., classic Rayban glasses)
- **Experiment:** After representing each hypothesis in our ontological framework of influence, we can measure cyclicalities between re-appearances of glasses in terms of their strength of presence within a given time window. If there is a clear time series strength based on social media celebrities wearing the different variations of “classical” or “retro” glasses, and that influence fades down in periods where they become non-fashionable, we can determine, based on a given seasonality strength parameter or visual similarity metric, which of the influence network models better fits our hypothesis.

Example 3: Fashion cyclicalities: e.g., predicting when bell trousers will be fashionable again. Bell trousers and wide trousers have been in fashion on and off all the time with different variations. Could we predict when they will be in fashion strongly again, looking back at the past? (see Fig. 9).

- **Question 1:** Is there a fixed interval which needs to happen in time for bell trousers to be out of fashion, for people to miss them again and they to come back?
- **Question 2:** Does every time that wide trousers became fashionable, they do it with a nuance or variation?
- **Question 3:** Is there a cyclic period where the exact classical bell trousers will always come back and if so, is this period predictable?
- **Hypothesis 1:** A minimum of 10 years need to happen for people to forget they already got enough of bell trousers.
- **Experiment:** Collect images from different decades where it was rare to find people wearing bell trousers and periods where most of people were wearing them, and annotate their main 1-year wide intervals when there significantly different peaks on these two wearing patterns. Execution examples:
 - Collecting as much data points (images of bell trousers wearers) as possible by recovering old black and white pictures since Google Trends only works from 2004 onwards.
 - Perform deep learning to separate silhouette representations from other trousers style features such as colour or material to gather a large dataset of bell and not bell trousers for each decade.
 - Perform a t-test to find if there is statistically significant difference in between periods where bell trousers were fashionable versus not fashionable. If there isn’t statistically significance among “bell trouser fashionable” periods, and we can repeat the experiment for min. 4 peaks of seasonality/decades we may conclude such cyclic behaviour occurs in time and be able to predict when the effect will occur next (better being able to prepare for manufacture and inventory optimization).

Broader examples integrating social media across channels, could validate hypotheses on the influence mechanism aspect (e.g., being the mechanism social media or counterfeiting). An example of ways to refute/validate the hypothesis




Fig. 8: Use case 2: **Quantifying Subjective Attributes:** Can influence networks characterize the differences between *retro* and *classic* glasses and their mechanisms of influence?




Fig. 9: Use case 3: **Predicting Fashion Cyclicity**: Can influence networks tell apart fashion cycles periodicity to predict when bell trousers will be fashionable again?

on the mechanism assumed could be by correlating search terms with social media co-occurring hashtags (see Fig. 10).

These were examples where machine learning problems could benefit of using ontology augmented computational intelligence techniques. Experiments in other domains, such as human activity recognition [6] [9], show that the coupling of data-driven approaches together with fuzzy ontologies versus the use of uniquely crisp (simple) ontologies, are successful approaches in computational modelling of real-life problems. The reason is that some domains require to handle uncertainty, vagueness, imprecision or missing data specifically accounting for context and in a rich expressive manner. Fuzzy ontologies are a good formalism candidate to model such domains, for instance in online matchmaking [15].

8 Business Development around knowledge graphs

A potential business development of tools built on top of the fashion knowledge graph, such as, e.g., style advice note writing assistant tools would be, for



Fig. 10: Use case 4: **Representing Fashion Evolution on Social Media:** Can we measure the "viral" nature of rapid style changes?

instance, an internal dashboard tool to aid decision making for stylists. Other than the in-house utility of such efficiency decision support tool, in this section potential business interests for e-commerce sites and personalization tools for commercializing B2B products are discussed. Clients also could benefit of this tool if it is robust enough, to avoid having the customer fill all the 50 question profile questionnaire (some customers, specially for men, do not finish completing).

8.1 Monetization of Fashion Knowledge Graphs

According to new reports, the virtual digital assistant market will reach \$15.8 Billion by 2021¹¹. Looking at a stylist note writing assistant, we can learn from

¹¹ <https://www.tractica.com/newsroom/press-releases/the-virtual-digital-assistant-market-will-reach-15-8-billion-worldwide-by-2021/>

similar models in the market on the fashion and also media analytics (both NLP and computer vision). How could a stylist note writing assistant help in other domains? Helping lowering the cognitive load of assembling an outfit or look within a *fix* box is the immediate benefit within SF organization. Estimating the benefits of the tool outside this concrete domain is difficult, since it is something that is not sold as a service and each domain would most probably require some fine-tuning. However, a B2B API can be provided for text analytics and insight extractions that can justify not the direct earnings or lift for the company (which may be difficult to directly measure/evaluate), but they may better justify the customer satisfaction, better improve the customers' requests and lead to less dissatisfaction and churn. For instance, a tested campaign is the so-called "Want her back" (the possibility to get the same stylist to assemble your look next time if a customer thought the stylist got his/her style). Measuring the increase in the amount of customers wanting to stick to the same stylist is an implicit feedback metric worth considering as for customer experience improvement.

8.2 Business cases on Fashion, NLP and dialogue-based personalization

There are some existing services that aim at personalizing fashion in one or other ways: *Cladwell*¹² (provides advice and a small good quality *capsule* seasonal set that provides a minimal wardrobe with your favourite pieces every season), *Bombfell* (the male version of Stitch Fix), *Trunk Club* (a higher price point version), and *Lookiero*¹³ (Spanish version of Stitch Fix) are similar services to Stitch Fix that also use fashion personalization and delivery. In the clothing delivery + renting space, *Rent the runway* and *Le Tote*¹⁴ are similar providers but for second hand clothing. *Boon + Gable*¹⁵, on the other side, include in the service the visit of the personal stylist with clothes to your home for you to get in-person advice and combine with your existing wardrobe. However, Stitch Fix is the oldest service, with 5 years of experience learning from data, and where personalization and human in the loop are key competitive advantages that make use of machine learning and internal crowd-sourcing tools.

On the information extraction and information extraction domain, some more general machine learning NLP tools as a service exists, such as *nuun.io*, *Fox-Type*¹⁶ (helps detecting politeness among others), or *Receptiviti*¹⁷ (psychology insights for analyzing tone and other aspects in natural language). Based on text-based interfaces, the bloom of chat bots is unstoppable (in Facebook Messenger, Slack, etc.); however, these are merely for customer service (*helloaida.io*) or informative (e.g., Immigration attorney *visabot.co*, *Polly.datalog.ai* and *getAsteria.com* for conversational/health/personalization purposes).

¹² <https://cladwell.com/>

¹³ <https://lookiero.com/>

¹⁴ <https://www.letote.com/>

¹⁵ <https://www.boonandgable.com/>

¹⁶ <https://foxttype.com/>

¹⁷ <http://www.receptiviti.ai/>

Other models having NLP analytics as the core value proposition are writing assistants such as *Grammarly*¹⁸, *Boomerang*¹⁹ or bot kits/conversational agents²⁰

Business models to consider for such developed tool are also in the NLP space: *claralabs*, *x.ai* or *Talla*²¹ are core NLP based meeting assistants. *Olivia.ai*, *MoneyStrands.com* and *Bond*²² are finance personal assistants, but there is also AI chat assistants for shopping such as *kip*²³. In the fashion annotation computer vision scene, examples such as *wide-eyes.it*, *21buttons.com* or *chicisimo.com* offer the counterpart to the text only annotation tools.

One time service could be as well of use for real-time outfit assessment, taking a model of a consultation fee, such as in *Remedy* remote health care assistance²⁴.

8.3 Case study survey: Potential business proposition on digital style advisors

A small survey done in social media shows the potential market of such one-time transaction style advice service. Results in Figures 11 and 12 show that 34% of people would be willing to pay for having some kind of last minute style advice based upon a *selfie* before leaving home, while 66% would not. This demonstrates that special occasions must be carefully studied and despite being a not large margin of potential benefit, it could be of business value (e.g. getting tips on mixing and matching garments for an interview, meeting, date, ceremony or other special events). Usage of existing platforms such as *Snapchat* or *Instagram*, which allow video social sharing immediately, are proposed as areas to explore.

Would you pay for a real-time outfit selfie assessment before leaving home to make sure you are wearing your best and most stylish look? This is the question to evaluate if a style advice agent would be worth building and if so, how much would people pay for it.

The transaction to pay for could be per picture asking advice for. Finding interest in the product idea involves having a prototype beta test tool providing advice from curated stylists with a karma and reputation feedback-based system. For stylists, it would let them build their reputation and grow their influence audience. This could also be, as a side effect, a way for companies like *Stitch Fix* to recruit stylists and augment the style advice knowledge base *bible*. The final product can be an app where the user send their closet pictures, or a picture of their look just before going out to that job interview, your fourth date, or your

¹⁸ <http://www.grammarly.com>

¹⁹ http://www.boomeranggmail.com/respondable/?utm_medium=email&utm_source=user+email&utm_campaign=respondable+launch&utm_content=bottle+banner

²⁰ <https://howdy.ai/botkit/>, <https://rundexter.com>

²¹ <http://talla.com/office-management>

²² https://medium.com/@uday_akkaraju/meet-bond-a-bot-thatll-make-you-richer-3edd2540bf06#.2mdudz88g

²³ <https://www.kipthis.com/>

²⁴ Remedy costs \$15 per consultation done by professionals remotely, instead of avg. \$80. <https://www.remedymedical.com/>

conference talk. Such a product can only be built after data showing interest is gathered. Key questions are: Are consumers really willing to pay for this service? How to guarantee there is going there a stylist there when you need it? How will the expert advice quality control be executed?

Because of these issues, the aim of this survey was to strictly find if users would be willing to pay for, e.g. an AI/chatbot-like service giving professional, *fashionista*, or peer advice on the go.

The Twitter survey reached 766 impressions, 41 total engagements, 24 votes, 7 detail expands, 3 retweets, 3 likes, 2 profile clicks and 1 reply in only 24h (see Figures 11, and 12). A next step would be finding what are the minimum requirements the users would want the *fashionista-bot* to have. The same survey conducted in Facebook without choice answers (i.e., open question) resulted in a larger proportion of negative comments such as the following (combined with 11 likes):

- *No, because there are more important matters in this world! ;)*
- *Yes, because what you wear is a perspective of your taste!*
- *Then why would you let anybody else decide on it , instead of yourself ?*
- *You cannot say that what you are now wearing is only decided by you!! before even you want to buy a piece of clothing, it is already decided by designers and fashion gurus. ;) you only pick and try to come up with a set of outfit. I would rather have a professional/intellectual idea on that before wear something hideous!*
- *No, I do not even use it for free.*
- *Nope, a mirror is enough*
- *I would pay not to get assessed*
- *No*
- *Personally I would not like anybody to "tell me" what I should be wearing. I like, though, an app that would suggest an outfit in the morning. I think there could be some use cases for this though: people that need to look professional each day but have no time or interest to decide their own outfits every time. In that case the app could be your "personal stylist", although I'm uncertain of how good the app could decide if something flatters you or not, because its not only a matter of proportions... If the clothes and outfits would already be "verified" by a stylist beforehand, then the risk of miscalculations gets smaller. Also, in the case that the app really would know what looks good on you or not, what are the chances of it being "creative"? A lot of inventions, art and fashion arises from accidents, miscalculations or just wild fantasy, so there is always the risk of the app being too safe and boring :) But interesting thought anyhow! I'm always interested in questions of fashion, although i prefer to ponder on them without a computer normally :).*
- *Yes! This kind of app would be my hero since I do not need to think about my outfit of the day. And believe me sometimes it takes half an hour to come up with a nice outfit! Personally speaking, I have a strong connection to my outfit :D It can change my mood easily and when I am not wearing based on my standards, I feel uncomfortable in my skin. I really do not care what others*

- think about me! I care what I feel about myself! Even more exciting option would be to have an app to suggest your outfit in the morning! Sounds cool!*
- *The correct attire is always combat boots, cargo pants and a black tee-shirt. Add sunglasses and dog tags if you're feeling fancy. Regardless of gender. :)*
 - *I would not. I think it's time consuming, i don't trust apps to do this yet... Maybe in 20 years.*
 - *Thanks but no thanks*
 - *I would use the money and buy a stylish brand!*
 - *No*
 - *Nope. I wouldn't want to be paid for having it. My outfit is my decision. How does this thing not sound like society telling you how you should dress? (especially for women I guess)*
 - *Nope*



Fig. 11: Social media survey results for potential monetization.

A "magic mirror" application for a last minute "Am I looking great?" before exiting home could be interesting, but according to the survey results, it is of doubtful value, since it appears to exist a very niche and limited prospect on acquiring paying users. The survey would require a final tool prototype for further exploration and business development refinement.

8.4 Target market

Customers of machine learning-targeted media (text and image) annotation tools can serve a B2B model for customers of a varied range. The range of potential




Fig. 12: Potential monetization for viability of business models: impressions

customers includes a very diverse range of business, considering the fashion world only initially, for instance, customers of subjective influence networks could include:

1. *Chicfy.com* A service for online second hand sale and swap. The system’s business model consists of taking a commission for each transaction made.
2. *Vinted.com*: An application similar to the previous system with no commissions.
3. *21Buttons.com* A service for users to rate other users outfits and find and buy with one click the clothes worn. Users uploading an outlook picture get a commission if other users buy any of their exhibited items.
4. *wide-eyes.it* A service that provides widgets for online shops to show visually similar elements to the searched ones.
5. *CrowdStar*²⁵ A fashion mobile game where users dress digital mannequins and receive credits if their looks have the highest ranking in n. of votes. The business model is in-game purchases (for a more varied wardrobe), clothes are real and can be bought.
6. *Betabrand*²⁶: A site where anyone can create and upload any garment design, crowd-fund or crowd-source garment patterns. Betabrand uses ”Crowd Funding Predictive Modeling” to achieve market research through crowd funding and help in accurately predicting inventory adoption rates and future demand of catalogues to minimize the risk of a poorly-performing product or excess inventory.
7. Media data-driven insight providers, such as Secret Sauce Partners²⁷, a Budapest and San Francisco-based company that offers a data-driven merchandising (DDM) platform for *fit predictor* (finds a shopper’s best fitting size in seconds), a *style finder* (apparel shopping by visual features) and a *outfit maker* (automated outfit recommendations based on stylistic matches)²⁸.

²⁵ www.crowdstar.com/

²⁶ www.betabrand.com/

²⁷ www.secretsaucepartners.com/

²⁸ <http://tech.gilt.com/2015/01/27/new-gilt-product-feature-fit-predictor>

8. *E-commerce sites* that are not fully tech and data-driven, such as *Gilt* or *Stylelect*²⁹ or *therealreal.com* (luxury consignment sales), and *Best Secret*³⁰ (invitation only luxury discounted sales).

9 Conclusions

In this paper we propose a new ontological augmentation in the fashion domain, which represents subjective feature information as an influence network. Because fashion (just like art, music or languages) strongly contains subjective information (cultural phenomena which are not designed nor engineered), we believe that such an augmentation might result in the construction of higher performing machine learning and data analysis systems.

Following the theoretical modeling, we suggest quantitative measures to assess the framework’s utility to machine learning systems. Especially we focused on quantifying how well the ontology can represent domain data, and how the features from an influence network could be integrated into machine learning systems.

9.1 Future Work

Future work on cognitive cities use cases and applications within the culture, fashion or leisure domains should focus on ubiquitous but effortless recommendation systems based on the user’s digital and off-line footprint. In fashion, an example along these lines is the aim of the application that results from H&M’s partnership with Google for creating a customized ”data dress” based on the places a person hangs out after work or the restaurants she visits³¹. Other example of fashion trend spotting on Google is using multiple markets focusing on apparel trends to enable a better understanding of how trends spread and behaviors emerge across markets³². Other ideas can extend the work of [21] as well as our illustrated examples (poster³³). Interestingly, the work in [21] uses translational topic coherence among some crowd-sourcing tasks in order to translate (abstract) style-language into (concrete) element-language and, in this way, generate recommendations from natural language requests or subjective description notes. This is a step that, against Stitch Fix’s philosophy, has as target to achieve a fully digital stylist (without the human in the loop). In this work, although PLTM (Polylingual topic model) were not originally intended to support direct

²⁹ gilt.com, <http://tech.gilt.com/>, www.stylelect.com

³⁰ www.bestsecret.com

³¹ <http://tinyurl.com/hxg6r7e>

³² <https://storage.googleapis.com/think/docs/twg-fashion-trends-2016.pdf>.

<http://www.businessinsider.com/google-partners-with-hm-ivyrevel-for-coded-couture-project-2017-2>

<http://www.ivyrevel.com/>

³³ <https://github.com/NataliaDiaz/PostersAndPresentations/blob/master/posters>

translation between languages, in domains where word order is unimportant, given a document in one language, they can be used to produce an equivalent document in a different *language* by identifying high probability words³⁴

However, the extraction of the most precious information from experts in a context-aware manner and its continuous integration in a broad dimensionality space for a timely recommendation is still a large part of the machine learning bottleneck. These areas are those where future resources can be well allocated in machine learning in general.

Future work on subjective influence networks will instantiate concrete machine learning problems into the proposed approach. For instance, an example can be quantifying fuzzy influence networks in social media opinions [24]. In this way we will validate our theoretical assumptions by incarnating and materializing different influence functions, distance and quality measures, scales and other parameters for our model assessment and evaluation in different machine learning problems. Future efforts should also be put into considering the integration of KGs into black box neural recommendation pipelines.

Integration of both text and images would be another area to enrich collective intelligence and a way of achieving data pooling for context awareness. For instance, *wide-eyes.it* focuses on producing visually similar images on a B2B manner. Integrating insights of text together with images is paramount for cross-referencing. Embeddings methods integrating both text, image and social influence are still in early stage. More automatic ways of integrating expert feedback and new approaches to having the human in the machine learning loop should be explored in the area of unsupervised reinforcement learning [1, 20]. Exploiting other social network channels' text feeds is also an under-explored area within most organizations (which are limited to analyze only the one or two most popular social feeds). Only in this way we will be able to assess properly reward functions that affect and evaluate the system's learning as a whole.

Ultimately, these actions will help refining the model's capability to effectively quantify influence and subjectivity in fashion, style and other subjective and more volatile domains.

10 Acknowledgments

This work was done within an internship program of the EIT Digital doctoral school from EU as part of the Innovation & Entrepreneurship (I&E) curriculum. We thank EIT Digital Doctoral School, advisors Prof. Johan Lilius and PhD. Jussi Autere for the financial and human support throughout the Innovation & Entrepreneurship education. From Stitch Fix team we thank Jay B. Martin for his generous editorial assistance and management, Ian Horn for mentoring,

³⁴ E.g., given a document in the element language, we can infer the topic distribution for that document under the trained model: *Since the topic distribution for a document will be the same in style language, we can produce an equivalent outfit in the style language by identifying high probability words in that language*[21].

and the rest of the Human Computation and Algorithm and Analytics teams at Stitch Fix for the inspiring and collaborative atmosphere (John McDonnell, Xian Li, Katherine Livins, John Clevenger, Roberto Sanchís Ojeda, Hoda Eydgahi, Jay Wang, Akshay Wadia, Sky Jin, Sonya Berg, etc.). Likewise, we thank the EU COST Action on

References

1. P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *Proceedings of the twenty-first international conference on Machine learning*, page 1. ACM, 2004.
2. S. Ajmani, H. Ghosh, A. Mallik, and S. Chaudhury. An ontology based personalized garment recommendation system. In *Web Intelligence (WI) and Intelligent Agent Technologies (IAT), 2013 IEEE/WIC/ACM International Joint Conferences on*, volume 3, pages 17–20. IEEE, 2013.
3. K. Bollacker, N. Díaz-Rodríguez, and X. Li. Beyond clothing ontologies: Modeling fashion with subjective influence networks. *KDD workshop on Machine Learning meets Fashion*, 2016.
4. K. Bollacker, C. Evans, P. Paritosh, T. Sturge, and J. Taylor. Freebase: A collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD '08, pages 1247–1250, New York, NY, USA, 2008. ACM.
5. J. G. Breslin, A. Harth, U. Bojars, and S. Decker. Towards semantically-interlinked online communities. In *The Semantic Web: Research and Applications*, pages 500–514. Springer, 2005.
6. N. Díaz Rodríguez, M. Cuéllar, J. Lilius, and M. Delgado Calvo-Flores. A fuzzy ontology for semantic modeling and recognition of human behavior. *Knowl. Based Syst*, 66:46–60, 2014.
7. N. Díaz-Rodríguez, M. P. Cuéllar, J. Lilius, and M. D. Calvo-Flores. A survey on ontologies for human behavior recognition. *ACM Comput. Surv.*, 46(4):43:1–43:33, Mar. 2014.
8. S. D’Onofrio and E. Portmann. Cognitive computing in smart cities. *Informatik-Spektrum*, 40(1):46–57, 2017.
9. N. Díaz-Rodríguez, O. L. Cadahía, M. P. Cuéllar, J. Lilius, and M. D. Calvo-Flores. Handling real-world context awareness, uncertainty and vagueness in real-time human activity tracking and recognition with a fuzzy ontology-based hybrid method. *Sensors*, 14(10):18131–18171, 2014.
10. R. Guha, D. Brickley, and S. Macbeth. Schema. org: Evolution of structured data on the web. *Communications of the ACM*, 59(2):44–51, 2016.
11. M. Hepp. Goodrelations: An ontology for describing products and services offers on the web. In *Knowledge Engineering: Practice and Patterns*, pages 329–346. Springer, 2008.
12. J. Li and Y. Li. Cognitive model based fashion style decision making. *Expert Systems with Applications*, 39(5):4972 – 4977, 2012.
13. J. Miller. *Fashion and Music*. Bloomsbury Publishing, 2011.
14. I. Novalija and G. Leban. Applying NLP for building domain ontology: Fashion collection. A:147–150, 2013.

15. A. Ragone, U. Straccia, F. Bobillo, T. Di Noia, E. Di Sciascio, and F. M. Donini. Fuzzy description logics for bilateral matchmaking in e-marketplaces. In *Description Logics*, 2008.
16. Y. Raimond, S. A. Abdallah, M. B. Sandler, and F. Giasson. The music ontology. In *ISMIR*, 2007.
17. Y. Raimond and M. Sandler. Evaluation of the music ontology framework. In *The Semantic Web: Research and Applications*, pages 255–269. Springer, 2012.
18. A. J. Scott. Beyond the creative city: Cognitive-cultural capitalism and the new urbanism. *Regional Studies*, 48(4):565–578, 2014.
19. A. Steinfeld, S. R. Bennett, K. Cunningham, M. Lahut, P.-A. Quinones, D. Wexler, D. Siewiorek, J. Hayes, P. Cohen, J. Fitzgerald, O. Hansson, M. Pool, and M. Drummond. Evaluation of an integrated multi-task machine learning system with humans in the loop. In *NIST Performance Metrics for Intelligent Systems Workshop (PerMIS)*, 2007.
20. A. L. Thomaz and C. Breazeal. Reinforcement learning with human teachers: Evidence of feedback and guidance with implications for learning performance. 2006.
21. K. Vaccaro, S. Shivakumar, Z. Ding, K. Karahalios, and R. Kumar. The elements of fashion style. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, UIST '16, pages 777–785, New York, NY, USA, 2016. ACM.
22. D. Vogiatis, D. Pierrakos, G. Paliouras, S. Jenkyn-Jones, and B. J. H. H. A. Posen. Expert and community based style advice. *Expert Syst. Appl.*, 39(12):10647–10655, Sept. 2012.
23. C. Wah, S. Branson, P. Perona, and S. Belongie. Multiclass recognition and part localization with humans in the loop. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 2524–2531. IEEE, 2011.
24. L. Wang and J. M. Mendel. Fuzzy opinion networks: A mathematical framework for the evolution of opinions and their uncertainties across social networks. *CoRR*, abs/1602.06508, 2016.
25. L. Wang, X. Zeng, L. Koehl, and Y. Chen. Intelligent fashion recommender system: Fuzzy logic in personalized garment design. *Human-Machine Systems, IEEE Transactions on*, 45(1):95–109, 2015.
26. S. Zoghbi, G. Heyman, J. C. Gomez, and M.-F. Moens. Cross-modal fashion search. In *MultiMedia Modeling*, pages 367–373. Springer International Publishing, 2016.
27. S. Zoghbi, I. Vulić, and M.-F. Moens. I pinned it. where can I buy one like it?: Automatically linking pinterest pins to online webshops. In *Proceedings of the 2013 workshop on Data-driven user behavioral modelling and mining from social media*, pages 9–12. ACM, 2013.