AIS SIGPRAG

Pre-ICIS Workshop 2015
"Practice-based Design
and Innovation of
Digital Artifacts"

Supporting Business Model Idea Generation Through Machine-generated Ideas: Towards a Design Theory

Thomas John

Chair of Information Management & E-Finance, University of Paderborn
thomas.john@uni-paderborn.de

Abstract

Successful business model innovation is impossible without innovative business model ideas. However, human capacity for generating ideas is limited in a number of ways, for example, in terms of the amounts of prior knowledge and cognitive flexibility that humans can possess. Therefore, with *business model idea generators* we propose a new class of information systems that can contribute to alleviating these limitations. We envision these idea generators to generate ideas that complement the ideas generated by humans, which we hope increases the overall quality of business model ideas available in a given context, and thereby leads to higher rates of successful business model innovation. Our contribution is a design theory that describes the high-level architecture of the idea generator systems that we propose.

Keywords: Business model innovation, business model idea generation, creativity, collective intelligence, machine learning

1 Introduction

A business model describes a firm's mechanisms for value creation, value delivery and value capture (Teece 2010), and as such is a detailed description of a firm's strategy (Adner et al. 2014; Casadesus-Masanell and Ricart 2010). The interest in business models and business model innovation is enormous – from researchers and practitioners alike. A recent global survey of some 3,000 executives in 26 countries finds that a majority of 60% consider "defin[ing] an effective business model" a major challenge in their firms' innovation activities (GE 2014, p. 40). IBM's global CEO studies (IBM 2006, 2008, 2010, 2012) consistently underline the importance of business models, with each study drawing on interviews with several hundreds to nearly 2,000 CEOs. Virtually all CEOs seek business model innovation at least *"moderate[ly]"*, and more than two-thirds aspire *"extensive"* business model innovation (IBM 2008, p. 48). In line with the interest among practitioners, academic attention to business models has increased rapidly in disciplines as varied as information systems, innovation management, and strategy (Zott et al. 2011).

High quality raw ideas are essential for successful innovation (Kornish and Ulrich 2014). Consequently, idea generation is an essential part of business model innovation processes (Schneider and Spieth 2013) or, put bluntly, *"ideas constitute the* *lifeblood for firms in generating […] new business models”* (Ende et al. 2014, p. 1). However, at the same time, *“idea generation […] is the step in business model innovation that is least understood”* (Martins et al. 2015, p. 8). A recent review of business model research (Schneider and Spieth 2013) underscores this point, and the authors identify the need to better understand how to support firms in business model idea generation as an important direction for future research.

The goal of this article is to respond to this call by proposing a new class of information systems that we refer to as *business model idea generators* (or, simply, *idea generators*). We intend these idea generators to improve idea generation by providing raw ideas (Kornish and Ulrich 2014) for business models – with the ideas being specifically tailored to the product or service for which business model ideas are sought. The generated ideas are intended to complement human idea generation, and to highlight opportunities that humans may overlook. The ideas thus have the potential to increase the overall quality of ideas that are available in a given business model context, and thereby help to increase the probability of successful business model innovation.

Our work represents an important departure from current research on business model idea generation, as current research without exception takes a human-centered approach (i.e., exclusively relies on human-generated ideas). For example, a number of modeling tools exist that support humans in recording intermediate ideas during idea generation, either physically (e.g., Osterwalder and Pigneur 2010) or digitally (e.g., Gordijn and Akkermans 2003). Moreover, idea stimuli have been proposed in the form of catalogs of business model patterns (e.g., (Abdelkafi et al. 2013; Gassmann et al. 2014) or concrete business models (e.g., Stampfl and Sniukas 2013). In contrast, we take a machine-centered approach that is inspired by blind variation/selective retention theory (BVSR, Campbell 1960; Simonton 2011), and draws on the premise that ideas generated by a machine can be valuable catalysts of the human idea generation process. BVSR states that, when humans produce creative ideas, these ideas result from iterations of idea creation and idea evaluation. In the course of these iterations, knowledge concerning the quality of partial ideas is accumulated, leading to better ideas being created in every iteration – and eventually to high-quality, creative ideas (Simonton 2011). We propose that a business model idea generator can quasi-automate the process implied by BVSR by iteratively performing the following three steps. Step 1 *idea creation* creates intermediate ideas by forming novel combinations of the knowledge stored in a business model knowledge base. Step 2 *idea evaluation* determines the quality of these ideas through crowd evaluation (Mollick and Nanda 2015). This step is performed quasi-automatically by connecting to an existing crowdsourcing platform (e.g., Crowdflower, http://www.crowdflower.com) through its API (application programming interface). Step 3 *knowledge accumulation* derives knowledge concerning promising partial ideas through supervised machine learning (Jordan and Mitchell 2015), and employs as training data the business model ideas created in step 1 and evaluated in step 2. After any given iteration, the knowledge built up in step 3 knowledge accumulation guides the choice of which business model ideas to create in step 1 of the subsequent iteration. Taken together, repeatedly running through steps 1-3 should lead to a gradual buildup of knowledge concerning the features that characterize promising business model ideas in a given context (i.e., for a given product or service), and hopefully, in the end leads to high-quality ideas.

The contribution of our paper is an information systems design theory that, as a *"systematic specification of design knowledge"* (Gregor and Jones 2007, p. 314), describes the high-level architecture of business model idea generators. With that design theory we intend to introduce a completely new perspective into research on business model idea generation, a perspective which is machine-centered rather than human-centered. Our approach draws on and integrates research in creativity, collective intelligence, and machine learning. It might help to alleviate limitations inherent in human-centered idea generation, such as limitations in cognitive flexibility and the available prior knowledge (Dane 2010), and thereby may contribute to better business model innovation ideas. While, from the perspective of business model research, we contribute to addressing an important problem, from the perspective of information systems (IS) research we contribute to grasping what has been called a *"unique opportunity"* for IS research, namely to leverage IS competences for the sake of exploring strategic objects such as business models (Osterwalder and Pigneur 2013, p. 239). In the following, we first summarize prior work in creativity research to motivate why a machine-centered approach to idea generation may have benefits compared to a human-centered approach. We then go on to describe our design theory, and finally outline plans for its empirical evaluation.

2 Why Machine-centered Creativity Is Worth Exploring

To better understand why machine-centered idea generation may carry benefits compared to human-centered idea generation, we first sketch the limitations that humans have when trying to be creative, and then go on to describe how an idea generator might contribute to alleviating these limitations. The presented reasoning might in a variety of aspects be considered simplistic and artificial. However, our intention is not to provide a full-fledged comparison of the relative advantages that humans and machines may have, but rather to provide a thought experiment that, grounded in cognitive theory, provides an intuition of the factors that limit human creativity. Keeping this in mind, we hope the presented reasoning is able to motivate why there is value in exploring how machines could support human business model idea generation endeavors.

Prior research has addressed the question of *'How does human creativity work?'* at a number of different levels. These include the neurological level (i.e., where is creative capacity located in human brains?), the cognitive level (i.e., abstracting from specific brain locations, how is creativity created subconsciously?), the individual level (e.g., how can creativity be promoted consciously, for example, through which creativity techniques?), the group level (e.g., what group composition makes a creative group?) and higher levels of analysis such as organizations and societies (Amabile and Hennessey 2010). Most relevant in our context is the cognitive level, because limitations at this level propagate to and affect all higher levels of analysis, and thus are central to human creative capacity. In line with this, recent research emphasizes the importance of a cognitive perspective for improving our understanding of how to promote business model idea generation (Martins et al. 2015).

At the cognitive level, creative performance is substantially affected by the two factors of *domain expertise* and *cognitive flexibility* (Dane 2010). By integrating a vast array of prior work, Dane (2010) was the first to establish the central role of these two factors (and their interaction) for determining creative performance. Domain expertise is domain-specific knowledge that can, for example, be acquired through deliberate practice (Ericsson and Charness 1994) or experiential learning (Armstrong and Mahmud 2008). Expertise is captured in the form of schemas, which are structures containing *"knowledge about a concept or type of stimulus, including its attributes and the relations among those attributes"* (Fiske and Taylor 1991, p. 98; see Figure 1 for an illustration). These schemas get relatively more detailed and accurate as an individual acquires expertise. Accordingly, knowledge of existing business models is mentally stored in schemas (Martins et al. 2015) that get more elaborate as one acquires more expertise concerning the central choices underlying a given business model. For example, a novice's schema of Zara's business model may comprise that Zara offers fashion and does so with extremely short turnaround times (it takes less than two weeks from the first idea to the final product being available in stores). A more expert schema might add that Zara procures especially from local – and more expensive – suppliers. This additional information would allow capturing as well that there are interrelations, for example, that buying from local suppliers enables Zara to offer such short turnaround times, because the choice of local rather than oversees suppliers accelerates communication and distribution processes. The expert schema of Zara's business model would comprise more attributes (Zara offers fashion, it does so with short turnaround times, and procures from local suppliers) and more interrelations between these attributes (local suppliers enable short turnaround times) (example adapted from Priem et al. (2013)).

Creative ideas arise from novel recombinations of the knowledge captured in a person's domain schemas (e.g., existing business models known by that person). Put differently, the schemas represent the *'raw material'* for new ideas, and the more raw material there is available for creating ideas, the higher is the potential for creative ideas to actually be created (see Figure 2). As new ideas arise from novel recombinations of existing knowledge, a prerequisite for creativity is that domain schemas are flexible so as to allow changing and combining them with the aim of deriving new ideas. Consequently, creativity does not only increase with domain expertise, but also increases with higher levels of cognitive flexibility. However, having invoked schemas numerous times (as it typically happens in the course of building up expertise) makes them inflexible. To illustrate, consider how you go about accelerating and braking while driving your car. Having driven a car for years makes us internalize that the gas pedal is on the right side, and the brake pedal is left of the gas pedal. The upside is that we do not have to think any more about which pedal is where. The cor-

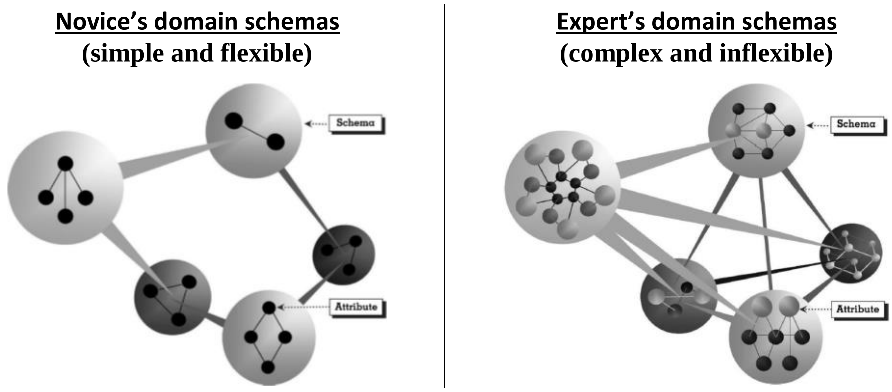

Figure 1: Illustration of domain schemas of novices and experts (Dane 2010).

responding schemas have been invoked numerous times and have become *'hard-wired'* in our brains – finding the correct pedal happens *'automatically'*. However, the downside of this automation is that if the pedals were switched, it would be rather tedious for us to get used to the new positions. In contrast, someone who has learned driving just recently would have substantially less trouble getting used to the new positions, because the corresponding schemas are still more flexible. The same effect applies to creative idea generation and results in a trade-off between expertise and flexibility (at least regarding radical ideas). Put differently, *"the relationship between domain expertise and radical idea generation takes the form of an inverted U"* (Dane 2010, p. 588). Translated to the business model context, this means that having been exposed to the business models prevalent in a given industry over years makes it cognitively more difficult to break free from these business models, which limits one's ability to create radically new business model ideas even though the acquired industry knowledge would put one into an expert position.

To summarize, when generating business model ideas, humans make use of two properties of the human cognitive system that enable their creative capacity and, at the same time, constrain it: First, they use the capacity to build up knowledge, that is, knowledge of existing business models. Second, they use the capacity to recombine that knowledge in novel ways to actually arrive at new business model ideas. However, these properties are designed in ways that constrain human creativity in three important ways. First, the capacity to build up knowledge is limited because learning takes time (Simon 1996) – no one can possibly know all business models that exist. Second, cognitive flexibility is limited because it is inherent in human memory that it is easier to retrieve information that is associated with each other than information that is not (Kohonen 2012). Third, as one acquires knowledge (which typically goes along with knowledge reuse), one loses the flexibility to form novel recombinations of that knowledge (Dane 2010). The centrality of these limitations can be seen in the fact that the vast majority of approaches for promoting creative capacity, in one way or another, seek to address one of these limitations (or both). For example, going from individual to group idea generation, or activating even more individuals through crowd sourcing and open innovation, are simply ways of broadening the available knowledge base (e.g., Kornish and Ulrich 2011). Employing creativity techniques that facilitate changing perspective and questioning assumptions (e.g., Smith 1998) are simply means to increase cognitive flexibility. Finally, employing the business model patterns and business model catalogues mentioned above broadens the knowledge base (if the business models are yet unknown to an individual) and, at the same time, increases cognitive flexibility (if deliberate effort is undertaken to apply a certain business model to a focal firm).

Having illustrated how human-centered idea generation proceeds, we now come to the benefits that machine-centered idea generation potentially has. Figure 2 compares the limitations of the human cognitive system (i.e., limited domain expertise, limited cognitive flexibility, the trade-off between domain expertise and cognitive

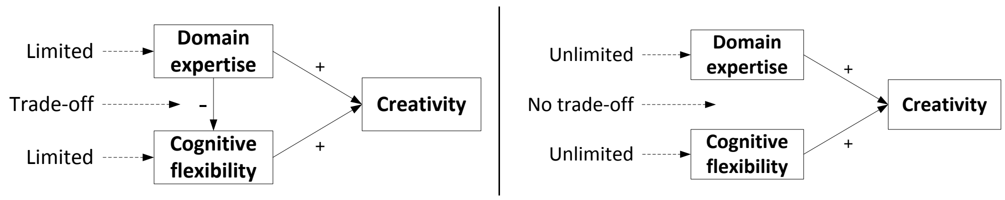

Figure 2: Simplified illustration of the interdependencies between expertise, flexibility, and creativity in human-centered and machine-centered creativity.

flexibility) with the characteristics that (at least theoretically) a machine could have with regards to creativity. First, a machine could possess unlimited *'domain expertise'* because its expertise could be built up cooperatively by many individuals, rather than one individual accumulating knowledge only on her own (i.e., a machine could possess more business model knowledge, or raw material for new ideas, than any individual possibly could). Second, a machine is also unlimited with regards to its *'cognitive flexibility'*. A machine is not by its nature constrained to more easily retrieve pieces of information that are closely associated with each other. Finally, machines, unlike humans, do not suffer from the trade-off between domain expertise and cognitive flexibility (obviously, performance may suffer as database size grows; however, that effect is negligible for the sake of our argument). Taken together, at least within the simplistic worldview that we adopted, our observations seem to suggest that machines can have advantages over humans when it comes to creativity. This does not contravene that outside that worldview humans have advantages. However, as we are not seeking to replace humans' ideas, but rather seeking to complement them, it should suffice to state that there may be circumstances under which machines have the potential to generate ideas that are different from those generated by humans, but still valuable. Therefore, in the following we explore how this potential may be tapped for the purpose of business model idea generation. We do so by proposing a theoretically grounded design theory that describes the high level architecture of machines that generate business model ideas, or business model idea generators.

3 A Machine-centered Approach for Increasing Creative Capacity

For documenting our design theory, we adopt the recommendations by Gregor and Jones (2007) who propose that a design theory should comprise the following eight components: purpose and scope, constructs, principles of form and function, artifact mutability, testable propositions, justificatory knowledge, principles of implementation, and expository instantiation (with the last two components being optional). In the following, we address each in turn.

3.1 Purpose and Scope

The purpose of the proposed design theory is to provide prescriptive and explanatory knowledge about how to design the high-level architecture of systems that we term business model idea generators. Such systems are intended to support business model idea generation processes of individuals or groups for a given product or service (the product/service may or may not already be existing). Such individuals or groups include anyone who might be confronted with the task of developing a business model, which includes entrepreneurs, innovation managers, product managers, and consultants. The systems resulting from the proposed design theory would not replace these individuals (groups) or replace their ideas. Rather, we expect these systems to provide high quality ideas that, while most likely being in need to be refined, increase the overall quality of available ideas in a given business model innovation context. The potential for idea generators to generate such ideas arises from that they are not bound to some of the limitations that humans have when trying to be creative (as described in the previous section). Put in the context of the general innovation process (Hansen and Birkinshaw 2007), we envision idea generators to be used before or parallel to the initial phase of idea generation. In that sense, idea generators either prepare human idea generation or complement it (see Figure 3).

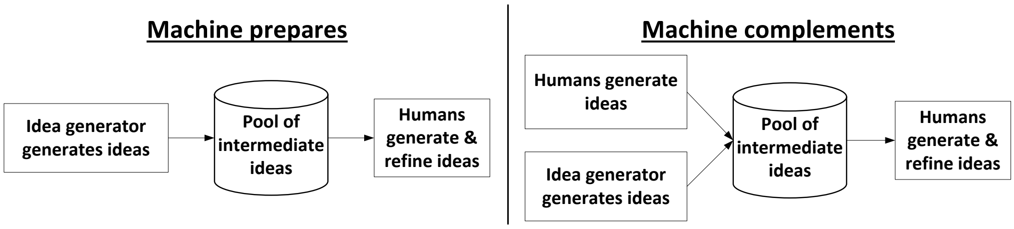

Figure 3: Use contexts for the business model idea generator.

3.2 Justificatory Knowledge

In the following, we derive six design principles from prior theoretical and empirical work in creativity, collective intelligence and machine learning. These design principles capture the justificatory knowledge for the proposed design theory for business model idea generator architectures.

Principle 1: Iterate idea creation, idea evaluation, and idea quality accumulation. Creativity research has found that humans generate creative ideas by going through numerous iterations of idea creation and idea evaluation. One of the most prominent proponents of this view is blind variation/selective retention theory, whose explanatory accounts have received wide-spread support (Campbell 1960; Simonton 2011). However, the idea of iterative idea creation/idea evaluation has been expressed similarly by others. For example, creative processes have been characterized to involve alternations of divergent thinking (i.e., idea creation) and convergent thinking (i.e., idea evaluation, Mumford et al. 1991). Likewise, creating random stimuli, and subsequently reinterpreting these stimuli has been shown to promote creativity (Finke et al. 1992). The underlying idea is that the step of idea creation creates novelty, while the step of idea evaluation ensures that novelty goes along with usefulness. Obviously, learning needs to take place from one iteration to the next, because otherwise every idea creation step would naïvely create ideas that expectedly are no better than the ideas created in the previous iteration. Therefore, idea creation and evaluation need to be complemented by idea quality accumulation.

Principle 2: Implement idea creation through applying cognitive procedures to domain knowledge. As outlined in section 2, idea creation is performed by recombining prior knowledge in novel ways. There are a variety of different procedures that can be employed for this purpose, with analogical reasoning and conceptual combination being examples recently highlighted in the business model context (Martins et al. 2015). Analogical reasoning involves transferring structural properties from business models in one domain to business models in another. This could, for example, involve applying the *freemium* business model pattern that is quite popular with smartphone apps to other industries not familiar with this type of business model (which would result in new business model ideas for the target industry). Conceptual combination involves combining two entities while selectively retaining properties of one or the other to create a new entity that is different from both the original ones. This could involve combining properties of search engine business model (e.g., ad-financing) with properties of car business models so as to arrive at car business models that involve subsidizing prices with ads printed onto the cars. As noted, other cognitive procedures exist. It is merely necessary that an idea generator can invoke at least one of these to create new ideas.

Principle 3: Implement idea evaluation quasi-automatically through a crowd (drawing on the crowd's knowledge of the use context). Currently (and for the foreseeable future), machines will not be able to evaluate the creativity of ideas at a level of sophistication that is comparable to that of humans (Colton and Wiggins 2012). Employing expert judges for rating creativity has been termed the *"gold standard"* for assessing creativity (Baer and McKool 2014). However, empirical evidence has emerged recently for that non-experts can assess creativity at a level comparable to that of experts (Magnusson et al. 2014; Mollick and Nanda 2015). This statement holds at least as long as the non-experts are able to understand the ideas that they are assessing (which, in a way, makes them experts again). To illustrate this point, while it would be reasonable to let crowds evaluate business models for *perfume,* it would most likely not be reasonable to let them evaluate business models for *enterprise resource planning* (ERP) *systems* (as the average crowd user is unlikely to know what an ERP is, not to mention what qualities a well-designed ERP business model should have). Crowd platforms such as Crowdflower (http://www.crowdflower.com) or Amazon Mechanical Turk (http://www.mturk.com) allow tapping hundreds of thousands of users and, through their APIs, make it possible to automatically create tasks and quickly retrieve results. In manual tests on Crowdflower we typically receive results in less than one hour (requiring ideas to be rated by 10 raters).

Principle 4: Implement idea quality accumulation through machine learning. If knowledge is to be built up over the course of multiple iterations, obviously in each iteration the newly acquired knowledge on idea quality needs to be integrated with the knowledge acquired in previous iterations. Put differently, what has been learned about *'what makes a good business model?'* in iterations 1 to n-1 needs to be updated (and potentially revised) using the knowledge acquired in iteration n. From a technical perspective, this can be achieved with a supervised machine learning algorithm (Jordan and Mitchell 2015) that takes pairs of business model and the corresponding evaluation as training data. With every iteration (i.e., every additional set of training data), the algorithm's ability to predict evaluations for a new business model improves – which corresponds to that the algorithm improves in being able to discern *'promising'* from *'not so promising'* business model ideas.

Principle 5: Guide idea creation through the accumulated idea quality knowledge. Creating all ideas that can be derived from the stored domain knowledge would minimize the risk of overseeing possibly promising business model ideas. It would also technically be possible to create all these ideas (even if they go into the millions). However, evaluating all these ideas would be problematic. Given a large enough crowd, evaluation would potentially still be feasible in acceptable amounts of time. But evaluating such great amounts of ideas would be rather costly and, given that most randomly created ideas will obviously be of no or little value, it would be rather inefficient to create all ideas. Therefore, there is the need to prioritize which ideas to create. This can be achieved on the basis of the accumulated idea quality knowledge, which allows selecting a subset of the most promising business model ideas for actual creation (e.g., 50 ideas). Obviously, in the first iteration accumulated idea quality knowledge does not exist yet, so in that iteration ideas need to be created at random.

Principle 6: Integrate idea evaluation by converting ideas into a "crowd-readable" format and back. On the one hand, ideas within the idea generator will likely be in a rather formal format. On the other hand, the crowd needs the ideas to be in a human-readable format, and with instructions that prescribe how to evaluate the ideas (e.g., on a scale from *1 – not at all creative* to *7 – very creative*). Therefore, the following two-step conversion is necessary. First, right after idea creation, ideas need to be converted from the rather formal format into a human-readable format (e.g., by means of natural language generation). Second, the ideas need to be converted into a format that is compatible with the API of the chosen crowd platform (e.g., meta information needs to be added to convey how the ideas are to be presented and on which scale they are to be rated). Right after the idea evaluation step, the conversion needs to be done the other way round.

3.3 Constructs

From the design principles presented above, the following constructs can be derived:

Table 1: Constructs of the proposed design theory (d = data, f = function/activity).

#	Name	Description
1	Domain (d)	Business model knowledge base (i.e., the raw material for new ideas)
2	Cognitive procedures (d)	Procedures that create new ideas by being applied to domain knowledge
3	Use context (d)	Knowledge on the context that ideas will be used in
4	Unevaluated ideas (d)	Newly created ideas before evaluation
5	Evaluated ideas (d)	Newly created ideas after evaluation
6	Accumulated idea quality (d)	Accumulated knowledge on what features characterize promising ideas
7	Idea creation (f)	Applies cognitive procedure(s) to domain knowledge to create new ideas, deliberately selects business model ideas to be created based on aggregated idea quality
8	Idea evaluation (f)	Employs use context knowledge to evaluate ideas
9	Idea quality accumulation (f)	Builds up knowledge about the characteristics of promising ideas
10	Conversion (f)	Converts ideas into a *"crowd-readable"* format and back

3.4 Principles of Form and Function

In the following, we describe the principles of form and function that result from the presented design principles and constructs (see Figure 4). The architecture shown in Figure 4 implies that, at the highest level, a user specifies a product or service as the input for the idea generator, and receives business model ideas that are suitable for the specified product or service. In the first iteration, *idea creation* randomly creates a number of business model ideas (e.g., 50) by applying one or several *cognitive procedures* to the *domain knowledge*. The resulting ideas, together with the product/service description, are stored in *unevaluated ideas*. Thereafter, the *conversion* prepares the *unevaluated ideas* for the crowd platform, which triggers *idea evaluation*. *Idea evalu-*

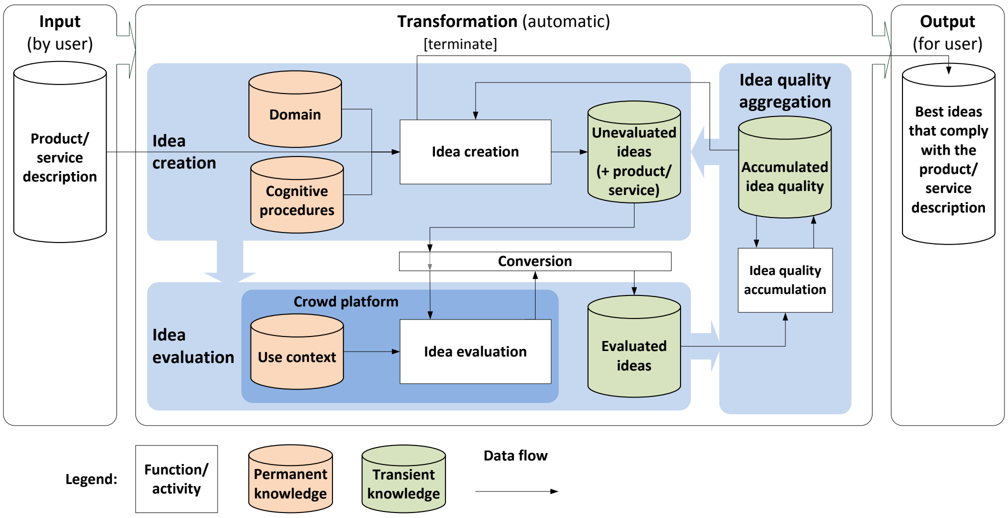

Figure 4: Use contexts for the business model idea generator.

ation is performed by the crowd by drawing onknowledge of the *use context*. In the perfume example mentioned earlier, this would involve that the crowd draws on its knowledge of how perfume is used so as to evaluate the business model ideas. *Conversion* therefore would need to ensure that the product/service description provided by the user of the idea generator is made available to the crowd (e.g., by automatically including it in the task description that crowd workers receive). After *idea evaluation*, *conversion* transforms the ideas back into the format of the idea generator and stores them together with the corresponding evaluations in *evaluated ideas*. In the first iteration, *idea quality accumulation* would simply hand over the information on evaluated ideas to *accumulated idea quality* (the actual role of *idea quality accumulation* will become clear from the second iteration onwards). In the second iteration, *idea creation* does not randomly create business model ideas (as it did in the first iteration). Rather, it prioritizes which ideas to create based on the knowledge stored in *accumulated idea quality*. That is, features that are present in ideas that have been evaluated rather positively in the first iteration are more likely to appear in ideas created in the second iteration than features that have been evaluated rather negatively. With this new set of created ideas, the second iteration continues with *conversion* and *idea evaluation*. In the second iteration, *idea quality accumulation* has a different role compared to the first one. Recall that in the first iteration it simply stored knowledge on the evaluated ideas in *accumulated idea quality*. However, from the second iteration onwards, it needs to integrate knowledge on idea quality from previous iterations with the knowledge generated in the current iteration. As such, *accumulated idea quality* is updated and revised in every iteration through the additional information that is provided by new sets of *evaluated ideas*. With *accumulated idea quality* getting better in every iteration, *idea creation* is likely to also create better ideas in every iteration. Once a predefined stop criterion is met (e.g., number of iterations = 4), the iterations stop and for the last time a set of ideas is created. That set, which hopefully contains only comparably good ideas, is handed over to the user as the output of the idea generation process.

3.5 Principles of Implementation

In the following, we sketch the actions that are needed to bring a design into being on the grounds of the defined design theory. For the most part, these actions relate to refining constructs defined in the architecture. This is a typical feature of design theories as *“a single construct in a [design] theory can represent a sub-system that has its own separate design theory”* (Gregor and Jones 2007, p. 325). The required actions include the following (not necessarily in that order):

1. *Domain* construct*:* Define data structures to capture business model knowledge in the domain construct and populate these data structures with business model knowledge.
2. *Cognitive procedures* construct*:* Define procedures that can be applied to the defined data structure.
3. *Idea creation* construct*:* Define heuristics that describe how ideas are selected for creation based on the aggregated idea quality knowledge.
4. *Idea evaluation and conversion* constructs*:* Select a crowd platform, define how ideas should be presented to the crowd on that platform, and then implement a component that converts ideas from the data structures defined in the domain construct into the format in which the ideas will be shown to the crowd.
5. *Idea quality aggregation* construct*:* Define what algorithms to use and how to parametrize them.
6. *Overall:* Define how many ideas should be created in every iteration, how many ideas should be contained in the output to the user, and which criteria to use for terminating idea creation (e.g., average quality of ideas in the last iteration, number of total iterations).

3.6 Expository Instantiation

In the following, we provide a simple example to illustrate the intuition behind our approach. As our research currently is at a conceptual stage, we have no running prototype yet. This Gregor and Jones (2007) as they recommend providing an instantiation to facilitate communicating the content of a design theory, but nonetheless define an instantiation as optional within their framework. Owing to space reasons, we only loosely follow the principles of implementation defined earlier.

A business model in our example is defined to consist of three components: revenue model, channel, and customer relationship (see Figure 5). Each of these components may take on exactly one value. For example, the revenue model may be *pay per use*. Other values are shown in Figure 5 under *sample design options*. The business model definition and sample design options in Figure 5 correspond to the knowledge captured in the *domain* construct. The *cognitive procedure* that creates new business model ideas in our example simply selects one value for each business model component. This in the first iteration may lead to the two sample business model ideas shown in the right-most column in Figure 5 (we assume that *perfume* is the product that a user of the idea generator has specified). An idea generator would in the first iteration create a number of such sample ideas, would feed them to the crowd for evaluation, and through that evaluation would iteratively learn about what features characterize high-quality business model ideas for the product perfume. On that basis, in the next iteration it could hopefully create better ideas, in the one that follows even better ones, and so on.

The presented example admittedly is very simple. However, note that even with the simple logic contained therein, a reasonable amount of ideas could be created that, depending on the amount and variety of design options being defined, could already be considerably creative. Users of an idea generator could specify arbitrary products/services (as long as the crowd is able to evaluate the corresponding ideas, see design principle 3). The quality assigned to a certain business model idea in the idea evaluation step would depend on the product or service specified by the user (as a business model idea that is suitable for one product might not be suitable for another). To illustrate this point, selecting a Tupperware-like *party at home* as the sales channel might be a good option if the product is perfume, but might be less reasonable if the product is car batteries.

3.7 Testable Propositions

Testable propositions are necessary to guide the evaluation of design theories. Depending on the level of abstraction and the purpose of the proposed designs, a proposition can take different forms which can range from rather general (e.g., *"If a system or method that follows certain principles is instantiated then it will work."*) to more specific (e.g., it will not only work, but *"…it will be better in some way than other systems or methods.")* (Gregor and Jones 2007, p. 327). As the proposed theory pertains to a high-level architecture, it is hardly feasible to derive specific propositions, as the quality of the ideas that result from an implemented idea generator largely depends on a variety of design decisions still to be made upon implementation (cf. 3.3 Principles of Implementation). Nonetheless, based on the reasoning on the limitations of the human cognitive system provided earlier, and given the relative advantages that machines may have, we feel confident to make the following proposition (which unfortunately is hardly falsifiable, but still seems to be the farthest we can get at this moment):

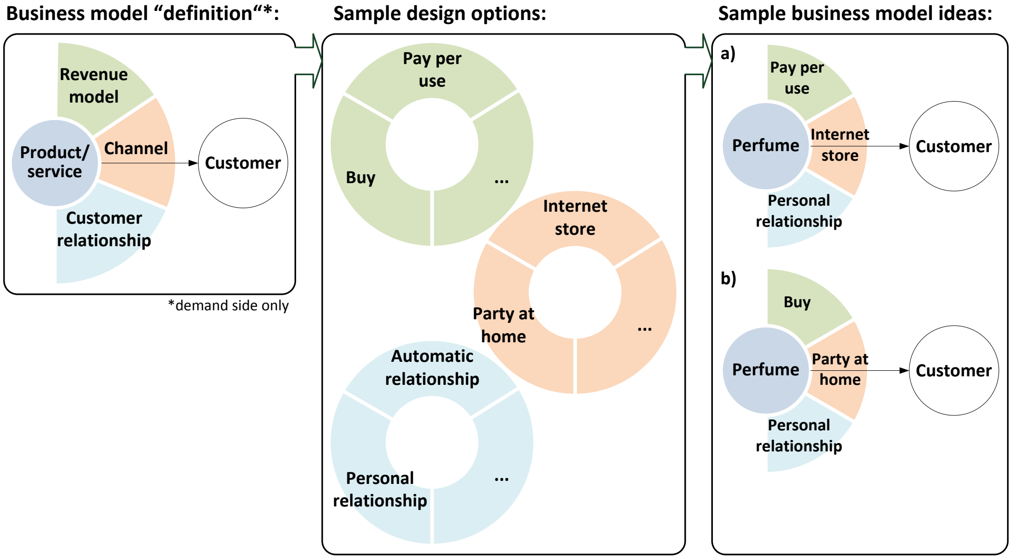

Figure 5: Illustratory example.

Proposition: Business model idea generators instantiated from the proposed high-level architecture have the potential to generate ideas that are of the same (or even better) quality as human-generated ideas.

3.8 Artifact Mutability

Artifact mutability describes possible changes in a design theory (Gregor and Jones 2007). The main types of mutability are *construct mutability*, *model mutability*, *method mutability*, and *instantiation mutability* (Pöppelbuß and Goeken 2015). In our view, there are mainly two types of change that could be foreseen for the proposed architecture. Both would involve changing the architecture in terms of its constructs as well as its principles of form and function. In that sense, these changes would involve a combination of construct, model, and method mutability.

A first type of change could arise from technological progress, which one day might allow to automatically evaluate ideas, and thus would dispense with the need to include a crowd in the architecture. However, as noted earlier, this is highly unlikely in the foreseeable future. The second type of change, however, is more realistic even in the short term. At the moment, the designed architecture prescribes that *accumulated idea quality* is cleared every time the user starts the idea generator with a new product or service. However, having used the idea generator many times probably reveals patterns that are stable across product/service contexts. That is, probably there are certain business models which never seem to be evaluated as *'good'*. Such knowledge could then help to improve the quality of created ideas. However, we would have to keep in mind that especially rare and highly unlikely combinations may be sources of radically new ideas. So we would need to exercise caution when preventing certain business models from being created altogether.

4 Proposed Empirical Evaluation

For evaluating the proposed theory, a system would need to be instantiated from our theory using the principles of implementation stated earlier. The theory could be evaluated by using a controlled experiment, as such experiments are widely accepted to test how a given tool or technique affects creative performance. Creative performance is typically operationalized through expert raters who blind-rate the ideas generated in each experimental condition (i.e., they are blind to the purpose of the experiment and do not know in which experimental condition a given idea has been created). For evaluating our theory, there needs to be a sample task, that is, a sample product or service for which business model ideas are to be created (e.g., perfume). This task would be worked on in a between-subjects design with the following conditions:

1. *Human only:* control condition,
2. *Machine prepares:* treatment 1, see Figure 3,
3. *Machine complements:* treatment 2, see Figure 3,
4. *Machine only:* treatment 3, this would use the ideas as they are generated by the idea generator, without them being refined by humans. This condition would not be externally valid because it is unrealistic that humans feel they have to use ideas generated by the machine as they are, not being able to refine these ideas. Still,

this condition would be interesting to, in a way, determine some *'base level'* performance of the idea generator.

Expert ratings of the ideas generated in these four conditions could shed light on how much (if any at all) value there is in using an idea generator, and how the idea generator should be used (i.e., rather for preparing or for complementing).

5 Discussion and Conclusion

Business model innovation has become a key factor for firm success (Chesbrough 2010). Nonetheless, especially the first phase of the business model innovation process, that of business model idea generation, is still poorly understood (Martins et al. 2015, p. 8), and more research has been called for to address this issue (Schneider and Spieth 2013). As a response to this call, we propose a machine-centered approach to support business model idea generation. We envision that business model idea generators can be built that develop business model ideas for a given product or service. These ideas would stimulate and complement idea generation performed by humans. The advantages that such idea generators may have arise from the idea generators' potential to alleviate the limitations that humans have when trying to be creative. These advantages mainly include the following: First, an idea generator may potentially have stored a lot more business model knowledge than a single human being (or even a group) can possess, because that knowledge could be compiled by many individuals who could draw on virtually any industry that exists (cf. Figure 2: unlimited domain expertise). Second, an idea generator would not be biased towards creating ideas from related industries as it is the case for humans (cf. Figure 2: unlimited cognitive flexibility). Third, an idea generator does not lose its ability to create ideas from distant industries as it acquires more domain knowledge (cf. Figure 2: no trade-off between domain expertise and cognitive flexibility).

Our contribution is a design theory that, drawing on research in creativity, collective intelligence, and machine learning, describes the high-level architecture of the business model idea generators that we envision. Our theory addresses all components of a design theory as proposed by Gregor and Jones (2007) (except for artifact mutability), and also comes with a sketch of a potential research design for empirical evaluation. We believe that our theory has the potential to inform researchers and practitioners in designing the idea generation systems that we propose. However, our research currently is at a conceptual stage, and still lacks empirical evaluation. In addition, we developed the idea generator architecture from a range of theoretical and empirical contributions that largely center around psychological creativity research. However, with computational creativity (e.g., Colton and Wiggins 2012; McCormack and d'Inverno 2014) there is a research field that has a complementary in that it has a more pragmatic and technical, rather than theoretical, perspective on machine-centered idea generation. In the next steps of our research, we plan to integrate contributions from that field into the base of justificatory knowledge of our design theory. Moreover, while our high-level approach is in line with the notion that design theories may contain constructs that in themselves contain sub-theories again (cf. 3.3 Principles of Implementation), we are aware that quite some effort will be needed to refine the high-level constructs that we propose in ways that allow implementing a running system. Still, we deem this effort justified as the resulting idea generators could take business model idea generation to a whole new level.

References

Abdelkafi, N., Makhotin, S., and Posselt, T. 2013. "Business Model Innovations For Electric Mobility—What Can Be Learned From Existing Business Model Patterns?" *International Journal of Innovation Management* (17:01).

Adner, R., Csaszar, F. A., and Zemsky, P. B. 2014. "Positioning on a Multiattribute Landscape," *Management Science* (60:11), pp. 2794–2815.

Amabile, M., and Hennessey, B. 2010. "Creativity," *Annual Review of Psychology* (61:1), pp. 569–598.

Armstrong, S. J., and Mahmud, A. 2008. "Experiential learning and the acquisition of managerial tacit knowledge," *Academy of Management Learning & Education* (7:2), pp. 189–208.

Baer, J., and McKool, S. S. 2014. "The Gold Standard for Assessing Creativity," *International Journal of Quality Assurance in Engineering and Technology Education (IJQAETE)* (3:1), pp. 81–93.

Campbell, D. T. 1960. "Blind variation and selective retentions in creative thought as in other knowledge processes," *Psychological review* (67:6), p. 380.

Casadesus-Masanell, R., and Ricart, J. E. 2010. "From strategy to business models and onto tactics," *Long Range Planning* (43:2-3), pp. 195–215.

Chesbrough, H. 2010. "Business model innovation: Opportunities and barriers," *Long Range Planning* (43:2-3), pp. 354–363.

Colton, S., and Wiggins, G. 2012. "Computational creativity: the final frontier?" *ECAI* .

Dane, E. 2010. "Reconsidering the trade-off between expertise and flexibility: A cognitive entrenchment perspective," *Academy of Management Review* (35:4), pp. 579–603.

Ende, J., Frederiksen, L., and Prencipe, A. 2014. "The Front End of Innovation: Organizing Search for Ideas," *Journal of Product Innovation Management* .

Ericsson, K. A., and Charness, N. 1994. "Expert performance: Its structure and acquisition," *American Psychologist* (49:8), p. 725.

Finke, R. A., Ward, T. B., and Smith, S. M. 1992. "Creative cognition: Theory, research, and applications," .

Fiske, S. T., and Taylor, S. E. 1991. "Social cognition, 2nd," *NY: McGraw-Hill* , p. 16.

Gassmann, O., Frankenberger, K., and Csik, M. 2014. *The Business Model Navigator: 55 Models That Will Revolutionise Your Business:* Pearson UK.

GE 2014. "GE Global Innovation Barometer 2014," .

Gordijn, J., and Akkermans, H. 2003. "Value-based requirements engineering: Exploring innovative e-commerce ideas," *Requirements Engineering* (8:2), pp. 114–134.

Gregor, S., and Jones, D. 2007. "The anatomy of a design theory," *Journal of the Association for Information Systems* (8:5), pp. 312–335.

Hansen, M. T., and Birkinshaw, J. 2007. "The innovation value chain," *Harvard Business Review* (85:6), p. 121.

IBM 2006. "Expanding the innovation horizon: Global CEO study 2006," IBM Business Consulting Services (ed.), Somers, NY.

IBM 2008. "The enterprise of the future: Global CEO study," IBM Institute for Business Value (ed.), Somers, NY.

IBM 2010. "Capitalizing on complexity: Insights from the global chief executive officer study," IBM Institute for Business Value (ed.), Somers, NY.

IBM 2012. "Leading through connections: Insights from the global chief executive officer study," IBM Institute for Business Value (ed.), Somers, NY.

Jordan, M. I., and Mitchell, T. M. 2015. "Machine learning: Trends, perspectives, and prospects," *Science* (349:6245), pp. 255–260.

Kohonen, T. 2012. *Self-organization and associative memory:* Springer.

Kornish, L. J., and Ulrich, K. T. 2011. "Opportunity spaces in innovation: Empirical analysis of large samples of ideas," *Management Science* (57:1), pp. 107–128.

Kornish, L. J., and Ulrich, K. T. 2014. "The importance of the raw idea in innovation: Testing the sow's ear hypothesis," *Journal of Marketing Research* (51:1), pp. 14–26.

Magnusson, P. R., Wästlund, E., and Netz, J. 2014. "Exploring Users' Appropriateness as a Proxy for Experts When Screening New Product/Service Ideas," *Journal of Product Innovation Management* .

Martins, L. L., Rindova, V. P., and Greenbaum, B. E. 2015. "Unlocking the hidden value of concepts: a cognitive approach to business model innovation," *Strategic Entrepreneurship Journal* (9:1), pp. 99–117.

McCormack, J., and d'Inverno, M. 2014. "On the Future of Computers and Creativity," .

Mollick, E., and Nanda, R. 2015. "Wisdom or madness? Comparing crowds with expert evaluation in funding the arts," *Management Science* (in press), pp. 1–21.

Mumford, M., Mobley, M., Reiter-Palmon, R., Uhlman, C., and Doares, L. 1991. "Process analytic models of creative capacities," *Creativity Research Journal* (4:2), pp. 91–122.

Osterwalder, A., and Pigneur, Y. 2010. *Business model generation: A handbook for visionaries, game changers, and challengers*, Hoboken, NJ: Wiley.

Osterwalder, A., and Pigneur, Y. 2013. "Designing business models and similar strategic objects: The contribution of IS," *Journal of the Association for Information Systems* (14:Special Issue), pp. 237–244.

Pöppelbuß, J., and Goeken, M. 2015. "Understanding the Elusive Black Box of Artifact Mutability," *Wirtschaftsinformatik Proceedings* .

Priem, R. L., Butler, J. E., and Li, S. 2013. "Toward reimagining strategy research: retrospection and prospection on the 2011 AMR decade award article," *Academy of Management Review* (38:4), pp. 471–489.

Schneider, S., and Spieth, P. 2013. "Business model innovation: Towards an integrated future research agenda," *International Journal of Innovation Management* (17:1).

Simon, H. 1996. *The sciences of the artificial*, Cambridge: MIT Press.

Simonton, D. K. 2011. "Creativity and discovery as blind variation: Campbell's (1960) BVSR model after the half-century mark," *Review of General Psychology* (15:2), p. 158.

Smith, G. F. 1998. "Idea generation techniques: A formulary of active ingredients," *The Journal of Creative Behavior* (32:2), pp. 107–133.

Stampfl, G., and Sniukas, M. 2013. *'Borrow' business models to reinvent your industry*. http://www.innovationmanagement.se/2013/11/04/borrow-business-models-to-reinvent-your-industry/. Accessed 22 September 2015.

Teece, D. J. 2010. "Business models, business strategy and innovation," *Long Range Planning* (43:2-3), pp. 172–194.

Zott, C., Amit, R., and Massa, L. 2011. "The business model: Recent developments and future research," *Journal of Management* (37:4), pp. 1019–1042.