

The logo features the Lucidworks brand mark, a red square with a white 'L' shape, followed by the word 'Lucidworks' in a white sans-serif font. Below this, the words 'Fusion Files' are written in a larger, white sans-serif font. The background of the top section is a dark blue and black space filled with a complex network of glowing red and blue lines and dots, resembling a data visualization or a neural network.

Lucidworks Fusion Files

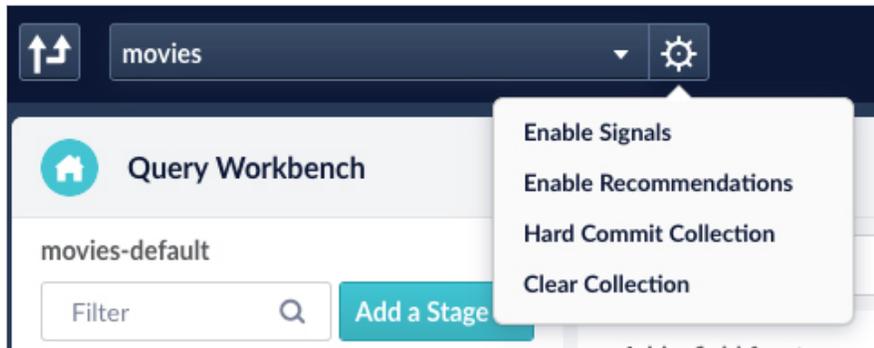
4 AI Quick Hits You Can Add to Your Enterprise Search App Right Now

Enterprise search is one of the most underappreciated and yet most critical part of any organization's infrastructure. It doesn't directly generate revenue, but if you get it wrong, it is a tax on every employee and business process. If you get it right, it becomes a catalyst to quicker decisions and growth. Conventional keyword search isn't enough anymore. New techniques and technology can find data for users faster and with less effort. Whether you call these techniques cognitive search, machine learning, artificial intelligence, or the more mundane "relevance tuning," you need to do your very best work here if you want to get the most out of your enterprise search investment.

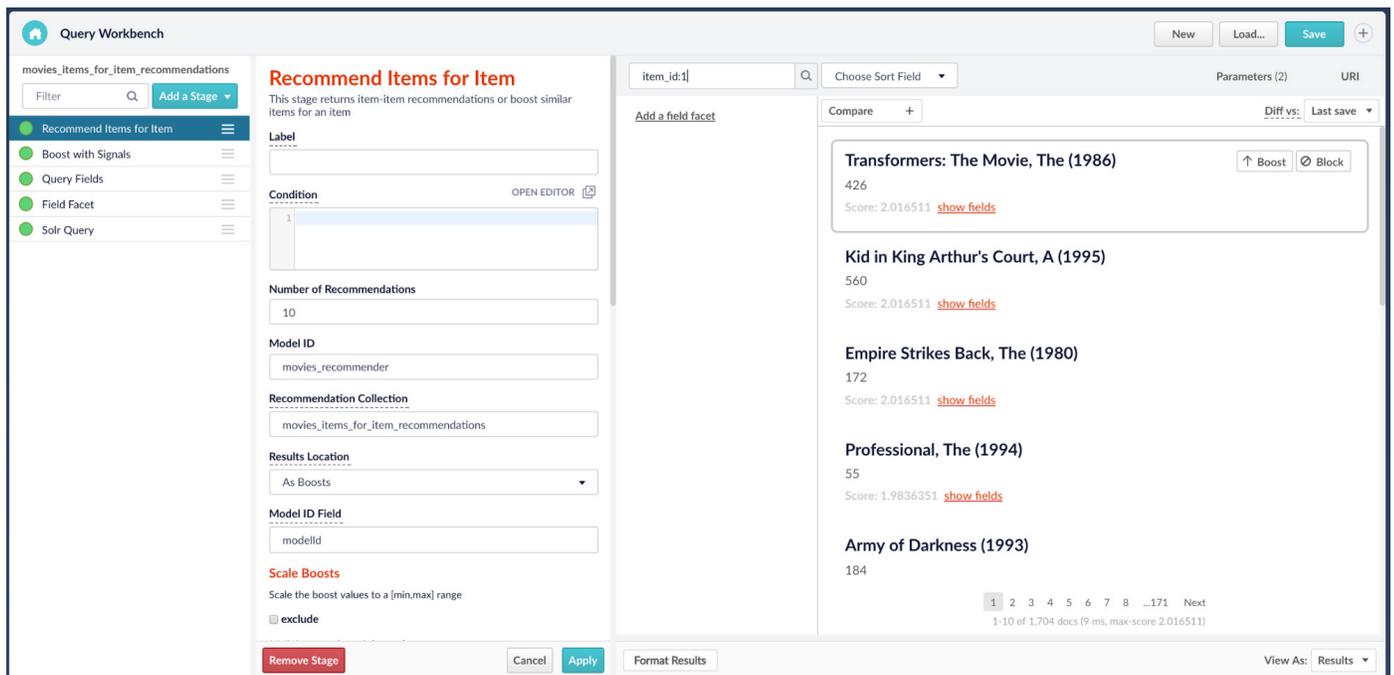
Luckily, our Lucidworks Fusion platform enables these capabilities out-of-the-box ready-to-go with little effort on your part.

Recommenders

You know how Netflix says “Because you watched” and then recommends shows for you? You can have this in your app too.



This is an example of recommendation. You can capture “signal” data such as which items a user clicks on then use that information to make recommendations to other users. When enough users who click on “Toy Story” and then click on “Transformers: The Movie” your search application can then recommend that to other users who click on “Toy Story.”

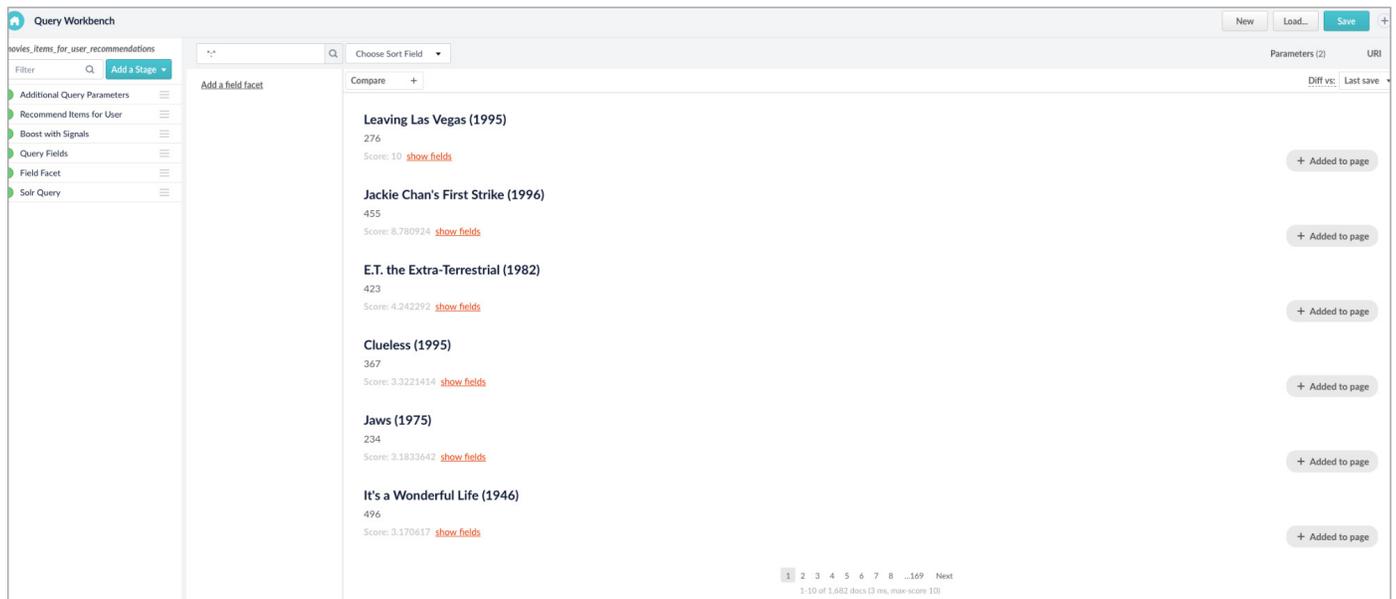


Fusion has different types of recommendations. One such recommender allows you to get **recommendations based on the item you have selected**. Fusion can capture which users click on one item and also click on a different item. For example, when a user clicks on “Toy Story” your application can then suggest the items that other

users clicked on who also clicked on Toy Story. This isn't just for movies, for Enterprise Search, consider that when users that click on one document in Sharepoint and tend to click on another document?

These recommendations are useful for a search application once a user has selected an item. One of the finest things you can do for your returning users is recommend things before they've clicked on anything.

Fusion can also **recommend items based on the user's click history** rather than a specific item. For Enterprise search, most users on a project or in a similar role are likely to look at similar things. Why not let search lead them there without the hunt?



Analyzing past behavior isn't the only way that you can "recommend" things to users. Fusion can also help first time users. **Recommend More Like This** does a similarity search on the actual item and recommends items with similar descriptions or phrases.

Phrase Detection

From politics to finance and everything in between, not every "thing" can be defined by just one word. If I'm looking for articles about "fake news" or "income tax," then content that has the words "fake" and "news" or "income" and "tax" do not meet my needs. While relevance may boost documents with the words in close proximity to the top, what I really want to do is exclude a document that includes the phrase, "Higher income clients in nicer neighborhoods may avoid paying city property taxes by moving outside the city limits." I'm strictly looking for documents about "income tax."

Phrase detection is built into the Solr component of the Fusion stack. It is a little

more effort than synonym detection, but the payoff is important, especially in jargon-rich industries like finance or manufacturing. Check out [this blog post about automatic phrasing](#).

However, doing this with just vanilla Solr isn't as easy as it should be. Configuration isn't really the hard part. Figuring out what phrases people might search on is the hard part. This is where Fusion's Query Explorer is critical.

Fusion can automatically identify statistically interesting phrases (SIP) in a collection of documents. You can use those to configure your automatic phrasing. Read up on [Fusion's Query Explorer](#).

4a0c736d-c8e4-4bfd-bfa9-f0be8b52fff6

inputSips
Score: 1 [hide fields](#) + Added to page

collection	inputSips
id	4a0c736d-c8e4-4bfd-bfa9-f0be8b52fff6
llr	23.23540324730675
ngram	["stylish","look"]
score	1
timestamp	2017-05-03T07:00:00Z
version	1566410020385456000

52b8cf0c-fe52-4474-a599-8c39e07acaf4

inputSips
Score: 1 [hide fields](#) + Added to page ↑ Boost ⊗ Block

collection	inputSips
id	52b8cf0c-fe52-4474-a599-8c39e07acaf4
llr	23.23540324730675
ngram	["holds","up"]
score	1
timestamp	2017-05-03T07:00:00Z
version	1566410020385456000

Spell Checking

Speaking of what users type, one of the great ironies of the modern age is that the predominant form of communication has become text-driven, while people have become less concerned with using good grammar or spelling words correctly. This can be a challenge for search applications. Users searching on something may spell words incorrectly in their queries but documents in the collections may also have spelling errors, despite the prevalence of spellcheck in word processing tools.

Lucidworks gives you multiple tools for analyzing your content and queries, and suggesting a corrected or at least matching spelling. The Solr component of Fusion offers spellchecking based on the popular [Levenshtein Distance](#) algorithm. Documentation for how to configure that is in the [Solr Reference Guide](#).

Fusion goes above and beyond that, allowing you to pull high-quality analytics from your index using [Fusion's Query Explorer](#).

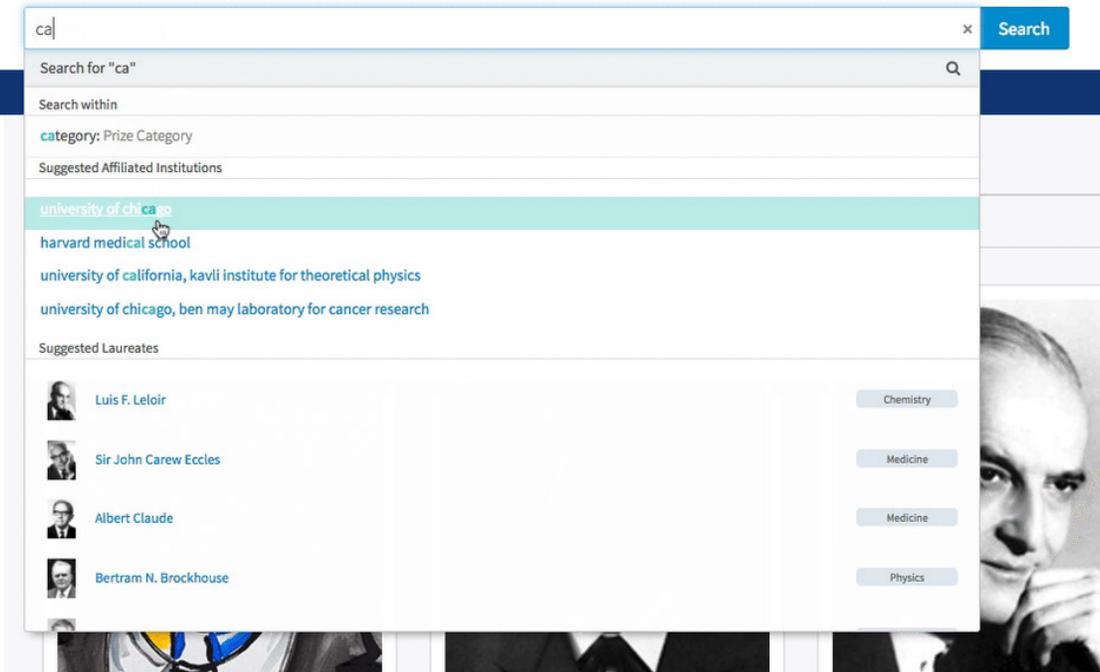
bdd96fb8-c35e-462a-8006-b915660788d7 ↑ Boost 🚫 Block

1

Score: 1 [hide fields](#) + Added to page

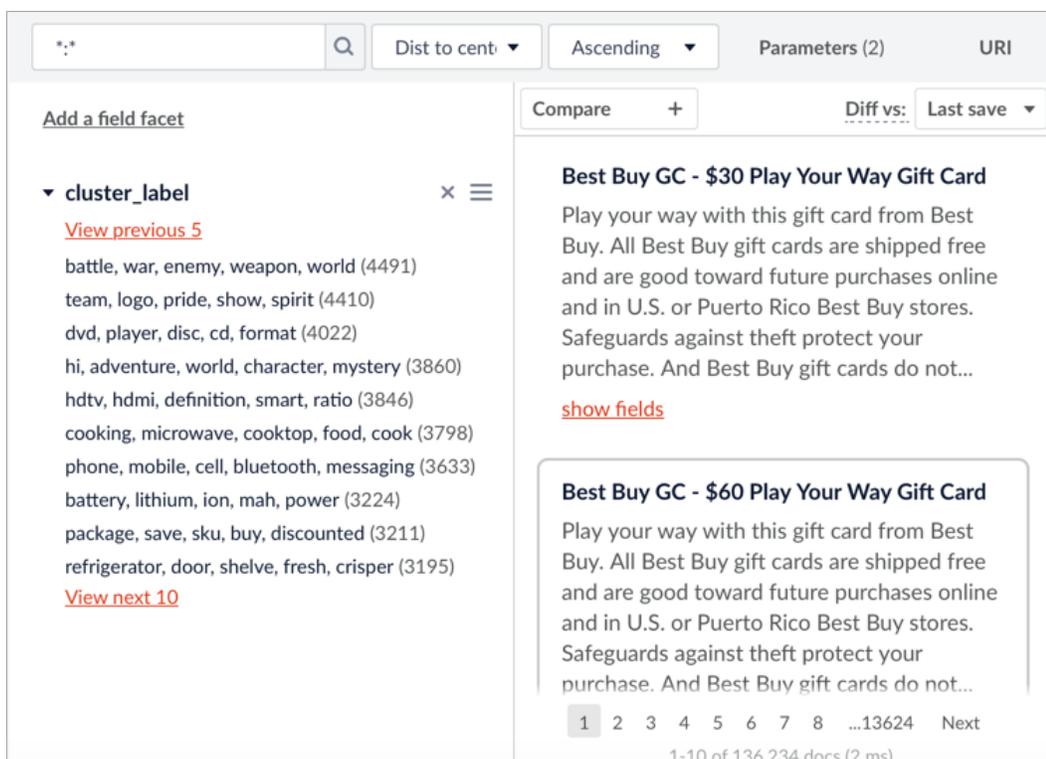
distance	1
headQuery	hp touchpad
headQueryCount	24182
id	bdd96fb8-c35e-462a-8006-b915660788d7
otherQuery	hp toucpad
otherQueryCount	26
score	1
version	1567797927890387000

While Solr and Fusion can correct spelling mistakes in queries, wouldn't it be better if users just didn't make spelling mistakes in the first place? They don't have to! Building upon [Solr's auto-complete functionality](#), Fusion [App Studio can integrate autosuggestion and typeahead](#) into your search application with no additional code required.



Document Clustering

Some things just go together. It isn't just pizza and beer, but your organization's documents and data probably have particular groupings and relationships. Companies used to spend mountains of cash on filing systems and trying to develop perfect hierarchical ontologies, but it never really captures all of the connections. Fortunately, Fusion can [algorithmically cluster documents in a collection](#).

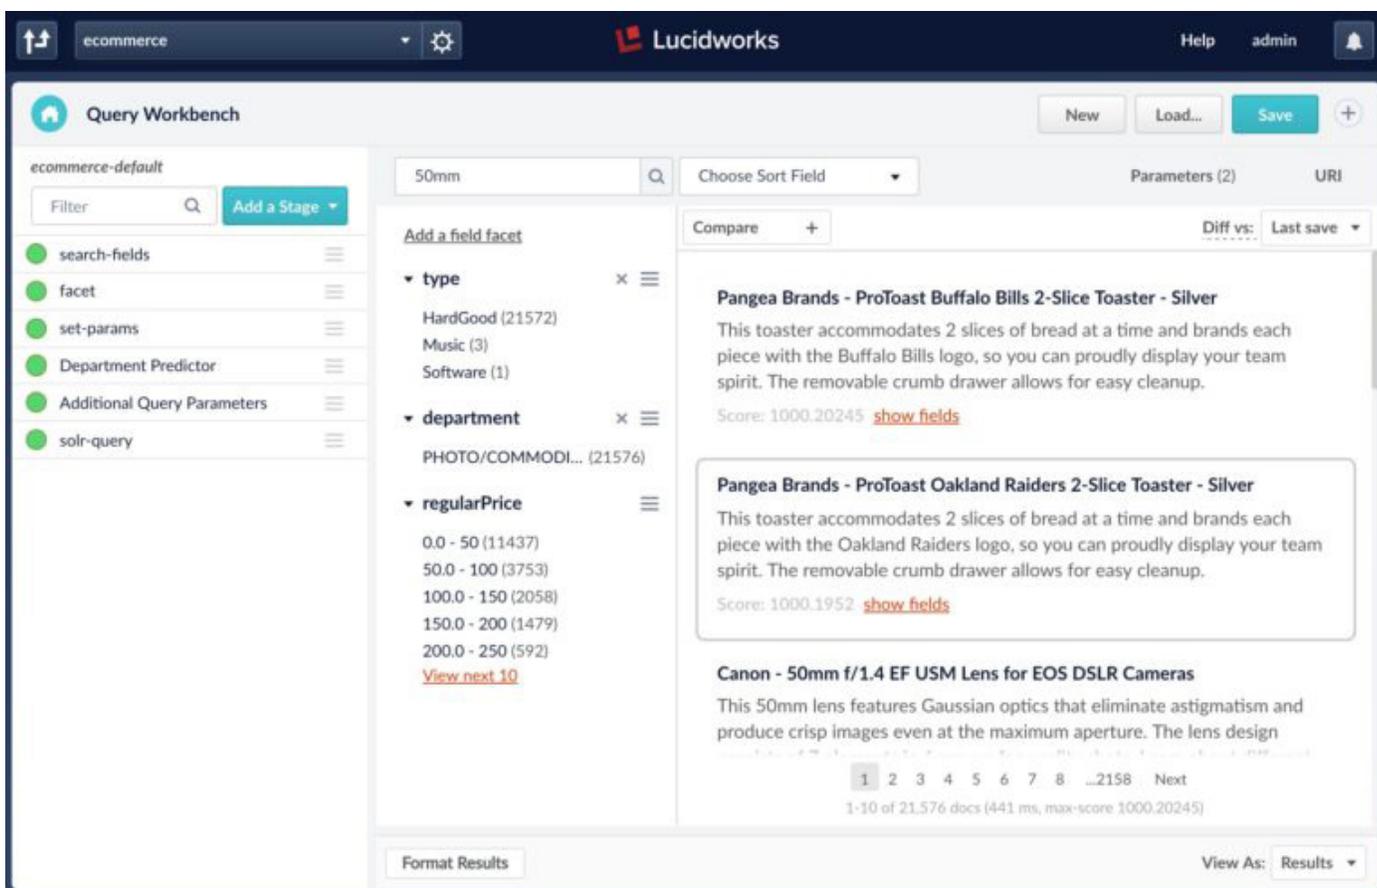


This has many uses from finding outliers to automatically finding categories to making suggestions to users based on other items they're searching on. Think about the last time you tried to figure out how your SharePoint is really organized. The search is bad and the organization of the material doesn't help you navigate to what you're looking for. This is one more way that Fusion can help automatically organize information.

Classification

Clustering is another method for finding groups. Unlike "classification," with clustering you don't know what the groups should be in advance.

Unlike clustering where Fusion just finds things that look like they are "close" together, with Fusion's Classification job you are in the driver's seat. You pick the categories.



If you already have categories and documents that are sorted into those categories, you have what we'd like to call "a training set for a Fusion Classification job" but we'll just call it a "training set" for short. Basically Fusion can finish that painstaking task of sorting things into categories for you. You point Fusion to a training set of categorized documents in order to create a "model." Then point that model at a new set of similar

documents to be classified. From this will come your facets (categories). There are step by step instructions on how to implement this [here](#).

If you need your SharePoint documents organized by project or functional group, but can't stand the thought of painstakingly doing it yourself, Fusion can do that as well.

While this is obviously really useful for enterprise search users trying to organize their world, any retailer looking to get their categories straight before the shopping season begins in earnest should probably be leveraging Fusion's classifiers right about now!

Next Steps

This isn't intended to be an exhaustive list, but just these four quick wins are simple enough that your team can probably execute them in the next few months. This short article barely scratches the surface of what Lucidworks Fusion has to offer. Fusion's broad palette of capabilities includes [recommendations](#), [analytics](#) (you can even [connect your BI tool and run SQL queries](#)), [entity extraction](#) or [tag parts-of-speech](#). You can even [talk to it if you want!](#)

Lucidworks has out-of-the-box functionality available that you can add to your application with little to no extra code. If you're not already using Fusion, the average implementation time for building a Fusion app is not very long. It's time to get more value out of your data and accelerate your employees. It is time for Fusion.

Get Started or Learn More

For more information or to start using Lucidworks Fusion, contact us today to learn more at <http://lucidworks.com/contact> or call 415-329-6515.